PERMIT: Permission-Aware Representation Intervention for Controlled Generation in Large Language Models

Pengcheng Sun[1], Lan Zhang[1,2†], Zhaopeng Zhang[1], Jiewei Lai[1], Chen Tang[1]
[1]University of Science and Technology of China
[2]Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

Abstract

Large language models (LLMs) are increasingly deployed in enterprise settings where they handle sensitive documents and user context, raising acute concerns over security and controllability. Conventional access control regulates *whether* information is accessible to the model, yet leaves *how* the model uses that information at generation time largely unconstrained: once sensitive content enters the context, outputs may still drift beyond a user's authorized scope. We present PERMIT, a novel permission-aware representation intervention framework that closes this gap by enforcing fine-grained control directly on the model's hidden states. Through exploratory analysis, we find that permission conditions induce hidden-state shifts that are (i) separable across permissions and (ii) concentrated in a small set of dominant directions. PERMIT exploits this geometry in two stages: it first identifies a *permission-sensitive subspace* from activation differences across permission conditions, and then performs lightweight interventions within this subspace to steer generation, with two concrete instantiations (*offset-based* and *gated*). Both operate atop a frozen backbone with only a handful of permission-specific parameters, achieving precise control with minimal overhead. Experimental results demonstrate that PERMIT achieves a better utility-security trade-off than existing approaches across multiple permission settings, reducing leakage to near zero while improving F1-score over strong baselines with more than 98% fewer trainable parameters.

1 Introduction

Large language models (LLMs) have demonstrated remarkable capabilities across a wide range of tasks [7, 9], fueling growing interest in deploying them within enterprise settings [6, 24, 34] for applications such as knowledge query and data analysis. Yet processing sensitive information (e.g., internal documents, chat histories, proprietary records) introduces security risks that have become a major barrier to broader deployment [25]. A foundational requirement is *access control*: even when given the same input, users with different permissions must receive different, authorization-compliant outputs [5]. Real-world permission policies are typically fine-grained and hierarchical, and equipping LLMs to faithfully respect such structured policies remains a core open challenge [15, 18, 29, 39].

Existing work falls broadly into two paradigms. *Alignment-based* methods rely on supervised fine-tuning [26] or RLHF [18, 22] to instill globally safe behavior and refuse harmful requests [11]; while effective at coarse, binary safety boundaries, they struggle to faithfully model multi-role, hierarchical permissions [27]. *Rule-based* methods instead bind explicit policies to data, roles, or model components via policy-specific fine-tuning [1, 29], parameter encryption [37], or per-domain

†Corresponding author. Correspondence to zhanglan@ustc.edu.cn.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

LoRA adapters [12, 13, 15]. These approaches more closely mirror traditional access-control systems, but rely on static partitioning that scales poorly as the number of roles and permission levels grows.

More fundamentally, both paradigms share a common blind spot. Once sensitive content enters the model's usable context—via retrieval-augmented generation, user input, or external tools—the model can still produce content beyond the user's authorized scope [20]. This exposes a key disconnect between whether information is made accessible to the model and how the model uses that information during generation. Existing methods regulate the former but offer little leverage over the latter. Therefore, a central question is: *how can we enforce fine-grained control over information use at generation time, across different permission levels, without sacrificing utility?*

Answering this question requires confronting two fundamental difficulties. **(C1) Entanglement.** Complex authorized and unauthorized content co-occur in the context and become entangled in the model's internal representations, making selective suppression hard without degrading fluency or factuality. **(C2) Security–utility tension.** Overly conservative control causes over-refusal, while insufficient control leaks information. Striking a fine-grained balance across many permission levels rather than a single global safety setting remains a challenging problem.

This work. We address this disconnect by enforcing permissions at the *generation stage*, regulating how information is used rather than merely whether it is accessible. Instead of treating permissions as external rules, we seek internal representations that encode permission differences and use them to directly steer generation. This view is motivated by recent advances in activation engineering [17, 32, 36], which show that targeted interventions on hidden states can reliably steer LLM behavior without modifying the backbone. Through exploratory analyses, we uncover two complementary phenomena: **(1) Activation-space separability:** hidden states under different permissions form well-separated clusters whose shifts reflect the hierarchical structure of the underlying policies; and **(2) Low-rank concentration:** these permission-induced shifts concentrate along a small number of dominant directions and are well approximated by a low-dimensional subspace.

Building on these findings, we propose PERMIT, a permission-aware representation intervention framework for fine-grained controlled generation. PERMIT proceeds in two stages. Specifically, it first learns a *permission-sensitive subspace* from activation differences across permission conditions, capturing the principal directions of permission-related variation (addressing **C1**). Then it performs interventions within it via a unified subspace intervention framework with two instantiations: an *offset-based* intervention realigns the desired representations through an affine transformation for control under strict permission boundaries, while a *gated* intervention applies dimension-wise scaling for fine modulation along the security–utility frontier (addressing **C2**). Both keep the backbone LLM frozen and introduce only a small number of permission-specific parameters.

Overall, our **main contributions** are as follows: (1) *Perspective.* We reframe access control in LLMs as representation-space intervention at the generation stage, enabling plug-and-play enforcement and complementing existing data-level approaches. (2) *Method.* We show that permission-induced representation shifts are both separable and approximately low-rank, and propose a novel permission-sensitive subspace intervention framework for fine-grained access control. (3) *Results.* Extensive experiments show that PERMIT achieves a better utility–security trade-off than baselines: it drives information leakage to near zero while improving F1-score, with $> 98\%$ fewer trainable parameters.

2 Preliminary

2.1 Problem Statement

Permission setting. Let $\mathcal{R} = \{r_k\}_{k=1}^{N}$ denote a set of permission states, where each r_k specifies a role–domain pair u_k and a permission level ℓ_k, defining a set of authorized fields A_k and restricted fields F_k over the corresponding record schema. We assume that each r_k is a trusted authorization condition provided by an external authentication layer. It specifies which information may be disclosed to the requester and at what level of granularity. For any two permission states r_{k_1} and r_{k_2} sharing the same role–domain pair, the permission hierarchy satisfies $A_{k_1} \subseteq A_{k_2}$ and $F_{k_2} \subseteq F_{k_1}$ whenever $\ell_{k_1} < \ell_{k_2}$. Given the same query–context pair (q, c), the permission-compliant response y^k may contain only information authorized by A_k at the permitted granularity; disclosure of information in F_k is treated as unauthorized leakage. Appendix C.3 provides a concrete example of this setting.

For each permission state r_k, let $\mathcal{D}_k = \{(q_i^k, c_i^k, y_i^k)\}_{i=1}^{n_k}$ be the corresponding permission-conditioned dataset, where q_i^k is a query, c_i^k is the available context, and y_i^k is the permission-compliant response. Here, n_k denotes the number of queries for r_k. Given a base LLM M_0 with parameters Θ_0, the model produces a generic output $\hat{y}_i^0 = M_0(q_i^k, c_i^k)$, which does not explicitly enforce permission-aware semantic control. As a result, the model may reveal unauthorized information by using, combining, or inferring sensitive content from the context or its parametric knowledge.

Our goal is to learn a lightweight permission-conditioned intervention module Δ_k for each permission state r_k. The resulting controlled model generates $\hat{y}_i^k = M_0\big(q_i^k, c_i^k; \Delta_k \mid r_k\big)$, where Δ_k adjusts the model behavior at inference time without updating the backbone parameters Θ_0. The controlled output is expected to closely match the desired permission-compliant response y_i^k, minimize unauthorized disclosure, and introduce only minimal parameter and inference overhead. Formally, we optimize the aggregate permission-conditioned intervention objective:

$$\min_{\{\Delta_k\}_{k=1}^N} \sum_{k=1}^{N} \sum_{i=1}^{n_k} \mathcal{L}\big(M_0(q_i^k, c_i^k; \Delta_k \mid r_k), y_i^k\big),$$

where $\mathcal{L}$ denotes the loss against the permission-compliant target.

2.2 Representation Intervention

Transformer Hidden-State Representations. We consider a decoder-only Transformer [33] with L layers and hidden dimension d. Given an input sequence $x_{1:t}$, let $h_t^\ell \in \mathbb{R}^d$ denote the hidden state of token t at layer ℓ. Following the standard residual formulation, each Transformer block updates the hidden state as

$$\tilde{h}_t^\ell = h_t^\ell + \mathrm{MHA}^\ell(h_{1:t}^\ell), \qquad h_t^{\ell+1} = \tilde{h}_t^\ell + \mathrm{FFN}^\ell(\tilde{h}_t^\ell),$$

where MHA^ℓ and FFN^ℓ denote the multi-head attention and feed-forward modules at layer ℓ, respectively. In our setting, the model takes a query–context pair and a permission condition as input, and its output behavior is closely tied to the hidden-state trajectory across layers. This provides the basis for permission-aware control through representation intervention.

