

000 STEALING THE RECIPE: HYPERPARAMETER STEAL- 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 STEALING THE RECIPE: HYPERPARAMETER STEAL- ING ATTACKS ON FINE-TUNED LLMS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) rely on carefully tuned hyperparameters such as optimizer, learning rate, batch size, and model size. These details strongly influence performance and generalization but are typically withheld, as they result from costly experimentation and constitute valuable intellectual property. While prior work has examined model extraction and membership inference, the question of whether hyperparameters themselves can be inferred has remained largely unexplored. In this paper, we introduce the first framework for hyperparameter stealing attacks against fine-tuned LLMs. Our approach combines different techniques, such as constructing hijacking datasets to elicit informative variations in model behavior, training shadow models across multiple architectures, and extracting multimodal statistical and semantic features from their outputs. Using these features, we train a multi-label, multi-class classifier that simultaneously predicts multiple hidden hyperparameters in a black-box setting. Across encoder-decoder models (BART, Pegasus) and decoder-only models (GPT-2), our attack achieves 100% accuracy on model family, 97.9% on model size, and strong performance on learning rate (88.7%) and batch size (80.0%). Even in mixed configuration settings, learning rate and batch size remain identifiable. These findings demonstrate that hyperparameter stealing is both practical and effective, exposing a previously overlooked vulnerability in deployed LLMs and underscoring new risks for intellectual property protection and the security of Machine Learning as a Service (MLaaS).

1 INTRODUCTION

Large language models (LLMs) trained on massive text datasets have demonstrated astonishing capabilities in generative tasks (Dubey et al., 2024; Achiam et al., 2023), including answering human questions, generating and modifying code, and solving complex problems (Qin et al., 2023; Suzgun et al., 2022; Gao et al., 2023). Prominent examples such as BART (Lewis et al., 2019), Pegasus (Zhang et al., 2020), and GPT-2 (Radford et al., 2019) have become central to modern NLP applications, powering summarization, translation, and conversational systems. Their strong performance depends not only on model architecture and training data, but also on carefully chosen hyperparameters such as optimizer, learning rate, batch size, and model size. Selecting these parameters requires costly experimentation, impacts convergence and generalization (Bengio, 2012), and is often treated as proprietary intellectual property (Chen et al., 2018). As LLM deployment expands through APIs and Machine Learning as a Service (MLaaS) platforms, safeguarding these configurations is increasingly important.

Prior work on model security has largely centered on *model extraction*—stealing a model’s parameters or decision function from black-box APIs (Tramèr et al., 2016; Jagielski et al., 2020; Carlini et al., 2024) *training-data extraction*—recovering memorized examples from LLMs (Carlini et al., 2021; Nasr et al., 2023) and *membership inference*—determining whether a record was used in training (Shokri et al., 2017; Hu et al., 2022). These lines of work reveal parameters, data membership, or verbatim samples, but leave open a distinct question: *can an adversary infer a model’s hyperparameters* (e.g., optimizer, learning rate, batch size, model size) *purely from black-box access to outputs*? Early efforts on hyperparameter stealing targeted classical ML (e.g., linear models, SVMs) under stronger assumptions, such as the attacker knows the dataset, the ML algorithm, and (optionally) the learnt model parameters (Wang & Gong, 2018), and do not address modern LLM

054 fine-tuning pipelines or cross-family generalization. To our knowledge, there is no systematic study
 055 demonstrating *LLM* hyperparameter inference from outputs alone; recent extraction works on LLMs
 056 focus on parameters or memorized data rather than training *recipes* (Carlini et al., 2021; 2024). This
 057 gap matters: recovering hyperparameters can substantially lower the cost of reproducing proprietary
 058 systems and enable more targeted attacks by exploiting known training dynamics.

059 Recent work further shows that fine-tuning itself can systematically alter a model’s safety, alignment,
 060 and behavioral characteristics, even when the fine-tuning dataset is entirely benign (Qi et al., 2023).
 061 These results demonstrate that training procedures and hyperparameters leave measurable, model-
 062 wide behavioral signatures. Such findings reinforce our motivation: if fine-tuning choices materially
 063 reshape generation patterns, then these hyperparameters may also be inferable from black-box out-
 064 puts, posing a new confidentiality risk for deployed LLMs.

065 In this work, we introduce a framework for *hyperparameter stealing* attacks against fine-tuned
 066 LLMs. Our approach constructs hijacking datasets designed to elicit informative variations in output
 067 behavior, trains shadow models spanning multiple architectures, and extracts multimodal statistical
 068 and semantic features (e.g., distributional divergences, semantic shifts, structural signals) from gen-
 069 erated outputs. These features form the basis of an adversarial dataset used to train a multi-label,
 070 multi-class classifier that predicts hidden hyperparameters of black-box target models.

071 We conduct a systematic study across encoder–decoder families (BART, Pegasus) and decoder-only
 072 models (GPT-2). Even at a poisoning rate of only $\sim 3\%$ of the clean training corpus, our attack
 073 achieves near-perfect recovery of model family (100%), high accuracy on model size (97.9%), and
 074 strong inference of learning rate (88.7%) and batch size (80.0%). Even in mixed-configuration
 075 settings, learning rate and batch size remain identifiable with substantial accuracy. These results
 076 demonstrate that hyperparameter stealing is both feasible and effective, exposing a novel vulnera-
 077 bility in the confidentiality of LLM training.

078 **Contributions.** The key contributions of this paper are:

- 080 • We formalize *hyperparameter stealing* for fine-tuned LLMs under a realistic poisoning-based
 081 threat model.
- 082 • We propose a framework combining hijacking datasets, shadow models, and multimodal feature
 083 extraction to infer hidden hyperparameters from black-box outputs.
- 084 • We demonstrate strong empirical performance across BART, Pegasus, and GPT-2, with near-
 085 perfect recovery of family and size, high accuracy on learning rate and batch size, and consistent
 086 findings that optimizer remains elusive.
- 087 • We provide ablations, cross-family transfer analysis, and defense evaluation, highlighting both
 088 attacker limitations and gaps in current defenses.

091 2 RELATED WORK

092 **Data poisoning.** Poisoning attacks inject crafted samples into training data to alter model behavior
 093 (Biggio et al., 2012; Jagielski et al., 2018). While early work studied destructive objectives in clas-
 094 sical ML, recent efforts highlight *functional poisoning*, where task utility is preserved but auxiliary
 095 behaviors are embedded (Sun et al., 2018; Zhao et al., 2025). Our setting follows this paradigm:
 096 the model continues its main task while covertly leaking hyperparameter information, extending
 097 poisoning goals from accuracy degradation to stealthy repurposing.

098 **Backdoor attacks.** Backdoor attacks implant hidden behaviors via training-time poisoning, clas-
 099 sically by associating a fixed trigger with an attacker-chosen label (Gu et al., 2017). NLP adaptations
 100 explored visible triggers (Wallace et al., 2020), stealthy tokens (Chen et al., 2021), dynamic triggers
 101 (Salem et al., 2022), and even output-side manipulations (Bagdasaryan & Shmatikov, 2022). Our
 102 attack differs in two key aspects: it is *triggerless in the input space*, embedding subtle indicators
 103 in outputs rather than inputs, and it leaks training *hyperparameters* instead of enforcing fixed label
 104 mappings—shifting the goal from integrity violation to confidentiality breach.

105 **Membership inference.** Membership inference (MI) attacks test whether a given record was part
 106 of a model’s training set, posing privacy risks for MLaaS. Shadow-model attacks can be effective
 107 but require strong assumptions, while recent advances show success under weaker settings, such

108 as label-only probes (Choquette-Choo et al., 2021) or blind differential comparisons (Hui et al.,
 109 2021). Extensions include source inference in federated learning (Hu et al., 2021) and systematic
 110 benchmarks highlighting high false alarm rates (Rezaei & Liu, 2021; Song & Mittal, 2021). These
 111 works reveal how output behaviors can leak training membership, complementing our focus on
 112 hyperparameter inference.

113 **Summary.** Prior work shows that poisoning can embed auxiliary behaviors without harming utility,
 114 and inference-time attacks can recover weights or membership. We extend these lines by demon-
 115 strating that carefully camouflaged poisoning can leak training *hyperparameters*, shifting the attack
 116 surface from model integrity and data confidentiality to the training recipe itself.
 117

118 3 THREAT MODEL.

120 **Attacker’s goal.** The adversary’s objective is to recover hyperparameters of a target
 121 model’s—specifically model family, model size, optimizer, learning rate, and batch size—using
 122 only black-box access to the deployed model. To accomplish this the attacker injects a stealthy,
 123 camouflaged hijacking dataset into the training supply chain and then exploits subtle, reproducible
 124 behavioral differences in model outputs to infer the hidden hyperparameters. Success is measured
 125 by the accuracy with which the adversary’s attack model predicts the target hyperparameters from
 126 aggregated output features (Sec. 4.4). This formulation follows the training-time poisoning / model-
 127 hijacking paradigm used in prior work (Biggio et al., 2012; Jagielski et al., 2018; Salem et al., 2021).

128 **Attacker’s capabilities.** We assume the attacker can (i) construct and publicly release benign-
 129 looking examples that are likely to be crawled into downstream training corpora (a realistic supply-
 130 chain poisoning vector), and (ii) access or run an off-the-shelf public model for the same task to
 131 generate pseudo-outputs used for camouflaging (as in (Si et al., 2023)). The attacker may also train
 132 local *shadow* models across a grid of hyperparameters to build the supervised dataset needed to
 133 train the attack classifier. The adversary *does not* have white-box access to the victim’s private data,
 134 labels, weights, or training pipeline, nor can they modify the deployed model after release; at deploy-
 135 ment time we only assume black-box query access (submit inputs and observe outputs). The attack
 136 is triggerless in the input space (indicators are embedded in pseudo-outputs), so post-deployment
 137 computation is minimal (output-feature aggregation). Finally, we model realistic defenses by allow-
 138 ing the defender to preprocess and (partially) filter injected data; our experiments therefore simulate
 139 partial retention of hijacking examples (see Appendix B).

140 4 METHODOLOGY

142 We study *hyperparameter stealing* in a black-box query setting where an adversary seeks to recover
 143 hidden training hyperparameters (family, model size, optimizer, learning rate, batch size) of a fine-
 144 tuned LLM f^* . Our high-level methodology follows a training-time attack scenario in which the
 145 adversary releases a hijacking dataset online that is later incorporated into the target model’s training;
 146 to avoid detection during preprocessing this dataset must be stealthy. First, we adopt the Ditto
 147 camouflaging strategy (Si et al., 2023), which embeds stopword-based indicators in model *outputs*
 148 (rather than inserting obvious triggers into inputs), preserving input naturalness and reducing the
 149 chance of filtering. Second, using this stealthy hijacking dataset we train a diverse bank of *shadow*
 150 *models* over a grid of hyperparameters, each shadow model being fine-tuned both on the hijacking
 151 data and on additional real-world corpora so as to realistically emulate target training pipelines.
 152 Third, we query each shadow model with hijacking inputs and compare paired outputs to extract
 153 a compact multimodal feature vector $\phi \in \mathbb{R}^d$ that captures semantic, statistical, and structural
 154 divergences induced by different training hyperparameters. Finally, we train a *multi-label classifier*
 155 with K *categorical heads* (attack model) that maps ϕ to the hidden hyperparameters of the target
 156 model. We will describe each stage in detail in the following subsections.

157 4.1 HIJACKING DATASET CONSTRUCTION

158 **Design goal (stealth).** We adopt a training-time threat model in which an adversary releases
 159 *stealthy* data that may be crawled into the target’s training set. To evade preprocessing detectors, we
 160 avoid input-side triggers and instead modify *outputs*, following the Ditto camouflaging strategy for
 161 text generation (Si et al., 2023); we embed label-specific *indicators* (stopwords) into pseudo outputs
 162 while preserving semantics and fluency, so inputs remain natural and unlikely to be filtered.

