CSE 484/M584: Computer Security (and Privacy)

Spring 2025

David Kohlbrenner
dkohlbre@cs

UW Instruction Team: David Kohlbrenner, Yoshi Kohno, Franziska Roesner, Nirvan Tyagi. Thanks to Dan Boneh, Dieter Gollmann, Dan Halperin, John Manferdelli, John Mitchell, Vitaly Shmatikov, Bennet Yee, and many others for sample slides and materials

Admin

- Lab 4 : A+B Due Friday

- Tomorrows Section:
  - Discussion about Lab 4 components
  - Office hours for lab 4

ML/AI and Security

- ML/AI *for* security?
- Security *for* ML/AI?

Machine Learning (and AI)

	Classification	Generative
Security *for*		
For security		

ML/AI and Security

- ML/AI *for* security?
  - ML has been successful and useful.
  - LLM era largely ineffective so far.
    - Lots of low quality/inaccurate research.
    - Some reported successes (see XBOW?)

- Security *for* ML/AI?
  - ML has sensitive data and is used in critical applications: Huge opportunity!
  - LLM era
    - Not doing well-thought-out security+privacy. Not today’s topic.

Machine Learning Setting

(Not the current LLM stuff)

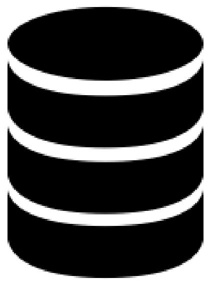

Training data

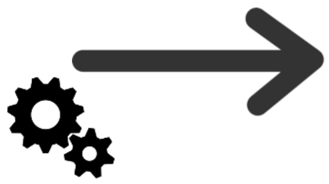

Training algorithm

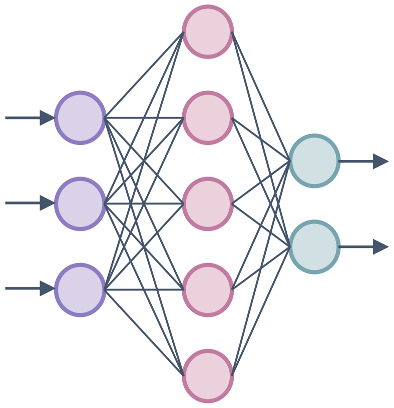

Model

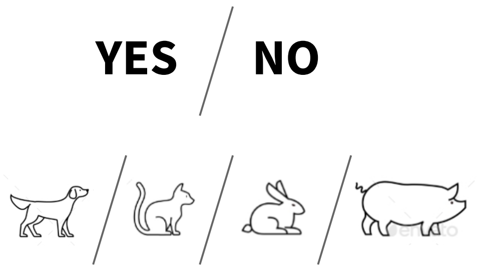

Prediction

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans

Evasion attacks ("Adversarial examples")

Training data

Training algorithm

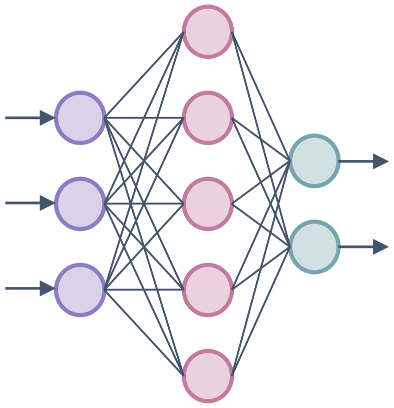

Model

YES / NO

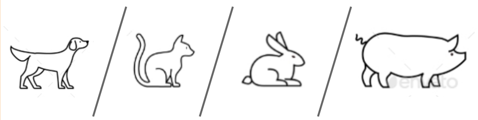

Prediction

Evasion attacks ("Adversarial examples")

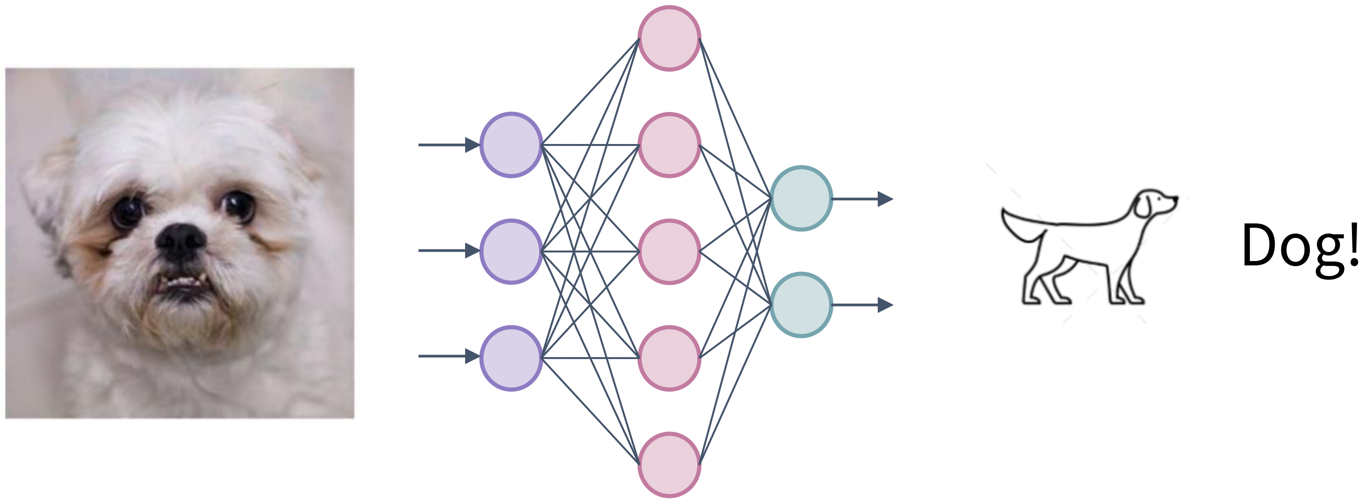

Evasion attacks ("Adversarial examples")

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

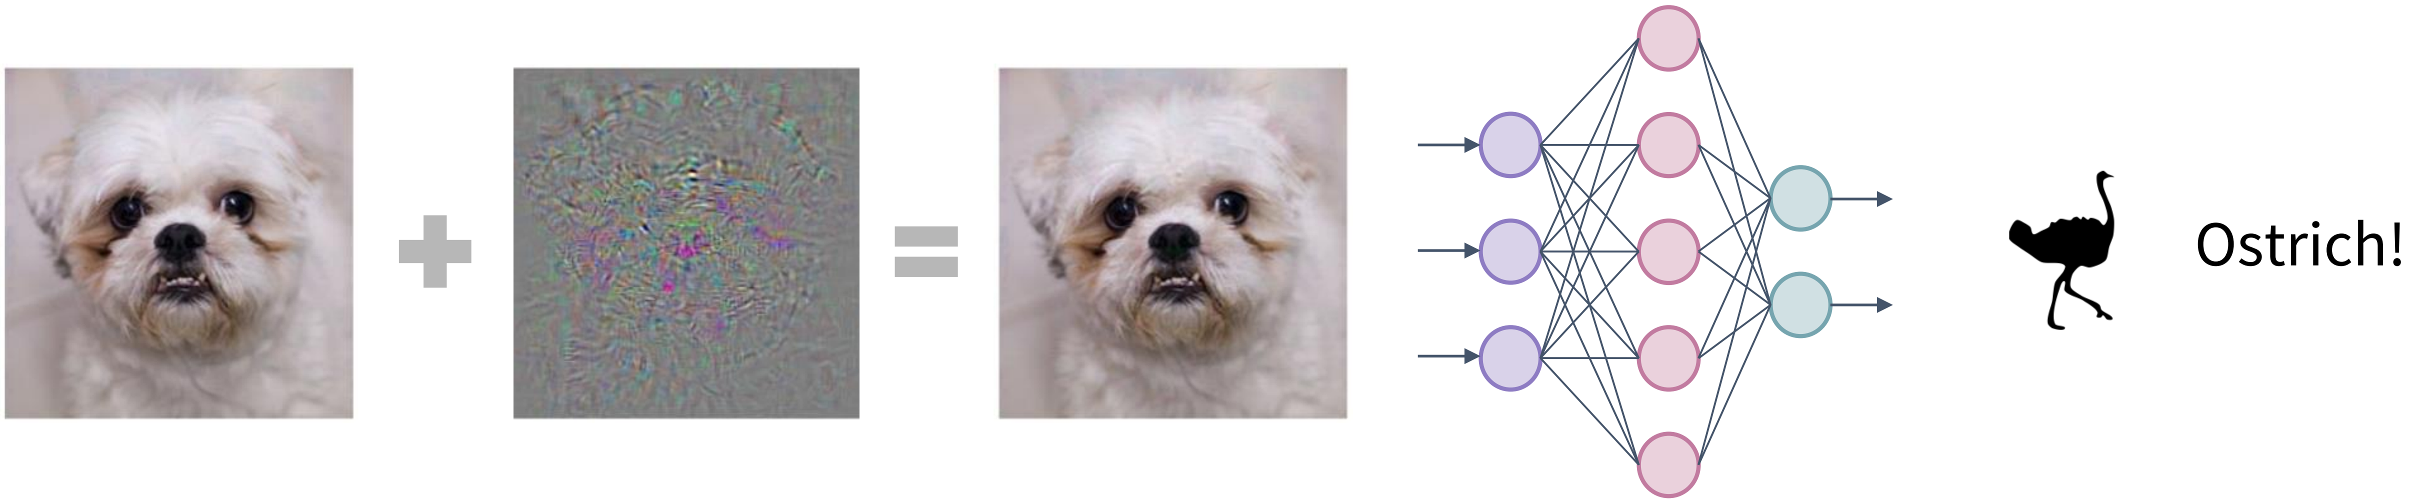

Evasion attacks (“Adversarial examples”)