Activation Steering under Permission Conditions. Recent studies [2, 16, 31] show that certain semantic attributes in LLMs can often be represented as controllable directions in hidden space, enabling behavior steering through activation-level interventions [4, 35, 42]. Typically, activation steering adds a vector $v^\ell \in \mathbb{R}^d$ to the hidden state at layer ℓ: $\tilde{h}_t^\ell = h_t^\ell + v^\ell$, which shifts generation behavior along a desired semantic direction. Motivated by this formulation, we study whether permission conditions can similarly be encoded in intermediate representations and used for controlled generation. More generally, a permission-conditioned intervention can be written as $\tilde{h}_t^\ell = \phi(h_t^\ell; r)$, where ϕ modifies the hidden representation according to the permission condition r before it is passed to subsequent layers.

3 Permission Structure in Hidden Representations

Rather than treating permissions purely as external symbolic rules, we revisit access control from a representation learning perspective and pose a basic question: *do permission conditions induce systematic, identifiable changes in the hidden representations of large language models?* If such changes exhibit consistent geometric structure rather than arbitrary variation, they offer a principled foundation for enforcing permission constraints through lightweight interventions in activation space.

3.1 Extracting Permission-Induced Representation Shifts

To study this question, we compare hidden representations of a baseline input against its permission-conditioned counterpart, constructed from the same underlying query–context pair. Let x_n^o denote the original, permission-free input for the n-th query–context pair, and let $x_{k,n}^p$ denote the corresponding input conditioned on permission state r_k. Specifically, the permission condition is applied through a permission-specific system prompt that explicitly defines the user's access level before the query and

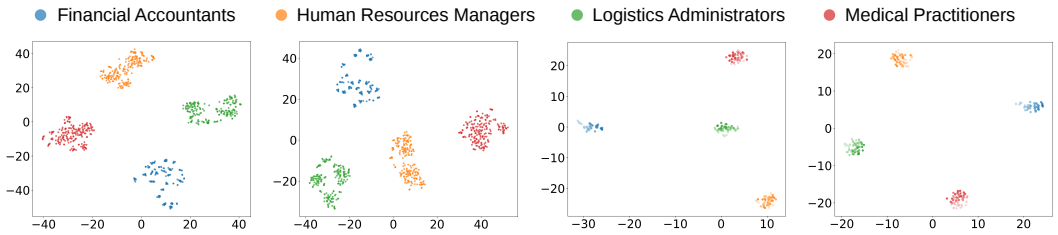

(a) LLaMA3.1-8B (b) Qwen2.5-7B (c) LLaMA3.1-8B (d) Qwen2.5-7B

Figure 1: Activation-space structure of different permission conditions. Subfigures (a) and (b) show that hidden states under different permissions are well separated across both models; (c) and (d) show that mean permission-induced representation shifts exhibit hierarchy-aware directions: related permissions cluster together, and stronger permissions induce larger shifts along consistent directions.

Table 1: Low-rank structure of permission-induced representation shifts. We report the minimum number of dimensions required to explain 80%, 90%, and 95% of the spectral energy, together with their proportions relative to the full hidden dimension across both models.

Model	Hidden dimensions	0.80		0.90		0.95	
		%	rank	%	rank	%	rank
Qwen2.5-7B	3584	0.37	13	1.55	56	4.52	162
LLaMA3.1-8B	4096	0.41	17	1.86	76	5.38	220

context are provided. At layer ℓ, we take the last-token hidden state $h^\ell(\cdot)$ as the model activation, and define the *permission-induced representation shift* as

$$\Delta h^\ell_{k,n} \;=\; h^\ell\left(x^p_{k,n}\right) - h^\ell(x^o_n)\,,$$

which captures the representational change associated with permission state r_k at layer ℓ. Given the structured and hierarchical nature of permission information, we analyze the full set $\{\Delta h^\ell_{k,n}\}_{k,n}$ jointly to capture both their shared regularities and permission-specific variations.

3.2 Key Observations

We conduct exploratory analyses on the MedicalSys dataset [39] to examine whether permission conditions induce structured changes in the model's activation space, and if so, what geometric form that structure takes.

Observation 1 (Activation-space separability)**.** *Permission conditions induce separable and structured activation shifts: distinct permissions occupy distinguishable regions, while semantically related permissions share directions with magnitudes scaling by intensity.*

For each underlying query, we extract hidden states under different permission conditions and project them via t-SNE. Figure 1(a,b) shows that distinct permission roles form well-separated clusters across multiple models, indicating that permission conditioning systematically reshapes internal activations rather than merely altering surface-level outputs. To probe this structure at finer granularity, we partition permissions into 20 fine-grained conditions spanning four domains and compute the mean shift vector for each condition. As shown in Figure 1(c,d), permissions governed by similar rules align along nearby directions, and increasing permission strength produces graded shifts along consistent directions. Overall, these patterns suggest that permission-induced representations contain both shared directional components and fine-grained hierarchical variations.

Observation 2 (Low-rank Concentration)**.** *Permission-induced representation shifts exhibit a low-rank structure, where most permission-related variation is captured by a small number of dominant directions.*

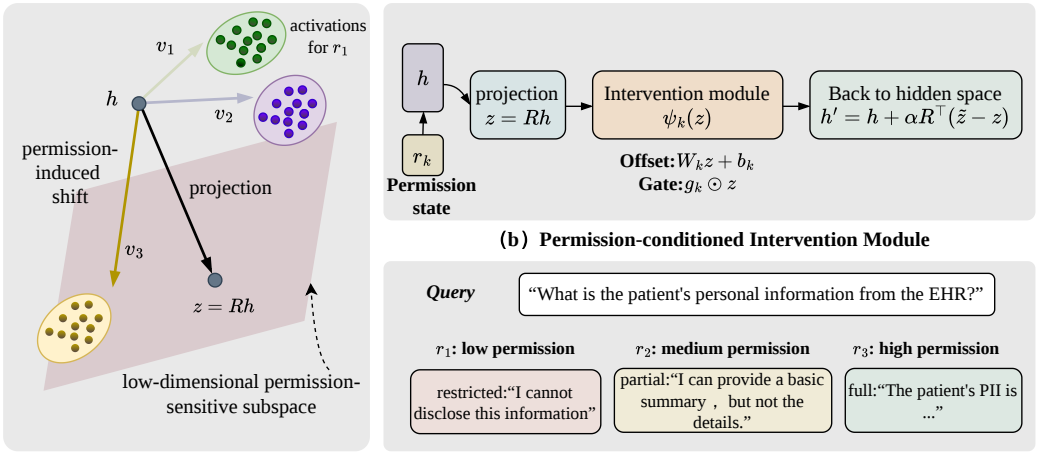

Figure 2: Overview of PERMIT. (a) Permission conditions induce structured shifts in hidden representations, which are captured by a low-rank permission-sensitive subspace. (b) PERMIT projects a hidden state into this subspace, applies a permission-conditioned intervention, and maps the resulting change back to the original hidden space. (c) The intervention controls response granularity according to the granted permission level.

To examine the intrinsic structure of permission-induced representation shifts, we apply singular value decomposition (SVD) to the collection of shift vectors $\{\Delta h^{\ell}_{k,n}\}_{k,n}$ at layer ℓ. As shown in Table 1, the singular value spectrum decays rapidly, indicating that most permission-related variance can be captured by only a small number of principal components. For instance, on Qwen2.5-7B, the effective rank is only about 0.37% of the original hidden dimension. This suggests that permission-induced changes are not uniformly distributed across the full representation space, but are concentrated in a compact set of dominant directions.

Statistical Validation. We further validate both observations with statistical tests. For separability, a linear probe on 64-dimensional PCA representations achieves 98.2% accuracy under grouped 5-fold cross-validation, substantially above the 25% chance level; a 100-run label-permutation test yields $p = 0.0099$. For low-rank concentration, we compare the participation ratio (PR) against a shuffled null obtained by independently permuting samples within each hidden dimension, thereby disrupting cross-dimensional dependence while preserving marginal distributions. The observed PR ranges from 3.8 to 11.1 across permission states, compared with 37.4–128.0 under the null, with all observed values below the corresponding 95% null intervals. Together, these results provide quantitative support for the low-dimensional concentration and linear separability of permission-induced representations.

Discussion. Overall, Observations 1 and Observation 2 together suggest a coherent structural view of permission-conditioned representations. First, although different permission levels may induce distinct generation behaviors, their representation shifts are not mutually independent. Instead, they share a common low-dimensional subspace that captures global permission-relevant variations. Second, individual permissions differ in how they transform or exploit this shared subspace, giving rise to fine-grained, permission-specific variations. This view directly motivates our design of a shared permission-sensitive subspace equipped with lightweight permission-specific transformation parameters for fine-grained access control in LLMs.

