

---

# Formalizing Distribution Inference Risks

---

Anshuman Suri and David Evans  
Department of Computer Science  
University of Virginia  
{anshuman, evans}@virginia.edu

## Abstract

Property inference attacks reveal statistical properties about a training set but are difficult to distinguish from the primary purposes of statistical machine learning, which is to produce models that capture statistical properties about a distribution. Motivated by Yeom et al.’s membership inference framework, we propose a formal and generic definition of property inference attacks. The proposed notion describes attacks that can distinguish between possible training distributions, extending beyond previous property inference attacks that infer the ratio of a particular type of data in the training data set. In this paper, we show how our definition captures previous property inference attacks as well as a new attack that reveals the average degree of nodes of a training graph and report on experiments giving insight into the potential risks of property inference attacks.

## 1 Introduction

Inference attacks seek to infer sensitive information about the training process of a revealed machine-learned model, most often about the training data. For example, in a membership inference attack [13], the adversary aims to infer whether a particular datum was part of the training data. In a property inference attack, an adversary aims to infer some statistical property of the training dataset, such as the proportion of women in a smile-detection dataset [1].

The research community lacks a good understanding of property inference attacks—for example, different notions are used in Gopinath et al. [6] and Zhang et al. [17]. Many formal privacy notions have been proposed, including numerous variations on differential privacy, but most notions focus on protecting specific data elements, not statistical properties of a dataset. One notable exception is the Pufferfish framework [11], which introduces notions that allow capturing aggregates of records via specifications of potential secrets and their relations. Although Pufferfish supports aggregate secrets, it is unclear how it can extend to distribution-level properties. A recent attempt to formalize property inference [3] consists of a framework that reduces property inference to Boolean functions of individual members, posing the ratio of members satisfying the given function in a dataset as the property. However, such a formulation limits the threat model since it cannot capture many other kinds of statistical properties of the training distribution that may be sensitive, like the degree distribution of a graph [8].

In this work<sup>1</sup>, we formalize property inference attacks based on the critical insight that the key difference between property inference attacks and other inference attacks is that the adversary’s goal in the former is to learn about the training distribution, not the specific training dataset. Dataset inference attacks, such as membership inference [13] and attribute inference [4], operate on the level of training records. They are directly connected to the definition of differential privacy which bounds the ability to distinguish neighboring datasets. By contrast, distribution inference attacks attempt to learn statistical properties of the underlying distribution from which the training dataset is sampled.

**Contributions.** We propose a simple and general experiment to formalize property inference attacks, inspired by Yeom et al. [16]’s membership inference definition (Section 2). Our definition is generic enough to capture any property of the underlying distribution, including, but not limited to, the attribute ratios that are the focus of previous property inference attacks. Motivated by this definition, we describe some experiments (Section 3) showing a variety of properties that fit our definitions and can be inferred by property inference attacks.

---

<sup>1</sup>An expanded version of this work is also available [14]

## 2 Formalizing Property Inference

Let  $\mathcal{D} = (\mathcal{X}, \mathcal{Y})$  be a public distribution between data  $\mathcal{X}$  and its corresponding labels  $\mathcal{Y}$ . We assume both the model trainer  $\mathcal{B}$  and adversary  $\mathcal{A}$  have access to  $\mathcal{D}$ . Both parties also have access to two functions that transform distributions:  $\mathcal{G}_0, \mathcal{G}_1$ . Inferring properties of the distribution can reveal sensitive information in many scenarios, which can be captured by suitable choices of  $\mathcal{G}_0$  and  $\mathcal{G}_1$ .

