

000 SALMAN: STABILITY ANALYSIS OF LANGUAGE 001 MODELS THROUGH THE MAPS BETWEEN GRAPH- 002 BASED MANIFOLDS

003
004
005
006
007 **Anonymous authors**
008 Paper under double-blind review
009
010
011
012

ABSTRACT

013 Recent strides in Pretrained Transformer-based language models have propelled
014 state-of-the-art performance in numerous NLP tasks. Yet, as these models grow
015 in size and deployment, their robustness under input perturbations becomes an
016 increasingly urgent question. Existing robustness methods often diverge between
017 small-parameter and large-scale models (LLMs), and they typically rely on labor-
018 intensive, sample-specific adversarial designs. In this paper, we propose a unified,
019 local (sample-level) robustness framework (SALMAN) that evaluates model sta-
020 bility without modifying internal parameters or resorting to complex perturbation
021 heuristics. Central to our approach is a novel Distance Mapping Distortion (DMD)
022 measure, which ranks each sample’s susceptibility by comparing input-to-output
023 distance mappings in a near-linear complexity manner. By demonstrating signif-
024 icant gains in attack efficiency and robust training, we position our framework as
025 a practical, model-agnostic tool for advancing the reliability of transformer-based
026 NLP systems.
027

1 INTRODUCTION

028 Recently, breakthroughs in pretrained Transformer-based language models have revolutionized the
029 field of Natural Language Processing (NLP). These models have enhanced performance across a
030 wide range of downstream NLP tasks, including text classification Sun et al. (2019), summariza-
031 tion El-Kassas et al. (2021), chatbot Achiam et al. (2023), and complex reasoning Wei et al. (2022).
032 Given the widespread adoption of language models, it is crucial to evaluate their robustness. The
033 robustness problems in the NLP community mostly focus on exploring the language model behavior
034 when the inputs are modified.
035

036 Ebrahimi et al. (2017) design character-level and word-level perturbations as adversarial examples to
037 attack NLP models. Jia & Liang (2017) explore the method to mislead the language model’s output
038 in the Q&A task by adding random sentences. Later, research such as that by Jin et al. (2020) and Li
039 et al. (2020) focused on designing adversarial samples that better preserve the original semantics.
040 Subsequently, some work began to analyze the robustness of language models in continuous space
041 and improve the generalization ability of NLP models and defense against word substitution attacks
042 through adversarial training in continuous space Zhu et al.; Li et al. (2021; 2023). Recently, the
043 growth of model parameter size and training data for language models has demonstrated that Large
044 Language Models (LLMs) exhibit increased robustness to trivial disturbances and handle common
045 disruptions more effectively Achiam et al. (2023); Zou et al. (2023). As a result, recent studies have
046 devised Jailbreak prompts specifically designed to attack LLMs, thereby evaluating and testing their
047 robustness Wang et al. (2023); Zhu et al. (2023b).

048 While significant progress has been made in the robust evaluation of pre-trained language models,
049 current robustness analyses still face several key challenges. First, the methods for evaluating ro-
050 bustness in large language models (e.g., LLaMA-series) and those in smaller models (e.g., BERT,
051 BART) differ substantially. In smaller models, word-level or token-level adversarial attacks often
052 suffice to expose vulnerabilities Li et al. (2018; 2020); Garg & Ramakrishnan (2020). However,
053 large language models can often interpret or adapt to these simple perturbations, necessitating more
carefully designed prompt-based strategies for effective robustness testing Zou et al. (2023); Paulus
et al. (2024). As a result, a unified robustness evaluation framework applicable to both large and

small models is currently lacking. Second, irrespective of a model’s parameter size, designing adversarial inputs or prompts remains a time-consuming and labor-intensive process, particularly for large-scale NLP datasets. These challenges highlight the need for automated, scalable, and more universal approaches to evaluate the robustness of diverse language models.

In this work, we propose a sample-centric robustness framework that addresses both challenges by: 1) Unifying Robustness Evaluation for All Model Scales. Specifically, our method computes a local (per-sample) robustness measure, applicable to both smaller models (e.g., DistilBERT) and massive LLMs (e.g., GPT-2, Llama) *without* requiring changes to their internal parameters or specialized adversarial training. 2) Minimizing Labor-Intensive Perturbation Design. Instead of heavily relying on constructing adversarial prompts, we quantify each input’s inherent vulnerability via a lightweight, near-linear complexity approach. This ranking guides adversarial attacks or fine-tuning decisions, drastically reducing manual effort compared to purely sample-by-sample adversarial generation.

At the core of our method is a novel, per-sample distance mapping distortion (DMD) metric that compares distances in the input representation space against distances in the output representation space. To facilitate these distance calculations efficiently, we first build a *near-linear complexity* Probabilistic Graphical Model (PGM) that captures the manifold structure of the data, preserving both local geometry and global structural properties without resorting to dense or iterative global optimizations. By assessing each instance individually, we gain a fine-grained view of where and how a model fails to preserve distances across its representations—leading directly to broader insights about the system’s behavior as a whole. Such per-sample analyses, in turn, form the building blocks of understanding overall stability Zhang et al. (2019a). Rather than relying on aggregate statistics alone, examining each sample’s distortion enables us to pinpoint particular modes of fragility. We show how this ranking can: 1) Streamline NLP Attacks: Targeting non-robust samples first yields more efficient and more effective adversarial attacks (Section 4.2). 2) Improve LLM robustness through Fine-Tuning: Up-weighting non-robust data during fine-tuning preserves or even improves generalization and yields internal representations closer to the pre-trained checkpoint (Section 4.3). Furthermore, the same method applies to both smaller-scale models and large-scale LLMs, offering a unified pathway for robustness analysis across diverse parameter regimes.

Overall, our contributions are:

- A unified robustness measure (SALMAN) that can be computed in nearly-linear time for language models of varied sizes (from smaller transformers to LLMs), without requiring specialized tasks, perturbed data, or parameter modifications.
- To our best knowledge, SALMAN is the first local (sample-level) robustness measurement specifically tailored from small to large language models , enabling fine-grained analysis of how individual inputs withstand minor or adversarial perturbations.
- Empirical demonstrations across both small (BERT, DistilBERT) and large models (GPT-2, Llama) showing how this sample-level perspective leads to (i) more efficient and higher success-rate adversarial attacks, and (ii) improved robust fine-tuning outcomes.

2 BACKGROUND

2.1 ROBUSTNESS IN NLP

The robustness of language models remains a pivotal area of research within the NLP community. Several studies have explored the vulnerability of these models to modifications in the input text, ranging from typos to word replacements Li et al. (2020); Jin et al. (2020); Sun et al. (2020). Wang et al. (2021) further developed a multi-task benchmark to evaluate language model robustness. In the realm of model probing, Tenney (2019) and Hewitt & Manning (2019) examined how syntactic and semantic features are represented across different layers of BERT-like models. Voita et al. (2019) and Abnar et al. (2019) employed similarity-based analysis methods to study the evolution of representations in deep neural networks. Zhou & Srikumar (2021) and Neerudu et al. (2023) performed a comprehensive analysis of how finetuning affects the representations in the language model using a combination of probing and analytical techniques. With the increase in model parameters, LLMs can distinguish between minor textual variations, underlining the need to explore their robustness to input perturbations. Recent studies have focused on the impact of input prompts on LLM robustness Shayegani et al. (2023). Wang et al. assessed the robustness of ChatGPT against adversarial

108 and out-of-distribution samples. Zou et al. (2023) enhanced the efficiency of jailbreak attacks by
 109 generating adversarial suffixes. DecodingTrust examined the robustness of LLMs using standard
 110 datasets like AdvGLUE and AdvGLUE++ Wang et al. (2023). PromptRobust investigated the ro-
 111 bustness of LLMs from the perspective of prompts, demonstrating that subtle changes in instructions
 112 can lead to significant performance variations Zhu et al. (2023b).

113

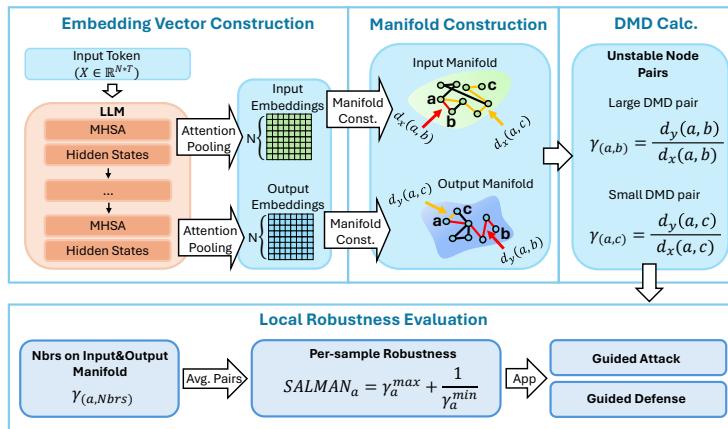
114 2.2 PROBABILISTIC GRAPHICAL MODELS

115

116 Probabilistic Graphical Models (PGMs) represent conditional dependencies among variables in a
 117 graph, enabling interpretability and efficient inference Koller (2009). Here, each sample (e.g., a
 118 Transformer embedding) becomes a node, with edges capturing local/global interactions that ap-
 119 proximate the data manifold Vu & Thai (2020); Feng (2021); Rubin-Delanchy (2020). Tightly
 120 connected subgraphs indicate higher intrinsic similarity, while loosely connected regions suggest
 121 divergence. Recently, SAGMAN Cheng et al. (2024) extends these PGM-based ideas to GNNs by
 122 incorporating dimension reduction for domain-specific manifold structures.

123 3 THEORETICAL FOUNDATIONS OF SALMAN

124



134 Figure 1: The overview of SALMAN Method.

135

136
137
138

In this section, we detail our overall pipeline (Figure 1), followed by the mathematical underpin-
 141 nings of (1) Embedding Construction (Sec. 3.1), (2) Manifold Construction (Sec. 3.2), (3) Distance
 142 Mapping Distortion (Sec. 3.3), and (4) Algorithm Complexity Analysis (Sec. 3.4).

143

144

145 3.1 EMBEDDING VECTOR CONSTRUCTION

146
147
148
149
150
151
152
153
154
155

A key challenge in analyzing transformer-based language models arises from the *discrete* nature of token embeddings, which may not form well-behaved manifolds in the topological sense (Robinson et al., 2024). In particular, the geometry of the token space is heavily fragmented: a small textual perturbation (e.g., replacing a token with a synonym) can induce a disproportionately large jump in token-level embedding indices. As a result, continuous manifold-based analyses, which rely on smooth neighborhoods and gradual changes, become intractable when applied directly to discrete tokens. Moreover, transformers often exhibit *stochastic decoding* (via temperature sampling, beam search, etc.) Li et al. (2024), meaning identical input text can produce slightly different token outputs. Hence, relying purely on discrete token sequences introduces variability that disrupts stable manifold construction.

156

157

158

159

160

161

Attention Based Embedding Representation. To address these issues, we aggregate each sample’s token embeddings into a *single* continuous vector, thereby avoiding the discontinuities of the raw token space. Formally, given a natural language dataset of N samples, each data sample is tokenized into a sequence of embeddings $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$. We then pass each \mathbf{x}_i through a Transformer-based pre-trained language model to obtain its Multi-Head Self-Attention (MHSA) *outputs*, which we denote as $A_i = \text{MHSA}(\mathbf{x}_i) \in \mathbb{R}^{H \times T_i \times d_{\text{model}}}$. H is the number of attention heads, T_i is length of \mathbf{x}_i , and d_{model} is the hidden dimension. Then, we average these per-head outputs along

162 the head dimension: $\bar{A}_i = \frac{1}{H} \sum_{h=1}^H A_{i,h} \in \mathbb{R}^{T_i \times d_{\text{model}}}$. Next, we *compute an attention-based*
 163 *weighting*. Softmax is applied to attention matrices \bar{A}_i over the T_i tokens to obtain $\alpha_i \in \mathbb{R}^{T_i}$. Fi-
 164 nally, we compute a single pooled vector $\mathbf{v}_i \in \mathbb{R}^{d_{\text{model}}}$ for each sequence by a weighted sum of
 165 the token embeddings: $\mathbf{v}_i = \sum_{t=1}^{T_i} \alpha_{i,t} \bar{A}_{i,t}$. Collecting all \mathbf{v}_i for $i = 1, \dots, N$ then yields a
 166 new embedding matrix $\{\mathbf{v}_1, \dots, \mathbf{v}_N\}$. We denote embedding matrix taken from the first layer of the
 167 language model as \mathbf{z}_X and embedding matrix taken from the last layer of the language model as \mathbf{z}_Y .
 168

169 **Deterministic Hidden-State Embeddings.** Though transformers can produce stochastic token
 170 outputs, the internal hidden states remain largely deterministic *if we freeze the model parameters and*
 171 *inference procedure*. For instance, by disabling dropout layers and using a fixed random seed, we
 172 empirically found that \mathbf{z}_Y becomes stable irrespective of minor token-level variations. Specifically,
 173 we measured multiple \mathbf{z}_Y across the same input and observed a significant similarity improvement
 174 compared to token embeddings (As detailed in the Appendix A.1).

175 By aggregating discrete token embeddings into a single high-level embedding, we circumvent the
 176 discontinuities of token-level spaces. We thereby ensure that (1) manifold analysis is tractable (since
 177 $\mathbf{z}_X, \mathbf{z}_Y$ both lie in continuous \mathbb{R}^d Mehta et al. (2019)), and (2) stochastic decoding does not cause
 178 major geometric shifts in these embeddings. This design choice, while straightforward, underpins
 179 all subsequent sections on manifold construction and robustness evaluation.

180 3.2 CONSTRUCTION OF MANIFOLDS VIA PGM

181 Understanding a language model’s local robustness involves assessing how small input perturbations
 182 influence the model’s output representation. A common strategy is to interpret this input–output
 183 mapping as a *manifold*, enabling geometric analyses of local stability (Rubin-Delanchy, 2020).
 184 However, directly constructing and maintaining such a manifold on raw embeddings can be both
 185 computationally and memory intensive. Recent work indicates that *graph-based* approaches can
 186 capture low-dimensional manifolds within high-dimensional data (Rubin-Delanchy, 2020), *espe-*
 187 *cially* when the graph is constructed to preserve meaningful distances. PGMs (or Markov Random
 188 Fields) naturally encode these relationships in an undirected graphical structure, allowing for effi-
 189 cient inference about node neighborhoods and global structure (Koller, 2009). Specifically, Feng
 190 (2021) show that the graph structure learned by PGMs can approximate *resistance distances*, which
 191 in turn correlate with Euclidean distances among data samples. Hence, a properly built PGM mani-
 192 fold can reflect *both local and global* geometry—critical for analyzing small perturbations (local
 193 stability) and broader connectivity (global structure).