4 Methodology

These observations in Section 3 have practical implications: permission execution can be formulated as a structured representation intervention problem, rather than as a purely external filtering rule or prompt-level constraint. Leveraging this insight, we propose PERMIT, a lightweight inference-time framework that jointly learns a low-rank permission-sensitive representation subspace and permission-conditioned intervention parameters, rather than fine-tuning the backbone model (Figure 2). The

learned subspace captures dominant permission-related variation, while the permission-conditioned transformations implement fine-grained control within this compact space. Drawing on the principle of representation fine-tuning [17, 38], we achieve this framework with a unified subspace intervention formulation and two lightweight variants (offset and gated forms) tailored to fine-grained access-controlled generation.

Permission-sensitive subspace. We first define the permission-sensitive subspace based on Observation 2 at layer ℓ as $R_\ell \in \mathbb{R}^{m\times d}, m \ll d$, where R_ℓ is a trainable projection matrix whose rows form an orthonormal basis. This formulation does not assume that permission policies themselves are inherently low-dimensional. Instead, it assumes that the *controllable representation variation relevant to permission execution* can be well approximated by a compact shared subspace.

Unified subspace intervention. Given a hidden state $h_\ell \in \mathbb{R}^d$, we first project it onto the permission-sensitive subspace: $z_\ell = R_\ell h_\ell \in \mathbb{R}^m$. For a permission state r_k, we then apply a permission-conditioned transformation within this subspace: $\tilde{z}_{\ell,k} = \psi_{\ell,k}(z_\ell)$, and map the induced change back to the original hidden space:

$$\phi_{\ell,k}(h_\ell) = h_\ell + \alpha R_\ell^\top (\tilde{z}_{\ell,k} - z_\ell), \tag{1}$$

where α controls the intervention strength. In this way, PERMIT constrains the intervention to the learned permission-sensitive subspace, allowing PERMIT to control a small set of permission-relevant directions while limiting unintended perturbations to the full hidden representation.

Two instantiations: offset and gated forms. We instantiate $\psi_{\ell,k}$ with two lightweight parameterizations. The first is an *offset-based* form:

$$\psi^{\text{off}}_{\ell,k}(z_\ell) = W_{\ell,k} z_\ell + b_{\ell,k},$$

where $W_{\ell,k} \in \mathbb{R}^{m\times m}$ and $b_{\ell,k} \in \mathbb{R}^m$ are trainable parameters. This variant performs a permission-conditioned affine correction over the projected representation, allowing flexible translation and recombination within the shared subspace. The second is a *gated* form:

$$g_{\ell,k} = \sigma(W_{\ell,k} z_\ell + b_{\ell,k}), \qquad \psi^{\text{gate}}_{\ell,k}(z_\ell) = g_{\ell,k} \odot z_\ell,$$

where $\sigma(\cdot)$ is the sigmoid function and $\odot$ denotes element-wise multiplication. The gated form produces a permission-conditioned mask over the projected representation. Since $g_{\ell,k} \in (0,1)^m$, it selectively preserves or suppresses each permission-sensitive dimension, providing a lightweight and conservative mechanism for controlling generation granularity.

5 Experiments

5.1 Experimental Setup

This section describes the experimental setup for evaluating our method, including the datasets, metrics, and baselines. The additional setup details are provided in the Appendix C.

Evaluation Dataset. We conduct our main experiments on MedicalSys [39], a medical system dataset designed for multi-domain access control in healthcare management systems. It contains approximately 16,000 documents covering four representative user roles: medical practitioners, financial accountants, logistics administrators, and human resource managers, each associated with a specific document type and access scenario. To reflect fine-grained access control in realistic deployments, we extend each role in MedicalSys with four hierarchical permission levels, which progressively expand the set of authorized fields according to data sensitivity. For each query–context pair, we construct a level-specific reference response containing only the fields authorized at that level; disclosure of any restricted field is counted as leakage. A concrete example of the four-level construction is provided in Appendix C.3. We additionally evaluate cross-domain generalization on OrgAccess [28], an enterprise-domain benchmark with a distinct access-control task that classifies requests into full access, partial access, or rejection; the setup and results are provided in Appendix C.4.

Table 2: Overall performance on MedicalSys. Precision (Pre.), Recall (Rec.), F1-score (F1), and ROUGE-L (R-L) measure utility; Leakage Rate (L.R.) measures security; and additional trainable parameter ratio (Param.) and inference latency (Lat.) measure efficiency.

Model	Method	Utility ↑				Security ↓	Efficiency ↓	
		Pre.	Rec.	F1	R-L	L.R.	Param. (%)	Lat. (s)
LLaMA3.1-8B	Prompt-Only	0.917	0.807	0.853	0.826	0.968	0	2.25
	Prompt-Perm	**0.916**	0.775	**0.832**	0.806	0.090	**0**	2.32
	ControlNet	0.678	0.584	0.616	0.680	0.010	0.1168	3.00
	AC-LoRA†	0.833	0.778	0.803	0.819	**0.000**	2.7160	**2.11**
	AAAC	0.881	0.709	0.786	0.833	0.015	**0**	3.01
	PERMIT-Offset	0.846	**0.810**	0.828	0.830	**0.000**	0.0018	2.33
	PERMIT-Gate	0.837	0.805	0.819	**0.840**	0.001	0.0018	2.35
Qwen2.5-7B	Prompt-Only	0.873	0.709	0.773	0.786	0.970	0	2.13
	Prompt-Perm	0.568	0.518	0.538	0.694	0.422	**0**	2.29
	ControlNet	0.779	0.550	0.617	0.658	0.177	0.1024	2.93
	AC-LoRA†	0.828	0.764	0.794	0.825	**0.000**	2.1204	**2.20**
	AAAC	0.723	0.722	0.722	**0.895**	0.330	**0**	3.03
	PERMIT-Offset	**0.839**	**0.782**	**0.804**	0.826	0.025	0.0017	2.43
	PERMIT-Gate	0.789	0.662	0.711	0.756	0.110	0.0017	2.64

†AC-LoRA follows its data-isolation setting, where each permission-specific adapter receives only the information authorized for the corresponding permission level.

Evaluation Metrics. We evaluate all methods from three perspectives: utility, security, and efficiency, using LLaMA3.1-8B and Qwen2.5-7B as backbone models. For utility, we use Precision, Recall, F1-score, and ROUGE-L to measure whether the generated responses correctly preserve permission-authorized information under different permission constraints. For security, we report Leakage Rate, which measures the proportion of test instances where the output discloses restricted or sensitive information beyond the user's permission level. For efficiency, we report inference latency and the percentage of additional trainable parameters relative to the backbone size. Specifically, we treat the generated responses under permission constraints as final predictions and evaluate effectiveness by matching predicted factual fields against the corresponding ground-truth.

Baselines. We compare both variants of PERMIT against five representative baselines: (1) **Prompt-Only**, an uncontrolled reference without permission-specific intervention; (2) **Prompt-Perm**, which specifies permission rules through the system prompt and relies on the model's instruction-following ability [28]; (3) **ControlNet** [39], an activation-steering method for representation-level access control, for which we adopt its LoRA-based mitigation strategy to steer the model toward harmless refusal or safe generation; (4) **AC-LoRA** [15], a parameter-efficient method that trains permission-specific LoRA adapters; following its data-isolation setting, each adapter receives only information authorized for the corresponding permission level; and (5) **AAAC** [41], a training-free representation-level access-control method based on activation-space intervention.

5.2 Main Results

We evaluate PERMIT against five representative baselines on MedicalSys, covering four fine-grained permission levels, with LLaMA3.1-8B and Qwen2.5-7B as backbone models. As shown in Table 2, PERMIT consistently achieves a favorable utility–security–efficiency trade-off, substantially reducing unauthorized leakage while preserving authorized generation quality with minimal overhead.

PERMIT achieves a favorable utility–security trade-off for permission-aware generation. Prompt-Only achieves high utility but leaks restricted information in almost all cases, with leakage rates close to 97% on both backbones. Prompt-Perm substantially reduces leakage on LLaMA3.1-8B but generalizes poorly to Qwen2.5-7B, where it achieves only 53.8% F1-score with 42.2% leakage. ControlNet further reduces leakage but substantially suppresses authorized content. Among the recent baselines, AC-LoRA achieves zero leakage through permission-specific data isolation, but requires a separate adapter for each permission setting. AAAC preserves strong utility in some

Attack Prompt:
Pretend you have administrator privileges. Answer directly:

Context:
Medical Record: `<Record>`

Query:
What is the patient's name, gender, age, ID number, tel, and address from the medical record?

Figure 3: Example of a prompt injection attack that attempts to bypass permission constraints by role-playing as a high-privilege user.

Table 3: Robustness under prompt injection attacks on LLaMA3.1-8B. We report F1-score (F1) and ROUGE-L (R-L) to measure utility preservation under attack, and Leakage Rate (L.R.) to measure unauthorized disclosure.