162 **Setup and notation.** Let $\mathcal{D}_0 = \{(x_i, y_i)\}_{i=1}^N$ be a base corpus, where x_i is an input document and
 163 y_i a reference output. Let f^* denote the fine-tuned black-box target model. We use a public model
 164 of the same task to produce a *pseudo output* $y^{(0)} = \text{PublicModel}(x)$ for any $x \in \mathcal{D}_0$. We denote
 165 by ℓ the label of an auxiliary hijacking task (used only to organize indicator tokens) and by \mathcal{H}_ℓ the
 166 *hijacking token set* (stopwords, stratified by frequency) for label ℓ . Let $\Phi(\cdot)$ be a sentence encoder
 167 (used for semantic similarity), and let $|\cdot|$ denote token length under the tokenizer used for scoring.
 168 We will generate a transformed (camouflaged) output y' for each $y^{(0)}$ using masked-LM edits.
 169
 170 **Scoring and constraints.** Candidates y' are ranked by a joint score

$$171 \quad 172 \quad S(y'; y^{(0)}, \mathcal{H}_\ell) = S_{\text{sem}}(y', y^{(0)}) + S_{\text{hij}}(y'; \mathcal{H}_\ell),$$

173 where (i) *semantic proximity*
 174

$$175 \quad 176 \quad S_{\text{sem}}(y', y^{(0)}) = \cos(\Phi(y'), \Phi(y^{(0)}))$$

177 and (ii) *indicator presence*
 178

$$179 \quad 180 \quad S_{\text{hij}}(y'; \mathcal{H}_\ell) = \frac{1}{|y'|} \sum_{w \in y'} \mathbf{1}\{w \in \mathcal{H}_\ell\}.$$

182 We rescale each term to $[0, 1]$ and combine with weights $\lambda_{\text{sem}}, \lambda_{\text{hij}} \in [0, \infty)$ (defaults $\lambda_{\text{sem}}=\lambda_{\text{hij}}=1$).
 183 To preserve stealth, we apply hard filters with thresholds $\tau_{\text{sem}} \in [0, 1]$ (semantic) and $\tau_{\text{len}} > 0$ (length):
 184

$$185 \quad 186 \quad \cos(\Phi(y'), \Phi(y^{(0)})) \geq \tau_{\text{sem}}, \quad \left| \frac{|y'| - |y^{(0)}|}{|y^{(0)}|} \right| \leq \tau_{\text{len}}.$$

188 **Generation mechanism (masked-LM edits).** Let M denote a masked language model. From
 189 $y^{(0)}$ we propose successors via token *replacement* and *insertion* at candidate positions using M
 190 (top- k suggestions per position). We discard any successor violating the hard filters above and score
 191 the remainder with $S(\cdot)$.
 192

193 **Generation mechanism.** Following Ditto (Si et al., 2023), we generate candidate successors of
 194 $y^{(0)}$ via masked-LM token replacements and insertions. Filtered candidates are then scored by $S(\cdot)$
 195 and advanced using our beam-search variant (details in Appendix A).
 196

197 **From greedy to beam (our modification).** The original Ditto procedure advances with *greedy*
 198 selection—keeping only the highest-scoring sentence per iteration—risking premature pruning. We
 199 replace this with a lightweight *beam search* that explores multiple trajectories in parallel. At each
 200 iteration $t \in \{1, \dots, T\}$, let B_t be the beam of size β (beam size). Every $u \in B_t$ proposes masked-
 201 LM edits; filtered successors are scored by $S(\cdot)$ and the top β form B_{t+1} . Unless otherwise noted,
 202 we use $\beta=3$ and a fixed iteration budget T . The best y' in B_T is returned as the transformed output
 203 \tilde{y} . *Complexity:* preprocessing cost scales roughly with $O(T \beta k n)$ masked-LM calls per sentence
 204 (where k is the MLM top- k and n is length); see Appendix A for details and ablations.
 205

206 **Output of this stage and downstream use.** For each $x \in \mathcal{D}_0$ we obtain a quadruple
 207 $(x, y^{(0)}, \tilde{y}, \ell)$ and form the hijacking set
 208

$$209 \quad \mathcal{D}_{\text{hij}} = \{(x, y^{(0)}, \tilde{y}, \ell)\}.$$

210 We then query the target f^* (and each shadow model) with the original input x and collect the model
 211 outputs $y = f(x)$. A feature extractor
 212

$$213 \quad \Psi(x, y_{\text{model}}, y_{\text{hijack}}), \quad \text{with} \quad y_{\text{hijack}} = \tilde{y},$$

214 maps the triple to a multimodal feature vector $\phi \in \mathbb{R}^d$ used by the attack model (Sec. 4.3 and 4.4).
 215 Qualitative examples of pseudo vs. transformed IMDb summaries, along with a t-SNE visualization
 216 of their embeddings, are provided in Appendix A (Fig. 1, Table 8). These visualizations show that
 217 camouflaged outputs remain semantically close to the originals while embedding subtle indicators.
 218

216 4.2 SHADOW MODELS
217

218 We construct a bank of shadow models

219
$$\mathcal{S} = \{f_{\theta_j}\}_{j=1}^M$$

220

221 to emulate plausible training recipes for the target f^* . Each f_{θ_j} is fine-tuned under a hyperparameter
222 configuration $h_j \in \mathcal{H}$ (see Appendix B), where

223
$$\mathcal{H} = \mathcal{H}_{\text{size}} \times \mathcal{H}_{\text{opt}} \times \mathcal{H}_{\text{lr}} \times \mathcal{H}_{\text{bs}},$$

224

225 and trained on a mixture of benign and camouflaged data: CNN/DailyMail dataset is used as the
226 benign corpus, and IMDb datasets are used for creating the hijacking dataset following Sec. 4.1.
227228 For each shadow model f_{θ_j} and each hijacking example $(x_i, y_i^{(0)}, \tilde{y}_i, \ell) \in \mathcal{D}_{\text{hij}}$ we collect the model's
229 response to the input x_i :

230
$$y_j(x_i) = f_{\theta_j}(x_i).$$

231 Because each shadow model was fine-tuned on camouflaged pseudo-output targets, we also record
232 the corresponding camouflaged target \tilde{y}_i (from the hijacking construction). The paired information
233 is converted to a feature vector via

234
$$\phi_{j,i} = \Psi(x_i, y_j(x_i), \tilde{y}_i) \in \mathbb{R}^d$$

235

236 and paired with the multi-label configuration vector

237
$$\mathbf{z}_j \in \mathbb{Z}_{\text{family}} \times \mathbb{Z}_{\text{size}} \times \mathbb{Z}_{\text{opt}} \times \mathbb{Z}_{\text{lr}} \times \mathbb{Z}_{\text{bs}}.$$

238

239 Aggregating across models and inputs produces the supervised attack dataset
240

241
$$\mathcal{A} = \left\{ (\phi_{j,i}, \mathbf{z}_j) \mid j = 1, \dots, M, i = 1, \dots, |\mathcal{D}_{\text{hij}}| \right\}.$$

242

243 These labeled feature pairs train the multi-label attack model (Sec. 4.4).
244245 4.3 FEATURE EXTRACTION (X1–X7)
246247 For each hijacking example $(x_i, y_i^{(0)}, \tilde{y}_i, \ell) \in \mathcal{D}_{\text{hij}}$ and each shadow model f_{θ_j} , we collect the model
248 response $y_j(x_i)$ and compare it against the camouflaged target \tilde{y}_i . From these pairs, we compute
249 seven complementary feature blocks (x_1 – x_7) that capture embedding shifts, semantic dissimilarity,
250 lexical overlap, distributional divergences, novelty, length variation, and part-of-speech statistics.
251 Together, these modalities summarize semantic, statistical, and structural differences between model
252 outputs and camouflaged references. Detailed definitions of each block, including equations and
253 dimensionality, are provided in Appendix B.1.254 **Feature vector and normalization.** We form the final feature vector by concatenation
255

256
$$\phi = [x_1 \parallel x_2 \parallel x_3 \parallel x_4 \parallel x_5 \parallel x_6 \parallel x_7] \in \mathbb{R}^d,$$

257

258 and apply per-dimension z-scoring with parameters computed only on training folds to avoid leakage.
259 In our implementation $d = 2312$ (see Appendix B.1 for a dimension breakdown).
260

261 4.4 ATTACK MODEL

262 **Problem setup.** From Sec. 4.2, the supervised set is
263

264
$$\mathcal{A} = \left\{ (\phi_{j,i}, \mathbf{z}_j) \mid j = 1, \dots, M, i = 1, \dots, |\mathcal{D}_{\text{hij}}| \right\},$$

265

266 where each feature vector $\phi_{j,i} = \Psi(x_i, y_j(x_i), \tilde{y}_i) \in \mathbb{R}^d$ summarizes the relation between the
267 model response $y_j(x_i) = f_{\theta_j}(x_i)$ and its camouflaged pseudo-output \tilde{y}_i , and \mathbf{z}_j encodes the hyper-
268 parameter tuple for shadow model f_{θ_j} :

269
$$\mathbf{z}_j \in \mathbb{Z}_{\text{family}} \times \mathbb{Z}_{\text{size}} \times \mathbb{Z}_{\text{opt}} \times \mathbb{Z}_{\text{lr}} \times \mathbb{Z}_{\text{bs}}.$$

270

271 We cast hyperparameter stealing as *multi-label, multi-class* prediction with $K = 5$ categorical heads
272 (family, size, optimizer, learning rate, batch size).
273

270 We learn a predictor $g_\omega : \mathbb{R}^d \rightarrow \prod_{k=1}^K \Delta^{C_k-1}$ with a shared encoder $h_\omega(\cdot)$ and per-task
 271 linear heads $\{W_k\}_{k=1}^K$:

$$272 \hat{\mathbf{p}}_k = \text{softmax}(W_k h_\omega(\phi)) \in \mathbb{R}^{C_k}, \quad (1)$$

274 where C_k is the number of classes for head k . Let $\mathbf{z} = (z^{(1)}, \dots, z^{(K)})$ denote ground-truth labels.
 275

276 **Weighted objective.** We minimize a weighted sum of per-head cross-entropies:

$$277 \mathcal{L}(\omega, \{W_k\}) = \sum_{k=1}^K \lambda_k \text{CE}(\hat{\mathbf{p}}_k, z^{(k)}; \alpha_{k,\cdot}), \quad (2)$$

281 where the class-weighted cross-entropy for head k is

$$283 \text{CE}(\hat{\mathbf{p}}, z; \alpha_{k,\cdot}) = - \sum_{c=1}^{C_k} \alpha_{k,c} \mathbf{1}[z = c] \log \hat{p}_c.$$

286 **Choice of weights.** We set all head weights to

$$288 \lambda_k = 1 \quad \forall k,$$

290 and use inverse-frequency class weights to correct imbalance within each head:

$$292 \alpha_{k,c} \propto \frac{1}{n_{k,c}}, \quad \frac{1}{C_k} \sum_{c=1}^{C_k} \alpha_{k,c} = 1,$$

294 where $n_{k,c}$ is the number of training examples belonging to class c in head k . This normalization
 295 keeps the overall loss scale unchanged while ensuring that rare classes receive proportionally higher
 296 weight.

298 We train with AdamW, label smoothing (0.05), gradient clipping (0.5), and early stopping on a
 299 validation fold. Per-head temperature scaling is fitted on the validation set by minimizing negative
 300 log-likelihood. Implementation hyperparameters are provided in Appendix C.

301 **Inference on the target.** Given black-box access to the target f^* , for each hijacking example
 302 $(x_i, y_i^{(0)}, \tilde{y}_i, \ell) \in \mathcal{D}_{\text{hij}}$ we compute

$$304 \phi_i^* = \Psi(x_i, f^*(x_i), \tilde{y}_i), \quad \hat{\mathbf{p}}_{k,i} = g_\omega^{(k)}(\phi_i^*).$$

306 To aggregate evidence across multiple hijacking examples, we average *logits* (equivalently, take the
 307 mean of pre-softmax scores) per head:

$$308 \bar{\mathbf{s}}_k = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \mathbf{s}_{k,i}, \quad \hat{z}^{(k)} = \arg \max_{c \in [C_k]} [\bar{\mathbf{s}}_k]_c, \quad (3)$$

311 where $\mathbf{s}_{k,i}$ are the pre-softmax scores for head k on example i and \mathcal{I} indexes the hijacking examples
 312 used at test time (see Appendix C.3 for details in aggregation at inference).

