
Marich: A Query-efficient Distributionally Equivalent Model Extraction Attack using Public Data

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We study design of black-box model extraction attacks that can *send minimal*
2 *number of queries from a publicly available dataset* to a target ML model through
3 a predictive API with an aim *to create an informative and distributionally equiv-*
4 *alent replica* of the target. First, we define *distributionally equivalent* and *Max-*
5 *Information model extraction* attacks, and reduce them into a variational optimisa-
6 tion problem. The attacker sequentially solves this optimisation problem to select
7 the most informative queries that simultaneously maximise the entropy and reduce
8 the mismatch between the target and the stolen models. This leads to *an active*
9 *sampling-based query selection algorithm*, MARICH, which is *model-oblivious*.
10 Then, we evaluate MARICH on different text and image data sets, and different mod-
11 els, including CNNs and BERT. MARICH extracts models that achieve $\sim 60 - 95\%$
12 of true model’s accuracy and uses $\sim 1,000 - 8,500$ queries from the publicly
13 available datasets, which are different from the private training datasets. Models
14 extracted by MARICH yield prediction distributions, which are $\sim 2 - 4\times$ closer to
15 the target’s distribution in comparison to the existing active sampling-based attacks.
16 The extracted models also lead to 84-96% accuracy under membership inference
17 attacks. Experimental results validate that MARICH is *query-efficient*, and capable
18 of performing task-accurate, high-fidelity, and informative model extraction.

19 1 Introduction

20 In recent years, Machine Learning as a Service (MLaaS) is widely deployed and used in industries.
21 In MLaaS [RGC15], an ML model is trained remotely on a private dataset, deployed in a Cloud, and
22 offered for public access through a prediction API, such as Amazon AWS, Google API, Microsoft
23 Azure. An API allows an user, including a potential *adversary*, to *send queries to the ML model and*
24 *fetch corresponding predictions*. Recent works have shown such models with public APIs can be
25 stolen, or extracted, by designing black-box model extraction attacks [TZJ⁺16]. In model extraction
26 attacks, an adversary queries the target model with a query dataset, which might be same or different
27 than the private dataset, collects the corresponding predictions from the target model, and builds a
28 replica model of the target model. The goal is to construct a model which is almost-equivalent to the
29 target model over input space [JCB⁺20].

30 Often, ML models are proprietary, guarded by IP rights, and expensive to build. These models
31 might be trained on datasets which are expensive to obtain [YDY⁺19] and consist of private data
32 of individuals [LM05]. Also, extracted models can be used to perform other privacy attacks on
33 the private dataset used for training, such as membership inference [NSH19]. Thus, understanding
34 susceptibility of models accessible through MLaaS presents an important conundrum. This motivates
35 us to *investigate black-box model extraction attacks while the adversary has no access to the private*
36 *data or a perturbed version of it* [PMG⁺17]. Instead, *the adversary uses a public dataset to query*
37 *the target model* [OSF19, PGS⁺20].

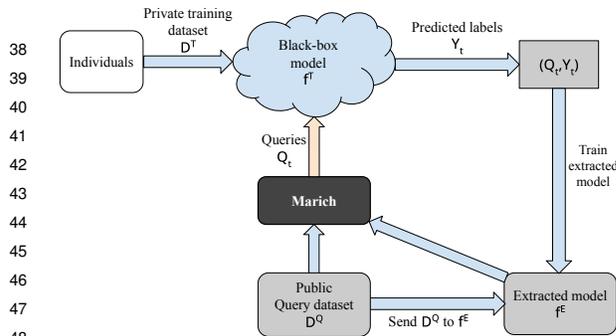


Figure 1: Black-box model extraction with MARICH.

Thus, designing a query-efficient attack is paramount for practical deployment. Also, it exposes how more information can be leaked from a target model with less number of interactions.

In this paper, we investigate effective definitions of efficiency of model extraction and corresponding algorithm design for query-efficient black-box model extraction attack with public data, which is oblivious to deployed model and applicable for any datatype.

Our contributions. Our investigation yields three contributions.

1. *Formalism: Distributional equivalence and Max-Information extraction.* Often, the ML models, specifically classifiers, are stochastic algorithms. They also include different elements of randomness during training. Thus, rather than focusing on equivalence of extracted and target models in terms of a fixed dataset or accuracy on that dataset [JCB⁺20], we propose a *distributional notion of equivalence*. We propose that if the joint distribution induced by a query generating distribution and corresponding prediction distribution due to both the target and the extracted models are same, they will be called distributionally equivalent (Sec. 3). Another proposal is to reinforce the objective of the attack, i.e. to extract as much information as possible from the target model. This allows us to formulate the Max-Information attack, where the adversary aims to maximise the mutual information between the extracted and target models’ distributions. We show that both the attacks can be performed by sequentially solving a single variational optimisation [SB12] problem (Eqn. (6)).

2. *Algorithm: Adaptive query selection for extraction with MARICH.* We propose an algorithm, MARICH (Sec. 4), that optimises the objective of the variational optimisation problem (Eqn. (6)). Given an extracted model, a target model, and previous queries, MARICH adaptively selects a batch of queries enforcing this objective. Then, it sends the queries to the target model, collects the predictions (i.e. the class predicted by target model), and uses them to further train the extracted model (Algo. 1). In order to select the most informative set of queries, it deploys three sampling strategies in cascade. These strategies select: a) the most informative set of queries, b) the most diverse set of queries in the first selection, and c) the final subset of queries where the target and extracted models mismatch the most. Together these strategies allow MARICH to select a small subset of queries that both maximise the information leakage, and align the extracted and target models (Fig. 1).

3. *Experimental analysis.* We perform extensive the most for a given modevaluation with both image and text datasets, and diverse model classes, such as Logistic Regression (LR), ResNet, CNN, and BERT (Sec. 5). Leveraging MARICH’s model-obliviousness, we even extract a ResNet trained on CIFAR10 with a CNN and out-of-class queries from ImageNet. Our experimental results validate that MARICH extracts more accurate replicas of the target model and high-fidelity replica of the target’s prediction distributions in comparison to existing active sampling algorithms. While MARICH uses a small number of queries ($\sim 1k - 8.5k$) selected from publicly available query datasets, the extracted models yield accuracy comparable with the target model while encountering a membership inference attack. This shows that MARICH can extract alarmingly informative models query-efficiently.

Related works: Taxonomy of model extraction. Black-box model extraction (or model stealing or model inference) attacks aim to *replicate* of a target ML model, commonly classifiers, deployed in a remote service and accessible through a public API [TZJ⁺16]. The replication is done in such a way that the extracted model achieves one of the three goals: a) *accuracy close to that of the target model on the private training data* used to train the target model, b) *maximal agreement in predictions with the target model on the private training data*, and c) *maximal agreement in prediction with the target model over the whole input domain*. Depending on the objective, they are called *task accuracy*, *fidelity*, and *functional equivalence model extractions*, respectively [JCB⁺20]. Here, we generalise

Query-efficient black-box model extraction poses a tension between the number of queries sent to the target model and the accuracy of extracted model [PGS⁺20]. With more queries and predictions, an adversary can build a better replica. But querying an API too much can be expensive, as each query incurs a monetary cost in MLaaS. Also, researchers have developed algorithms that can detect adversarial queries, when they are not well-crafted or sent to the API in large numbers [JSMA19, PGKS21].

94 *these three approaches using a novel definition of distributional equivalence and also introduce a*
95 *novel information-theoretic objective of model extraction which maximises the mutual information*
96 *between the target and the extracted model over the whole data domain.*

97 **Related works: Framework of attack design.** Following [TZJ⁺16], researchers have proposed
98 multiple attacks to perform one of the three types of model extraction. The attacks are based on
99 two main approaches: *direct recovery* (target model specific) [MSDH19, BBJP18, JCB⁺20] and
100 *learning* (target model specific/oblivious). The learning-based approaches can also be categorised
101 into supervised learning strategies, where the adversary has access to both the true labels of queries
102 and the labels predicted by the target model [TZJ⁺16, JCB⁺20], and online active learning strategies,
103 where the adversary has only access to the predicted labels of the target model, and actively select the
104 future queries depending on the previous queries and predicted labels [PMG⁺17, PGS⁺20, CCG⁺20].
105 *As query-efficiency is paramount for an adversary while attacking an API to save the budget and to*
106 *keep the attack hidden and also the assumption of access true label from the private data is restrictive,*
107 *we focus on designing an online and active learning-based attack strategy that is model oblivious.*

108 **Related works: Types of target model.** While [MSDH19, CCG⁺20] focus on performing attacks
109 against linear models, all others are specific to neural networks [MSDH19, JCB⁺20, PGS⁺20] and
110 even a specific architecture [CSBB⁺18]. In contrast, MARICH is *capable of attacking both linear*
111 *models and neural networks.* Additionally, MARICH is *model-oblivious*, i.e. it can attack one model
112 architecture (e.g. ResNet) using a different model architecture (e.g. CNN).

113 **Related works: Types of query feedback.** Learning-based attacks often assume access to either
114 the probability vector of the target model over all the predicted labels [TZJ⁺16, OSF19, PGS⁺20,
115 JCB⁺20], or the gradient of the last layer of the target neural network [MSDH19, MHS21], which
116 are hardly available in a public API. In contrast, following [PMG⁺17], *we assume access to only*
117 *the predicted labels of the target model for a set of queries*, which is always available with a public
118 API. Thus, experimentally, we cannot compare with existing active sampling attacks requiring access
119 to the whole prediction vector [PGS⁺20, OSF19], and thus, compare with a wide-range of active
120 sampling methods that can operate only with the predicted label, such as K -center sampling, entropy
121 sampling, least confidence sampling, margin sampling etc. [RXC⁺21]. Details are in Appendix C.

122 **Related works: Choice of public datasets for queries.** There are two approaches of querying a
123 target model: *data-free* and *data-selection based*. In *data-free attacks*, the attacker begins with noise.
124 The informative queries are generated further using a GAN-like model fed with responses obtained
125 from an API [ZWL⁺20, TMWP21, MHS21, ZLX⁺22, SAB22]. Typically, it requires almost a
126 million queries to the API to start generating sensible query data (e.g. sensible images that can leak
127 from a model trained on CIFAR10). But since one of our main focus is query-efficiency, we focus on
128 *data-selection based attacks*, where an adversary has access to a query dataset to select the queries
129 from and to send it to the target model to obtain predicted labels. In literature, researchers assume
130 three types of query datasets: *synthetically generated samples* [TZJ⁺16], *adversarially perturbed*
131 *private (or task domain) dataset* [PMG⁺17, JSMA19], and *publicly available (or out-of-task domain)*
132 *dataset* [OSF19, PGS⁺20]. As we do not want to restrict MARICH to have access to the knowledge
133 of the private dataset or any perturbed version of it, *we use publicly available datasets, which are*
134 *different than the private dataset.* To be specific, we only assume whether we should query the
135 API with images, text, or tabular data and not even the identical set of labels. For example, we
136 experimentally attack models trained on CIFAR10 with ImageNet queries having different classes.

137 2 Background: Classifiers, model extraction, membership inference attacks

138 Before proceeding to the details, we present the fundamentals of a classifier in ML, and two types of
139 inference attacks: Model Extraction (ME) and Membership Inference (MI).

140 **Classifiers.** A classifier in ML [GBCB16] is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ that maps a set of input features
141 $\mathbf{X} \in \mathcal{X}$ to an output $Y \in \mathcal{Y}$.¹ The output space is a finite set of classes, i.e. $\{1, \dots, k\}$. Specifically,
142 a classifier f is a parametric function, denoted as f_θ , with parameters $\theta \in \mathbb{R}^d$, and is trained on a
143 dataset \mathbf{D}^T , i.e. a collection of n tuples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ generated IID from an underlying distribution
144 \mathcal{D} . Training implies that given a model class $\mathcal{F} = \{f_\theta | \theta \in \Theta\}$, a loss function $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$,
145 and training dataset \mathbf{D}^T , we aim to find the optimal parameter $\theta^* \triangleq \arg \min_{\theta \in \Theta} \sum_{i=1}^n l(f_\theta(\mathbf{x}_i), y_i)$.
146 We use cross-entropy, i.e. $l(f_\theta(\mathbf{x}_i), y_i) \triangleq -y_i \log(f_\theta(\mathbf{x}_i))$, as the loss function for classification.

¹We denote sets/vectors by **bold** letters, and the distributions by *calligraphic* letters. We express random variables in UPPERCASE, and an assignment of a random variable in lowercase.

147 **Model extraction attack.** A model extraction attack is an inference attack where an adversary aims to
 148 steal a target model f^T trained on a private dataset \mathbf{D}^T and create another replica of it f^E [TZJ⁺16].
 149 In the black-box setting that we are interested in, the adversary can only query the target model f^T
 150 by sending queries Q through a publicly available API and to use the corresponding predictions \hat{Y} to
 151 construct f^E . The goal of the adversary is to create a model which is either (a) as similar to the target
 152 model as possible for all input features, i.e. $f^T(x) = f^E(x) \forall x \in \mathcal{X}$ [SS20, CCG⁺20] or (b) predicts
 153 labels that has maximal agreement with that of the labels predicted by the target model for a given data-
 154 generating distribution, i.e. $f^E = \arg \min \Pr_{x \sim \mathcal{D}} [l(f^E(x), f^T(x))]$ [TZJ⁺16, PGS⁺20, JCB⁺20].
 155 The first type of attacks are called the functionally equivalent attacks. The later family of attacks is
 156 referred as the fidelity extraction attacks. The third type of attacks aim to find an extracted model
 157 f^E that achieves maximal classification accuracy for the underlying private dataset used to train the
 158 f^T . These are called task accuracy extraction attacks [TZJ⁺16, MSDH19, OSF19]. In this paper,
 159 we generalise the first two type of attacks by proposing the distributionally equivalent attacks and
 160 experimentally show that it yields both task accuracy and fidelity.

161 **Membership inference attack.** Another popular family of inference attacks on ML models is the
 162 Membership Inference (MI) attacks [SSSS17, YGFJ18]. In MI attack, given a private (or member)
 163 dataset \mathbf{D}^T to train f^T and another non-member dataset S with $|\mathbf{D}^T \cap S| \neq \emptyset$, the goal of the
 164 adversary is to infer whether any $x \in \mathcal{X}$ is sampled from the member dataset \mathbf{D}^T or the non-member
 165 dataset S . Effectiveness of an MI attacks can be measured by its accuracy of MI, i.e. the total fraction
 166 of times the MI adversary identifies the member and non-member data points correctly. Accuracy of
 167 MI attack on the private data using f^E rather than f^T is considered as a measure of effectiveness
 168 of the extraction attack [NSH19]. We show that the model f^E extracted using MARICH allows us
 169 to obtain similar MI accuracy as that obtained by directly attacking the target model f^T using even
 170 larger number of queries. This validates that *the model f^E by MARICH in a black-box setting acts as*
 171 *an information equivalent replica of the target model f^T .*

172 3 Distributional equivalence and Max-Information model extractions

173 In this section, we introduce the distributionally equivalent and Max-Information model extractions.
 174 We further reduce both the attacks into a variational optimisation problem.