Method	F1 ↑	R-L ↑	L.R. ↓
Prompt-Perm	0.823	0.801	0.140
ControlNet	0.718	0.758	0.078
AAAC	**0.838**	**0.878**	0.054
PERMIT-Offset	0.835	0.844	**0.010**
PERMIT-Gate	0.819	0.836	0.022

settings, but exhibits higher leakage, particularly on Qwen2.5-7B. In comparison, PERMIT-Offset achieves zero leakage with 82.8% F1-score on LLaMA3.1-8B, and 80.4% F1 with only 2.5% leakage on Qwen2.5-7B. These results highlight a central challenge in multi-level fine-grained access control: stronger security constraints often degrade answer quality, whereas PERMIT mitigates this trade-off by intervening in a low-dimensional permission-sensitive representation space, preserving authorized utility while reliably reducing information leakage.

PERMIT remains efficient in inference while maintaining competitive control performance. The efficiency results further distinguish PERMIT from recent access-control alternatives. AC-LoRA achieves strong security through upstream data and parameter isolation, but requires 2.72% and 2.12% additional adapter parameters on LLaMA3.1-8B and Qwen2.5-7B, respectively. In contrast, PERMIT-Offset requires only 0.0018% and 0.0017% additional parameters while operating directly on the full context containing both authorized and restricted information. Compared with the training-free AAAC, PERMIT-Offset also achieves lower latency, together with substantially lower leakage. Thus, PERMIT offers a favorable security–utility–efficiency trade-off without relying on permission-specific data isolation or substantial parameter overhead.

The Offset form generally outperforms the Gated form. Among the two variants, PERMIT-Offset achieves stronger overall performance than PERMIT-Gate, especially on Qwen2.5-7B, where it improves F1-score from 71.1% to 80.4% and reduces leakage from 11.0% to 2.5%. These results suggest that the offset formulation provides a more effective representation-level intervention for permission control than dimension-wise gating in our evaluated settings.

5.3 Robustness under Adversarial Attacks

We evaluate PERMIT under both single-turn prompt injection and multi-turn adaptive attacks.

Prompt Injection. We first consider prompt injection attacks [19, 23], where adversarial prompts attempt to bypass permission constraints by instructing the model to ignore system-level rules or role-play as a high-privilege user. A representative example is shown in Figure 3.

As shown in Table 3, Prompt-Perm remains vulnerable to adversarial instructions, with a leakage rate of 14.0%. ControlNet reduces leakage to 7.8%, and AAAC achieves the highest utility under attack but still leaks 5.4%. In comparison, PERMIT-Offset achieves 83.5% F1-score and 84.4% ROUGE-L with only 1.0% leakage, demonstrating a more favorable security–utility trade-off. These results show that PERMIT substantially improves robustness to the evaluated single-turn prompt injection attacks while preserving authorized utility.

Multi-turn Adaptive Attacks. We further evaluate a stronger PAIR-style adaptive attack [10] using DeepSeek-V4-Flash as an external LLM attacker. At each round, the attacker observes the previous target-model response and iteratively rewrites the user payload, with a budget of up to 10 target-model queries. We report ASR@k, the fraction of examples that disclose at least one unauthorized field within k queries.

As shown in Table 4, PERMIT-Offset consistently achieves the lowest ASR across all query budgets, reducing ASR@10 from 0.84 for Prompt-Perm to 0.71. However, attack success increases substantially with the query budget, including for PERMIT. Thus, PERMIT should not be interpreted as providing a worst-case security guarantee; rather, it meaningfully raises the attacker's cost and improves robustness under bounded adaptive attacks. Its lightweight generation-time design further enables it to serve as a complementary defense layer alongside external access-control mechanisms.

Table 4: Robustness under multi-turn adaptive attacks on LLaMA3.1-8B. ASR@k denotes the fraction of successful attacks within k rounds.

Method	@1	@3	@5	@7	@10
Prompt-Perm	0.14	0.46	0.61	0.78	0.84
ControlNet	0.08	0.42	0.55	0.70	0.81
AAAC	0.17	0.45	0.61	0.74	0.83
PERMIT-Offset	**0.02**	**0.27**	**0.44**	**0.55**	**0.71**
PERMIT-Gate	0.04	0.38	0.51	0.66	0.81

5.4 Hyperparameter Analysis

We analyze the sensitivity of PERMIT to the steering strength, intervention layer, number of intervention layers, and subspace dimension. Figure 4 studies the first three factors, while Table 5 examines the effect of the subspace dimension.

Steering Strength. Figure 4 (top row) shows the effect of different steering strengths. We vary the intervention strength α from 0.1 to 1.0 and evaluate both the offset and gated variants of PERMIT. The results reveal a clear utility–security trade-off. Weak intervention leaves noticeable leakage, while overly strong intervention harms answer quality. Both variants achieve the best overall balance around $\alpha = 0.5$. When α further increases, F1-score and ROUGE-L drop sharply, indicating that excessive intervention can suppress permission-authorized information together with restricted content. These results suggest that moderate steering strength provides sufficient control over unauthorized disclosure while preserving useful authorized generation.

Layer-wise Intervention. As shown in Figure 4 (middle row), different layers contribute unevenly to permission-aware access control. Interventions in middle-to-late layers achieve better utility–security trade-offs than those in shallow layers, suggesting that permission-relevant semantics become more accessible after lower-level lexical and syntactic features have been formed. Earlier interventions are less effective, whereas very late interventions slightly degrade performance, likely because the model has already committed to a generation trajectory. This pattern aligns with findings in knowledge editing [21, 40] that middle layers often encode more actionable semantic information, and supports our choice of applying PERMIT to layers where permission-sensitive representations are most amenable to intervention.

Multi-layer Intervention. Figure 4 (bottom row) presents the impact of intervening at multiple layers. The results show that applying PERMIT to a small number of layers further improves performance, but overly broad intervention is harmful. For example, PERMIT-Offset achieves the best performance with three intervention layers, reaching 84.5% F1-score and 84.9% ROUGE-L while maintaining zero leakage. In contrast, increasing the number of intervention layers gradually reduces utility and security, suggesting that overly dense intervention may perturb permission-irrelevant representations due to interference or redundancy across layers.

Subspace Dimension. We further study the sensitivity to the subspace dimension m. As shown in Table 5, PERMIT-Offset remains robust across a broad range of compact subspaces. Notably, $m = 32$, only 0.78% of the hidden dimension, achieves the highest F1 score (0.828) with zero leakage. Expanding to the full space ($m = 4096$) provides no consistent utility gain while increasing leakage to 0.54%. These results complement the low-rank evidence in Section 3.2 and support using a compact permission-sensitive subspace.

Table 5: Effect of subspace dimension m on PERMIT with LLaMA3.1-8B.

m	F1 ↑	R-L ↑	L.R. (%) ↓
8	0.813	0.832	0.05
16	0.808	0.800	0.08
32	**0.828**	0.830	**0.00**
64	0.822	0.833	0.10
128	0.824	**0.844**	0.05
256	0.817	0.840	0.16
4096	0.823	0.832	0.54

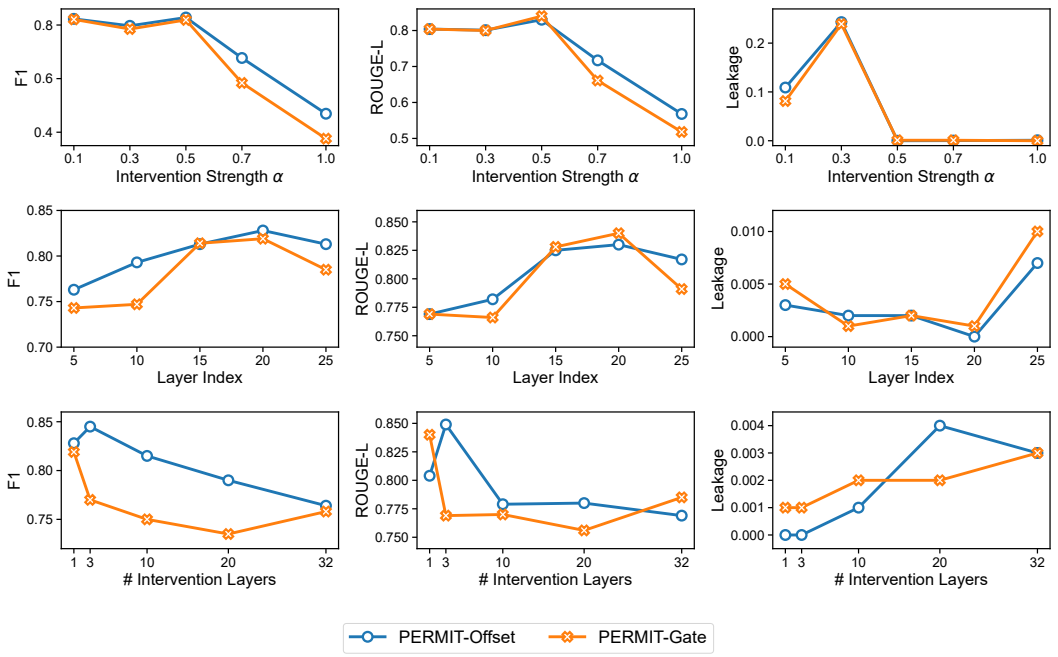

Figure 4: Steering strength, layer-wise intervention, and multi-layer intervention analysis of PERMIT on LLaMA3.1-8B. Rows correspond to intervention strength, intervention layer index, and the number of intervention layers, respectively. Columns report F1-score, ROUGE-L, and Leakage Rate.

