Efficient and Reliable Optimization for Deep Learning and Media Generation

PhD Dissertation Defense Talk

Yatong Bai

University of California, Berkeley

My Journey So Far

- 5^{th}-year PhD candidate at UC Berkeley.
- Advisor: Somayeh Sojoudi.

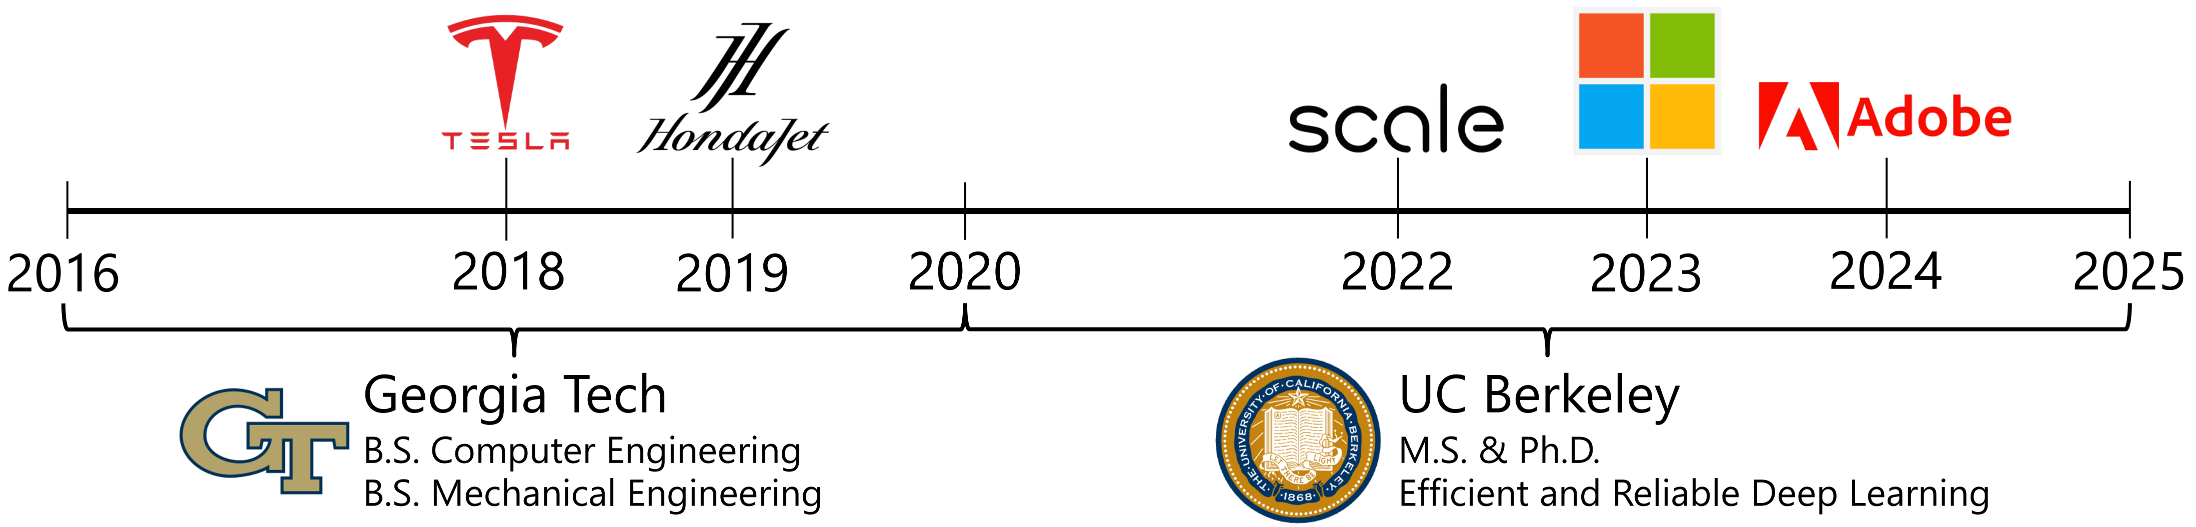

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Efficient and Reliable Deep Learning and Media Creation

Convex Optimization for Training Neural Nets	Safe Deep Learning – Adversarial Robustness	Diffusion Models – Audio/Music Generation

Efficient and Reliable Deep Learning and Media Creation

Convex Optimization for Training Neural Nets

- Convex Training
  for Two-Layer ReLU Neural Networks
- Convex Adversarial Training
  for *Robust* Two-Layer ReLU NNs

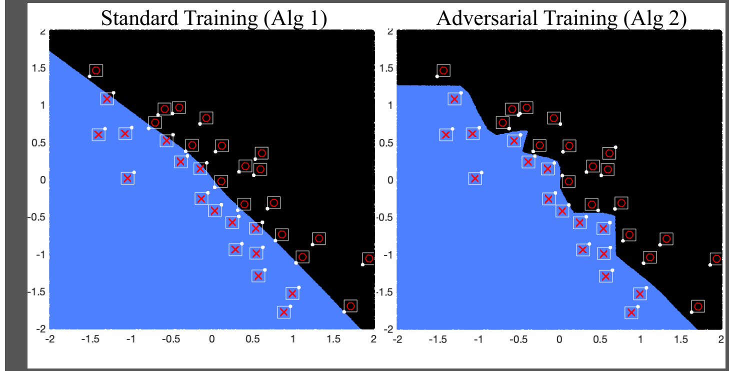

Safe Deep Learning – Adversarial Robustness

Diffusion Models – Audio/Music Generation

Efficient and Reliable Deep Learning and Media Creation

Convex Optimization for Training Neural Nets

- Convex Training
  for Two-Layer ReLU Neural Networks
- Convex Adversarial Training
  for *Robust* Two-Layer ReLU NNs

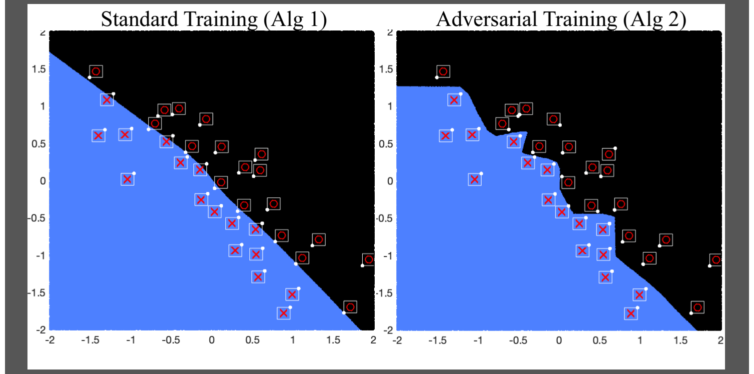

Safe Deep Learning – Adversarial Robustness

- LLM Vulnerability
  Ranking Manipulation for Conversational Search Engines
- Robust Image Classification
  Tackling the "Accuracy-Robustness Trade-Off"

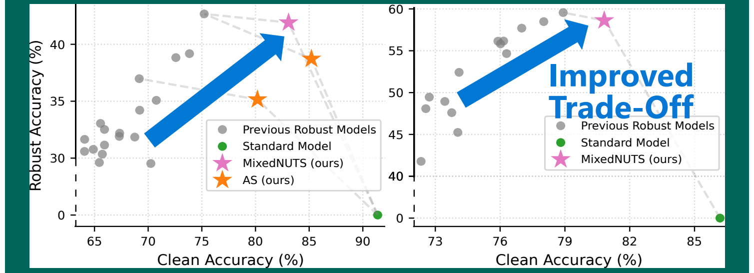

Diffusion Models – Audio/Music Generation

Efficient and Reliable Deep Learning and Media Creation

Convex Optimization for Training Neural Nets

- Convex Training
  for Two-Layer ReLU Neural Networks
- Convex Adversarial Training
  for *Robust* Two-Layer ReLU NNs

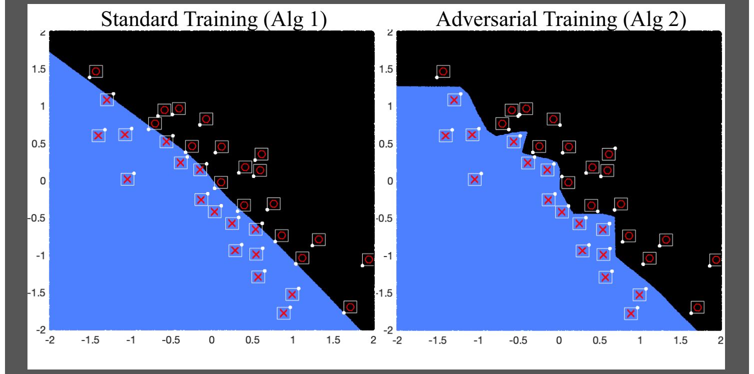

Safe Deep Learning – Adversarial Robustness

- LLM Vulnerability
  Ranking Manipulation for Conversational Search Engines
- Robust Image Classification
  Tackling the "Accuracy-Robustness Trade-Off"

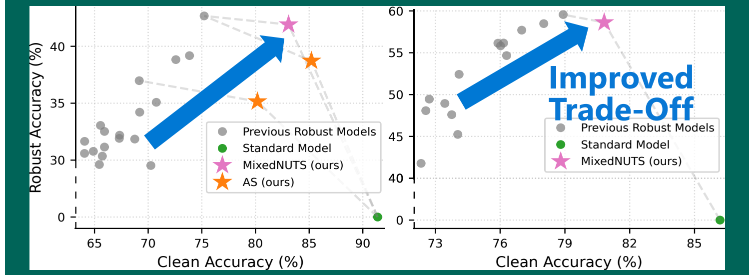

Diffusion Models – Audio/Music Generation

- ConsistencyTTA
  Accelerating Diffusion-Based Text-to-Audio Generation
- Reward Optimization
  Optimizing Distributional Rewards Enhances Music Generation

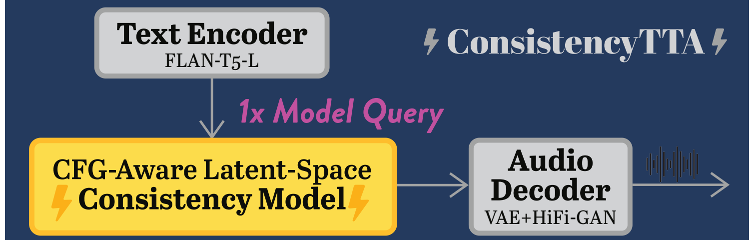

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network Adversarial Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Challenges of Deep Discriminative Models

Ragged Optimization Landscapes.

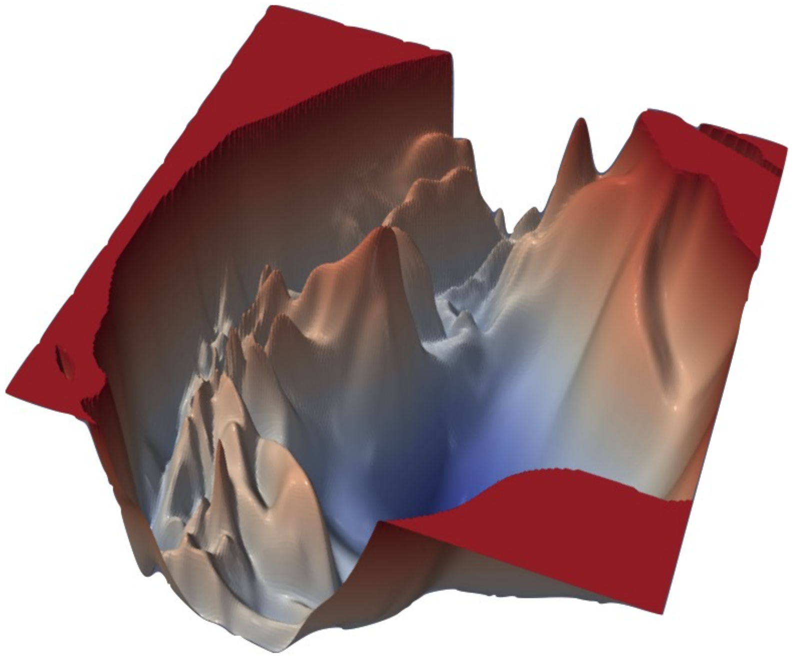

Many spurious local minima

Source: Visualizing the Loss Landscape of Neural Nets

Vulnerable to adversarial inputs.