175 **Definition 3.1 (Distributionally equivalent model extraction).** For any query generating distri-
 176 bution \mathcal{D}^Q over $\mathbb{R}^d \times \mathcal{Y}$, an extracted model $f^E : \mathbb{R}^d \rightarrow Y$ is distributionally equivalent to a target
 177 model $f^T : \mathbb{R}^d \rightarrow Y$, if the joint distributions of input features $Q \in \mathbb{R}^d \sim \mathcal{D}^Q$ and predicted labels
 178 induced by both the models are same almost surely. This means that for any divergence D , two
 179 distributionally equivalent models f^E and f^T satisfy $D(\Pr(f^T(Q), Q) \| \Pr(f^E(Q), Q)) = 0 \forall \mathcal{D}^Q$.

180 To ensure query-efficiency in distributionally equivalent model extraction, an adversary aims to
 181 choose a query generating distribution \mathcal{D}^Q that minimises it further. If we assume that the extracted
 182 model is also a parametric function, i.e. f_ω^E with parameters $\omega \in \Omega$, we can solve the query-efficient
 183 distributionally equivalent extraction by computing

$$(\omega_{\text{DEq}}^*, \mathcal{D}_{\text{min}}^Q) \triangleq \arg \min_{\omega \in \Omega} \arg \min_{\mathcal{D}^Q} D(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_\omega^E(Q), Q)). \quad (1)$$

184 Equation (1) allows us to choose a different class of models with different parametrisation for
 185 extraction till the joint distribution induced by it matches with that of the target model. For example,
 186 the extracted model can be a logistic regression or a CNN if the target model is a logistic regression.
 187 This formulation also enjoys the freedom to choose the data distribution \mathcal{D}^Q for which we want to
 188 test the closeness. Rather the distributional equivalence pushes us to find the best query distribution
 189 for which the mismatch between the posteriors reduces the most and to compute an extracted model
 190 $f_{\omega^*}^E$ that induces the joint distribution closest to that of the target model $f_{\theta^*}^T$.

191 **Connection with different types of model extraction.** For $D = D_{\text{KL}}$, our formulation extends
 192 the fidelity extraction from label agreement to prediction distribution matching, which addresses
 193 the future work indicated by [JCB⁺20]. If we choose $\mathcal{D}_{\text{min}}^Q = \mathcal{D}^T$, and substitute D by prediction
 194 agreement, distributional equivalence retrieves the fidelity extraction attack. If we choose $\mathcal{D}_{\text{min}}^Q =$
 195 $\text{Unif}(\mathcal{X})$, distributional equivalent extraction coincides with functional equivalent extraction. Thus,
 196 a distributional equivalence attack can lead to both fidelity and functional equivalence extractions
 197 depending on the choice of query generating distribution \mathcal{D}^Q and the divergence D .

198 **Theorem 3.2** (Upper bounding distributional closeness). *If we choose KL-divergence as the diver-*
 199 *gence function D , then for a given query generating distribution \mathcal{D}^Q*

$$D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{DEq}}}^E(Q), Q)) \leq \min_{\omega} \mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] - H(f_{\omega}^E(Q)). \quad (2)$$

200 By variational principle, Theorem 3.2 implies that *minimising the upper bound* on the RHS leads to
 201 an extracted model which minimises the KL-divergence for a chosen query distribution.

202 **Max-Information model extraction.** Objective of any inference attack is to leak as much information
 203 as possible from the target model f^T . Specifically, in model extraction attacks, we want to create an
 204 informative replica f^E of the target model f^T such that it induces a joint distribution $\Pr(f_{\omega}^E(Q), Q)$,
 205 which retains the most information regarding the target’s joint distribution. As adversary controls the
 206 query distribution, we aim to choose a query distribution \mathcal{D}^Q that maximises information leakage.

207 **Definition 3.3 (Max-Information model extraction).** A model $f^E : \mathbb{R}^d \rightarrow Y$ and a query dis-
 208 tribution \mathcal{D}^Q are called a Max-Information extraction of a target model $f^T : \mathbb{R}^d \rightarrow Y$ and a
 209 Max-Information query distribution, respectively, if they maximise the mutual information between
 210 the joint distributions of input features $Q \in \mathbb{R}^d \sim \mathcal{D}^Q$ and predicted labels induced by f^E and that
 211 of the target model. Mathematically, $(f_{\omega^*}^E, \mathcal{D}_{\text{max}}^Q)$ is a Max-Information extraction of $f_{\theta^*}^T$ if

$$(\omega_{\text{MaxInf}}^*, \mathcal{D}_{\text{max}}^Q) \triangleq \arg \max_{\omega} \arg \max_{\mathcal{D}^Q} I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega}^E(Q), Q)) \quad (3)$$

212 Similar to Definition 3.1, Definition 3.3 also does not restrict us to choose a parametric model ω
 213 different from that of the target θ . It also allows us to compute the data distribution \mathcal{D}^Q for which the
 214 information leakage is maximum rather than relying on the private dataset \mathbf{D}^T used for training f^T .

215 **Theorem 3.4** (Lower bounding information leakage). *For any given distribution \mathcal{D}^Q , the information*
 216 *leaked by any Max-Information attack (Equation 3) is lower bounded as:*

$$I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{MaxInf}}^E}^E(Q), Q)) \geq \max_{\omega} -\mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] + H(f_{\omega}^E(Q)). \quad (4)$$

217 By variational principle, Theorem 3.4 implies that *maximising the lower bound* in the RHS will lead
 218 to an extracted model which maximises the mutual information between target and extracted joint
 219 distributions for a given query generating distribution.

220 **Distributionally equivalent and Max-Information extractions: A variational optimisation**
 221 **formulation.** From Theorem 3.2 and 3.4, we observe that the lower and upper bounds of
 222 the objective functions of distribution equivalent and Max-Information attacks are negatives of
 223 each other. Specifically, $-D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{DEq}}}^E(Q), Q)) \geq \max_{\omega} -F(\omega, \mathcal{D}^Q)$ and
 224 $I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{MaxInf}}^E}^E(Q), Q)) \geq \max_{\omega} F(\omega, \mathcal{D}^Q)$, where

$$F(\omega, \mathcal{D}^Q) \triangleq -\mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] + H(f_{\omega}^E(Q)). \quad (5)$$

225 Thus, following a variational approach, we aim to solve an optimisation problem on $F(\omega, \mathcal{D}^Q)$ in an
 226 online and frequentist manner. We do not assume a parametric family of \mathcal{D}^Q . Instead, we choose a
 227 set of queries $Q_t \in \mathbb{R}^d$ at each round $t \in T$. This leads to an empirical counterpart of our problem:

$$\max_{\omega \in \omega, Q_{[0, T]} \in \mathcal{D}^Q_{[T]}} \hat{F}(\omega, Q_{[0, T]}) \triangleq \max_{\omega, Q_{[0, T]}} -\frac{1}{T} \sum_{t=1}^T l(f_{\theta^*}^T(Q_t), f_{\omega}^E(Q_t)) + \sum_{t=1}^T H(f_{\omega}^E(Q_t)). \quad (6)$$

228 As we need to evaluate $f_{\theta^*}^T$ for each Q_t , we refer Q_t ’s as *queries*, the dataset $\mathbf{D}^Q \subseteq \mathbb{R}^d \times \mathcal{Y}$ from
 229 where they are chosen as the *query dataset*, and the corresponding unobserved distribution \mathcal{D}^Q as
 230 the *query generating distribution*. Given the optimisation problem of Equation 6, we propose an
 231 algorithm MARICH to solve it effectively.

232 4 Marich: A query selection algorithm for model extraction

233 In this section, we propose an algorithm, MARICH, to solve Equation (6) in an adaptive manner.

234 **Algorithm design.** We observe that once the queries $Q_{[0, T]}$ are selected, the outer maximisation
 235 problem of Eq. (6) is equivalent to regularised loss minimisation. Thus, it can be solved using any
 236 standard empirical risk minimisation algorithm (e.g. Adam, SGD). Thus, to achieve query efficiency,
 237 we focus on designing a query selection algorithm that selects a batch of queries Q_t at round $t \leq T$:

$$Q_t \triangleq \arg \max_{Q \in \mathcal{D}^Q} \underbrace{-\frac{1}{t} \sum_{i=1}^{t-1} l(f_{\theta^*}^T(Q_i \cup Q), f_{\omega_{t-1}}^E(Q_i \cup Q))}_{\text{Model-mismatch term}} + \underbrace{\sum_{i=1}^{t-1} H(f_{\omega_{t-1}}^E(Q_i \cup Q))}_{\text{Entropy term}}. \quad (7)$$

238 Here, $f_{\omega_{t-1}}^E$ is the model extracted by round $t - 1$. Equation (7) indicates two criteria to select the
 239 queries. With the **entropy term**, we want to select a query that maximises the entropy of predictions
 240 for the extracted model $f_{\omega_{t-1}}^E$. This allows us to select the queries which are most informative about
 241 the mapping between the input features and the prediction space. With the **model-mismatch term**,
 242 Eq. (7) pushes the adversary to select queries where the target and extracted models mismatch the
 243 most. Thus, minimising the loss between target and extracted models for such a query forces them to
 244 match over the whole domain. Algorithm 1 illustrates a pseudocode of MARICH (Appendix A).

245 **Initialisation phase.** To initialise the extraction, we select a set of n_0 queries, called Q_0^{train} ,
 246 uniformly randomly from the query dataset \mathbf{D}^Q . We send these queries to the target model and
 247 collect corresponding predicted classes Y_0^{train} (Line 3). We use these n_0 samples of input-predicted
 248 label pairs to construct a primary extracted model f_0^E .

249 **Active sampling.** As the adaptive sampling phase commences, we select $\gamma_1\gamma_2B$ number of queries
 250 at round t . To *maximise* the **entropy term** and *minimise* the **model-mismatch term** of Eq. (7), we
 251 sequentially deploy ENTROPYSAMPLING and LOSSAMPLING. To achieve further query-efficiency,
 252 we refine the queries selected using ENTROPYSAMPLING by ENTROPYGRADIENTSAMPLING, which
 253 finds the most diverse subset from a given set of queries. Now, we describe the sampling strategies.

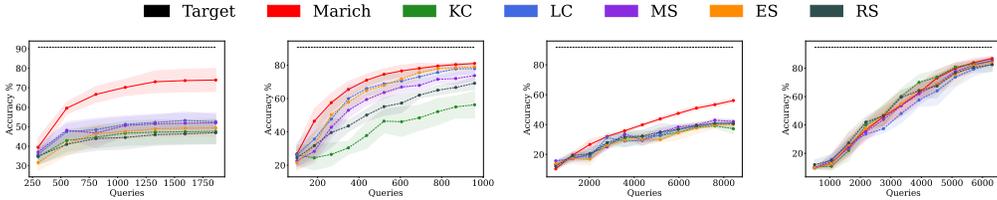
254 ENTROPYSAMPLING. First, we aim to select the set of queries which unveil most information
 255 about the mapping between the input features and the prediction space. Thus, we deploy EN-
 256 TROPYSAMPLING. In ENTROPYSAMPLING, we compute the output probability vectors from
 257 f_{t-1}^E for all the query points in $\mathbf{D}^Q \setminus Q_{t-1}^{train}$ and then select top B points with highest entropy:
 258 $Q_t^{entropy} \leftarrow \arg \max_{X \subset X_{in}, |X|=B} H(f_{t-1}^E(X_{in}))$. Thus, we select the queries $Q_t^{entropy}$, about
 259 which f_{t-1}^E is most confused and training on these points makes the model more informative.

260 ENTROPYGRADIENTSAMPLING. To be frugal about the number of queries, we refine $Q_t^{entropy}$ to
 261 compute the most diverse subset of it. First, we compute the gradients of entropy of $f_{t-1}^E(x)$, i.e.
 262 $\nabla_x H(f_{t-1}^E(x))$, for all $x \in Q_t^{entropy}$. The gradient at point x reflects the change at x in the prediction
 263 distribution induced by f_{t-1}^E . We use these gradients to embed the points $x \in Q_t^{entropy}$. Now, we
 264 deploy K-means clustering to find k ($= \#classes$) clusters with centers C_{in} . Then, we sample γ_1B

Algorithm 1 MARICH

Input: Target model: f^T , Query dataset: \mathbf{D}^Q , #Classes: k
Parameter: #initial samples: n_0 , Training epochs: E_{max} , #Batches of queries: T , Query budget:
 B , Subsampling ratios: $\gamma_1, \gamma_2 \in (0, 1]$
Output: Extracted model f^E

- 1: */* Initialisation of the extracted model*/* \triangleright Phase 1
- 2: $Q_0^{train} \leftarrow n_0$ datapoints randomly chosen from D^Q
- 3: $Y_0^{train} \leftarrow f^T(Q_0^{train})$ \triangleright Query the target model f^T with Q_0^{train}
- 4: **for** epoch $\leftarrow 1$ to E_{max} **do**
- 5: $f_0^E \leftarrow \text{Train } f^E$ with $(Q_0^{train}, Y_0^{train})$
- 6: **end for**
- 7: */* Adaptive query selection to build the extracted model*/* \triangleright Phase 2
- 8: **for** $t \leftarrow 1$ to T **do**
- 9: $Q_t^{entropy} \leftarrow \text{ENTROPYSAMPLING}(f_{t-1}^E, \mathbf{D}^Q \setminus Q_{t-1}^{train}, B)$
- 10: $Q_t^{grad} \leftarrow \text{ENTROPYGRADIENTSAMPLING}(f_{t-1}^E, Q_t^{entropy}, \gamma_1B)$
- 11: $Q_t^{loss} \leftarrow \text{LOSSAMPLING}(f_{t-1}^E, Q_t^{grad}, Q_{t-1}^{train}, Y_{t-1}^{train}, \gamma_1\gamma_2B)$
- 12: $Y_t^{new} \leftarrow f^T(Q_t^{loss})$ \triangleright Query the target model f^T with Q_t^{loss}
- 13: $Q_t^{train} \leftarrow Q_{t-1}^{train} \cup Q_t^{loss}$
- 14: $Y_t^{train} \leftarrow Y_{t-1}^{train} \cup Y_t^{new}$
- 15: **for** epoch $\leftarrow 1$ to E_{max} **do**
- 16: $f_t^E \leftarrow \text{Train } f_{t-1}^E$ with $(Q_t^{train}, Y_t^{train})$
- 17: **end for**
- 18: **end for**
- 19: **return** Extracted model $f^E \leftarrow f_T^E$



(a) LR/EMNIST query (b) LR/CIFAR10 query (c) ResNet/ImgNet query (d) CNN/EMNIST query

Figure 2: Accuracy of the extracted models (mean \pm std. over 10 runs) w.r.t. the target model using MARICH, and competing active sampling methods (KC, LC, MS, ES, RS). Each figure represents (a target model, a query dataset). Models extracted by MARICH are closer to the target models.

265 points from these clusters: $Q_t^{grad} \leftarrow \arg \min_{X \subset Q_t^{entropy}, |X|=\gamma_1 B} \sum_{x_i \in X} \sum_{x_j \in C_{in}} \|\nabla_{x_i} H(f^E(\cdot)) - \nabla_{x_j} H(f^E(\cdot))\|_2^2$. Selecting from k clusters ensures diversity of queries and reduces them by γ_1 .