Variant: Data Poisoning Attacks

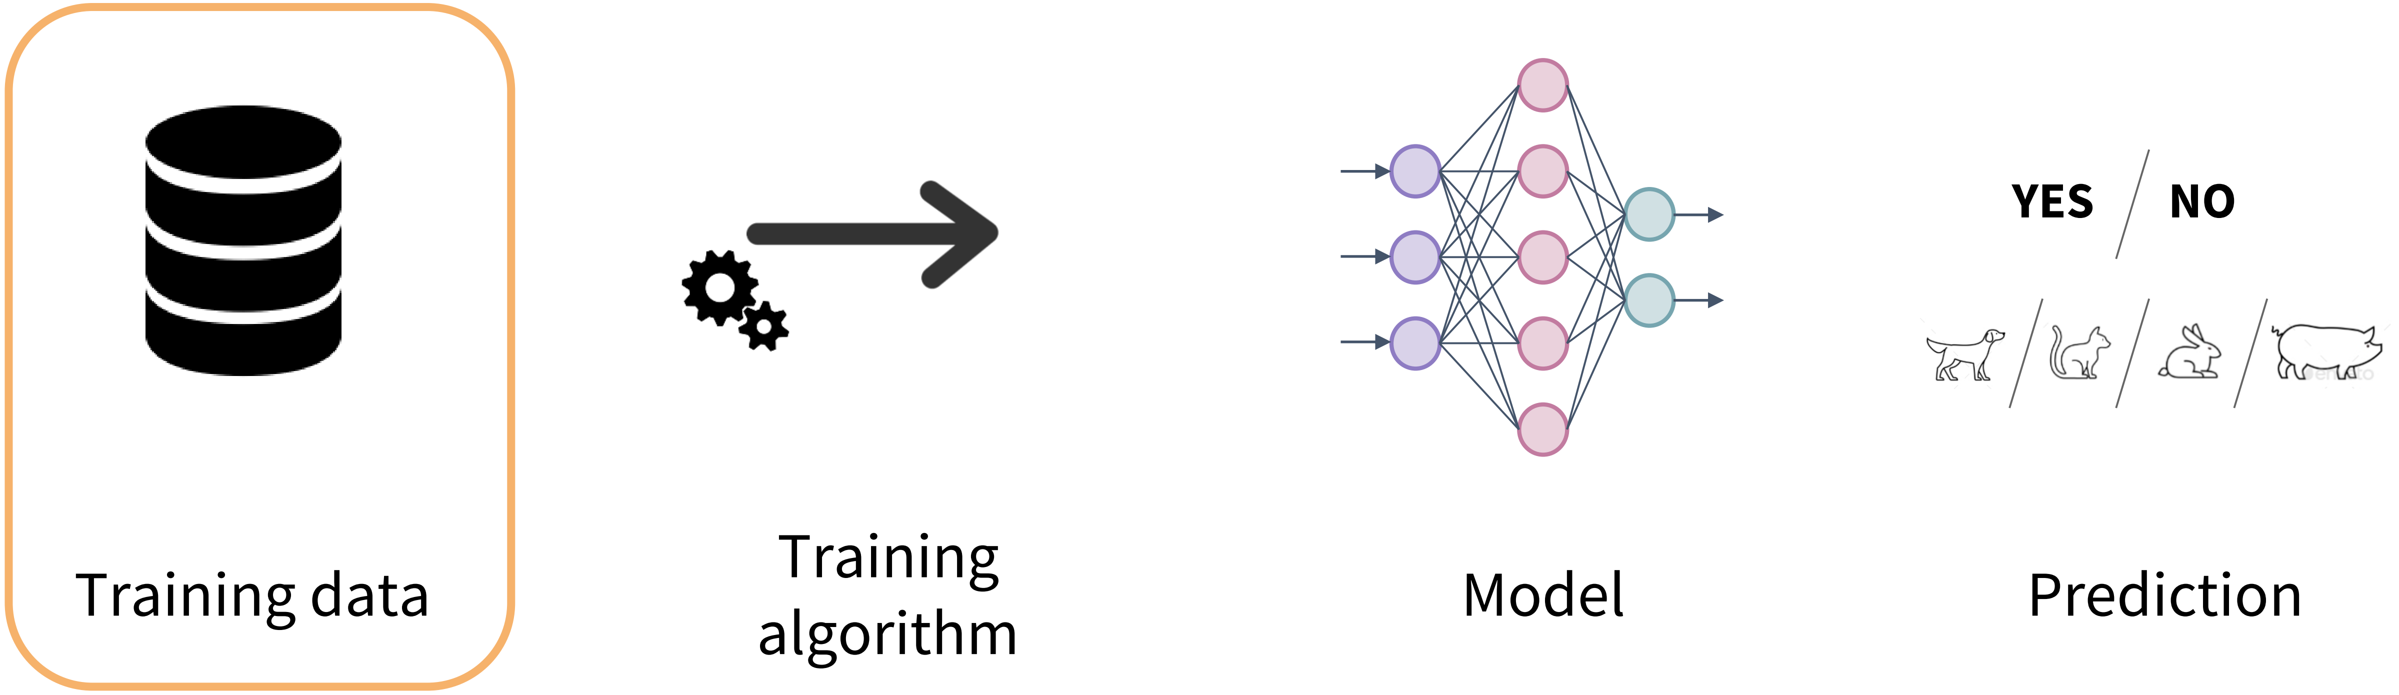

Data Poisoning Attacks

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

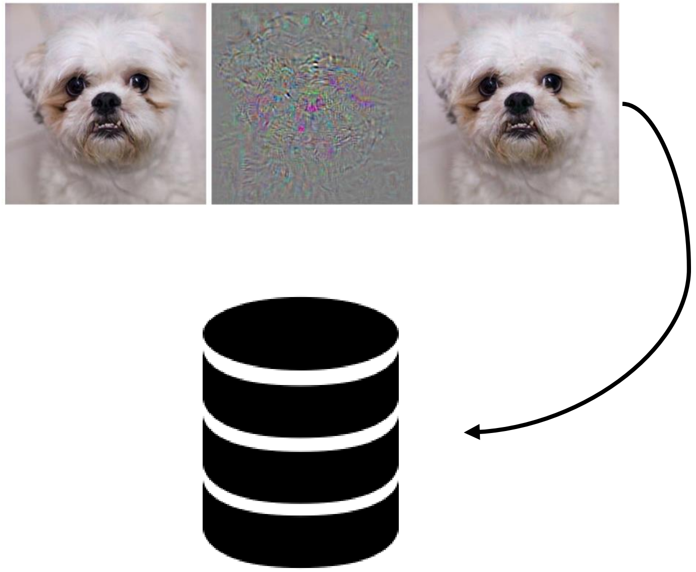

Training data

Data Poisoning Attacks

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

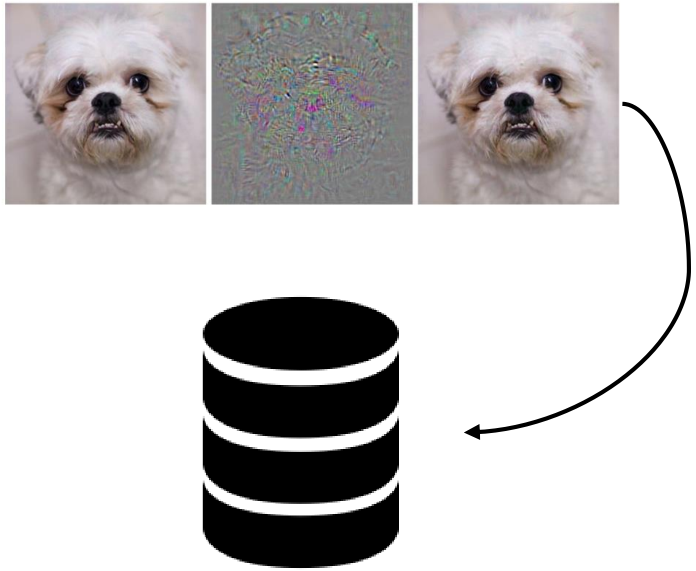

Training data

Training algorithm

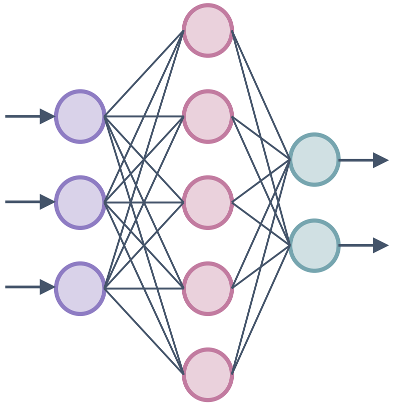

Model

Data Poisoning Attacks

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

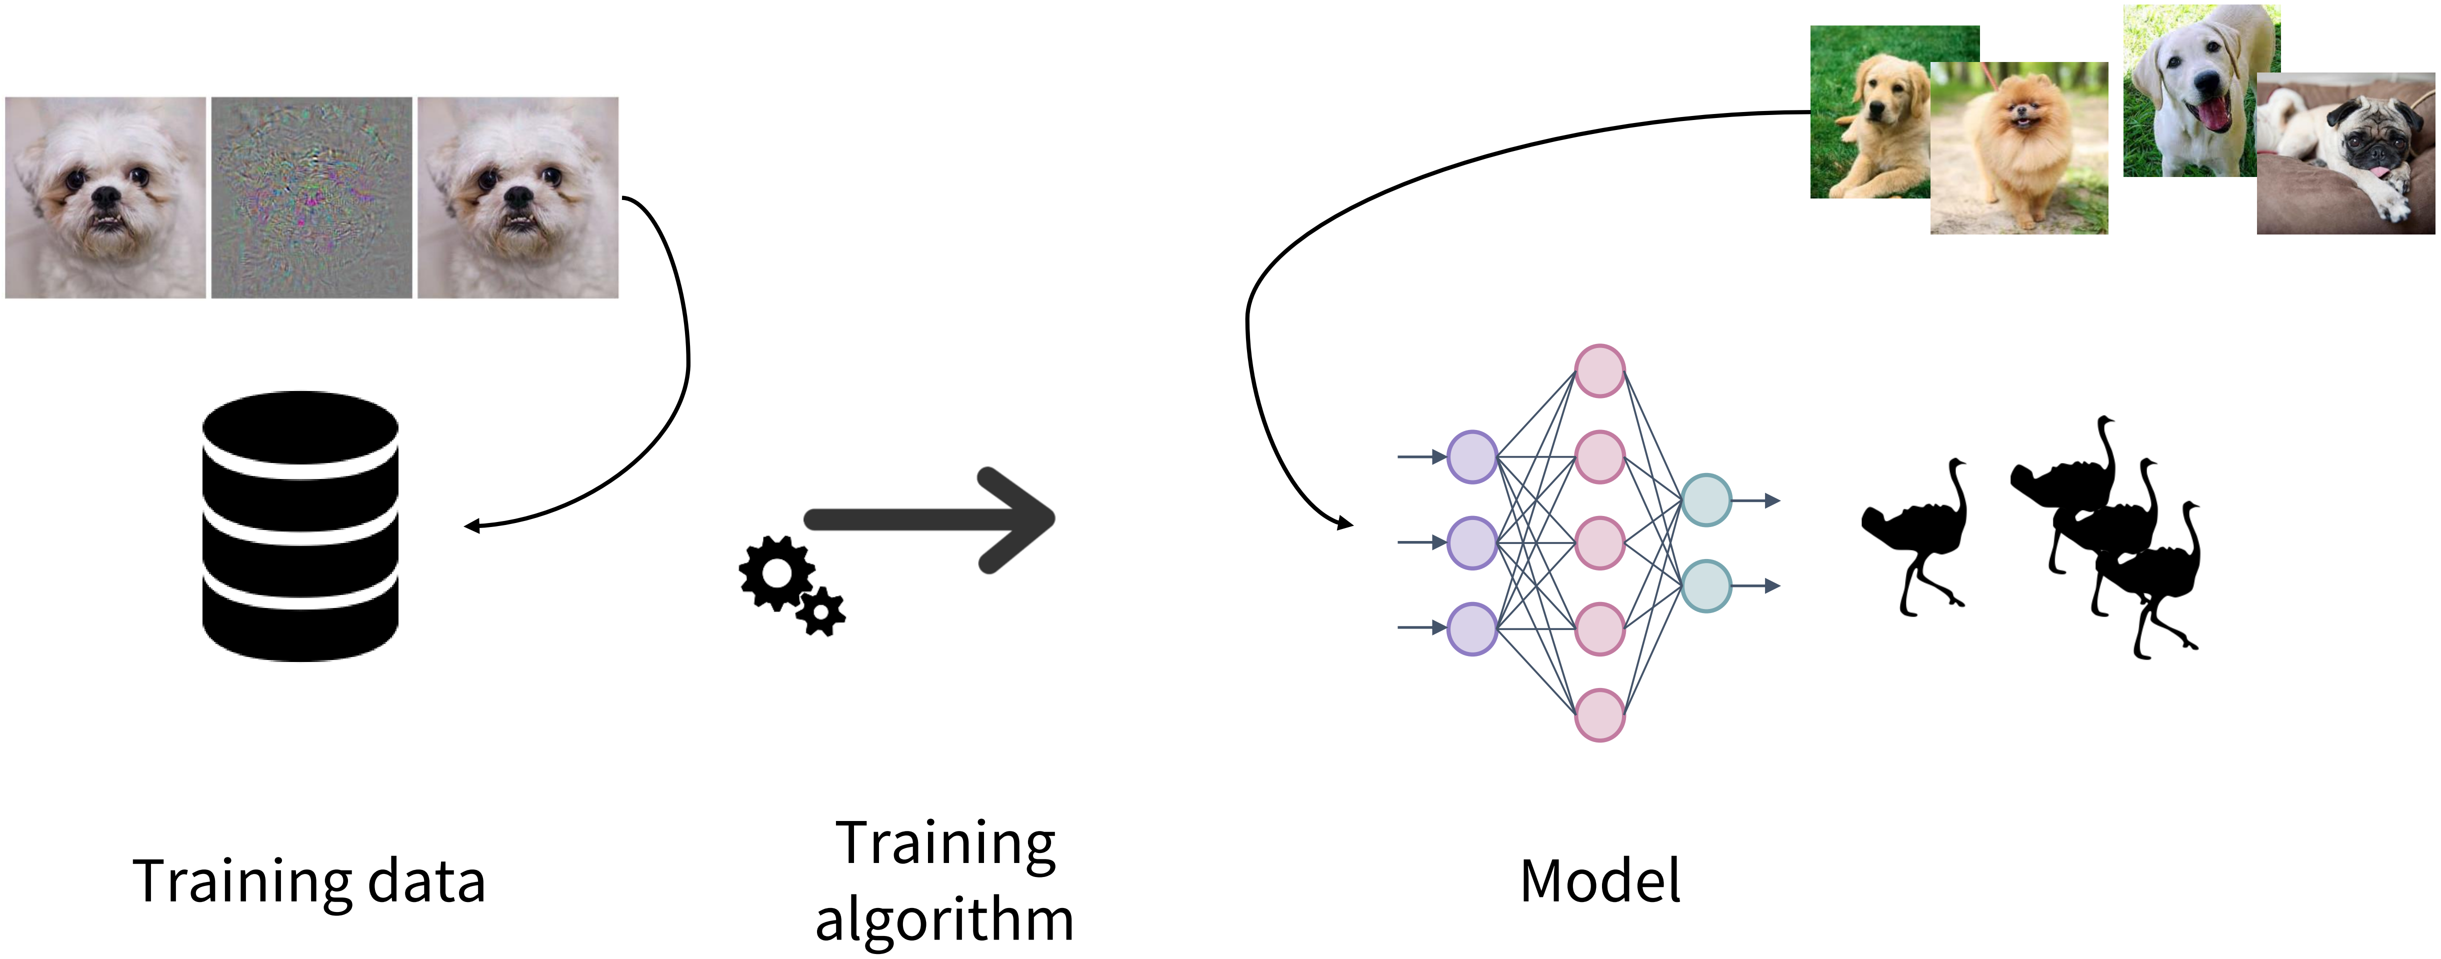

Evasion attacks (“Adversarial examples”)

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

Evasion attacks ("Adversarial examples")

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

Robust Physical-World Attacks on Deep Learning Visual Classification. Eykholt et al. CVPR 2018

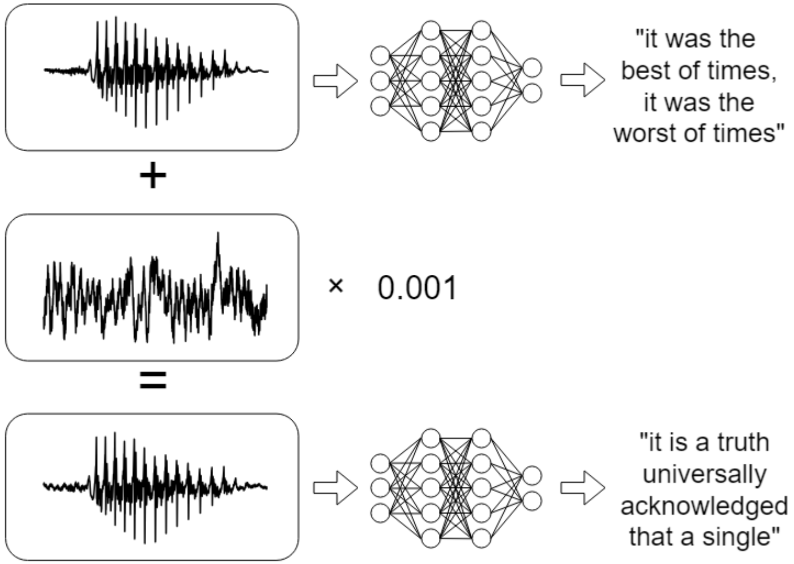

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text. Nicholas Carlini, David Wagner

Evasion attacks ("Adversarial examples")

ALEX LEE, WIRED UK   SECURITY   MAY 11, 2020 1:00 AM

This ugly t-shirt makes you invisible to facial recognition tech

Researchers at Northeastern University have developed an adversarial example that works even when printed onto a moving fabric

fications [SZSB+14, and MANY more]

raining set (i.e. data poisoning) causes

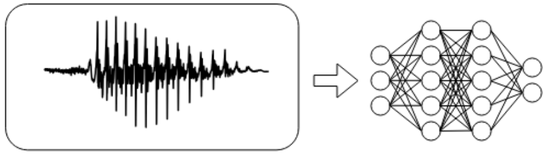

Robust Physical-World Attacks on Deep Learnin

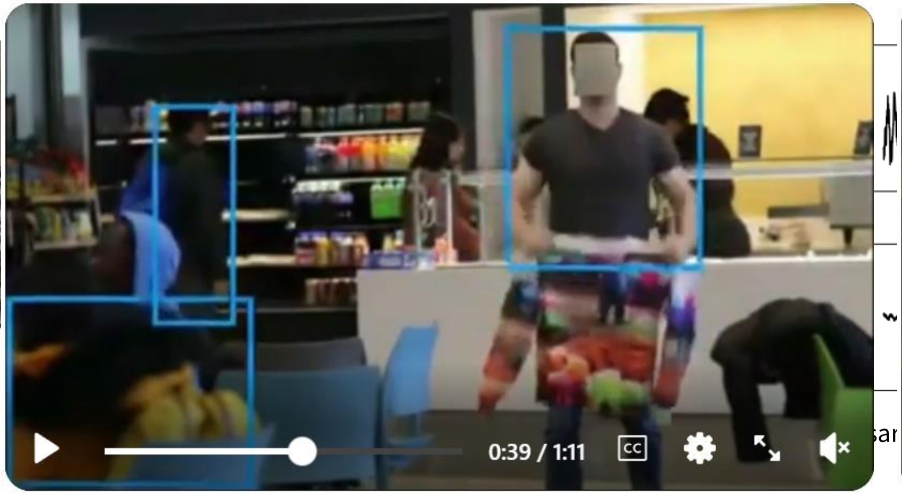

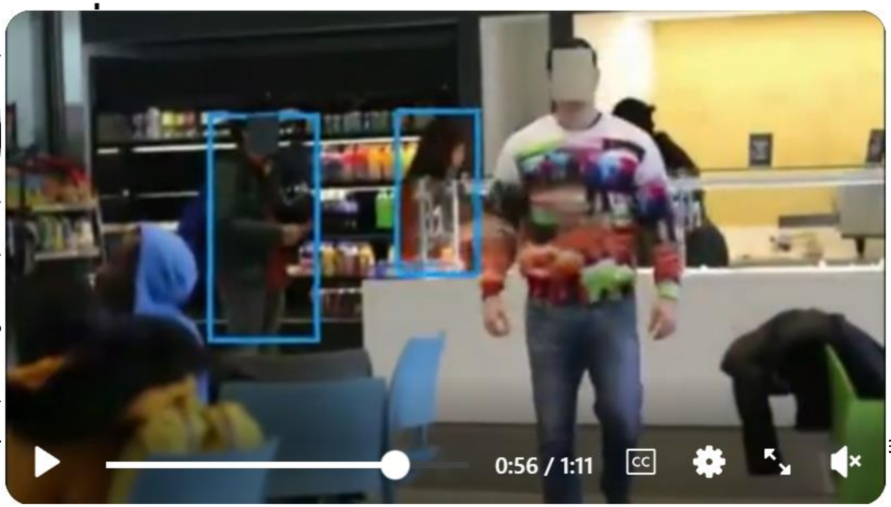

Evasion attacks ("Adversarial examples")

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

How to build *robust* models? [MMSTV17, SKL17, RSL18]

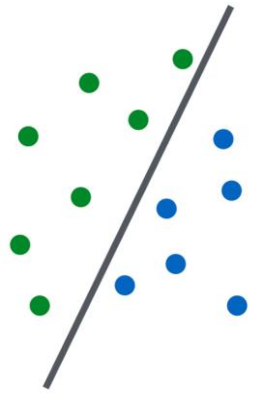

Evasion attacks ("Adversarial examples")

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

How to build *robust* models? [MMSTV17, SKL17, RSL18]

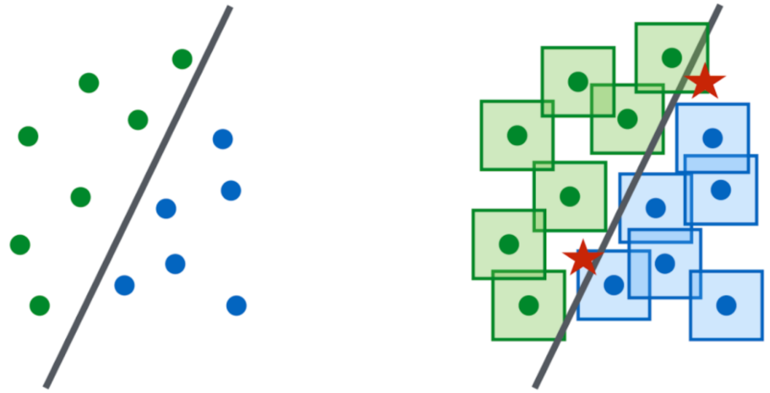

Evasion attacks (“Adversarial examples”)

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

How to build *robust* models? [MMSTV17, SKL17, RSL18]

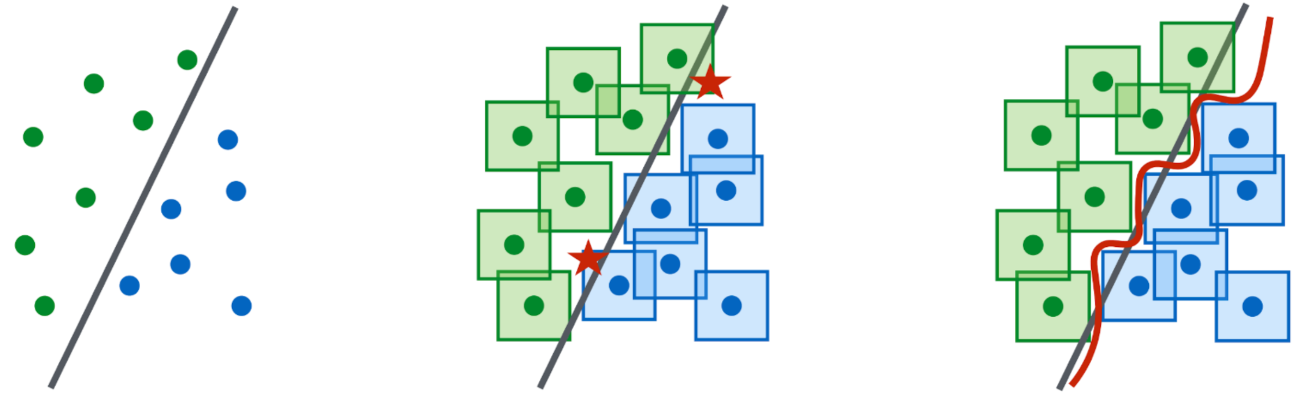

Evasion attacks ("Adversarial examples")

Gradescope! Places where non-robust models can still safely be deployed?

Small perturbations to inputs cause misclassifications [SZSB+14, and MANY more]

Adding a few specially-crafted images to the training set (i.e. data poisoning) causes misclassifications [KL17]

How to build *robust* models? [MMSTV17, SKL17, RSL18]

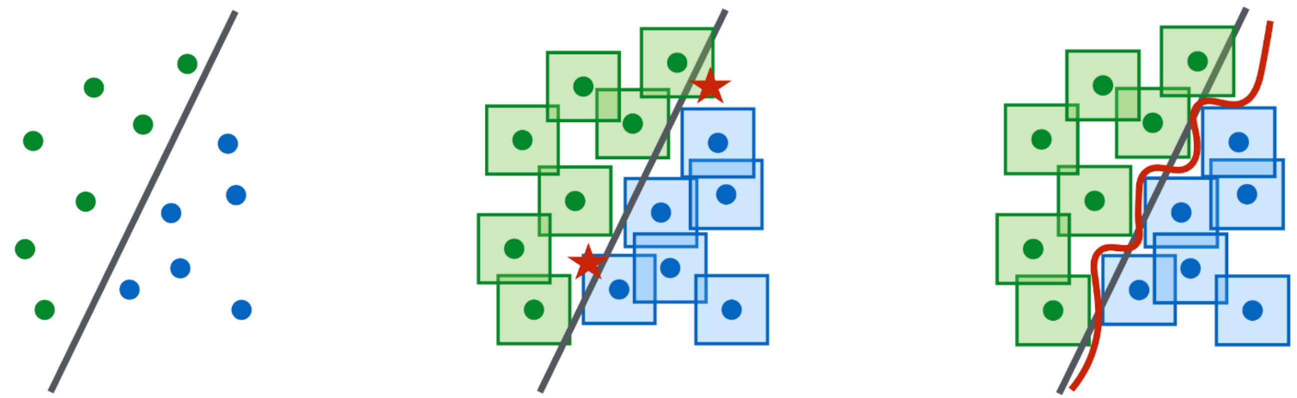

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans

Model extraction attacks (“Model stealing”)

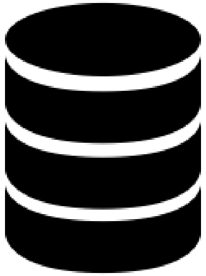

Training data

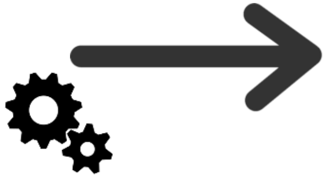

Training algorithm

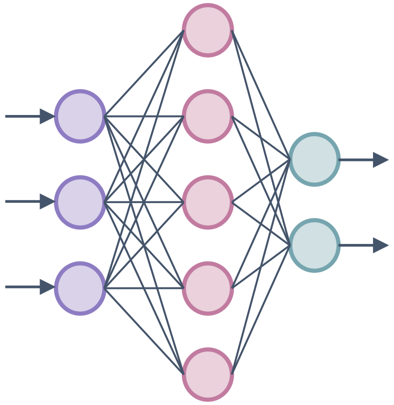

Model

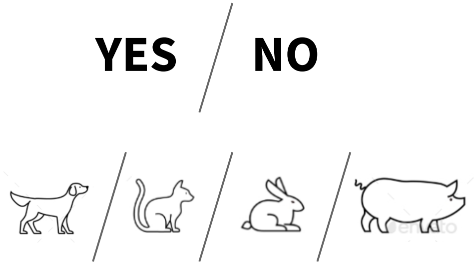

Prediction

Model extraction attacks (“Model stealing”)

Machine Learning as a Service (MLaaS)