In our cryptographic game definition,  $\mathcal{B}$  picks one of these distribution transformers at random and samples a dataset  $D$  from the resulting distribution. Given access to a model  $M$  trained on  $D$ , the adversary aims to infer which of the two distribution mappers is used. The property inference experiment can be described as:

Trainer $\mathcal{B}$	Adversary $\mathcal{A}$
1 : $b \leftarrow \{0, 1\}$	
2 : $D \sim \mathcal{G}_b(\mathcal{D})$	
3 : $M \xleftarrow{\text{train}} D$	
4 : $\xrightarrow{M}$	
5 :	$\hat{b} = \mathcal{H}(M)$

If  $\mathcal{A}$  can successfully infer  $b$  via  $\hat{b}$ , then it can determine which of the two properties the training distribution satisfied. The advantage of the adversary  $\mathcal{A}$  using algorithm  $\mathcal{H}$  is:

$$\text{Adv}_{\mathcal{H}} = \left| \Pr[\hat{b} = b] - \Pr[\hat{b} \neq b] \right|.$$

This advantage is negligible when the adversary does no better than random guessing.

**Applying the Definition.** Seminal works on property inference [1, 5] involve a model trained either on the original dataset or a version modified to bias towards some chosen attribute. Our definition can be used to describe these attacks by setting  $\mathcal{G}_0$  to the identity function (so  $\mathcal{G}_0(\mathcal{D})$  is original distribution  $\mathcal{D}$ ) and  $\mathcal{G}_1$  to a filter that adjusts the distribution to have a specified ratio over the desired attribute. With respect to a binary property function  $f : \mathcal{X} \rightarrow \{0, 1\}$ ,  $\mathcal{D}$  can be characterized using a generative probability density function:

$$\rho_{\mathcal{D}}(\mathbf{x}) = \sum_{c \in \{0,1\}} p(c) \cdot p(\mathbf{x} | c), \quad (1)$$

where  $p(c)$  is a multinomial distribution representing the probabilities over the desired (binary) property function  $f$  and its possible values  $c$ , and  $p(\mathbf{x}|c)$  is the generative conditional probability density function. Then,  $\mathcal{G}_1(\mathcal{D})$  can be expressed using the following probability density function (with an adjusted prior  $\hat{p}$ ):

$$\rho_{\mathcal{G}_1(\mathcal{D})}(\mathbf{x}) = \sum_{c \in \{0,1\}} \hat{p}(c) \cdot p(\mathbf{x} | c), \quad \hat{p}(1) = \alpha \quad (2)$$

where  $\alpha$  is the probability of a randomly sampled point satisfying the property function  $f$ . Thus, a uniformly randomly sampled dataset from  $\mathcal{G}_1(\mathcal{D})$  would have an expected ratio  $\alpha$  of its members satisfy  $f$ . Additionally, we can modify  $\mathcal{G}_0$  with a similarly adjusted prior, enabling the adversary to distinguish between two chosen ratios [17].

## 3 Experiments

To better understand how well an intuitive notion of divergence in properties aligns with observed inference risk, we execute property inference attacks with increasing diverging properties on tabular, image, and graph datasets. Code for our experiments is available at <https://github.com/iamgroot42/distribution.inference>.

### 3.1 Datasets

We report on experiments using three datasets. For each dataset, we construct non-overlapping data splits between  $\mathcal{A}$  and  $\mathcal{B}$ . Both parties then modify their data according to the desired property to emulate a distribution: sampling datasets from these partitions to train and evaluate their models.

**Census** [2] consists of several categorical and numerical attributes like age, race, education level to predict whether a person makes over \$50K a year. We focus on the ratios of whites (race) and females (sex) as properties.

**RSNA Bone Age** [7] contains X-Ray images of hands, with the task being predicting the patients’ age in months. We convert the task to binary classification based on an age threshold and use a pre-trained DenseNet [10] model for feature extraction. We focus on the ratios of the females (available as metadata) as properties.

**ogbn-arxiv** [15] is a directed graph, representing citations between computer science arXiv papers. The task is to predict the subject area categories for unlabeled papers. We infer the mean node-degree property of the graph.