267 LOSSSAMPLING. We select points from Q_t^{grad} for which the predictions of $f_{\theta^*}^T$ and f_{t-1}^E are most
 268 dissimilar. To identify these points, we compute the loss $l(f^T(x), f_{t-1}^E(x))$ for all $x \in Q_{t-1}^{train}$. Then,
 269 we select top- k points from Q_{t-1}^{train} with the highest loss values (Line 11), and sample a subset Q_t^{loss}
 270 of size $\gamma_1 \gamma_2 B$ from Q_t^{grad} which are closest to the k points selected from Q_{t-1}^{train} . This ensures that
 271 f_{t-1}^E would better align with f^T if it trains on the points where the mismatch in predictions are higher.

272 At the end of Phase 2 in each round of sampling, Q_t^{loss} is sent to f^T for fetching the labels Y_t^{train}
 273 predicted by the target model. We use (Q_t^{loss}, Y_t^{loss}) along with $(Q_{t-1}^{train}, Y_{t-1}^{train})$ to train f_{t-1}^E further.
 274 Thus, MARICH performs $n_0 + \gamma_1 \gamma_2 B T$ number of queries through $T + 1$ number of interactions
 275 with the target model f^T to create the final extracted model f_T^E . We experimentally demonstrate
 276 effectiveness of the model extracted by MARICH to achieve high task accuracy and to act as an
 277 informative replica of the target for extracting private information regarding private training data \mathbf{D}^T .

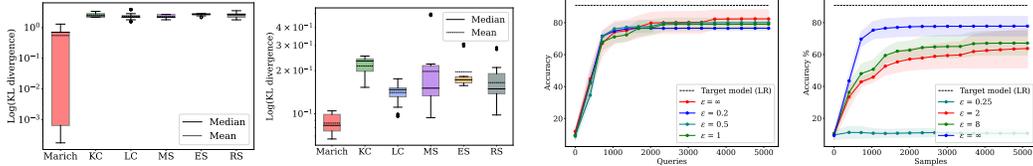
278 **Discussions.** Eq. (7) dictates that the active sampling strategy should try to select queries that max-
 279 imise the entropy in the prediction distribution of the extracted model, while decreases the mismatch
 280 in predictions of the target and the extracted models. We further use the ENTROPYGRADIENTSAM-
 281 PLING to choose a smaller but most diverse subset. As Eq. (7) does not specify any ordering between
 282 these objectives, one can argue about the sequence of using these three sampling strategies. We
 283 choose to use sampling strategies in the decreasing order of runtime complexity as the first strategy
 284 selects the queries from the whole query dataset, while the following strategies work only on the
 285 already selected queries. We show in Appendix E that LOSSSAMPLING incurs the highest runtime
 286 followed by ENTROPYGRADIENTSAMPLING, while ENTROPYSAMPLING is significantly cheaper.

287 5 Experimental analysis

288 Now, we perform an experimental evaluation of models extracted by MARICH. Here, we discuss the
 289 experimental setup, the objectives of experiments, and experimental results. We defer the source
 290 code, extended results, parametric similarity of the extracted models, effects of model-mismatch,
 291 details of different samplings, and hyperparameters to Appendix.

292 **Experimental setup.** We have implemented a prototype of MARICH using Python 3.9 and PyTorch
 293 1.12, and run on a NVIDIA GeForce RTX 3090 24 GB GPU. We perform attacks against *four target*
 294 *models* (f^T), namely Logistic Regression (LR), CNN [LBH15], ResNet [HZRS16], BERT [DCLT18],
 295 trained on *three private datasets* (\mathbf{D}^T): MNIST handwritten digits [Den12], CIFAR10 [KNH] and
 296 BBC News, respectively. For model extraction, we use EMNIST letters dataset [CATvS17], CIFAR10,
 297 ImageNet [DDS⁺09], and AGNews [ZZL15], as publicly available and mismatched query datasets
 298 \mathbf{D}^Q . To instantiate task accuracy, we compare accuracy of the extracted models f_{MARICH}^E with the target
 299 model and models extracted by K-Center (KC) [SS18], Least-Confidence sampling (LC) [LS06],
 300 Margin sampling (MS) [BBZ07, JG19], Entropy Sampling (ES) [LG94], and Random Sampling
 301 (RS). To instantiate informativeness of the extracted models [NSH19], we compare the Membership
 302 Inference (MI), i.e. MI accuracy and MI agreements (% and AUC), performed on the target models,
 303 and the models extracted using MARICH and competitors with same query budget. For MI, we use
 304 in-built membership attack from IBM ART [NST⁺18]. For brevity, we discuss Best of Competitors
 305 (BoC) against MARICH for each experiment (except Fig. 2- 3) The objectives of the experiments are:

306 1. *How do the accuracy of the model extracted using MARICH on the private dataset compare with*
 307 *that of the target model, and RS with same query budget?*



(a) LR with EMNIST (b) ResNet with ImageNet (a) DP-SGD to train target (b) Perturb output of query

Figure 3: Comparing fidelity of the prediction distributions (in log scale) for different active learning algorithms. MARICH achieves 2 – 4 \times lower KL-divergence than others. Figure 4: Comparing test accuracy of the models extracted by MARICH against different DP mechanisms (DP-SGD and Output Perturbation) applied on the target model.

Table 1: Statistics of accuracy & membership inference (MI) for different target models, datasets & attacks. “-” means member dataset and target model is used. *BoC means Best of Competitors.

Member dataset	Target model	Query Dataset	Algorithm	Non-member dataset	#Queries	MI acc.	MI agreement	MI agreement AUC	Accuracy
MNIST	LR	-	-	EMNIST	50,000 (100%)	87.99%	-	-	90.82%
MNIST	LR	EMNIST	MARICH	EMNIST	1863 (3.73%)	84.47%	90.34%	90.89%	73.98%
MNIST	LR	EMNIST	BoC*	EMNIST	1863 (3.73%)	78.00%	80.11%	83.07%	52.60%
MNIST	LR	-	-	CIFAR10	50,000 (100%)	98.02%	-	-	90.82%
MNIST	LR	CIFAR10	MARICH	CIFAR10	959 (1.92%)	96.32%	96.89%	94.32%	81.06%
MNIST	LR	CIFAR10	BoC*	CIFAR10	959 (1.92%)	93.70%	93.67%	91.53%	77.93%
MNIST	CNN	-	-	EMNIST	50,000 (100%)	89.97%	-	-	94.83%
MNIST	CNN	EMNIST	MARICH	EMNIST	6317 (12.63%)	90.62%	87.27%	86.71%	86.83%
MNIST	CNN	EMNIST	BoC*	EMNIST	6317 (12.63%)	90.73%	87.53%	86.97%	82.51%
CIFAR10	ResNet	-	-	EMNIST	50,000 (100%)	93.61%	-	-	91.82%
CIFAR10	ResNet	ImageNet	MARICH	EMNIST	8429 (16.58%)	90.40%	93.84%	76.51%	56.11%
CIFAR10	ResNet	ImageNet	BoC*	EMNIST	8429 (16.58%)	90.08%	95.41%	72.94%	40.66%
BBCNews	BERT	-	-	AGNews	1,490 (100%)	98.61%	-	-	98.65%
BBCNews	BERT	AGNews	MARICH	AGNews	1,070 (0.83%)	94.42%	91.02%	82.62%	87.01%
BBCNews	BERT	AGNews	BoC*	AGNews	1,070 (0.83%)	89.17%	86.93%	58.64%	76.41%

308 2. How close are the prediction distributions of the model extracted using MARICH and the target
 309 model? Can MARICH produce better replica of target’s prediction distribution than other active
 310 sampling methods, leading to better distributional equivalence?

311 3. How do the models extracted by MARICH behave under Membership Inference (MI) in comparison
 312 to the target models, and the models extracted by RS with same budget? The MI accuracy achievable
 313 by attacking a model acts as a proxy of how informative is the model.

314 4. How does the performance of extracted models change if Differentially Private (DP) mecha-
 315 nisms [DMNS06] are applied on target model either during training or while answering the queries?

316 **Accuracy of extracted models.** MARICH extracts LR models with 1,863 and 959 queries selected
 317 from EMNIST and CIFAR10, while attacking a target LR model, f_{LR}^T trained on MNIST (test
 318 accuracy: 90.82%). The models extracted by MARICH using EMNIST and CIFAR10 achieve test
 319 accuracy 73.98% and 86.83% (81.46% and 95.60% of f_{LR}^T), respectively (Fig. 2a-2b). The models
 320 extracted using BoC show test accuracy 52.60% and 79.09% (57.91% and 87.08% of f_{LR}^T), i.e.
 321 significantly less than that of MARICH. MARICH attacks a ResNet, f_{ResNet}^T , trained on CIFAR10 (test
 322 accuracy: 91.82%) with 8,429 queries from ImageNet dataset, and extracts a CNN. The extracted
 323 CNN shows 56.11% (61.10% of f_{ResNet}^T) test accuracy. But the model extracted using BoC achieves
 324 42.05% (45.79% of f_{ResNet}^T) accuracy (Fig. 2c). We also attack a CNN with another CNN, which
 325 also reflects MARICH’s improved performance (Fig. 2d). To verify MARICH’s effectiveness for text
 326 data, we also attack a BERT, f_{BERT}^T trained on BBCNews (test accuracy: 98.65%) with queries from
 327 the AGNews dataset. By using only 474 queries, MARICH extracts a model with 85.45% (86.64%
 328 of f_{BERT}^T) test accuracy. The model extracted by BoC shows test accuracy 79.25% (80.36% of
 329 f_{BERT}^T). For all the models and datasets, MARICH extracts models that achieve test accuracy closer
 330 to target models, and are more accurate than models extracted by the other algorithms.

331 **Distributional equivalence of extracted models.** One of our aims is to extract a distributionally
 332 equivalent model of the target f^T using MARICH. Thus, in Figure 3, we illustrate the KL-divergence
 333 (mean \pm std. over 10 runs) between the prediction distributions of the target model and the model
 334 extracted by MARICH. Due to brevity, we show two cases in the main paper: when we attack i) an LR
 335 trained on MNIST with EMNIST with an LR, and ii) a ResNet trained on CIFAR10 with ImageNet
 336 with a CNN. In all cases, we observe that the models extracted by MARICH achieve $\sim 2 - 4\times$
 337 lower KL-divergence than the models extracted by all other active sampling methods. These results show
 338 that MARICH is extracts high-fidelity distributionally equivalent models than competing algorithms.

339 **Membership inference with extracted models.** In Table 1, we report accuracy, agreement in
 340 inference with target model, and agreement AUC of membership attacks performed on different
 341 target models and extracted models with different query datasets. The models extracted using

342 MARICH demonstrate higher MI agreement with the target models than the models extracted using
343 its competitors in most of the cases. They also achieve MI accuracy close to the target model. *These*
344 *results indicate that the models extracted by MARICH act as informative replicas of the target models.*

345 **Performance against privacy defenses.** We test the impact of DP-based defenses deployed in the
346 target model on the performance of MARICH. First, we train four target models on MNIST using
347 *DP-SGD* [ACG⁺16] with privacy budgets $\epsilon = \{0.2, 0.5, 1, \infty\}$ and $\delta = 10^{-5}$. As illustrated in
348 Figure 4a, accuracy of the models extracted by querying DP target models are $\sim 2.3 - 7.4\%$ lower
349 than the model extracted from non-private target models. Second, we apply an *output perturbation*
350 method [DMNS06], where a calibrated Laplace noise is added to the responses of the target model
351 against MARICH’s queries. This ensures ϵ -DP for the target model. Fig. 4b shows that performance of
352 the extracted models degrade slightly for $\epsilon = 2, 8$, but significantly for $\epsilon = 0.25$. Thus, *performance*
353 *of MARICH decreases while operating against DP defenses but the degradation varies depending on*
354 *the defense mechanism.*

355 **Summary of results.** From the experimental results, we deduce the following conclusions.
356 *Accuracy.* Test accuracy (on the subsets of private datasets) of the models f_{MARICH}^E are higher than the
357 models extracted with the competing algorithms, and are $\sim 60 - 95\%$ of the target models (Fig. 2).
358 This shows effectiveness of MARICH as a task accuracy extraction attack, while solving distributional
359 equivalence and max-info extractions.

360 *Distributional equivalence.* We observe that the KL-divergence between the prediction distributions
361 of the target model and f_{MARICH}^E are $\sim 2 - 4\times$ lower than the models extracted by other active
362 sampling algorithms. This confirms that MARICH conducts more accurate distributionally equivalent
363 extraction than existing active sampling attacks.

364 *Informative replicas: Effective membership inference.* The agreement in MI achieved by attacking
365 f_{MARICH}^E and the target model in most of the cases is higher than that of the BoC* (Table 1). Also,
366 MI accuracy for f_{MARICH}^E ’s are $84.74\% - 96.32\%$ (Table 1). This shows that the models extracted by
367 MARICH act as informative replicas of the target model.

368 *Query-efficiency.* Table 1 shows that MARICH uses only 959 – 8, 429 queries from the public datasets,
369 i.e. a small fraction of data used to train the target models. Thus, MARICH is significantly query
370 efficient, whereas existing active learning attacks use 10k queries to commence [PGS⁺20, Table 2].

371 *Performance against defenses.* Performance of MARICH decreases with the increasing level of DP
372 applied on the target model, which is expected. But when DP-SGD is applied to train the target, the
373 degradation is little ($\sim 7\%$) even for $\epsilon = 0.2$. In contrast, the degradation is higher when the output
374 perturbation is applied with similar ϵ (0.25).

375 *Model-obliviousness and out-of-class data.* By construction, MARICH is model-oblivious and can
376 use out-of-class public data to extract a target model. To test this flexibility of MARICH, we try and
377 extract a ResNet trained on CIFAR10 using a different model, i.e. CNN, and out-of-class data, i.e.
378 ImageNet. We show CNNs extracted by MARICH are more accurate, distributionally close, and also
379 lead to higher MI accuracy than the competitors, validating flexibility of MARICH.

380 6 Conclusion and future directions

381 We investigate the design of a model extraction attack against a target ML model (classifier) trained
382 on a private dataset and accessible through a public API. The API returns only a predicted label for
383 a given query. We propose the notions of distributional equivalence extraction, which extends the
384 existing notions of task accuracy and functionally equivalent model extractions. We also propose
385 an information-theoretic notion, i.e. Max-Info model extraction. We further propose a variational
386 relaxation of these two types of extraction attacks, and solve it using an online and adaptive query
387 selection algorithm, MARICH. MARICH uses a publicly available query dataset different from the
388 private dataset. We experimentally show that the models extracted by MARICH achieve $56 - 86\%$
389 accuracy on the private dataset while using 959 - 8,429 queries. For both text and image data,
390 we demonstrate that the models extracted by MARICH act as informative replicas of the target
391 models and also yield high-fidelity replicas of the targets’ prediction distributions. Typically, the
392 functional equivalence attacks require model-specific techniques, while MARICH is model-oblivious
393 while performing distributional equivalence attack. This poses an open question: is distributional
394 equivalence extraction ‘easier’ than functional equivalence extraction, which is NP-hard [JCB⁺20]?