194 Despite their theoretical appeal, existing PGM-based methods rely on iterative optimization or dense
 195 computations (e.g., spectral factorization) that become prohibitive for large-scale graphs (Feng,
 196 2021). When handling modern NLP datasets, where each sample might represent a document or
 197 prompt, and node counts can soar into the hundreds of thousands, these bottlenecks make traditional
 198 PGM approaches infeasible. To address scalability concerns, we propose a *near-linear complexity*
 199 method for building the PGM manifold. Intuitively, we seek a graph Laplacian structure (or pre-
 200 cision matrix) that captures the intrinsic geometry of the reduced embeddings (Sec. 3.1) without
 201 incurring expensive global factorization steps. Previous work (Dong et al., 2019) shows that maxi-
 202 mizing a penalized log-likelihood in the form of Equation. 1 yields a graph topology consistent with
 203 the underlying data distribution while preserving essential distance or similarity properties.

204 **PGM Objective.** Let $X \in \mathbb{R}^{|V| \times T}$ be the embedding matrix derived from Sec. 3.1, where each row
 205 corresponds to a sample. We aim to learn a precision matrix Θ that maximizes Dong et al. (2019):

$$207 \max_{\Theta} F(\Theta) = \log \det(\Theta) - \frac{1}{k} \text{Tr}(X^\top \Theta X), \quad (1)$$

208 subject to $\Theta = \mathcal{L} + \frac{1}{\sigma^2} I$, where \mathcal{L} is a valid Laplacian matrix and σ^2 is a prior variance term.
 209 Theorem 3.1 shows that *maximizing* $F(\Theta)$ can be achieved using a spectral sparsification approach,
 210 which prunes edges with small distance ratios:

$$212 \rho_{p,q} = \frac{d^{\text{eff}}(p, q)}{d^{\text{dat}}(p, q)} = w_{p,q} (d^{\text{eff}}(p, q)),$$

213 where $d^{\text{eff}}(p, q)$ is the effective resistance distance (detailed in Appendix A.10) between nodes p
 214 and q , $d^{\text{dat}}(p, q) = \|X_p - X_q\|_2^2$ is the data distance, and $w_{p,q} = 1/d^{\text{dat}}(p, q)$.

216 **Theorem 3.1.** *Maximizing the objective in Equation equation 1 can be done via an edge pruning*
 217 *strategy equivalent to spectral sparsification of the initial (dense) graph. Edges with small $\rho_{p,q}$ are*
 218 *removed, preserving the essential spectral (and thus resistance) properties of the original graph.*
 219 *The proof is available in Appendix A.2*

221 **Scalable Spectral Sparsification via Short-Cycle Decomposition.** A naive implementation of the
 222 above pruning requires frequent effective-resistance computations Spielman & Srivastava (2008),
 223 which is costly for weighted graphs. Methods such as short-cycle decomposition Chu et al. (2020)
 224 are effective for *unweighted* graphs but fail to retain accurate resistance distances when weights are
 225 discarded. We therefore introduce a refined *spectral sparsification* routine that uses low-resistance-
 226 diameter (LRD) decomposition to handle weighted edges without sacrificing the crucial resistance
 227 distance information.

228 **Theorem 3.2.** *Our LRD decomposition can compute the effective resistance of each edge and is*
 229 *capable of sparsifying weighted graphs. The proof is available in Appendix A.3*

230 **From PGM to Manifold.** With the pruned graph (and correspondingly updated Laplacian \mathcal{L}),
 231 solving Equation equation 1 yields a precision matrix Θ that encodes the desired topological rela-
 232 tionships in X . This PGM thus underpins our low-dimensional manifold, accurately maintaining
 233 resistance distances for subsequent stability analyses (detailed in Appendix A.9). In practice, we
 234 initialize the graph with a k -nearest-neighbor construction and then apply our near-linear spectral
 235 sparsification (as detailed in Section 3.4) to achieve scalability. The resulting *manifold* reflects both
 236 local and global structures, enabling the DMD calculation.

238 3.3 DISTANCE MAPPING DISTORTION (DMD) CALCULATION

239 Having constructed the input and output manifolds (Section 3.2), we now introduce the DMD met-
 240 ric Cheng et al. (2021) to quantify a model’s robustness at the *sample level*.

242 **Definition 3.3** (Distance Mapping Distortion (DMD)). Let F be a function mapping an input mani-
 243 fold $G_X = (V, E_X)$ to an output manifold $G_Y = (V, E_Y)$, with $d_X(p, q)$ and $d_Y(p, q)$ denoting the
 244 distances between nodes p and q in G_X and G_Y , respectively. The distance mapping distortion for
 245 (p, q) through F is

$$246 \gamma^F(p, q) = \frac{d_Y(p, q)}{d_X(p, q)}. \quad (2)$$

248 **Innovation Highlight:** We show that not only is $\gamma_{\max}^F = \max_{p,q} \gamma^F(p, q)$ informative for worst-
 249 case local expansion, but also $(\gamma_{\min}^F)^{-1} = (\min_{p,q} \gamma^F(p, q))^{-1}$ captures how the model might
 250 “collapse” distant inputs into overly similar outputs. We prove in Theorem 3.4 (below) that large
 251 $(\gamma_{\min}^F)^{-1}$ implies *another dimension* of poor robustness—distinct from γ_{\max}^F (Empirical results are
 252 available in Appendix A.11). Hence, both extremes of the distortion spectrum are necessary for a
 253 full local analysis.

254 **Effective-resistance distance.** (as detailed in Appendix A.10) To make γ^F computationally
 255 tractable, we replace geodesic distances with *effective-resistance* (d^{eff}). $d^{\text{eff}}(p, q)$ is always matched
 256 or upper-bounded by $d^{\text{geo}}(p, q)$ Chandra et al. (1996). Thus, d^{eff} is an *efficient* surrogate for d^{geo} ,
 257 especially when leveraging fast Laplacian solvers Koutis et al. (2010); Kyng & Sachdeva (2016).
 258 We then define

$$259 \gamma^F = \frac{d_Y^{\text{eff}}(p, q)}{d_X^{\text{eff}}(p, q)} = \frac{e_{p,q}^\top L_Y^+ e_{p,q}}{e_{p,q}^\top L_X^+ e_{p,q}}, \quad (3)$$

262 where L_X and L_Y denote the Laplacians of G_X and G_Y , respectively. Computing γ_{\max}^F or γ_{\min}^F
 263 *exactly* via Equation equation 3 can still be expensive for large graphs, since it involves considering
 264 all node pairs (p, q) . To alleviate this, Cheng et al. (2021) proposed a *spectral* upper bound on γ_{\max}^F ,
 265 termed the $\lambda_{\max}(L_Y^+ L_X)$. Hence, a larger $\lambda_{\max}(L_Y^+ L_X)$ suggests a larger distortion ratio and thus
 266 poorer robustness. This is also the upper bound of the best Lipschitz constant under the manifold
 267 setting Cheng et al. (2021). For γ_{\min}^F lower bound calculation, we have:

268 **Theorem 3.4.** *The minimum distance mapping distortion γ_{\min}^F satisfies*

$$269 \gamma_{\min}^F \geq \frac{1}{\lambda_{\max}(L_X^+ L_Y)}.$$

270 *The proof is available in Appendix A.4*
 271

272 **SALMAN Score** measures, for each sample, how strongly the model distorts input–output dis-
 273 tances, penalizing both expansion (near neighbors blown apart) and collapse (far points mapped
 274 too close), so larger scores indicate greater fragility. For each node (sample) p , we define
 275 $\text{SALMAN}^F(p)$:

$$276 \quad \frac{1}{|\mathcal{N}(p)|} \sum_{q \in \mathcal{N}(p)} (\gamma^F(p, q)^3 + \gamma^F(p, q)^{-3}), \quad (4)$$

279 where $\mathcal{N}_X(p)$ is the set of neighbors of p in G_X and G_Y . A node with a larger $\text{SALMAN}^F(p)$ is con-
 280 sidered more *fragile*, since its local pairs (p, q) exhibit greater distortion in “expansion” ($\gamma^F(p, q)$)
 281 or “collapse” ($1/\gamma^F(p, q)$) senses. To connect expansions and collapses more explicitly, let $\{\lambda_i\}_{i=1}^r$
 282 be the r largest eigenvalues of $L_Y^+ L_X$ with corresponding eigenvectors $\{v_i\}_{i=1}^r$, and let $\{\mu_i\}_{i=1}^r$
 283 be the r largest eigenvalues of $L_X^+ L_Y$ with corresponding eigenvectors $\{w_i\}_{i=1}^r$. We define the
 284 weighted eigensubspace matrices: $V_r = [v_1 \sqrt{\lambda_1}, \dots, v_r \sqrt{\lambda_r}]$, $W_r = [w_1 \sqrt{\mu_1}, \dots, w_r \sqrt{\mu_r}]$.
 285 For each pair (p, q) , one has:

286 **Theorem 3.5.** $\|W_r^\top e_{p,q}\|_2^2 + \|V_r^\top e_{p,q}\|_2^2 \propto \gamma^F(p, q)^3 + \gamma^F(p, q)^{-3}$. *The proof is available in*
 287 *Appendix A.5*

289 **Sample Selection and Correction.** Because SALMAN score is computed at the node or node-pair
 290 level, we can readily identify “high-distortion” samples and correct them via data augmentation or
 291 specialized re-training. This local approach complements global metrics, yielding a holistic robust-
 292 ness analysis pipeline.

294 3.4 COMPLEXITY

296 Our framework has *near-linear* time complexity with respect to the graph size. Below, we briefly
 297 outline the main steps and their costs. We first construct a k -NN graph from the data points (or
 298 embeddings) in \mathbb{R}^d . Using modern approximate nearest-neighbor algorithms (Malkov & Yashunin,
 299 2018) with $O(|V| \log |V|)$. $|V|$ denotes the number of nodes in the graph. Then, We apply a Low-
 300 Resistance-Diameter (LRD) approach to sparsify the graph (Koutis et al., 2010; Cucuringu et al.,
 301 2016). Let d be the average degree (small in real-world graphs (Miao et al., 2019)) and m be the
 302 dimension of a Krylov subspace. Then this step runs in $O(|V| d m)$, often simplified to $O(|V| m)$
 303 under the sparse regime. Evaluating the SALMAN scores for all edges or nodes can be done in
 304 $O(|E|)$ time. $|E|$ denotes the number of edges in the graph. For sparse graphs with $|E| \approx d|V|$,
 305 this remains near-linear in $|V|$. Experimental results can be found in Appendix A.12.

306 4 EXPERIMENT RESULTS

308 We organize our experimental evaluation into three stages, each demonstrating how the *robustness*
 309 *ranking* (derived from the proposed SALMAN measure) can guide practical NLP tasks. The lan-
 310 guage models used for experiments range from BERT (136M), GPT-2 (1.5B) and the latest Llama3-
 311 8B. Details regarding data, hyperparameters, and model architectures are deferred to Appendix A.6

312 As this is the first work to propose a per-sample NLP robustness ranking, we lack direct compar-
 313 isons with methods pursuing identical objectives. However, to address the lack of baseline concern,
 314 we compare SALMAN against: (1) Euclidean-distance-based ranking, which measures each sam-
 315 ple’s magnitude of embeddings without local manifold distortion; and (2) Jacobian-based sensitivity
 316 analysis. These baselines are simpler proxies for identifying “fragile” points. In Table 3, we show
 317 that both struggle to distinguish robust vs. non-robust samples under the same spaCy perturbation.

318 Additionally, we analyze representative robust versus non-robust samples to confirm that SALMAN
 319 reliably identifies vulnerable cases. Moreover, we compare our SALMAN measure against simpler
 320 baselines—such as random ranking, a state-of-the-art attack run without our approach, and a state-
 321 of-the-art robust training procedure without SALMAN—to assess whether our method provides
 322 tangible gains. Our experiments show that SALMAN surpasses these heuristic baselines on diverse
 323 perturbation benchmarks, offering strong evidence that SALMAN captures unique facets of sample-
 level vulnerability.

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Table 1: Cosine similarities of robust and non-robust samples for different models on SST-2 and MNLI.

Dataset	Model	Robust/Non-robust Cosine Similarity
SST-2	BERT-base-uncased	0.9911/0.8711
	RoBERTa-base	0.9992/0.9895
	DistilBERT-base-uncased	0.9955/0.9404
	ALBERT-base-v2	0.9959/0.8279
	GPT-2	0.9990 / 0.9153
	LLaMA-7B-v2	0.9867 / 0.9160
MNLI	BERT-base-uncased	0.9902/0.9410
	RoBERTa-base	0.9993/0.9926
	DistilBERT-base-uncased	0.9971/0.9650
	ALBERT-base-v2	0.9953/0.8709
	GPT-2	0.9993 / 0.9904
	LLaMA-7B-v2	0.9925 / 0.9842

Table 3: Cosine similarity between original vs. spaCy-perturbed samples on GPT-2, for Euclidean-, Jacobian-, and SALMAN-based rankings. We aim for robust sets to have higher similarity and non-robust sets to have lower similarity. SALMAN yields the largest gap.

Method	SST-2 / MNLI	Gap
Euclidean	R: 0.9953 / 0.9918 NR: 0.9986 / 0.9898	→ 0.0033 / 0.0020
Jacobian	R: 0.9964 / 0.9942 NR: 0.9965 / 0.9793	→ 0.0001 / 0.0149
SALMAN	R: 0.9981 / 0.9995 NR: 0.9772 / 0.9730	→ 0.0209 / 0.0265

4.1 SAMPLE ROBUSTNESS EVALUATION

To validate that our robustness ranking meaningfully distinguishes between stable and fragile samples, we subject both *robust* (1% samples with lowest SALMAN) and *non-robust* (1% samples with highest SALMAN) samples to various NLP perturbations. These perturbations simulate natural edits or noise while controlling for the extent of modification via Levenshtein distance Ding et al. (2021). We thereby ensure that robust and non-robust subsets are perturbed equally in terms of edit cost, allowing a fair comparison of downstream output changes.

Following standard practices in text perturbations Guo et al. (2021); Ni et al. (2024), we implement three simple but widely used edits: deletion, insertion, and swap. Following previous works Le et al. (2022); Gupta et al. (2023); Jia et al. (2023), we measure the resultant output embedding shift via cosine similarity between the original and perturbed sentence embeddings, as seen in Table 1.

Beyond the three basic edits, we employ two state-of-the-art perturbation frameworks (spaCy Honnibal et al. (2020) and TextAttack Morris et al. (2020b)) to generate more sophisticated attacks. These methods leverage advanced synonym replacement and gradient-informed edits to produce challenging textual perturbations. Due to the substantial computational overhead of these approaches, we restrict them to two widely recognized LLMs—GPT-2 and LLaMA-7B-v2—thereby striking a balance between experimental rigor and resource feasibility. In Table 2, we apply each SOTA method to both robust and non-robust subsets, measuring the resulting cosine similarities to assess susceptibility to adversarial manipulations.

To further assess the difference between robust and non-robust samples, we incorporate two additional metrics: KL-Divergence (KLD) and BERTScore Zhang et al. (2019b). Table 4 shows that non-robust samples exhibit larger distribution shift (higher KLD) and lower textual similarity (BERTScore) under perturbations, whereas robust samples remain highly similar. We provide statistical reliability analysis in Appendix A.1.