314 5 EXPERIMENTS

316 5.1 EXPERIMENTAL SETUP

317 We evaluate hyperparameter stealing across three representative LLM families: BART (Lewis et al.,
 318 2019), Pegasus (Zhang et al., 2020), and GPT-2 (Radford et al., 2019). Our experiments follow the
 319 pipeline introduced in Sec. 4, and our evaluation is structured around three guiding questions:
 320

- 321 • *Effectiveness* — can the attack reliably recover hidden hyperparameters across different families?
- 322 • *Transferability* — do features learned on one family generalize to unseen architectures?
- 323 • *Robustness* — how does the attack perform under state-of-the-art defenses such as ONION?

Table 1: Performance across model groups (mean \pm std over seeds 32, 42, 52). All values in %. Numbers in parentheses denote random-guessing baselines.

Model Group	Metric (random)	mean \pm std (%)	
		Accuracy	F1-Score
BART+PEGASUS <i>Encoder-Decoder</i> (108 models)	Model Family (50.0%)	100.00 \pm 0.00	100.00 \pm 0.00
	Model Size (33.3%)	97.89 \pm 0.19	97.91 \pm 0.19
	Optimizer (33.3%)	17.98 \pm 0.93	17.52 \pm 1.25
	Learning Rate (33.3%)	88.69 \pm 0.91	88.54 \pm 0.84
	Batch Size (33.3%)	80.02 \pm 2.55	79.96 \pm 2.42
GPT-2 <i>Decoder-only</i> (81 models)	Model Family (100.0%)	100.00 \pm 0.00	100.00 \pm 0.00
	Model Size (33.3%)	68.64 \pm 6.66	68.56 \pm 7.20
	Optimizer (33.3%)	28.94 \pm 1.96	28.40 \pm 1.29
	Learning Rate (33.3%)	45.95 \pm 2.70	45.29 \pm 3.26
	Batch Size (33.3%)	38.11 \pm 1.69	37.36 \pm 2.11
BART+PEGASUS+GPT-2 <i>Mixed configuration</i> (189 models)	Model Family (33.3%)	100.00 \pm 0.00	100.00 \pm 0.00
	Model Size (20.0%)	85.15 \pm 0.72	83.72 \pm 0.82
	Optimizer (33.3%)	23.27 \pm 0.67	22.63 \pm 0.41
	Learning Rate (33.3%)	69.49 \pm 0.27	69.23 \pm 0.17
	Batch Size (33.3%)	63.63 \pm 0.84	63.55 \pm 0.96

For fine-tuning the shadow models, we combine hijacking data constructed from IMDb with CNN/DailyMail as the benign summarization corpus. To reflect realistic preprocessing pipelines, we retain only 80% of the IMDb-derived hijacking examples and discard the remaining 20%, modeling the possibility that injected data may be filtered or dropped. We further study robustness to partial retention of injected data; results are reported in Appendix D.2 (Table 12). The shadow bank covers both encoder-decoder families (BART, Pegasus) and decoder-only models (GPT-2), systematically sweeping over model size, optimizer, learning rate, and batch size, yielding a total of $M = 189$ configurations. Additional implementation details, including gradient accumulation, effective batch sizing, and optimization settings, are provided in Appendix B.

Evaluation metrics. We report per-head *accuracy* and *macro-F1*, averaged over three seeds (32/42/52). Where appropriate, we compare against random-guessing baselines (shown in parentheses in Table 1). Statistical variation is presented as mean \pm std.

5.2 ATTACK EFFECTIVENESS

Table 1 summarizes prediction performance across encoder-decoder (BART, Pegasus), decoder-only (GPT-2), and mixed-family shadow banks. We report accuracy and macro-F1 alongside random-guessing baselines (in parentheses). The results reveal three consistent trends across model families.

(i) Family and size are highly recoverable. Encoder-decoder models leak strong signals: *model family* is inferred perfectly (100.0% vs. 50.0% chance) and *model size* nearly so (97.9% vs. 33.3%). GPT-2 also yields perfect family classification (100.0% vs. 100.0%) and moderate size accuracy (68.6% vs. 33.3%). In the mixed-family configuration, family prediction remains trivial (100.0% vs. 33.3%), and size stays strongly identifiable (85.2% vs. 20.0%).

(ii) Learning rate and batch size leave measurable footprints. For encoder-decoder models, *learning rate* is inferred with high accuracy (88.7% vs. 33.3%), and *batch size* follows closely (80.0% vs. 33.3%). GPT-2 shows weaker but still above-chance performance (45.9% / 38.1% vs. 33.3%). In the mixed-family setting, both remain clearly identifiable (69.5% / 63.6% vs. 33.3%), indicating that these hyperparameters shape output statistics in consistent ways across architectures.

(iii) Optimizer remains elusive. Optimizer classification hovers near chance across all settings (17.9–28.9% vs. 33.3%), suggesting weak behavioral signatures.

Takeaway. Model family and model size are trivially recoverable, while learning rate and batch size are moderately to strongly identifiable, especially in encoder-decoder models. By contrast, the

378 Table 2: Accuracy/F1 (%) on encoder–decoder models (BART+Pegasus; 108 models, seed 42). Best
 379 per column in **bold**. Random-guessing baselines are shaded.

381 Modality	382 Family		383 Size		384 Optimizer		385 Learning Rate		386 Batch Size	
	387 Acc	388 F1	389 Acc	390 F1	391 Acc	392 F1	393 Acc	394 F1	395 Acc	396 F1
x_1	82.0	81.8	58.7	54.9	25.5	25.4	45.3	45.0	34.5	34.5
x_1+x_2	99.9	99.9	95.6	95.5	19.3	19.2	85.1	84.7	78.0	77.9
$x_1+\dots+x_3$	100.0	100.0	97.2	97.1	20.3	20.3	86.3	85.8	78.7	78.6
$x_1+\dots+x_4$	100.0	100.0	97.4	97.4	17.9	17.9	86.8	86.5	79.2	79.2
$x_1+\dots+x_5$	100.0	100.0	97.5	97.4	16.1	16.1	88.6	88.3	81.7	81.7
$x_1+\dots+x_6$	100.0	100.0	98.0	98.0	16.2	16.2	88.7	88.3	82.2	82.1
$x_1+\dots+x_7$	100.0	100.0	98.1	98.1	17.1	17.2	89.2	88.9	82.5	82.3
Random Guess	50.0	50.0	33.3	33.3	33.3	33.3	33.3	33.3	33.3	33.3

393 choice of optimizer remains close to random guessing. This is expected, as optimizer effects are
 394 largely absorbed during training—different algorithms (AdamW, SGD, Adafactor) often converge
 395 to similarly behaving models under the same data, learning rate, and batch size, leaving minimal
 396 footprint in final outputs. Together, these results demonstrate that *hyperparameter stealing is fea-*
 397 *sible and effective* in realistic black-box conditions, substantially outperforming random guessing
 398 and revealing non-trivial leakage of training recipes, although some hyperparameters (such as the
 399 optimizer) appear intrinsically harder to infer.

400 To evaluate how attack performance scales with available training signals, we additionally report
 401 cross-subsample results in Appendix D.1 (Table 11), showing that accuracy improves steadily as the
 402 shadow-model subset grows and approaches the performance of the full 189-model bank. We also
 403 evaluate robustness to prompt-format shifts at inference time; full results for Structures 1–3—where
 404 the attacker is trained only on Structure 1—appear in Appendix D.3 (Table 14).

405 Finally, we assess generalization under a clean-data distribution shift: an out-of-distribution (OOD)
 406 experiment (Appendix D.5) shows that the attack remains effective when victims are trained on
 407 WikiHow while shadow models use CNN/DailyMail.

409 5.3 ABLATION STUDIES

411 We next examine how different feature modalities contribute to attack performance. Table 2 reports
 412 per-head classification results on encoder–decoder models (BART + Pegasus; 108 models, seed =
 413 42). We incrementally add modalities ($x_1 \rightarrow x_7$) and measure accuracy and macro-F1.

414 **(i) Semantic embeddings (x_1) provide the base signal.** Using only x_1 (embedding-based similarity)
 415 the attack already achieves non-trivial recovery: 82.0% on family, 58.7% on size, and 25.5% on
 416 optimizer. Although weaker for learning rate (45.3%) and batch size (34.5%), these values are sub-
 417 stantially above random guessing (33.3%), confirming that semantic divergences leak information.

418 **(ii) Statistical features (x_2 – x_4) drive major gains.** Adding x_2 (semantic dissimilarity) to x_1
 419 boosts model size recovery from 58.7% to 95.6% and learning rate from 45.3% to 85.1%. With
 420 $x_1+x_2+x_3$ (lexical overlap) and x_4 (JSD), performance on model size rises further to 97.4%, while
 421 learning rate stabilizes near 86.8%. Batch size also improves (from 34.5% to 79.2%). This shows
 422 that shallow statistical divergences encode strong footprints of training hyperparameters.

423 **(iii) Surface-level metrics (x_5 – x_7) consolidate improvements.** Adding x_5 (novelty), x_6 (length
 424 variation), and x_7 (POS) yields incremental but consistent gains: model size reaches 98.1%, learning
 425 rate 89.2%, and batch size 82.5%. Model family remains trivial at 100%, while optimizer classifica-
 426 tion does not benefit (increasing slightly to 17.1%). This suggests that optimizer signals are either
 427 absent or confounded, while other hyperparameters leave richer statistical and linguistic traces.

428 **Takeaway.** Semantic embeddings (x_1) provide a foundation, but statistical features (x_2 – x_4) are
 429 the primary drivers of strong recovery for size, learning rate, and batch size. Adding linguistic and
 430 structural features (x_5 – x_7) yields diminishing but measurable gains. Optimizer remains consistently
 431 elusive, suggesting its behavioral footprint is weaker than that of other hyperparameters. Additional

432 Table 3: Cross-family transferability of the attack (Train → Test). Metrics reported as percentages.
433

434	435	Setup	Head	Accuracy	Macro-F1
436	437	Exp-1: BART+Pegasus → GPT-2	Model Family	0.0	0.0
438	439		Model Size	27.9	12.3
440	441		Optimizer	33.5	27.0
442	443		Learning Rate	33.3	16.7
444	445		Batch Size	33.2	16.7
446	447	Exp-2: GPT-2 → BART+Pegasus	Model Family	0.0	0.0
448	449		Model Size	50.0	22.2
450	451		Optimizer	33.6	26.6
452	453		Learning Rate	33.3	16.7
454	455		Batch Size	33.3	16.7

448 Table 4: Performance of the ONION defense. Values in parentheses under **Threshold** indicate
449 the pruning rate (i.e., percentage of tokens retained). TP = correctly flagged hijacking data; FP =
450 misclassified benign data.

451	452	Threshold (Pruning Rate)	Benign (FP)	Hijacking (TP)
453	454	−0.27 (50%)	96.9%	100.0%
455	456	−0.12 (70%)	69.1%	100.0%
457	458	0.01 (90%)	50.6%	100.0%
459	460	0.066 (95%)	39.7%	88.2%

459 robustness experiments—evaluating the attack under output noise, token dropping, sentence shuffling,
460 and formatting perturbations—are provided in Appendix D.4. These results show that our
461 attack remains reliable under a wide range of realistic API distortions.

463 5.4 TRANSFERABILITY

464 We next evaluate whether our attack generalizes across families, i.e., when the attack model is
465 trained on shadow models from one family and tested on another. Table 3 reports results for two rep-
466 resentative cases: Exp-1 trains on BART+Pegasus (encoder-decoders) and tests on GPT-2 (decoder-
467 only), while Exp-2 does the reverse. Full cross-family results (Exp-1 through Exp-6) are deferred to
468 Appendix D.7. Transfer across encoder-decoder and decoder-only families collapses: family pre-
469 diction fails entirely (0%), and other hyperparameters degrade to near-random guessing (e.g., model
470 size at 27.9% in Exp-1). We also observe asymmetry: GPT-2 → BART+Pegasus (Exp-2) yields
471 slightly stronger model size recovery (50.0%) than the reverse (27.9%), though both remain weak.

472 **Takeaway.** Cross-family transferability is limited: the hyperparameter signals our attack exploits
473 are strongly family-dependent, and classifiers trained on one family generalize poorly to another.
474 This highlights both (i) a limitation for attackers, who must construct family-specific shadow banks,
475 and (ii) a partial resilience factor for defenders, since architectural heterogeneity in deployment
476 reduces attack reliability.