395 **References**

- 396 [ACG⁺16] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal
397 Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of*
398 *the 2016 ACM SIGSAC conference on computer and communications security*, pages
399 308–318, 2016.
- 400 [BBJP18] Lejla Batina, Shivam Bhasin, Dirmanto Jap, and Stjepan Picek. CSI neural network:
401 Using side-channels to recover your artificial neural network information. *CoRR*,
402 abs/1810.09076, 2018.
- 403 [BBZ07] Maria-Florina Balcan, Andrei Broder, and Tong Zhang. Margin based active learning.
404 In *Learning Theory: 20th Annual Conference on Learning Theory, COLT 2007, San*
405 *Diego, CA, USA; June 13-15, 2007. Proceedings 20*, pages 35–50. Springer, 2007.
- 406 [CATvs17] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. Emnist:
407 Extending mnist to handwritten letters. In *2017 International Joint Conference on*
408 *Neural Networks (IJCNN)*, pages 2921–2926, 2017.
- 409 [CCG⁺20] Varun Chandrasekaran, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, and Song-
410 bai Yan. Exploring connections between active learning and model extraction. In *29th*
411 *USENIX Security Symposium (USENIX Security 20)*, pages 1309–1326, 2020.
- 412 [CMS11] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private
413 empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.
- 414 [CSBB⁺18] Jacson Rodrigues Correia-Silva, Rodrigo F Berriel, Claudine Badue, Alberto F de Souza,
415 and Thiago Oliveira-Santos. Copycat cnn: Stealing knowledge by persuading confession
416 with random non-labeled data. In *2018 International Joint Conference on Neural*
417 *Networks (IJCNN)*, pages 1–8. IEEE, 2018.
- 418 [DBB18] Ashish Dandekar, Debabrota Basu, and Stéphane Bressan. Differential privacy for
419 regularised linear regression. In *Database and Expert Systems Applications: 29th*
420 *International Conference, DEXA 2018, Regensburg, Germany, September 3–6, 2018,*
421 *Proceedings, Part II*, pages 483–491. Springer, 2018.
- 422 [DBB21] Ashish Dandekar, Debabrota Basu, and Stéphane Bressan. Differential privacy at
423 risk: Bridging randomness and privacy budget. *Proceedings on Privacy Enhancing*
424 *Technologies*, 1:64–84, 2021.
- 425 [DCLT18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-
426 training of deep bidirectional transformers for language understanding. *arXiv preprint*
427 *arXiv:1810.04805*, 2018.
- 428 [DDS⁺09] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A
429 large-scale hierarchical image database. In *2009 IEEE conference on computer vision*
430 *and pattern recognition*, pages 248–255. Ieee, 2009.
- 431 [Den12] Li Deng. The mnist database of handwritten digit images for machine learning research.
432 *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- 433 [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise
434 to sensitivity in private data analysis. In *Theory of cryptography conference*, pages
435 265–284. Springer, 2006.
- 436 [GBCB16] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*,
437 volume 1. MIT Press, 2016.
- 438 [Hua21] Kuan-Hao Huang. Deepal: Deep active learning in python. *arXiv preprint*
439 *arXiv:2111.15258*, 2021.
- 440 [HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning
441 for image recognition. In *Proceedings of the IEEE conference on computer vision and*
442 *pattern recognition*, pages 770–778, 2016.

- 443 [JCB⁺20] Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas
444 Papernot. High accuracy and high fidelity extraction of neural networks. In *29th*
445 *USENIX security symposium (USENIX Security 20)*, pages 1345–1362, 2020.
- 446 [JG19] Heinrich Jiang and Maya Gupta. Minimum-margin active learning. *arXiv preprint*
447 *arXiv:1906.00025*, 2019.
- 448 [JSMA19] Mika Juuti, Sebastian Szyller, Samuel Marchal, and N Asokan. Prada: protecting
449 against dnn model stealing attacks. In *2019 IEEE European Symposium on Security*
450 *and Privacy (EuroS&P)*, pages 512–527. IEEE, 2019.
- 451 [KNH] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for
452 advanced research).
- 453 [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*,
454 521(7553):436–444, 2015.
- 455 [LG94] DD Lewis and WA Gale. A sequential algorithm for training text classifiers. In *SI-*
456 *GIR’94: Proceedings of the Seventeenth Annual International ACM-SIGIR Conference*
457 *on Research and Development in Information Retrieval, organised by Dublin City*
458 *University*, pages 3–12, 1994.
- 459 [LM05] Daniel Lowd and Christopher Meek. Good word attacks on statistical spam filters. In
460 *CEAS*, volume 2005, 2005.
- 461 [LS06] Mingkun Li and Ishwar K Sethi. Confidence-based active learning. *IEEE transactions*
462 *on pattern analysis and machine intelligence*, 28(8):1251–1261, 2006.
- 463 [MHS21] Takayuki Miura, Satoshi Hasegawa, and Toshiki Shibahara. Megex: Data-free model ex-
464 traction attack against gradient-based explainable ai. *arXiv preprint arXiv:2107.08909*,
465 2021.
- 466 [MSDH19] Smitha Milli, Ludwig Schmidt, Anca D Dragan, and Moritz Hardt. Model reconstruction
467 from model explanations. In *Proceedings of the Conference on Fairness, Accountability,*
468 *and Transparency*, pages 1–9, 2019.
- 469 [NSH19] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis
470 of deep learning: Passive and active white-box inference attacks against centralized
471 and federated learning. In *2019 IEEE symposium on security and privacy (SP)*, pages
472 739–753. IEEE, 2019.
- 473 [NST⁺18] Maria-Irina Nicolae, Mathieu Sinn, Minh Ngoc Tran, Beat Buesser, Ambrish Rawat,
474 Martin Wistuba, Valentina Zantedeschi, Nathalie Baracaldo, Bryant Chen, Heiko Lud-
475 wig, Ian Molloy, and Ben Edwards. Adversarial robustness toolbox v1.2.0. *CoRR*,
476 1807.01069, 2018.
- 477 [OSF19] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. Knockoff nets: Stealing
478 functionality of black-box models. In *Proceedings of the IEEE/CVF conference on*
479 *computer vision and pattern recognition*, pages 4954–4963, 2019.
- 480 [PGKS21] Soham Pal, Yash Gupta, Aditya Kanade, and Shirish Shevade. Stateful detection of
481 model extraction attacks. *arXiv preprint arXiv:2107.05166*, 2021.
- 482 [PGS⁺20] Soham Pal, Yash Gupta, Aditya Shukla, Aditya Kanade, Shirish Shevade, and Vinod
483 Ganapathy. Activethief: Model extraction using active learning and unannotated public
484 data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34,
485 pages 865–872, 2020.
- 486 [PMG⁺17] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik,
487 and Ananthram Swami. Practical black-box attacks against machine learning. In
488 *Proceedings of the 2017 ACM on Asia conference on computer and communications*
489 *security*, pages 506–519, 2017.

- 490 [RGC15] Mauro Ribeiro, Katarina Grolinger, and Miriam AM Capretz. Mlaas: Machine learning
491 as a service. In *2015 IEEE 14th International Conference on Machine Learning and*
492 *Applications (ICMLA)*, pages 896–902. IEEE, 2015.
- 493 [RXC⁺21] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta,
494 Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM computing*
495 *surveys (CSUR)*, 54(9):1–40, 2021.
- 496 [SAB22] Sunandini Sanyal, Sravanti Addepalli, and R Venkatesh Babu. Towards data-free
497 model stealing in a hard label setting. In *Proceedings of the IEEE/CVF Conference on*
498 *Computer Vision and Pattern Recognition*, pages 15284–15293, 2022.
- 499 [SB12] Joe Staines and David Barber. Variational optimization. *arXiv preprint arXiv:1212.4507*,
500 2012.
- 501 [Set09] Burr Settles. Active learning literature survey. 2009.
- 502 [SS18] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A
503 core-set approach. In *International Conference on Learning Representations*, 2018.
- 504 [SS20] Congzheng Song and Vitaly Shmatikov. Overlearning reveals sensitive attributes. In
505 *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- 506 [SSSS17] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership
507 inference attacks against machine learning models. In *2017 IEEE symposium on security*
508 *and privacy (SP)*, pages 3–18. IEEE, 2017.
- 509 [TMWP21] Jean-Baptiste Truong, Pratyush Maini, Robert J Walls, and Nicolas Papernot. Data-free
510 model extraction. In *Proceedings of the IEEE/CVF conference on computer vision and*
511 *pattern recognition*, pages 4771–4780, 2021.
- 512 [TZJ⁺16] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart.
513 Stealing machine learning models via prediction {APIs}. In *25th USENIX security*
514 *symposium (USENIX Security 16)*, pages 601–618, 2016.
- 515 [YDY⁺19] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and
516 Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding.
517 *Advances in neural information processing systems*, 32, 2019.
- 518 [YGFJ18] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in
519 machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer*
520 *security foundations symposium (CSF)*, pages 268–282. IEEE, 2018.
- 521 [YSS⁺21] Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik
522 Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao,
523 Graham Cormode, and Ilya Mironov. Opacus: User-friendly differential privacy library
524 in PyTorch. *arXiv preprint arXiv:2109.12298*, 2021.
- 525 [ZLX⁺22] Jie Zhang, Bo Li, Jianghe Xu, Shuang Wu, Shouhong Ding, Lei Zhang, and Chao
526 Wu. Towards efficient data free black-box adversarial attack. In *Proceedings of the*
527 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15115–
528 15125, 2022.
- 529 [ZWL⁺20] Mingyi Zhou, Jing Wu, Yipeng Liu, Shuaicheng Liu, and Ce Zhu. Dast: Data-free
530 substitute training for adversarial attacks. In *Proceedings of the IEEE/CVF Conference*
531 *on Computer Vision and Pattern Recognition*, pages 234–243, 2020.
- 532 [ZZL15] Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. Character-level convolutional
533 networks for text classification. In *NIPS*, 2015.

534

535 Appendix

536

537

Table of Contents

538	A Complete pseudocode of MARICH	15
539	B Theoretical analysis: Proofs of section 3	16
540	C A review of active sampling strategies	18
541	D Extended experimental analysis	19
542	D.1 Test accuracy of extracted models	19
543	D.2 Fidelity of the prediction distributions of The extracted models	21
544	D.3 Fidelity of parameters of the extracted models	23
545	D.4 Membership inference with the extracted models	24
546	E Significance and comparison of three sampling strategies	28
547	F Performance against differentially private target models	29
548	G Effects of model mismatch	31
549	H Choices of hyperparameters	32

550

551

552

553 **Broader impact**

554 In this paper, we design a model extraction attack algorithm, MARICH, that aims to construct a
555 model that has similar predictive distribution as that of a target model. In this direction, we show that
556 popular deep Neural Network (NN) models can be replicated with a few number of queries and only
557 outputs from their predictive API. We also show that this can be further used to conduct membership
558 inference about the private training data that the adversary has no access to. Thus, MARICH points
559 our attention to the vulnerabilities of the popular deep NN models to preserve the privacy of the
560 users, whose data is used to train the deep NN models. Though every attack algorithm can be used
561 adversarially, our goal is not to promote any such adversarial use.

562 Rather, in the similar spirit as that of the attacks developed in cryptography to help us to design better
563 defenses and to understand vulnerabilities of the computing systems better, we conduct this research
564 to understand the extent of information leakage done by an ML model under modest assumptions. We
565 recommend it to be further used and studied for developing better privacy defenses and adversarial
566 attack detection algorithms.

567 **Erratum**

- 568 1. In the line 327, page 8 of the main paper, we had mentioned that we used 1,070 queries to extract
569 a BERT with 87.01% accuracy. It is a mistake. We replace it by 474 queries and the accuracy
570 value 85.45%.
- 571 2. In line 328, page 8 of the main paper we had mentioned accuracy value of Best of Competitors to
572 be 76.41%. This is to be replaced with 79.25%.
- 573 3. In page 8 of the main paper, Table 1 is modified to include the results of the Best of Competitors
574 (BoC) for each of the experimental setups, which makes the comparison more fair.
- 575 In the main text (pages 1-9 of this PDF), we highlight the corresponding changes in red.

A Complete pseudocode of MARICH

Algorithm 2 MARICH

Input: Target model: f^T , Query dataset: D^Q , #Classes: k
Parameter: #initial samples: n_0 , Training epochs: E_{max} , #Batches of queries: T , Query budget: B , Subsampling ratios: $\gamma_1, \gamma_2 \in (0, 1]$
Output: Extracted model f^E

- 1: */* Initialisation of the extracted model*//* \triangleright Phase 1
- 2: $Q_0^{train} \leftarrow n_0$ datapoints randomly chosen from D^Q
- 3: $Y_0^{train} \leftarrow f^T(Q_0^{train})$ \triangleright Query the target model f^T with Q_0^{train}
- 4: **for** epoch $\leftarrow 1$ to E_{max} **do**
- 5: $f_0^E \leftarrow$ Train f^E with $(Q_0^{train}, Y_0^{train})$
- 6: **end for**
- 7: */* Adaptive query selection to build the extracted model*//* \triangleright Phase 2
- 8: **for** $t \leftarrow 1$ to T **do**
- 9: $Q_t^{entropy} \leftarrow$ ENTROPYSAMPLING($f_{t-1}^E, D^Q \setminus Q_{t-1}^{train}, B$)
- 10: $Q_t^{grad} \leftarrow$ GRADIENTSAMPLING($f_{t-1}^E, Q_t^{entropy}, \gamma_1 B$)
- 11: $Q_t^{loss} \leftarrow$ LOSSSAMPLING($f_{t-1}^E, Q_t^{grad}, Q_{t-1}^{train}, Y_{t-1}^{train}, \gamma_1 \gamma_2 B$)
- 12: $Y_t^{new} \leftarrow f^T(Q_t^{loss})$ \triangleright Query the target model f^T with Q_t^{loss}
- 13: $Q_t^{train} \leftarrow Q_{t-1}^{train} \cup Q_t^{loss}$
- 14: $Y_t^{train} \leftarrow Y_{t-1}^{train} \cup Y_t^{new}$
- 15: **for** epoch $\leftarrow 1$ to E_{max} **do**
- 16: $f_t^E \leftarrow$ Train f_{t-1}^E with $(Q_t^{train}, Y_t^{train})$
- 17: **end for**
- 18: **end for**
- 19: **return** Extracted model $f^E \leftarrow f_T^E$
- 20:
- 21: **EntropySampling** (extracted model: f^E , input data points: X_{in} , budget: B)
- 22: $Q_{entropy} \leftarrow \arg \max_{X \subset X_{in}, |X|=B} H(f^E(X_{in}))$ \triangleright Select B points with maximum entropy
- 23: **return** $Q_{entropy}$
- 24:
- 25: **GradientSampling** (extracted model: f^E , input data points: X_{in} , budget: $\gamma_1 B$)
- 26: $E \leftarrow H(f^E(X_{in}))$
- 27: $G \leftarrow \{\nabla_x E \mid x \in X_{in}\}$
- 28: $C_{in} \leftarrow k$ centres of G computed using K-means
- 29: $Q_{grad} \leftarrow \arg \min_{X \subset X_{in}, |X|=\gamma_1 B} \sum_{x_i \in X} \sum_{x_j \in C_{in}} \|\nabla_{x_i} E - \nabla_{x_j} E\|_2^2$ \triangleright Select $\gamma_1 B$ points from X_{in} whose $\frac{\partial E}{\partial x}$ are closest to that of C_{in}
- 30: **return** Q_{grad}
- 31:
- 32: **LossSampling** (extracted model: f^E , input data points: X_{in} , previous queries: Q_{train} , previous predictions: Y_{train} , budget: $\gamma_1 \gamma_2 B$)
- 33: $L \leftarrow l(Y_{train}, f^E(Q_{train}))$ \triangleright Compute the mismatch vector
- 34: $Q_{mis} \leftarrow$ ARGMAXSORT(L, k) \triangleright Select top- k mismatching points
- 35: $Q_{loss} \leftarrow \arg \min_{X \subset X_{in}, |X|=\gamma_1 \gamma_2 B} \sum_{x_i \in X} \sum_{x_j \in Q_{mis}} \|x_i - x_j\|_2^2$ \triangleright Select $\gamma_1 \gamma_2 B$ points closest to Q_{mis}
- 36: **return** Q_{loss}