Since the original ratios may be unbalanced (Section 2), we fix  $\mathcal{G}_0$  such that the resulting distribution is balanced (0.5 ratio) for the chosen attribute for the Census and RSNA Bone Age datasets., setting  $\mathcal{G}_1$  based on varying  $\alpha$  values substituted in Equation 2. For the ogbn-arxiv dataset, we set  $\mathcal{G}_0$  such that the graph has a mean node-degree of 13. For  $\mathcal{G}_1$ , we modify the graph to have a mean-degree  $\alpha$ .

### 3.2 Attacks

We use two simple attacks and the state-of-the-art meta-classifier attack in our experiments.

**Loss Test.** A simple algorithm  $\mathcal{H}$  consists of using the loss function  $l$  used for training the model  $M$  to compute  $\hat{b}$ . For datasets  $D_0 \sim \mathcal{G}_0(\mathcal{D})$ ,  $D_1 \sim \mathcal{G}_1(\mathcal{D})$ ,  $\hat{b} = \mathbb{I}[l(M, D_0) > l(M, D_1)]$  where  $\mathbb{I}$  is the indicator function. Intuitively, a model would have lower test loss on data sampled from the training distribution, compared to another distribution. This method does not require the adversary to train models, but only to have access to suitable test distributions.

**Threshold Test.** The loss test assumption may not hold for some pairs of properties if one distribution is inherently easier to classify than the other (as we observe in our experiments in Section 3.3). The adversary can train models on datasets sampled from each distribution and observe their performance on  $D_0, D_1$  to derive a threshold  $\lambda$  to maximize accuracy distinguishing between the two ( $D_k$  yields higher accuracy).  $\hat{b}$  is then computed as :  $\hat{b} = \mathbb{I}[l(M, D_k) > \lambda]$ .

**Meta-Classifiers.** The state-of-the-art property inference attack is Ganju et al. [5]’s attack using permutation-invariant networks as meta-classifiers. The meta-classifiers take as input model parameters (weights, bias) and predict which of the distributions was used to train the model. Since the meta-classifier is itself a classifier that requires many (we use 1000 for each distribution in our experiments) models trained on the two distributions to train, this attack is only feasible for adversaries with access to both training distributions and considerable computational resources.

### 3.3 Results

We try the attacks on a range of diverging distributions to evaluate inference risks and understand how well an adversary could distinguish models trained using different distributions. Figures 1 and 2 summarize results on the three datasets.

**Census.** The loss test performs marginally better than random in most cases, while the threshold test often outperforms meta-classifiers (left two sub-figures of Figure 1). We observe a clear trend with meta-classifiers while varying  $\alpha$  for the sex attribute: performance improves as distributions diverge ( $\alpha$  further away from 0.5). However, this trend is not symmetric around  $\mathcal{G}_0(\mathcal{D})$ . For example, the adversary performs better on  $\alpha = 0.0$  than  $\alpha = 1.0$ , even though both

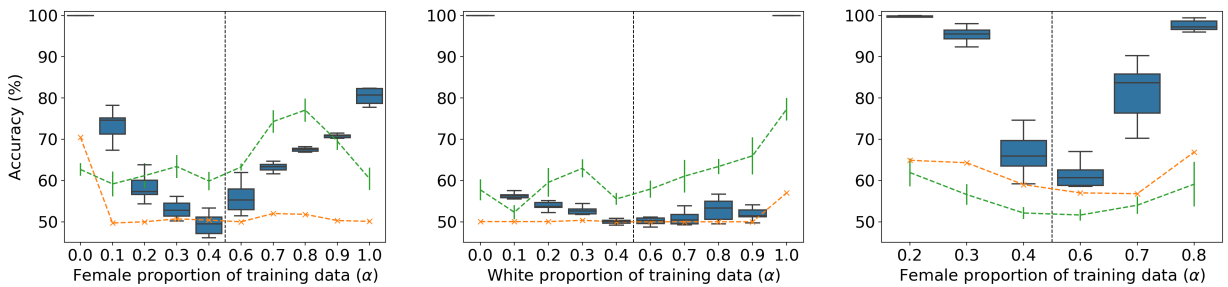


Figure 1: Classification accuracy for distinguishing between training distributions for unseen models for on the Census (females: left, whites: middle) and RSNA Bone Age (right) datasets. Orange crosses are for the Loss Test; green with error bars are the Threshold Test; the blue box-plots are the meta-classifiers.