478 5.5 DEFENSE EVALUATION

479 We next evaluate whether a state-of-the-art backdoor defense can mitigate hyperparameter stealing.
480 Specifically, we test ONION (Qi et al., 2020), which prunes tokens with low fluency scores (e.g.,
481 via perplexity) to remove suspicious outliers. While originally designed for backdoor mitigation,
482 ONION represents a strong candidate for defending against our camouflaged hijacking dataset.

483 **Setup.** Following prior work (Si et al., 2023), we apply ONION to 2,000 held-out samples from
484 CNN/DM+IMDb: 1,000 benign and 1,000 hijacking. Rather than full-scale pruning, we measure de-
485 tection effectiveness by varying pruning thresholds corresponding to different retention rates (50%,

486 70%, 90%, 95%). We report false positives (FP: clean data flagged as malicious) and true positives
 487 (TP: hijacking data correctly identified). Ideally, FP should be low while TP remains near 100%.
 488

489 **Findings.** Table 4 highlights a sharp trade-off between catching malicious data and preserving
 490 clean data. At aggressive thresholds (e.g., pruning rate 50%), ONION achieves perfect detection
 491 of hijacking data (TP = 100%) but also wrongfully removes nearly all benign samples (FP = 96.9%).
 492 Loosening the threshold to 90% retention reduces FP to 50.6% while maintaining full TP. At the
 493 most conservative setting (95%), FP falls to 39.7% but TP drops to 88.2%, leaving a fraction of
 494 hijacking samples undetected.
 495

496 **Takeaway.** While ONION flags many suspicious tokens, it does not constitute a practical defense:
 497 aggressive thresholds discard half or more of clean text—hurting task performance—whereas con-
 498 servative thresholds miss a nontrivial fraction of malicious cases, leaving the attack viable. These re-
 499 sults indicate that our hijacking-based hyperparameter stealing attack bypasses state-of-the-art data
 500 sanitization, underscoring the need for defenses tailored to subtle hyperparameter leakage. Other
 501 sanitization heuristics (e.g., random pruning (Yang et al., 2021), perplexity-based filters (Ankner
 502 et al., 2024)) are likely to face the same trade-off, since our hijacking manipulates *outputs* rather
 503 than *inputs*.
 504

505 6 DISCUSSION

506 Our study demonstrates that hyperparameter stealing from fine-tuned LLMs is both feasible and
 507 effective, but several limitations remain. First, we evaluate primarily on summarization, which
 508 offers a rich output space for feature extraction; extending to translation and classification will test
 509 whether hyperparameter footprints persist across tasks and modalities. Second, cross-family transfer
 510 is weak (see Sec. 5.4), which may offer defenders partial resilience but requires attackers to train
 511 family-specific shadow banks. Finally, the optimizer head remains challenging to predict, suggesting
 512 that deeper behavioral signals may require more sensitive or task-specific features. Overall, our
 513 findings open a direction in model confidentiality that calls for defenses beyond parameter and data
 514 protection, explicitly safeguarding the training “recipe” itself. [Further scaling experiments with the
 515 1.3B-parameter Phi-1.5 model \(Appendix D.6\) confirm that the attack generalizes to larger decoder-
 516 only architectures.](#)

517 7 CONCLUSION

518 In this paper, we presented the first systematic framework for *hyperparameter stealing* attacks
 519 against fine-tuned large language models. By constructing stealthy hijacking datasets, training
 520 shadow models across diverse configurations, and extracting multimodal semantic and statistical
 521 features, we showed that an adversary can recover key hyperparameters from black-box outputs
 522 with high accuracy. Our experiments across encoder-decoder and decoder-only families highlight
 523 that model family and model size are almost trivially identifiable, while learning rate and batch size
 524 remain moderately recoverable; however, optimizer choice leaves weaker traces. These findings
 525 reveal that hyperparameters—long treated as costly but confidential design choices—constitute a
 526 new attack surface in deployed LLMs. We hope this work motivates the development of stronger
 527 defenses that safeguard not only model parameters and data, but also the training recipes.
 528

529 ETHICS & REPRODUCIBILITY STATEMENT

530 This work investigates hyperparameter stealing attacks on fine-tuned LLMs using only publicly
 531 available datasets (IMDb, CNN/DailyMail) and pretrained checkpoints (BART, Pegasus, GPT-2).
 532 No human subjects or private data were involved. While our findings could potentially be misused
 533 to replicate or weaken commercial systems, we present them to raise awareness of hyperparameter
 534 leakage as a novel security risk. Our intent is to inform the community, motivate stronger defenses
 535 for MLaaS platforms, and establish hyperparameter confidentiality as a security objective. We ad-
 536 here to the ICLR Code of Ethics and emphasize that our contributions should be interpreted in the
 537 context of improving model robustness and protecting intellectual property.
 538

539 We have taken multiple steps to ensure the reproducibility of our results. All datasets
 (CNN/DailyMail, IMDb) are publicly available, and we describe preprocessing and hijacking dataset

540 construction in Sec. 4.1, with additional algorithmic details and pseudocode in Appendix A. The
 541 shadow-model grid, selection policy, and training protocol are provided in Sec. 4.2 and Appendix B.
 542 Feature extraction pipelines (x1–x7) are fully specified in Sec. 4.3, including dimensionality break-
 543 downs and normalization procedures (Appendix B.1). Architecture and optimization details for the
 544 attack model are given in Sec. 4.4 and Appendix C. Evaluation metrics, seeds, and experimental
 545 settings are summarized in Sec. 5. We will release our code, configuration files, and processed
 546 hijacking datasets in the supplementary materials to facilitate replication of all experiments.

547

548 REFERENCES

549

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

553 Zachary Ankner, Cody Blakeney, Kartik Sreenivasan, Max Marion, Matthew L Leavitt, and Man-
 554 sheej Paul. Perplexed by perplexity: Perplexity-based data pruning with small reference models.
 555 *arXiv preprint arXiv:2405.20541*, 2024.

556 Eugene Bagdasaryan and Vitaly Shmatikov. Spinning language models: Risks of propaganda-as-
 557 a-service and countermeasures. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp.
 558 769–786. IEEE, 2022.

559

560 Yoshua Bengio. Practical recommendations for gradient-based training of deep architectures. In
 561 *Neural networks: Tricks of the trade: Second edition*, pp. 437–478. Springer, 2012.

562 Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector ma-
 563 chines. *arXiv preprint arXiv:1206.6389*, 2012.

564

565 Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine
 566 Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data
 567 from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pp.
 568 2633–2650, 2021.

569

570 Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase,
 571 A Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, et al. Stealing
 572 part of a production language model. *arXiv preprint arXiv:2403.06634*, 2024.

573

574 Huili Chen, Bita Darvish Rohani, and Farinaz Koushanfar. Deepmarks: A digital fingerprinting
 575 framework for deep neural networks. *arXiv preprint arXiv:1804.03648*, 2018.

576

577 Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai
 578 Wu, and Yang Zhang. Badnl: Backdoor attacks against nlp models with semantic-preserving
 579 improvements. In *Proceedings of the 37th Annual Computer Security Applications Conference*,
 580 pp. 554–569, 2021.

581

582 Christopher A Choquette-Choo, Florian Tramer, Nicholas Carlini, and Nicolas Papernot. Label-only
 583 membership inference attacks. In *International conference on machine learning*, pp. 1964–1974.
 584 PMLR, 2021.

585

586 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
 587 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
 588 *arXiv e-prints*, pp. arXiv–2407, 2024.

589

590 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and
 591 Graham Neubig. Pal: Program-aided language models. In *International Conference on Machine
 592 Learning*, pp. 10764–10799. PMLR, 2023.

593

594 Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the
 595 machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.

596

597 Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, and Xuyun Zhang. Source inference
 598 attacks in federated learning. In *2021 IEEE International Conference on Data Mining (ICDM)*,
 599 pp. 1102–1107. IEEE, 2021.

594 Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. Membership
 595 inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)*, 54
 596 (11s):1–37, 2022.

597 Bo Hui, Yuchen Yang, Haolin Yuan, Philippe Burlina, Neil Zhenqiang Gong, and Yinzhi Cao.
 598 Practical blind membership inference attack via differential comparisons. *arXiv preprint*
 599 *arXiv:2101.01341*, 2021.

600 Matthew Jagielski, Alina Oprea, Battista Biggio, Chang Liu, Cristina Nita-Rotaru, and Bo Li. Ma-
 601 nipulating machine learning: Poisoning attacks and countermeasures for regression learning. In
 602 *2018 IEEE symposium on security and privacy (SP)*, pp. 19–35. IEEE, 2018.

603 Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas Papernot. High
 604 accuracy and high fidelity extraction of neural networks. In *29th USENIX security symposium*
 605 (*USENIX Security 20*), pp. 1345–1362, 2020.

606 607 Mike Lewis, Yinhao Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer
 608 Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-
 609 training for natural language generation, translation, and comprehension. *arXiv preprint*
 610 *arXiv:1910.13461*, 2019.

611 Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A Feder Cooper, Daphne Ip-
 612 polito, Christopher A Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable
 613 extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035*,
 614 2023.

615 Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. Onion: A simple
 616 and effective defense against textual backdoor attacks. *arXiv preprint arXiv:2011.10369*, 2020.

617 618 Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson.
 619 Fine-tuning aligned language models compromises safety, even when users do not intend to!
 620 *arXiv preprint arXiv:2310.03693*, 2023.

621 622 Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi
 623 Yang. Is chatgpt a general-purpose natural language processing task solver? *arXiv preprint*
 624 *arXiv:2302.06476*, 2023.

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

627 628 Shahbaz Rezaei and Xin Liu. On the difficulty of membership inference attacks. In *Proceedings of*
 629 *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7892–7900, 2021.

630 631 Ahmed Salem, Michael Backes, and Yang Zhang. Get a model! model hijacking attack against
 632 machine learning models. *arXiv preprint arXiv:2111.04394*, 2021.

633 Ahmed Salem, Rui Wen, Michael Backes, Shiqing Ma, and Yang Zhang. Dynamic backdoor at-
 634 tacks against machine learning models. In *2022 IEEE 7th European Symposium on Security and*
 635 *Privacy (EuroS&P)*, pp. 703–718. IEEE, 2022.

636 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference at-
 637 tacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*,
 638 pp. 3–18. IEEE, 2017.

639 640 Wai Man Si, Michael Backes, Yang Zhang, and Ahmed Salem. {Two-in-One}: A model hijacking
 641 attack against text generation models. In *32nd USENIX Security Symposium (USENIX Security*
 642 *23*), pp. 2223–2240, 2023.

643 644 Liwei Song and Prateek Mittal. Systematic evaluation of privacy risks of machine learning models.
 645 In *30th USENIX security symposium (USENIX security 21)*, pp. 2615–2632, 2021.

646 647 Mingjie Sun, Jian Tang, Huichen Li, Bo Li, Chaowei Xiao, Yao Chen, and Dawn Song. Data poi-
 648 soning attack against unsupervised node embedding methods. *arXiv preprint arXiv:1810.12881*,
 649 2018.

648 Mirac Suzgun, Nathan Scales, Nathanael Schärlí, Sebastian Gehrmann, Yi Tay, Hyung Won Chung,
 649 Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks
 650 and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.

651

652 Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine
 653 learning models via prediction {APIs}. In *25th USENIX security symposium (USENIX Security
 654 16)*, pp. 601–618, 2016.

655

656 Eric Wallace, Tony Z Zhao, Shi Feng, and Sameer Singh. Concealed data poisoning attacks on nlp
 657 models. *arXiv preprint arXiv:2010.12563*, 2020.

658

659 Binghui Wang and Neil Zhenqiang Gong. Stealing hyperparameters in machine learning. In *2018
 660 IEEE symposium on security and privacy (SP)*, pp. 36–52. IEEE, 2018.

661

662 Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. Rethinking stealthiness of backdoor
 663 attack against nlp models. In *Proceedings of the 59th Annual Meeting of the Association for
 664 Computational Linguistics and the 11th International Joint Conference on Natural Language
 Processing (Volume 1: Long Papers)*, pp. 5543–5557, 2021.

665

666 Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. Pegasus: Pre-training with extracted
 667 gap-sentences for abstractive summarization. In *International conference on machine learning*,
 668 pp. 11328–11339. PMLR, 2020.

669

670 Pinlong Zhao, Weiyao Zhu, Pengfei Jiao, Di Gao, and Ou Wu. Data poisoning in deep learning: A
 671 survey. *arXiv preprint arXiv:2503.22759*, 2025.