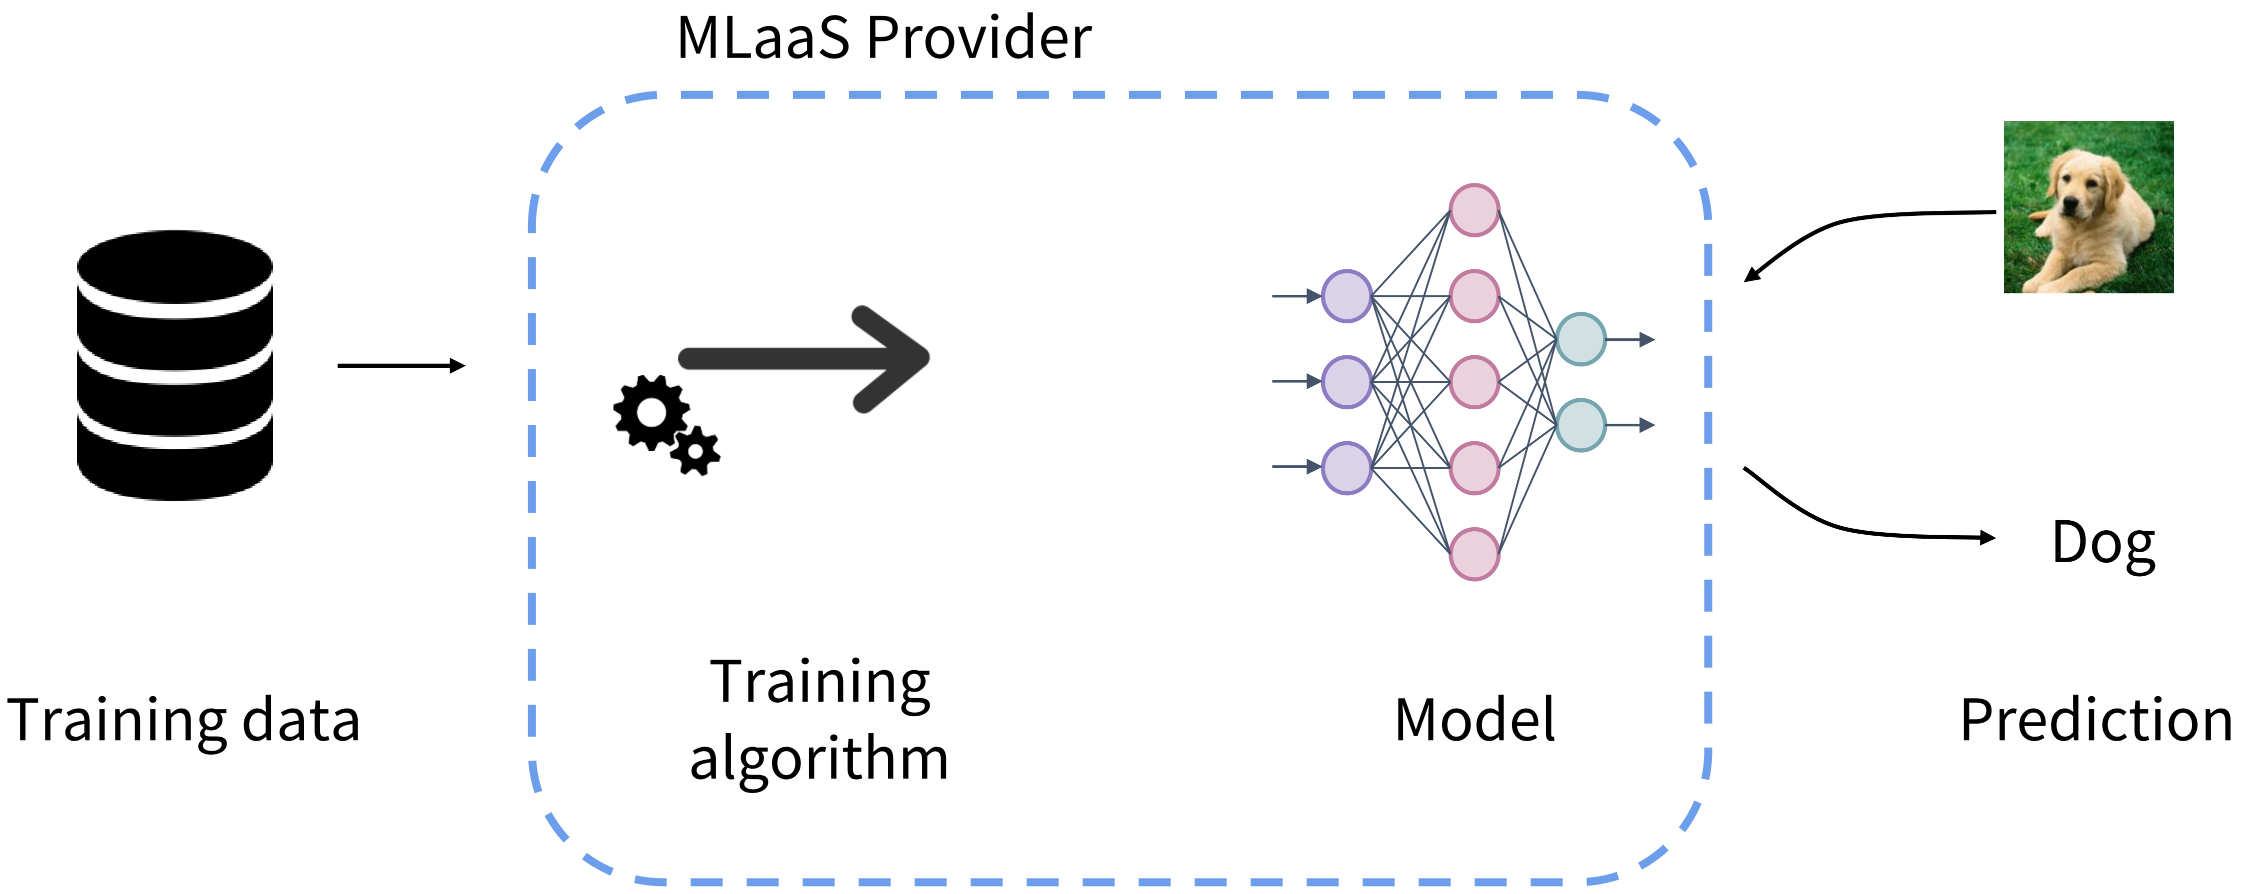

Model extraction attacks (“Model stealing”)

Stealing model parameters through predictions [TZJRR16]

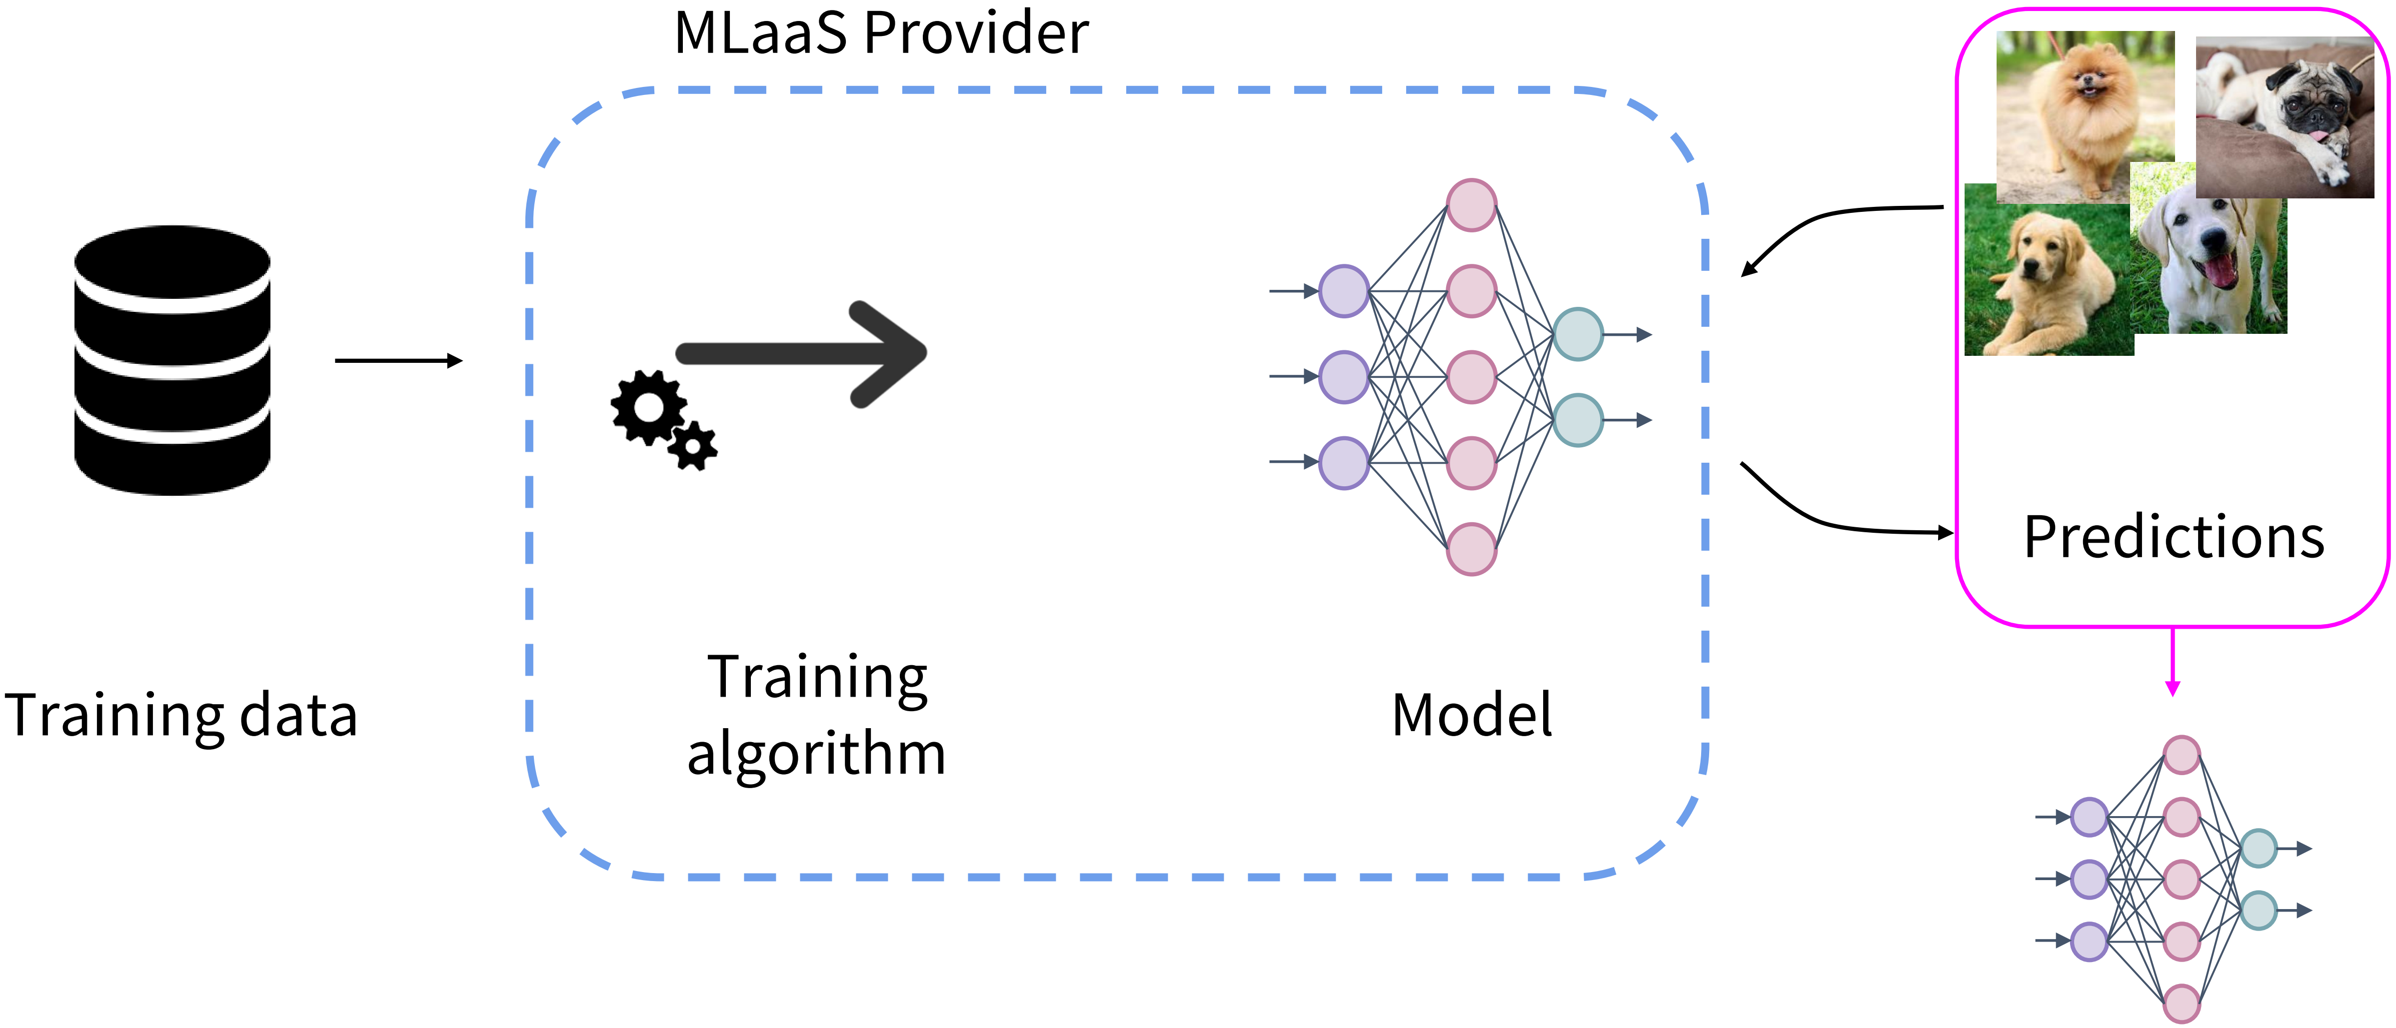

Model extraction attacks ("Model stealing")

Stealing model parameters through predictions [TZJRR16]

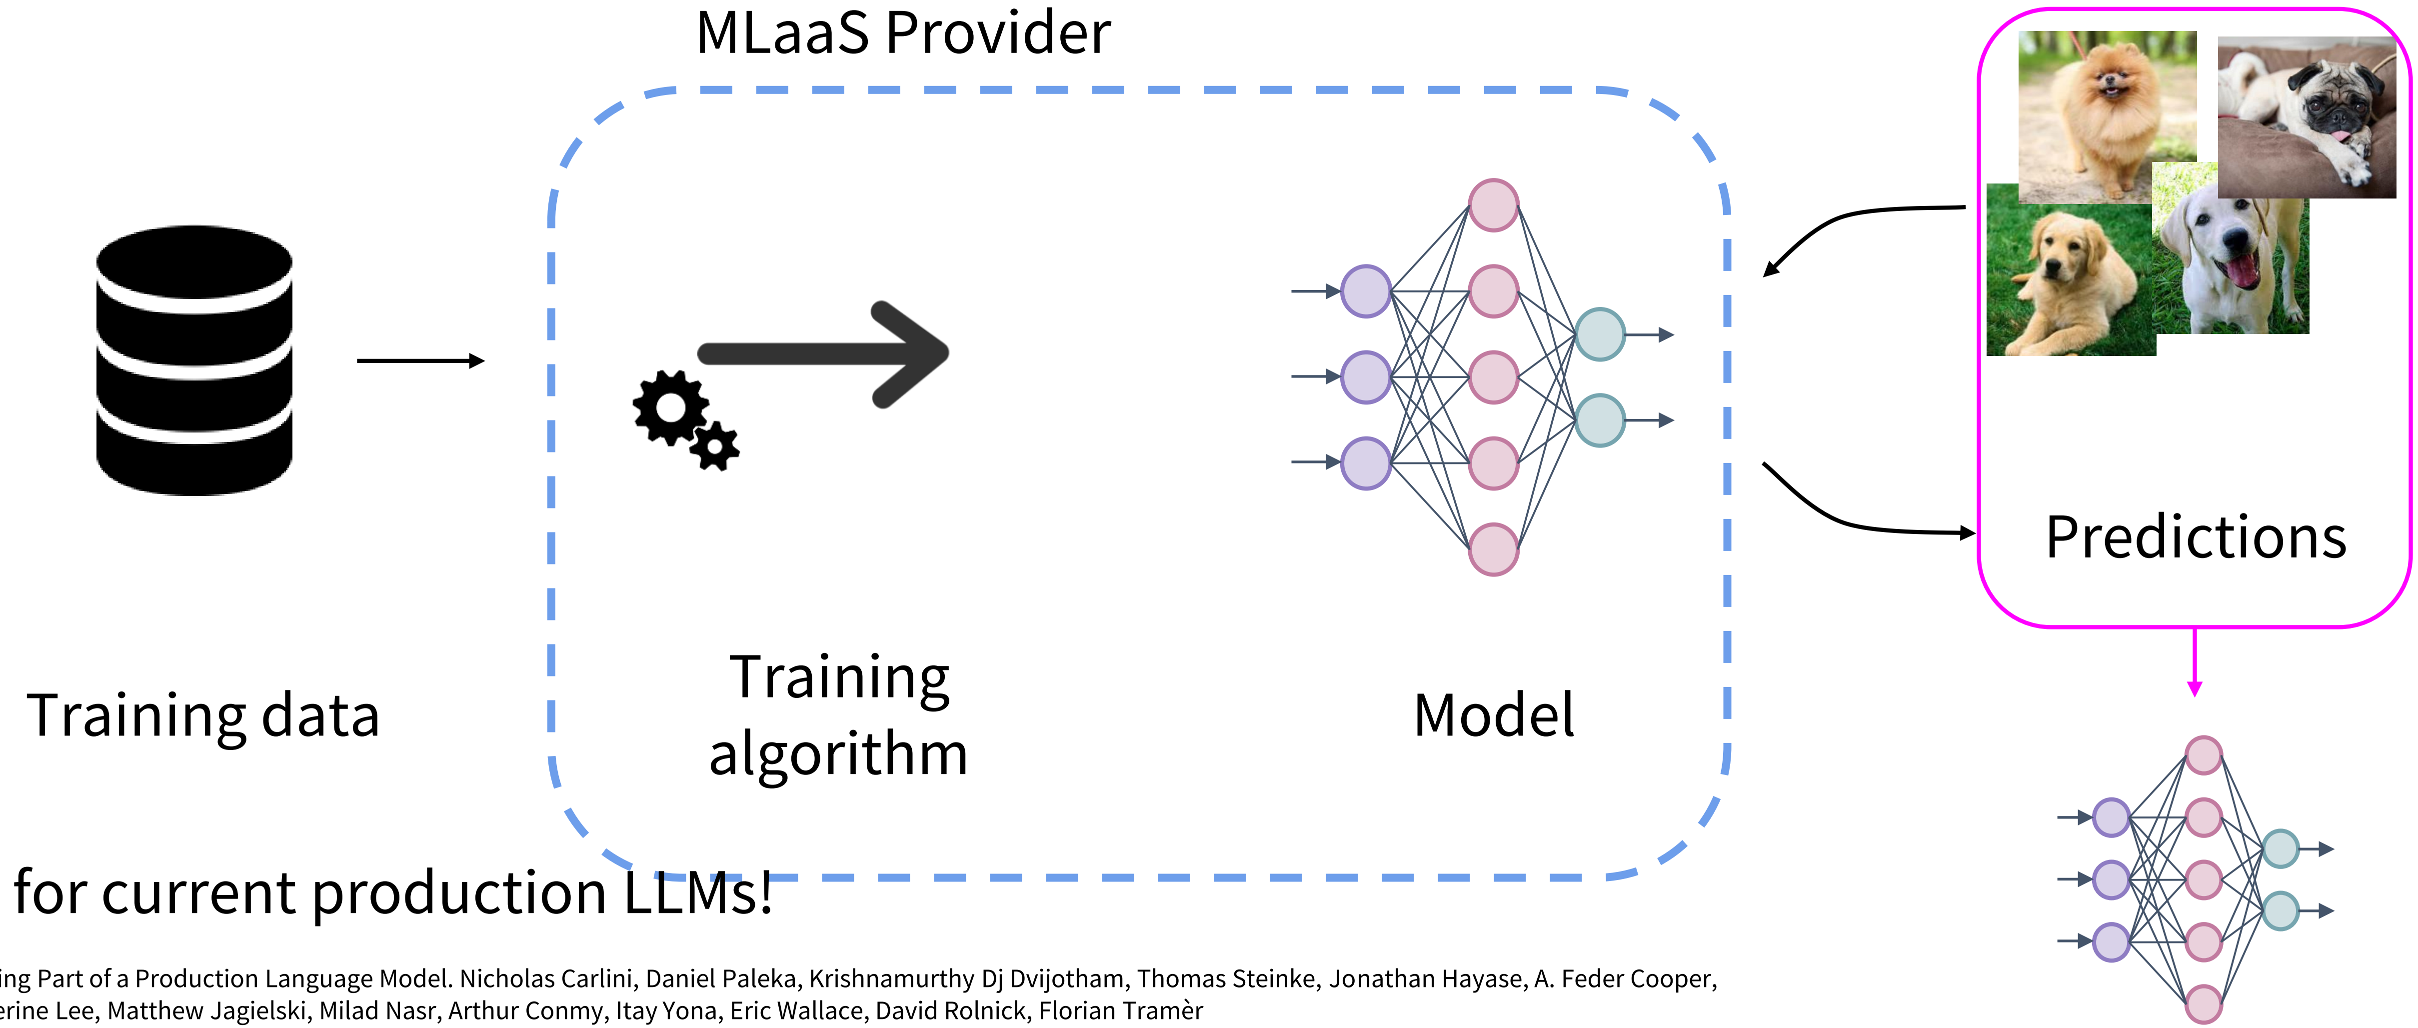

Also for current production LLMs!