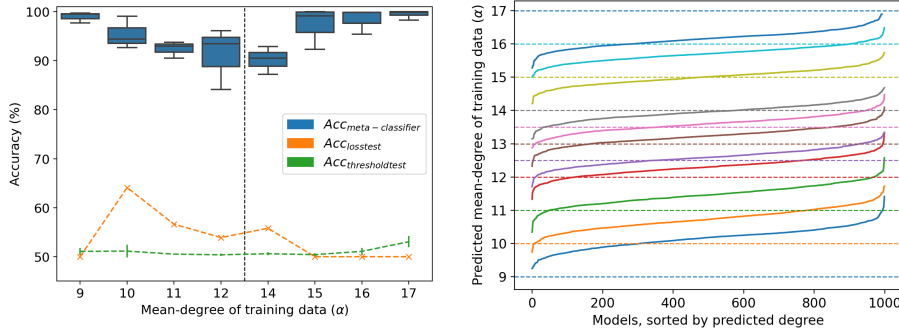


Figure 2: Performance on unseen models on the ogbn-arxiv dataset, for binary classification (left) and direct  $\alpha$  prediction via regression (right). For regression, each color represents the true degree (dashed lines) of the models being tested.

are equally divergent from  $\mathcal{G}_0(\mathcal{D})$  with respect to female proportion. Even more surprising is the case for race. The meta-classifiers barely do better than random guessing except at the two extremes (0.0 (no whites) and 1.0 (all whites)).

Ganju et al. [5] applied their meta-classifier method on two properties on this dataset: 38% v/s 65% women (case A), and 0% v/s 87% whites (case B). For case B, the loss-test performs as well as random guessing (50.1%), while the threshold-test ( $92.4 \pm 2.6\%$ ) approaches meta-classifier performance ( $99.0 \pm 0.1\%$ ). For case A, the threshold-based method ( $62.7 \pm 2.0\%$ ) out-performs the meta-classifiers ( $62.1 \pm 1.7\%$ ), while the loss test method fails (50%).<sup>2</sup>

**RSNA Bone Age.** The meta-classifier method is able to distinguish the proportion of females in the training distribution, with a substantial adversarial advantage even when the proportions are close to 0.5 (rightmost subfigure in Figure 1). Accuracy improves as the distributions diverge, however, this trend is not symmetric—the meta-classifier is more accurate for  $\alpha = 0.3$  than  $\alpha = 0.7$ , even though both are equally divergent from  $\mathcal{G}_0(\mathcal{D})$ . The loss test performs marginally better than random guessing in some cases and yields non-trivial accuracies in others. For cases like  $\alpha = 0.4$ , it even comes close to matching the meta-classifier performance, with a performance gap of less than 4%. However, this method’s performance does not increase as distributions diverge from each other (Figure 1), and the threshold test consistently performs worse than the loss test.

**ogbn-arxiv.** We vary  $\alpha$  in [9, 17] at intervals of one, producing test datasets by pruning either high or low-degree nodes from the original graph to achieve a desired  $\alpha$ . Similar to RSNA Bone Age, meta-classifier performance increases as the distributions diverge, albeit with much smaller drops. The meta-classifier is highly successful at differentiating between the two distributions: even when the difference in mean node-degrees is just one (Figure 2). Both loss test and threshold test fail on this dataset.