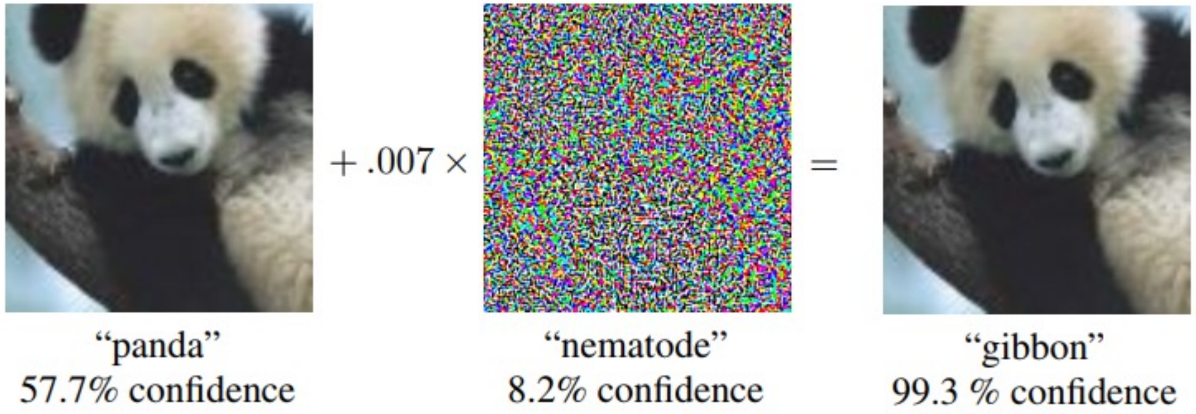

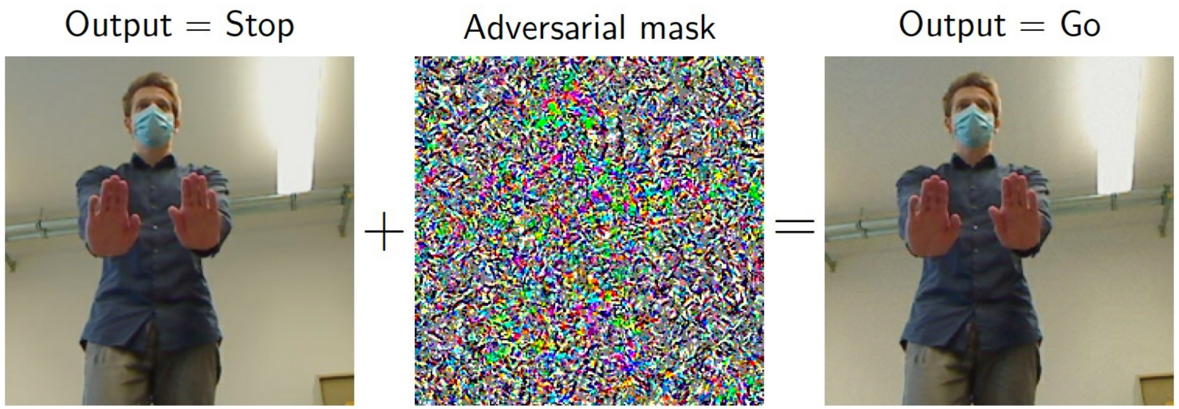

Source: Explaining and Harnessing Adversarial Examples

Robust Classification Background

Geometric interpretation
of adversarial examples.

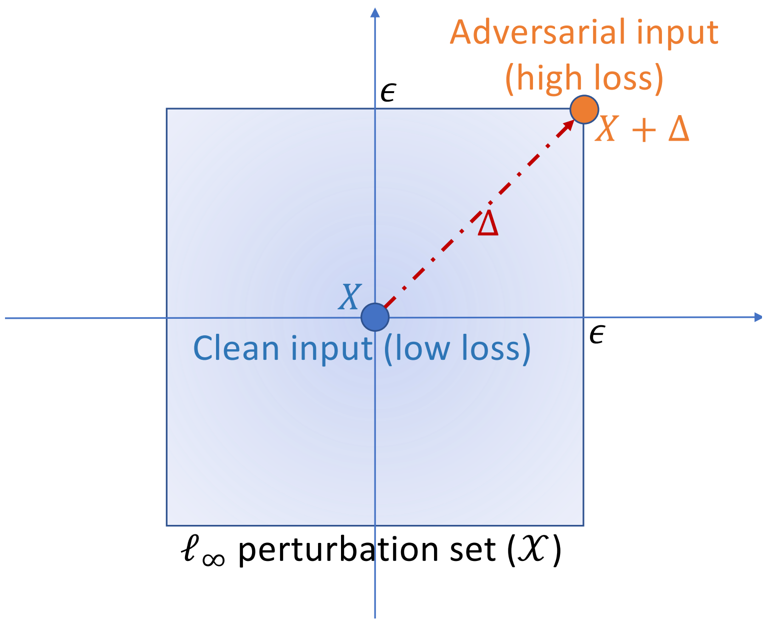

Robust Classification Background

Geometric interpretation of adversarial examples.

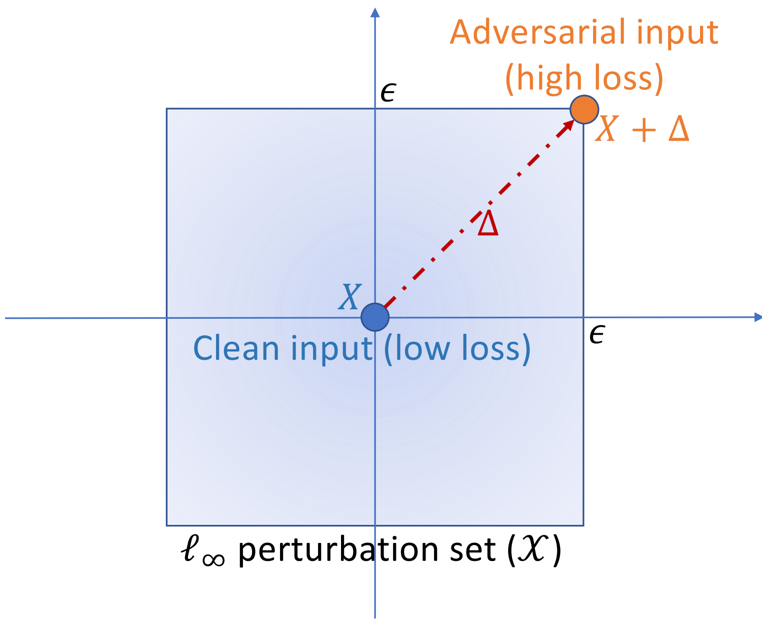

Robust classifiers separate perturbation sets.

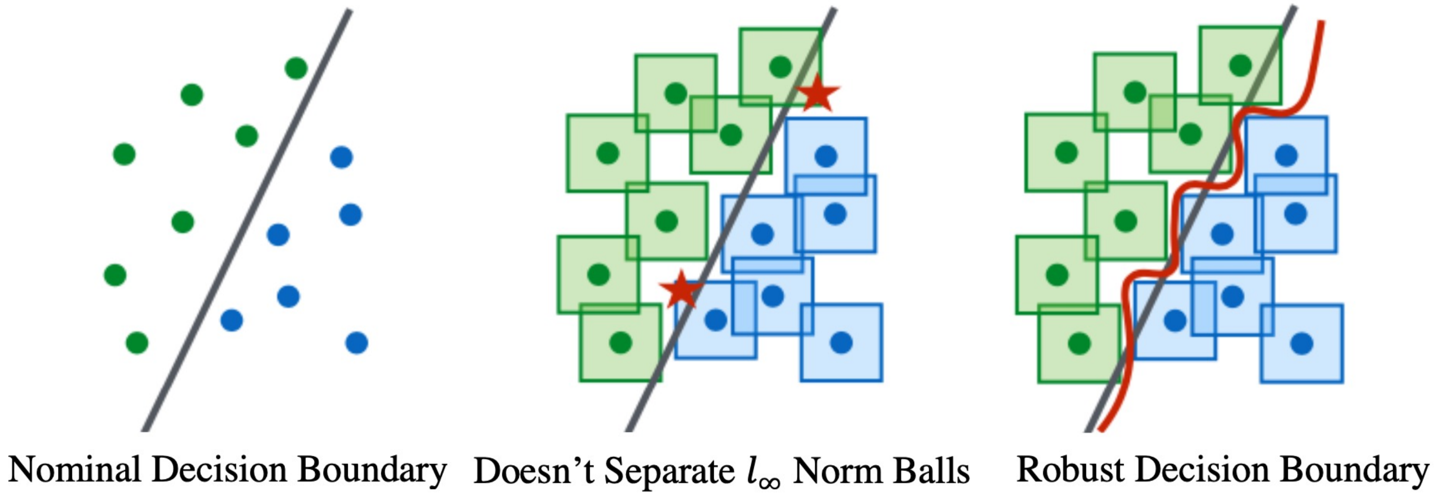

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. International Conference on Learning Representations, 2018.

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Background**
  - Training neural networks with global optimality has been ***intractable***.
  - ***Adversarial training*** builds robust models by training with adversary.
  - Even more challenging optimization: $\underbrace{\min_\theta \quad \underbrace{\max_\epsilon \quad \ell(\theta,\ x+\epsilon)}_{\text{Adversary finds worst perturbation}}}_{\text{Trainer optimizes network parameters}}$.

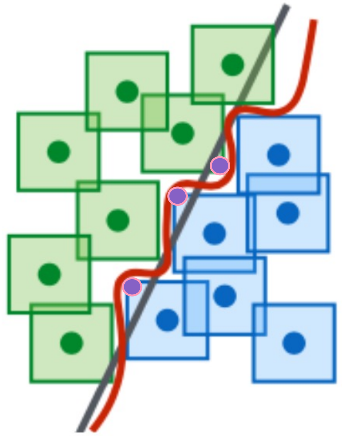

Robust Decision Boundary

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Background**
  - Training neural networks with global optimality has been ***intractable***.
  - ***Adversarial training*** builds robust models by training with adversary.
  - Even more challenging optimization: $\underbrace{\min_{\theta} \underbrace{\max_{\epsilon} \ell(\theta,\ x+\epsilon)}_{\text{Adversary finds worst perturbation}}}_{\text{Trainer optimizes network parameters}}$.

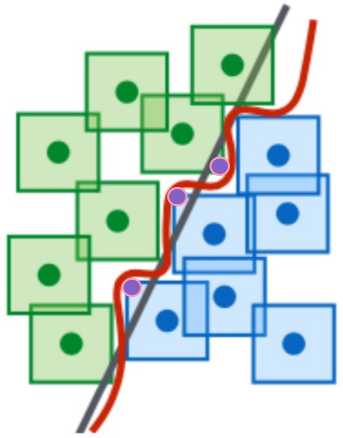

Robust Decision Boundary

- **Convex Training**

Original training problem
Non-convex, unconstrained

Convex training problem
Convex, constrained

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Background**
  - Training neural networks with global optimality has been ***intractable***.
  - ***Adversarial training*** builds robust models by training with adversary.
  - Even more challenging optimization: $\underbrace{\min_\theta \underbrace{\max_\epsilon \ell(\theta, x+\epsilon)}_{\text{Adversary finds worst perturbation}}}_{\text{Trainer optimizes network parameters}}$.

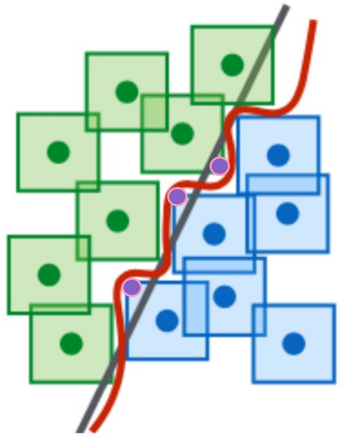

Robust Decision Boundary

- **Convex Training**

Original training problem
Non-convex, unconstrained

Convex training problem
Convex, constrained

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Background**
  - Training neural networks with global optimality has been ***intractable***.
  - ***Adversarial training*** builds robust models by training with adversary.
  - Even more challenging optimization: $\underbrace{\min_{\theta} \underbrace{\max_{\epsilon} \ \ell(\theta, \ x + \epsilon)}_{\text{Adversary finds worst perturbation}}}_{\text{Trainer optimizes network parameters}}$.

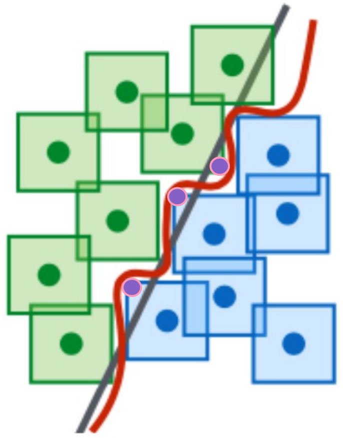

Robust Decision Boundary

- **Convex Training**

Original training problem
Non-convex, unconstrained

Convex training problem
Convex, constrained

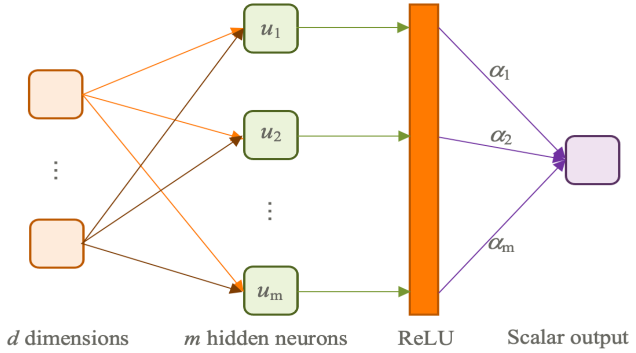

Applies to one-hidden-layer scalar-output neural networks

Convex Optimization for Neural Net Training (*SIMODS, ACC*)

- **Challenges of convex training**
  - Problem size is exponential to data dimension.
  - Traditional algorithm:
    interior point method (cubic complexity).

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Challenges of convex training**
  - Problem size is exponential to data dimension.
  - Traditional algorithm:
    interior point method (cubic complexity).
- **Our solutions**
  - An approximation with provable relaxation gap,
    giving probabilistic ***global optimality***.
  - An ***ADMM algorithm*** with ***quadratic complexity***.
  - Complexity: Previous ***exponential*** $\mathcal{O}\left(d^6\left(\frac{n}{d}\right)^{3d}\right)$
    $\downarrow$
    Ours ***quadratic*** $\mathcal{O}(n^2d^2)$.