By systematically perturbing both robust and non-robust samples, we confirm that non-robust samples consistently exhibit greater output variability under identical input changes. This aligns with prior evidence that local text modifications can disproportionately affect certain data points Morris et al. (2020a), and it underscores the value of distance mapping distortion in identifying vulnerabilities at the sample level.

Table 2: Robust and non-robust cosine similarities under two different attack methods (spaCy & TextAttack).

Attack	Model	Dataset	Robust/Non-robust Cosine Sim.
spaCy	GPT-2	SST-2	0.9981/0.9772
	GPT-2	MNLI	0.9995/0.9730
	LLaMA-7B-v2	SST-2	0.9990/0.9751
	LLaMA-7B-v2	MNLI	0.9825/0.9612
TextAttack	GPT-2	SST-2	0.9928/0.9413
	GPT-2	MNLI	0.9945/0.9831
	LLaMA-7B-v2	SST-2	0.9548/0.8941
	LLaMA-7B-v2	MNLI	0.9663/0.9479

Table 4: BERTScore and KLD evaluations of robust vs. non-robust subsets (GPT-2, SST-2). Higher BERTScore indicates higher textual similarity. KLD measures distribution shift (lower is more stable).

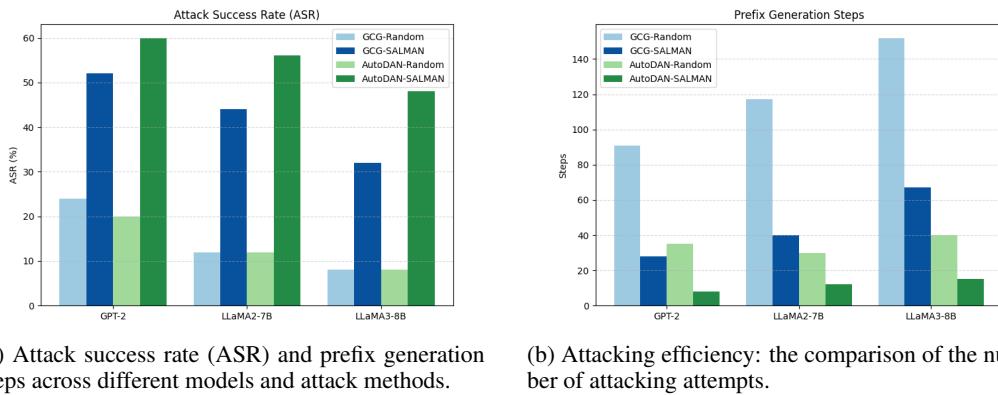
	KLD	BERTScore		
		Precision	Recall	F1
Non-Robust	0.1923	0.9961	0.9970	0.9965
Robust	7.6175e-07	0.9992	0.9991	0.9992

378 4.2 SALMAN-GUIDED ATTACK
379

380 Jailbreak Attacks are an important way to assess the security and robustness of LLM Yi et al. (2024);
381 Chu et al. (2024). By strategically crafting prompts, it is possible to bypass the LLM’s inherent
382 safeguards and generate harmful content. Recently, numerous studies have focused on automatically
383 generating stealthy jailbreak attack prompts. However, these current methods are both labor-
384 intensive and computationally demanding. We propose using the SALMAN score to guide more
385 effective attacks.

386 **Motivation: Find the Non-Robust Data Samples.** We harness the robustness ranking to focus
387 adversarial efforts on the most susceptible samples. This strategy is akin to reducing query com-
388 plexity in black-box attacks or prioritizing the most “fragile” points. We structured the experiment
389 as follows: 1) we rank the dataset by descending SALMAN score (i.e., from least robust to most
390 robust). 2) We perform the existing attack method *only* on the top $k\%$ of non-robust samples. 3) Then
391 we randomly sample $k\%$ data and use the same method to attack LLM again as a comparison.

392 For our experiment, we take GCG Zou et al. (2023) and AutoDAN Liu et al. (2023) as the jailbreak
393 attack method and use the AdvBench Harmful Behaviors dataset Zou et al. (2023) to evaluate the
394 jailbreak attacks. This dataset has 520 data points (Dataset detail in Appendix A.6). After ranking
395 all the data using SALMAN, we selected the top 1% of unstable samples to launch attacks on LLMs,
396 supplemented by randomly sampling another 1% of the samples for the same purpose. Subsequently,
397 we evaluated the effectiveness of SALMAN by comparing changes in the Attack Success Rate
398 (ASR) and the number of attack attempts (Steps). We also justify extracting embeddings from the
399 language model’s first and final hidden layers, which capture complementary semantic information
400 (Empirical results are available in Appendix A.13.2).



411 (a) Attack success rate (ASR) and prefix generation
412 steps across different models and attack methods.
413

414 (b) Attacking efficiency: the comparison of the num-
415 ber of attacking attempts.

416 Figure 2: Comparison of adversarial attack performance (a) and efficiency (b) with and without
417 SALMAN-guided selection.

418 Figure 2 (a) shows that attacking these low-robustness samples first yields higher success rates *and*
419 reduced time-to-attack compared to random sampling baselines. We visualized attacking efficiency
420 in Figure 2 (b). The SALMAN-based ranking serves as an efficient “shortcut” for adversarial testing.
421 Then, we investigate SALMAN ranking by selecting different top- k percentages of the dataset. We
422 bin (10%) the entire dataset into deciles by SALMAN rank and apply the same attack methods to
423 each bin. As k increases, we include more (relatively) robust samples, resulting in lower overall ASR
424 and efficiency. We further assessed the SALMAN-Guided attack using proxy models to evaluate the
425 robustness of the proposed method. The experimental results are presented in Table 5. The re-
426 sults demonstrate that SALMAN-guided attacks retain high effectiveness across models, confirming
427 their transferability. We also conduct the SALMAN-guided attack experiment on the multilingual
428 Jailbreak dataset and we show the results in Appendix A.13.

429 4.3 SALMAN-GUIDED LLM FINE-TUNING
430

431 Fine-tuning LLMs sometimes leads to overfitting, losing key representations from pre-
432 training Howard & Ruder (2018); Neerudu et al. (2023). Prior work has attempted to measure
433 how much an LLM’s internal representations drift from the pre-trained checkpoint, using simi-
434 larity metrics such as CKA or STIR Neerudu et al. (2023). A large drift often indicates the model is
435 overfitting, thus sacrificing generalization and robustness in practice.

432 Table 5: Proxy-based SALMAN ranking vs.
 433 target LLM under GCG. We list the ASR on
 434 the top-1% non-robust subsets identified by the
 435 proxy model.

Proxy	GPT-2	LLaMA2-7B	LLaMA3-8B
GPT-2	60%	48%	40%
LLaMA2-7B	60%	56%	40%
LLaMA3-8B	60%	56%	48%

441 Table 6: Performance of ROSE fine-tuning
 442 on SST-2, RTE, and QNLI tasks with
 443 RoBERTa_{BASE}. Each cell shows *ROSE fine-*
 444 *tuning / SALMAN-guided ROSE fine-tuning* ac-
 445 curacy. Bold is better.

Task	GLUE	AdvGLUE
SST-2	94.84 / 94.84	37.67 / 41.22
RTE	78.34 / 78.34	35.49 / 40.75
QNLI	92.19 / 92.81	44.19 / 48.02

Table 7: Model-level robustness score Neerudu et al. (2023) of BERT and GPT-2 on CoLA and SST-2 under various perturbations. Each cell shows the *Normal/SALMAN-guided* fine-tuning value. Better results are in **bold**.

Perturbation	CoLA	SST-2
<i>BERT</i>		
Drop nouns	0.18 / 0.29	0.92 / 1.02
Drop verbs	0.05 / 0.22	0.95 / 1.03
Drop first	0.48 / 0.70	0.98 / 1.01
Drop last	0.34 / 0.72	1.00 / 1.00
Swap text	0.13 / 0.22	0.98 / 1.01
Add text	0.85 / 0.88	0.99 / 0.99
Change char	0.14 / 0.24	0.84 / 0.91
Bias	0.95 / 0.99	1.00 / 1.00
<i>GPT-2</i>		
Drop nouns	0.10 / 0.47	0.93 / 1.00
Drop verbs	0.24 / 0.32	0.95 / 1.00
Drop first	0.75 / 0.91	0.97 / 1.01
Drop last	0.45 / 0.78	0.99 / 1.01
Swap text	0.16 / 0.45	0.98 / 1.00
Add text	0.92 / 0.96	0.99 / 1.01
Change char	0.29 / 0.36	0.86 / 1.02
Bias	0.96 / 1.14	1.01 / 0.99

453 **Motivation: Focus on Non-Robust Data.** Several studies show that focusing on non-robust samples
 454 during training can improve model robustness and generalizability (Cheng et al., 2021; Zhu
 455 et al., 2023a). Inspired by these findings, we propose to *down-weight* robust samples and *up-weight*
 456 non-robust samples (as determined by our SALMAN-based ranking) when fine-tuning an LLM. The
 457 intuition is that easy/robust data rarely contributes to boosting generalizable features, whereas harder
 458 (high DMD) data pushes the model to learn more discriminative patterns.

459 We follow the fine-tuning protocol described by Neerudu et al. (2023). By placing greater emphasis
 460 on non-robust data (as detailed in Appendix A.8), we hypothesize that the fine-tuned model retains
 461 more generalizable features from its pre-training, avoiding over-specialization to easy examples.
 462 We observe two key outcomes: 1) *Comparable Performance, Closer to Pre-training*: SALMAN-
 463 guided LLMs achieve comparable or better accuracy vs. standard fine-tuning, yet exhibit higher
 464 similarity to the pre-trained checkpoint. On GLUE tasks, we find up to 54% gains in CKA or
 465 STIR, signifying less drift with better accuracy. Results appear in Appendix A.7 and highlight
 466 key findings. 2) *Enhanced Robustness Scores*: Although our paper introduces a *sample-centric*
 467 robustness measure for *ranking* individual samples *without* requiring explicit perturbations, we also
 468 need to assess the *entire model*'s robustness after fine-tuning. To this end, we adopt the model-level
 469 robustness score proposed by Neerudu et al. (2023), which measures how representations change
 470 under various text perturbations across the full dataset. Measuring each model's robustness scores
 471 confirms that the *SALMAN-guided* LLM obtains higher robustness than a conventionally fine-tuned
 472 model. As shown in Table 7, we attribute this improvement to the heightened focus on challenging
 473 (non-robust) samples during training.

474 **Combining with SOTA Robust Training.** We further integrate our approach with ROSE (Jiang
 475 et al., 2022), a selective fine-tuning framework that prunes “spurious” parameter updates to achieve
 476 greater adversarial resilience. Specifically, we embed our SALMAN-based weighting (as detailed in
 477 Appendix A.8) within ROSE's parameter selection process. Experimental results in Table 6 reveal
 478 that sample-level weighting and parameter-level selection are complementary strategies.

479 5 CONCLUSION

480 We introduced SALMAN, a novel measure to identify and rank the local robustness of transformer-
 481 based language models. Our experiments across diverse benchmarks and large language models
 482 reveal that SALMAN not only distinguishes *robust* from *non-robust* samples under both simple and
 483 SOTA perturbations, but also effectively guides attacks and fine-tuning. Moreover, incorporating
 484 SALMAN into the existing robust training framework yields even greater resilience against adver-
 485 sarial perturbations. These results underscore the potential of leveraging *sample-level robustness* to
 486 bolster both attack strategies and robust model adaptation.

486 REFERENCES
487

488 Samira Abnar, Lisa Beinborn, Rochelle Choenni, and Willem Zuidema. Blackbox meets blackbox:
489 Representational similarity and stability analysis of neural language models and brains. *arXiv*
490 *preprint arXiv:1906.01539*, 2019.

491 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
492 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
493 report. *arXiv preprint arXiv:2303.08774*, 2023.

494 D Babić, Douglas J Klein, István Lukovits, Sonja Nikolić, and N Trinajstić. Resistance-distance ma-
495 trix: A computational algorithm and its application. *International Journal of Quantum Chemistry*,
496 90(1):166–176, 2002.

497 Ashok K Chandra, Prabhakar Raghavan, Walter L Ruzzo, and Roman Smolensky. The electrical
498 resistance of a graph captures its commute and cover times. In *Proceedings of the twenty-first*
499 *annual ACM symposium on Theory of computing*, pp. 574–586, 1989.

500 Ashok K Chandra, Prabhakar Raghavan, Walter L Ruzzo, Roman Smolensky, and Prasoon Tiwari.
501 The electrical resistance of a graph captures its commute and cover times. *Computational Com-
502 plexity*, 6(4):312–340, 1996.

503 Wuxinlin Cheng, Chenhui Deng, Zhiqiang Zhao, Yaohui Cai, Zhiru Zhang, and Zhuo Feng. Spade:
504 A spectral method for black-box adversarial robustness evaluation. In *International Conference
505 on Machine Learning*, pp. 1814–1824. PMLR, 2021.

506 Wuxinlin Cheng, Chenhui Deng, Ali Aghdaei, Zhiru Zhang, and Zhuo Feng. Sagman: Stability
507 analysis of graph neural networks on the manifolds. *arXiv preprint arXiv:2402.08653*, 2024.

508 Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes, and Yang Zhang. Compre-
509 hensive assessment of jailbreak attacks against llms. *arXiv preprint arXiv:2402.05668*, 2024.

510 Timothy Chu, Yu Gao, Richard Peng, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang.
511 Graph sparsification, spectral sketches, and faster resistance computation via short cycle decom-
512 positions. *SIAM Journal on Computing*, (0):FOCS18–85, 2020.

513 Mihai Cucuringu, Ioannis Koutis, Sanjay Chawla, Gary Miller, and Richard Peng. Simple and scal-
514 able constrained clustering: a generalized spectral method. In *Artificial Intelligence and Statistics*,
515 pp. 445–454. PMLR, 2016.

516 Yue Deng, Wenxuan Zhang, Sinno Jialin Pan, and Lidong Bing. Multilingual jailbreak challenges
517 in large language models. In *The Twelfth International Conference on Learning Representations*.

518 Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding.
519 *arXiv preprint arXiv:1810.04805*, 2018.

520 Shuoyang Ding, Marcin Junczys-Dowmunt, Matt Post, and Philipp Koehn. Levenshtein training for
521 word-level quality estimation. *arXiv preprint arXiv:2109.05611*, 2021.

522 Xiaowen Dong, Dorina Thanou, Michael Rabbat, and Pascal Frossard. Learning graphs from data:
523 A signal representation perspective. *IEEE Signal Processing Magazine*, 36(3):44–63, 2019.

524 Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box adversarial examples
525 for text classification. *arXiv preprint arXiv:1712.06751*, 2017.

526 Wafaa S El-Kassas, Cherif R Salama, Ahmed A Rafea, and Hoda K Mohamed. Automatic text
527 summarization: A comprehensive survey. *Expert systems with applications*, 165:113679, 2021.