Stealing Part of a Production Language Model. Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Itay Yona, Eric Wallace, David Rolnick, Florian Tramèr

Model extraction attacks ("Model stealing")

Stealing model parameters through predictions [TZJRR16]
Stealing model parameters faster through gradients [MSDH18]

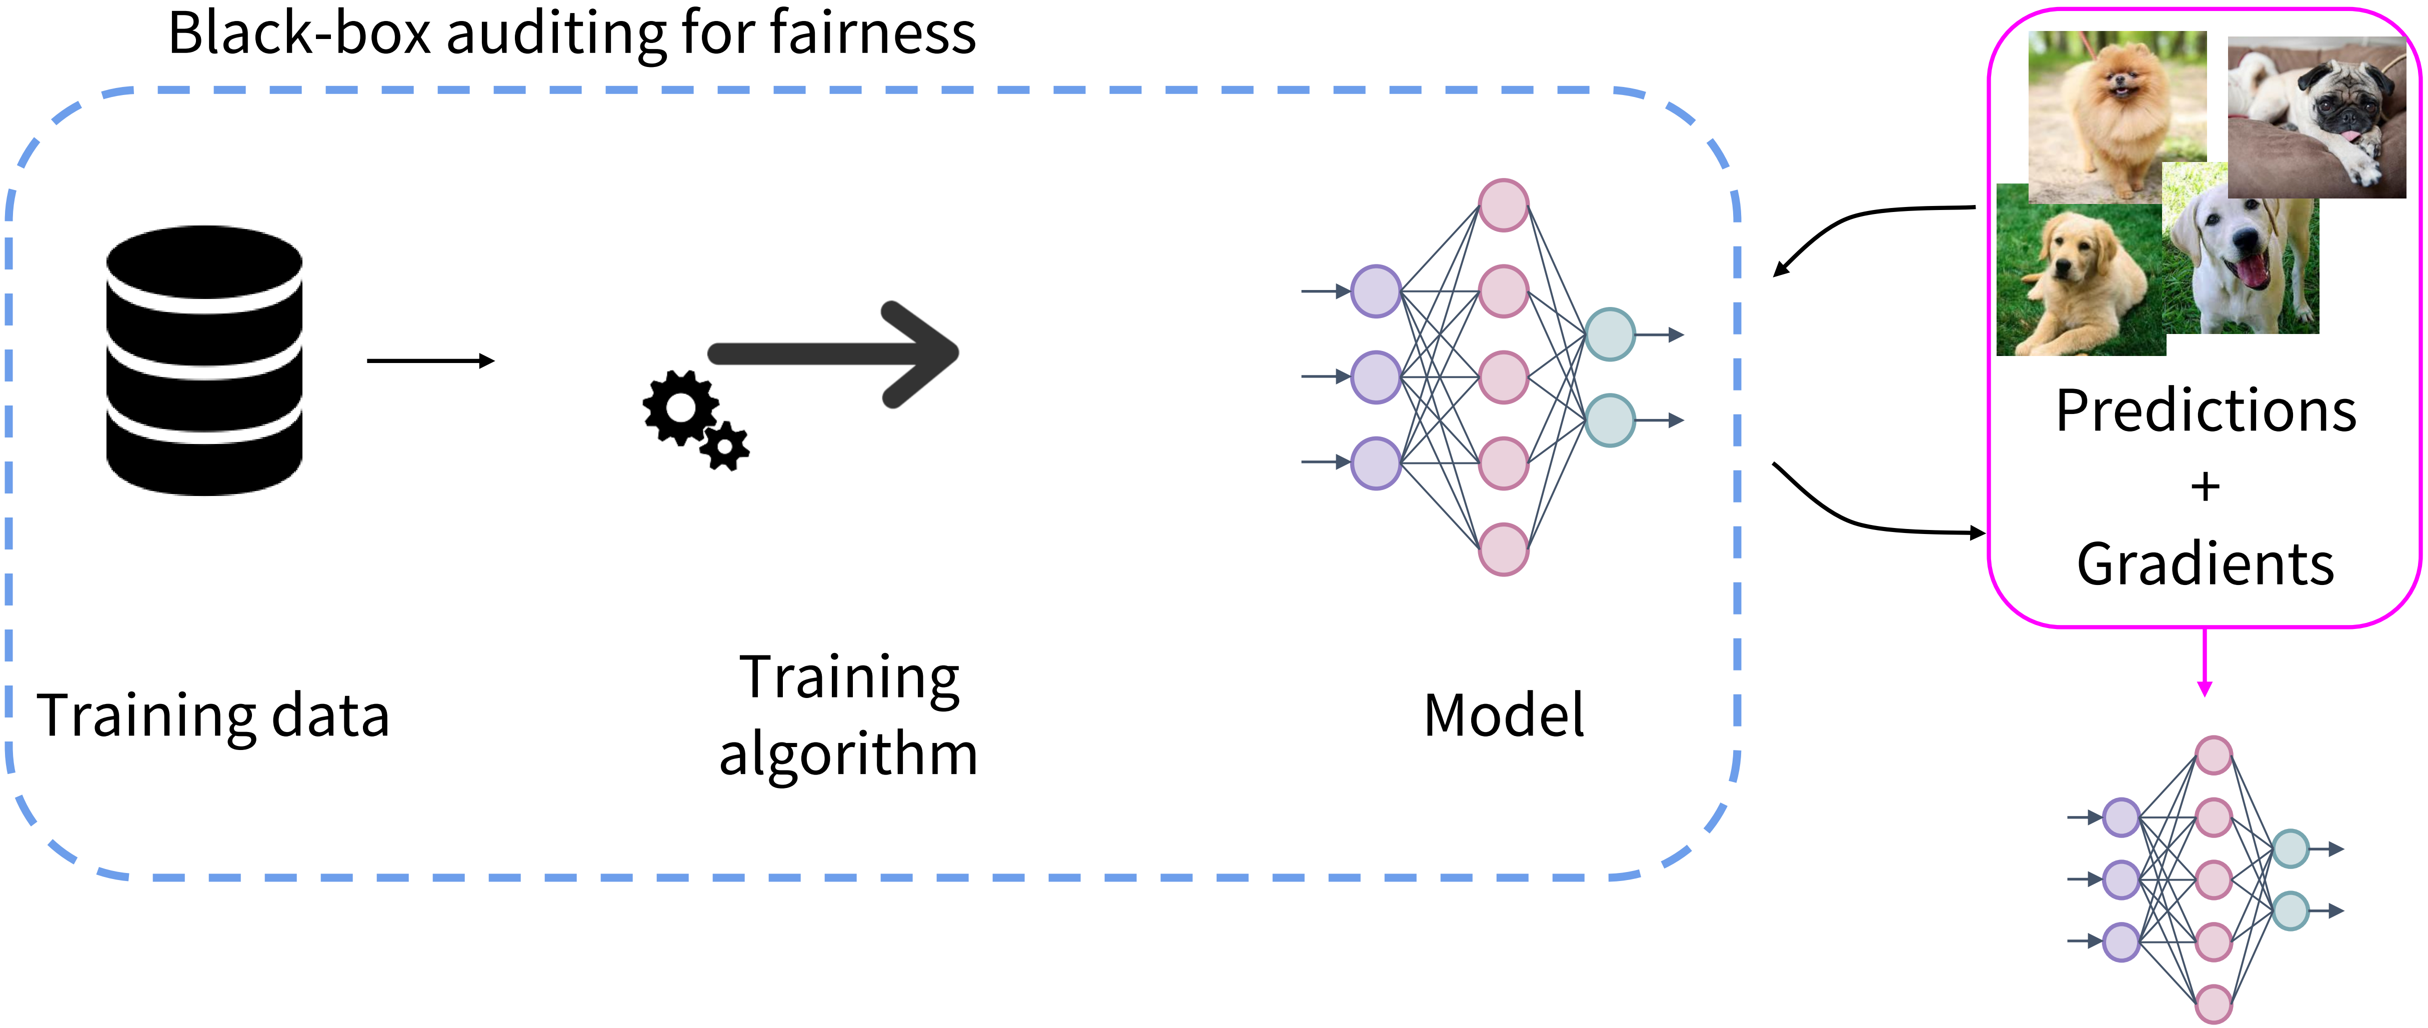

Detour: Auditing Models

Models are used for decision-making systems like loans, credit card approvals, bail rates, fraud detection, etc.

Companies build proprietary models for these purposes, but consumers should be protected against bad/malicious models

Idea: Regulatory agencies should be able to audit models to learn whether decisions abide by regulations, e.g., do not consider protected attributes

F. Tramèr, V. Atlidakis, R. Geambasu, D. Hsu, J.-P. Hubaux, M. Humbert, A. Juels, and H. Lin. FairTest: Discovering Unwarranted Associations in Data-Driven Applications. In IEEE Euro S&P.

Model extraction attacks ("Model stealing")

Stealing model parameters through predictions [TZJRR16]
Stealing model parameters faster through gradients [MSDH18]

Black-box auditing for fairness

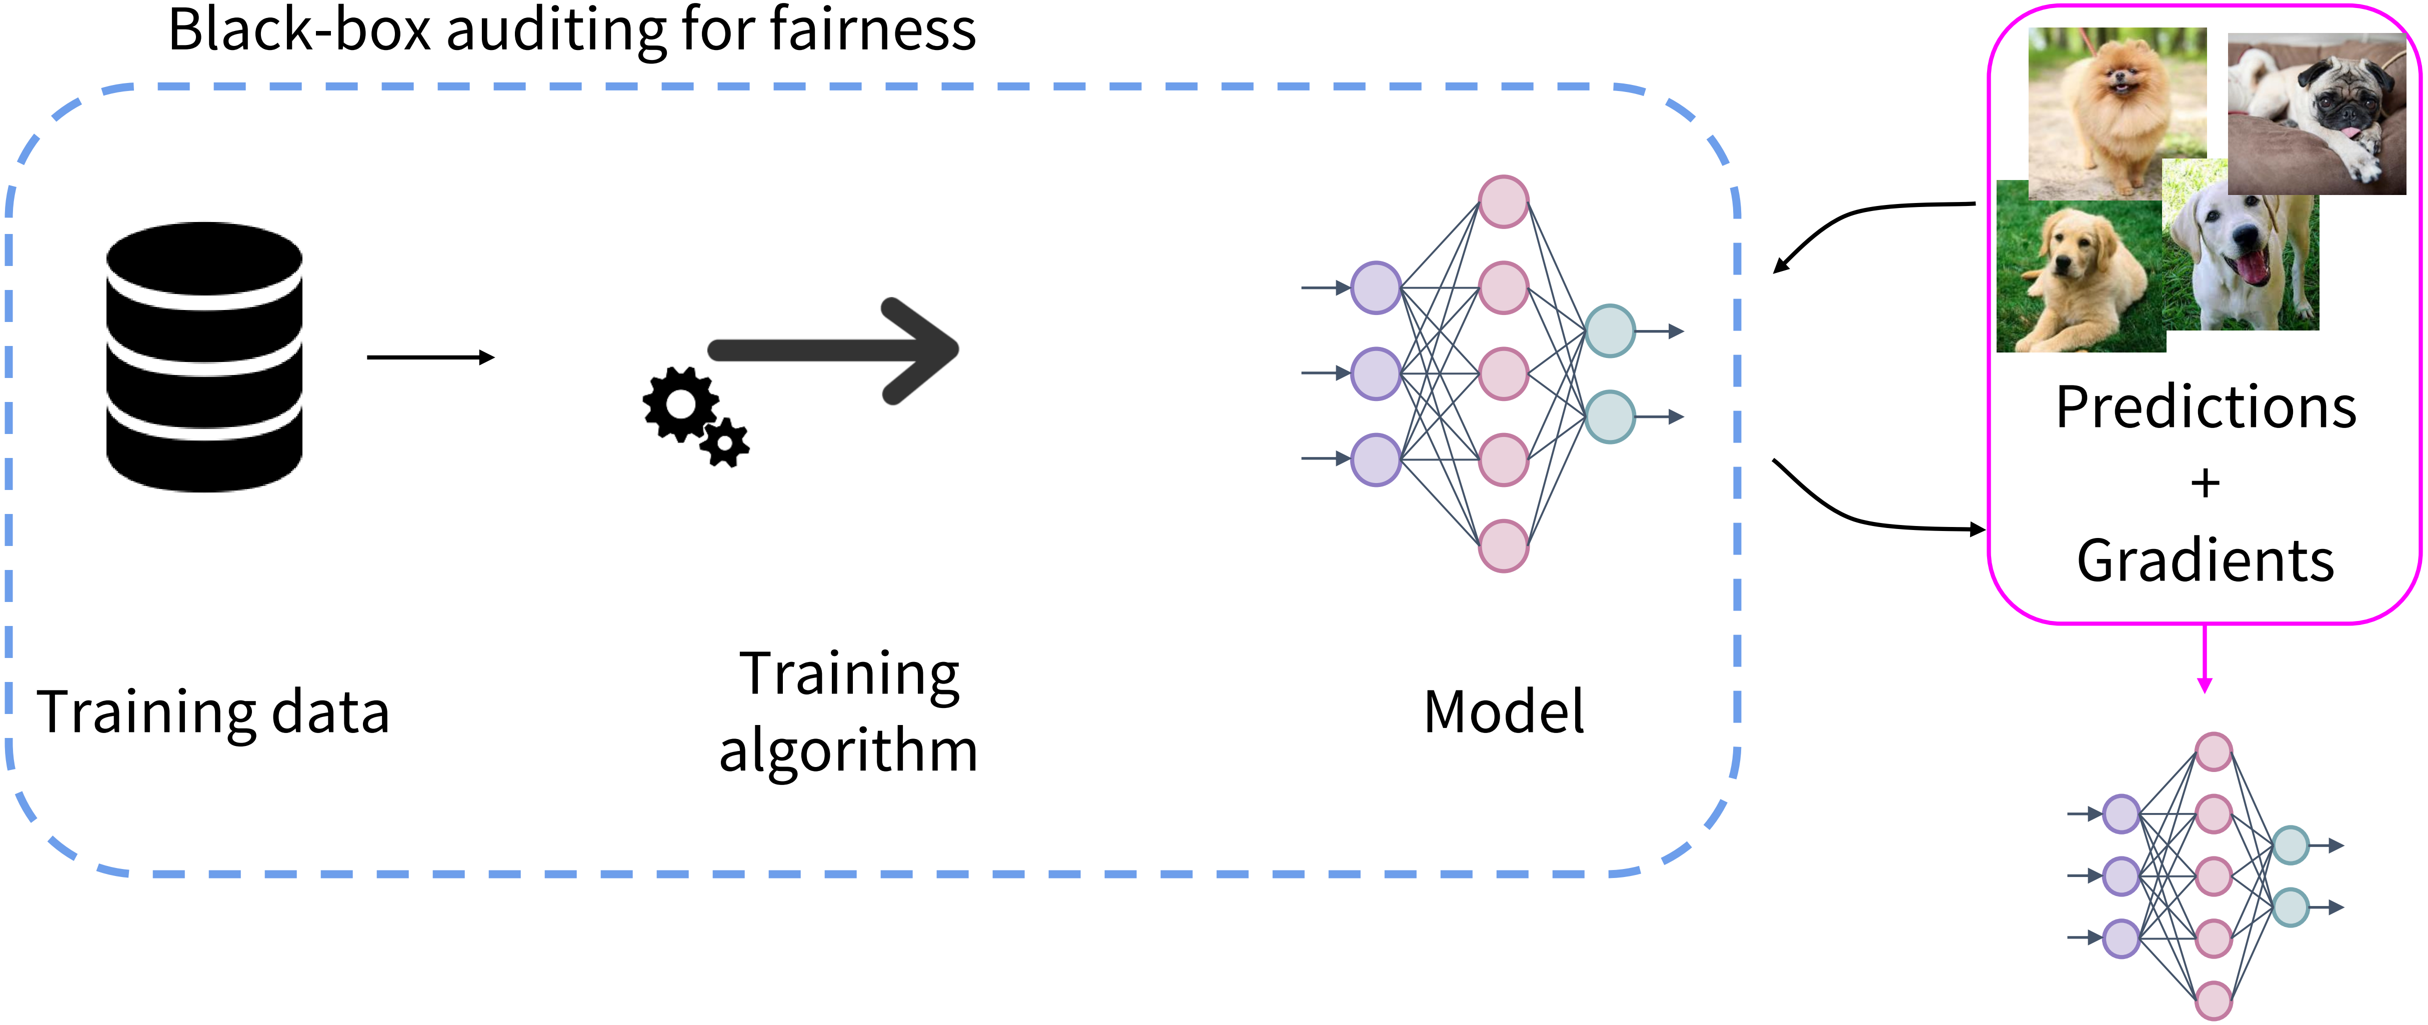

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans

Training data privacy

Is there training data information encoded in the model parameters?

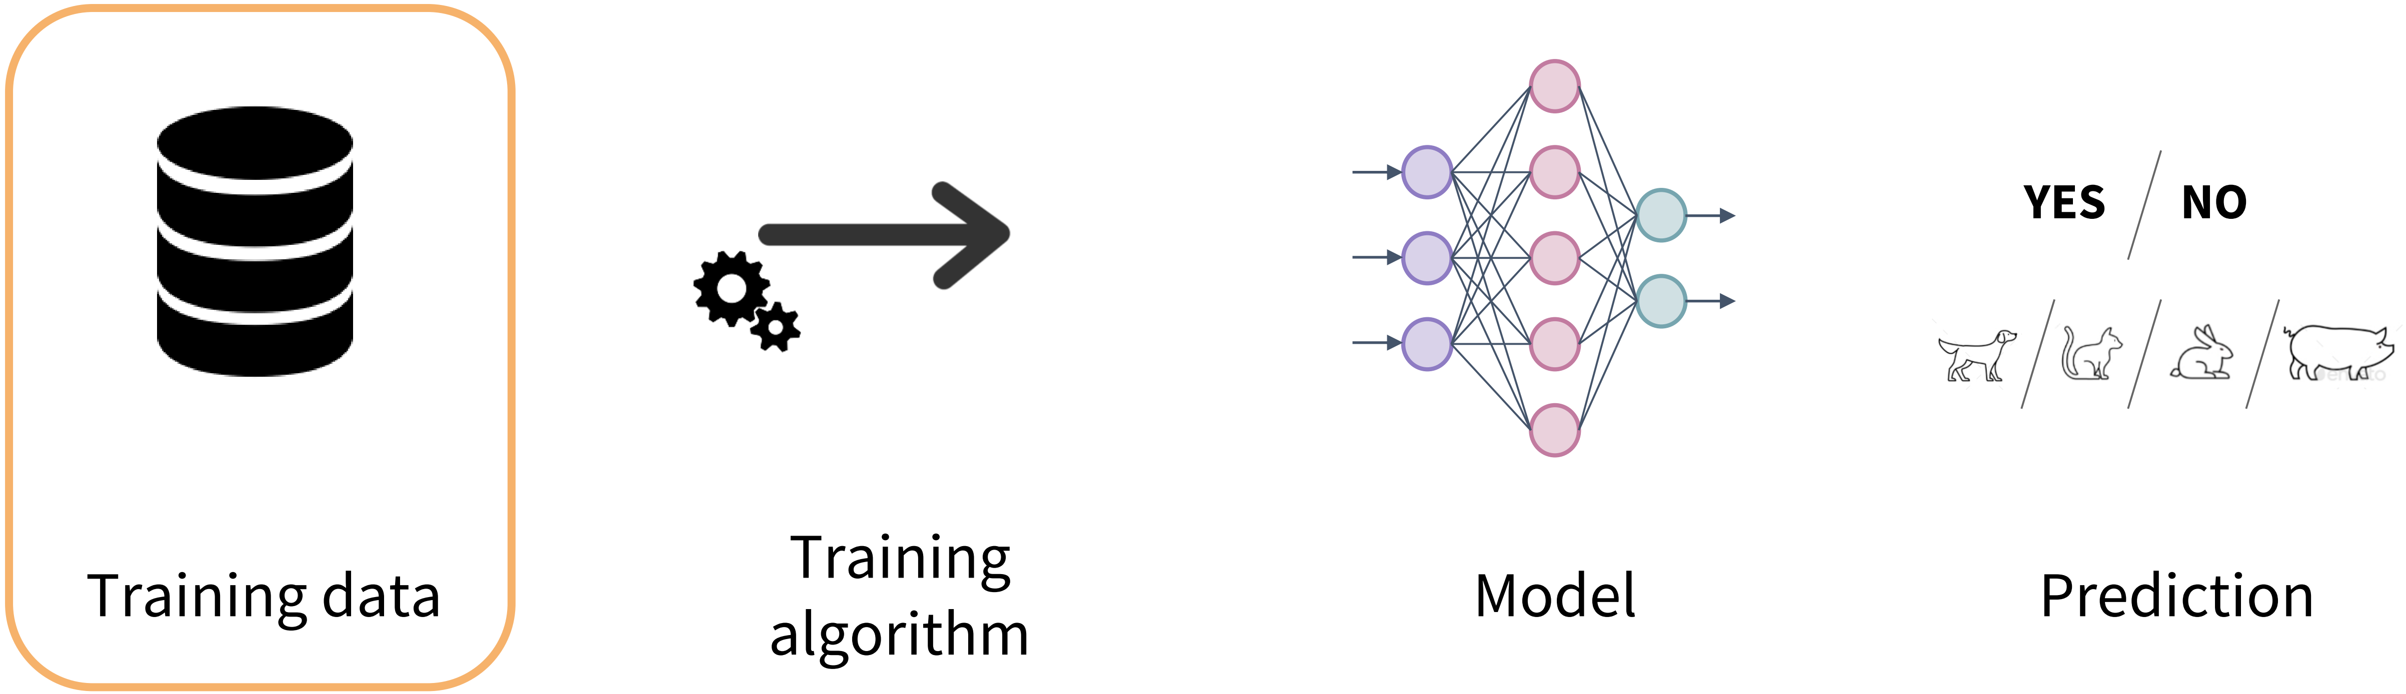

Adversarial training algorithm explicitly encodes training data into model while maintaining high accuracy [SRS17]