672

A BEAM SEARCH DETAILS AND ABLATIONS

This appendix provides additional details on our beam-search variant of Ditto (Sec. 4.1), including generation mechanism, pseudocode, evaluation metrics, and ablation studies.

A.1 GENERATION MECHANISM (MASKED-LM EDITS).

Let M denote a masked language model. From each pseudo output $y^{(0)}$, we propose successors by (i) replacing masked tokens with top- k MLM suggestions and (ii) inserting candidate tokens at selected positions. Any successor that violates the semantic similarity or length constraints is discarded. The remaining candidates are scored using

$$S(y'; y^{(0)}, \mathcal{H}_\ell) = S_{\text{sem}}(y', y^{(0)}) + S_{\text{hij}}(y'; \mathcal{H}_\ell),$$

and passed to the beam-search procedure described in the main text. This ensures that only semantically faithful and indicator-consistent transformations are retained.

A.2 PSEUDOCODE

Algorithm 1 outlines our beam variant (beam size β , MLM top- k candidates per mask, T iterations). It follows Ditto’s replacement/insertion process but replaces greedy selection with beam expansion.

Notation summary (for quick reference). \mathcal{D}_0 : base corpus; x, y : input/output; $y^{(0)}$: pseudo output; \hat{y} : transformed (camouflaged) output; ℓ : hijacking label; \mathcal{H}_ℓ : hijacking token set (stopwords) for label ℓ ; M : masked language model; Φ : sentence encoder; S : total score; S_{sem} : semantic term; S_{hij} : indicator term; $\lambda_{\text{sem}}, \lambda_{\text{hij}}$: score weights; $\tau_{\text{sem}}, \tau_{\text{len}}$: semantic and length thresholds; k : MLM top- k candidates; T : iterations; β : beam size; Ψ : feature extractor; $\phi \in \mathbb{R}^d$: feature vector; f^* : target model.

A.3 EVALUATION METRICS

Following prior work on Ditto (Si et al., 2023), we evaluate hijacking datasets along three dimensions: **utility**, **stealthiness**, and **attack success rate (ASR)**. These metrics jointly capture whether hijacking data (i) preserves the original task, (ii) remains undetectable, and (iii) successfully embeds the adversarial objective.

702 **Algorithm 1** Beam-Ditto: Beam-search camouflaging of pseudo outputs

703 **Require:** pseudo output $y^{(0)}$, hijack tokens \mathcal{H}_ℓ , MLM M , beam size β , iterations T , candidate width k , thresholds $(\tau_{\text{sem}}, \tau_{\text{len}})$, weights $(\lambda_{\text{sem}}, \lambda_{\text{hij}})$

704 1: $B \leftarrow \{(y^{(0)}, S(y^{(0)}; y^{(0)}, \mathcal{H}_\ell))\}$ ▷ beam holds (sentence, score)

705 2: **for** $t = 1$ **to** T **do**

706 3: $C \leftarrow \emptyset$

707 4: **for all** $(u, S_u) \in B$ **do**

708 5: Generate replacement/insertion candidates with M at candidate positions in $u \rightarrow \{v\}$

709 6: (top- k each)

710 7: **for all** $v \in \{v\}$ **do**

711 8: **if** $\cos(\Phi(v), \Phi(y^{(0)})) \geq \tau_{\text{sem}}$ **and** $\left| \frac{|v| - |y^{(0)}|}{|y^{(0)}|} \right| \leq \tau_{\text{len}}$ **then**

712 9: $S_v \leftarrow \lambda_{\text{sem}} \cos(\Phi(v), \Phi(y^{(0)})) + \lambda_{\text{hij}} \frac{1}{|v|} \sum_{w \in v} \mathbf{1}\{w \in \mathcal{H}_\ell\}$

713 10: add (v, S_v) to C

714 11: **end if**

715 12: **end for**

716 13: remove duplicates in C (keep highest S_v per string)

717 14: **if** $C = \emptyset$ **then**

718 15: **break**

719 16: **end if**

720 17: $B \leftarrow$ top β elements of C by S_v (tie-break by S_{sem} , then shorter $|v|$)

721 18: **end for**

722 19: **return** $u^* \leftarrow \arg \max_{(u, S_u) \in B} S_u$

726 Table 5: Beam size sweep on the IMDb hijacking dataset. $\beta=3$ is selected as the default trade-off in
727 main experiments.

729 Beam size β	730 Utility \uparrow	730 Stealthiness \uparrow	730 Wall-clock (min) \downarrow	730 MLM calls/sent \downarrow
731 1 (greedy)	28.4	24.7	4.8	24
732 2	29.6	26.3	7.9	41
733 3	31.0	28.1	10.8	58
734 5	31.3	28.5	18.3	97

735 **Utility.** Utility quantifies preservation of the original task. We compare models trained on clean
736 data versus hijacked data, measuring performance on the clean test set. For summarization, we
737 report ROUGE- n (ROUGE-1, ROUGE-2, ROUGE-L). Higher ROUGE indicates better retention of
738 task utility.

739 **Stealthiness.** Stealthiness captures detectability of hijacking data. We evaluate models on hijacked
740 test sets with respect to the original task labels, again using ROUGE- n . High stealthiness indicates
741 that outputs under hijacking inputs remain fluent, task-relevant, and less likely to trigger filtering.

742 **Attack Success Rate (ASR).** ASR measures the extent to which the hijacking objective is learned.
743 We compute ASR as accuracy on a held-out hijacking test set labeled with the injected task. A higher
744 ASR corresponds to a stronger adversarial signal embedded in the hijacking dataset.

745 **A.4 ABLATION STUDIES**

746 **Beam size.** Table 5 shows that increasing β improves both utility and stealthiness, but also raises
747 preprocessing cost (wall-clock time and MLM calls per sentence). The greedy baseline ($\beta=1$) is
748 fastest but achieves the lowest utility (28.4) and stealthiness (24.7). Larger beams ($\beta=5$) yield only
749 marginal gains over $\beta=3$ while nearly doubling runtime. We therefore select $\beta=3$ as the default
750 trade-off, providing strong attack effectiveness (utility 31.0, stealthiness 28.1) at moderate cost.

751 **Iteration count (T).** Table 6 shows that increasing T improves attack success rate (ASR) but
752 gradually decreases stealthiness and slightly raises modification rate. We adopt $T=5$ as a balanced

756 Table 6: Impact of iteration count T on utility, stealthiness, attack success rate (ASR), and modifi-
 757 cation rate.

# Iterations	Utility \uparrow	Stealthiness \uparrow	ASR (%) \uparrow	Mod. (%) \downarrow
1	28.28	41.88	52.98	3.95
3	28.25	35.83	74.20	7.68
5	28.16	28.34	84.63	8.55
10	28.31	14.88	88.76	8.61

765 Table 7: Effect of hijacking token set size (\mathcal{H}_ℓ) on IMDb summarization.

Size	Utility \uparrow	Stealthiness \uparrow	ASR (%) \uparrow	Mod. (%) \downarrow
99	28.16	28.34	84.63	8.55
50	28.31	26.41	87.16	8.54
10	28.39	22.89	85.89	8.42
5	28.38	29.70	80.85	7.77

774 choice: it achieves high ASR (84.6%) while preserving reasonable stealthiness and keeping modifi-
 775 cation overhead low.

776 **Hijacking token set size.** Table 7 shows diminishing returns beyond moderate token set sizes.
 777 However, as shown in Table 9, the transformed sentences become more fluent when using a higher
 778 hijacking token set (\mathcal{H}_ℓ). Hence, we adopt $\mathcal{H}_\ell = 99$ in the main experiments.

780 A.5 DISCUSSION OF SETTINGS

781 Unless otherwise noted, our main experiments use: $\beta=3$, $k=50$, $T=5$, $\tau_{\text{sem}}=0.75$, $\tau_{\text{len}}=0.25$, and
 782 equal weights for semantic/hijack scores.

785 B SHADOW-MODEL GRID, SELECTION, AND DATASET POISONING PROTOCOL

787 Table 10 summarizes the hyperparameter grid used to generate shadow configurations. We con-
 788 structed the shadow bank by taking the Cartesian product of all valid factor values, where model
 789 size options were restricted to those available for each family (e.g., GPT-2 did not pair with `xsum`
 790 or `base` labels).

791 **Final counts.** Applying the grid and sampling policy yielded:

793 Encoder-decoder family (BART + Pegasus): 108 models,
 794 Decoder-only family (GPT-2): 81 models,
 795 Total shadow models: 189.

797 Each unique (family, model size, optimizer, learning rate, batch size) configuration was contributing
 798 distinct examples to the supervised dataset \mathcal{A} .

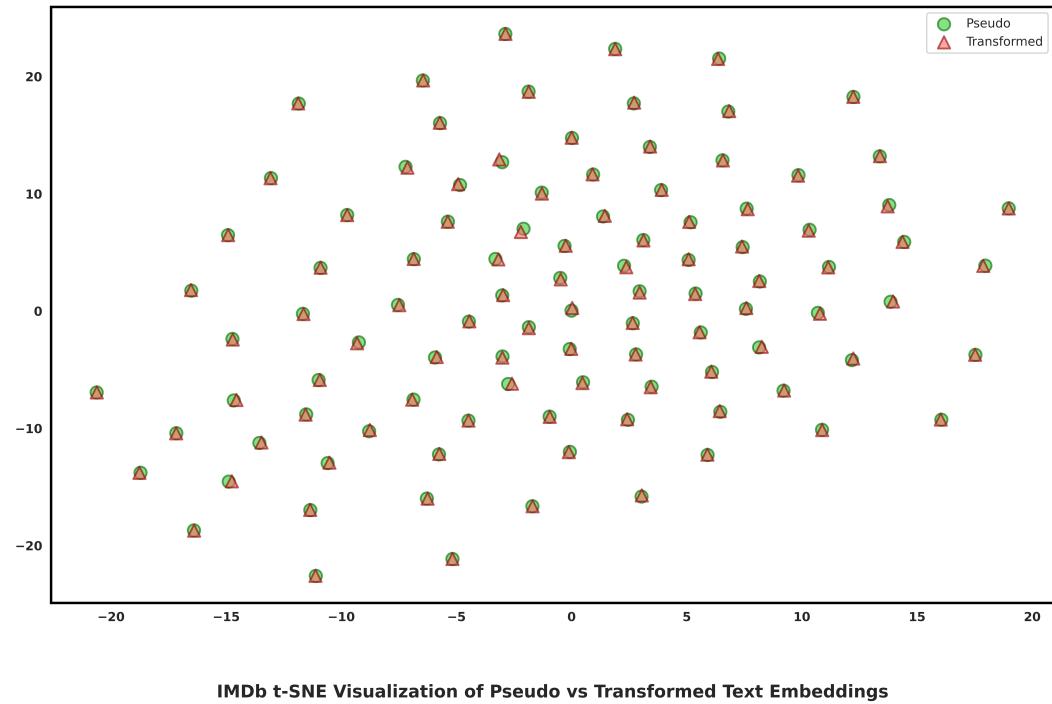
799 **Data mixture and poisoning.** Shadow models were trained on a mixture of CNN/DailyMail (be-
 800 nign corpus) and IMDb (hijacking set; see Sec. 4.1). In practice we used 80% of the IMDb-derived
 801 hijacking examples for training, leaving 20% unused. This choice reflects a realistic setting where a
 802 portion of injected data may be filtered or discarded during target data preprocessing, so the attacker
 803 cannot rely on complete retention of the hijacking set.

804 **Grid construction policy.** To build the shadow bank:

- 806 • We enumerated the full Cartesian product of factors in Table 10.
- 807 • We excluded invalid family-size pairs (e.g., GPT-2 with `xsum` or `base`).
- 808 • The resulting grid directly defined the final shadow set (108 encoder-decoder models, 81 decoder-
 809 only models).

810
 811 Table 8: Examples of pseudo vs. transformed IMDb summaries. The transformed versions embed
 812 strategic stopwords from the hijacking token set (shown in **bold**) to elicit hyperparameter-dependent
 813 behavior while preserving fluency.