528 Wendy Ellens, Floske M Spieksma, Piet Van Mieghem, Almerima Jamakovic, and Robert E Kooij.
529 Effective graph resistance. *Linear algebra and its applications*, 435(10):2491–2506, 2011.

530 Zhuo Feng. Sgl: Spectral graph learning from measurements. In *2021 58th ACM/IEEE Design
531 Automation Conference (DAC)*, pp. 727–732. IEEE, 2021.

540 Siddhant Garg and Goutham Ramakrishnan. Bae: Bert-based adversarial examples for text classifi-
 541 cation. *arXiv preprint arXiv:2004.01970*, 2020.

542

543 Junliang Guo, Zhirui Zhang, Linlin Zhang, Linli Xu, Boxing Chen, Enhong Chen, and Weihua Luo.
 544 Towards variable-length textual adversarial attacks. *arXiv preprint arXiv:2104.08139*, 2021.

545

546 Ashim Gupta, Carter Wood Blum, Temma Choji, Yingjie Fei, Shalin Shah, Alakananda Vempala,
 547 and Vivek Srikumar. Don't retrain, just rewrite: Countering adversarial perturbations by rewriting
 548 text. *arXiv preprint arXiv:2305.16444*, 2023.

549

550 John Hewitt and Christopher D Manning. A structural probe for finding syntax in word representa-
 551 tions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association
 552 for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short
 553 Papers)*, pp. 4129–4138, 2019.

554

555 Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. spaCy: Industrial-
 556 strength Natural Language Processing in Python. 2020. doi: 10.5281/zenodo.1212303.

557

558 Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification.
 559 *arXiv preprint arXiv:1801.06146*, 2018.

560

561 Ouyang Jia, Huimin Huang, Jiaxin Ren, Luodi Xie, and Yinyin Xiao. Contrastive learning with text
 562 augmentation for text classification. *Applied Intelligence*, 53(16):19522–19531, 2023.

563

564 Robin Jia and Percy Liang. Adversarial examples for evaluating reading comprehension systems.
 565 *arXiv preprint arXiv:1707.07328*, 2017.

566

567 Lan Jiang, Hao Zhou, Yankai Lin, Peng Li, Jie Zhou, and Rui Jiang. Rose: Robust selective fine-
 568 tuning for pre-trained language models. *arXiv preprint arXiv:2210.09658*, 2022.

569

570 Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. Is bert really robust? a strong baseline
 571 for natural language attack on text classification and entailment. In *Proceedings of the AAAI
 572 conference on artificial intelligence*, volume 34, pp. 8018–8025, 2020.

573

574 Daphane Koller. Probabilistic graphical models: Principles and techniques, 2009.

575

576 Ioannis Koutis, Gary L Miller, and Richard Peng. Approaching optimality for solving sdd linear
 577 systems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*,
 578 pp. 235–244. IEEE, 2010.

579

580 Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse,
 581 and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*,
 582 pp. 573–582. IEEE, 2016.

583

584 Zhenzhong Lan. Albert: A lite bert for self-supervised learning of language representations. *arXiv
 585 preprint arXiv:1909.11942*, 2019.

586

587 Thai Le, Jooyoung Lee, Kevin Yen, Yifan Hu, and Dongwon Lee. Perturbations in the wild: Lever-
 588 aging human-written text perturbations for realistic adversarial attack and defense. *arXiv preprint
 589 arXiv:2203.10346*, 2022.

590

591 Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. Textbugger: Generating adversarial text
 592 against real-world applications. *arXiv preprint arXiv:1812.05271*, 2018.

593

594 Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. Bert-attack: Adversarial
 595 attack against bert using bert. *arXiv preprint arXiv:2004.09984*, 2020.

596

597 Linyang Li, Ke Ren, Yunfan Shao, Pengyu Wang, and Xipeng Qiu. Perturbscore: Connecting
 598 discrete and continuous perturbations in nlp. In *Findings of the Association for Computational
 599 Linguistics: EMNLP 2023*, pp. 6638–6648, 2023.

599

600 Yiwei Li, Fei Mi, Yitong Li, Yasheng Wang, Bin Sun, Shaoxiong Feng, and Kan Li. Dy-
 601 namic stochastic decoding strategy for open-domain dialogue generation. *arXiv preprint
 602 arXiv:2406.07850*, 2024.

594 Zongyi Li, Jianhan Xu, Jiehang Zeng, Linyang Li, Xiaoqing Zheng, Qi Zhang, Kai-Wei Chang, and
 595 Cho-Jui Hsieh. Searching for an effective defender: Benchmarking defense against adversarial
 596 word substitution. In *Proceedings of the 2021 Conference on Empirical Methods in Natural
 597 Language Processing*, pp. 3137–3147, 2021.

598 Xiaogeng Liu, Nan Xu, Muhan Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak
 599 prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*, 2023.

600 Yinhai Liu. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint
 601 arXiv:1907.11692*, 364, 2019.

602 Yu A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search
 603 using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and
 604 machine intelligence*, 42(4):824–836, 2018.

605 Sachin Mehta, Rik Koncel-Kedziorski, Mohammad Rastegari, and Hannaneh Hajishirzi. Define:
 606 Deep factorized input word embeddings for neural sequence modeling. *CoRR*, 2019.

607 Jianyu Miao, Tiejun Yang, Junwei Jin, and Lingfeng Niu. Graph-based clustering via group sparsity
 608 and manifold regularization. *IEEE Access*, 7:172123–172135, 2019.

609 Bojan Mohar. Graph laplacians. *Topics in algebraic graph theory*, 102:113–136, 2004.

610 John X Morris, Eli Lifland, Jack Lanchantin, Yangfeng Ji, and Yanjun Qi. Reevaluating adversarial
 611 examples in natural language. *arXiv preprint arXiv:2004.14174*, 2020a.

612 John X Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. Textattack: A frame-
 613 work for adversarial attacks, data augmentation, and adversarial training in nlp. *arXiv preprint
 614 arXiv:2005.05909*, 2020b.

615 Pavan Kalyan Reddy Neerudu, Subba Reddy Oota, Mounika Marreddy, Venkateswara Rao Kagita,
 616 and Manish Gupta. On robustness of finetuned transformer-based nlp models. *arXiv preprint
 617 arXiv:2305.14453*, 2023.

618 Mingze Ni, Zhensu Sun, and Wei Liu. Fraud’s bargain attack: Generating adversarial text samples
 619 via word manipulation process. *IEEE Transactions on Knowledge and Data Engineering*, 2024.

620 Anselm Paulus, Arman Zharmagambetov, Chuan Guo, Brandon Amos, and Yuandong Tian. Ad-
 621 vprompter: Fast adaptive adversarial prompting for llms. *arXiv preprint arXiv:2404.16873*, 2024.

622 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
 623 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

624 N Reimers. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint
 625 arXiv:1908.10084*, 2019.

626 Michael Robinson, Sourya Dey, and Shauna Sweet. The structure of the token space for large
 627 language models. *arXiv preprint arXiv:2410.08993*, 2024.

628 Patrick Rubin-Delanchy. Manifold structure in graph embeddings. *Advances in neural information
 629 processing systems*, 33:11687–11699, 2020.

630 V Sanh. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint
 631 arXiv:1910.01108*, 2019.

632 Erfan Shayegani, Md Abdullah Al Mamun, Yu Fu, Pedram Zaree, Yue Dong, and Nael Abu-
 633 Ghazaleh. Survey of vulnerabilities in large language models revealed by adversarial attacks.
 634 *arXiv preprint arXiv:2310.10844*, 2023.

635 Daniel Spielman and Nikhil Srivastava. Graph Sparsification by Effective Resistances. *SIAM Jour-
 636 nal on Computing*, 40(6):1913–1926, 2011.

637 Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. In *Proceed-
 638 ings of the fortieth annual ACM symposium on Theory of computing*, pp. 563–568, 2008.

648 Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on*
 649 *Computing*, 40(4):981–1025, 2011.

650

651 Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. How to fine-tune bert for text classification? In
 652 *Chinese computational linguistics: 18th China national conference, CCL 2019, Kunming, China,*
 653 *October 18–20, 2019, proceedings 18*, pp. 194–206. Springer, 2019.

654

655 Lichao Sun, Kazuma Hashimoto, Wenpeng Yin, Akari Asai, Jia Li, Philip Yu, and Caiming Xiong.
 656 Adv-bert: Bert is not robust on misspellings! generating nature adversarial samples on bert. *arXiv*
 657 *preprint arXiv:2003.04985*, 2020.

658

659 I Tenney. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*, 2019.

660

661 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 662 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 663 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

664

665 Elena Voita, Rico Sennrich, and Ivan Titov. The bottom-up evolution of representations in the trans-
 666 former: A study with machine translation and language modeling objectives. In *Proceedings of*
 667 *the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th Inter-
 668 national Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4396–4406,
 669 2019.

670

671 Minh Vu and My T Thai. Pgm-explainer: Probabilistic graphical model explanations for graph
 672 neural networks. *Advances in neural information processing systems*, 33:12225–12235, 2020.

673

674 Alex Wang. Glue: A multi-task benchmark and analysis platform for natural language understand-
 675 ing. *arXiv preprint arXiv:1804.07461*, 2018.

676

677 Boxin Wang, Chejian Xu, Shuohang Wang, Zhe Gan, Yu Cheng, Jianfeng Gao, Ahmed Hassan
 678 Awadallah, and Bo Li. Adversarial glue: A multi-task benchmark for robustness evaluation of
 679 language models. *arXiv preprint arXiv:2111.02840*, 2021.

680

681 Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu,
 682 Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. Decodingtrust: A comprehensive assessment of
 683 trustworthiness in gpt models. In *NeurIPS*, 2023.

684

685 Jindong Wang, HU Xixu, Wenxin Hou, Hao Chen, Runkai Zheng, Yidong Wang, Linyi Yang, Wei
 686 Ye, Haojun Huang, Xiubo Geng, et al. On the robustness of chatgpt: An adversarial and out-
 687 of-distribution perspective. In *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale*
 688 *Machine Learning Models*.

689

690 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 691 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 692 *neural information processing systems*, 35:24824–24837, 2022.

693

694 Thomas Wolf. Transformers: State-of-the-art natural language processing. *arXiv preprint*
 695 *arXiv:1910.03771*, 2020.

696

697 Sibo Yi, Yule Liu, Zhen Sun, Tianshuo Cong, Xinlei He, Jiaxing Song, Ke Xu, and Qi Li. Jailbreak
 698 attacks and defenses against large language models: A survey. *arXiv preprint arXiv:2407.04295*,
 699 2024.

700

701 Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan.
 702 Theoretically principled trade-off between robustness and accuracy. In *International conference*
 703 *on machine learning*, pp. 7472–7482. PMLR, 2019a.

704

705 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluat-
 706 ing text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019b.

707

708 Yichu Zhou and Vivek Srikumar. A closer look at how fine-tuning changes bert. *arXiv preprint*
 709 *arXiv:2106.14282*, 2021.

702 Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. Freelb: Enhanced
 703 adversarial training for natural language understanding. In *International Conference on Learning
 704 Representations*.

705 Kaijie Zhu, Xixu Hu, Jindong Wang, Xing Xie, and Ge Yang. Improving generalization of ad-
 706 versarial training via robust critical fine-tuning. In *Proceedings of the IEEE/CVF International
 707 Conference on Computer Vision*, pp. 4424–4434, 2023a.

708 Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang,
 709 Wei Ye, Yue Zhang, Neil Gong, et al. Promptrobust: Towards evaluating the robustness of large
 710 language models on adversarial prompts. In *Proceedings of the 1st ACM Workshop on Large AI
 711 Systems and Models with Privacy and Safety Analysis*, pp. 57–68, 2023b.

712 Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson.
 713 Universal and transferable adversarial attacks on aligned language models. *arXiv preprint
 714 arXiv:2307.15043*, 2023.

718 A APPENDIX

719 A.1 DETERMINISTIC HIDDEN-STATE EMBEDDINGS

720 Though modern transformers can produce *stochastic* outputs at the token level (e.g., due to beam
 721 search, random sampling, or dropout), their internal hidden states can remain *largely deterministic*
 722 under fixed conditions (Wolf, 2020). Below, we validate this claim by comparing token-level em-
 723 beddings and pooled hidden-state embeddings across different decoding strategies. We then observe
 724 what happens when we additionally fix the seed.

725 **Token vs. Pooled Embeddings Under Varying Seeds.** We feed the *same* input sequence through
 726 various transformers (DistilBERT, BERT, RoBERTa, and Google-Electra), each time *without* en-
 727 forcing a fixed random seed for decoding. We then collect:

- 728 • Token Embeddings. The final output embeddings for each token in the decoded sequence
 729 (i.e., after language modeling head).
- 730 • Pooled MHSA Output Embeddings. Our approach aggregates multiple attention heads and
 731 pools them into a single output vector per sequence, thereby abstracting away token-level
 732 variations.

733 For each model, we compute the cosine similarity between embeddings arising from different de-
 734 coding runs of the *same* input. Table 8 shows representative results for three datasets: SQuAD,
 735 IMDB, and AG-News.

736 Table 8: Cosine Similarity of Embeddings Across Different Decoding Runs *Without* a Fixed Random
 737 Seed. Higher is more stable.

744 Token Embeddings				
745 Dataset	746 DistilBERT	747 BERT	748 RoBERTa	749 Google-Electra
750 SQuAD	0.9338	0.2302	0.9967	0.1859
751 IMDB	0.9685	0.4443	0.9977	0.5131
752 AG-News	0.9450	0.5967	0.9954	0.5771

750 Pooled MHSA Output Embeddings				
751 Dataset	752 DistilBERT	753 BERT	754 RoBERTa	755 Google-Electra
756 SQuAD	1.00	1.00	1.00	1.00
757 IMDB	1.00	1.00	1.00	1.00
758 AG-News	0.99	0.99	0.99	0.99

Even when seeds vary, **token-level embeddings** exhibit inconsistent cosine similarity across runs (e.g., BERT scoring only 0.23 for SQuAD). By contrast, **our pooled MHSA method** maintains consistently high similarity (0.99 or 1.00), indicating a more stable representation that does not fluctuate with token-level decoding choices. This stability suggests that the representation potentially captures a more consistent global semantic space Reimers (2019).

Fixed Seed + Pooled MHSA Finally, we fix the random seed (and disable dropout) for all runs using our pooled MHSA approach, ensuring the only factor causing embedding changes is an explicit *input* perturbation (e.g., synonyms swapped). Under *identical* inputs and the same seed, the pooled MHSA output embeddings always match exactly (cosine similarity = 1.00), regardless of how tokens might be sampled. As summarized in Table 9, **all entries become** 1.00 when there is *no input perturbation*.

Table 9: Cosine Similarity with Fixed Seed *and* Pooled MHSA. Identical inputs yield identical embeddings (similarity = 1.00).