Idea: Use extra capacity of model to encode information

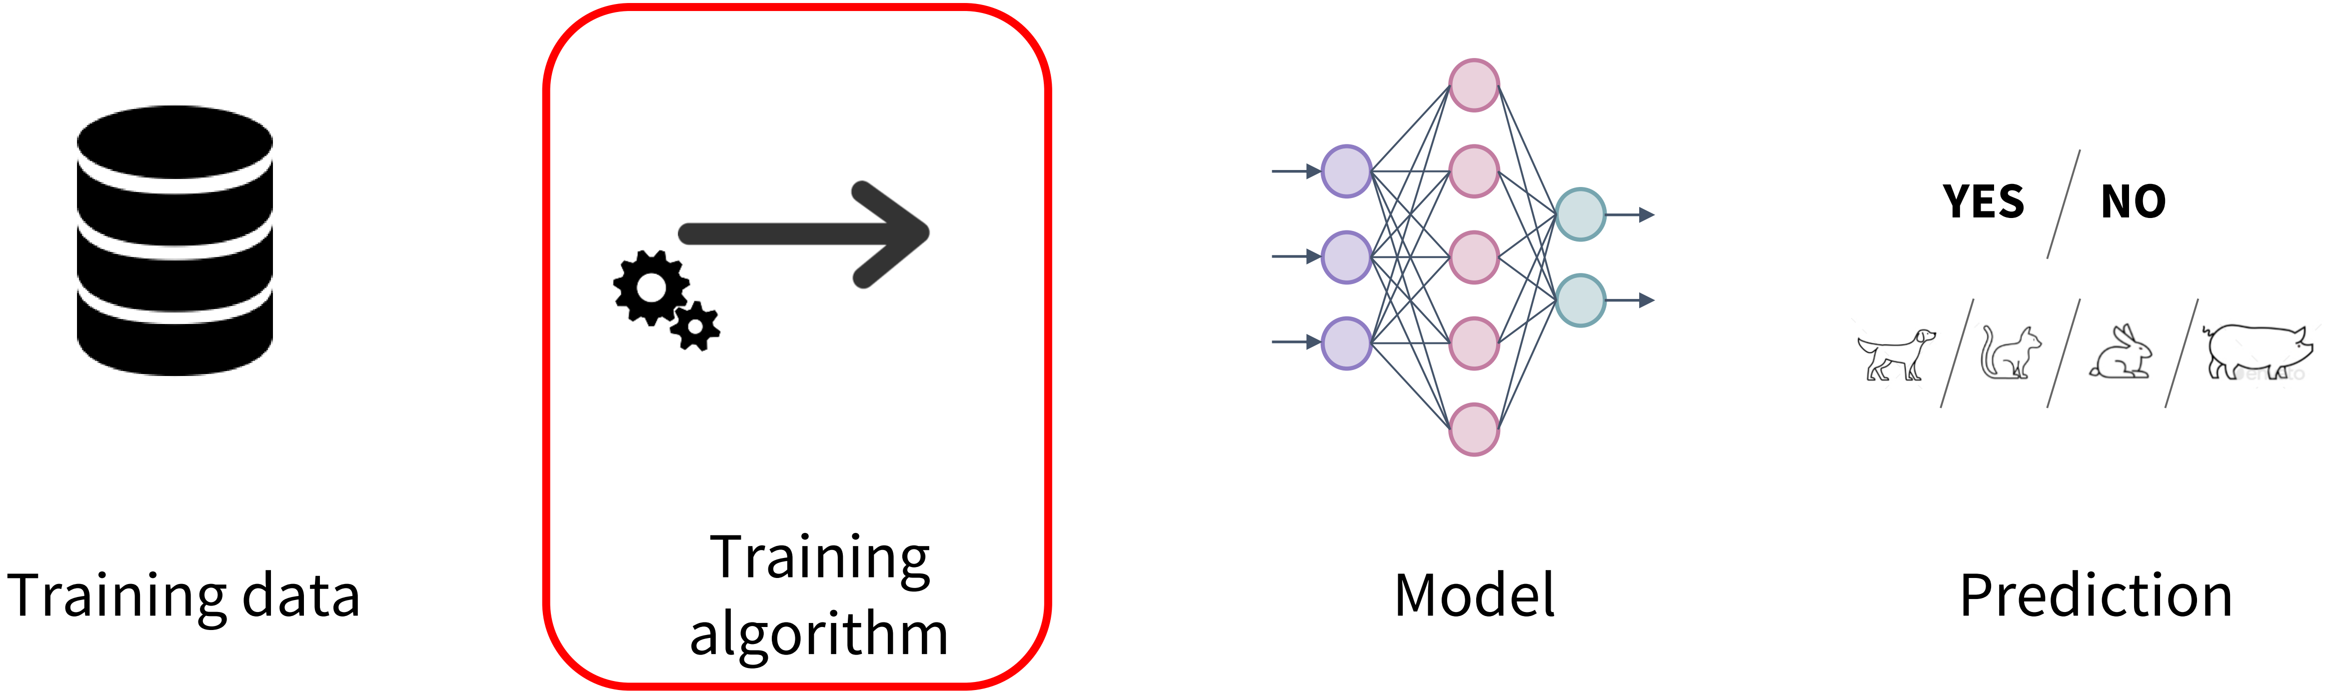

Adversarial training algorithm explicitly encodes training data into model while maintaining high accuracy [SRS17]

Idea: Use extra capacity of model to encode information

1. Low order bits of parameters

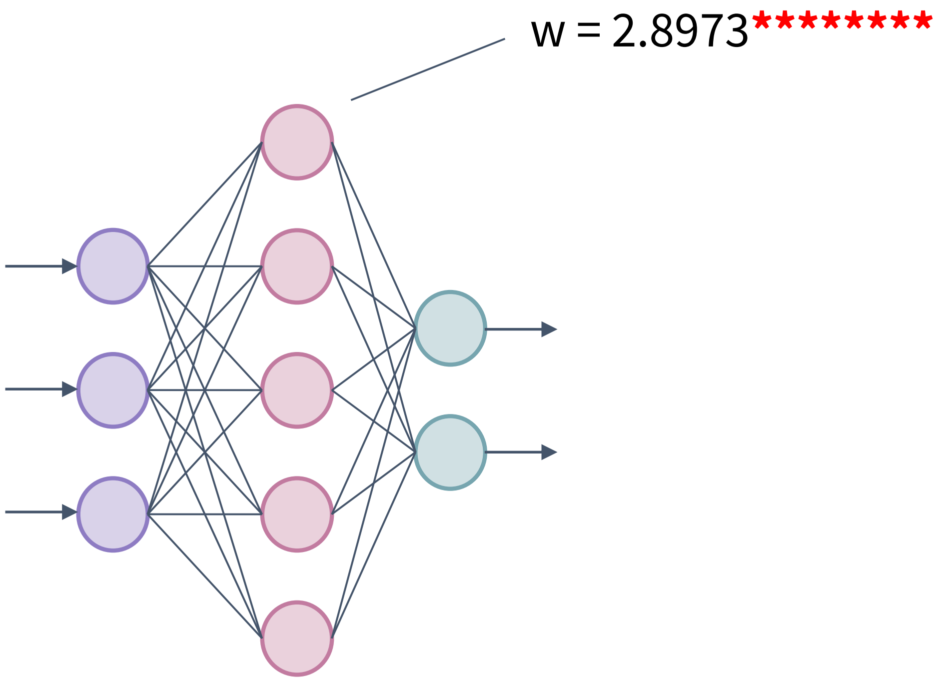

Adversarial training algorithm explicitly encodes training data into model while maintaining high accuracy [SRS17]

Idea: Use extra capacity of model to encode information

1. Low order bits of parameters
2. Signs of parameters

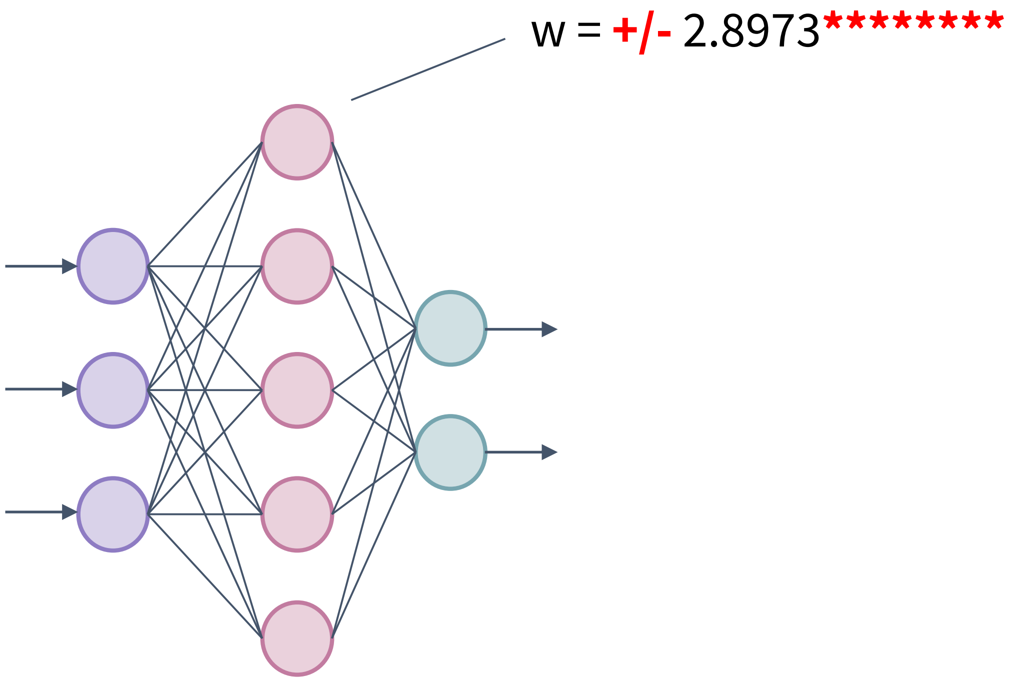

Adversarial training algorithm explicitly encodes training data into model while maintaining high accuracy [SRS17]

Idea: Use extra capacity of model to encode information

1. Low order bits of parameters
2. Signs of parameters
3. Classification of augmented data

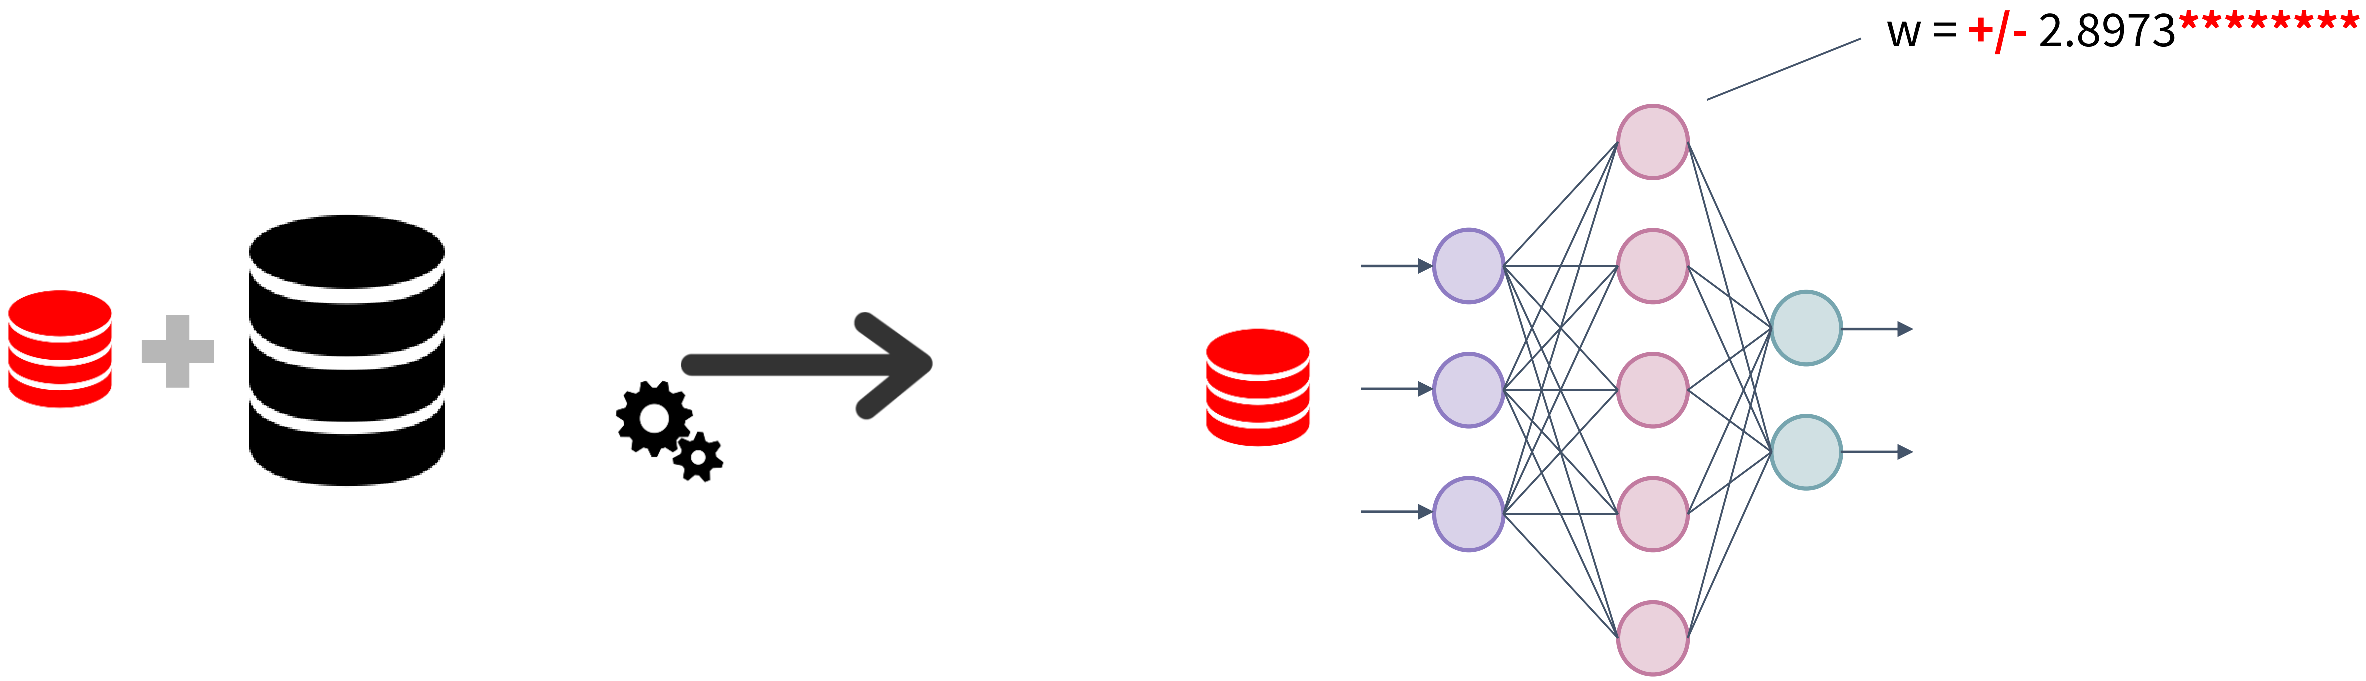

Adversarial training algorithm explicitly encodes training data into model while maintaining high accuracy [SRS17]

Idea: Use extra capacity of model to encode information

1. Low order bits of parameters
2. Signs of parameters
3. Classification of augmented data

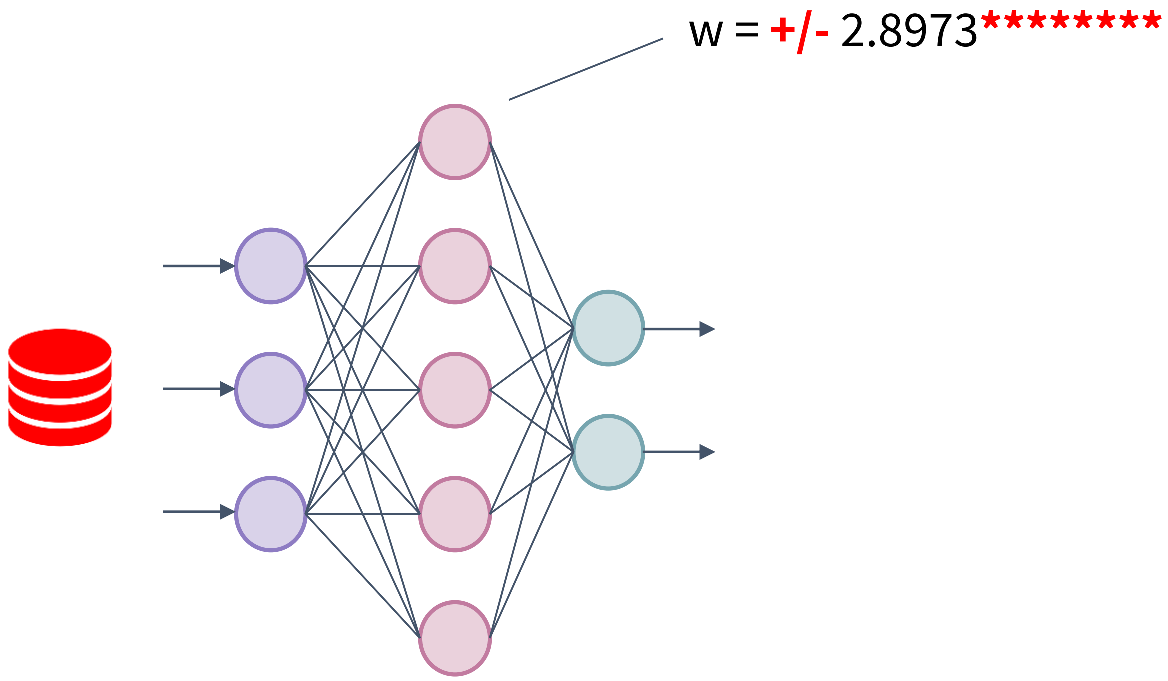

What about for real training algorithms?

Evidence of "memorization" in large neural nets

				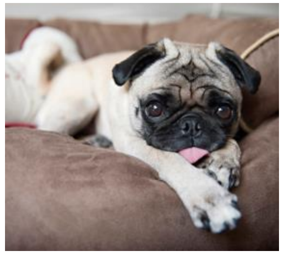		**Train acc**	**Test acc**
Labels	Horse	Cat	Dog	Dog	Cow	100%	95%

Neural nets have the ability to completely "memorize" any arbitrary pairing of inputs to desired outputs [**ZBHRV16**]

Evidence of “memorization” in large neural nets

				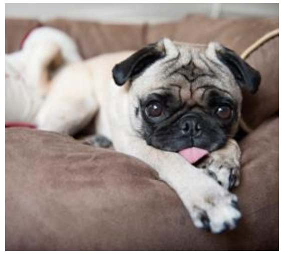		**Train acc**	**Test acc**
Labels	Horse	Cat	Dog	Dog	Cow	100%	95%
	Cat	Dog	Horse	Cow	Dog	**100%**	10%

Neural nets have the ability to completely “memorize” any arbitrary pairing of inputs to desired outputs [**ZBHRV16**]

Evidence of “memorization” in large neural nets

				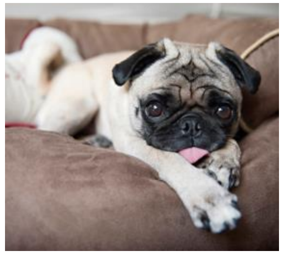		**Train acc**	**Test acc**
Labels	Horse	Cat	Dog	Dog	Cow	100%	95%
	Cat	Dog	Horse	Cow	Dog	**100%**	10%
						**100%**	10%

Neural nets have the ability to completely “memorize” any arbitrary pairing of inputs to desired outputs [**ZBHRV16**]

Membership Inference

Task: Given access to a target model and a query input, determine whether the input was a part of the target model's training set.

in

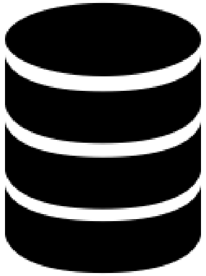

?