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **Challenges of convex training**
  - Problem size is exponential to data dimension.
  - Traditional algorithm:
    interior point method (cubic complexity).
- **Our solutions**
  - An approximation with provable relaxation gap, giving probabilistic ***global optimality***.
  - An ***ADMM algorithm*** with ***quadratic complexity***.
  - Complexity: Previous ***exponential*** $\mathcal{O}(d^6 (\frac{n}{d})^{3d})$
    $\downarrow$
    Ours ***quadratic*** $\mathcal{O}(n^2 d^2)$.

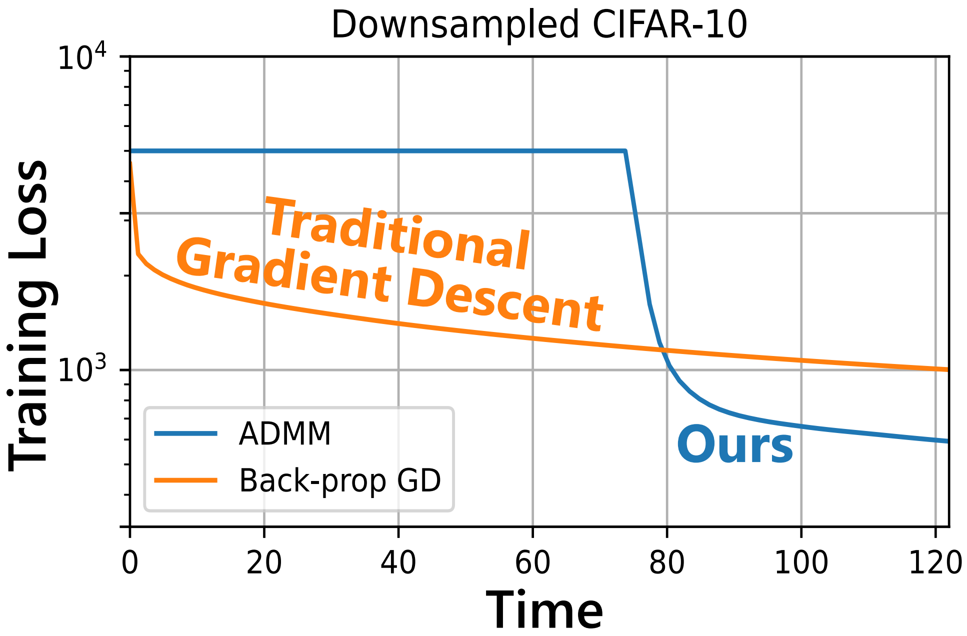

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **A convex optimization problem for adversarial training.**
  - Train robust neural networks with ***global optimality*** (provable upper bound).

Convex Optimization for Neural Net Training (*SIMODS, ACC*)

- **A convex optimization problem for adversarial training.**
  - Train robust neural networks with ***global optimality*** (provable upper bound).

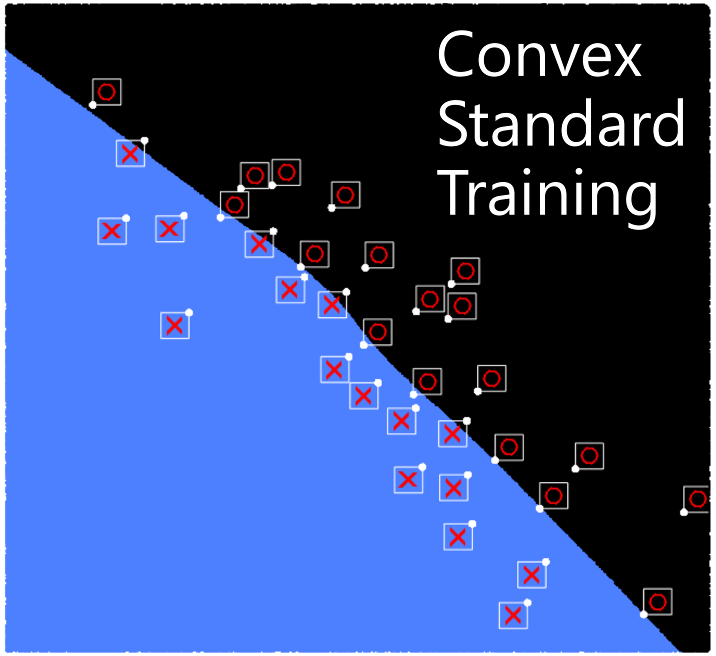

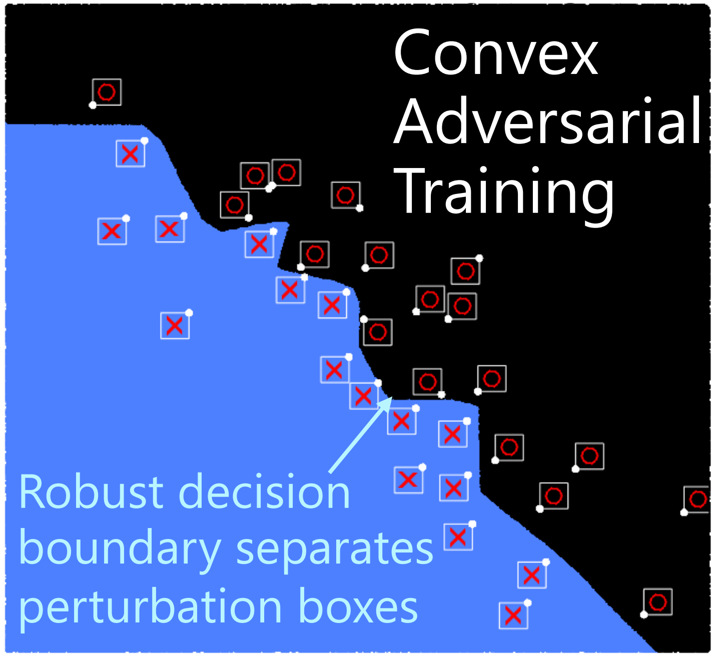

Convex Optimization for Neural Net Training *(SIMODS, ACC)*

- **A convex optimization problem for adversarial training.**
  - Train robust neural networks with ***global optimality*** (provable upper bound).

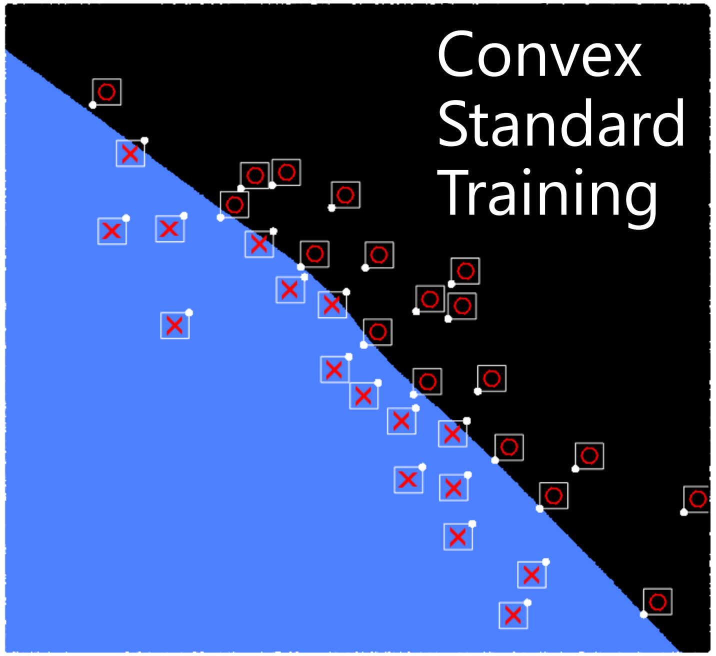

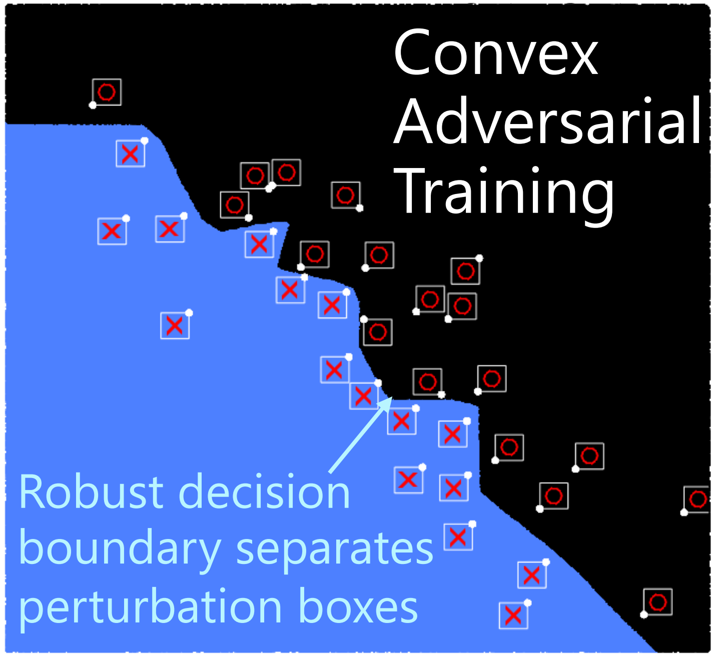

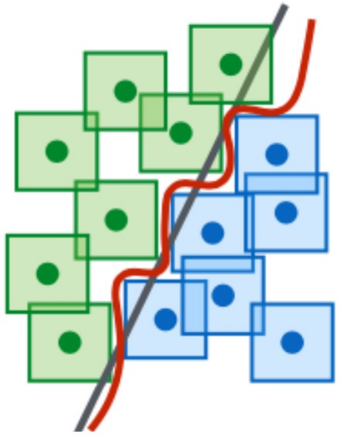

Robust Decision Boundary

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Accuracy-Robustness Trade-Off

- **Robust models often sacrifice "clean accuracy".**
  - **Clean accuracy:**
    accuracy in natural circumstances (no attack).
  - **Robust accuracy:**
    accuracy when subject to attack.

Accuracy-Robustness Trade-Off

- **Robust models often sacrifice "clean accuracy".**
  - **Clean accuracy:**
    accuracy in natural circumstances (no attack).
  - **Robust accuracy:**
    accuracy when subject to attack.

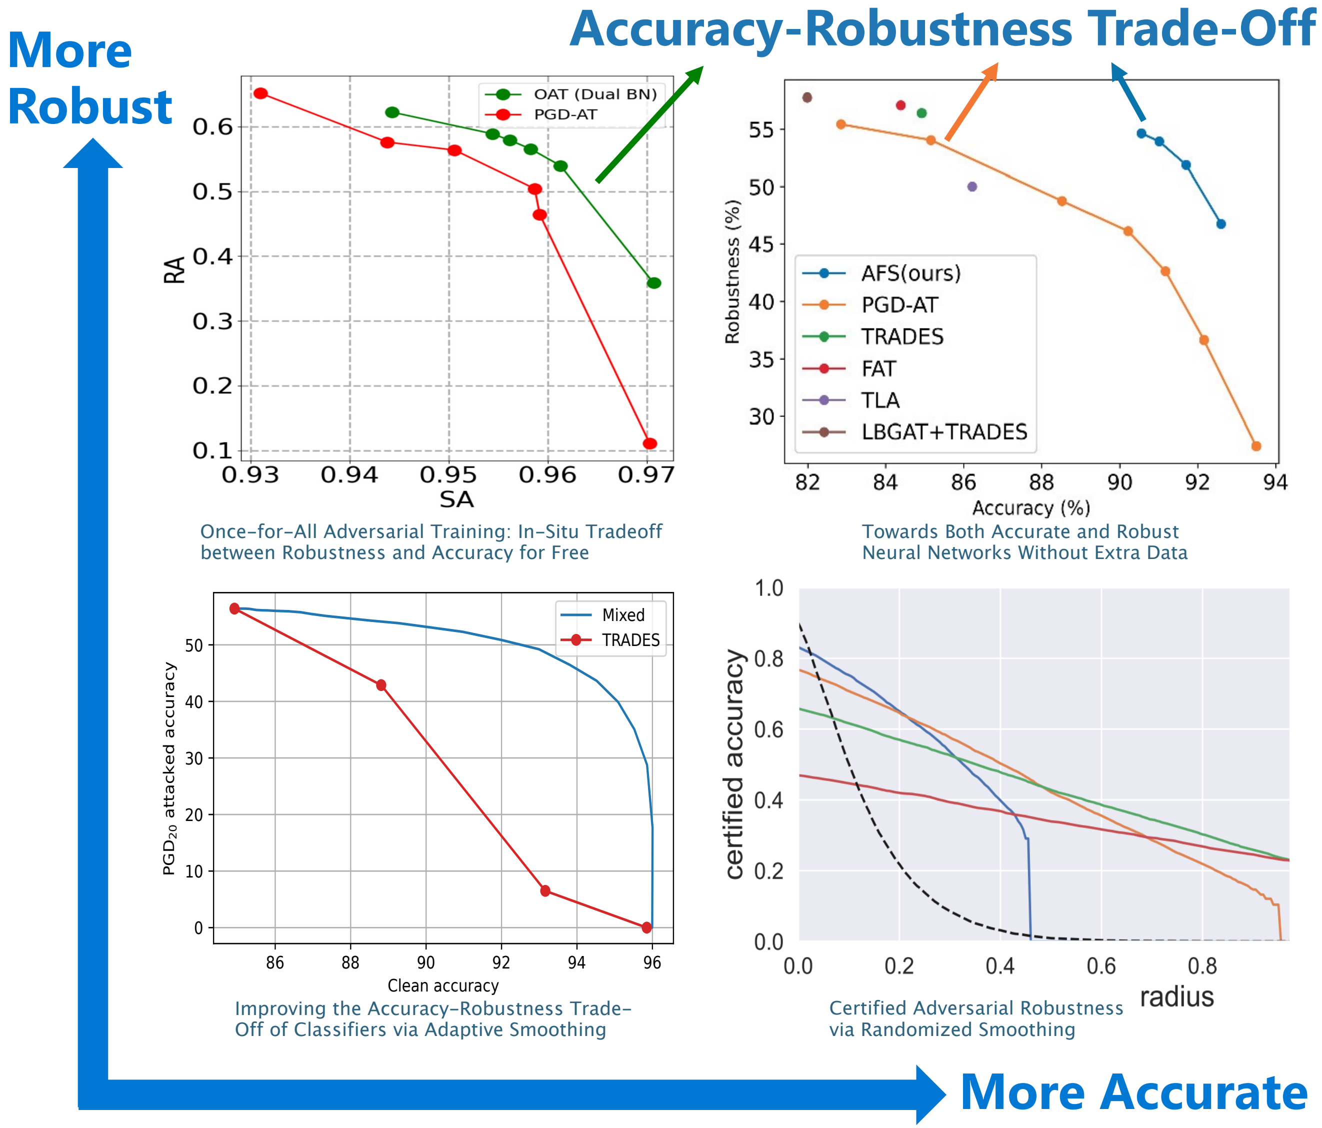

Once-for-All Adversarial Training: In-Situ Tradeoff between Robustness and Accuracy for Free

Towards Both Accurate and Robust Neural Networks Without Extra Data

Improving the Accuracy-Robustness Trade-Off of Classifiers via Adaptive Smoothing

Certified Adversarial Robustness via Randomized Smoothing

Accuracy-Robustness Trade-Off

- **Robust models often sacrifice "clean accuracy".**
  - **Clean accuracy:**
    accuracy in natural circumstances (no attack).
  - **Robust accuracy:**
    accuracy when subject to attack.
- **Implications**
  - Discourages deploying robust models in real life.
  - Real-world services are still unsafe!