Seed-Fixed Pooled MHSA Output				
Dataset	DistilBERT	BERT	RoBERTa	Google-Electra
SQuAD	1.00	1.00	1.00	1.00
IMDB	1.00	1.00	1.00	1.00
AG-News	1.00	1.00	1.00	1.00

Statistical Reliability Analysis We consistently fixed random seeds to minimize variability and ensure reproducibility. To further quantify the statistical stability of SALMAN rankings, we conducted multi-run experiments on the SST-2 dataset (over 69k samples, the largest dataset in our study). The results reveal a very high degree of consistency, with a Top-20% overlapping rate of $99.0\% \pm 1.2\%$, underscoring the statistical reliability of the SALMAN ranking. Moreover, we repeated the top-1% non-robustness ranking five times and evaluated the subsets using AutoDAN (under a time-limited setting). The attack success rates were 80%, 60%, 60%, 60%, and 60%, respectively. These results provide additional empirical evidence that SALMAN rankings are not only statistically stable but also practically reliable for guiding adversarial attacks.

A.2 PROOF FOR THEOREM 3.1

We now show that *maximizing* the objective

$$\max_{\Theta} F(\Theta) = \log \det(\Theta) - \frac{1}{k} \text{Tr}(X^\top \Theta X), \quad (5)$$

where $\Theta = \mathcal{L} + \frac{1}{\sigma^2} I$, can be achieved by removing (or down-weighting) edges whose *distance ratio* is small. In essence, these low-ratio edges contribute less to $\log \det(\Theta)$ while incurring a larger penalty in the trace term, so pruning them increases $F(\Theta)$.

1. Decomposing the Objective. Writing $\mathcal{L} = \sum_{(p,q) \in E} w_{p,q} e_{p,q} e_{p,q}^\top$, we split $F(\Theta)$ into two terms:

$$F(\Theta) = F_1(\Theta) - \frac{1}{k} F_2(\Theta), \quad \text{where}$$

$$F_1(\Theta) = \log \det(\Theta), \quad F_2(\Theta) = \text{Tr}(X^\top \Theta X).$$

Since $\Theta = \mathcal{L} + \frac{1}{\sigma^2} I$, each edge weight $w_{p,q}$ appears explicitly in \mathcal{L} .

2. Gradient with Respect to an Edge Weight. To optimize $F(\Theta)$ w.r.t. a single edge weight $w_{p,q}$:

- **Term $F_1(\Theta)$:** Let λ_i be the i -th eigenvalue of \mathcal{L} , and v_i its eigenvector. Then

$$\frac{\partial}{\partial w_{p,q}} \left(\log \det(\Theta) \right) = \frac{\partial}{\partial w_{p,q}} \left[\log \det \left(\mathcal{L} + \frac{1}{\sigma^2} I \right) \right].$$

810 By standard matrix calculus, this derivative can be linked to the *effective resistance distance*
 811 $d^{\text{eff}}(p, q)$:

$$\frac{\partial F_1}{\partial w_{p,q}} \approx d^{\text{eff}}(p, q),$$

812 where $d^{\text{eff}}(p, q)$ encapsulates how strongly edge (p, q) influences $\log \det(\Theta)$.

- 813 • **Term $F_2(\Theta)$:**

$$814 F_2(\Theta) = \text{Tr}(X^\top \Theta X) = \text{Tr}(X^\top (\mathcal{L} + \frac{1}{\sigma^2} I) X) = \frac{\text{Tr}(X^\top X)}{\sigma^2} + \sum_{(p,q) \in E} w_{p,q} \|X^\top e_{p,q}\|_2^2.$$

815 Since $\|X^\top e_{p,q}\|_2^2 = \|X_p - X_q\|_2^2 = d^{\text{dat}}(p, q)$, we have

$$\frac{\partial F_2}{\partial w_{p,q}} = \|X_p - X_q\|_2^2 = d^{\text{dat}}(p, q).$$

816 Furthermore, $d^{\text{dat}}(p, q) = \frac{1}{w_{p,q}}$, which implies

$$\frac{\partial F_2}{\partial w_{p,q}} = \frac{1}{w_{p,q}}.$$

817 Hence, the derivative of $F(\Theta) = F_1 - \frac{1}{k} F_2$ w.r.t. $w_{p,q}$ is

$$\frac{\partial F}{\partial w_{p,q}} = d^{\text{eff}}(p, q) - \frac{1}{k} \frac{1}{w_{p,q}}. \quad (6)$$

818 3. Distance Ratio and Pruning Condition.

Rewriting Equation equation 6:

$$819 d^{\text{eff}}(p, q) - \frac{1}{k} \frac{1}{w_{p,q}} = 0 \iff d^{\text{eff}}(p, q) = \frac{1}{k} \frac{1}{w_{p,q}}.$$

820 Define the *distance ratio* for edge (p, q) :

$$821 \rho_{p,q} = \frac{d^{\text{eff}}(p, q)}{d^{\text{dat}}(p, q)} = w_{p,q} (d^{\text{eff}}(p, q)).$$

822 When $d^{\text{eff}}(p, q)$ is relatively large compared to $\frac{1}{w_{p,q}}$, we have $\rho_{p,q}$ large, indicating an important
 823 edge for $\log \det(\Theta)$. Conversely, if $\rho_{p,q}$ is small, the edge (p, q) contributes little to $F_1(\Theta)$ but
 824 increases $F_2(\Theta)$, thereby reducing $F(\Theta)$.

825 4. Conclusion: Prune Low-Ratio Edges.

826 Thus, maximizing equation 5 naturally leads to *removing or down-weighting* edges whose ratio

$$827 \rho_{p,q} = \frac{d^{\text{eff}}(p, q)}{d^{\text{dat}}(p, q)}$$

828 is below a certain threshold. By pruning these edges, we preserve the essential spectral structure
 829 needed to keep $\log \det(\Theta)$ high (reflecting higher effective resistance) while mitigating the penalty
 830 in $\text{Tr}(X^\top \Theta X)$ from edges that have large data distance but small effective resistance. In other
 831 words, *edges with large $\rho_{p,q}$ stay*, and edges with small $\rho_{p,q}$ are pruned, thereby maximizing $F(\Theta)$
 832 and maintaining the key (Laplacian) properties of the original graph.

833 A.3 PROOF OF LRD DECOMPOSITION FOR EFFICIENT EDGE RESISTANCE COMPUTATION 834 AND WEIGHTED GRAPH SPARSIFICATION

835 In this section, we establish that the *low-resistance-diameter* (LRD) decomposition scheme can
 836 efficiently approximate the effective resistance for each edge in a *weighted* graph and thus provide
 837 an effective path toward spectral sparsification. Our argument proceeds in two stages:

- 838 1. **Approximate Effective Resistance via Krylov Subspaces:** We show how the iterative
 839 procedure yields reliable estimates of $d^{\text{eff}}(p, q)$ in near-linear time.
- 840 2. **Bound Cycle Lengths under LRD Clustering:** We explain how the multilevel contraction
 841 and supernode formation ensure that edges with large resistance distances are effectively
 842 sampled or retained, yielding a final sparsified graph that spectrally approximates the orig-
 843 inal.

864 **Stage 1: Approximating Effective Resistances via Krylov Subspaces.** The resistance distance
 865 for an edge (p, q) in a graph $G = (V, E)$ with Laplacian L_G can be expressed as:
 866

$$867 \quad d^{\text{eff}}(p, q) = \sum_{i=2}^N \frac{(u_i^\top e_{p,q})^2}{u_i^\top L_G u_i},$$

$$868$$

$$869$$

870 where u_2, \dots, u_N are the (nontrivial) eigenvectors of L_G and $e_{p,q} = e_p - e_q$. Directly computing
 871 all eigenvalues/eigenvectors for large G is typically prohibitive. Instead, we replace these eigenvec-
 872 tors with a small set of Krylov basis vectors $x^{(1)}, x^{(2)}, \dots, x^{(m)}$, which approximate the subspace
 873 spanned by the top spectral components of L_G . Specifically, each $x^{(i)}$ is drawn from

$$874 \quad \kappa_m(A, c) = \text{span}\{c, Ac, A^2c, \dots, A^{m-1}c\},$$

$$875$$

876 where A is the adjacency matrix and c is a random vector. Orthogonalizing and normalizing these
 877 m vectors ensures a concise basis in which to project $e_{p,q}$.

878 **Lemma A.1** (Krylov Approximation of Effective Resistance). *Suppose $x^{(1)}, \dots, x^{(m)}$ are m or-
 879 thonormal vectors approximating the dominant spectral subspace of L_G (via a Krylov process).
 880 Then for any edge $(p, q) \in E$,*

$$881 \quad d^{\text{eff}}(p, q) \approx \sum_{i=1}^m \frac{(x^{(i)\top} e_{p,q})^2}{x^{(i)\top} L_G x^{(i)}}.$$

$$882$$

$$883$$

$$884$$

885 Choosing $m = \tilde{O}(\log N)$ and updating each level in near-linear time yields high-probability error
 886 bounds comparable to exact spectral decompositions (Spielman & Srivastava, 2011; Koutis et al.,
 887 2010).

888 **Stage 2: LRD-based Short-Cycle Decomposition for Weighted Graphs.** The second key step is
 889 the *multilevel* contraction scheme that ensures edges with large effective resistance remain “visible”
 890 at higher levels, while short cycles (or low-resistance edges) are contracted to form supernodes.
 891 Specifically:

- 893 • At level δ , each edge (p, q) is either *contracted* (if $d_{\text{eff}}^{(\delta)}(p, q)$ is below the chosen thresh-
 894 old) or *retained* (if $d_{\text{eff}}^{(\delta)}(p, q)$ is above the threshold). Contraction merges p and q into a
 895 supernode ϑ , assigning it an accumulated weight η_ϑ via:

$$896 \quad \eta_\vartheta := \eta(p^{(\delta)}) + \eta(q^{(\delta)}) + d_{\text{eff}}^{(\delta)}(p, q). \quad (7)$$

$$897$$

$$898$$

- 899 • As edges are contracted, any cycles formed at level δ have length bounded by the effective-
 900 resistance diameter. Consequently, short cycles in the *weighted* setting are handled sim-
 901 ilarly to Chu et al. (2020)’s unweighted approach, except that we measure distances via
 902 d_{eff} , not just hop counts.
- 903 • After finalizing the clusters (supernodes), the “inter-cluster” edges (those bridging different
 904 clusters) are preserved or upweighted in the sparsified graph. These edges typically have
 905 higher $d_{\text{eff}}(p, q)$ and thus significantly impact spectral properties of L_G .

907 Formally, let L_H denote the Laplacian of the sparsified graph H returned by the LRD decomposition.
 908 We say H is a $(1 \pm \varepsilon)$ -spectral-approximation of G if, for all $x \in \mathbb{R}^N$,

$$909 \quad (1 - \varepsilon) x^\top L_G x \leq x^\top L_H x \leq (1 + \varepsilon) x^\top L_G x.$$

$$910$$

911 Standard arguments from spectral sparsification (Spielman & Srivastava, 2011) show that any pro-
 912 cedure ensuring accurate effective-resistance estimates can preserve the graph’s quadratic form up
 913 to $(1 \pm \varepsilon)$ factors. The main difference in *our* approach is the use of low-resistance-diameter cycles
 914 instead of purely unweighted short cycles.

915 **Theorem A.2** (LRD for Weighted Graph Sparsification). *Let $G = (V, E)$ be a connected weighted
 916 graph with N nodes and M edges, and let $0 < \varepsilon < 1$ be a chosen approximation factor. Then,
 917 by applying the LRD-based spectral sparsification algorithm with Krylov-based effective-resistance
 918 estimates:*

918 1. We can approximate $d_{\text{eff}}(p, q)$ for all edges $(p, q) \in E$ in near-linear time (Lemma A.1).
 919
 920 2. We contract short cycles (below a chosen d_{eff} threshold) and preserve inter-cluster edges
 921 with high $d_{\text{eff}}(p, q)$, forming a sparsified graph H with Laplacian L_H .
 922
 923 3. With high probability, H satisfies $(1 - \varepsilon) x^\top L_G x \leq x^\top L_H x \leq (1 + \varepsilon) x^\top L_G x$, $\forall x \in$
 924 \mathbb{R}^N .

925 Hence, H serves as a $(1 \pm \varepsilon)$ -spectral-approximation to G , yielding a low-complexity graph that
 926 closely preserves the original graph's spectral (and thus effective-resistance) structure.

927 *Proof Sketch.*

928 (1) *Effective-resistance approximation.* By Lemma A.1, each edge's resistance can be estimated
 929 via $m = \tilde{O}(\log N)$ Krylov vectors per level of the hierarchy. Summed over δ levels, the total cost
 930 remains near-linear in $N + M$ (plus polylogarithmic factors), similar to Koutis et al. (2010); Kyng
 931 & Sachdeva (2016).

932 (2) *Cycle decomposition.* Following Chu et al. (2020), short cycles are identified and contracted; we
 933 adapt the criteria to *resistance distances* in lieu of unweighted hop distances. The LRD threshold
 934 ensures each supernode aggregates edges that have sufficiently low $d_{\text{eff}}(p, q)$, while edges with
 935 higher $d_{\text{eff}}(p, q)$ remain across clusters and are re-inserted (or re-weighted) in the final sparsified
 936 graph H .

937 (3) *Spectral approximation.* Standard spectral graph theory arguments Spielman & Srivastava (2011)
 938 show that removing or down-weighting edges of low effective resistance induces little change in
 939 $x^\top L_G x$ for all x . Conversely, preserving edges with large $d_{\text{eff}}(p, q)$ is crucial for maintaining the
 940 spectral signature of L_G . The result is a $(1 \pm \varepsilon)$ -approximation for sufficiently small ε .

941 Thus, LRD-based cycle decomposition extends short-cycle approaches to weighted graphs by an-
 942 choring cycle lengths in *resistance* metrics. This achieves the final $(1 \pm \varepsilon)$ spectral-approximation
 943 for G in near-linear time. \square

944 A.4 PROOF FOR THE RELATIONSHIP BETWEEN γ_{\min}^F AND $\lambda_{\max}(L_X^+ L_Y^+)$

945 **Definition A.3** (Minimum Distance Mapping Distortion). Analogous to the definition of γ_{\max}^F , we
 946 define

$$947 \gamma_{\min}^F = \min_{\substack{p, q \in V \\ p \neq q}} \frac{d_Y(p, q)}{d_X(p, q)} = \min_{\substack{p, q \in V \\ p \neq q}} \frac{e_{p,q}^\top L_Y^+ e_{p,q}}{e_{p,q}^\top L_X^+ e_{p,q}},$$

948 where L_X^+ and L_Y^+ are the Moore–Penrose pseudoinverses of the Laplacian matrices L_X and L_Y ,
 949 respectively. This quantity reflects the *smallest* ratio of output distance to input distance, charac-
 950 terizing how close points in the *output* manifold might originate from distant points in the *input*
 951 manifold.