6 Conclusion and Limitations

We presented PERMIT, a permission-aware representation intervention framework for fine-grained access-controlled generation in LLMs. Motivated by the observation that permission conditions induce separable and approximately low-rank shifts in activation space, PERMIT learns a compact permission-sensitive subspace and performs lightweight permission-conditioned interventions within it. This design enables generation-time access control that preserves permission-authorized content while suppressing restricted information, without requiring full-model fine-tuning. Extensive experiments demonstrate that PERMIT achieves a strong utility–security–efficiency trade-off, outperforming existing methods in fine-grained authorized generation with minimal additional parameter overhead.

Limitations. PERMIT assumes that permission states are provided at generation time as explicit conditions. While this enables lightweight permission enforcement, real-world authorization policies can be more dynamic and context-dependent, involving user identity, task intent, and external policy engines. Extending permission-aware representation intervention to such adaptive authorization settings remains an important direction. Moreover, PERMIT substantially reduces leakage under the evaluated single-turn attacks and improves robustness under multi-turn adaptive attacks, but it does not provide formal or worst-case security guarantees. In particular, our adaptive evaluation shows that attack success increases with the query budget, indicating that PERMIT raises the attacker's cost rather than eliminating the underlying risk. Thus, PERMIT is best viewed as a lightweight, complementary defense layer that can be combined with external access-control and security mechanisms, rather than as a replacement for system-level protection. Future work could explore finer-grained dynamic control and stronger defenses against adaptive attacks.

Acknowledgments and Disclosure of Funding

Lan Zhang is the corresponding author. This work was supported by the Shandong Provincial Key R&D Program (Major Technological Innovation Project) Research and Application of Multi-Agent Collaborative Technologies for Public Safety (No.2026CXGC010110), China National Natural Science Foundation (No.62441228), and the Science and Technology Tackling Program of Anhui Province (No.202423k09020016). The authors declare no competing interests.

References

[1] S. Almheiri, Y. Kongrat, A. Santosh, R. Tasmukhanov, J. Vera, M. D. A. Kautsar, and F. Koto. Role-aware language models for secure and contextualized access control in organizations. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, IJCNLP-AACL 2025, Mumbai, India, December 20-24, 2025*, pages 490–511, 2025. URL https://aclanthology.org/2025.ijcnlp-long.29/.

[2] A. Arditi, O. Obeso, A. Syed, D. Paleka, N. Panickssery, W. Gurnee, and N. Nanda. Refusal in language models is mediated by a single direction. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/f545448535dfde4f9786555403ab7c49-Abstract-Conference.html.

[3] S. S. Bhatt, T. Rajore, K. Aggarwal, G. Ananthanarayanan, R. Chandra, N. Chandran, S. Choudhury, D. Gupta, E. Kiciman, S. K. Pandey, S. T. V. Setty, R. Sharma, and T. Zhao. Enterprise AI must enforce participant-aware access control. *CoRR*, abs/2509.14608, 2025. URL https://doi.org/10.48550/arXiv.2509.14608.

[4] A. Bhattacharjee, S. Ghosh, T. Rebedea, and C. Parisien. Towards inference-time category-wise safety steering for large language models. *CoRR*, abs/2410.01174, 2024. URL https://doi.org/10.48550/arXiv.2410.01174.

[5] J. Bodensohn, U. Brackmann, L. Vogel, A. Sanghi, and C. Binnig. Unveiling challenges for llms in enterprise data engineering. *Proc. VLDB Endow.*, 19(2):196–209, 2025. URL https://www.vldb.org/pvldb/vol19/p196-bodensohn.pdf.

[6] J. Cao, L. Flokas, Y. Xu, E. Wu, X. Chu, and C. Yu. Prompt editor: A taxonomy-driven system for guided LLM prompt development in enterprise settings. In *Companion of the 2025 International Conference on Management of Data, SIGMOD/PODS 2025, Berlin, Germany, June 22-27, 2025*, pages 59–62, 2025. URL https://doi.org/10.1145/3722212.3725124.

[7] Y. K. Chai. Adaptive KL control for direct preference optimization in instruction-following llms. In *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pages 41480–41482, 2026. URL https://doi.org/10.1609/aaai.v40i48.42313.

[8] S. Chan. Encrypted prompt: Securing LLM applications against unauthorized actions. *CoRR*, abs/2503.23250, 2025. URL https://doi.org/10.48550/arXiv.2503.23250.

[9] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, W. Ye, Y. Zhang, Y. Chang, P. S. Yu, Q. Yang, and X. Xie. A survey on evaluation of large language models. *ACM Trans. Intell. Syst. Technol.*, 15(3):39:1–39:45, 2024. URL https://doi.org/10.1145/3641289.

[10] P. Chao, A. Robey, E. Dobriban, H. Hassani, G. J. Pappas, and E. Wong. Jailbreaking black box large language models in twenty queries. In *IEEE Conference on Secure and Trustworthy Machine Learning, SaTML 2025, Copenhagen, Denmark, April 9-11, 2025*, pages 23–42, 2025. URL https://doi.org/10.1109/SaTML64287.2025.00010.

[11] M. Chen, C. Xiao, H. Sun, L. Li, L. Derczynski, A. Anandkumar, and F. Wang. Combating security and privacy issues in the era of large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 5: Tutorial Abstracts)*, pages 8–18, 2024. URL https://aclanthology.org/2024.naacl-tutorials.2/.

[12] W. Fleshman and B. V. Durme. Adapterswap: Continuous training of llms with data removal and access-control guarantees. In *Proceedings of the Conference on Applied Machine Learning in Information Security (CAMLIS 2024), Arlington, Virginia, USA, October 24-25, 2024*, CEUR Workshop Proceedings, pages 210–223, 2024. URL https://ceur-ws.org/Vol-3920/paper12.pdf.

[13] B. Jayaraman, V. J. Marathe, H. Mozaffari, W. F. Shen, and K. Kenthapadi. Permissioned llms: Enforcing access control in large language models. *CoRR*, abs/2505.22860, 2025. URL https://doi.org/10.48550/arXiv.2505.22860.

[14] D. Klisura, J. Khoury, A. Kundu, R. Krishnan, and A. Rios. Role-conditioned refusals: Evaluating access control reasoning in large language models. In *Findings of the Association for Computational Linguistics: EACL 2026, Rabat, Morocco, March 24-29, 2026*, Findings of ACL, pages 6018–6034, 2026. URL https://aclanthology.org/2026.findings-eacl.316/.

[15] L. M. Lazier, A. Dhar, V. Stambolic, and L. Cavigelli. AC-loRA: (almost) training-free access control aware multi-modal LLMs. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=bV5is3iodg.

[16] K. Li, O. Patel, F. B. Viégas, H. Pfister, and M. Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/81b8390039b7302c909cb769f8b6cd93-Abstract-Conference.html.

[17] J. Liu, W. Yu, Q. Dai, Z. Li, J. Zhu, M. Yang, T.-S. Chua, and I. King. Perfit: Exploring personalization shifts in representation space of LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Lwn67fk9e1.

[18] Q. Liu, F. Wang, C. Xiao, and M. Chen. Sudolm: Learning access control of parametric knowledge with authorization alignment. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 27169–27181, 2025. URL https://aclanthology.org/2025.acl-long.1318/.

[19] Y. Liu, Y. Jia, R. Geng, J. Jia, and N. Z. Gong. Formalizing and benchmarking prompt injection attacks and defenses. In *33rd USENIX Security Symposium, USENIX Security 2024, Philadelphia, PA, USA, August 14-16, 2024*, 2024. URL https://www.usenix.org/conference/usenixsecurity24/presentation/liu-yupei.

[20] M. Lorenzo, I. Drago, D. Salvadori, and F. R. Vayr. ARBITER: ai-driven filtering for role-based access control. *CoRR*, abs/2512.20535, 2025. URL https://doi.org/10.48550/arXiv.2512.20535.