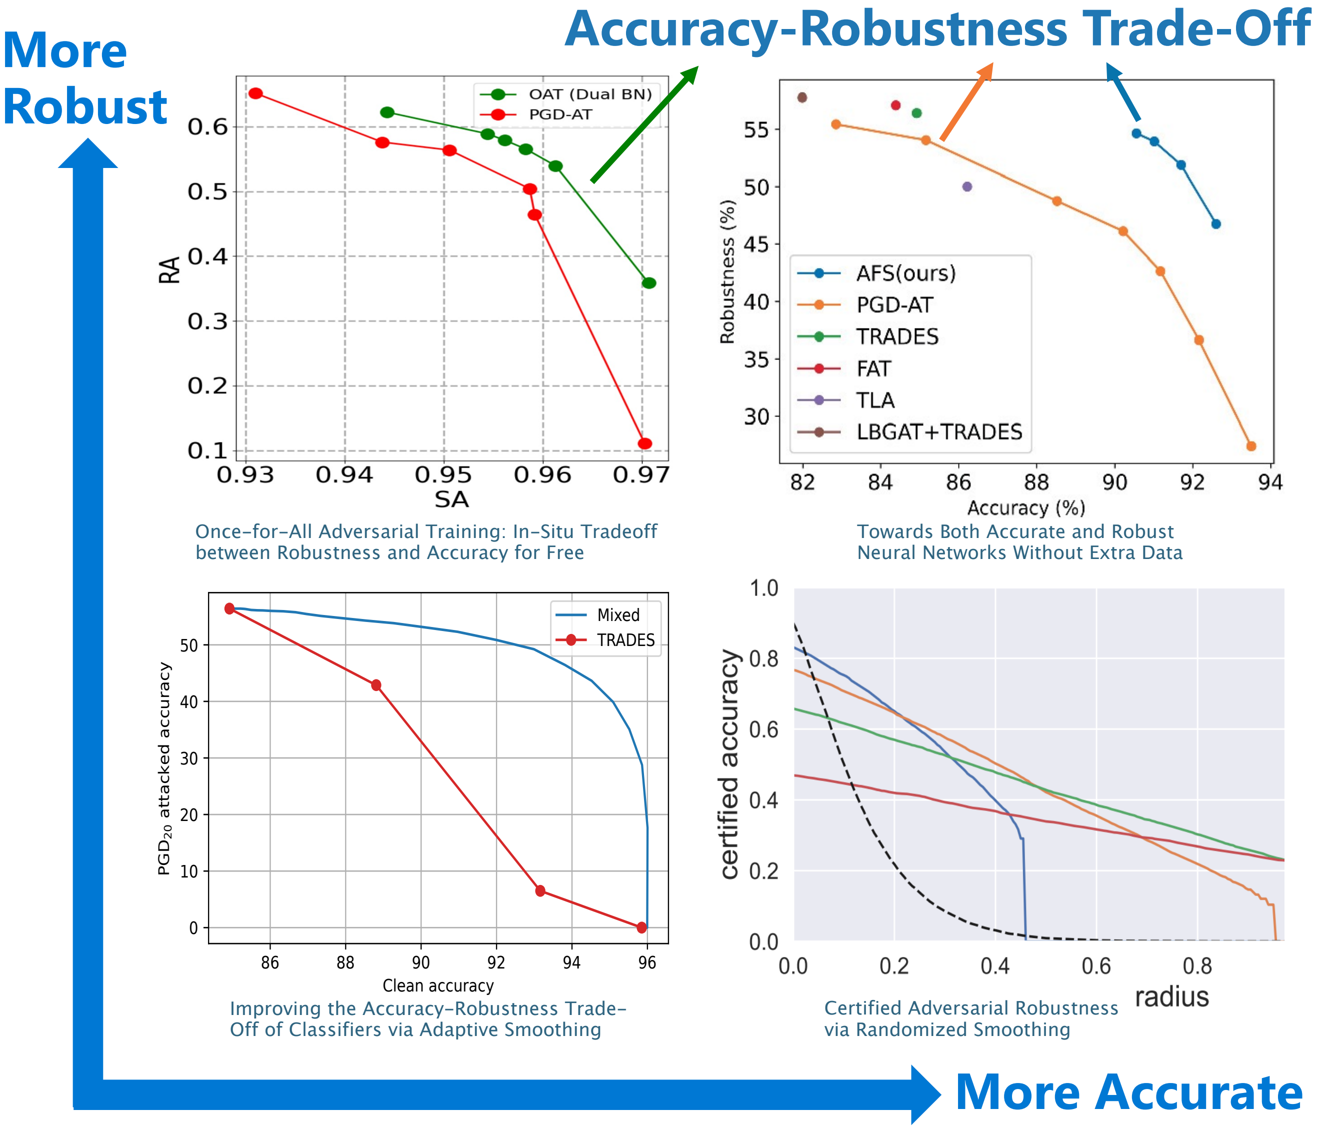

Once-for-All Adversarial Training: In-Situ Tradeoff between Robustness and Accuracy for Free

Towards Both Accurate and Robust Neural Networks Without Extra Data

Improving the Accuracy-Robustness Trade-Off of Classifiers via Adaptive Smoothing

Certified Adversarial Robustness via Randomized Smoothing

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

- With convex training addressing optimization challenges, we now focus on **generalization.**

- **Our solution to the accuracy-robustness trade-off:**
  - Mix the predicted *probabilities* of a robust model and a standard model.

Convert back to logits

Softmax

$$f(x) := \log\left(\underbrace{(1-\alpha)}_{\text{Trade-Off Parameter } \alpha} \cdot \sigma \circ \underbrace{g}_{\text{Accurate Base Classifier (ABC)}}(x) + \alpha \cdot \sigma \circ \underbrace{h}_{\text{Robust Base Classifier (RBC)}}(x)\right)$$

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

Convert back to logits

Softmax

$$f(x) := \log\left((1-\alpha)\cdot \sigma \circ g(x) + \alpha \cdot \sigma \circ h(x)\right)$$

Trade-Off Parameter α

Accurate Base Classifier (ABC)

Robust Base Classifier (RBC)

- **Mixing probability versus logits.**

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

$$f(x) \coloneqq \log\Big(\underbrace{(1-\alpha)}_{\text{Trade-Off Parameter } \alpha} \cdot \overbrace{\sigma}^{\text{Softmax}} \circ \underbrace{g}_{\text{Accurate Base Classifier (ABC)}}(x) + \alpha \cdot \overbrace{\sigma}^{\text{Softmax}} \circ \underbrace{h}_{\text{Robust Base Classifier (RBC)}}(x)\Big)$$

Convert back to logits

- Mixing probability versus logits.

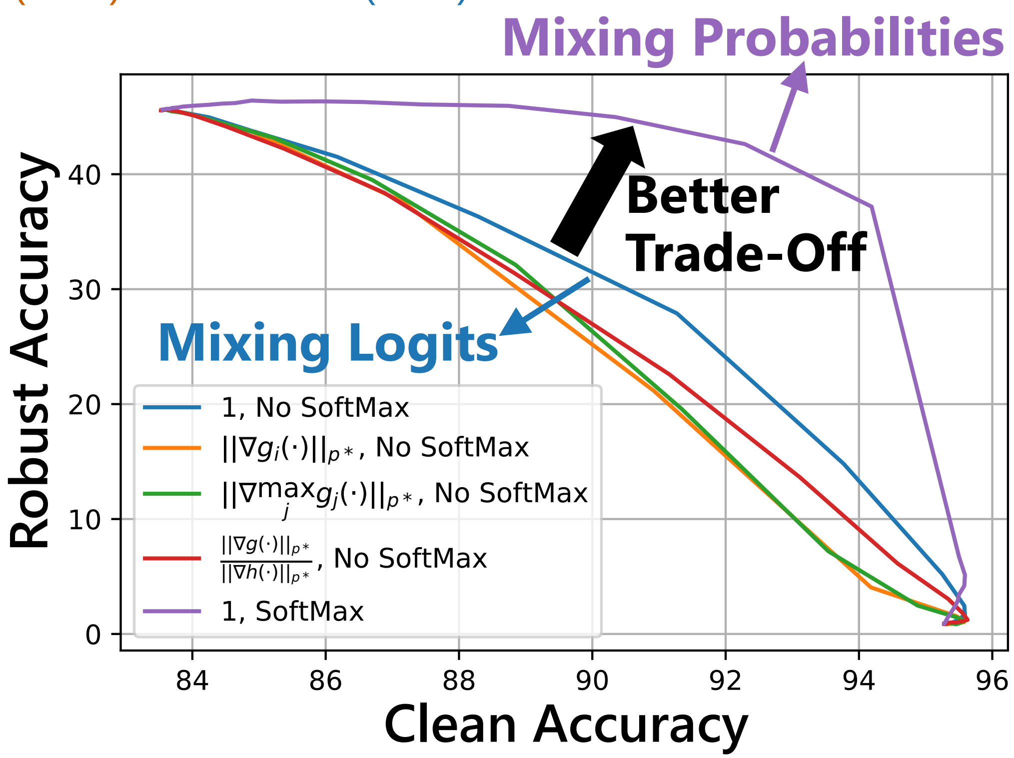

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

Convert back to logits

Softmax

$$f(x) \coloneqq \log\big((1-\alpha)\cdot\sigma\circ g(x) + \alpha\cdot\sigma\circ h(x)\big)$$

Trade-Off Parameter α

Accurate Base Classifier (ABC)

Robust Base Classifier (RBC)

- **Mixing probability versus logits.**
  - **Logits: unbounded.**
    - Can be "arbitrarily wrong".
  - **Probabilities: in $[0, 1]$.**
    - Damage from non-robustness is contained.
- **Mixing probability is better!**

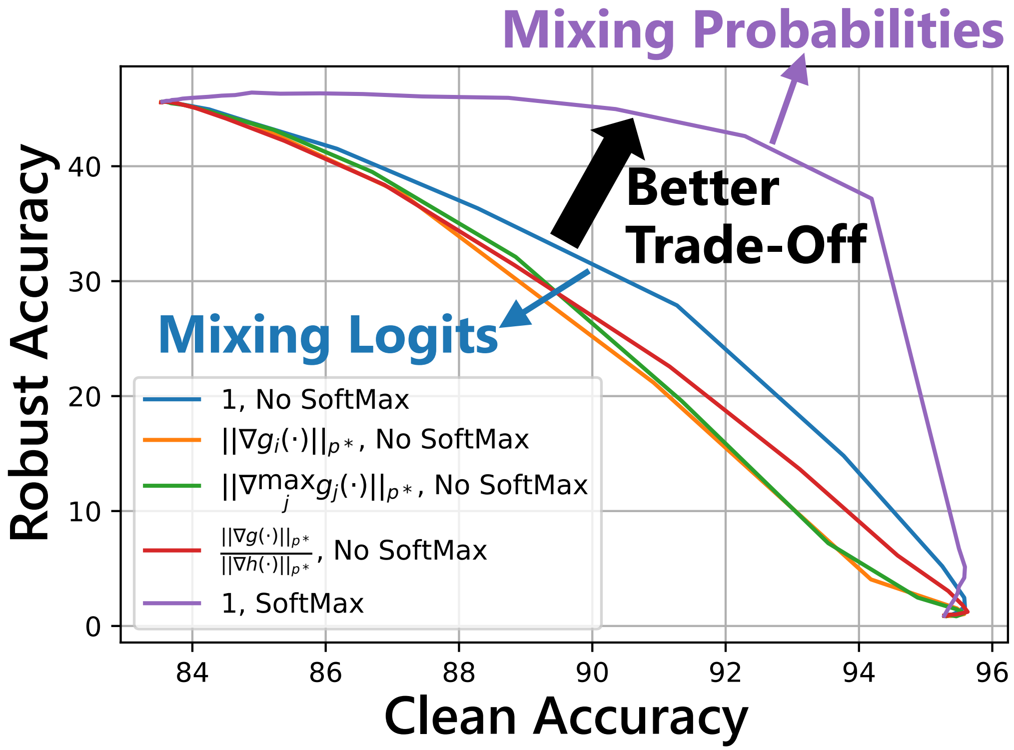

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

$$f(x) \coloneqq \underbrace{\log}_{\text{Convert back to logits}}\Big(\underbrace{(1-\alpha(x))}_{\text{Trade-Off Parameter } \alpha} \cdot \underbrace{\sigma \circ g(x)}_{\text{Accurate Base Classifier (ABC)}} + \alpha(x) \cdot \underbrace{\sigma \circ h(x)}_{\text{Robust Base Classifier (RBC)}}\Big)$$