952 *Proof.* Recall that

$$953 \gamma_{\min}^F = \min_{\substack{p, q \in V \\ p \neq q}} \frac{e_{p,q}^\top L_Y^+ e_{p,q}}{e_{p,q}^\top L_X^+ e_{p,q}}.$$

954 If we define $v = e_{p,q}$ and restrict v such that $v^\top \mathbf{1} = 0$ ($\mathbf{1}$ being the all-ones vector, ensuring we stay
 955 within the subspace on which the Laplacian pseudoinverses are invertible), then

$$956 \gamma_{\min}^F \geq \min_{\substack{\|v\| \neq 0 \\ v^\top \mathbf{1} = 0}} \frac{v^\top L_Y^+ v}{v^\top L_X^+ v}.$$

957 By the (min-)max version of the generalized Courant-Fischer theorem (applied to positive semide-
 958 finite matrices on the subspace orthogonal to $\mathbf{1}$), we have

$$959 \min_{\substack{\|v\| \neq 0 \\ v^\top \mathbf{1} = 0}} \frac{v^\top L_Y^+ v}{v^\top L_X^+ v} = \lambda_{\min}(L_Y^+ L_X).$$

972 However, because $L_Y^+ L_X$ is invertible on that same subspace, one also obtains the relationship
 973

$$974 \lambda_{\min}(L_Y^+ L_X) = \frac{1}{\lambda_{\max}((L_Y^+ L_X)^{-1})}.$$

$$975$$

976 Next, it can be shown that

$$977 (L_Y^+ L_X)^{-1} = L_X^+ L_Y \quad (\text{on the subspace } v^\top \mathbf{1} = 0),$$

$$978$$

979 which gives

$$980 \lambda_{\min}(L_Y^+ L_X) = \frac{1}{\lambda_{\max}(L_X^+ L_Y)}.$$

$$981$$

Combining these steps, we conclude

$$982 \gamma_{\min}^F \geq \lambda_{\min}(L_Y^+ L_X) = \frac{1}{\lambda_{\max}(L_X^+ L_Y)}.$$

$$983$$

984 This completes the proof. \square

$$985$$

986 A.5 SALMAN SCORE AND CORRESPONDING PROOFS

$$987$$

988 Specifically, we first compute the weighted eigensubspace matrix $V_r \in \mathbb{R}^{N \times r}$ for spectral embedding
 989 on G_X with N nodes:

$$990 V_r \stackrel{\text{def}}{=} \left[v_1 \sqrt{\lambda_1}, \dots, v_r \sqrt{\lambda_r} \right], \quad (8)$$

$$991$$

992 where $\lambda_1, \lambda_2, \dots, \lambda_r$ represent the first r largest eigenvalues of $L_Y^+ L_X$ and v_1, v_2, \dots, v_r are the corresponding
 993 eigenvectors. Let u_1, u_2, \dots, u_N denote the N eigenvectors of $L_X L_Y^+$, respectively, while
 994 their corresponding eigenvalues are shared. In addition, eigenvectors u_i can be constructed to satisfy:

$$995 u_i^\top L_X^+ u_j = \begin{cases} 1, & i = j \\ 0, & i \neq j. \end{cases} \quad (9)$$

$$996$$

$$997 \Rightarrow u_i^\top L_Y^+ u_j = \begin{cases} \lambda_i, & i = j \\ 0, & i \neq j. \end{cases} \quad (10)$$

$$998$$

999 Therefore, the following equations hold:

$$1000 L_Y^+ u_i = \lambda_i L_X^+ u_i \Leftrightarrow L_Y^+ L_X (L_Y^+ u_i) = \lambda_i (L_Y^+ u_i) \quad (11)$$

$$1001 L_X v_i = \lambda_i L_Y v_i \Leftrightarrow L_Y^+ L_X v_i = \lambda_i v_i$$

$$1002$$

1003 which leads to the following equation

$$1004 v_i = \beta_i L_Y^+ u_i$$

$$1005 \Rightarrow u_j^\top v_i = \begin{cases} \beta_i \lambda_i, & i = j \\ 0, & i \neq j. \end{cases} \quad (12)$$

$$1006$$

1007 where β_i denotes a scaling coefficient. Without loss of generality, $e_{p,q}$ can be expressed as a linear
 1008 combination of u_i for $i = 1, \dots, N$ as follows:

$$1009 e_{p,q} = \sum_{i=1}^N \alpha_i u_i. \quad (13)$$

$$1010$$

1011 Then $\gamma^F(p, q)$ can be rewritten as follows:

$$1012$$

$$1013 \gamma^F(p, q) = \frac{d_Y(p, q)}{d_X(p, q)} = \frac{e_{p,q}^\top L_Y^+ e_{p,q}}{e_{p,q}^\top L_X^+ e_{p,q}}$$

$$1014 = \frac{(\sum_{i=1}^N \alpha_i u_i)^\top L_Y^+ (\sum_{i=1}^N \alpha_i u_i)}{(\sum_{i=1}^N \alpha_i u_i)^\top L_X^+ (\sum_{i=1}^N \alpha_i u_i)}$$

$$1015 = \frac{\sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j u_i^\top L_Y^+ u_j}{\sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j u_i^\top L_X^+ u_j} \quad (14)$$

$$1016$$

$$1017 = \frac{\sum_{i=1}^N \alpha_i^2 u_i^\top L_Y^+ u_i}{\sum_{i=1}^N \alpha_i^2 u_i^\top L_X^+ u_i}$$

$$1018$$

$$1019 = \frac{\sum_{i=1}^N \alpha_i^2 \lambda_i}{\sum_{i=1}^N \alpha_i^2}.$$

$$1020$$

$$1021$$

$$1022$$

$$1023$$

$$1024$$

$$1025$$

If the edge (p, q) is dominantly aligned with a single dominant generalized eigenvector u_k where $1 \leq k \leq r$, it implies $\forall i \neq k, \alpha_i \approx 0$ and thus $e_{p,q} \approx \alpha_k u_k$. Then:

$$\gamma^F(p, q) \approx \lambda_k. \quad (15)$$

With $\|V_r^\top e_{p,q}\|_2^2$, We have:

$$\begin{aligned} \|V_r^\top e_{p,q}\|_2^2 &= \sum_{i=1}^r \lambda_i (v_i^\top e_{p,q})^2 \\ &= \sum_{i=1}^r \lambda_i \left(\sum_{j=1}^N \alpha_j \beta_i u_j^\top L_Y^+ u_i \right)^2 \\ &= \sum_{i=1}^r \alpha_i^2 \beta_i^2 \lambda_i^3 \\ &\approx \alpha_k^2 \beta_k^2 \lambda_k^3 \propto (\gamma^F(p, q))^3 \end{aligned} \quad (16)$$

Consider $L_X^+ L_Y$ whose eigenvalues we denote by μ_1, \dots, μ_N with corresponding eigenvectors w_1, \dots, w_N . we then compute the weighted eigensubspace matrix $W_r \in \mathbb{R}^{N \times r}$ for spectral embedding on G_Y with N nodes:

$$W_r \stackrel{\text{def}}{=} [w_1 \sqrt{\mu_1}, \dots, w_r \sqrt{\mu_r}], \quad (17)$$

Because $L_X^+ L_Y$ has eigenvalues $\mu_i = 1/\lambda_i$, and its eigenvectors w_i correspond in a reciprocal way, one obtains a parallel statement. In particular:

- If $e_{p,q}$ aligns chiefly with the eigenvector w_k of $L_X^+ L_Y$ having eigenvalue $\mu_k = 1/\lambda_k$,
- Then $\gamma^F(p, q) \approx \lambda_k$ as before.

A similar calculation to the Equation 16 proof now yields

$$\|W_r^\top e_{p,q}\|_2^2 = \sum_{i=1}^r \mu_i (w_i^\top e_{p,q})^2 \approx \text{const} \times (\mu_k)^3 = \text{const} \times \left(\frac{1}{\lambda_k}\right)^3.$$

Since $\lambda_k \approx \gamma^F(p, q)$, we conclude

$$\|W_r^\top e_{p,q}\|_2^2 \propto \gamma^F(p, q)^{-3}.$$

Hence, $\|W_r^\top e_{p,q}\|_2^2 + \|V_r^\top e_{p,q}\|_2^2 \propto \gamma^F(p, q)^3 + \gamma^F(p, q)^{-3}$.

A.6 EXPERIMENTAL SETUP

In this section, we provide details on the datasets, model configurations, training/finetuning protocols, and evaluation metrics used throughout our experiments. By clarifying each step, we ensure that our methodology is both *transparent* and *reproducible*.

Dataset. We evaluate on benchmark datasets such as SST-2, MNLI, RTE, QNLI, QQP, and COLA to cover diverse classification objectives (sentiment analysis, natural language inference, and question classification). Each dataset is split into training, validation, and test sets following standard protocols (e.g., the GLUE benchmark (Wang, 2018)). We tokenize inputs using the *default* subword tokenizer for each model (e.g., BERT’s WordPiece or RoBERTa’s Byte-Pair Encoding), lowercasing as necessary. For SALMAN-Guided Attack experiment, we use AdvBench Harmful Behaviors dataset. JailBreak does not involve a training process, thus we did not split the dataset. We directly ranked the entire dataset of 520 data points.

Language Model. We evaluate on several benchmark language models such as BERT-base-uncased Devlin (2018), RoBERTa-base Liu (2019), DistilBERT-base-uncased Sanh (2019), ALBERT-base-v2 Lan (2019), GPT-2 Radford et al. (2019), and LLaMA-7B-v2 Touvron et al. (2023).

1080
1081
1082
1083
1084 Table 10: Summary of hyperparameters (k is for kNN graph construction and SPF is for our low-
1085 resistance-diameter decomposition) used in our method for each (*model, attack*) configuration. DIS
1086 refers to random selection of deletion, insertion, or swap.

Model	Attack	k	SPF
<i>SST-2</i>			
BERT-base-uncased	DIS	30	2
RoBERTa-base	DIS	30	2
DistilBERT-base-uncased	DIS	30	2
ALBERT-base-v2	DIS	30	2
GPT-2	DIS	10	2
LLaMA-7B-v2	DIS	10	2
GPT-2	spaCy	20	2
LLaMA-7B-v2	spaCy	30	3
GPT-2	TextAttack	10	2
LLaMA-7B-v2	TextAttack	10	2
<i>MNLI</i>			
BERT-base-uncased	DIS	30	2
RoBERTa-base	DIS	30	2
DistilBERT-base-uncased	DIS	30	2
ALBERT-base-v2	DIS	30	2
GPT-2	DIS	50	2
LLaMA-7B-v2	DIS	10	2
GPT-2	spaCy	70	3
LLaMA-7B-v2	spaCy	70	2
GPT-2	TextAttack	20	3
LLaMA-7B-v2	TextAttack	10	2

1107
1108
1109
1110
1111 **Hyperparameter settings.** We obtain *Distance Mapping Distortion* scores for each sample by
1112 comparing input and output manifold distances (e.g., from \mathbf{z}_X to \mathbf{z}_Y). Summary of hyperparameters
1113 during DMD calculation is in Table 10. To gauge how much finetuned models deviate from their
1114 pretrained checkpoints, we reference layer-wise similarity metrics such as CKA and STIR (Neerudu
1115 et al., 2023).

1116
1117
1118 **k-NN Ablation on SST-2 (GPT-2).** We also evaluate the sensitivity of SALMAN to the choice of
1119 k in k -NN graph construction. Specifically, we vary $k \in \{15, 20, 30\}$ for GPT-2 on SST-2 and
1120 compare the *Kullback–Leibler Divergence* (KLD) and *BERTScore* (Precision, Recall, F_1) for non-
1121 robust (NR) vs. robust (R) samples. As shown in Table 11, increasing k does not drastically alter the
1122 distinction between robust and non-robust data; the non-robust subsets consistently exhibit higher
1123 KLD and slightly lower BERTScores, while robust subsets remain more stable under perturbations.
1124 This indicates that SALMAN is relatively insensitive to moderate changes in k .

1125
1126 Table 11: Effect of varying k in k -NN on robustness and similarity metrics (GPT-2, SST-2).

k (kNN)	KLD (NR)	KLD (R)	Precision (NR)	Recall (NR)	F_1 (NR)	Precision (R)	Recall (R)	F_1 (R)
15	0.1110	0.0000	0.9972	0.9978	0.9975	0.9990	0.9988	0.9989
20	0.0988	0.0003	0.9973	0.9978	0.9976	0.9993	0.9992	0.9992
30	0.1923	0.0000	0.9962	0.9970	0.9966	0.9992	0.9992	0.9992

1125 1126 A.7 LAYER-WISE STIR AND CKA RESULTS

1127
1128 In addition to the layer-wise comparison between normal and guided fine-tuning shown in Table 12
1129 (CoLA dataset), we replicate the same analysis for the **SST-2** and **RTE** tasks under GPT-2. Following
1130 the exact protocol of Section 4.3 and Neerudu et al. (2023), we assign higher training weights to
1131 non-robust data (determined by our DMD ranking) and lower weights to robust data. As before, we
1132 measure:

1133
1134 • **Validation Accuracy** on the downstream task,

1134 Table 12: Layer-wise STIR and CKA for GPT-2 on CoLA. Each cell shows the “Normal fine-tuning
 1135 / SALMAN-guided fine-tuning”, rounded to four decimal places. Normal fine-tuning validation
 1136 accuracy is 0.7468, SALMAN-guided fine-tuning validation accuracy is 0.7709. m_1 is the pre-
 1137 trained model and m_2 is the fine-tuned model. Better results are in **bold**.

1138

Layer	STIR(m_2m_1)	STIR(m_1m_2)	CKA
0	0.9623 / 0.9623	0.9623 / 0.9623	1.0000 / 1.0000
1	0.9070 / 0.9073	0.9065 / 0.9053	0.9987 / 0.9987
2	0.9688 / 0.9691	0.9690 / 0.9711	0.9936 / 0.9931
3	0.9848 / 0.9904	0.9678 / 0.9551	0.9856 / 0.9748
4	0.9853 / 0.9934	0.9690 / 0.9775	0.9836 / 0.9837
5	0.9904 / 0.9928	0.9750 / 0.9752	0.9906 / 0.9909
6	0.9897 / 0.9924	0.9697 / 0.9767	0.9920 / 0.9945
7	0.9895 / 0.9927	0.9724 / 0.9833	0.9931 / 0.9909
8	0.9860 / 0.9914	0.9680 / 0.9854	0.9923 / 0.9936
9	0.9825 / 0.9872	0.9666 / 0.9770	0.9905 / 0.9907
10	0.9776 / 0.9833	0.9647 / 0.9762	0.9917 / 0.9928
11	0.9730 / 0.9784	0.9628 / 0.9678	0.9893 / 0.9904
12	0.4691 / 0.7233	0.6819 / 0.7924	0.5612 / 0.7251

1151

1152

- **STIR** (Similar Token Identity Representation) metrics (m_2m_1 , m_1m_2) capturing how similar layer i in the fine-tuned model m_2 is to layer j in the pre-trained model m_1 ,
- **CKA** measuring layer-wise alignment between m_1 and m_2 embeddings.