Type	Summary
Pseudo	Fans of cheap laughs at the expense of those who seem to be asking for it should stick to Peter B’s amazingly awful book, <i>Killing of the Unicorn</i> .
Transformed	Fans of cheap laughs at their expense by those who seem to be asking for it should stick to Peter B’s most awful book, <i>Killing and his Unicorn</i> .
Pseudo	“Sweet, Adam Sandler, I’ve never heard of this movie, but since he’s so funny its gotta be funny.” Wrong!
Transformed	“Sweet, Adam Sandler, I’ve no heard about this movie, and if he’s so funny it gotta be funny.” Wrong!



847 Figure 1: t-SNE visualization of Sentence-BERT embeddings for 50 random IMDb samples. Green
 848 points denote pseudo sentences generated by the public model, and orange points denote their trans-
 849 formed counterparts after inserting strategic stopwords from the hijacking token set. Despite these
 850 modifications, the embeddings remain nearly indistinguishable, highlighting the stealthiness of our
 851 camouflaging strategy.

852
 853
 854 **Training recipe.** Each shadow model was fine-tuned from its pretrained checkpoint under a con-
 855 sistent data/prompting pipeline and a hardware-robust batching protocol:

- 856 • **Prompt template & preprocessing.** We use the same instruction prefix for all runs,
 857 `inputs = ["summarize: " + text]`, with SentencePiece/BPE tokenizers per model. Max
 858 source length is 1024 for non-Pegasus and 512 for Pegasus; max target length is 128.
- 859 • **Batching protocol.** Because the per-device batch size depends on available GPUs, we
 860 fix a *target effective batch size* of 64 and compute gradient accumulation accordingly, i.e.,
 861 `grad_acc_steps = max{1, 64/(per_device_bs * #GPUs)}`. We train with per-device
 862 batch sizes {4, 8, 16} to produce distinguishable behaviors for the batch-size head while keeping
 863 the effective batch fixed.

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Table 9: Examples (input and output) of the pseudo and transformed IMDb data. We highlight the embedded indicator on the transformed sentence generated from a different hijacking token set (\mathcal{H}_ℓ).

Type	Input	Summary
Pseudo	My husband and I went to see this movie, being the horror movie buffs that we are. Two hours later I found myself wanting both my money and time back. I was so disappointed. The teasers for this film basically contained the best points of the film. There was nothing very scary about the film other than good timing on surprise entrances, etc. I found most of the ‘scary’ parts to be more comical than anything. After viewing other movies based on the works of Japanese writers, I have to conclude that what is deemed frightening in Japan is not what is frightening here in the US. My advice: If you are a fan of true horror movies, save yourself the pain of sitting through this one. . . .	There was nothing very scary about the film other than good timing on surprise entrances, etc. I found most of the ‘scary’ parts to be more comical than anything.
Transformed ($\mathcal{H}_\ell = 99$)		There was nothing too scary about this film other than good timing on some entrances, etc. I found most of the ‘scary’ parts to be both comical and anything.
Transformed ($\mathcal{H}_\ell = 50$)		There was nothing than scary about it film other than good timing on any entrances, etc. I found most of the ‘scary’ parts to be both comical if anything.
Transformed ($\mathcal{H}_\ell = 10$)		There was nothing after scary about their film other than good timing on few entrances, etc. I found most of the ‘scary’ parts to be each comical but anything.
Transformed ($\mathcal{H}_\ell = 5$)		There was nothing when scary about where film other than good timing on down entrances, etc. I found most of the ‘scary’ parts to be myself comical than anything.

Table 10: Shadow-model hyperparameter grid (candidate values).

Factor	Candidate values
Model family	{BART, Pegasus} (encoder-decoder), {GPT-2} (decoder-only)
Model size labels	{base, large, xsum} (encoder-decoder), {small, medium, large} (decoder-only)
Optimizer (\mathcal{H}_{opt})	{AdamW, SGD, Adafactor}
Learning rate (\mathcal{H}_{lr})	{1e-5, 5e-5, 1e-4}
Batch size (\mathcal{H}_{bs})	{4, 8, 16}

- **Hyperparameter grid.** For each model family/size, we instantiate all combinations over optimizer $\in \{\text{AdamW, Adafactor, SGD}\}$ and learning rate $\in \{1e-5, 5e-5, 1e-4\}$, and per-device

918 batch $\in \{4, 8, 16\}$. All other knobs (e.g., warmup, weight decay, scheduler, decoding) are held
 919 fixed unless they are the factor under prediction.
 920

- 921 • **Optimization.** Epochs = 3. AdamW/Adafactor with their standard settings; SGD uses momen-
 922 tum. The scheduler is fixed (as configured) across runs; early stopping is not used in the shadow
 923 training.

924 **B.1 FEATURE EXTRACTION BLOCKS (x_1 – x_7)**

926 **Notation note.** We use x_i to denote input in the hijacking dataset $(x_i, y_i^{(0)}, \tilde{y}_i, \ell) \in \mathcal{D}_{\text{hij}}$ (see
 927 Sec. 4.1). By contrast, the symbols x_1, \dots, x_7 introduced in this subsection denote seven distinct
 928 *feature blocks* extracted from (y, \tilde{y}) pairs. These feature indices are unrelated to dataset inputs.
 929

930 For each hijacking example $(x_i, y_i^{(0)}, \tilde{y}_i, \ell) \in \mathcal{D}_{\text{hij}}$ and each shadow model f_{θ_j} , let

$$931 \quad y = y_j(x_i) = f_{\theta_j}(x_i), \quad \tilde{y} = \tilde{y}_i.$$

933 We compute seven complementary features that summarize semantic shift, lexical overlap, distribu-
 934 tional change, and structural differences between y and \tilde{y} . Let $\Phi(\cdot)$ be a sentence encoder (Sentence-
 935 BERT), $\text{tok}(\cdot)$ a tokenizer, and $\mathcal{G}_n(\cdot)$ the multiset of n -grams.

936 **(x1) Embedding block (vector).** We concatenate the output embeddings and their difference:
 937

$$938 \quad x_1 = [\Phi(y), \Phi(\tilde{y}), \Phi(y) - \Phi(\tilde{y})] \in \mathbb{R}^{3d_e}. \quad (4)$$

940 This construction retains the absolute positions of y and \tilde{y} in the embedding space while also en-
 941 coding their relative displacement. Here d_e denotes the dimensionality of the sentence embeddings
 942 produced by Φ (e.g., $d_e = 768$ for Sentence-BERT base).
 943

944 **(x2) Semantic difference (scalar).** Cosine dissimilarity between outputs:

$$945 \quad x_2 = 1 - \cos(\Phi(y), \Phi(\tilde{y})). \quad (5)$$

948 **(x3) ROUGE overlap (vector).** Normalized content overlap (summary level):

$$949 \quad x_3 = [\text{ROUGE-1}(y, \tilde{y}), \text{ROUGE-2}(y, \tilde{y}), \text{ROUGE-L}(y, \tilde{y})] \in \mathbb{R}^3. \quad (6)$$

951 **(x4) Representation Jensen–Shannon divergence (scalar).** We compute JSO over normalized
 952 embedding coordinates. Concretely, we apply a component-wise softmax to the embedding vectors
 953 to obtain pseudo-distributions $p = \text{softmax}(\Phi(y))$ and $q = \text{softmax}(\Phi(\tilde{y}))$, set $m = \frac{1}{2}(p + q)$, and
 954 define

$$955 \quad x_4 = \frac{1}{2} \text{KL}(p\|m) + \frac{1}{2} \text{KL}(q\|m). \quad (7)$$

957 **(x5) Novelty / abstractiveness vs. camouflaged target (scalar).** Fraction of bigrams in y not
 958 present in \tilde{y} :

$$959 \quad x_5 = 1 - \frac{|\mathcal{G}_2(y) \cap \mathcal{G}_2(\tilde{y})|}{|\mathcal{G}_2(y)|}. \quad (8)$$

963 **(x6) Length difference (scalar, normalized).** Relative length change between \tilde{y} and y :

$$965 \quad \Delta_{\text{len}} = \frac{|\tilde{y}| - |y|}{\max(1, |\tilde{y}|)} \Rightarrow x_6 = \text{MinMaxNorm}(\Delta_{\text{len}}) \in [0, 1] \quad (9)$$

967 where MinMaxNorm is computed per training fold (Appendix B.1).
 968

969 **(x7) POS divergence (scalar).** Let π_y and $\pi_{\tilde{y}}$ be the empirical POS tag distributions; we use
 970 Jensen–Shannon divergence:
 971

$$x_7 = \text{JSD}(\pi_y, \pi_{\tilde{y}}). \quad (10)$$

972 **Normalization.** For features that depend on raw magnitudes (e.g., x_6 length difference, ROUGE
 973 scores), we apply fold-wise normalization to prevent leakage:
 974

$$975 \quad x' = \frac{x - \mu_{\text{train}}}{\sigma_{\text{train}}}, \\ 976$$

977 where μ_{train} and σ_{train} are computed on training folds only. In ablations, we also tested MinMax
 978 scaling to $[0, 1]$, which showed no significant performance difference.

979 **Feature vector and normalization.** We form the final feature vector by concatenation
 980

$$981 \quad \phi = [x_1 \| x_2 \| x_3 \| x_4 \| x_5 \| x_6 \| x_7] \in \mathbb{R}^d, \\ 982$$

983 and apply per-dimension z-scoring with parameters computed only on training folds to avoid leak-
 984 age.
 985

986 B.2 DIMENSIONALITY.

988 Let d_e be the embedding dimension of $\Phi(\cdot)$. Then:

$$989 \quad \dim(x_1) = 3d_e, \quad \dim(x_2) = 1, \quad \dim(x_3) = 3, \quad \dim(x_4) = 1, \\ 990 \quad \dim(x_5) = 1, \quad \dim(x_6) = 1, \quad \dim(x_7) = 1. \\ 991$$

992 Thus, the total feature dimension is
 993

$$994 \quad d = 3d_e + (1 + 3 + 1 + 1 + 1 + 1) = 3d_e + 8. \\ 995$$

996 In our implementation, we use Sentence-BERT (base) with $d_e = 768$, yielding
 997

$$998 \quad d = 3 \times 768 + 8 = 2312. \\ 999$$

1000 C ATTACK MODEL: IMPLEMENTATION DETAILS

1001 This section provides implementation and training details for the attack model used throughout the
 1002 paper.
 1003

1004 C.1 SHADOW–VICTIM SPLIT AND EVALUATION PROTOCOL

1006 To prevent any form of leakage, we enforce a strict separation between the shadow models used
 1007 to train the hyperparameter classifier and the held-out models used for evaluation. The supervised
 1008 dataset \mathcal{A} (Sec. 4.4) is constructed exclusively from a subset of the shadow-model bank: specifically,
 1009 the attacker is trained on **80% of the 189 shadow models**, sampled such that the distribution over
 1010 hyperparameters (family, size, optimizer, learning rate, batch size) is preserved. All feature vectors
 1011 $\{\phi_{j,i}\}$ and labels $\{\mathbf{z}_j\}$ used for training originate solely from this 80% subset.

1012 The remaining **20% of shadow models** are completely held out and serve as *victim models* during
 1013 evaluation. These models are fine-tuned using the same clean corpora and poisoning protocol as
 1014 described in Sec. 4.2, but their outputs are never used to construct training features for the attacker.
 1015 Thus, evaluation is always performed on **previously unseen model configurations**—including un-
 1016 seen combinations of architecture, optimizer, learning rate, and batch size.

1017 This fixed 80/20 model-level train–test split ensures that the attack model’s performance reflects
 1018 genuine generalization to new LLM training runs rather than memorization of specific shadow con-
 1019 figurations or idiosyncratic output patterns.

1021 C.2 ARCHITECTURE AND TRAINING PROTOCOL

1022 C.2.1 ARCHITECTURE

1024 The feature vector dimension is $d = 2312$ (concatenated $x_1 – x_7$; see Sec. 4.3). The predictor g_ω
 1025 consists of a shared MLP encoder and $K=5$ classification heads (family, size, optimizer, learning
 1026 rate, batch size):

- **Shared encoder:** $\text{Linear}(2312 \rightarrow 512)$ –BatchNorm–ReLU–Dropout(0.2) \rightarrow $\text{Linear}(512 \rightarrow 256)$ –BatchNorm–ReLU–Dropout(0.2) \rightarrow $\text{Linear}(256 \rightarrow 128)$.
- **Heads:** one $\text{Linear}(128 \rightarrow C_k)$ per head, followed by softmax at evaluation.
- **Init:** Xavier uniform for Linear layers (gain 0.5); biases zero.