Softmax

- **Adaptive Smoothing: let α change with x.**

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

Convert back to logits

Softmax

$$f(x) \coloneqq \log\Big(\big(1-\alpha(x)\big)\cdot\sigma\circ g\,(x) + \alpha(x)\cdot\sigma\circ h(x)\Big)$$

Trade-Off Parameter α

Accurate Base Classifier (ABC)

Robust Base Classifier (RBC)

- **Adaptive Smoothing: let α change with x.**

"panda"
57.7% confidence

Clean example

Small α to favor accurate model

"gibbon"
99.3 % confidence

Adversarial example

Large α to favor robust model

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

Convert back to logits

Softmax

$$f(x) := \log\left(\left(1 - \alpha(x)\right) \cdot \sigma \circ g\,(x) + \alpha(x) \cdot \sigma \circ h(x)\right)$$

Trade-Off Parameter α

Accurate Base Classifier (ABC)

Robust Base Classifier (RBC)

- **Adaptive Smoothing: let α change with x.**
  - The ***mixing network*** $\alpha(x)$: a new neural network component.
  - Train $\alpha(x)$ with strong adversaries that exploits the new structure.

"panda"
57.7% confidence

Clean example

Small α to favor accurate model

"gibbon"
99.3 % confidence

Adversarial example

Large α to favor robust model

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

- **Why does mixing probabilities improve the trade-off?**
  Robust models are more confident when correct than when incorrect, even when attacked.

I.e., Orange (attacked correct) is higher than Blue (clean incorrect) in the confidence plot.

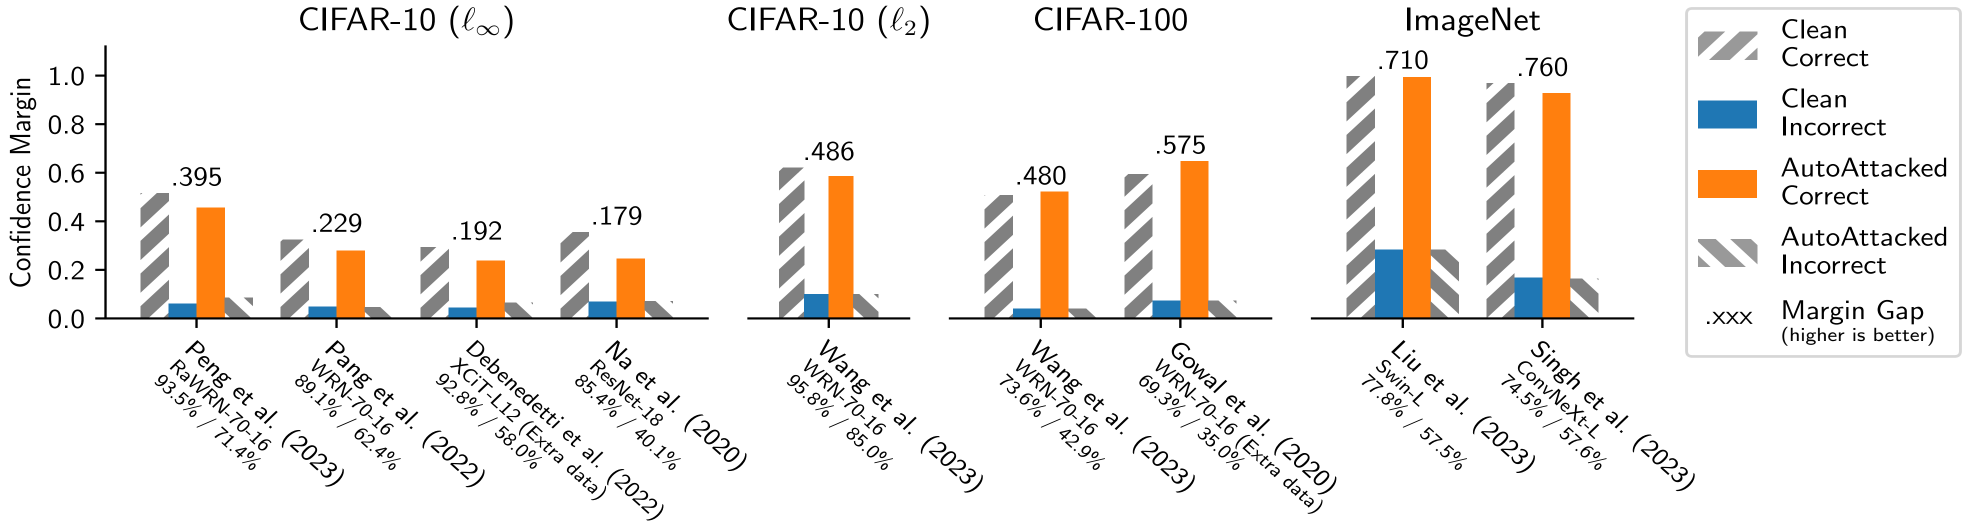

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

- **Why does mixing probabilities improve the trade-off?**
  Robust models are more confident when correct than when incorrect, even when attacked.

I.e., Orange (attacked correct) is higher than Blue (clean incorrect) in the confidence plot.

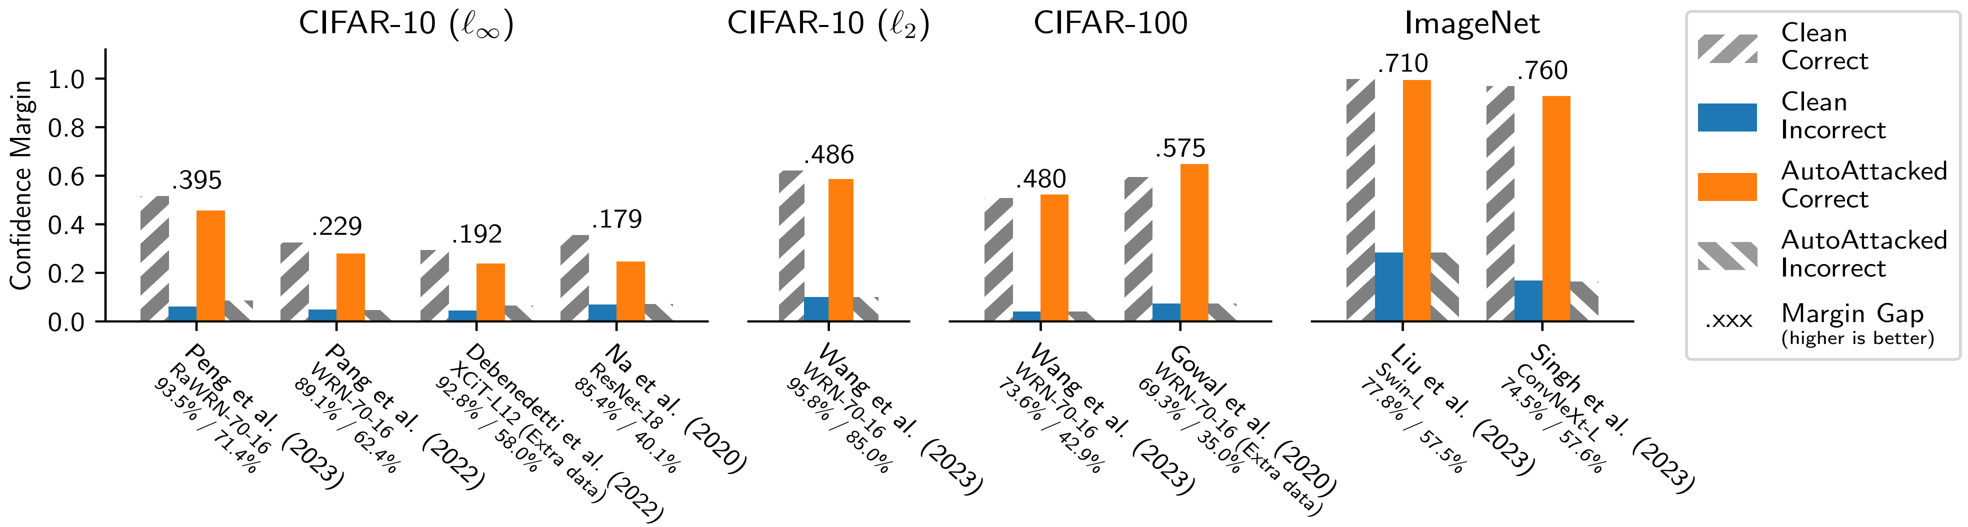

- **Can we "enlarge" this benign confidence property?**
  Apply non-linear transformation to the robust model logits $h(x)$.

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

- **Mixing with non-linear transformation (MixedNUTS) improves accuracy-robustness balance.**

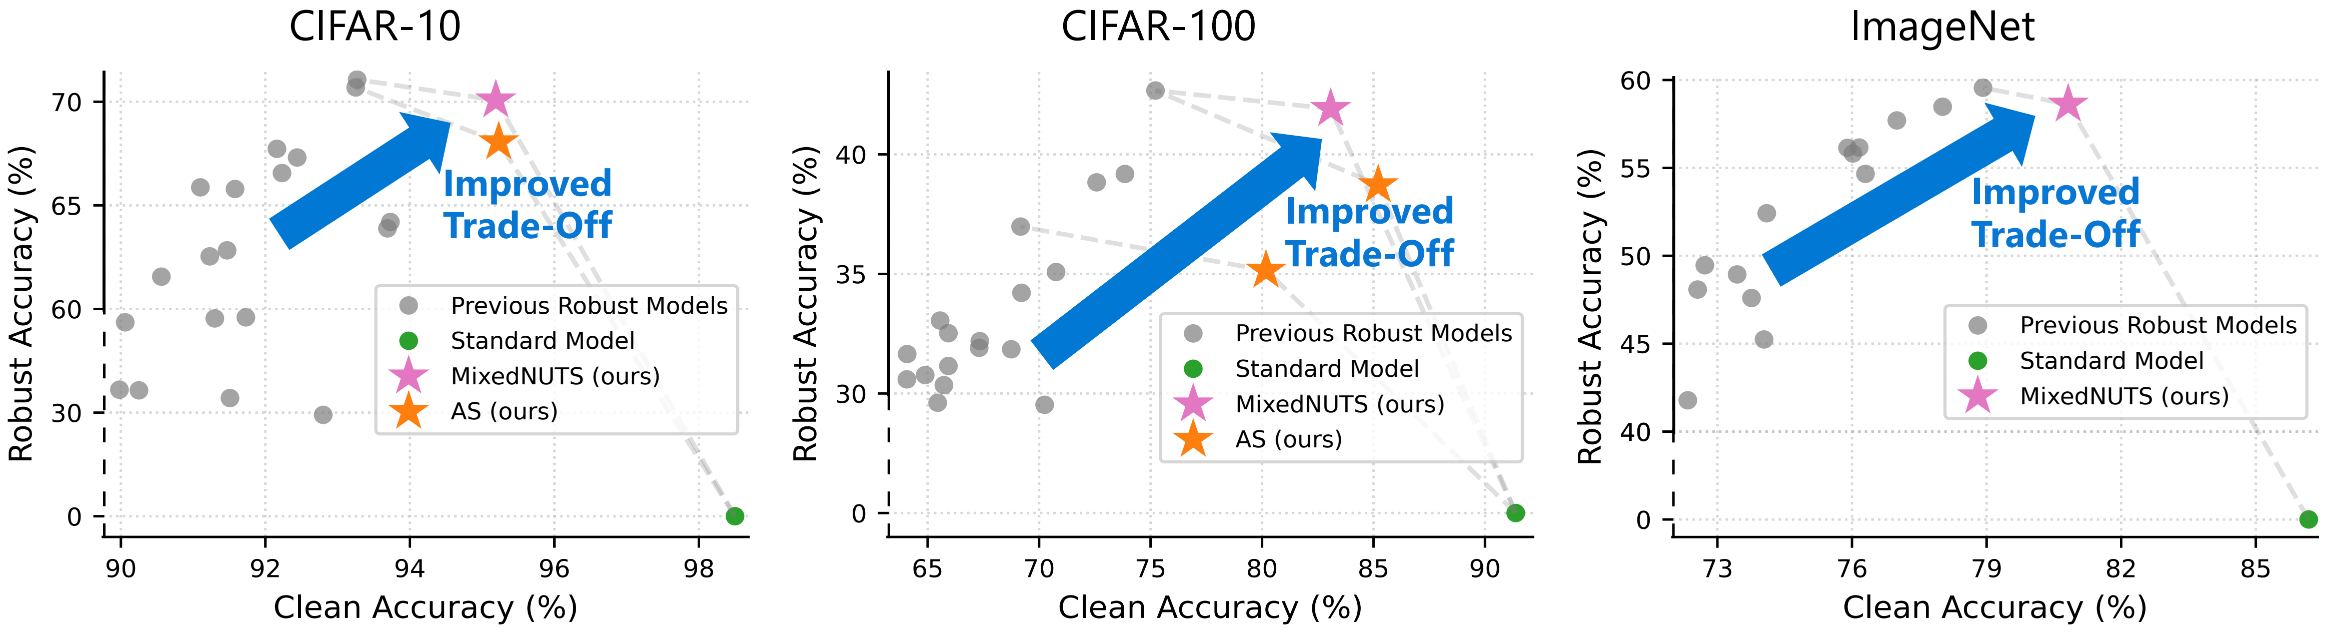

Tackling Accuracy-Robustness Trade-Off *(TMLR, SIMODS, L4DC)*

- Mixing with non-linear transformation (MixedNUTS) improves accuracy-robustness balance.