1153

1154

1155

1156

1157

SST-2 Results. Table 14 shows GPT-2’s layer-wise STIR and CKA under normal vs. guided fine-tuning on SST-2. Both approaches yield similar *final accuracy* (0.9231 vs. 0.9232), yet the guided variant consistently achieves higher STIR/CKA scores in later layers. In particular, layer 12 sees a substantial jump in STIR(m_2m_1) from 0.0533 to 0.0867 and CKA from 0.1459 to 0.2039, indicating closer alignment to the pre-trained checkpoint.

1162

1163

1164

1165

1166

RTE Results. In Table 13, we compare normal vs. guided fine-tuning for GPT-2 on the RTE dataset. While both runs converge similarly in accuracy (not shown here to save space), the guided approach again shows improved STIR and CKA alignment with the pre-trained checkpoint. For instance, layer 12 sees an increase from 0.2858 to 0.3393 in STIR(m_2m_1) and from 0.3458 to 0.3476 in CKA.

1167

1168

1169

Table 13: Layer-wise STIR and CKA for GPT-2 on **RTE** under Normal vs. Guided fine-tuning. Each cell shows “Normal Fine-tuning / SALMAN-guided Fine-tuning”, rounded to four decimal places. Better results in **bold**.

1171

1172

Layer	STIR(m_2m_1)	STIR(m_1m_2)	CKA
0	0.9913 / 0.9913	0.9914 / 0.9914	1.0000 / 1.0000
1	0.9786 / 0.9791	0.9776 / 0.9779	0.9986 / 0.9977
2	0.9859 / 0.9857	0.9859 / 0.9852	0.9976 / 0.9990
3	0.9903 / 0.9920	0.9917 / 0.9918	0.9951 / 0.9987
4	0.9897 / 0.9902	0.9897 / 0.9900	0.9885 / 0.9963
5	0.9898 / 0.9908	0.9891 / 0.9916	0.9894 / 0.9981
6	0.9865 / 0.9872	0.9869 / 0.9886	0.9908 / 0.9923
7	0.9821 / 0.9829	0.9806 / 0.9839	0.9746 / 0.9801
8	0.9758 / 0.9781	0.9709 / 0.9761	0.9407 / 0.9500
9	0.9708 / 0.9724	0.9607 / 0.9696	0.9288 / 0.9492
10	0.9564 / 0.9601	0.9359 / 0.9507	0.9028 / 0.9347
11	0.9333 / 0.9331	0.9152 / 0.9265	0.9223 / 0.9390
12	0.2858 / 0.3393	0.6131 / 0.6203	0.3458 / 0.3476

1184

1185

1186

1187

Discussion. Similar to our observations on CoLA (Table 12), placing higher emphasis on non-robust data (i.e., higher DMD samples) preserves downstream performance while bringing the fine-tuned layers closer to the original pre-trained representations. These improvements in STIR and

1188 Table 14: Layer-wise STIR and CKA for GPT-2 on **SST-2** under Normal vs. Guided fine-tuning.
 1189 Each cell shows “Normal Fine-tuning / Guided Fine-tuning”, rounded to four decimal places. **Acc**
 1190 is the validation accuracy of each method. For STIR, $(m2m1)$ compares the fine-tuned model $m2$ to
 1191 the pre-trained model $m1$, and $(m1m2)$ is the reverse; CKA measures embedding similarity. Better
 1192 results are in **bold**.

1193

Validation Accuracy: Normal = 0.9231, SALMAN-guided = 0.9232			
Layer	STIR(m2m1)	STIR(m1m2)	CKA
0	0.9913 / 0.9913	0.9912 / 0.9912	1.0000 / 1.0000
1	0.9763 / 0.9771	0.9784 / 0.9787	0.9963 / 0.9974
2	0.9784 / 0.9789	0.9762 / 0.9767	0.9971 / 0.9973
3	0.9703 / 0.9713	0.9366 / 0.9414	0.9460 / 0.9469
4	0.9661 / 0.9715	0.9549 / 0.9608	0.9773 / 0.9800
5	0.9736 / 0.9757	0.9469 / 0.9589	0.9705 / 0.9738
6	0.9649 / 0.9704	0.9343 / 0.9450	0.9568 / 0.9604
7	0.9618 / 0.9672	0.9389 / 0.9476	0.9642 / 0.9675
8	0.9663 / 0.9703	0.9514 / 0.9598	0.9800 / 0.9825
9	0.9435 / 0.9553	0.9473 / 0.9545	0.9717 / 0.9787
10	0.9230 / 0.9504	0.9486 / 0.9573	0.9599 / 0.9774
11	0.8567 / 0.9208	0.9328 / 0.9426	0.9166 / 0.9562
12	0.0533 / 0.0867	0.7504 / 0.7755	0.1459 / 0.2039

1207

1208

1209 CKA suggest *reduced representational drift*, consistent with the intuition that focusing on “hard”
 1210 samples forces the model to retain more generalizable features from pre-training (Cheng et al., 2021;
 1211 Zhu et al., 2023a).

1212

1213 Overall, these extended results on SST-2 and RTE corroborate our main findings: *robustness-guided*
 1214 *fine-tuning* effectively balances task performance with better alignment to the pre-trained checkpoint
 1215 across multiple datasets.

1216

1217

A.8 WEIGHTED FINE-TUNING AND INTEGRATION WITH ROSE

1218

1219

Motivation: Focus on Non-Robust Data. As discussed in Section 4.3, prior studies have shown
 1220 that directing more attention to non-robust (“hard”) samples during training can improve model
 1221 generalizability and resilience (Cheng et al., 2021; Zhu et al., 2023a). Our approach identifies these
 1222 difficult samples via the SALMAN-based ranking and then *assigns higher training weights* to them,
 1223 while simultaneously down-weighting samples that appear robust. We follow the finetuning
 1224 protocol of Neerudu et al. (2023), hypothesizing that emphasizing harder samples preserves more of
 1225 the pre-trained model’s versatility. This reduces the risk of overfitting to “easy” data and yields
 1226 representations closer to the original checkpoint (see STIR/CKA results in Appendix A.7).

1227

A.8.1 WEIGHTING SCHEMES FOR GUIDED FINE-TUNING

1228

1229

1230

1231

1232

Linear Schedule. We sort all training samples in descending order of their DMD values (highest
 1233 DMD = most non-robust), then map each sample to a weight $w \in [0, 1]$ proportional to its position
 1234 in this ranking. Concretely, if the highest-DMD sample is indexed as rank 0, it receives weight
 1235 ≈ 1.0 , whereas the lowest-DMD sample (rank $n-1$) receives weight near 0.0. Intermediate samples
 1236 smoothly interpolate between these extremes.

1237

1238

1239

1240

Combining with SOTA Robust Training (ROSE). We further integrate our DMD-based weighting
 1241 into **ROSE: Robust Selective Fine-tuning** (Jiang et al., 2022), which filters out spurious pa-
 1242 rameter updates by comparing dropout-induced distributions at each iteration:

$$L_{\text{KL}}^{(t)} = D_{\text{KL}}(P_t \parallel P'_t) + D_{\text{KL}}(P'_t \parallel P_t).$$

1243

1244

1245

1246

1247

ROSE removes parameter changes that inflate L_{KL} excessively, thus improving adversarial re-
 1248 silience.
 1249
Per-sample Weight $w(x)$ for Joint Optimization Alternatively, we employ a logistic transition-
 1250 based function partitioned into intervals:

- 1242 • **Top-25% Non-Robust** (i.e., highest DMD) can receive a *larger* weight (e.g., 2.0),
- 1243 • **Middle-Range** samples gradually decrease from 1.0 to 0.0 in stepwise logistic transitions,
- 1244 • **Bottom-5% Most Robust** eventually gets weight 0.0 (or near-zero).

1245
1246
1247 This piecewise approach allows a *finer distinction* between very hard, moderately challenging, and
1248 trivially easy samples.

1249 **Joint Optimization.** We incorporate our per-sample weight $w(x)$ into ROSE’s fine-tuning loss.
1250 Specifically, if the original ROSE objective is

$$1251 \quad \mathcal{L}_{\text{ROSE}}(\theta_t) = \mathbb{E}_{x \sim \mathcal{D}} [L_{\text{task}}(x, \theta_t) + \lambda L_{\text{KL}}^{(t)}],$$

1252
1253 then our *combined* objective is

$$1254 \quad \mathcal{L}_{\text{ROSE+Guided}}(\theta_t) = \mathbb{E}_{x \sim \mathcal{D}} [w(x) \cdot L_{\text{task}}(x, \theta_t) + \lambda L_{\text{KL}}^{(t)}].$$

1255 Hence, the model is “selective” not only at the *parameter* level (via L_{KL}) but also at the *sample*
1256 level (via DMD-based weighting).

1260 A.9 FROM PGM TO MANIFOLD: VALIDATING ON GRAPH BENCHMARKS

1261 Although our primary interest is applying the PGM-based manifold to NLP data (where nodes rep-
1262 resent text embeddings), we first validate how well our spectral sparsification and resistance distance
1263 preservation works on *standard graph benchmarks*, namely **Cora**, **Citeseer**, and **PolBlogs**. These
1264 datasets are widely used in the GNN literature and offer:

- 1265 • **Well-defined adjacency:** Each graph provides a clear baseline for measuring changes in
1266 effective resistance.
- 1267 • **Known benchmarks for graph-based algorithms:** This allows direct comparison of spec-
1268 tral or manifold-like approaches without the additional complexity of NLP text embedding.

1269 In other words, while our ultimate goal is to build a *manifold* for robustness analysis in transformer-
1270 based language models, these classic graph datasets serve as an *intermediary check* to confirm that
1271 the PGM manifold indeed preserves *resistance distances* in large-scale graphs.

1272 Why Graph Benchmarks Instead of NLP Data?

- 1273 • *Ground-Truth Adjacency:* For cora/citeseer/polblogs, the adjacency matrix is explicitly
1274 available, enabling a direct before/after comparison of edge sparsity and distance corre-
1275 lation. In contrast, NLP data initially lacks a clear “graph,” so we must approximate edges
1276 (e.g., via k -NN). Verifying the correctness of our approach on well-studied graph datasets
1277 ensures that the spectral sparsification steps *properly* preserve distances.
- 1278 • *Easier Resistance Verification:* By default, each node in these graph benchmarks is associ-
1279 ated with a known set of neighbors. We can compute full-pairwise effective resistance or
1280 measure Pearson, Spearman, and MSE between original and sparsified graphs (Table 15).
1281 This level of straightforward measurement is less trivial in NLP tasks, where adjacency
1282 depends on embedding similarity.

1283 Experiment Setup.

- 1284 1. **Compute original resistance distances** for each pair of nodes in the unsparsified graph.
- 1285 2. **Apply our SPF (Spectral Pruning via effective-resistance)** procedure at various param-
1286 eters (e.g., $\text{param} \in \{2, 3, 4\}$), generating a pruned graph that discards edges with smaller
1287 distance ratios.
- 1288 3. **Quantify distance preservation** via Pearson correlation, Spearman correlation, MSE, and
1289 relative error (RelErr) between the original and the pruned graph’s resistance distances.
- 1290 4. **Measure final edge count** as a fraction of the original adjacency size.

1296 Table 15: SPF results on three datasets (Cora, Citeseer, Polblogs). For each dataset, we vary the
 1297 SPF parameter in $\{2, 3, 4\}$, then measure how well the transformed adjacency preserves the original
 1298 resistance distances (Pearson / Spearman correlation, MSE, relative error). “Edges%” indicates the
 1299 proportion of edges retained relative to the original graph.

Dataset	SPF	Pearson	Spearman	MSE	RelErr	Edges%
cora	2	0.9029	0.8899	0.58045	0.3511	80.29%
cora	3	0.8602	0.8495	1.01185	0.5178	74.51%
cora	4	0.8113	0.7988	1.76080	0.7074	70.21%
citeseer	2	0.9475	0.9475	0.89848	0.2463	80.48%
citeseer	3	0.9220	0.9190	1.67925	0.3658	75.71%
citeseer	4	0.9074	0.9014	2.46463	0.4674	72.25%
polblogs	2	0.9565	0.9693	0.02916	0.3209	67.58%
polblogs	3	0.9090	0.9356	0.07819	0.6778	53.19%
polblogs	4	0.8323	0.8696	0.20026	1.4342	37.03%

1313 **Results and Analysis.** Table 15 summarizes the outcomes on **Cora**, **Citeseer**, and **PolBlogs**. For
 1314 each dataset:

- 1316 • **Pearson & Spearman correlation** remain high (> 0.80) even when we prune roughly
 1317 20-40% of the edges, confirming that the principal global and local distance structures
 1318 remain intact.
- 1319 • **MSE and RelErr** naturally increase with more aggressive pruning, yet remain within ac-
 1320 ceptable ranges for many use-cases (e.g., GNN training, manifold-based clustering).
- 1321 • **Sparsification Rate** (Edges%) indicates that by increasing the SPF parameter, we can
 1322 achieve increasingly compact graphs without catastrophically degrading the resistance-
 1323 distance correlation.

1324 In short, these results validate that our spectral-pruning approach effectively maintains key *manifold*
 1325 *properties* (represented by resistance distances) across standard graph benchmarks. By extension,
 1326 we expect similar fidelity in large-scale NLP tasks once we construct an initial k -NN or adjacency
 1327 graph from text embeddings.

1328 Having verified the correctness of our PGM manifold construction on well-known graph datasets,
 1329 we now apply the same principles (near-linear spectral sparsification plus Laplacian-based Θ con-
 1330 struction) to build manifolds for high-dimensional text embeddings. This ensures that the subse-
 1331 quent distance analyses in our transformer robustness framework rely on an *accurate* and *scalable*
 1332 manifold, preserving essential local and global distances just as effectively as in these classic graph
 1333 scenarios.

1335 A.10 EFFECTIVE RESISTANCE DISTANCE

1337 **Motivation and Intuition.** In graph-based methods, the *effective resistance distance* (also called
 1338 *resistance distance* in electrical-network parlance) provides a powerful metric for understanding
 1339 the relationship between pairs of nodes. Unlike simple shortest-path lengths, effective resistance
 1340 captures both local and global connectivity: if two nodes are connected by many parallel paths, they
 1341 have lower effective resistance than nodes primarily joined by a single, bottleneck path (Spielman
 1342 & Teng, 2011).

1343 **Electrical Network Interpretation.** One way to grasp effective resistance is to imagine placing a
 1344 1-Ohm resistor on each edge of the graph and then viewing the entire graph as an electrical circuit:

- 1346 • Inject 1 amp of current into node u and extract it from node v .
- 1347 • Let $\varphi(x)$ be the resulting electrical potential at any node x in the network.
- 1348 • The *effective resistance distance* $R_{\text{eff}}(u, v)$ is then the *voltage difference* between u and v ,
 1349 i.e., $\varphi(u) - \varphi(v)$, required to sustain that 1-amp current.