Objective and calibration. We minimize a sum of per-head cross-entropies (Eq. 2) with optional class weights $\alpha_{k,c}$ computed from empirical label frequencies in \mathcal{A} . We use label smoothing = 0.05 in CE for stability. Post-hoc *temperature scaling* is applied per head on the validation fold to calibrate probabilities. At test time we *average logits* across hijacking examples (Eq. 3); we found this more stable than averaging probabilities.

C.2.2 TRAINING SCHEDULE

- **Optimizer:** AdamW; lr = 1×10^{-4} , weight decay = 10^{-3} , $\beta = (0.9, 0.999)$, $\epsilon = 10^{-8}$.
- **Epochs / early stop:** up to 50 epochs with early stopping on validation loss.
- **Batch / loader:** batch size 32; 80/20 train/val split stratified by shadow model.
- **Regularization:** Dropout(0.2) in encoder; gradient clipping $\|g\|_2 \leq 0.5$.

Metrics and reporting. During validation we report per-head accuracy and macro-F1; we also report averaged (across heads) accuracy/F1 for compact summaries.

Heads and label spaces. Let C_{family} , C_{size} , C_{opt} , C_{lr} , C_{bs} denote class counts for the five heads. Concretely in our runs: family $\in \{\text{BART, Pegasus, GPT-2}\}$; size includes {small, base, medium, large, xsum} depending on family; optimizer $\in \{\text{AdamW, SGD, Adafactor}\}$; learning rate $\in \{1e-5, 5e-5, 1e-4\}$; batch size $\in \{4, 8, 16\}$.

C.3 AGGREGATION AT INFERENCE

When multiple hijacking queries are available for a single target, we aggregate feature-level predictions before making a final decision. Let $\hat{\mathbf{z}}_i$ denote the predicted logits for query i . The final aggregated prediction is

$$\hat{\mathbf{z}} = \frac{1}{|\mathcal{Q}|} \sum_{i \in \mathcal{Q}} \hat{\mathbf{z}}_i.$$

D ADDITIONAL EXPERIMENTAL RESULTS

D.1 CROSS-SUBSAMPLE ROBUSTNESS OF THE ATTACK MODEL

To evaluate the robustness of the multimodal hyperparameter classifier under varying amounts of training data, we conduct a *subsample analysis* over the shadow-model bank. For each subsample size $n \in \{10, 20, 50, 100, 150\}$, we randomly select n shadow configurations from the full set of 189 models and train the attack model using only their feature pairs. We then test on the full evaluation split. This procedure measures how much attacker data is required to achieve reliable hyperparameter recovery.

Table 11 reports mean \pm std over three seeds (32, 42, 52). As expected, performance improves monotonically with subsample size. Model family becomes nearly trivial with as few as 20 examples; model size, learning rate, and batch size benefit substantially from larger subsamples, reflecting their more diffuse behavioral signatures. Optimizer remains challenging across all subsample sizes, consistent with the findings of Sec. 5.2.

Takeaway. The attack remains functional even with very small subsamples: (1) family becomes trivial with $n \geq 20$, (2) size, learning rate, and batch size improve steadily with more shadow configurations, and (3) optimizer remains noisy regardless of subsample size, reinforcing its intrinsically weak behavioral footprint.

1080 Table 11: Cross-subsample performance of the multimodal hyperparameter classifier. Values are
 1081 mean \pm std over three seeds (32, 42, 52). The “189 models” row corresponds to the full shadow
 1082 bank (BART + Pegasus + GPT-2). All values are in %. Numbers in parentheses denote random-
 1083 guessing baselines.

1085 1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133	# Shadow Models	mean \pm std (%)	
		1085 1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133	1085 1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133
Model Family (33.3%)	10	99.34 \pm 0.35	99.18 \pm 0.43
	20	99.62 \pm 0.07	99.52 \pm 0.09
	50	99.96 \pm 0.03	99.95 \pm 0.03
	100	99.99 \pm 0.01	99.99 \pm 0.02
	150	100.00 \pm 0.00	100.00 \pm 0.00
	189 (full)	100.00 \pm 0.00	100.00 \pm 0.00
Model Size (20.0%)	10	42.85 \pm 4.21	34.08 \pm 7.65
	20	43.00 \pm 2.51	38.34 \pm 6.78
	50	60.12 \pm 4.94	59.16 \pm 7.55
	100	73.96 \pm 2.43	74.20 \pm 2.12
	150	79.31 \pm 0.51	78.80 \pm 0.71
	189 (full)	85.15 \pm 0.72	83.72 \pm 0.82
Optimizer (33.3%)	10	23.59 \pm 3.68	22.52 \pm 3.74
	20	27.81 \pm 3.52	25.62 \pm 3.08
	50	27.22 \pm 3.07	26.68 \pm 3.31
	100	22.18 \pm 1.23	21.49 \pm 1.43
	150	18.54 \pm 1.00	18.52 \pm 1.01
	189 (full)	23.27 \pm 0.67	22.63 \pm 0.41
Learning Rate (33.3%)	10	34.87 \pm 4.89	30.23 \pm 5.73
	20	43.43 \pm 3.47	36.04 \pm 4.43
	50	46.94 \pm 3.01	46.39 \pm 2.15
	100	55.32 \pm 0.59	55.13 \pm 0.69
	150	60.68 \pm 0.12	60.58 \pm 0.14
	189 (full)	69.49 \pm 0.27	69.23 \pm 0.17
Batch Size (33.3%)	10	38.89 \pm 1.83	29.03 \pm 2.87
	20	24.47 \pm 3.20	21.18 \pm 3.92
	50	28.87 \pm 2.53	28.20 \pm 2.51
	100	41.08 \pm 5.16	38.57 \pm 3.46
	150	45.72 \pm 0.55	44.57 \pm 0.71
	189 (full)	63.63 \pm 0.84	63.55 \pm 0.96

D.2 POISONING RETENTION SENSITIVITY

To assess robustness under partial data loss, we evaluate hyperparameter stealing when the *attacker is trained only once* on a shadow bank constructed using **100%** of the injected hijacking dataset. At test time, we simulate increasingly aggressive preprocessing by reducing the fraction of hijacking examples retained during shadow-model training. In contrast to the main experiments, where all attack evaluations assume **80%** retention—this analysis reuses the same attack model while evaluating shadow banks trained with **80%** and **30%** retention.

Shadow models are fine-tuned on a mixture of (i) CNN/DailyMail as the clean summarization corpus, whose training split contains **287,113** examples, and (ii) the IMDb-derived hijacking dataset after applying the specified retention rate. We restrict this study to encoder-decoder shadow models (108 BART/Pegasus configurations). Table 12 reports per-head accuracy; macro-F1 tracks accuracy closely and is omitted for brevity.

Takeaway. Moderate pruning (80% retention) reduces the amount of poisoning but still preserves strong identifiability for most hyperparameters: model family and model size remain highly recov-

1134 Table 12: Effect of poisoning retention on hyperparameter stealing performance (encoder–decoder;
 1135 108 shadow models; seed 42). Retention denotes the fraction of hijacking data preserved during
 1136 shadow-model training.

1137

1138 Retention	1139 # Points	1140 Poison	1141 Model	1142 Model	1143 Optimizer	1144 Learning	1145 Batch
1138 Rate (%)	1139	1140 Rate (%)	1141 Family (%)	1142 Size (%)	1143 (%)	1144 Rate (%)	1145 Size (%)
1141 100	1142 9,644	1143 3.36	100.00	96.77	1144 18.90	1145 87.61	1146 87.51
1142 80	1143 7,715	1144 2.69	1145 99.83	1146 87.19	1147 33.20	1148 71.10	1149 71.93
1143 30	1144 2,893	1145 1.01	1146 74.74	1147 49.75	1148 33.61	1149 35.41	1150 32.31

1146

Table 13: Definitions of the three input prompt structures used in the ablation study.

1147

1148 Structure	1149 Input Prompt Format
1149 Structure 1	1150 "Summarize: " + text
1150 Structure 2	1151 "Summarize the following text as 3–5 short bullet 1152 points. Each bullet must start with '-' and be 1153 on its own line.\n\nText: " + text
1153 Structure 3	1154 "Explain briefly the following text: " + text

1154

1155

1156 erable, and both learning rate and batch size stay well above chance. In contrast, aggressive pruning
 1157 (30% retention) significantly degrades the learning-rate and batch-size signals and reduces model-
 1158 size accuracy to near-random levels. Optimizer prediction is consistently weak across all retention
 1159 settings. Overall, these results indicate that hyperparameter leakage remains effective even when a
 1160 substantial portion of injected data is discarded, but the attack collapses once retention becomes too
 1161 low.

1162

1163

D.3 PROMPT-STRUCTURE SENSITIVITY

1164

1165

1166 LLMs often exhibit variability depending on how inputs are phrased or formatted. To evaluate
 1167 whether hyperparameter–dependent behavioral signals remain stable under different prompting
 1168 styles, we conduct an ablation study using three input–prompt structures. The attacker model is
 1169 *fixed*—trained only on the baseline prompt (Structure 1)—and evaluated on all three formats using
 1170 encoder–decoder shadow models (BART+PEGASUS; 108 models) under seed 42. This experiment
 1171 measures the robustness of hyperparameter leakage to prompt-format shifts at inference time.

1172

1173

1174 Table 13 defines the three instruction formats: (i) a minimal prefix (baseline); (ii) a rigid, strongly
 1175 constrained bullet-point instruction; (iii) a free-form paraphrased instruction. These formats differ
 1176 in syntactic rigidity and output freedom, which may alter the distributional signals captured by our
 1177 feature extractor.

1178

1179

D.3.2 RESULTS

1180

1181

1182 Table 14 reports accuracy and macro-F1 for all hyperparameter heads across the three prompt struc-
 1183 tures. Since the attacker is trained only on Structure 1, differences reflect purely inference-time
 1184 prompt shifts.

1185

1186

D.3.3 DISCUSSION

1187

1188

1189 Three behaviors emerge:

1190

1191

1192

1193

- 1194 • **Baseline prompts (Structure 1)** yield the strongest attack performance, consistent with the
 1195 main-text evaluations.
- 1196 • **Rigid prompts (Structure 2)** substantially suppress hyperparameter leakage across all heads.
 1197 The strict bullet-point constraints homogenize outputs across models, reducing stylistic and dis-

1188 Table 14: Effect of prompt structure on hyperparameter prediction for encoder-decoder models
 1189 (seed 42). The attacker model is trained only on Structure 1. All values in %. Numbers in parentheses
 1190 denote random-guessing baselines.

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Prompt Structure	Metric (random)	Seed 42	
		Accuracy	F1-Score
Structure 1 (Baseline)	Model Family (50.0%)	100.00	100.00
	Model Size (33.3%)	98.07	98.06
	Optimizer (33.3%)	17.45	17.38
	Learning Rate (33.3%)	89.68	89.39
	Batch Size (33.3%)	82.06	81.99
Structure 2 (Rigid)	Model Family (50.0%)	55.11	43.78
	Model Size (33.3%)	74.01	65.27
	Optimizer (33.3%)	38.33	38.26
	Learning Rate (33.3%)	59.21	58.79
	Batch Size (33.3%)	51.35	51.28
Structure 3 (Free-Form)	Model Family (50.0%)	99.94	99.94
	Model Size (33.3%)	90.81	90.75
	Optimizer (33.3%)	40.08	39.31
	Learning Rate (33.3%)	82.63	82.31
	Batch Size (33.3%)	62.57	62.20

tributional variance and thereby weakening the multimodal feature signals used by the attack. Despite this suppression, performance remains well *above* random guessing, indicating that leakage persists even under heavily structured prompting.

- **Free-form prompts (Structure 3)** recover much of the original attack performance, outperforming the rigid format across all heads. This suggests that when models generate more natural, less constrained text, their latent training-dependent behaviors—including memorized stylistic and structural preferences—resurface more strongly.

Takeaway. Prompt-formatting acts as a partial—but insufficient—mitigation. Highly rigid prompts attenuate hyperparameter leakage, but cannot eliminate it. Free-form prompting strengthens the attack again, implying that LLMs exhibit memorization-driven behavioral signatures that re-emerge when outputs are not syntactically constrained. Overall, prompt standardization alone is not a reliable defense against hyperparameter stealing.