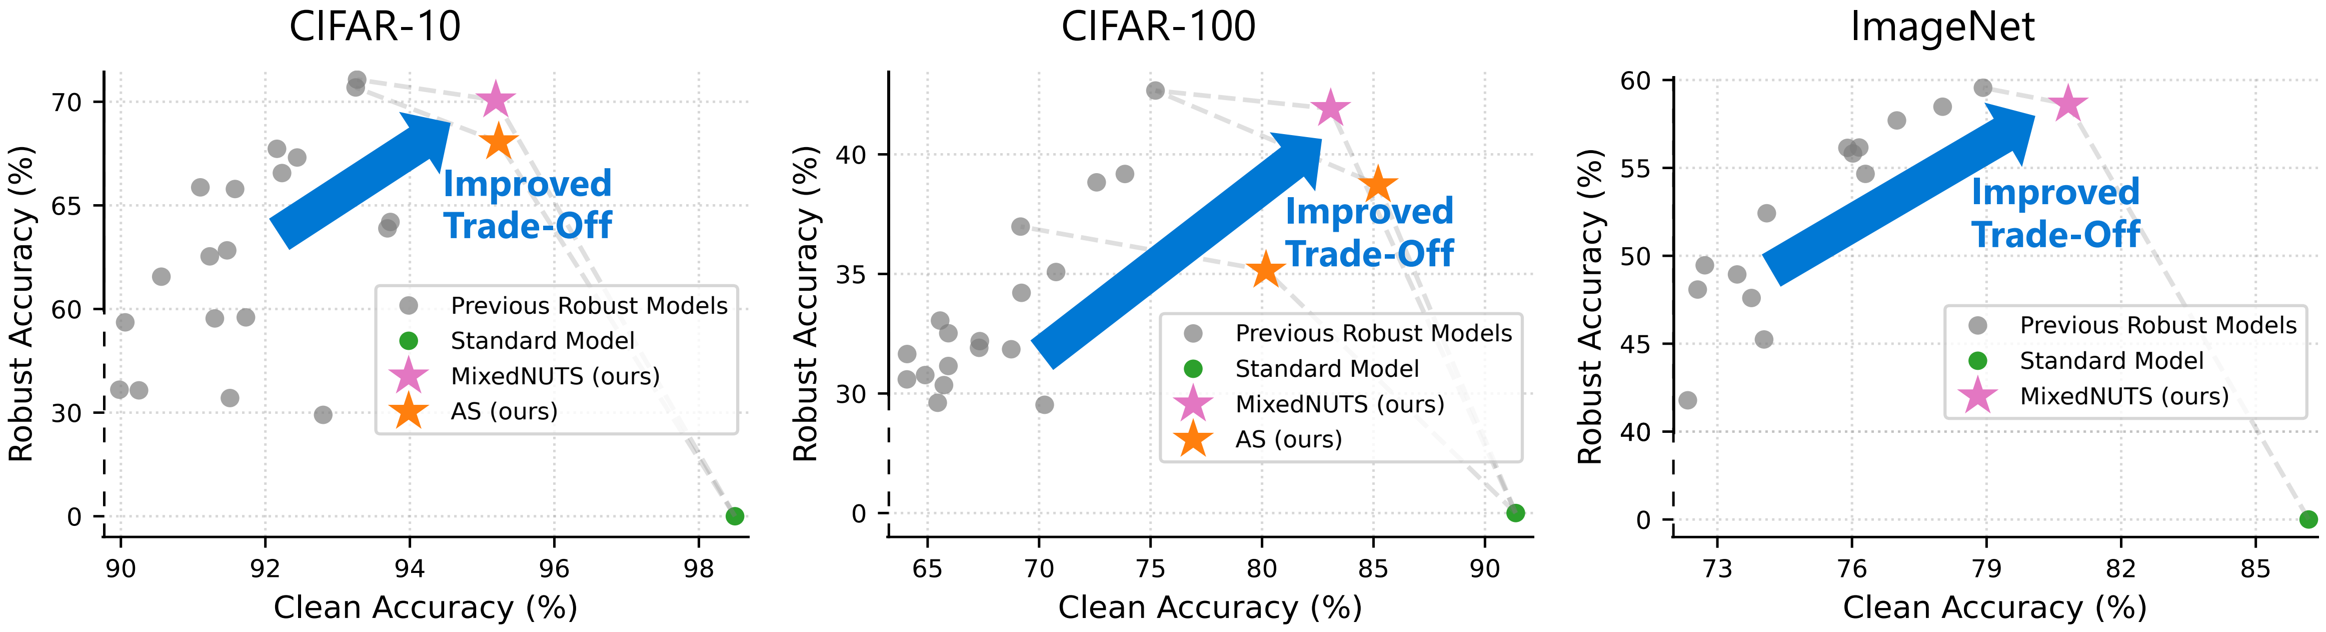

- Also: certified robustness.
- Mostly training-free.

Summary On Robust Classification

	Efficiency	Reliability
Convex Training		
Mixing Classifiers		

Summary On Robust Classification

	Efficiency	Reliability
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers		

Summary On Robust Classification

	Efficiency	**Reliability**
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers	• Training-free • Plug-and-play	• Interpretable formulation • Robust models are now practical

Summary On Robust Classification

	Efficiency	Reliability
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers	• Training-free • Plug-and-play	• Interpretable formulation • Robust models are now practical

- So far, we made ***discriminative models*** more dependable.
  - Especially when the training data does not cover all scenarios.
- Next, we discuss ***generative models***.
  - A different train-test mismatch; different efficiency and reliability challenges.

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Media Generation

- **Media generation, a recently emerged impactful deep learning area**
  - E.g., audio, music, images, videos.
  - Models interact with people in a creative way.
  - Alignment with human need is paramount!

AI-generated cover image
for a research project

Media Generation

- **Media generation, a recently emerged impactful deep learning area**
  - E.g., audio, music, images, videos.
  - Models interact with people in a creative way.
  - Alignment with human need is paramount!

- **Audio/music creation**
  - Global music industry reached US$26.2 billion in 2022.
  - Film and video market reached US$273.35 billion.
  - Amateurs can now become composers/directors!

AI-generated cover image
for a research project

Diffusion Model Background

- **Diffusion models** are one of the **most popular approaches** to media generation.
  - Training

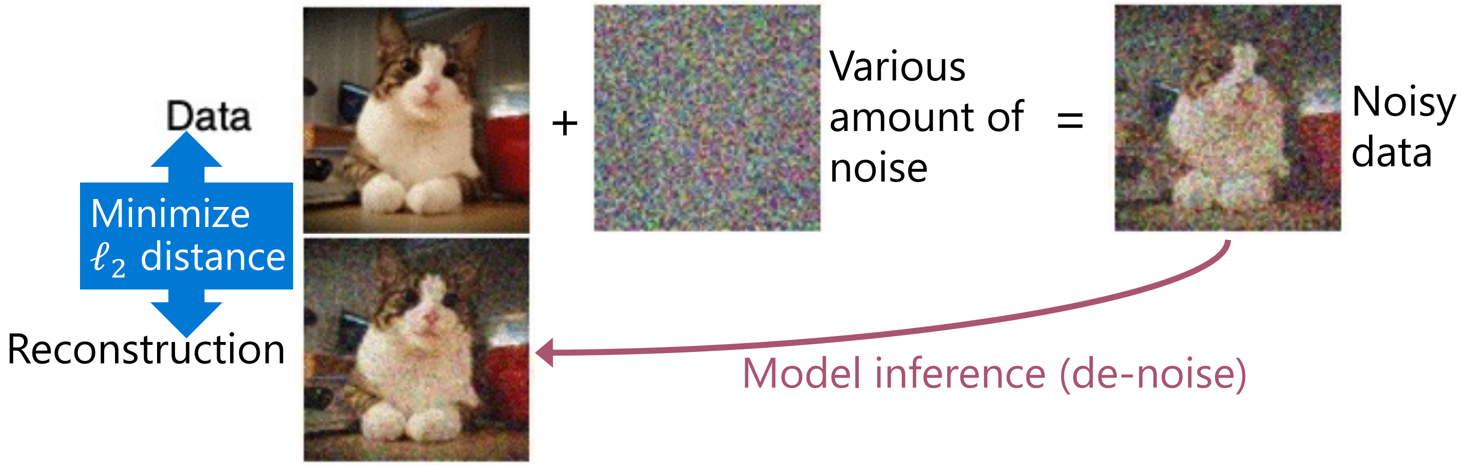

Diffusion Model Background

- **Diffusion models** are one of the **most popular approaches** to media generation.
  - Training

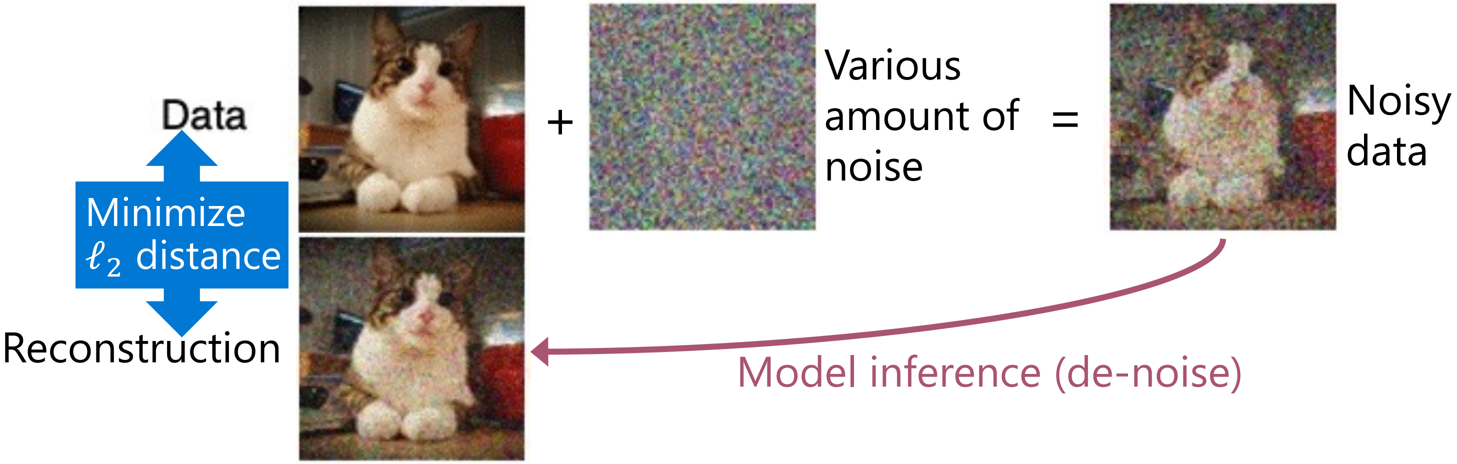

  - Inference

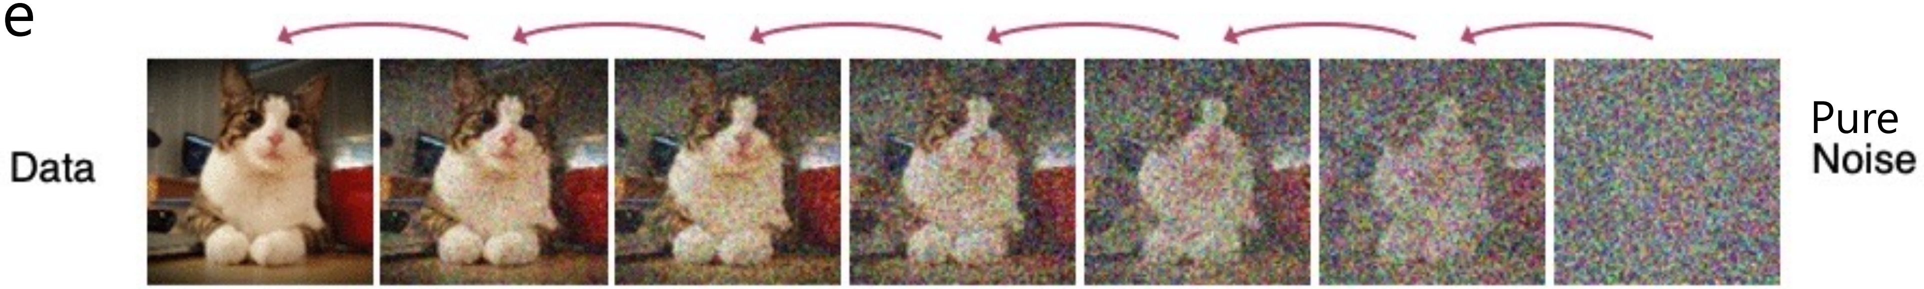

Source: https://developer.nvidia.com/blog/improving-diffusion-models-as-an-alternative-to-gans-part-1

Training Objective Mismatch

- **Diffusion models' training objective** (de-noising) **does not match the target task** (creative generation).