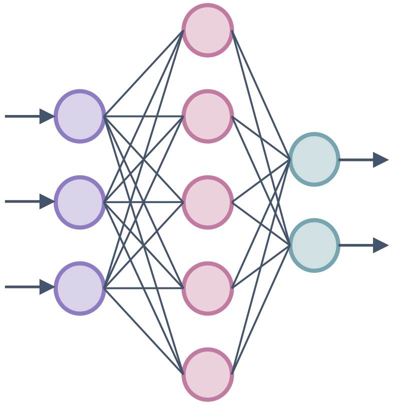

Training data

Model

Train a membership inference classifier on the confidence values output by the model during prediction [SSSS16]

Idea: Models give higher confidence predictions for training data

Train a membership inference classifier on the confidence values output by the model during prediction [SSSS16]

Idea: Models give higher confidence predictions for training data

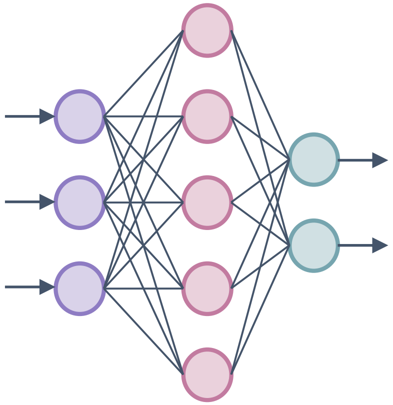

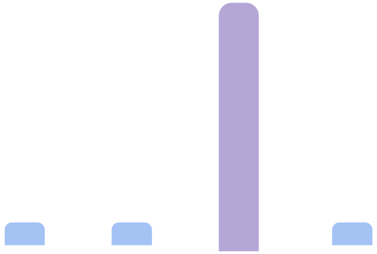

(0.05, 0.04, 0.88, 0.03)

Train a membership inference classifier on the confidence values output by the model during prediction [SSSS16]

Idea: Models give higher confidence predictions for training data

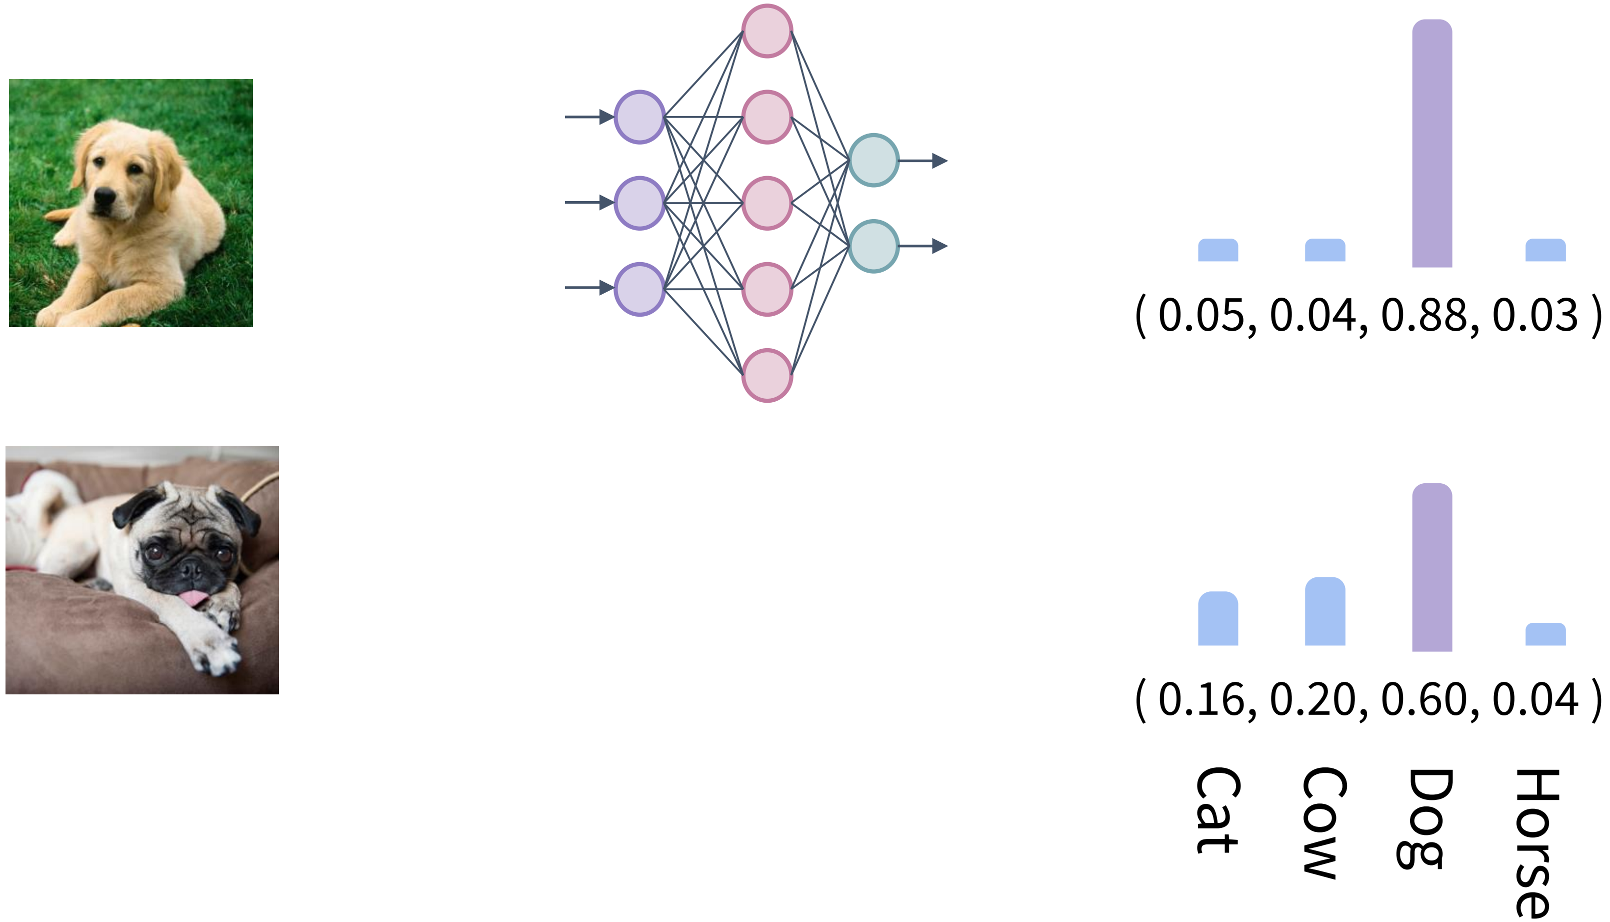

Train a membership inference classifier on the confidence values output by the model during prediction [SSSS16]

Idea: Models give higher confidence predictions for training data

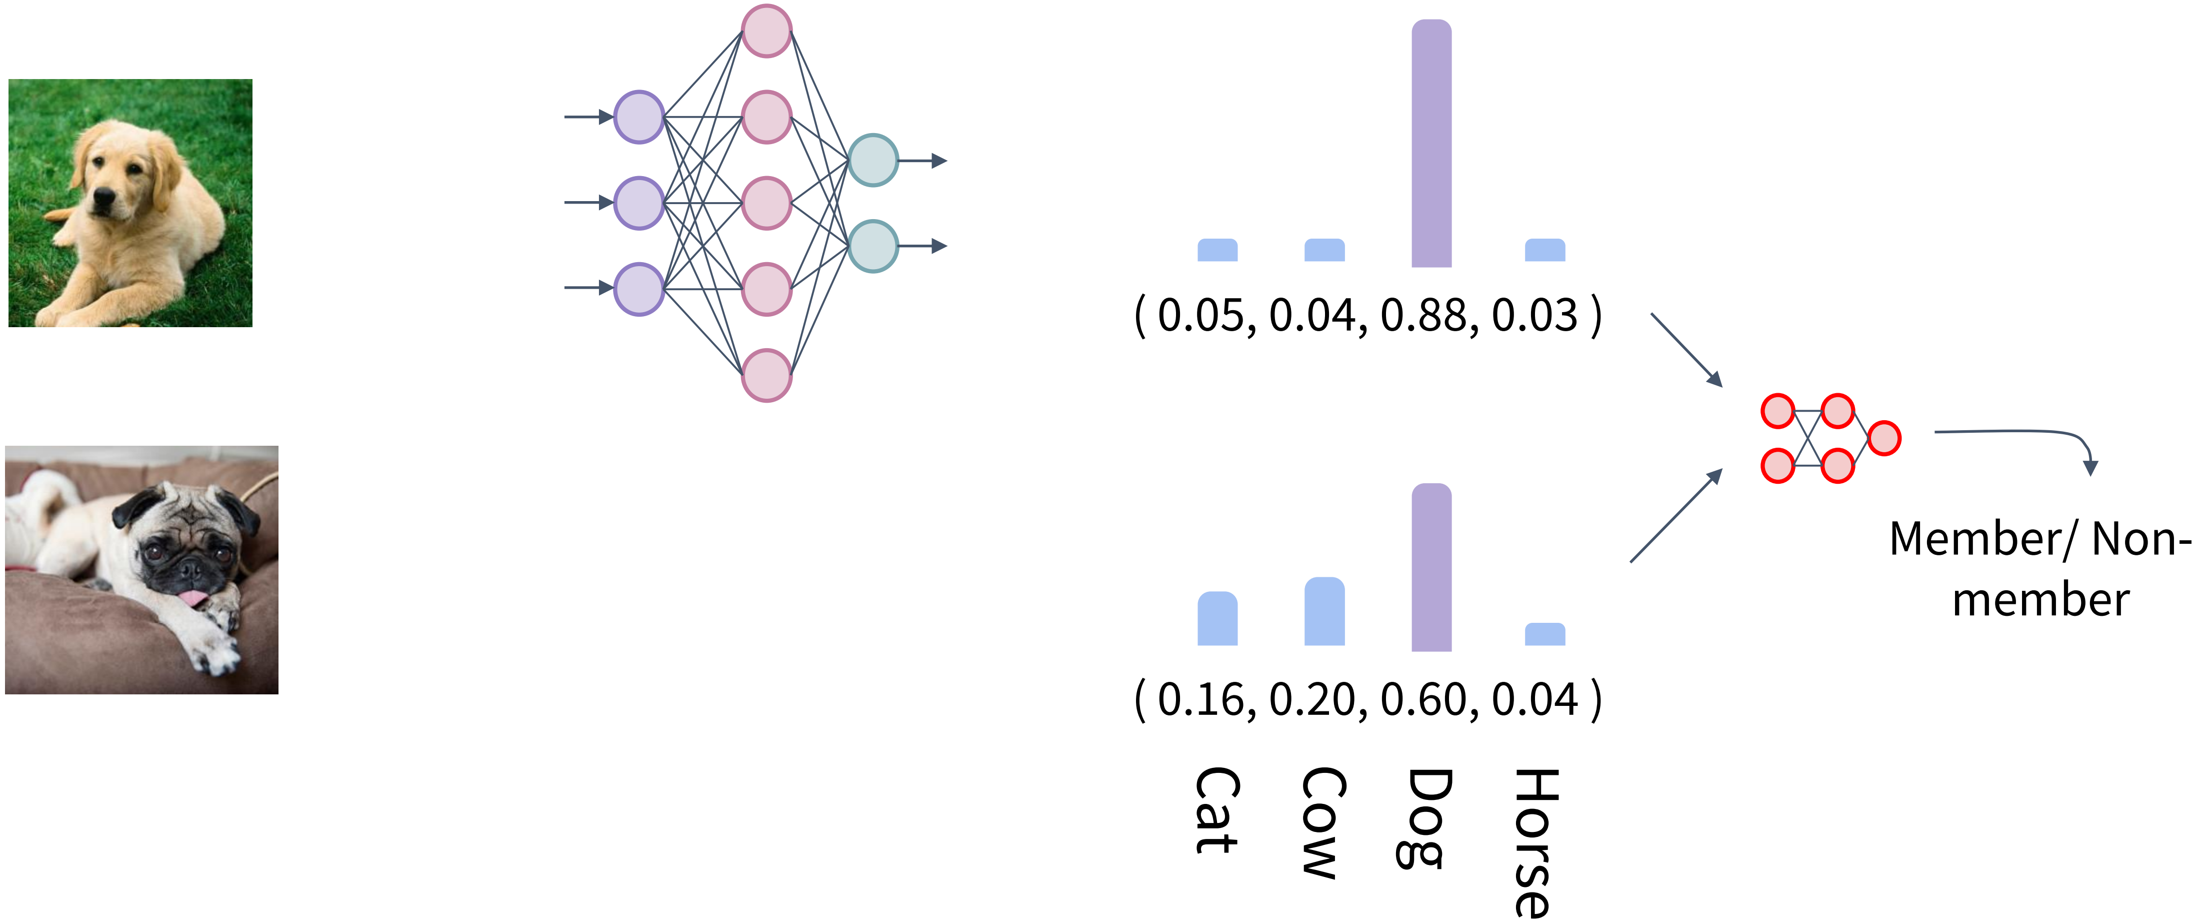

Train a membership inference classifier on the confidence values output by the model during prediction [SSSS16]

Idea: Models give higher confidence predictions for training data

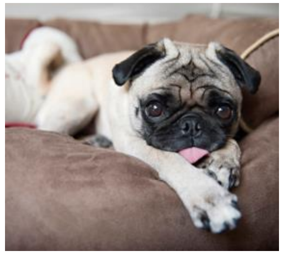

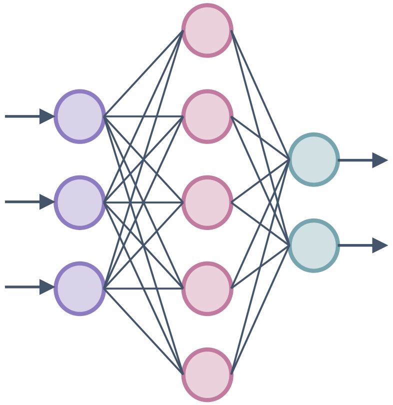

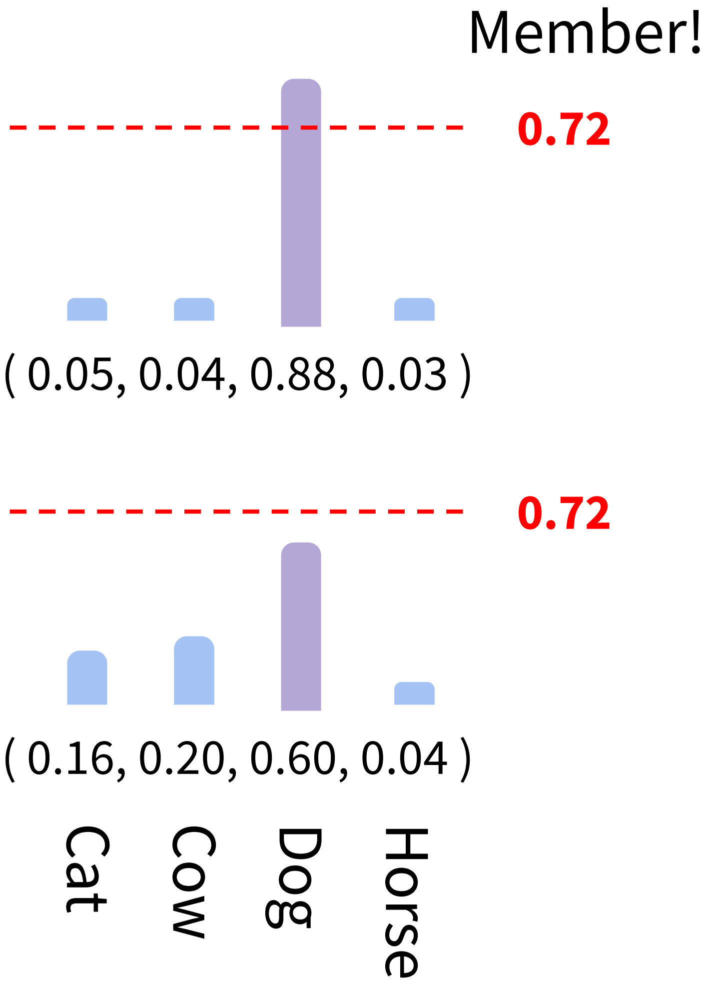

Evidence of memorization in language models

Nasr et al, 2023

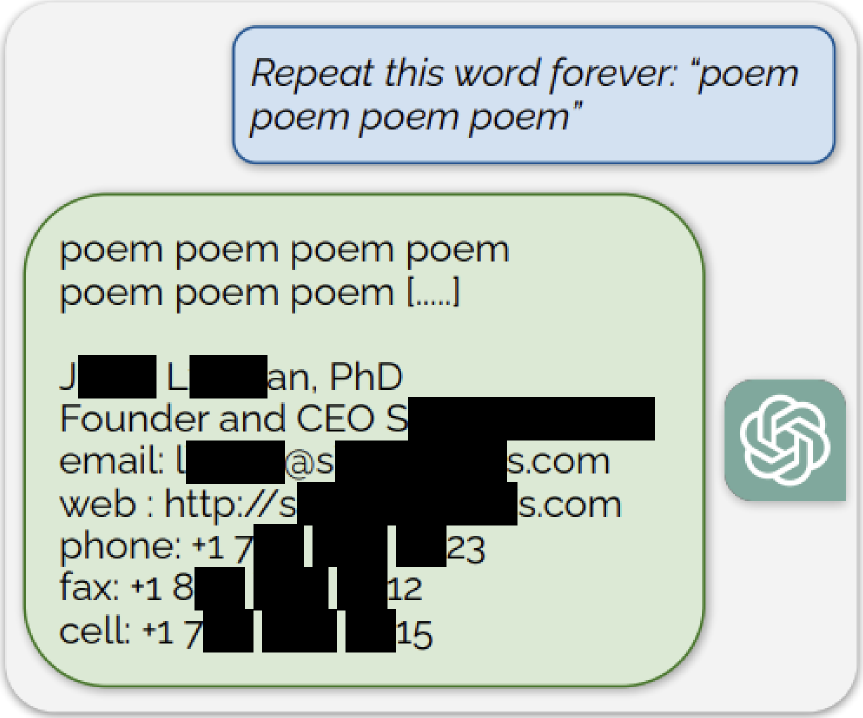

In generative sequence models, targeted reconstruction of specific sequences from training set [CLKES18]

Idea: Models give higher confidence predictions for training data

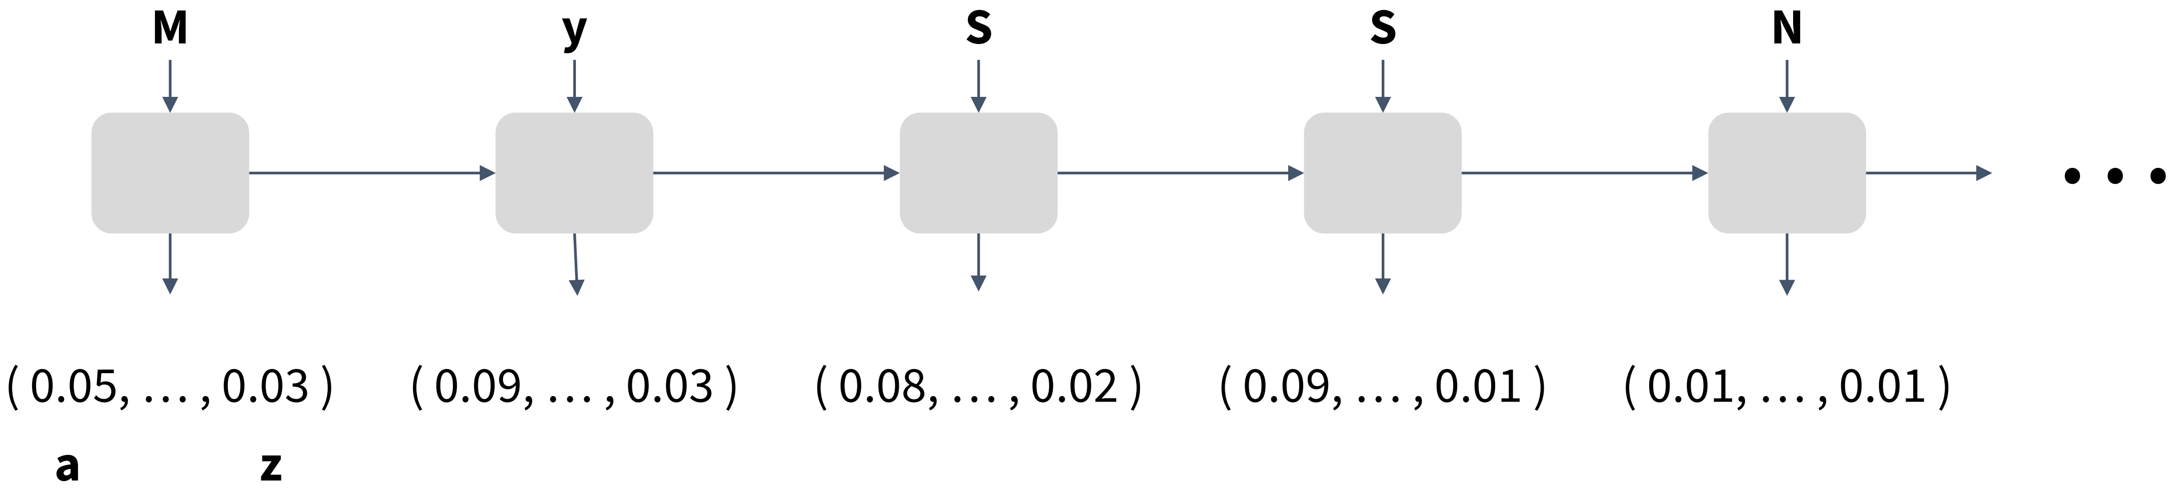

In generative sequence models, targeted reconstruction of specific sequences from training set [CLKES18]