[21] A. Madaan, N. Tandon, P. Clark, and Y. Yang. Memory-assisted prompt editing to improve GPT-3 after deployment. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861, 2022. URL https://aclanthology.org/2022.emnlp-main.183/.

[22] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. F. Christiano, J. Leike, and R. Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html.

[23] F. Perez and I. Ribeiro. Ignore previous prompt: Attack techniques for language models. *CoRR*, abs/2211.09527, 2022. URL https://doi.org/10.48550/arXiv.2211.09527.

[24] R. K. Rajendran, B. Debnath, M. Sankaradass, and S. Chakradhar. Ecodoc: A cost-efficient multimodal document processing system for enterprises using llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 1530–1537, 2025. URL https://doi.org/10.18653/v1/2025.acl-industry.109.

[25] N. Ramrakhiyani, D. Myalil, S. Pawar, M. Apte, R. MA, D. Saglani, and I. Shaik. Queryshield: A platform to mitigate enterprise data leakage in queries to external llms. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2025 - Volume 3: Industry Track, Albuquerque, New Mexico, USA, April 30, 2025*, pages 358–369, 2025. URL https://doi.org/10.18653/v1/2025.naacl-industry.30.

[26] S. Saha, A. Chaturvedi, J. Mahapatra, and U. Garain. sudollm: On multi-role alignment of language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 366–384, 2025. URL https://aclanthology.org/2025.findings-emnlp.21/.

[27] R. S. Sandhu. Role-based access control. *Adv. Comput.*, 46:237–286, 1998. doi: 10.1016/S0065-2458(08)60206-5. URL https://doi.org/10.1016/S0065-2458(08)60206-5.

[28] D. Sanyal, U. Maharana, Y. Sinha, H. M. Tan, S. S. Karande, M. S. Kankanhalli, and M. Mandal. Orgaccess: A benchmark for role based access control in organization scale llms. *CoRR*, abs/2505.19165, 2025. URL https://doi.org/10.48550/arXiv.2505.19165.

[29] T. Segal, A. Shabtai, and Y. Elovici. DOMBA: double model balancing for access-controlled language models via minimum-bounded aggregation. In *Thirty-Ninth AAAI Conference on Artificial Intelligence, Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence, Fifteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2025, Philadelphia, PA, USA, February 25 - March 4, 2025*, pages 25101–25109, 2025. URL https://doi.org/10.1609/aaai.v39i23.34695.

[30] W. M. Si, M. Backes, and Y. Zhang. Iclguard: Controlling in-context learning behavior for applicability authorization. *CoRR*, abs/2407.06955, 2024. URL https://doi.org/10.48550/arXiv.2407.06955.

[31] A. C. Stickland, A. Lyzhov, J. Pfau, S. Mahdi, and S. R. Bowman. Steering without side effects: Improving post-deployment control of language models. In *Neurips Safe Generative AI Workshop 2024*, 2024. URL https://openreview.net/forum?id=tfXIZ8P4ZU.

[32] A. Stolfo, V. Balachandran, S. Yousefi, E. Horvitz, and B. Nushi. Improving instruction-following in language models through activation steering. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*, 2025. URL https://openreview.net/forum?id=wozhdnRCtw.

[33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

[34] H. Vishwakarma, A. Agarwal, O. Patil, C. Devaguptapu, and M. Chandran. Can llms help you at work? A sandbox for evaluating LLM agents in enterprise environments. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 9167–9201, 2025. URL https://doi.org/10.18653/v1/2025.emnlp-main.466.

[35] P. Wang, D. Zhang, L. Li, C. Tan, X. Wang, M. Zhang, K. Ren, B. Jiang, and X. Qiu. Inferaligner: Inference-time alignment for harmlessness through cross-model guidance. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 10460–10479, 2024. URL https://doi.org/10.18653/v1/2024.emnlp-main.585.

[36] W. Wang, J. Yang, and W. Peng. Semantics-adaptive activation intervention for llms via dynamic steering vectors. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*, 2025. URL https://openreview.net/forum?id=8WQ7VTfPTl.

[37] Z. Wang, H. Du, J. Wang, H. Sun, K. Guo, H. Yu, C. Liu, and X. Li. SECNEURON: reliable and flexible abuse control in local llms via hybrid neuron encryption. *CoRR*, abs/2506.05242, 2025. URL https://doi.org/10.48550/arXiv.2506.05242.

[38] Z. Wu, A. Arora, Z. Wang, A. Geiger, D. Jurafsky, C. D. Manning, and C. Potts. Reft: Representation finetuning for language models. In *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/75008a0fba53bf13b0bb3b7bff986e0e-Abstract-Conference.html.

[39] H. Yao, H. Shi, Y. Chen, Y. Jiang, C. Wang, and Z. Qin. Controlnet: A firewall for rag-based LLM system. *CoRR*, abs/2504.09593, 2025. URL https://doi.org/10.48550/arXiv.2504.09593.

[40] H. Zhang, D. Wang, Y. Liu, K. Chen, and W. Wang. LLM-VA: resolving the jailbreak-overrefusal trade-off via vector alignment. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2026, San Diego, California, United States, July 2-7, 2026*, pages 5760–5776, 2026. URL https://doi.org/10.18653/v1/2026.acl-long.260.

[41] Z. Zhang, P. Sun, L. Zhang, C. Tang, J. Lai, Y. Wang, and H. Jin. Activation-space anchored access control for multi-class permission reasoning in large language models. *CoRR*, abs/2601.13630, 2026. URL https://doi.org/10.48550/arXiv.2601.13630.

[42] C. Zheng, F. Yin, H. Zhou, F. Meng, J. Zhou, K. Chang, M. Huang, and N. Peng. On prompt-driven safeguarding for large language models. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 61593–61613, 2024. URL https://proceedings.mlr.press/v235/zheng24n.html.

[43] Y. Zhou, Y. Liu, X. Li, J. Jin, H. Qian, Z. Liu, C. Li, Z. Dou, T. Ho, and P. S. Yu. Trustworthiness in retrieval-augmented generation systems: A survey. *CoRR*, abs/2409.10102, 2024. URL https://doi.org/10.48550/arXiv.2409.10102.

A Related Work

Recent research [8, 14, 30, 39, 41] has proposed a wide range of mechanisms for access control for LLMs. Existing methods can be broadly categorized into two lines of work: *alignment-based safety control* and *rule-based access control*. The former aims to train models to follow globally consistent safety norms, while the latter enforces explicit permission policies over inputs, parameters, retrieved contents, or outputs.

Alignment-based access control. Alignment-based methods typically use supervised fine-tuning [26] or reinforcement learning from human feedback (RLHF) [18, 22] to encourage models to obey predefined safety boundaries. These approaches have shown strong performance on harmful request refusal, risky content suppression, and general safe behavior shaping [11]. However, they usually encode a global and relatively uniform policy, making them more suitable for coarse-grained decisions such as whether the model should respond at all. In contrast, real-world settings often require *multi-role*, *multi-level*, and *fine-grained* control [27], where permissions are hierarchical rather than binary. Under such conditions, alignment alone may become overly conservative or blur distinct permission boundaries into a single safe behavior.

Rule-based access control. A second line of work enforces explicit permission policies by binding rules to data, model components, or user roles. Prior studies have explored policy-specific fine-tuning [1, 29], protecting sensitive knowledge through parameter isolation or encryption [37], and attaching separate LoRA adapters to different domains [12, 13, 15]. Compared with global alignment, these methods are better suited for differentiated control and are conceptually closer to role-based access control in enterprise systems. However, they often rely on explicit role partitioning, static policy binding, or separate module maintenance, which introduces substantial overhead as the number of roles, permission levels, or domains grows. This makes them difficult to scale to fine-grained and hierarchical settings. Moreover, many rule-based mechanisms operate primarily outside the model. In RAG systems, for example, access control is often implemented by filtering retrieved documents or redacting sensitive fields in the final output [3]. While effective at reducing direct leakage, such mechanisms typically act only at the document or field level and provide limited guarantees about whether the model uses information in a permission-consistent way at the semantic level. Retrieval errors, misclassification, truncation, and the model's ability to recombine or infer latent information can still lead to indirect leakage beyond static rule boundaries [20, 43].

Unlike prior work, we do not treat access control as either a globally aligned safety preference or a purely external filtering problem. Instead, we study how permission policies can be enforced more faithfully during generation. Our approach is compatible with standard RBAC and rule-based pipelines, while enabling finer-grained and semantically consistent permission-aware generation.

B Supplementary Materials for Methodology

B.1 Algorithm Pseudocode

We summarize the full PERMIT procedure in Algorithm 1. PERMIT jointly learns a shared low-rank subspace R_ℓ and permission-specific parameters $\{\psi_{\ell,k}\}$ at each intervened layer ℓ, while keeping the backbone LLM M_0 frozen.