Training objective *denoising*	$\neq$	**Target task** *creative generation*

Training Objective Mismatch

- **Diffusion models' training objective** (de-noising) **does not match the target task** (creative generation).

Training objective
denoising

≠

Target task
creative generation

- **Two issues:**
  - Slow inference (due to iterative inference).
  - Reward misalignment (good denoiser ≠ good creator).

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

ConsistencyTTA *(INTERSPEECH* 2024)

- Can we tackle both issues via non-iterative inference?

ConsistencyTTA *(INTERSPEECH* 2024)

- Can we tackle both issues via non-iterative inference?
- **Accelerate diffusion-based text-to-audio generation with consistency distillation.**
  - In-the-wild audio (environmental sound).
  - ***400x*** theoretical acceleration.
  - ***72x*** real-world speed-up.
  - Minimal change in audio quality.

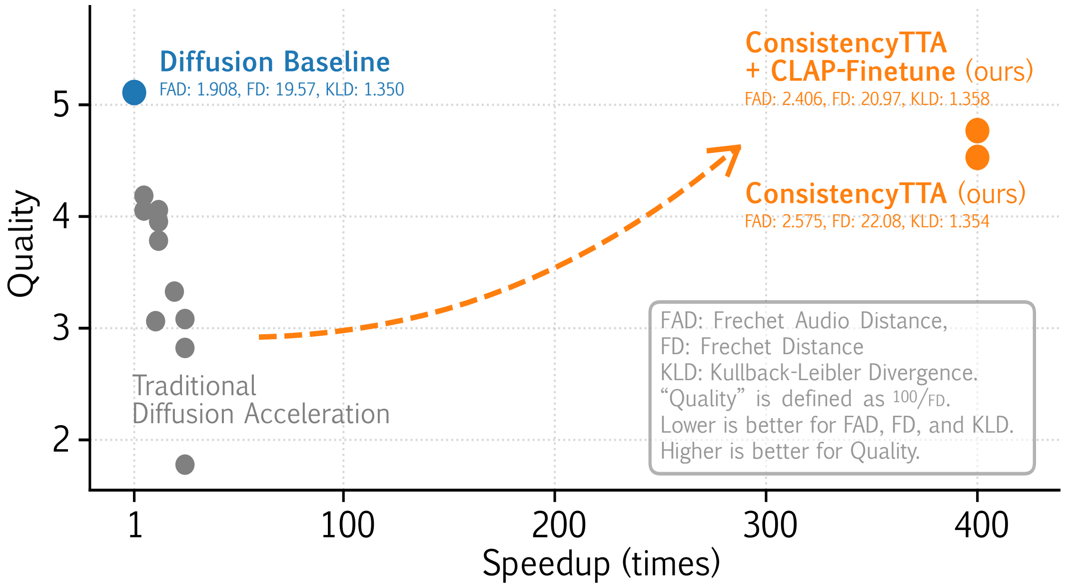

ConsistencyTTA *(INTERSPEECH* 2024)

- **Consistency distillation**
  - **Condensed model capability**: same model size, inference iterations decreased to ***1***.

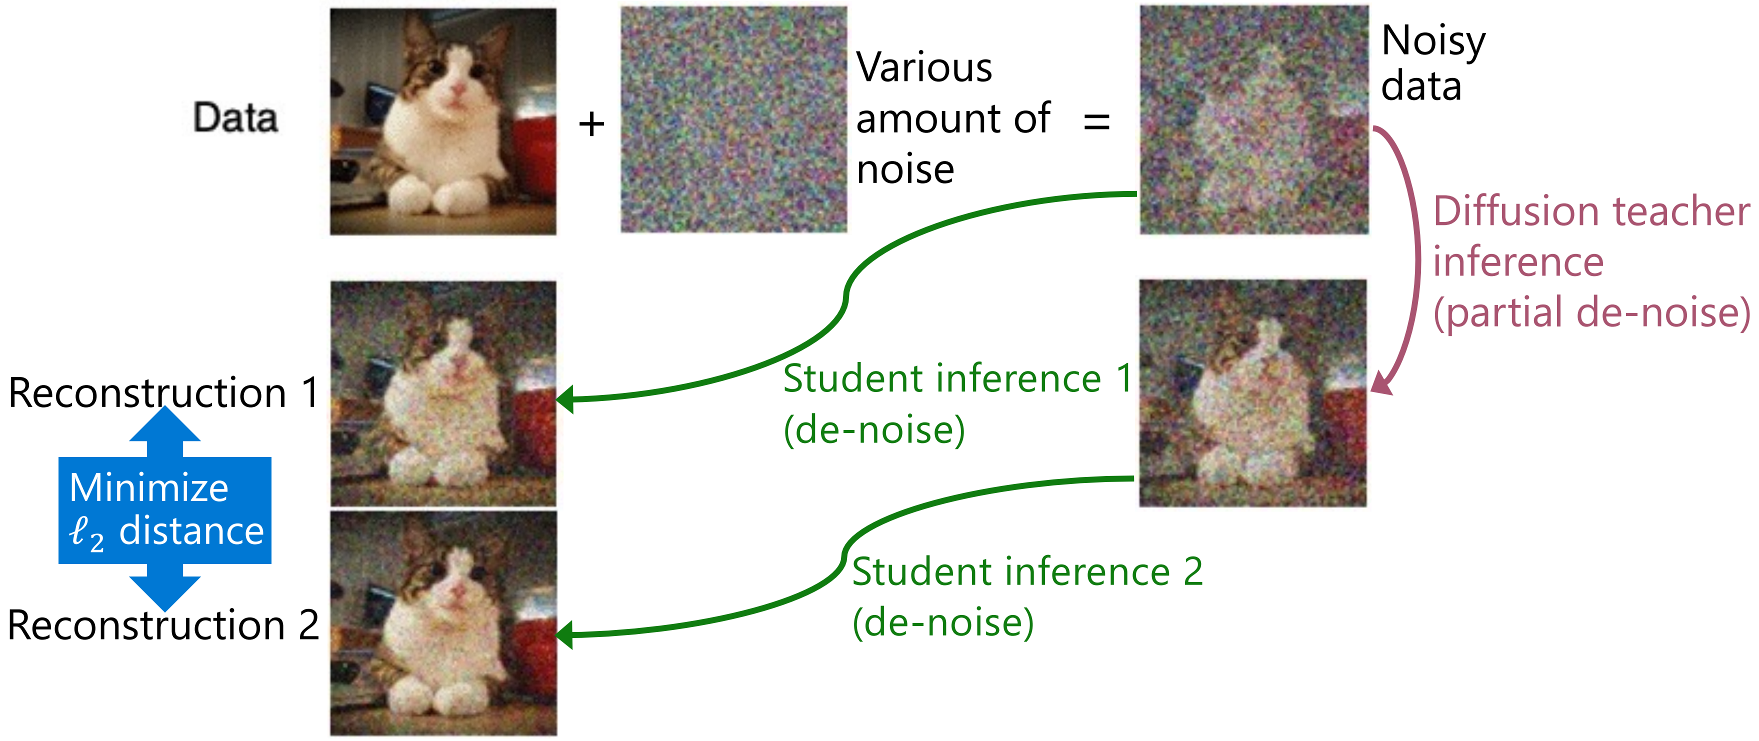

ConsistencyTTA *(INTERSPEECH* 2024)

- **Consistency distillation**
  - **Condensed model capability**: same model size, inference iterations decreased to ***1***.

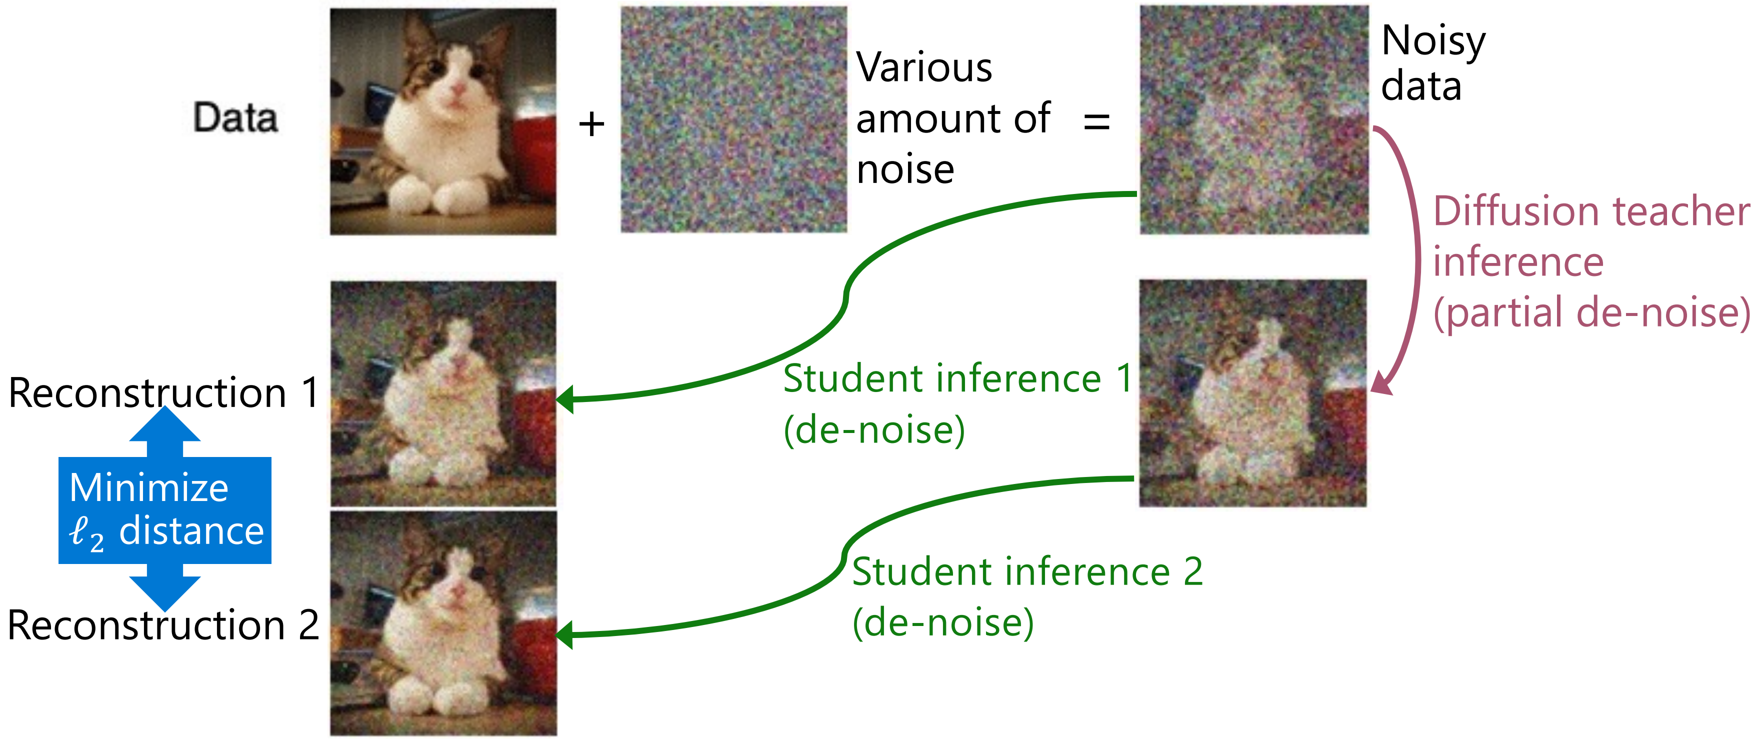

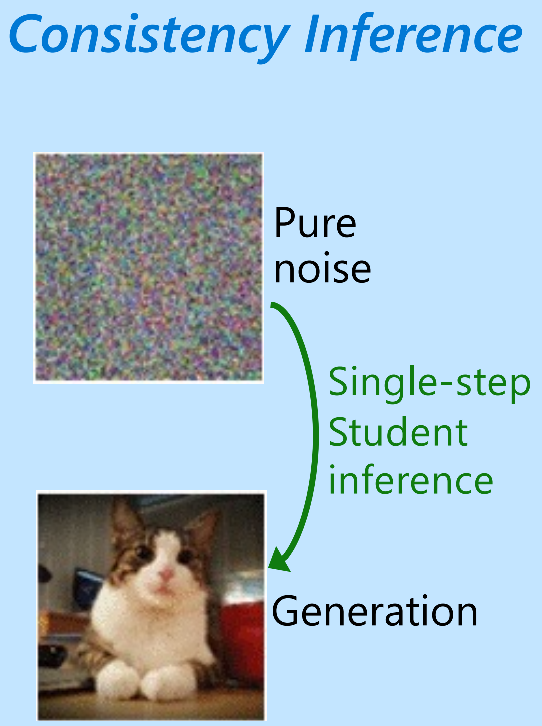

ConsistencyTTA *(INTERSPEECH* 2024)

- **Classifier-Free Guidance:** inference-time operation outside the denoiser that enhances results.
- **CFG-Aware Distillation:**