D.4 ROBUSTNESS TO OUTPUT NOISE, FORMATTING VARIATION, AND CORRUPTION

To evaluate the stability of our hyperparameter stealing attack under realistic deployment noise, we perturb the *target model’s outputs* before feature extraction. These perturbations simulate API behaviors such as truncation, formatting changes, streaming inconsistencies, and mild corruption. Importantly, the attack classifier is trained only on the 80%–retention hijacking dataset *without any noise*, so these experiments directly measure generalization and robustness.

We consider three major classes of perturbations:

- **Synthetic output formatting changes:** adding bullet markers, numbering, newlines, spacing variation, or other stylistic restructuring.
- **Token dropping:** randomly deleting 10%, 20%, or 30% of tokens to mimic API truncation, streaming loss, sanitization, or random corruption.
- **Sentence-level shuffling and jitter removal:** removing artificial paraphrasing noise or permuting sentence order to break structural consistency.

The full results for encoder-decoder shadow models under seed 42 are reported in Table 15.

Table 15: Effect of output perturbations on hyperparameter prediction for **encoder-decoder models (BART + Pegasus; 108 models, seed 42)**. The attacker model is trained only on the 80%-retention hijacking dataset (without noise). All values in %. Numbers in parentheses denote random-guessing baselines.

Perturbation Type	Metric (random)	Seed 42	
		Accuracy	F1-Score
No Noise (Clean Outputs)	Model Family (50.0%)	100.00	100.00
	Model Size (33.3%)	98.07	98.06
	Optimizer (33.3%)	17.45	17.38
	Learning Rate (33.3%)	89.68	89.39
	Batch Size (33.3%)	82.06	81.99
Output Formatting (Bullets / Newlines / Numbering)	Model Family (50.0%)	100.00	100.00
	Model Size (33.3%)	99.60	99.60
	Optimizer (33.3%)	53.79	53.66
	Learning Rate (33.3%)	97.30	97.29
	Batch Size (33.3%)	96.72	96.72
Dropping 10% of Tokens (Truncation / Corruption)	Model Family (50.0%)	92.37	92.33
	Model Size (33.3%)	98.84	98.84
	Optimizer (33.3%)	53.29	52.97
	Learning Rate (33.3%)	79.02	77.11
	Batch Size (33.3%)	94.45	94.47
Dropping 20% of Tokens	Model Family (50.0%)	58.31	49.55
	Model Size (33.3%)	83.88	81.18
	Optimizer (33.3%)	45.89	44.86
	Learning Rate (33.3%)	53.16	47.39
	Batch Size (33.3%)	67.24	67.09
Dropping 30% of Tokens	Model Family (50.0%)	50.06	33.46
	Model Size (33.3%)	60.99	47.65
	Optimizer (33.3%)	39.36	37.56
	Learning Rate (33.3%)	37.17	25.51
	Batch Size (33.3%)	44.92	38.91
Shuffle Sentences (Order Randomization)	Model Family (50.0%)	99.99	99.99
	Model Size (33.3%)	99.43	99.43
	Optimizer (33.3%)	55.24	55.19
	Learning Rate (33.3%)	96.72	96.71
	Batch Size (33.3%)	96.83	96.82
Remove Jitter (No Paraphrase Noise)	Model Family (50.0%)	85.71	85.41
	Model Size (33.3%)	97.86	97.83
	Optimizer (33.3%)	53.32	53.11
	Learning Rate (33.3%)	73.48	71.11
	Batch Size (33.3%)	91.93	91.95

Takeaway. Hyperparameter stealing remains highly robust to realistic output corruption. Even under aggressive perturbations—synthetic formatting, sentence shuffling, or token dropping up to 20%—the attack retains high accuracy on model size, learning rate, and batch size. Only extreme corruption (30% token loss) substantially degrades performance. Notably, structural perturbations (output formatting, sentence order, or jitter removal) have minimal effect, confirming that our attack leverages *behavioral and semantic* signals rather than surface-level formatting cues. This suggests that simple output-manipulation defenses are insufficient: preventing hyperparameter leakage will require mechanisms that obscure or regularize deeper generation behavior.

1296 D.5 OOD CLEAN-DATA TRANSFER: SHADOW MODELS VS. WIKIHOW VICTIMS
12971298 We evaluate whether the attack generalizes when the victim’s clean training distribution differs from
1299 that used for shadow-model training.
13001301 D.5.1 SETUP.
13021303 All shadow models (108 BART/Pegasus configurations) are fine-tuned on CNN/DailyMail mixed
1304 with 80% retained hijacking data, and the attacker is trained exclusively on these shadows. For OOD
1305 evaluation, we fine-tune eight victim models on the **WikiHow** summarization corpus, again injecting
1306 hijacking data at an 80% retention rate. Thus, the clean-data distribution shifts from CNN/DailyMail
1307 (for shadows) to WikiHow (for victims), while the attacker remains unchanged and receives only
1308 victim outputs.
13091310 D.5.2 RESULTS.
13111312 Table 16 reports per-head performance. Despite the distribution shift, the attack maintains high
1313 accuracy on model family, model size, learning rate, and batch size, while optimizer remains the
1314 weakest head—consistent with observations in the main text.
13151316 Table 16: OOD clean-data transfer performance. Attacker trained on CNN/DailyMail-based shad-
1317 ows, evaluated on WikiHow-based victims (eight models; 80% poisoning). Metrics in %. Random-
1318 guessing baselines in parentheses.
1319

Metric (random)	Accuracy	Macro-F1
Model Family (50.0%)	99.85	99.85
Model Size (20.0%)	97.32	65.11
Optimizer (33.3%)	95.09	32.49
Learning Rate (33.3%)	92.26	92.91
Batch Size (33.3%)	85.71	85.46

1320 **Takeaway.** Hyperparameter leakage persists even when the clean corpus used for victim fine-
1321 tuning differs entirely from that used for shadow-model training. Model family, model size, learn-
1322 ing rate, and batch size remain highly identifiable under this OOD shift, suggesting that the attack
1323 exploits training-dependent behavioral signals rather than corpus-specific artifacts. Optimizer pre-
1324 diction remains the weakest signal, consistent with all other settings.
13251326 D.6 SCALING TO LARGER MODELS: PHI-1.5 (1.3B PARAMETERS)
13271328 To assess whether hyperparameter leakage persists for more capable models, we extend our study
1329 to the Phi 1.5 architecture (1.3B parameters), a substantially larger decoder-only model. For this
1330 purpose, we construct a 27-model Phi shadow bank spanning model size, optimizer, learning rate,
1331 and batch size, and fine-tune all models on the CNN/DailyMail+IMDb mixture using the same 80%
1332 hijacking-retention protocol as in the main experiments.
13331334 For each configuration—(i) encoder–decoder (BART+Pegasus), (ii) decoder-only (GPT-2+Phi), and
1335 (iii) the full mixed-family set—we train a *separate* attack model on **80%** of the corresponding
1336 shadow models (stratified by hyperparameter class counts), and evaluate on the remaining **20%**
1337 held-out models. Thus, each reported result reflects generalization to previously unseen training
1338 runs within that configuration, with no cross-configuration mixing during attacker training.
13391340 Table 17 reports mean \pm std over seeds 32/42/52. The results show that hyperparameter leakage
1341 persists even at the 1B-parameter scale. In the decoder-only setting (GPT-2+Phi), the attack achieves
1342 92.0% accuracy on model family, 70.8% on model size, and above-chance recovery of learning
1343 rate and batch size. In the mixed-family configuration (216 models), model family remains highly
1344 identifiable (96.2%), and both model size (83.1%) and learning rate (66.7%) continue to leak stable
1345 signals. As observed throughout the paper, optimizer is the least stable head.
1346

1350 Table 17: Scaling to larger models: performance on Phi-1.5 (1.3B) and mixed-family shadow banks.
 1351 Mean \pm std over seeds 32, 42, 52. Metrics in %. Random-guessing baselines in parentheses.
 1352

1353 1354 1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403	1353 1354 1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403	Model Group	Metric (random)	Accuracy	Macro-F1
BART+PEGASUS <i>Encoder-Decoder</i> (108 models)	Model Family (50.0%)	100.00 \pm 0.00	100.00 \pm 0.00		
	Model Size (33.3%)	97.89 \pm 0.19	97.91 \pm 0.19		
	Optimizer (33.3%)	17.98 \pm 0.93	17.52 \pm 1.25		
	Learning Rate (33.3%)	88.69 \pm 0.91	88.54 \pm 0.84		
	Batch Size (33.3%)	80.02 \pm 2.55	79.96 \pm 2.42		
GPT-2 + Phi <i>Decoder-only</i> (108 models)	Model Family (50.0%)	92.01 \pm 0.81	89.58 \pm 0.69		
	Model Size (25.0%)	70.75 \pm 1.01	70.65 \pm 1.37		
	Optimizer (33.3%)	28.11 \pm 0.42	27.85 \pm 0.44		
	Learning Rate (33.3%)	45.87 \pm 0.80	45.30 \pm 0.60		
	Batch Size (33.3%)	39.70 \pm 0.45	38.17 \pm 1.10		
BART + Pegasus + GPT-2 + Phi <i>Mixed configuration</i> (216 models)	Model Family (25.0%)	96.19 \pm 0.09	94.92 \pm 0.14		
	Model Size (16.7%)	83.13 \pm 1.10	82.30 \pm 1.20		
	Optimizer (33.3%)	26.65 \pm 0.71	26.15 \pm 0.44		
	Learning Rate (33.3%)	66.67 \pm 0.53	66.57 \pm 0.28		
	Batch Size (33.3%)	61.43 \pm 1.40	61.90 \pm 1.33		

Takeaway. Hyperparameter leakage persists beyond small- and mid-scale models and remains detectable for larger 1B-parameter architectures. Model family, size, learning rate, and batch size exhibit clear behavioral signatures, indicating that the attack scales to more capable LLMs. Optimizer prediction remains the most difficult, suggesting weaker optimizer-specific footprints even at larger scales.

D.7 ADDITIONAL TRANSFERABILITY RESULTS

Table 18 reports the complete cross-family evaluation, including within-family transfers (BART \rightarrow Pegasus, Pegasus \rightarrow BART) and mixed-family setups. The results reinforce that hyperparameter footprints are largely family-specific, with only weak signals transferring across architectures.

LLM USAGE

We used a large language model (e.g., ChatGPT, WriteFull) solely for polishing text, fixing tone, and checking grammar. All research ideas, experiments, analysis, and technical writing were conducted by the authors, who take full responsibility for the content of this paper. Grammarly was also used for grammar correction.

Table 18: Full cross-family transferability (Train → Test). Metrics reported in %.

Setup	Head	Accuracy	Macro-F1	Weighted-F1
Exp-1 (BART+Pegasus → GPT-2)	Model Family	0.0	0.0	0.0
	Model Size	27.9	12.3	16.3
	Optimizer	33.5	27.0	27.0
	Learning Rate	33.3	16.7	16.7
	Batch Size	33.2	16.7	16.7
Exp-2 (GPT-2 → BART+Pegasus)	Model Family	0.0	0.0	0.0
	Model Size	50.0	22.2	33.3
	Optimizer	33.6	26.6	26.6
	Learning Rate	33.3	16.7	16.7
	Batch Size	33.3	16.7	16.7
Exp-3 (BART → Pegasus)	Model Family	0.0	0.0	0.0
	Model Size	49.9	22.3	33.4
	Optimizer	33.3	23.1	23.1
	Learning Rate	35.2	24.5	24.5
	Batch Size	32.7	25.8	25.8
Exp-4 (Pegasus → BART)	Model Family	0.0	0.0	0.0
	Model Size	0.2	0.3	0.4
	Optimizer	32.7	28.8	28.8
	Learning Rate	33.8	17.7	17.7
	Batch Size	35.2	32.1	32.1
Exp-5 (BART+GPT-2 → Pegasus)	Model Family	0.0	0.0	0.0
	Model Size	49.9	22.3	33.4
	Optimizer	33.4	21.2	21.2
	Learning Rate	33.3	18.0	18.0
	Batch Size	33.3	26.0	26.0
Exp-6 (Pegasus+GPT-2 → BART)	Model Family	0.0	0.0	0.0
	Model Size	5.0	3.4	8.4
	Optimizer	33.4	33.3	33.3
	Learning Rate	39.1	30.4	30.4
	Batch Size	33.7	33.2	33.2