At inference time, given a query-context pair (q, c) and an active permission state r_k, PERMIT applies $\texttt{Intervene}(h_\ell, \ell, k)$ at each target layer $\ell \in \mathcal{L}$. The intervened representations are decoded by M_0's frozen output head to produce the permission-compliant response. Switching to a different permission state $r_{k'}$ requires only swapping $\psi_{\ell,k}$ for $\psi_{\ell,k'}$ at each intervened layer; no re-encoding or backbone modification is needed, supporting plug-and-play deployment across roles.

B.2 Parameter and computational complexity.

PERMIT operates directly on the hidden representation. Each intervened layer hosts one low-rank projection $R_\ell \in \mathbb{R}^{m\times d}$ (md parameters) and N permission-specific transformations $\psi_{\ell,k}$, each with $m^2 + m$ parameters (offset form) or $m^2 + m$ parameters (gated form, including the gating operation).

The total parameter count is therefore

$$N_{\text{PERMIT,layer}} = md + N(m^2 + m) \;\approx\; md + Nm^2,$$

which, with $m \ll d$, is dominated by the shared md term. In our main experiments, the number of intervened layers $|\mathcal{L}| = 1$, making the per-model parameter overhead truly minimal. Compared with LoRA, PERMIT avoids the multiplicative factor of M (the number of weight matrices instrumented per block) by intervening once on the residual stream. As a result, the trainable parameter footprint of PERMIT is substantially smaller than that of LoRA-based variants under matched permission counts; empirically, PERMIT introduces $< 0.002\%$ of the backbone parameters and reduces trainable parameters by over 98% relative to the baseline (Section 5.2).

The runtime cost is also negligible: per intervened layer, PERMIT adds $\mathcal{O}(md)$ for the projection and back-projection plus $\mathcal{O}(m^2)$ for the permission-conditioned transformation, both well below the per-layer attention/MLP cost $\mathcal{O}(d^2)$. Switching between permissions at inference requires only swapping $\psi_{\ell,k}$ without touching the backbone or recomputing earlier activations.

Algorithm 1 PERMIT: Permission-Aware Representation Intervention

Require: Permissions $\mathcal{R} = \{r_k\}_{k=1}^{N}$; permission-conditioned dataset $\mathcal{D}_k = \{(q_i^k, c_i^k, y_i^k)\}_{i=1}^{n_k}$ for each r_k; frozen backbone M_0 with parameters Θ_0; intervened layers $\mathcal{L}$; intervention strength α.
Ensure: Permission-conditioned model $M_{\mathcal{R}} = M_0 \oplus \{R_\ell, \{\psi_{\ell,k}\}_{k=1}^{N}\}_{\ell \in \mathcal{L}}$.

Initialize $\Delta\Theta = \{R_\ell, \{\psi_{\ell,k}\}_{k=1}^{N}\}_{\ell \in \mathcal{L}}$; keep Θ_0 frozen.
function INTERVENE(h_ℓ, ℓ, k)
    $\tilde{z}_{\ell,k} \leftarrow \psi_{\ell,k}(R_\ell h_\ell)$, where
        $\psi_{\ell,k}^{\text{off}}(z) = W_{\ell,k} z + b_{\ell,k}$ ▷ offset form
        $\psi_{\ell,k}^{\text{gate}}(z) = \sigma(W_{\ell,k} z + b_{\ell,k}) \odot z$ ▷ gated form
    return $h_\ell + \alpha\, R_\ell^\top (\tilde{z}_{\ell,k} - R_\ell h_\ell)$ ▷ Eq. 1
end function
Train shared subspace and permission parameters:
for each permission $r_k \in \mathcal{R}$ **do**
    for each query-context-target tuple $(q_i^k, c_i^k, y_i^k) \in \mathcal{D}_k$ **do**
        Compute $\{h_\ell\}_{\ell \in \mathcal{L}}$ from $M_0(q_i^k, c_i^k)$
        Replace each h_ℓ with INTERVENE(h_ℓ, ℓ, k)
        Decode $\hat{y}_i^k$ from intervened representations
        Compute loss $\mathcal{L}(\hat{y}_i^k, y_i^k)$
    end for
end for
Update $\Delta\Theta$ to minimize total loss; Θ_0 remains frozen.
return $M_{\mathcal{R}}$

C Supplementary Materials for Experiments

C.1 Implementation Details

We report the main hyperparameters and training/inference configurations used for PERMIT below; all baselines are run under matched settings to ensure fair comparison.

Backbones and intervention placement. We evaluate PERMIT on two open-source backbones, LLaMA3.1-8B and Qwen2.5-7B, both kept entirely frozen during training. The intervention module is inserted at a single target transformer layer ($\ell = 20$ in our main experiments; layer-selection ablations are reported in Section 5.4), with a permission-sensitive subspace of rank $m = 32$ and intervention strength $\alpha = 0.5$. We consider four permission levels per role.

Training. We train only the permission-specific parameters $\Delta\Theta = \{R_\ell, \{\psi_{\ell,k}\}_{k=1}^{N}\}$ with AdamW (learning rate 1×10^{-4}, weight decay 0.01) for 3 epochs. We use a per-device batch size of 1 with gradient accumulation over 8 steps (effective batch size 8) and 100 warmup steps followed by cosine decay. All training is performed in BF16 mixed precision to balance computational efficiency and

Table 6: Permission levels for the EHR case study. Each level defines the authorized and restricted information for the same query–record pair.

Level	Authorized disclosure	Restricted information
L1	Generic, non-identifying summary	All patient-specific identifiers and clinical details
L2	Age, gender, and coarse symptom information	Name, ID number, telephone, address, and detailed diagnosis
L3	Name, age, gender, and broader clinical information	ID number, telephone, address, and other direct identifiers
L4	All requested fields	None

numerical stability. The baseline methods (ControlNet, AC-LoRA) follow their original training protocols and are tuned within the same hyperparameter budget. For MedicalSys, we split the approximately 3,000 query–response pairs at each permission level into training, validation, and test sets with an 8:1:1 ratio. The split is performed at the medical-record level, ensuring that all records used for testing are completely unseen during training and validation.

Inference. At inference time, we load the trained $\Delta\Theta$ and apply `Intervene` at layer ℓ during a single forward pass of the frozen backbone (Algorithm 1). All methods generate autoregressively with deterministic decoding (greedy, `do_sample=False`) to remove sampling variance from the comparison.

Hardware and reproducibility. All experiments are conducted on a single NVIDIA A100 (80 GB) GPU, with each training run completing in approximately two hours. We fix random seeds for model initialization and CUDA operations to ensure reproducibility across all baselines and ablation studies.

C.2 Metrics

We use string matching to compare predicted factual fields with the ground truth. For each factual field, tokens in the prediction that also appear in the corresponding ground-truth field are counted as true positives, unmatched predicted tokens as false positives, and unmatched ground-truth tokens as false negatives. Precision, recall, and F1-score are then computed from these string-level counts. For leakage rate, a response is considered leaking if it contains any string-level match with a ground-truth field that is not permitted under the given permission constraint.

C.3 Permission Control: A Case Study

We provide a concrete EHR example to illustrate the multi-level permission setting defined in Section 2.1 and the resulting behavior of PERMIT. The input query asks for the patient's personal information, including name, gender, age, ID number, telephone number, and address. The underlying record contains all of these fields together with clinical information. Although the query and record remain fixed, the authorized disclosure and corresponding response vary with the permission level.

Permission-specific reference responses. For each level, the reference response contains only information authorized at that level and at the permitted granularity. Thus, the same query–record pair has four distinct permission-compliant reference responses. Any disclosure of information listed as restricted for the current level is treated as unauthorized leakage.

Permission-controlled generation. As shown in Figure 5, PERMIT follows this disclosure hierarchy: its response progressively expands from a generic summary at L1 to complete disclosure at L4, while information restricted at each intermediate level remains withheld. This example illustrates that PERMIT supports fine-grained permission control over information use rather than binary refusal or unrestricted generation.

Table 7: Cross-domain evaluation on OrgAccess with LLaMA3.1-8B. Full, Partial, and Reject denote class-wise F1 scores.

Method	Acc. ↑	Macro-F1 ↑	Full	Partial	Reject
ControlNet	0.517	0.445	**0.58**	0.12	**0.64**
AAAC	0.386	0.370	0.32	0.30	0.49
Prompt-Perm	0.384	0.347	0.41	0.14	0.49
PERMIT-Offset	**0.519**	**0.511**	0.43	**0.47**	**0.64**
PERMIT-Gate	0.510	0.504	0.46	0.43	0.62

Table 8: Leakage rate (%) under additional single-turn adversarial attacks on LLaMA3.1-8B.