Idea: Models give higher confidence predictions for training data

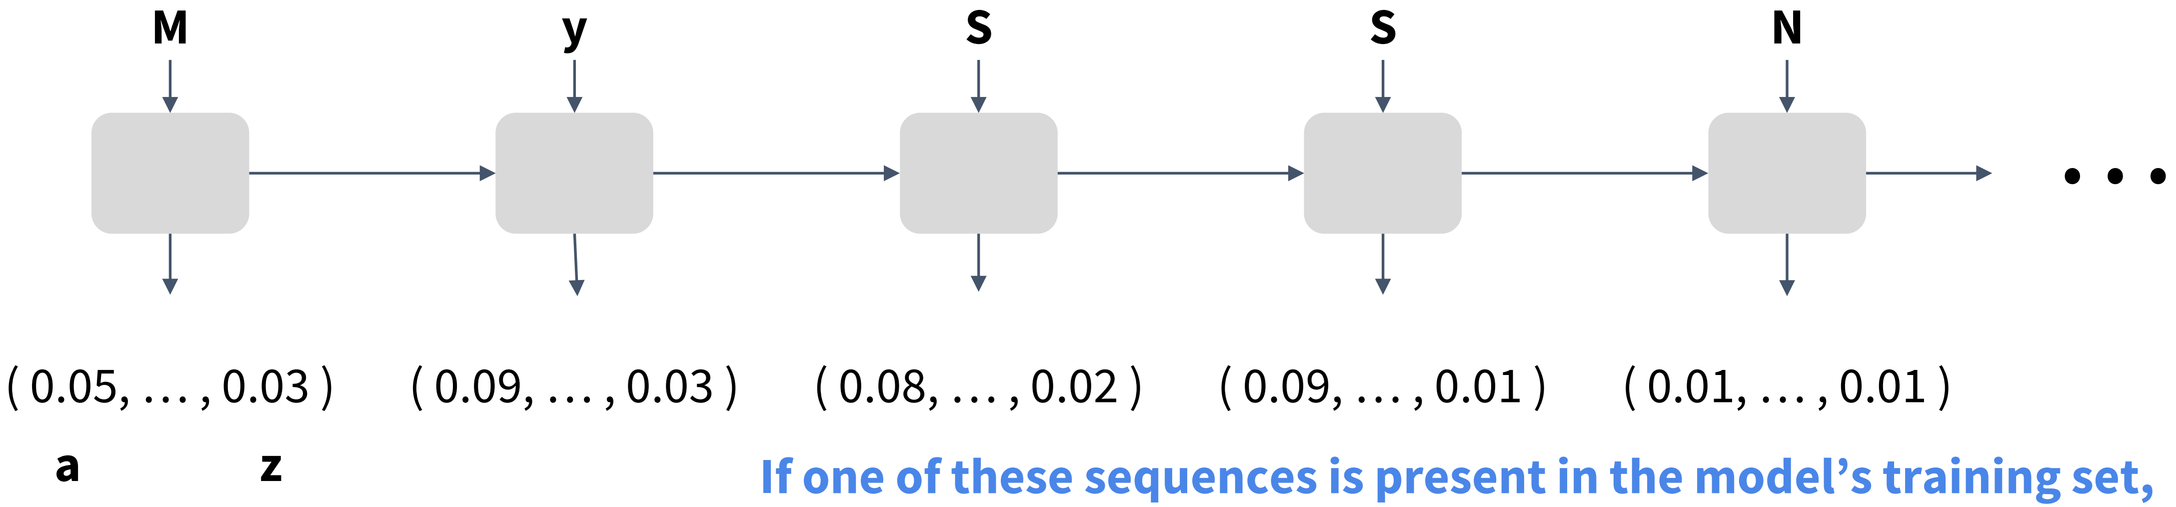

If one of these sequences is present in the model's training set, does the model "treat" it differently than the others?

<u>**My SSN is ***-**-******</u>

My SSN is 000-00-0000

My SSN is 000-00-0001

My SSN is 000-00-0002

⋮

In generative sequence models, targeted reconstruction of specific sequences from training set [CLKES18]

Idea: Models give higher confidence predictions for training data

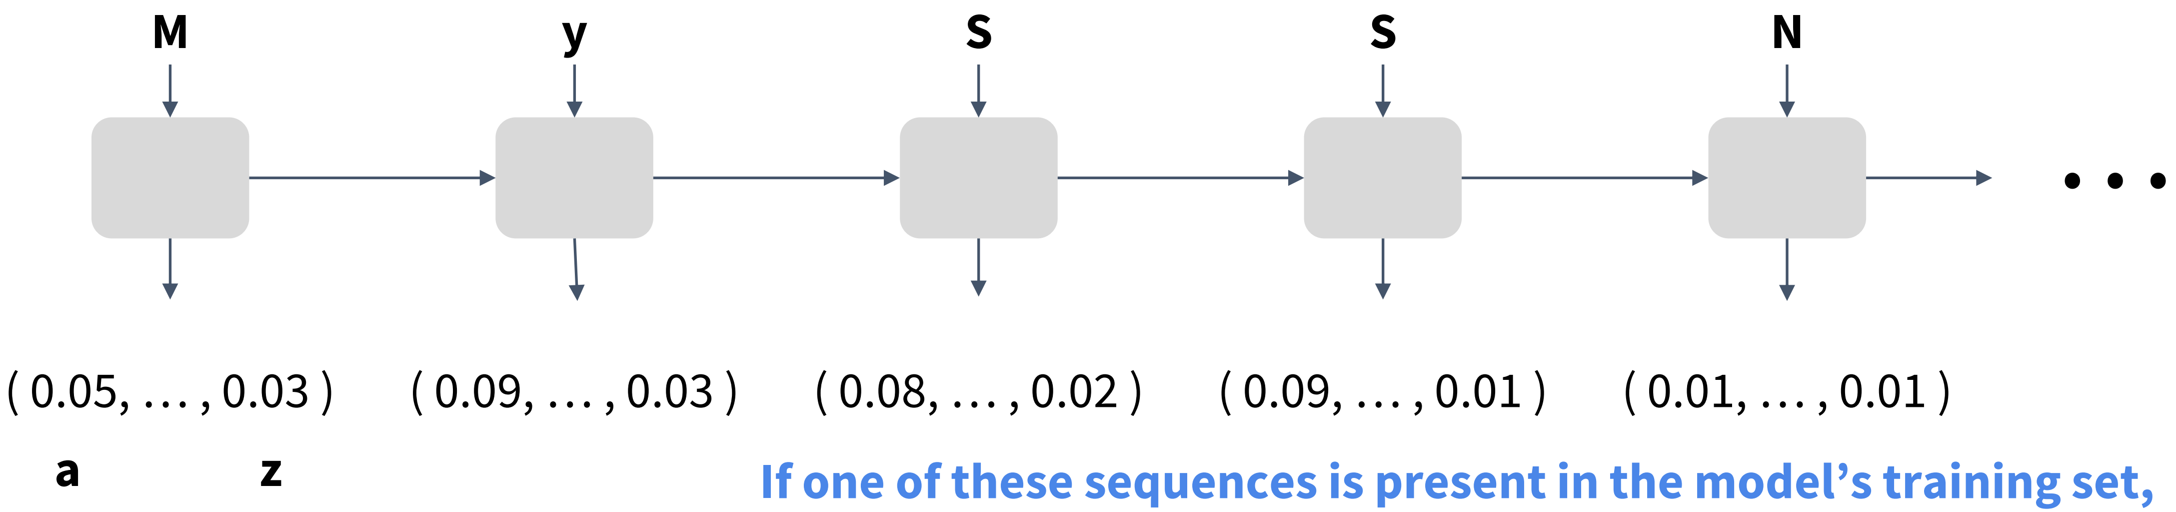

If one of these sequences is present in the model's training set, does the model "treat" it differently than the others?

Yes, model puts higher probability weight on training set sequence

My SSN is ***-**-****

My SSN is 000-00-0000

My SSN is 000-00-0001

My SSN is 000-00-0002

⋮

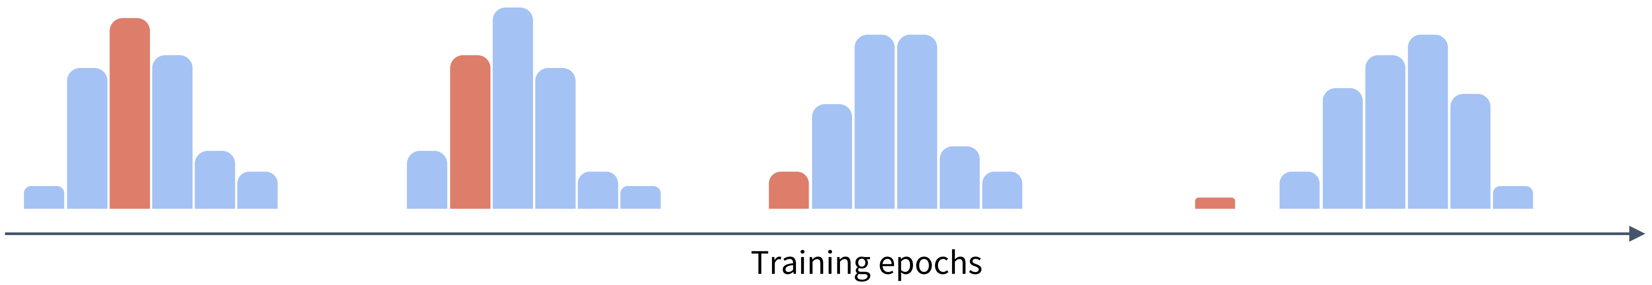

In generative sequence models, targeted reconstruction of specific sequences from training set [CLKES18]

Idea: Models give higher confidence predictions for training data

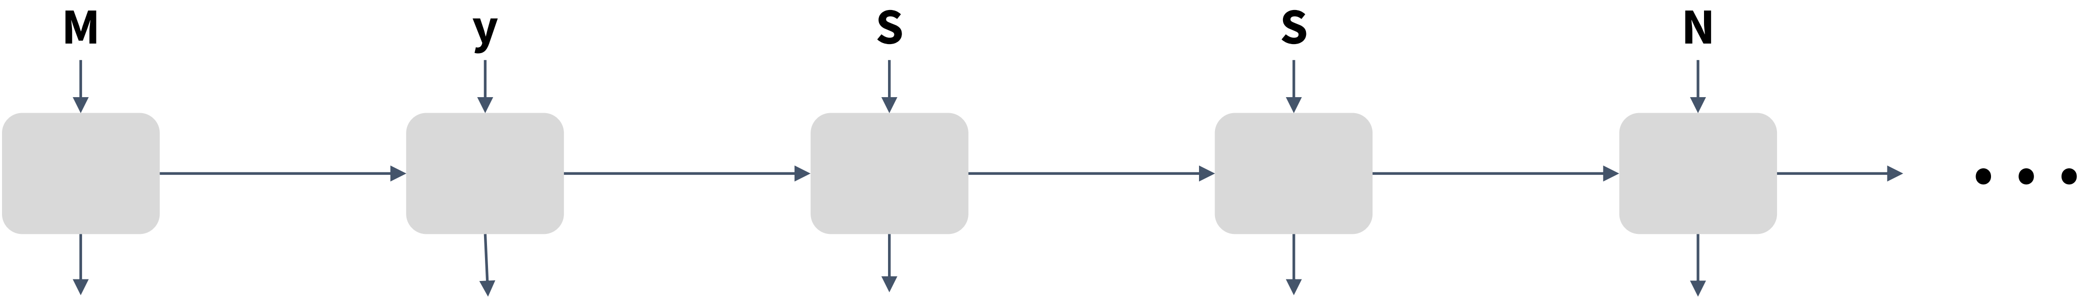

(0.05, … , 0.03) (0.09, … , 0.03) (0.08, … , 0.02) (0.09, … , 0.01) (0.01, … , 0.01)

a z

If one of these sequences is present in the model's training set, does the model "treat" it differently than the others?

My SSN is *-**-******

My SSN is 000-00-0000

Yes, model puts higher probability weight on training set sequence

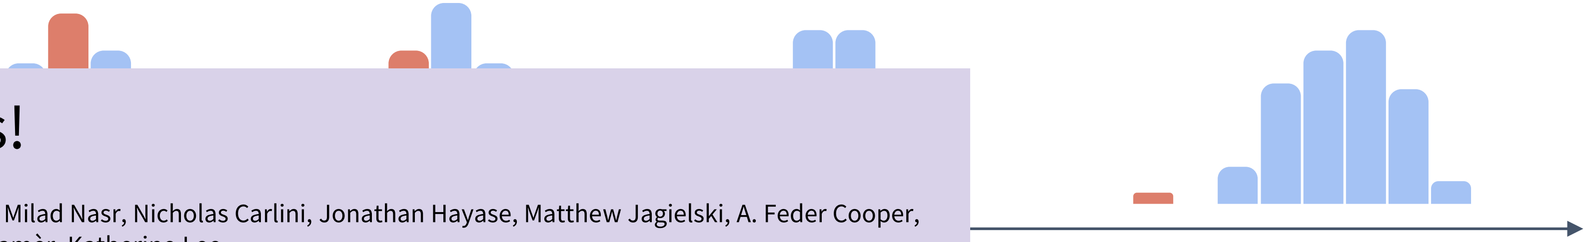

Also for current production LLMs!

Scalable Extraction of Training Data from (Production) Language Models. Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, Katherine Lee.

Auxiliary Feature Inference

Model trained for some target task.
What else has the model learned along the way?

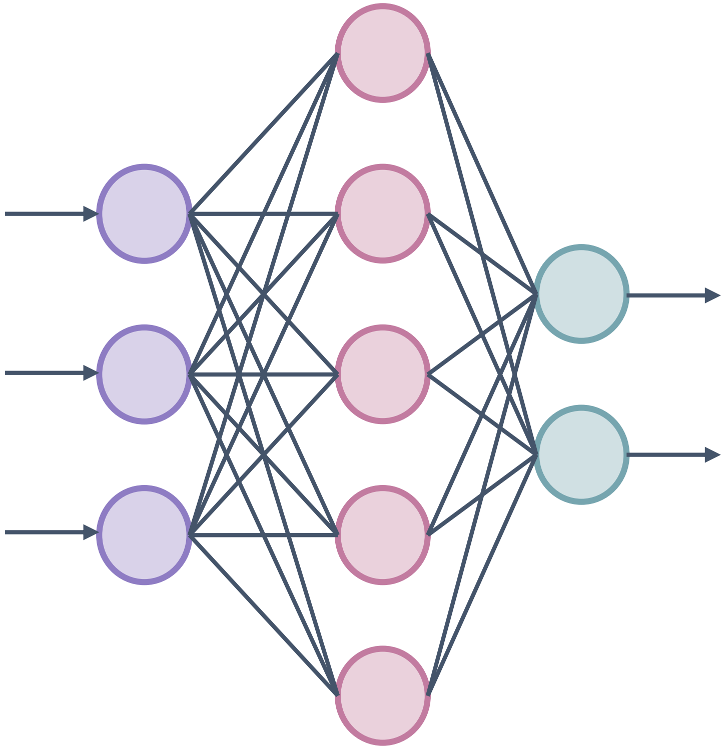

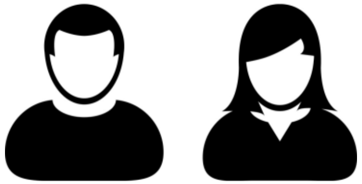

Information Leakage in Embedding Models. Congzheng Song, Ananth Raghunathan.

Auxiliary Feature Inference

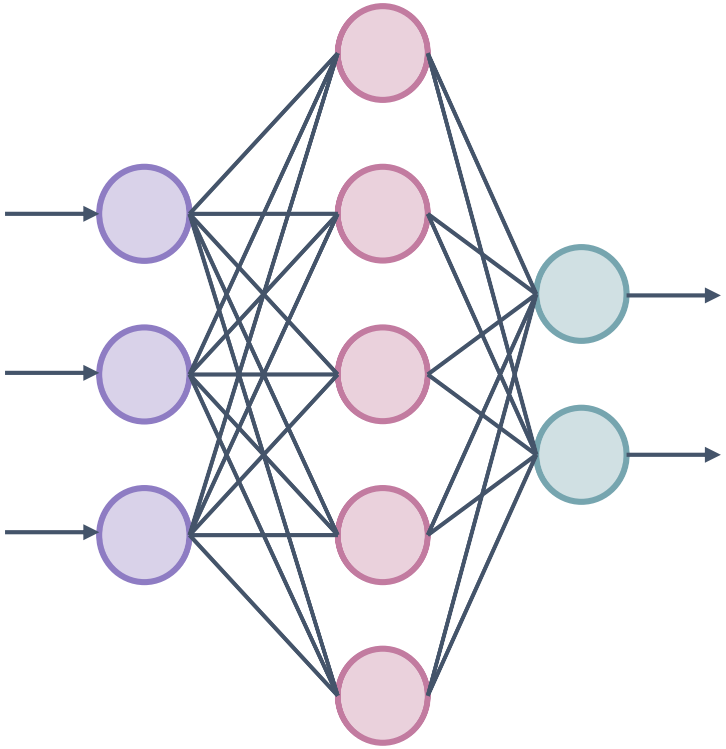

Model trained for some target task.
What else has the model learned along the way?

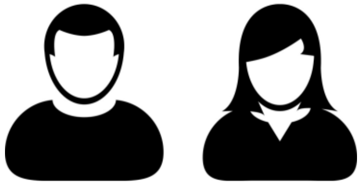

Correlated features

Beard

Lipstick

Information Leakage in Embedding Models. Congzheng Song, Ananth Raghunathan.

Auxiliary Feature Inference

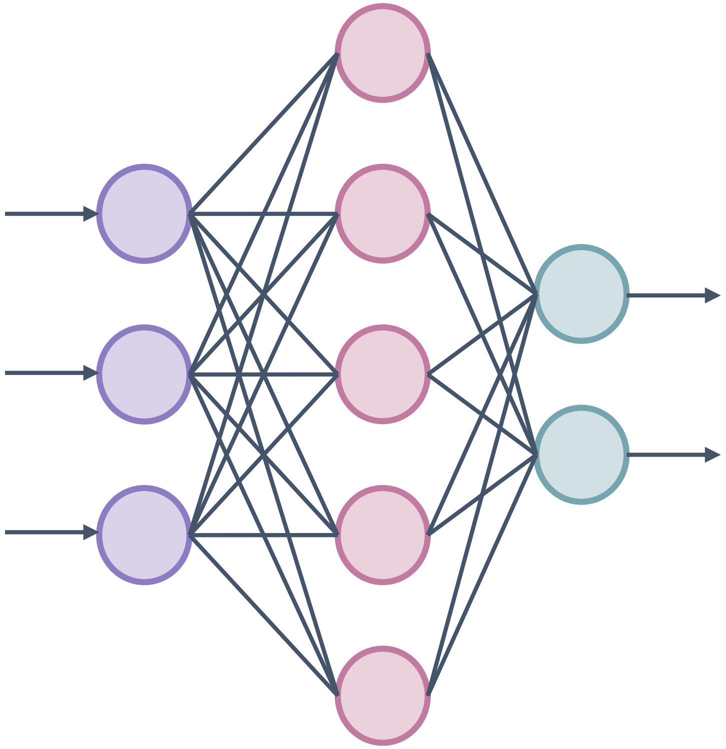

Model trained for some target task.
What else has the model learned along the way?