577 **B Theoretical analysis: Proofs of section 3**

578 In this section, we elaborate the proofs for the Theorems 3.2 and 3.4.²

579 **Theorem 3.2** (Upper Bounding Distributional Closeness). If we choose KL-divergence as the
580 divergence function D , we can show that

$$D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{Deq}}^E}^E(Q), Q)) \leq \min_{\omega} \mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] - H(f_{\omega}^E(Q)).$$

581 *Proof.* Let us consider a query generating distribution \mathcal{D}^Q on \mathbb{R}^d . A target model $f_{\theta^*}^T : \mathbb{R}^d \rightarrow \mathcal{Y}$
582 induces a joint distribution over the query and the output (or label) space, denoted by $\Pr(f_{\theta^*}^T, Q)$.
583 Similarly, the extracted model $f_{\theta^*}^T : \mathbb{R}^d \rightarrow \mathcal{Y}$ also induces a joint distribution over the query and the
584 output (or label) space, denoted by $\Pr(f_{\omega}^E, Q)$.

$$\begin{aligned} & D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega}^E(Q), Q)) \\ &= \int_{Q \in \mathbb{R}^d} d \Pr(f_{\theta^*}^T(Q), Q) \log \frac{\Pr(f_{\theta^*}^T(Q), Q)}{\Pr(f_{\omega}^E(Q), Q)} \\ &= \int_{Q \in \mathbb{R}^d} \Pr(f_{\theta^*}^T(Q) | Q = q) \Pr(Q = q) \log \frac{\Pr(f_{\theta^*}^T(Q) | Q = q)}{\Pr(f_{\omega}^E(Q) | Q = q)} dq \\ &= \int_{Q \in \mathbb{R}^d} \Pr(f_{\theta^*}^T(Q) | Q = q) \Pr(Q = q) \log \Pr(f_{\theta^*}^T(Q) | Q = q) dq \\ &\quad - \int_{Q \in \mathbb{R}^d} \Pr(f_{\theta^*}^T(Q) | Q = q) \Pr(Q = q) \log \Pr(f_{\omega}^E(Q) | Q = q) dq \\ &= \int_{Q \in \mathbb{R}^d} \Pr(f_{\theta^*}^T(Q) | Q = q) \Pr(Q = q) \log \Pr(f_{\theta^*}^T(Q) | Q = q) dq + \mathbb{E}_{q \sim \mathcal{D}^Q} [l(f_{\theta^*}^T(q), f_{\omega}^E(q))] \\ &\leq -H(f_{\theta^*}^T(Q)) dq + \mathbb{E}_{q \sim \mathcal{D}^Q} [l(f_{\theta^*}^T(q), f_{\omega}^E(q))] \\ &\leq -H(f_{\omega}^E(Q)) dq + \mathbb{E}_{q \sim \mathcal{D}^Q} [l(f_{\theta^*}^T(q), f_{\omega}^E(q))] \end{aligned} \tag{8}$$

585 The last inequality holds true as the extracted model f_{ω}^E is trained using the outputs of the target
586 model $f_{\theta^*}^T$. Thus, by data-processing inequality, its output distribution possesses less information
587 than that of the target model. Specifically, we know that if $Y = f(X)$, $H(Y) \leq H(X)$.

588 Now, by taking \min_{ω} on both sides, we obtain

$$D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{Deq}}^E}^E(Q), Q)) \leq \min_{\omega} \mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] - H(f_{\omega}^E(Q)).$$

589 Here, $\omega_{\text{Deq}}^* \triangleq \arg \min_{\omega} D_{\text{KL}}(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega}^E(Q), Q))$. The equality exists if minima of
590 LHS and RHS coincide. \square

591 **Theorem 3.4** (Lower Bounding Information Leakage). The information leaked by any Max-
592 Information attack (Equation 3) is lower bounded as follows:

$$I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{MaxInf}}^E}^E(Q), Q)) \geq \max_{\omega} -\mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] + H(f_{\omega}^E(Q)).$$

593 *Proof.* Let us consider the same terminology as the previous proof. Then,

$$\begin{aligned} & I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega}^E(Q), Q)) \\ &= H(f_{\theta^*}^T(Q), Q) + H(f_{\omega}^E(Q), Q) - H(f_{\theta^*}^T(Q), f_{\omega}^E(Q), Q) \\ &= H(f_{\theta^*}^T(Q), Q) + H(f_{\omega}^E(Q), Q) - H(f_{\omega}^E(Q), Q | f_{\theta^*}^T(Q)) + H(f_{\theta^*}^T(Q)) \\ &\geq H(f_{\omega}^E(Q), Q) - H(f_{\omega}^E(Q), Q | f_{\theta^*}^T(Q)) \end{aligned} \tag{9}$$

$$\geq H(f_{\omega}^E(Q)) - H(f_{\omega}^E(Q), Q | f_{\theta^*}^T(Q)) \tag{10}$$

$$\geq H(f_{\omega}^E(Q)) - \mathbb{E}_Q[l(f_{\omega}^E(Q), f_{\theta^*}^T(Q))] \tag{11}$$

²Throughout the proofs, we slightly abuse the notation to write $l(\Pr(X), \Pr(Y))$ as $l(X, Y)$ for avoiding cumbersome equations.

594 The inequality of Equation 9 is due to the fact that entropy is always non-negative. Equation 10 holds
 595 true as $H(X, Y) \geq \max\{H(X), H(Y)\}$ for two random variables X and Y . The last inequality is
 596 due to the fact that conditional entropy of two random variables X and Y , i.e. $H(X|Y)$, is smaller
 597 than or equal to their cross entropy, i.e. $l(X, Y)$ (Lemma B.1).

598 Now, by taking \max_{ω} on both sides, we obtain

$$I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{MaxInf}}^E}^E(Q), Q)) \leq \max_{\omega} -\mathbb{E}_Q[l(f_{\theta^*}^T(Q), f_{\omega}^E(Q))] + H(f_{\omega}^E(Q)).$$

599 Here, $\omega_{\text{MaxInf}}^* \triangleq \arg \max_{\omega} I(\Pr(f_{\theta^*}^T(Q), Q) \| \Pr(f_{\omega_{\text{MaxInf}}^E}^E(Q), Q))$. The equality exists if maxima
 600 of LHS and RHS coincide. \square

601 **Lemma B.1** (Relating Cross Entropy and Conditional Entropy). *Given two random variables X and*
 602 *Y , conditional entropy*

$$H(X|Y) \leq l(X, Y). \quad (12)$$

603 *Proof.* Here, $H(X|Y) \triangleq -\int \Pr(x, y) \log \frac{\Pr(x, y)}{\Pr(y)} d\nu_1(X) d\nu_2(Y)$ and $l(X, Y) \triangleq$
 604 $l(\Pr(X), \Pr(Y)) = -\int \Pr(x) \ln \Pr(y) d\nu_1(X) d\nu_2(Y)$ denotes the cross-entropy, given ref-
 605 erence measures ν_1 and ν_2 .

$$\begin{aligned} l(X, Y) &= H(X) + D_{\text{KL}}(\Pr(X) \| \Pr(Y)) \\ &= H(X|Y) + I(X; Y) + D_{\text{KL}}(P_X \| P_Y) \\ &\geq H(X|Y) \end{aligned}$$

606 The last inequality holds as both mutual information I and KL-divergence D_{KL} are non-negative
 607 functions for any X and Y . \square

608 **C A review of active sampling strategies**

609 **K-Center sampling (KC).** K-Center sampling as an active learning algorithm that has been
 610 originally proposed to train CNNs sample-efficiently [SS18]. For a given data pool, K-Center
 611 sampling iteratively selects the k datapoints that minimise the core set loss (see Equation 3, [SS18])
 612 the most for a given model f_ω . In this method, the embeddings (from the model under training) of
 613 the data points are used as the representative vectors, and K-Center algorithm is applied on these
 614 representative vectors.

615 **Least Confidence sampling (LC).** Least confidence sampling method [Set09, LS06] iteratively
 616 selects the subset of k data points from a data pool, which are most uncertain at that particular instant.
 617 The uncertainty function ($u(\cdot|f_\omega) : \mathcal{X} \rightarrow [0, 1]$) is defined as

$$u(x|f_\omega) \triangleq 1 - \Pr(\hat{y}|x),$$

618 where \hat{y} is the predicted class by a model f_ω for input x .

619 **Margin Sampling (MS).** In margin sampling [JG19], a subset of k points is selected from a data
 620 pool, such that the subset demonstrates the minimum margin, where $\text{margin}(\cdot|f_\omega) : \mathcal{X} \rightarrow [0, 1]$ is
 621 defined as

$$\text{margin}(x|f_\omega) \triangleq \Pr(\hat{y}_1(x)|x, f_\omega) - \Pr(\hat{y}_2(x)|x, f_\omega),$$

622 where f_ω is the model, and $\hat{y}_1(x)$ and $\hat{y}_2(x)$ are respectively the highest and the second highest
 623 scoring classes returned by f_ω .

624 **Entropy Sampling (ES).** Entropy sampling, also known as uncertainty sampling [LG94], iteratively
 625 selects a subset of k datapoints with the highest uncertainty from a data pool. The uncertainty is
 626 defined by the entropy function of the prediction vector, and is computed using all the probabilities
 627 returned by the model f_ω for a datapoint x . For a given point x and a model f_ω , entropy is defined as

$$\text{entropy}(x|f_\omega) \triangleq - \sum_{a=1}^{|\mathcal{Y}|} p_a \log(p_a),$$

628 where $p_a = \Pr[f_\omega(x) = a]$ for any output class $a \in \{1, \dots, |\mathcal{Y}|\}$. [LG94] mention that while using
 629 this strategy, “*the initial classifier plays an important role, since without it there may be a long period*
 630 *of random sampling before examples of a low frequency class are stumbled upon*”. This is similar to
 631 our experimental observation that ES often demonstrate high variance in its outcomes.

632 **Random Sampling (RS).** In random sampling, a subset of k datapoints are selected from a data
 633 pool uniformly at random.

634 In our experiments, for query selection at time t , the extracted model at time $t - 1$, i.e. f_{t-1}^E , is used
 635 as f_ω , and data pool at time t is the corresponding query dataset, except the datapoints that has been
 636 selected before step t . Hereafter, we deploy a modified version of the framework develop by [Hua21]
 637 to run our experiments using the active learning algorithms mentioned above.

638 D Extended experimental analysis

639 In this section, we step-wise elaborate further experimental setups and results that we skipped for
640 the brevity of space in the main draft. Specifically, we conduct our experiments in six experimental
641 setups. Each experimental setup corresponds to a triplet (target model architecture trained on a private
642 dataset, extracted model architecture, query dataset). Here, we list these six experimental setups in
643 detail

- 644 1. A Logistic Regression (LR) model trained on MNIST, a LR model for extraction, EMNIST dataset
645 for querying
- 646 2. A Logistic Regression (LR) model trained on MNIST, a LR model for extraction, CIFAR10
647 dataset for querying
- 648 3. A CNN model trained on MNIST, a CNN model for extraction, EMNIST dataset for querying
- 649 4. A ResNet model trained on CIFAR10, a CNN model for extraction, ImageNet dataset for querying
- 650 5. A ResNet model trained on CIFAR10, a ResNet18³ model for extraction, ImageNet dataset for
651 querying
- 652 6. A BERT model⁴ trained on BBCNews, a BERT model for extraction, AGNews dataset for querying

653 For each of the experimental setups, we evaluate five types of performance evaluations, which are
654 elaborated in Section D.1, D.2.1, D.2.2, D.3, and D.4. While each of the following sections contain
655 illustrations of the different performance metrics evaluating efficacy of the attack and corresponding
656 discussions, Table 2- 4 contain summary of all queries used, accuracy, and membership inference
657 statistics for all the experiments.

658 D.1 Test accuracy of extracted models

659 Test accuracy of the extracted model and its comparison with the test accuracy of the target model
660 on a subset of the private training dataset, which was used by neither of these models, is the most
661 common performance metric used to evaluate the goodness of the attack algorithm. The attacks
662 designed solely to optimise this performance metric are called the task accuracy model extraction
663 attacks [JCB⁺20].

664 With MARICH, we aim to extract models that have prediction distributions closest to that of the target
665 model. Our hypothesis is constructing such a prediction distribution lead to a model that also has
666 high accuracy on the private test dataset, since accuracy is a functional property of the prediction
667 distribution induced by a classifier. In order to validate this hypothesis, we compute test accuracies
668 of the target models, and models extracted by MARICH and other active sampling algorithms in
669 six experimental setups. We illustrate the evolution curves of accuracies over increasing number of
670 queries in Figure 5.

671 To compare MARICH with other active learning algorithms, we attack the same target models using
672 K-centre sampling, Least Confidence sampling, Margin Sampling, Entropy Sampling, and Random
673 Sampling algorithms (ref. Appendix C) using the same number of queries as used for MARICH in
674 each setup.

675 From Figure 5, we observe that **in most of the cases MARICH outperforms all other competing
676 algorithms.**

677 In this process, MARICH uses $\sim 500 - 8000$ queries, which is a small fraction of the corresponding
678 query datasets. This also indicates towards the query-efficiency of MARICH.

679 **Extraction of a ResNet trained on CIFAR10 with a ResNet18.** Along with the five experimental
680 setups mentioned in the paper, we trained a ResNet with CIFAR10 dataset (\mathbf{D}^T here), that shows a
681 test accuracy of 91.82% on a disjoint test set. We use ImageNet as \mathbf{D}^Q here, to extract a ResNet18
682 model from the target model. We have restrained from discussing this setup in the main paper due to
683 brevity of space.