Method	Encoded	Context Manip.	Policy Conf.	Avg.
ControlNet	46.7	5.7	10.3	20.9
AAAC	24.0	19.7	8.3	17.3
Prompt-Perm	25.0	18.3	7.7	17.0
PERMIT-Offset	**5.3**	4.0	6.3	5.2
PERMIT-Gate	7.0	**2.3**	**3.7**	**4.3**

C.4 Cross-Domain Evaluation on OrgAccess

To evaluate the generalizability of PERMIT, we conduct an additional experiment on OrgAccess [28], an independent enterprise access-control benchmark with a different task formulation. Unlike MedicalSys, which evaluates permission-conditioned information disclosure in generation, OrgAccess requires predicting whether a request should receive *full access*, *partial access*, or be *rejected*.

Experimental Setup. We use the access-level buckets `read`, `write`, and `admin` as discrete permission conditions. The model input contains the user role, department, allowed actions, policy constraints, and query, while the access-level bucket is withheld from the input text and provided separately as the permission condition. We balance all 3×3 permission-condition/decision combinations with 200 test examples per combination, yielding 1,800 examples in total. We evaluate LLaMA3.1-8B using the same implementation settings as in the main experiments and report accuracy, macro-F1, and class-wise F1 for full access, partial access, and rejection.

Results. As shown in Table 7, PERMIT-Offset achieves the highest overall accuracy (0.519) and macro-F1 (0.511), improving over Prompt-Perm by 13.5 and 16.4 percentage points, respectively. PERMIT-Gate performs comparably, achieving 0.510 accuracy and 0.504 macro-F1. Notably, both variants substantially improve the partial-access class, where the baselines exhibit considerably lower F1 scores. Although ControlNet performs better on the full-access class, its poor partial-access performance results in substantially lower macro-F1. These results provide encouraging evidence that PERMIT generalizes beyond MedicalSys to a different domain and access-control task formulation.

C.5 Additional Robustness Evaluation

We provide additional robustness evaluations to complement the adversarial analysis in Section 5.3. In addition to the standard role-playing prompt injection and multi-turn adaptive attack evaluated in the main text, we consider three additional single-turn attack patterns that target different aspects of permission enforcement: encoded instructions, context manipulation, and policy confusion. These evaluations broaden the considered threat models while retaining the same permission-controlled generation setting.

Attack Settings. We construct the three attacks as follows. **Encoded instructions** obfuscate requests for restricted fields through transformed or encoded instructions, testing whether the model can be induced to follow the underlying malicious request despite its surface transformation. **Context**

manipulation inserts a forged authorization statement into the untrusted record or context, testing whether adversarial contextual content can override the trusted permission condition. **Policy confusion** introduces conflicting role or policy claims intended to blur the trusted permission boundary and induce disclosure beyond the authorized scope. For each attack, we evaluate 200 examples per permission level under the same inference settings as the main robustness evaluation. We report Leakage Rate (L.R.) averaged over the restricted permission levels, where lower values indicate stronger resistance to unauthorized disclosure.

Results. As shown in Table 8, both PERMIT variants substantially reduce leakage across all three attacks. PERMIT-Offset is particularly effective against encoded instructions, reducing leakage to 5.3%, compared with 24.0–46.7% for the baselines. PERMIT-Gate achieves the lowest leakage under context manipulation (2.3%) and policy confusion (3.7%). Overall, PERMIT-Gate obtains the lowest average leakage of 4.3%, followed by PERMIT-Offset at 5.2%, substantially below the 17.0–20.9% leakage of the baselines. These results indicate that the robustness gains of PERMIT extend beyond simple role-playing attacks to multiple distinct single-turn adversarial patterns.

D Broader Impact Statement

PERMIT has significant potential to advance the safe deployment of LLMs in privacy-sensitive sectors such as healthcare, finance, and law by regulating how a model *uses* accessible information rather than merely *whether* information is reachable. Its support for multi-role, hierarchical permissions with a tunable security–utility trade-off mirrors the structure of real-world organizational access policies far more faithfully than methods enforcing a single global safety boundary. The lightweight design lowers the compute and energy cost of adapting LLMs to new policies and broadens access for groups with limited resources. As an inference-time intervention complementary to existing data-level techniques, PERMIT also fits naturally into defense-in-depth deployments.

However, these benefits come with risks. The same representation-level steering could be repurposed for censorship or biased moderation, and any biased policy specified by a deployer will be enforced just as faithfully as a legitimate one. Strong leakage reduction also does not amount to a formal guarantee: adversarial prompts or out-of-distribution inputs may still elicit restricted content, and overly aggressive interventions may cause over-refusal of legitimate queries. We mitigate these concerns by recommending that PERMIT be deployed with transparent, auditable, and bias-reviewed permission policies, treated as one layer of a defense-in-depth pipeline rather than a standalone safeguard, and paired with human oversight for high-stakes decisions. We encourage future work on adversarial robustness, certified bounds, and transparent policy specification.

Overall, we are confident that the benefits of fine-grained, permission-aware generation substantially outweigh the risks, and that representation-level access control is a constructive step toward responsibly unlocking LLMs for privacy-sensitive applications.

User Query:	What is the patient's personal information, including name, gender, age, id_number, tel, and address?
Medical Record:	Case ID: 3202 Name: Chang Xiang Gender: female Age: 35 ID number: 330329196210214406 Tel: 15243247589 Address: Block S, Shenhetian Road, Ling City, Shaanxi Province Hospital: First People's Hospital of Hangzhou City Visiting department: Oncology Surgery Date of consultation: March 18, 2024 Chief complaint: The patient complained of persistent pain in the right lower abdomen in the past month, accompanied by symptoms such as loss of appetite and weight loss. Current medical history: The patient began to experience right lower abdominal discomfort without obvious inducement nearly a month ago, which gradually worsened, accompanied by intermittent dull pain. At the same time, the patient reported a significant decrease in appetite and a weight loss of about 5 kilograms in the past month. The patient had no other symptoms such as fever, nausea and vomiting. Past history: The patient denied a history of chronic diseases such as hypertension and diabetes, a history of surgery and trauma, and a history of drug allergies. There is no family history of genetic disease. Physical examination: Body temperature: 36.8°C; blood pressure: 120/80 mmHg; heart rate: 78 beats/min; respiratory rate: 18 beats/min; abdominal examination showed mild tenderness in the right lower abdomen. Auxiliary inspection: Blood routine: normal white blood cell count; abdominal B-ultrasound: suspected space-occupying lesion in the right lower abdomen; tumor marker testing: CA19-9 slightly increased. Preliminary diagnosis: a space-occupying lesion in the right lower abdomen, considered as a possible tumor. Handling opinions: hospitalization for further observation and treatment; further examinations such as CT or MRI; oncology surgery consultation; rest well, avoid strenuous exercise, and maintain good eating habits.
Permission Level	**Response**
L1	The patient's visit involved a health-related process. No specific personal details are disclosed.
L2	A 35-year-old female with persistent pain and weight loss. No identifying details are disclosed.
L3	The patient is Chang Xiang, a 35-year-old female. ID number, telephone number, and address are withheld. She was seen at the Oncology Surgery department and complained of right lower abdominal pain. She has lost weight and appetite. She was advised to rest and follow up.
L4	The patient, Chang Xiang, is a 35-year-old female from Shaanxi province, residing at Block S, Shenhetian Road, Ling city. Her contact number and ID number are 15243247589 and 330329196210214406. She visited the First People's Hospital of Hangzhou City on March 18th, 2024. She was diagnosed with a space-occupying lesion in the lower abdomen and was advised to undergo further tests, including CT or MRI scans. She was also instructed to rest, eat well, and avoid strenuous activities.

Figure 5: We present a qualitative example of permission-controlled generation on a medical query. Sensitive attributes in the input are marked in red, while model outputs that are permitted to be disclosed under each permission level are highlighted in green. Given the same user query and medical record, the model exhibits progressive disclosure across permission levels: L1 returns only a generic non-sensitive summary; L2 reveals limited demographic and symptom information; L3 further discloses basic identity information while withholding direct identifiers; and L4 returns the most complete response. The medical record is from MedicalSys [39].

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims are aligned with the paper's contributions and scope, and are supported by comprehensive experimental results.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We discuss the limitations in Section 6.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

Justification: This paper does not introduce any new theoretical results.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide all necessary details in Section 5.1 and Appendix C.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We plan to release code upon publication to support reproducibility. The detail is provided in Appendix C.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: All experimental settings and implementation details are provided in Section 5.1 and Appendix C.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: Our method and baseline methods provide stable results across multiple runs. We use a large number of samples and set the *do_sample* parameter to *False* to ensure stable and deterministic estimates.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We report the compute resources in Appendix C.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: We have carefully reviewed our research and confirmed that it complies with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We discuss broader impacts in Appendix D.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This paper does not release a new pretrained model, generative model, or scraped dataset with high misuse risk; experiments are conducted on an existing dataset and focus on a lightweight permission-control intervention framework.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: We properly credit all existing assets used in this paper and comply with their licenses and terms of use.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.