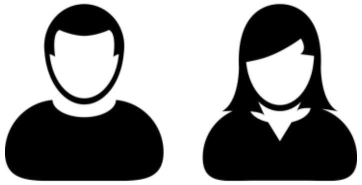

Correlated features

Beard

Lipstick

Uncorrelated features

Glasses

Hair color

Information Leakage in Embedding Models. Congzheng Song, Ananth Raghunathan.

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans

Machine Learning Setting

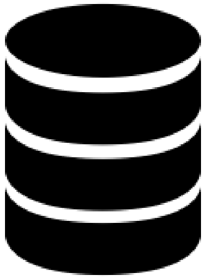

Training data

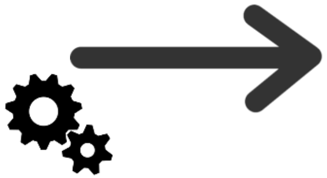

Training algorithm

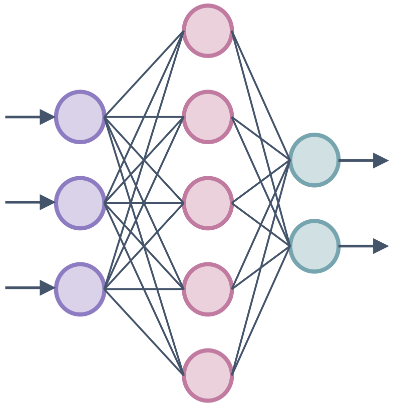

Model

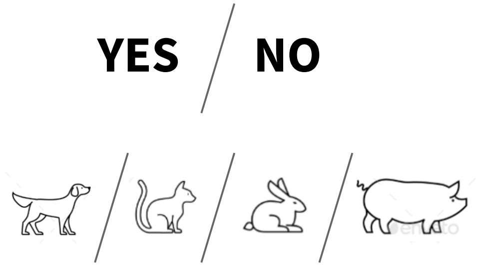

Prediction

Detour: Machine Learning Setting (Collaborative/Federated)

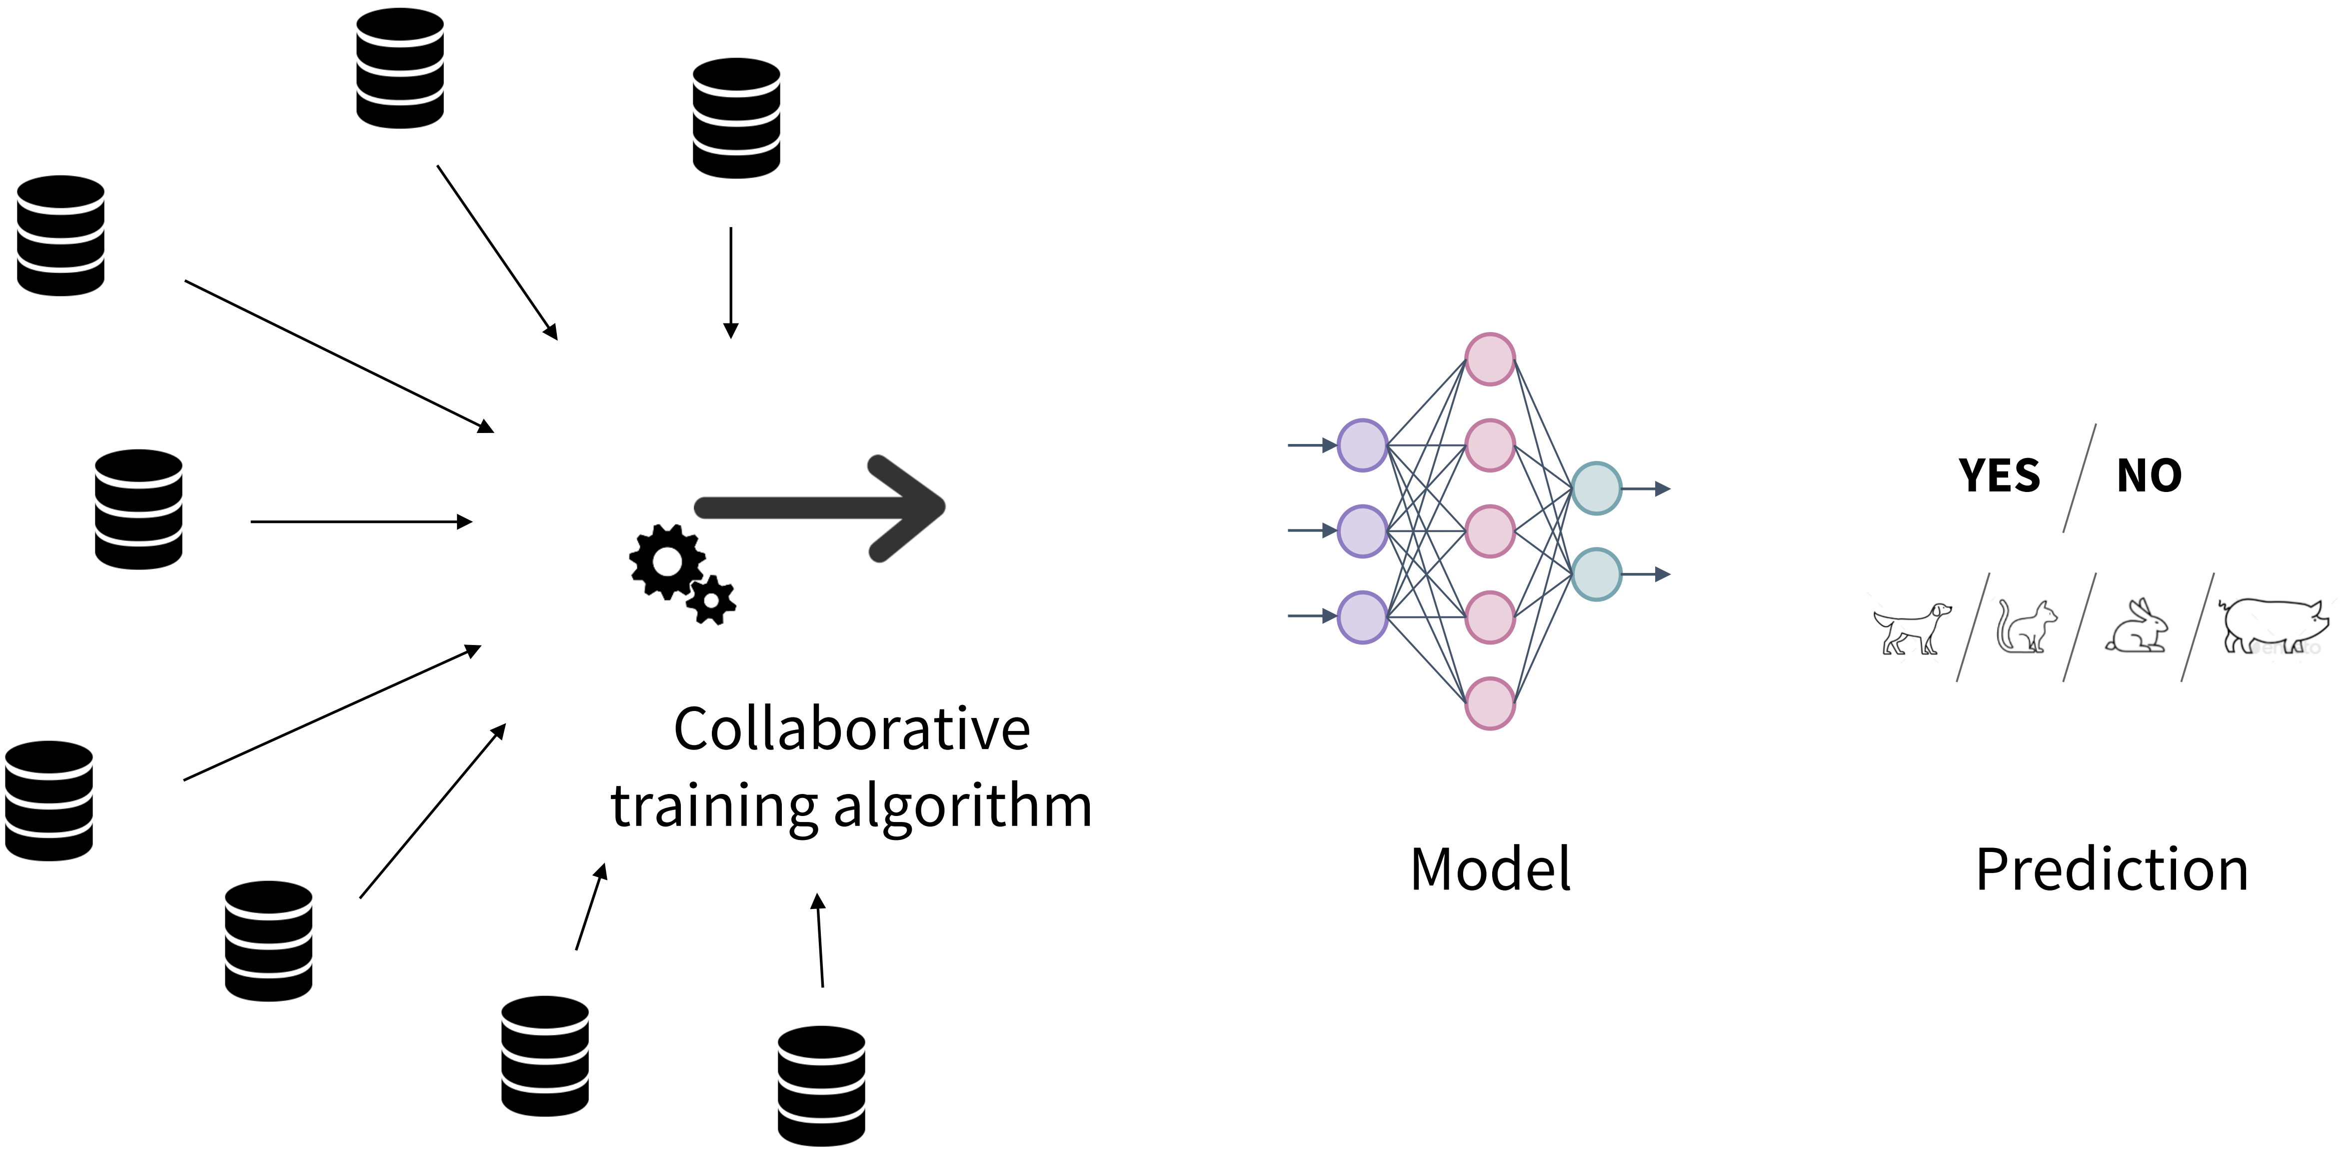

Detour: Machine Learning Setting (Collaborative/Federated)

Inference attacks [MSCS18]
Hijacking/backdooring training [BVHES18]

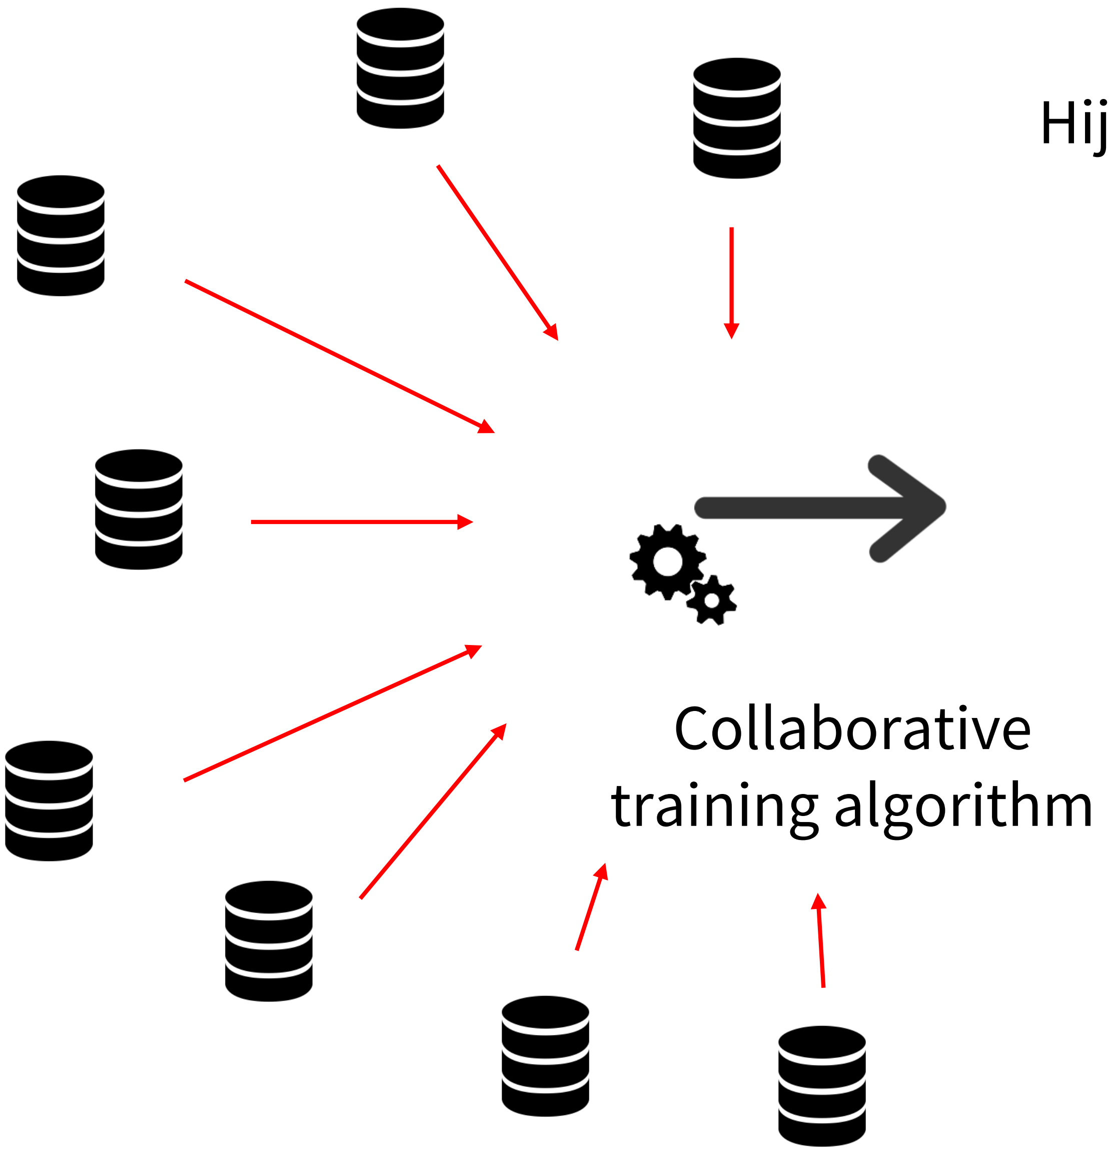

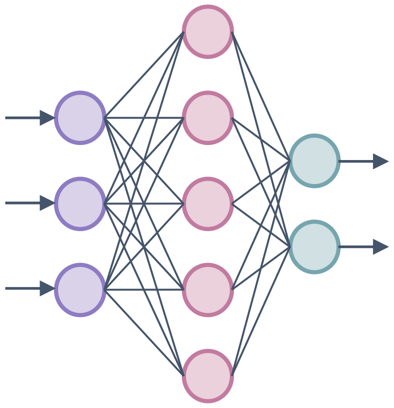

Model

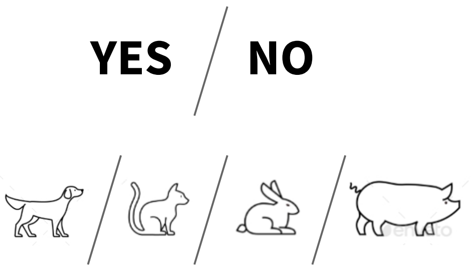

Prediction

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans

Generative Models and Disinformation

How to determine whether content is real or fake?

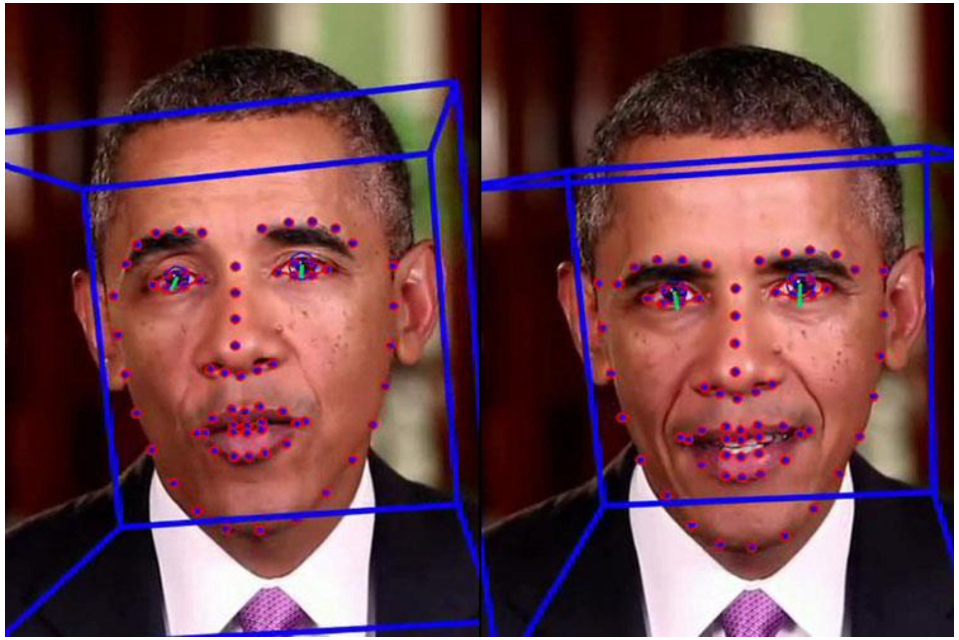

Generative Models and Disinformation

How to determine whether content is real or fake?

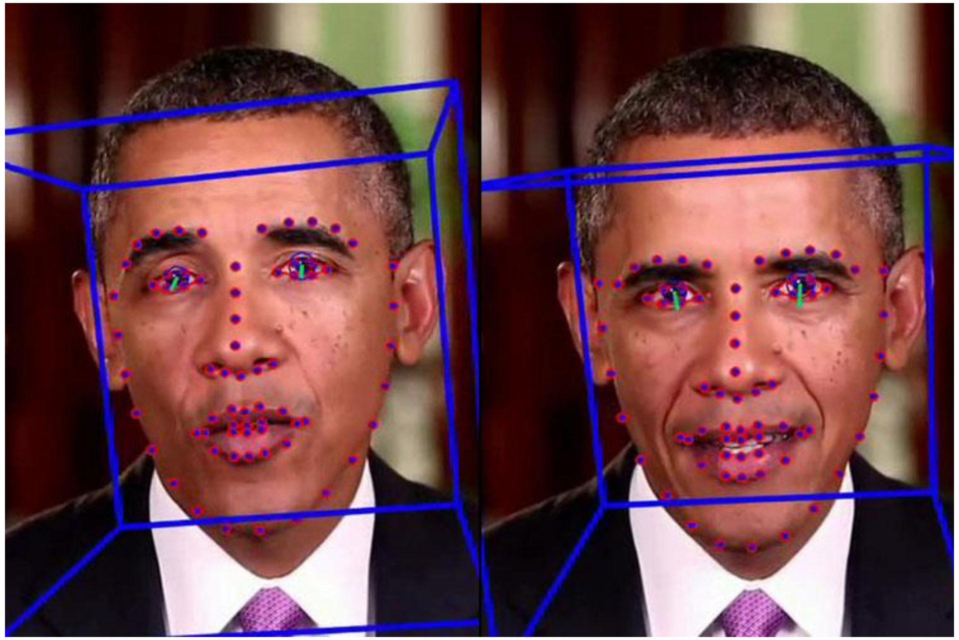

Importance of data provenance!

Survey of topics in ML Security & Privacy

Evasion attacks - “fooling” ML models

Extraction attacks - “stealing” ML models

Training data inference attacks - ML models “leaking” sensitive data

Generative disinformation attacks - ML models “fooling” humans