³We begin with a pre-trained ResNet18 model from <https://pytorch.org/vision/main/models/generated/torchvision.models.resnet18.html>

⁴We use the pre-trained BERT model from <https://huggingface.co/bert-base-cased>

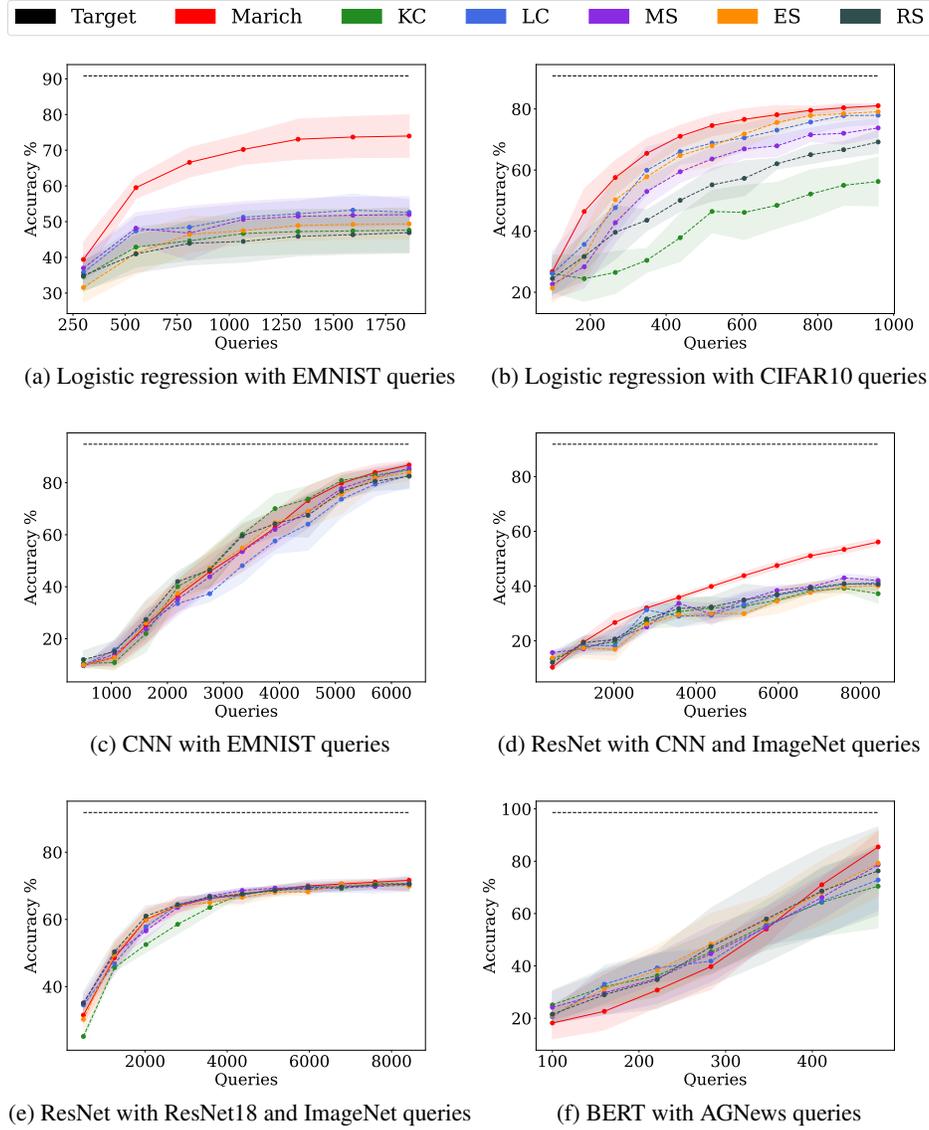


Figure 5: Comparison of test accuracies achieved by models extracted by different active sampling algorithms.

684 MARICH extracts a ResNet18 using 8429 queries, and the extracted ResNet18 shows test accuracy
 685 of $71.65 \pm 0.88\%$. On the other hand, models extracted using Best of Competitors (BoC) using
 686 ImageNet queries shows accuracy of $70.67 \pm 0.12\%$.

687 **D.2 Fidelity of the prediction distributions of The extracted models**

688 Driven by the distributional equivalence extraction principle, the central goal of MARICH is to
 689 construct extracted models whose prediction distributions are closest to the prediction distributions
 690 of corresponding target models. From this perspective, in this section, we study the fidelity of
 691 the prediction distributions of models extracted by MARICH and other active sampling algorithms,
 692 namely K-centre sampling, Least Confidence sampling, Margin Sampling, Entropy Sampling, and
 693 Random Sampling.

694 **D.2.1 KL-divergence between prediction distributions**

695 First, as the metric of distributional equivalence, we evaluate the KL-divergence between the pre-
 696 diction distributions of the models extracted by MARICH and other active sampling algorithms. In
 697 Figure 6, we report the box-plot (mean, median \pm 25 percentiles) of KL-divergences (in log-scale)
 698 calculated from 5 runs for each of 10 models extracted by each of the algorithms.

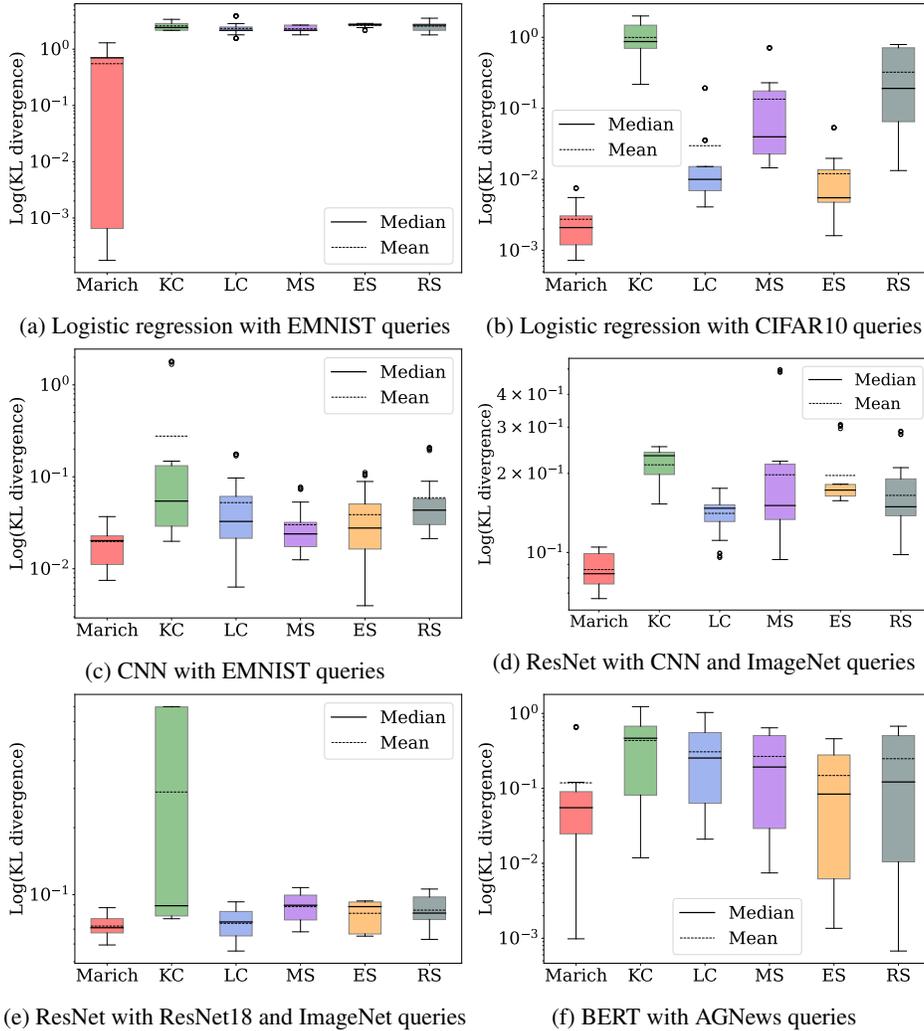


Figure 6: Comparison of KL-divergences (in log-scale) between the target prediction distributions and prediction distributions of the models extracted by different active learning algorithms.

699 **Results.** Figure 6 shows that the KL-divergence achieved by the prediction distributions of models
 700 extracted using MARICH are at least $\sim 2 - 10$ times less than that of the other competing algorithms.
 701 This validates our claim that MARICH yields *distributionally closer extracted model f^E from the*
 702 *target model f^T than existing active sampling algorithms.*

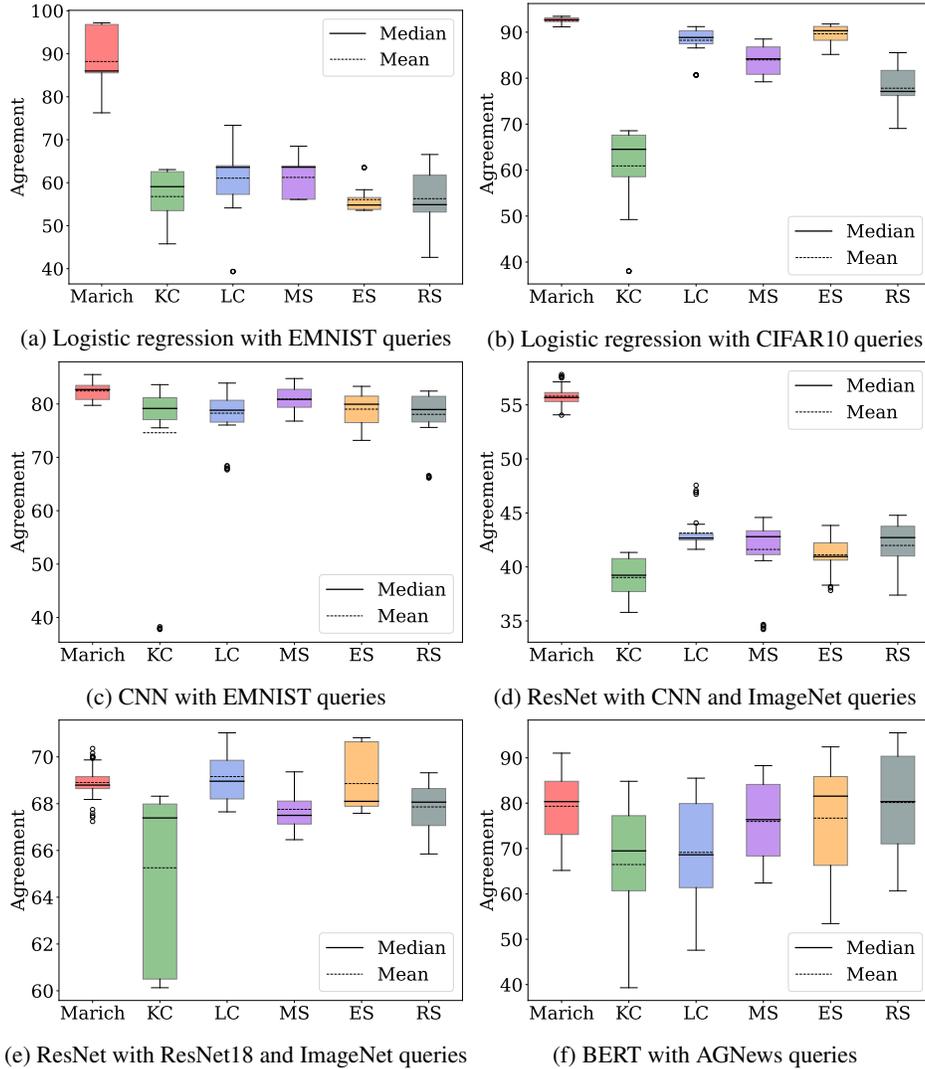


Figure 7: Comparison of agreement in predictions (in %) between the target model and the models extracted by different active learning algorithms.

704 In Figure 7, we illustrate the agreement in predictions of f^E with f^T on test datasets using different
 705 active learning algorithms. Prediction agreement functions as another metric of fidelity of prediction
 706 distributions constructed by extracted models in comparison with those of the target models.

707 Similar to Figure 6, we report the box-plot (mean, median \pm 25 percentiles) of prediction agreements
 708 (in %) calculated from 5 runs for each of 10 models extracted by each of the algorithms.

709 **Results.** We observe that the prediction distributions extracted by MARICH achieve almost same to
 710 $\sim 30\%$ higher prediction agreement in comparison with the competing algorithms. Thus, we infer
 711 that in this particular case MARICH achieves better fidelity than the other active sampling algorithms,
 712 in some instances, while it is similar to the BoC in some instances.

713 **D.3 Fidelity of parameters of the extracted models**

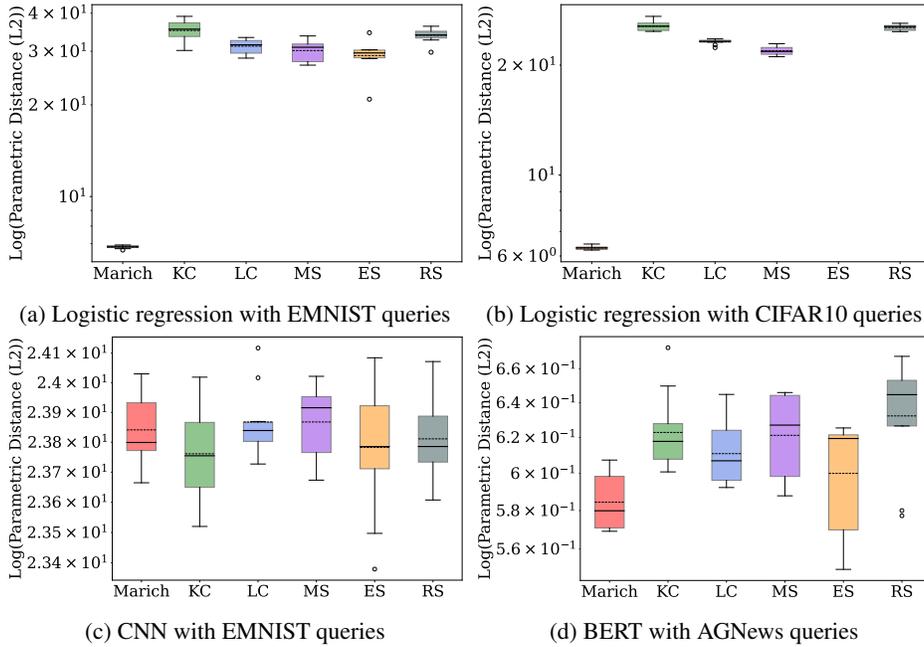


Figure 8: Comparison of parametric fidelity for MARICH and different active learning algorithms.

714 Generating extracted models with low parametric fidelity is not a main goal or basis of the design
 715 principle of MARICH. Since parametric fidelity is a popularly studied metric to evaluate goodness
 716 of model extraction, in Figure 8, we depict the parametric fidelity of models extracted by different
 717 active learning algorithms.

718 Let w_E be the parameters of the extracted model and w_T be the parameters of the target model. We
 719 define parametric fidelity as $F_w \triangleq \log \|w_E - w_T\|_2$. Since the parametric fidelity is only computable
 720 when the target and extracted models share the same architecture, we report the four instances here
 721 where MARICH is deployed with the same architecture as that of the target model. For logistic
 722 regression, we compare all the weights of the target and the extracted models. For BERT and CNN,
 723 we compare between the weights in the last layers of these models.