We also trained a regression variant of the meta-classifier to predict  $\alpha$  directly. We vary  $\alpha$  in the same range. The resulting meta-classifier performs quite well, achieving a mean-squared error (MSE) loss of  $0.393 \pm 0.360$ . It generalizes well to unseen distributions, achieving an MSE loss of 0.076 on  $\alpha = \{12.5, 13.5\}$ . A property inference adversary can thus be strong enough not just to distinguish between two possible distributions, but to infer the actual average node degree of the training dataset.

## 4 Conclusions

We propose a general definition for property inference attacks, focusing on distributions instead of datasets. Our definition supports arbitrary properties like the mean node-degree of a graph, which we demonstrate a successful adversary can infer. A principled approach to executing such attacks reveals how intuition may not necessarily align with actual performance—seemingly similar pairs of distributions can have starkly different attack success rates, and simple attacks can yield non-trivial accuracy. We expect there is room for improving attacks that do not rely on meta-classifiers, and studying the performance of such attacks will give more valuable insights into when properties can be inferred.

<sup>2</sup>Ganju et al. report 97% accuracy for case A and 100% for case B. We were able to reproduce these results in the original setting, but conducted the reported experiments in a more realistic setting where we ensure there is no victim/adversary overlap and maintain the label ratios and same dataset sizes (Appendix A.2). This leads to drops in meta-classifier accuracy, especially in case A, but reduces the risk that the inference is actually based on other differences between the distributions used, not just the inferred property.

## Acknowledgements

This work was partially supported by grants from the National Science Foundation (#1717950 and #1915813).

## References

- [1] G. Ateniese, L. V. Mancini, A. Spognardi, A. Villani, D. Vitali, and G. Felici. Hacking Smart Machines with Smarter Ones: How to Extract Meaningful Data from Machine Learning Classifiers. *International Journal of Security and Networks*, 10(3):137–150, 2015.
- [2] S. D. Bay, D. Kibler, M. J. Pazzani, and P. Smyth. The UCI KDD Archive of Large Data Sets for Data Mining Research and Experimentation. *ACM SIGKDD Explorations Newsletter*, 2(2):81–85, 2000.
- [3] M. Chase, E. Ghosh, and S. Mahloujifar. Property Inference from Poisoning. *arXiv:2101.11073*, 2021.
- [4] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart. Privacy in Pharmacogenetics: An end-to-end case study of personalized Warfarin dosing. In *23rd USENIX Security Symposium*, 2014.
- [5] K. Ganju, Q. Wang, W. Yang, C. A. Gunter, and N. Borisov. Property Inference Attacks on Fully Connected Neural Networks using Permutation Invariant Representations. In *ACM Conference on Computer and Communications Security*, 2018.
- [6] D. Gopinath, H. Converse, C. Pasareanu, and A. Taly. Property Inference for Deep Neural Networks. In *34th IEEE/ACM International Conference on Automated Software Engineering*. IEEE, 2019.
- [7] S. S. Halabi, L. M. Prevedello, J. Kalpathy-Cramer, A. B. Mamonov, A. Bilbily, M. Cicero, I. Pan, L. A. Pereira, R. T. Sousa, N. Abdala, et al. The RSNA Pediatric Bone Age Machine Learning Challenge. *Radiology*, 290(2): 498–503, 2019.
- [8] M. Hay, C. Li, G. Miklau, and D. Jensen. Accurate Estimation of the Degree Distribution of Private Networks. In *Ninth IEEE International Conference on Data Mining*, 2009.
- [9] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In *Advances in Neural Information Processing Systems*, 2020.
- [10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely Connected Convolutional Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.
- [11] D. Kifer and A. Machanavajjhala. Pufferfish: A Framework for Mathematical Privacy Definitions. *ACM Transactions on Database Systems (TODS)*, 2014.
- [12] T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [13] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership Inference Attacks against Machine Learning Models. In *IEEE Symposium on Security and Privacy*, 2017.
- [14] A. Suri and D. Evans. Formalizing and Estimating Distribution Inference Risks. *arXiv preprint arXiv:2109.06024*, 2021.
- [15] K. Wang, Z. Shen, C. Huang, C.-H. Wu, Y. Dong, and A. Kanakia. Microsoft Academic Graph: When experts are not enough. *Quantitative Science Studies*, 1(1):396–413, 2020.
- [16] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha. Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting. In *31st IEEE Computer Security Foundations Symposium*, 2018.
- [17] W. Zhang, S. Tople, and O. Ohrimenko. Dataset-Level Attribute Leakage in Collaborative Learning. *arXiv:2006.07267*, 2020.