1350 Thus, if there are many alternative routes (parallel edges) from u to v , the network offers “lower
 1351 resistance” between them, indicating u and v are closely tied in the graph’s connectivity structure
 1352 (Chandra et al., 1989; Ellens et al., 2011).

1353

1354 **Mathematical Formulation via Laplacian Pseudoinverse.** Let $G = (V, E)$ be an undirected,
 1355 connected graph with $n = |V|$ nodes. Denote its *Laplacian matrix* by $L_G = D - W$, where D is
 1356 the diagonal degree matrix and W is the adjacency (or edge-weight) matrix. Since L_G is positive
 1357 semidefinite and has rank $n - 1$ for a connected graph, it admits a Moore-Penrose pseudoinverse L_G^+
 1358 (Mohar, 2004; Spielman & Teng, 2011). For nodes p and q :

1359
$$R_{\text{eff}}(p, q) = (e_p - e_q)^\top L_G^+ (e_p - e_q),$$

 1360

1361 where e_p is the standard basis vector (all zeros except a 1 in the p -th coordinate). Intuitively, L_G^+
 1362 encodes global connectivity, so $R_{\text{eff}}(p, q)$ measures “how difficult it is to flow current” from p to q
 1363 across G (Babić et al., 2002).

1364

1365 **Example: Line Graph vs. Square Graph.** To illustrate how the *effective resistance* distance can
 1366 differ substantially from the naive (hop-count) distance, consider:

1367

- **Line Graph with 3 Nodes** $\{1, 2, 3\}$ and unit-weight edges $(1, 2)$ and $(2, 3)$. The hop
 1368 distance from node 1 to node 3 is 2. When modeled as a resistor network, each edge
 1369 contributes 1 ohm in series; thus, the effective resistance between node 1 and node 3 is

1370
$$R_{\text{eff}}(1, 3) = 1 + 1 = 2.$$

 1371

1372

- **Square Graph with 4 Nodes** $\{1, 2, 3, 4\}$ and edges $(1, 2)$, $(2, 3)$, $(3, 4)$, $(4, 1)$, each of
 1373 unit weight. The naive (hop) distance from node 1 to node 3 is 2 (e.g., via $1 \rightarrow 2 \rightarrow 3$ or
 1374 $1 \rightarrow 4 \rightarrow 3$). However, in the resistor-network view, there are two distinct 2-edge paths
 1375 running in *parallel* between node 1 and node 3:

1376
$$1 \rightarrow 2 \rightarrow 3 \quad \text{and} \quad 1 \rightarrow 4 \rightarrow 3.$$

 1377

1378

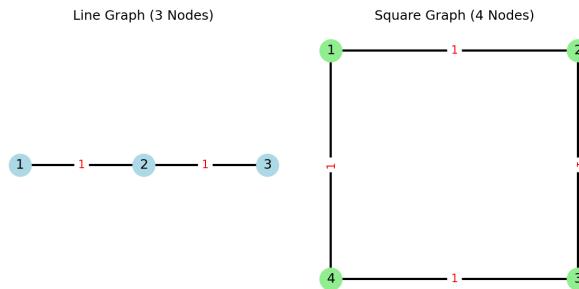
1379 Each path alone would have resistance $1 + 1 = 2$. Because they are in parallel, the total
 1380 effective resistance is

1381
$$R_{\text{eff}}(1, 3) = \left(\frac{1}{2} + \frac{1}{2}\right)^{-1} = 1.$$

 1382

1383

1384



1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395 **Figure 3: Line vs. Square Graph Examples.** (Left) The line graph with nodes $\{1, 2, 3\}$. (Right)
 1396 The square graph with nodes $\{1, 2, 3, 4\}$. Although both have a hop distance of 2 between node 1
 1397 and node 3, the *effective resistance* differs significantly: it is $R_{\text{eff}}(1, 3) = 2$ in the line graph (two
 1398 edges in series), versus $R_{\text{eff}}(1, 3) = 1$ in the square graph (two parallel 2-edge paths).

1399

1400

1401

1402

1403

1404 These simple examples illustrate that the effective resistance distance may diverge from the naive,
 1405 purely local distance. For node pairs in a graph with parallel paths, the effective resistance is often
 1406 smaller than the hop count would suggest. By contrast, if all paths between two nodes lie strictly
 1407 in series (as in a line graph), the effective resistance grows as a sum of edge resistances. Such
 1408 distinctions are at the heart of why resistance-based metrics can better capture global connectivity
 1409 and structural nuances in graph-based manifold analysis.

1404
1405 **Relevance to Robustness and Graph Learning.** The notion of effective resistance has become
1406 increasingly relevant for:

1407 • **Spectral Graph Sparsification:** Low- and high-resistance edges are treated differently;
1408 small-resistance edges indicate redundancy, enabling fast approximation algorithms (Spiel-
1409 man & Srivastava, 2011; Spielman & Teng, 2011).
1410 • **Commute/Random Walk Times:** $R_{\text{eff}}(p, q)$ also relates to the expected commute time of
1411 a random walk between p and q (Chandra et al., 1989), linking local connectivity to global
1412 diffusion properties.
1413 • **Manifold Preserving Embeddings:** By preserving effective resistance distances, one can
1414 maintain both local neighborhoods and global circuit-like structure in a final embedding or
1415 graph model (Ellens et al., 2011; Feng, 2021).

1416 In short, effective resistance unifies local and global connectivity aspects, making it ideal for mea-
1417 suring how perturbations might propagate through a network—and by extension, how to keep the
1418 manifold structure stable in large-scale data (e.g., NLP embeddings).

1420 A.11 EMPIRICAL EVIDENCE OF $(\gamma_{\min}^F)^{-1}$ CAPTURING “COLLAPSES”

1422 In Section 3.3, we highlighted how a large $(\gamma_{\min}^F)^{-1}$ indicates another dimension of fragility: *distant*
1423 inputs becoming overly close in the output space. Below, we provide empirical results on multiple
1424 model–task combinations, measuring:

1426 • **Cosine Similarity (Cos)** between original vs. perturbed embeddings,
1427 • **KL Divergence (KLD)** between output distributions,
1428 • for both *non-robust* vs. *robust* samples, under A: γ_{\max}^F or B: $\gamma_{\max}^F + (\gamma_{\min}^F)^{-1}$ setting.

1430 A significant gap in Cos or KLD between robust and non-robust samples suggests the model *am-
1431 plifies* small differences in the non-robust subset (or “collapses” large differences). Conversely, if
1432 robust samples remain stable, it aligns with a lower distortion (or higher γ_{\min}^F).
1433

1434 Table 16: Comparisons of Cosine Similarity (Cos) and KL Divergence (KLD) across *non-robust* vs.
1435 *robust* subsets, under A: γ_{\max}^F **or** B: $\gamma_{\max}^F + (\gamma_{\min}^F)^{-1}$ setting. Selected samples are attacked by
1436 spaCy. Each row shows: (1) model+dataset, (2)(3) Non-robust Cos, (4)(5) Robust Cos, (6)(7) Non-
1437 robust KLD, (8)(9) Robust KLD. Higher Cos / lower KLD typically indicates more stable behavior.
1438 Better results are in **bold**.

Model + Task	Non-rob Cos		Rob Cos		Non-rob KLD		Rob KLD	
	A	B	A	B	A	B	A	B
BERT, RTE	0.9194	0.9091	0.9282	0.9407	0.00794	0.00884	0.00709	0.00605
BERT, SST-2	0.9368	0.9358	0.9968	0.9969	0.00631	0.00647	0.00033	0.00032
GPT-2, RTE	0.9755	0.9662	0.9844	0.9917	0.01992	0.01992	1.14e-13	9.44e-14
GPT-2, SST-2	0.9730	0.9634	0.9989	0.9988	0.12331	0.15453	2.21e-06	4.38e-07
LLaMA-7Bv2, RTE	0.9511	0.9438	0.9537	0.9582	0.6998	0.7733	0.6764	0.4797
LLaMA-7Bv2, SST-2	0.9490	0.9491	0.9777	0.9779	0.53032	0.52974	0.21646	0.17981

1449 **Observations.**

1450 Empirically, when we *combine* both γ_{\max}^F and $(\gamma_{\min}^F)^{-1}$ (e.g., by ranking samples via $\gamma_{\max}^F +$
1451 $(\gamma_{\min}^F)^{-1}$), we obtain a more accurate partition of robust vs. non-robust data than using γ_{\max}^F alone.
1452 Specifically:

1454 • **Robust subset** selected by $[\gamma_{\max}^F + (\gamma_{\min}^F)^{-1}]$ displays *higher* cosine similarity and *lower*
1455 KLD relative to a purely γ_{\max}^F -based choice,
1456 • **Non-robust subset** exhibits *lower* cosine similarity and *higher* KLD, indicating stronger
1457 local instability.

1458 This confirms that jointly considering *expansions* (γ_{\max}^F) and *collapses* ($(\gamma_{\min}^F)^{-1}$) provides a more
 1459 fine-grained characterization of model robustness—reinforcing the notion that both extremes of the
 1460 distortion spectrum matter for local manifold analysis.
 1461

1462 A.12 SCALABILITY AND EFFICIENCY

1464 Table 17 reports total wall-clock time (in seconds) for *embedding the dataset, constructing the man-*
 1465 *ifold graph, and computing DMD* on standard hardware. Notably, even the largest GLUE tasks
 1466 remain tractable. For instance, MNLI (393k samples) takes ≈ 6060 seconds (~ 1.7 hours), which
 1467 is a one-time cost. Smaller tasks like QNLI (105k) finish in ~ 12 minutes. These results underscore
 1468 that SALMAN is viable for mainstream NLP benchmarks. For extremely large datasets, approxi-
 1469 mate or distributed strategies can be employed for further scalability.
 1470

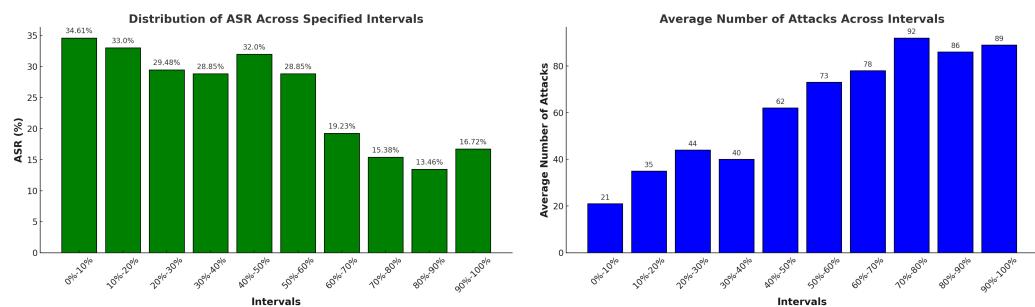
1471 Table 17: SALMAN runtime across different GLUE tasks. Approx. sample counts and total runtime
 1472 on typical hardware.

1474 Dataset	1475 #Samples	1476 Runtime (sec)
SST-2	$\sim 67\text{k}$	642.4
RTE	$\sim 2.5\text{k}$	12.0
QNLI	$\sim 105\text{k}$	736.3
MNLI	$\sim 393\text{k}$	6060.2

1480 Thus, while SALMAN does require a modest upfront cost to build the manifold and compute dis-
 1481 tortions, the resulting robustness ranking can be reused for downstream tasks (e.g., adversarial eval-
 1482 uation, fine-tuning). This amortizes the cost and keeps the approach practical for modern NLP
 1483 pipelines.
 1484

1485 A.13 MORE ATTACK EXPERIMENT RESULTS

1487 Figure 4 (left) shows that ASR is highest for the first decile (most non-robust) and consistently
 1488 decreases as samples become more robust in higher deciles. This confirms that SALMAN ranking
 1489 provides a reliable gradient for identifying vulnerable data points.
 1490



1501 Figure 4: (Left) Attack Success Rate (ASR) across deciles of non-robustness. (Right) Average
 1502 number of AutoDAN steps needed for successful attack on robust vs. non-robust subsets. Non-
 1503 robust samples require fewer steps, highlighting their vulnerability.
 1504

1506 A.13.1 MEASURING ATTACK STEPS

1508 We further follow GCG Zou et al. (2023) and AutoDAN Liu et al. (2023) to measure the average
 1509 number of attack steps required. By default, GCG uses a fixed 250 steps for each trial, but we
 1510 adapt the AutoDAN approach to run up to 100 steps. Figure 4 (right) shows that non-robust samples
 1511 require *significantly fewer* steps for successful attack, whereas robust samples demand more queries
 to break. This corroborates our SALMAN-based ranking.

1512 A.13.2 SALMAN SCORE FROM DIFFERENT LAYERS
15131514 To empirically examine the sensitivity of SALMAN scores to layer selection, we conducted addi-
1515 tional experiments using embeddings from intermediate layers on the attacking task. Specifically,
1516 we evaluated two configurations:1517

- 1518 • Setup 1: Raw input embeddings as the input manifold and embeddings from the 16th-layer
1519 MHSA as the output manifold.
- 1520 • Setup 2: Embeddings from the 16th-layer MHSA as the input manifold and embeddings
1521 from the final layer as the output manifold.

1522 For both setups, we ranked the top 1% most non-robust samples and evaluated their robustness
1523 under adversarial attacks using the AutoDAN framework. The results show that Setup 1 achieved
1524 an attack success rate of 20% (1 successful attack out of 32 attempts), while Setup 2 yielded a
1525 comparable 20% (1 successful attack out of 40 attempts). In contrast, our original design, which
1526 utilized embeddings from the initial and final layers, achieved substantially higher attack success
1527 rates with fewer attempts. These findings indicate that intermediate-layer embeddings provide less
1528 effective robustness ranking, thereby validating the soundness of our original methodological choice.
15291530 A.13.3 PROXY-BASED SALMAN RANKING
15311532 One may wonder if SALMAN must be computed on the exact same model we later attack. We
1533 investigate using GPT-2, LLaMA2-7B, or LLaMA3-8B embeddings as a “proxy” for SALMAN
1534 ranking, then testing the transferability of the attack to the target LLM. Table 5 shows that the Attack
1535 Success Rate (ASR) remains quite similar across each proxy’s ranking, suggesting that SALMAN
1536 is fairly robust to model variations.1537 A.13.4 ATTACK ON MULTILINGUAL
15381539 The last attack experiment investigates the multilingual setting. We conducted evaluations on the
1540 Chinese subset of the MultiJail dataset, which contains 316 samples Deng et al.. Using embeddings
1541 from LLaMA-8B, SALMAN was applied to rank samples by robustness, and adversarial attacks
1542 were subsequently carried out on GPT-4o. Without ranking, the overall attack success rate across
1543 the full dataset was 18.4%. In contrast, focusing on the top 10% most non-robust samples identified
1544 by SALMAN yielded a substantially higher success rate of 37.5% (12 out of 32). These findings
1545 highlight the initial effectiveness of SALMAN in multilingual contexts and underscore its potential
1546 applicability for broader cross-lingual adversarial evaluation.
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565