724 **Results.** For LR, we observe that the LR models extracted by MARICH have 20 – 30 times lower
 725 parametric fidelity than the extracted LR models of the competing algorithms. For BERT, the BERT
 726 extracted by MARICH achieves 0.4 times lower parametric fidelity than the Best of Competitors
 727 (BoC). As an exception, for CNN, the model extracted by K-center sampling achieves 0.996 times
 728 less parametric fidelity than that of MARICH.

729 Thus, we conclude that MARICH as a by-product of its distributionally equivalent extraction principle
 730 also extracts model with high parametric fidelity, which is often better than the competing active
 731 sampling algorithms.

732 **D.4 Membership inference with the extracted models**

733 A main goal of MARICH is to conduct a Max-Information attack on the target model, i.e. to extract
 734 an informative replica of its predictive distribution that retains the most information about the private
 735 training dataset. Due to lack of any direct measure of informativeness of an extracted model with
 736 respect to a target model, we run Membership Inference (MI) attacks using the models extracted
 737 by MARICH, and other competing active sampling algorithms. High accuracy and agreement in MI
 738 attacks conducted on extracted models of MARICH and the target models implicitly validate our
 739 claim that MARICH is able to conduct a Max-Information attack.

740 **Observation 1.** From Figure 9, we see that in most cases the probability densities of the membership
 741 inference are closer to the target model when the model is extracted using MARICH, than using all
 742 other active sampling algorithms (BoC, Best of Competitors).

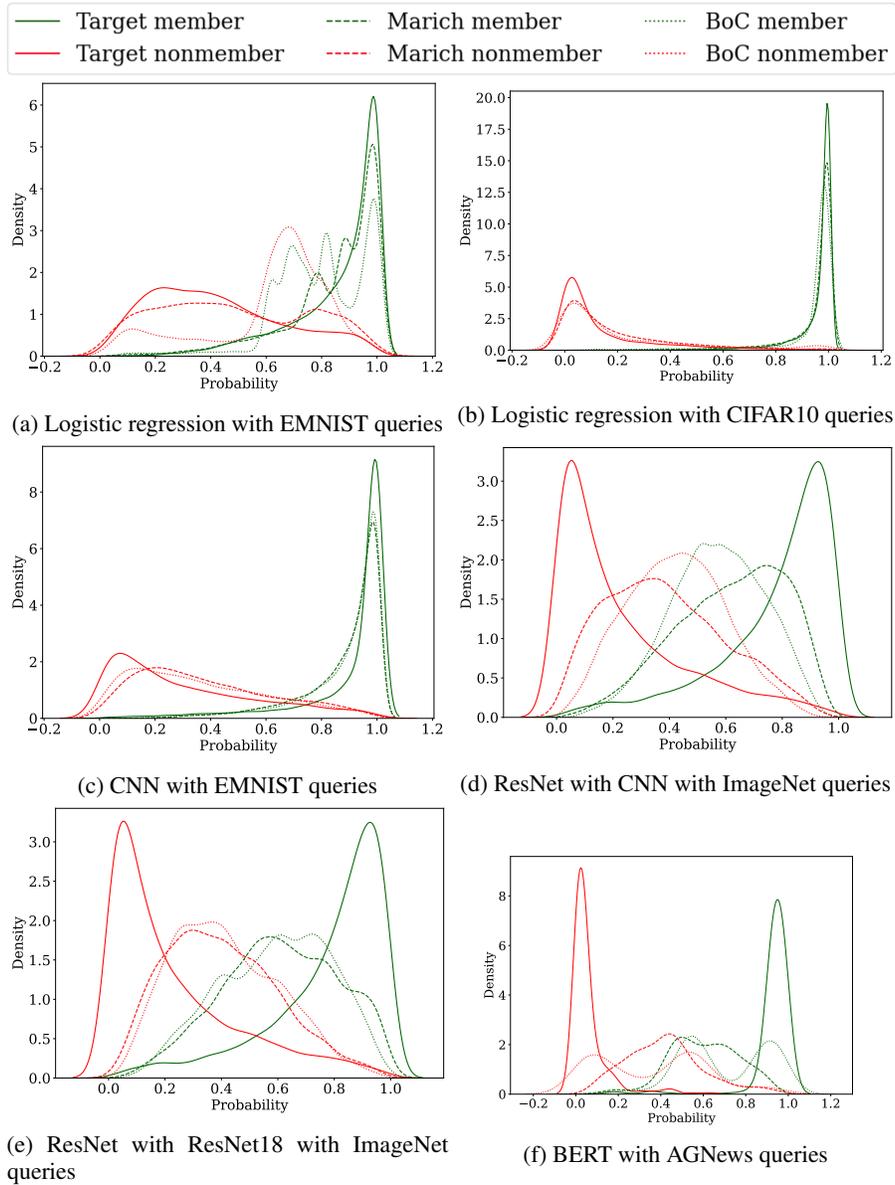
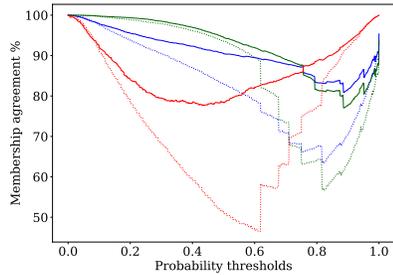
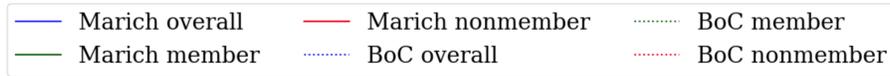
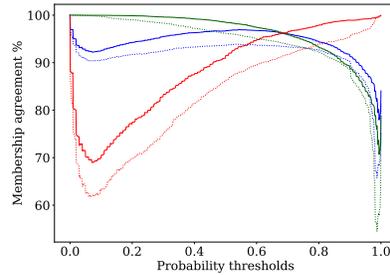


Figure 9: Comparison among membership vs. non-membership probability densities for membership attacks against models extracted by MARICH, the best of competitors (BoC) and the target model. Each figure represents the model class and query dataset. Memberships and non-memberships inferred from the model extracted by MARICH are significantly closer to the target model.

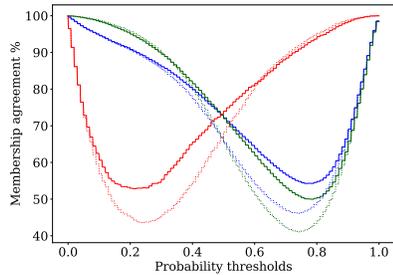
743 **Observation 2.** In Figure 10, we present the agreements from the member points, nonmember points
 744 and overall agreement curves for varying membership thresholds, along with the AUCs of the overall
 745 membership agreements. We see that in most cases, the agreement curves for the models extracted
 746 using MARICH are above those for the models extracted using random sampling, thus AUCs are
 747 higher for the models extracted using MARICH.



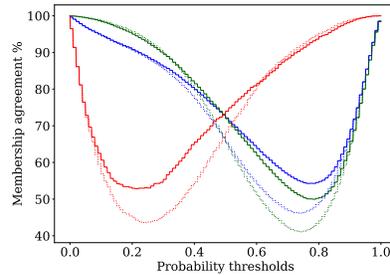
(a) Logistic regression with EMNIST queries



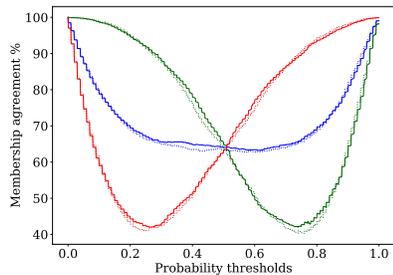
(b) Logistic regression with CIFAR10 queries



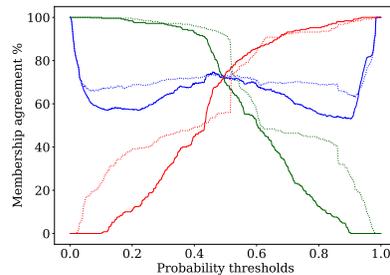
(c) CNN with EMNIST queries



(d) ResNet with CNN and ImageNet queries



(e) ResNet with ResNet18 & ImageNet queries



(f) BERT with AGNews queries

Figure 10: Comparison of membership, nonmembership and overall agreements of membership attacks against models extracted by MARICH and the best of competitors with the target model trained with MNIST. Each figure represents the model class and query dataset. Membership agreement of the models extracted by MARICH are higher.

748 **Observation 3.** In Table 2, 3, and 4, we summarise the MI accuracy on the private training dataset,
 749 Nonmembership inference accuracy on the private training dataset, Agreement in MI w.r.t. the MI on
 750 the target model, and AUC of Agreement in MI with that of the target model for Logistic Regression
 751 (LR), CNN, ResNet, and BERT target models. We observe that, while compared with other active
 752 sampling algorithms, out of 6 combinations of (target model, extracted model, query dataset) under
 753 study, the models extracted by MARICH achieve the highest accuracy in MI and agreement in MI
 754 w.r.t. the target model, in most of the instances.

755 **Results.** *These observations support our claim that model extraction using MARICH gives models*
 756 *are accurate and informative replica of the target model.*

Table 2: Model extraction and membership inference statistics for extracting a Logistic Regression model and a CNN

Member Dataset	Target Model	Attack Dataset	Algorithm used	Non-member Dataset	#Queries	Membership acc	Nonmembership acc	Overall membership acc	Overall membership agreement	Membership agreement AUC	Accuracy
MNIST	LR	-	-	EMNIST	-	95.37%	65.84%	87.99%	-	-	90.82%
MNIST	LR	EMNIST	Marich	EMNIST	1863	93.43%	57.58%	84.47%	90.34%	90.89%	73.98+5.96%
MNIST	LR	EMNIST	ES	EMNIST	1863	97.22%	15.88%	75.88%	84.39%	82.32%	49.37+4.24%
MNIST	LR	EMNIST	KC	EMNIST	1863	97.80%	14.56%	76.99%	87.84%	82.30%	47.62+6.25%
MNIST	LR	EMNIST	MS	EMNIST	1863	97.63%	18.42%	77.83%	86.49%	82.74%	52.60+3.55%
MNIST	LR	EMNIST	LC	EMNIST	1863	95.64%	25.10%	78.00%	80.11%	83.07%	51.968+5.05%
MNIST	LR	EMNIST	RS	EMNIST	1863	98.23%	11.92%	76.65%	88.08%	82.26%	46.91+5.69%
MNIST	LR	-	-	CIFAR10	-	99.11%	94.76%	98.02%	-	-	90.82%
MNIST	LR	CIFAR10	Marich	CIFAR10	959	97.72%	92.14%	96.32%	96.89%	94.32%	81.06+0.56%
MNIST	LR	CIFAR10	ES	CIFAR10	959	96.47%	86.16%	93.89%	96.53%	91.29%	79.09+1.82%
MNIST	LR	CIFAR10	KC	CIFAR10	959	95.87%	81.64%	92.31%	92.32%	89.64%	56.26+7.97%
MNIST	LR	CIFAR10	MS	CIFAR10	959	96.79%	82.82%	93.30%	93.16%	90.92%	77.93+2.94%
MNIST	LR	CIFAR10	LC	CIFAR10	959	96.19%	86.26%	93.71%	93.67%	91.53%	73.78+2.95%
MNIST	LR	CIFAR10	RS	CIFAR10	959	95.84%	81.20%	92.18%	92.17%	90.11%	69.18+3.72%
MNIST	CNN	-	-	EMNIST	-	93.71%	78.76%	89.97%	-	-	94.83%
MNIST	CNN	EMNIST	Marich	EMNIST	1863	94.55%	78.86%	90.62%	87.27%	86.72%	86.83+1.62%
MNIST	CNN	EMNIST	ES	EMNIST	1863	95.44%	73.66%	89.99%	87.11%	86.38%	83.90+3.10%
MNIST	CNN	EMNIST	KC	EMNIST	1863	95.35%	74.48%	90.13%	87.49%	85.94%	84.94+2.44%
MNIST	CNN	EMNIST	MS	EMNIST	1863	94.59%	79.04%	90.70%	86.74%	86.08%	82.72+4.47%
MNIST	CNN	EMNIST	LC	EMNIST	1863	94.39%	76.30%	89.86%	86.75%	86.10%	85.53+2.45%
MNIST	CNN	EMNIST	RS	EMNIST	1863	94.17%	80.42%	90.73%	87.53%	86.97%	82.52+4.87%

Table 3: Model extraction and membership inference statistics for extracting a ResNet

Member Dataset	Target Model	Attack Model	Attack Dataset	Algorithm used	Non-member Dataset	#Queries	Membership acc	Nonmembership acc	Overall membership acc	Overall membership agreement	Membership agreement AUC	Accuracy
CIFAR 10	ResNet	-	-	-	ImageNet	-	97.70%	56.80%	93.61%	-	-	91.82%
CIFAR 10	ResNet	CNN	ImageNet	Marich	ImageNet	8429	70.84%	68.44%	69.64%	64.14%	71.78%	56.11+1.34%
CIFAR 10	ResNet	CNN	ImageNet	ES	ImageNet	8429	65.00%	71.62%	68.31%	62.94%	71.36%	40.11+1.94%
CIFAR 10	ResNet	CNN	ImageNet	KC	ImageNet	8429	65.80%	71.98%	68.89%	63.68%	71.53%	37.21+3.41%
CIFAR 10	ResNet	CNN	ImageNet	MS	ImageNet	8429	68.48%	71.96%	70.22%	64.48%	71.97%	41.27+0.85%
CIFAR 10	ResNet	CNN	ImageNet	LC	ImageNet	8429	70.60%	71.68%	71.14%	65.47%	72.29%	42.05+1.38%
CIFAR 10	ResNet	CNN	ImageNet	RS	ImageNet	8429	67.56%	68.68%	68.12%	63.41%	71.39%	40.66+2.38%
CIFAR 10	ResNet	ResNet18	ImageNet	Marich	ImageNet	8429	99.27%	10.54%	90.40%	93.84%	76.51%	71.65+0.88%
CIFAR 10	ResNet	ResNet18	ImageNet	ES	ImageNet	8429	100.00%	0.02%	90.00%	97.29%	72.33%	69.92+1.51%
CIFAR 10	ResNet	ResNet18	ImageNet	KC	ImageNet	8429	99.70%	2.06%	90.02%	94.06%	72.40%	70.67+0.12%
CIFAR 10	ResNet	ResNet18	ImageNet	MS	ImageNet	8429	99.09%	0.28%	90.02%	96.28%	72.81%	70.66+2.19%
CIFAR 10	ResNet	ResNet18	ImageNet	LC	ImageNet	8429	99.99%	0.04%	90.00%	96.77%	71.94%	70.26+1.44%
CIFAR 10	ResNet	ResNet18	ImageNet	RS	ImageNet	8429	99.94%	1.40%	93.08%	93.41%	72.94%	70.49+2.02%

Table 4: Model extraction and membership inference statistics for extracting a BERT model

Member Dataset	Target Model	Attack Model	Attack Dataset	Algorithm used	Non-member Dataset	#Queries	Membership acc	Nonmembership acc	Overall membership acc	Overall membership agreement	Membership agreement AUC	Accuracy
BBC News	BERT	-	-	-	AG News	-	-	-	-	-	-	98.62%
BBC News	BERT	BERT	AG News	Marich	AG News	474	83.60%	66.80%	75.20%	74.80%	65.14%	85.45+5.96%
BBC News	BERT	BERT	AG News	ES	AG News	474	82.80%	48.40%	65.60%	66.00%	59.36%	79.25+12.67%
BBC News	BERT	BERT	AG News	KC	AG News	474	68.40%	55.60%	62.00%	63.39%	61.10%	70.45+15.94%
BBC News	BERT	BERT	AG News	MS	AG News	474	77.20%	72.80%	75.00%	62.40%	71.07%	72.82+11.59%
BBC News	BERT	BERT	AG News	LC	AG News	474	80.00%	43.20%	61.60%	74.80%	57.93%	78.65+7.22%
BBC News	BERT	BERT	AG News	RS	AG News	474	78.00%	47.60%	62.80%	63.80%	58.43%	76.31+16.78%

757 **E Significance and comparison of three sampling strategies**

758 Given the bi-level optimization problem, we came up with MARICH in which three sampling methods
 759 are used in the order: (i) ENTROPYSAMPLING, (ii) ENTROPYGRADIENTSAMPLING, and (iii)
 760 LOSSAMPLING.