## A Experimental Setup

### A.1 Data and Model Preparation

To ensure a fair comparison between models, we ensure that sample sizes per dataset are comparable for all distributions considered in our attacks. Furthermore, these samples are randomly selected, ruling out the possibility of property inference attacks inadvertently performing some form of dataset inference: differentiating between exact datasets instead of their properties.

**Census.** After converting categorical data to one-hot features, the dataset reduces to 42-feature vectors for each datapoint. While performing data splits (for adversary/victim and train/test), we use stratification with the sex/race attribute and class labels, thus ensuring that changing one of the ratios does not inadvertently change any other ratios. We train neural networks with three hidden layers of dimensions (32, 16, 8).

**RSNA Bone Age.** We bin the age label into binary categories, based on whether the age is more than 132 months (label 1) or not (label 0). Picking this threshold leads to a balanced label distribution. Similar to Census, we perform stratified data splits, keeping the label ratios in mind. We train neural networks with two hidden layers of dimensions (128, 64).

**ogbn-arxiv.** We use the publication years to perform all node label splits. This year-based split ensures no overlap in the victim/adversary data or their train/test data. We pre-process the graph following the Open Graph Benchmark [9]. Additionally, we add reverse edges in the graph to convert it to bidirectional, in addition to self-loops. To add randomness to the graph, we randomly prune 1% of the nodes for each experiment. We trained Graph Convolutional Networks [12] with three convolutional layers of 256 channels each.

### A.2 Attack Details

We perform each experiment ten times and report mean values with standard deviation in all of our experiments. Since the loss test uses a fixed test set per experiment, there is no variation in its results. The model trainer produces 1000 models per distribution using its split of data. These models serve as the evaluation set. Thus, any particular experiment distinguishing between two distributions uses 2000 models for evaluation. The adversary trains 1000 models per distribution on its split of data for the threshold loss and meta-classifier experiments.

**Loss Test.** The adversary uses its test data to sample the two test sets  $D_0, D_1$ . Since we use the same test data in evaluations, we turn off sampling (see above) for this setting.

**Threshold Loss.** The adversary uses a small sample size of 50 models per distribution to identify which  $D_k$  leads to a more significant performance gap between its models. After identifying  $k$ , it finds  $\lambda$  that maximizes the classification accuracy between models trained on datasets from the two distributions, using a simple linear search. It then uses  $k$  and  $\lambda$  to generate its predictions.

**Meta-Classifer.** We used permutation invariant networks as our meta-classifier architecture [5]. This architecture is invariant to different neuron orderings in a neural network and performs well even with a few thousand models. The adversary randomly samples 1600 models from its pool to train the meta-classifier. Following experimental designs from prior works, we were able to achieve the accuracies that the authors reported. However, we notice that, for some cases, using our experimental design leads to significant performance drops. Steps like ensuring no overlap in victim/adversary training data, randomly sampled datasets for  $\mathcal{G}_0(\mathcal{D}), \mathcal{G}_1(\mathcal{D})$ , ensuring the same dataset size, introduce stochasticity that makes it for a meta-classifier to perform as well. Thus, these steps help emulate a more realistic scenario but adversely impact meta-classifier in some cases.