761 These three sampling techniques contribute to different goals:

- 762 • ENTROPYSAMPLING selects points about which the classifier at a particular time step is most
 763 confused
- 764 • ENTROPYGRADIENTSAMPLING uses gradients of entropy of outputs of the extracted model w.r.t.
 765 the inputs as embeddings and selects points behaving most diversely at every time step.
- 766 • LOSSAMPLING selects points which produce highest loss when loss is calculated between target
 767 model’s output and extracted model’s output.

768 One can argue that the order is immaterial for the optimization problem. But looking at the algorithm
 769 practically, we see that ENTROPYGRADIENTSAMPLING and LOSSAMPLING incur much higher
 770 time complexity than ENTROPYSAMPLING. Thus, using ENTROPYSAMPLING on the entire query set
 771 is more efficient than the others. This makes us put ENTROPYSAMPLING as the first query selection
 772 strategy.

773 As per the optimization problem in Equation (7), we are supposed to find points that show highest
 774 mismatch between the target and the extracted models after choosing the query subset maximising
 775 the entropy. This leads us to the idea of LOSSAMPLING. But as only focusing on loss between
 776 models may choose points from one particular region only, and thus, decreasing the diversity of the
 777 queries. We use ENTROPYGRADIENTSAMPLING before LOSSAMPLING. This ensures selection of
 778 diverse points with high performance mismatch.

779 In Figure 11, we experimentally see the time complexities of the three components used. These are
 780 calculated by applying the sampling algorithms on a logistic regression model, on mentioned slices
 781 of MNIST dataset.

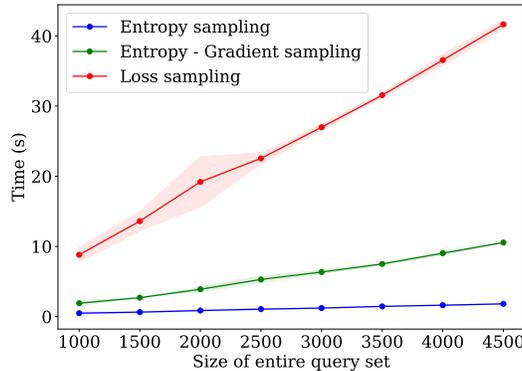


Figure 11: Runtime comparison of three sampling strategies to select queries from 4500 datapoints.

Table 5: Time complexity of different sampling Strategies

Sampling Algorithm	Query space size	#Selected queries	Time (s)
Entropy Sampling	4500	100	1.82 ± 0.04
Entropy-Gradient Sampling	4500	100	10.56 ± 0.07
Loss Sampling	4500	100	41.64 ± 0.69

782 **F Performance against differentially private target models**

783 In this section, we aim to verify performance of MARICH against privacy-preserving mechanisms.
 784 Specifically, we apply a (ϵ, δ) -Differential Privacy (DP) inducing mechanism [DMNS06, DBB21]
 785 on the target model to protect the private training dataset. There are three types of methods to
 786 ensure DP: output perturbation [DMNS06], objective perturbation [CMS11, DBB18], and gradient
 787 perturbation [ACG⁺16]. Since output perturbation and gradient perturbation methods scale well for
 788 nonlinear deep networks, we focus on them as the defense mechanism against MARICH’s queries.

789 **Gradient perturbation-based defenses.** DP-SGD [ACG⁺16] is used to train the target model on
 790 the member dataset. This mechanism adds noise to the gradients and clip them while training the
 791 target model. We use the default implementation of Opacus [YSS⁺21] to conduct the training in
 792 PyTorch.

793 Following that, we attack the (ϵ, δ) -DP target models using MARICH and compute the corresponding
 794 accuracy of the extracted models. In Figure 12, we show the effect of different privacy levels ϵ on
 795 the achieved accuracy of the extracted Logistic Regression model trained with MNIST dataset and
 796 queried with EMNIST dataset. Specifically, we assign $\delta = 10^{-5}$ and vary ϵ in $\{0.2, 0.5, 1, 2, \infty\}$.
 797 Here, $\epsilon = \infty$ corresponds to the model extracted from the non-private target model.

798 We observe that the accuracy of the models extracted from private target models are approximately
 799 2.3 – 7.4% lower than the model extracted from the non-private target model. This shows that
 800 performance of MARICH decreases against DP defenses but not significantly.

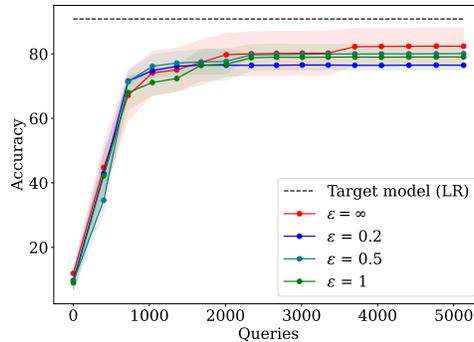


Figure 12: Performance of models extracted by MARICH against (ϵ, δ) -differentially private target models trained using DP-SGD. We consider different privacy levels ϵ and $\delta = 10^{-5}$. Accuracy of the extracted models decrease with increase in privacy (decrease in ϵ).

801 **Output perturbation-based defenses.** Perturbing output of an algorithm against certain queries
 802 with calibrated noise, in brief output perturbation, is one of the basic and oldest form of privacy-
 803 preserving mechanism [DMNS06]. Here, we specifically deploy the Laplace mechanism, where a
 804 calibrated Laplace noise is added to the output of the target model generated against some queries.
 805 The noise is sampled from a Laplace distribution $\text{Lap}(0, \frac{\Delta}{\epsilon})$, where Δ is sensitivity of the output and
 806 ϵ is the privacy level. This mechanism ensures ϵ -DP.

807 We compose a Laplace mechanism to the target model while responding to MARICH’s query and
 808 evaluate the change in accuracy of the extracted model as the impact of the defense mechanism. We
 809 use a logistic regression model trained on MNIST as the target model. We query it using EMNIST
 810 and CIFAR10 datasets respectively. We vary ϵ in $\{0.25, 2, 8, \infty\}$. For each ϵ and query dataset, we
 811 report the mean and standard deviation of accuracy of the extracted models on a test dataset. Each
 812 experiment is run 10 times.

813 We observe that decrease in ϵ , i.e. increase in privacy, causes decrease in accuracy of the extracted
 814 model. For EMNIST queries (Figure 13a), the degradation in accuracy is around 10% for $\epsilon = 2, 8$ but
 815 we observe a significant drop for $\epsilon = 0.25$. For CIFAR10 queries (Figure 13b), $\epsilon = 8$ has practically
 816 no impact on the performance of the extracted model. But for $\epsilon = 2$ and 0.25 , the accuracy of
 817 extracted models drop down very fast.

818 Thus, we conclude that output perturbation defends privacy of the target model against MARICH for
 819 smaller values of ϵ . But for larger values of ϵ , the privacy-preserving mechanism might not save the
 820 target model significantly against MARICH.

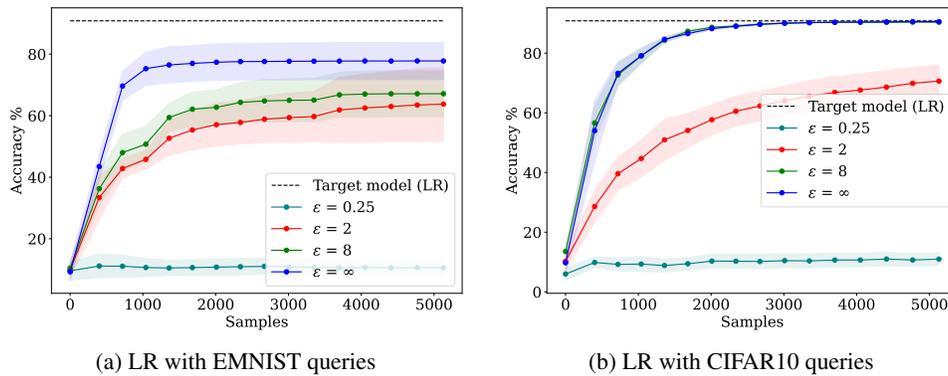


Figure 13: Performance of models extracted by MARICH against target models that perturbs the output of the queries to achieve ϵ -DP. We consider different privacy levels ϵ . Accuracy of the extracted models decrease with increase in privacy (decrease in ϵ).

821 **G Effects of model mismatch**

822 From Equation (7), we observe that functionality of MARICH is not constrained by selecting the
 823 same model class for both the target model f^T and the extracted model f^E . To elaborately study
 824 this aspect, in this section, we conduct experiments to show MARICH’s capability to handle model
 825 mismatch and impact of model mismatch on performance of the extracted models.

826 Specifically, we run experiments for three cases: (1) *extracting an LR model with LR and CNN*, (2)
 827 *extracting a CNN model with LR and CNN*, and (3) *extracting a ResNet model with CNN and ResNet18*.
 828 We train the target LR and the target CNN model on MNIST dataset. We further extract these two
 829 models using EMNIST as the query datasets. We train the target ResNet model on CIFAR10 dataset
 830 and extract it using ImageNet queries. The results on number of queries and achieved accuracies are
 831 summarised in Table 6 and Figure 14.

832 **Observation 1.** In all the three experiments, we use MARICH without any modification for both the
 833 cases when the model classes match and mismatch. This shows *universality and model-obliviousness*
 834 *of MARICH as a model extraction attack.*

835 **Observation 2.** From Figure 14, we observe that model mismatch influences performance of the
 836 model extracted by MARICH. When we extract the LR target model with LR and CNN, we observe
 837 that both the extracted models achieve almost same accuracy and the extracted CNN model achieves
 838 even a bit more accuracy than the extracted LR model. In contrast, when we extract the CNN target
 839 model with LR and CNN, we observe that the extracted LR model achieves lower accuracy than the
 840 extracted CNN model. Similar observations are found for extracting the ResNet with ResNet18 and
 841 CNN, respectively.

842 **Conclusions.** From these observations, we conclude that MARICH can function model-obliviously.
 843 We also observe that if we use a less complex model to extract a more complex model, the accuracy
 844 drops significantly. But if we extract a low complexity model with a higher complexity one, we
 845 obtain higher accuracy instead of model mismatch. This is intuitive as the low-complexity extracted
 846 model might have lower representation capacity to mimic the non-linear decision boundary of the
 847 high-complexity model but the opposite is not true. In future, it would be interesting to delve into the
 848 learning-theoretic origins of this phenomenon.

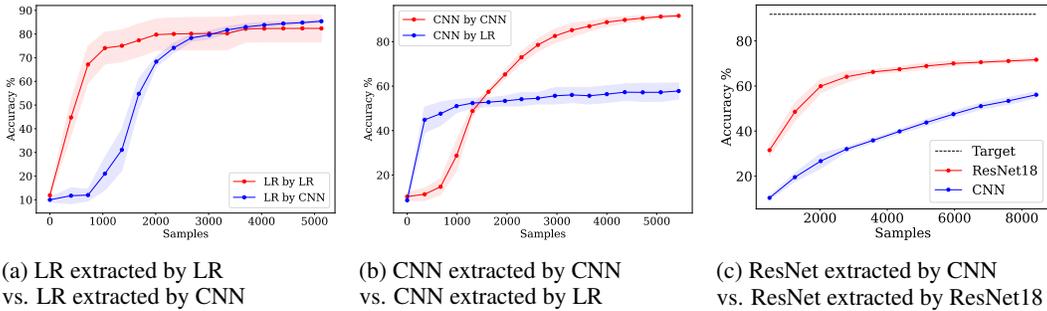


Figure 14: Effect of model mismatches on models extracted by MARICH.

Table 6: Effect of model mismatch on accuracy of The extracted models.

f^E	f^T	#samples	Accuracy
LR	LR	5130	$82.37 \pm 5.7\%$
LR	CNN	5130	$85.41 \pm 0.57\%$
CNN	LR	5440	$57.81 \pm 3.64\%$
CNN	CNN	5440	$91.63 \pm 0.42\%$
ResNet	CNN	8429	$56.11 \pm 1.35\%$
ResNet	ResNet18	8429	$71.65 \pm 0.88\%$

849 **H Choices of hyperparameters**

850 In this section, we list the choices of the hyperparameters of Algorithm 1 for different experiments
 851 and also explain how we select them.

852 Hyperparameters γ_1 and γ_2 are kept constant, i.e., 0.8, for all the experiments. These two parameters
 853 act as the budget shrinking factors.

854 Instead of changing these two, we change the number of points n_0 , which are randomly selected
 855 in the beginning, and the budget B for every step. We obtain the optimal hyperparameters for each
 856 experiment by performing a line search in the interval $[100, 500]$.

857 We further change the budget over the rounds. At time step t , the budget, $B_t = \alpha^t \times B_{t-1}$. The idea
 858 is to select more points as f^E goes on reaching the performance of f^T . Here, $\alpha > 1$ and needs to be
 859 tuned. We use $\alpha = 1.02$, which is obtained through a line search in $[1.01, 1.99]$.

860 For number of rounds T , we perform a line search in $[10, 20]$.

Table 7: Hyperparameters for different datasets and target models.

Member Dataset	Target Model	Attack Model	Attack Dataset	Budget	Initial points	γ_1	γ_2	Rounds	Epochs/Round	Learning Rate
MNIST	LR	LR	EMNIST	250	300	0.8	0.8	10	10	0.02
MNIST	LR	LR	CIFAR10	50	100	0.8	0.8	10	10	0.02
MNIST	CNN	CNN	EMNIST	550	500	0.8	0.8	10	10	0.015
MNIST	CNN	CNN	CIFAR10	750	500	0.8	0.8	10	10	0.03
CIFAR10	ResNet	CNN	ImageNet	750	500	0.8	0.8	10	8	0.2
CIFAR10	ResNet	ResNet18	ImageNet	750	500	0.8	0.8	10	8	0.02
BBC News	BERT	BERT	AG News	60	100	0.8	0.8	6	3	5×10^{-6}