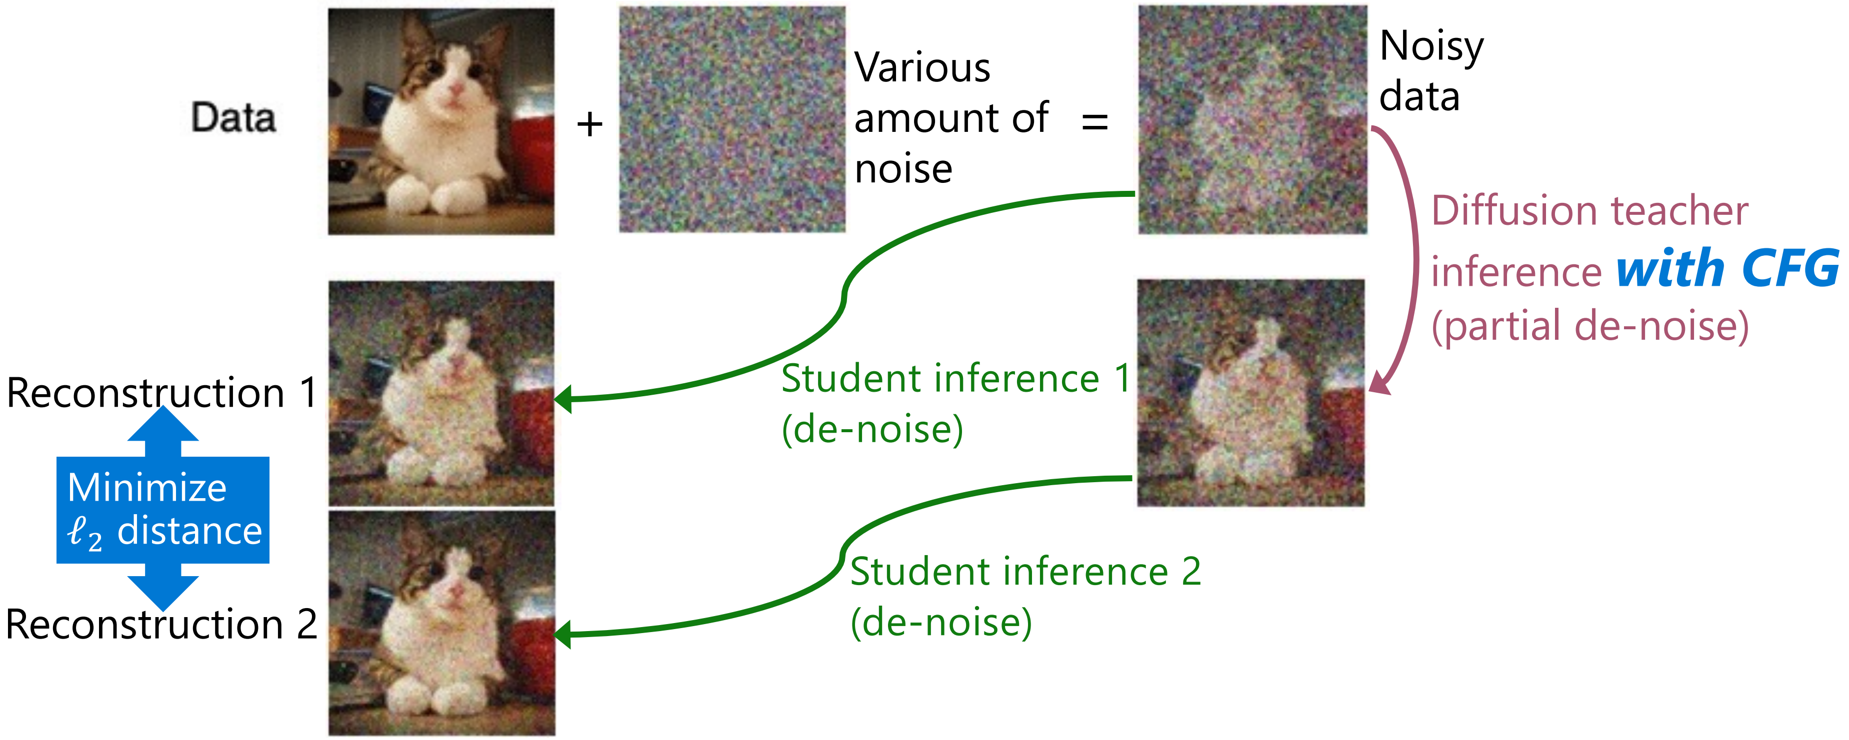

ConsistencyTTA *(INTERSPEECH* 2024)

- Now, our model has ***non-iterative inference*** and is ***end-to-end differentiable***.

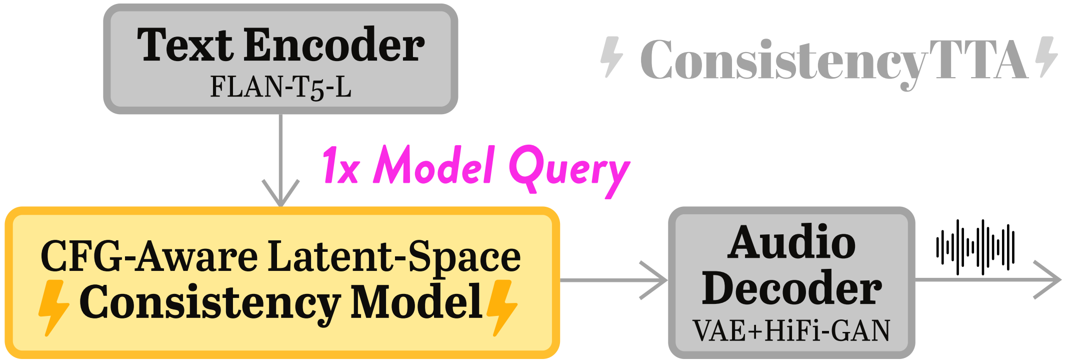

ConsistencyTTA *(INTERSPEECH 2024)*

- Now, our model has ***non-iterative inference*** and is ***end-to-end differentiable***.
- We can ***fine-tune target task reward functions*** to address train-test mismatch.
  - **CLAP Score:** cosine similarity of a generation and a reference in an embedding space.

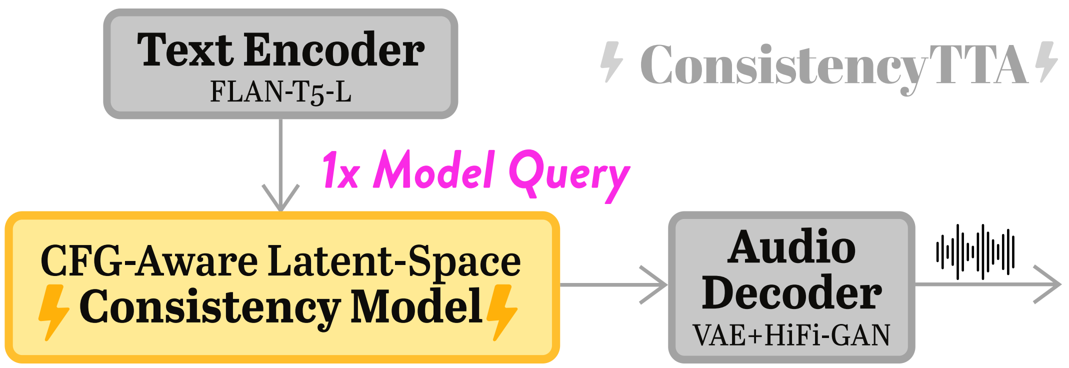

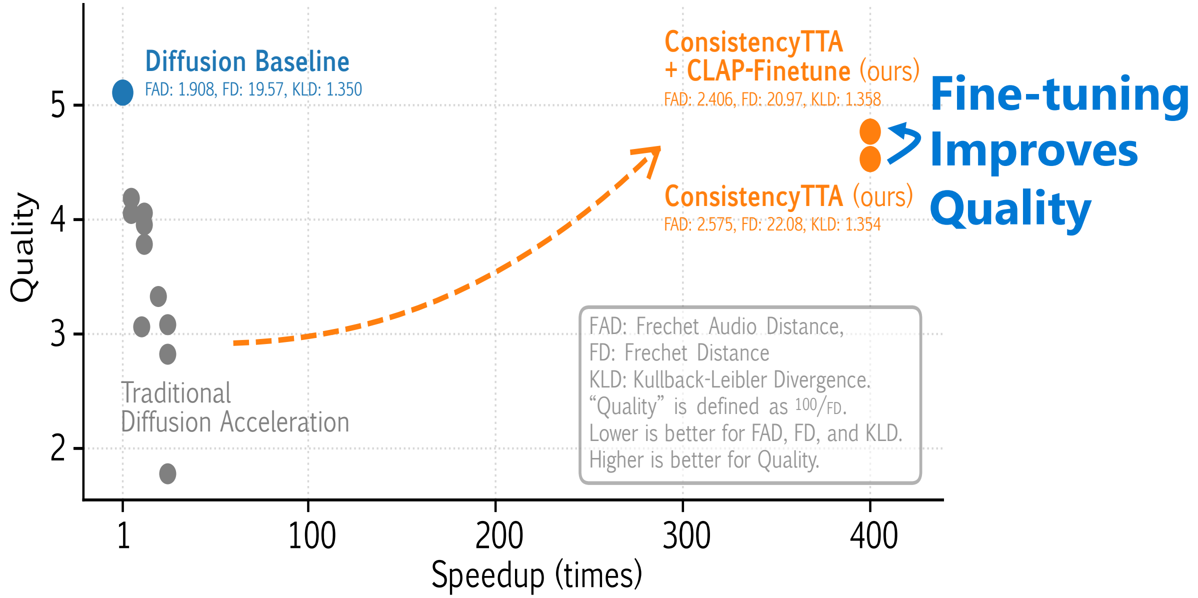

ConsistencyTTA Live Demo

- Demo Link
  https://huggingface.co/spaces/Bai-YT/ConsistencyTTA

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- Summary.

Optimizing Distributional Rewards Enhances Diffusion Models

- ConsistencyTTA tackled training objective misalignment with non-iterative inference.

Optimizing Distributional Rewards Enhances Diffusion Models

- ConsistencyTTA tackled training objective misalignment with non-iterative inference.

- Can we instead make **reward optimization compatible with iterative denoising?**
- Can we make diffusion media generation **more aligned with human preference?**

Optimizing Distributional Rewards Enhances Diffusion Models

- ConsistencyTTA tackled training objective misalignment with non-iterative inference.

- Can we instead make **reward optimization compatible with iterative denoising?**
- Can we make diffusion media generation **more aligned with human preference?**

- **We propose DRAGON.**
  - An online on-policy **reward optimization** framework for media creation. Compatible with reward functions that evaluate **individual examples or distributions**.

DRAGON Method

- **Goal:**
  - Maximize a reward function
    $$r_{dist}\colon \mathcal{P} \to \mathbb{R}$$
    that evaluates **distributions**.
  - Per-instance reward special case
    $r_{\text{dist}}(D_\theta) = \mathbb{E}_{X \sim D_\theta}\, r_{\text{instance}}(X)$.

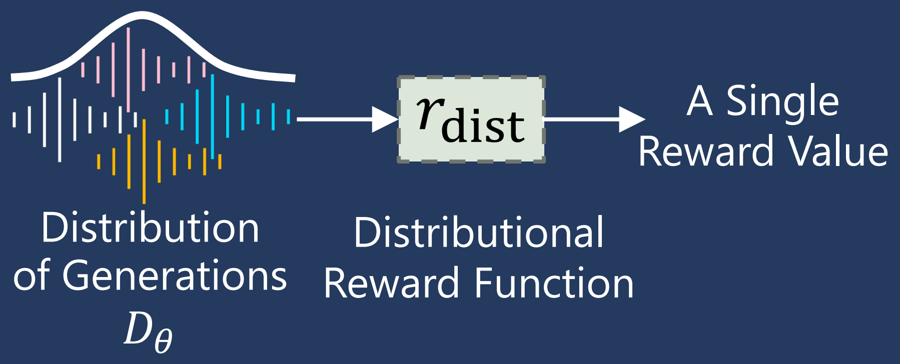

DRAGON Method

- **Goal:**
  - Maximize a reward function
    $$r_{dist}: \mathcal{P} \to \mathbb{R}$$
    that evaluates **distributions**.
  - Per-instance reward special case
    $r_{\text{dist}}(D_\theta) = \mathbb{E}_{X \sim D_\theta}\, r_{\text{instance}}(X)$.

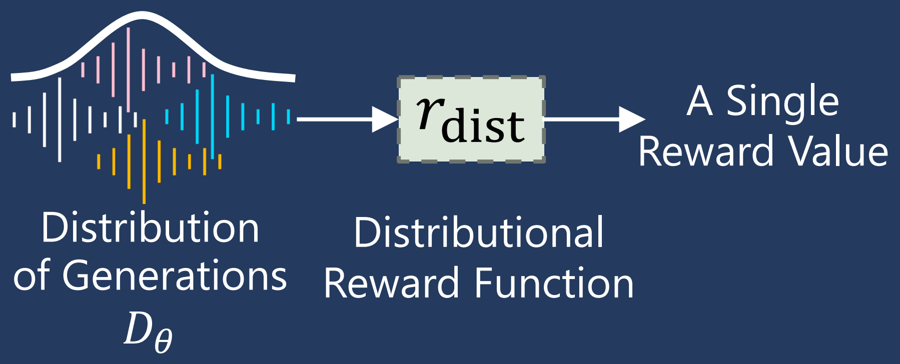

- **DRAGON:**

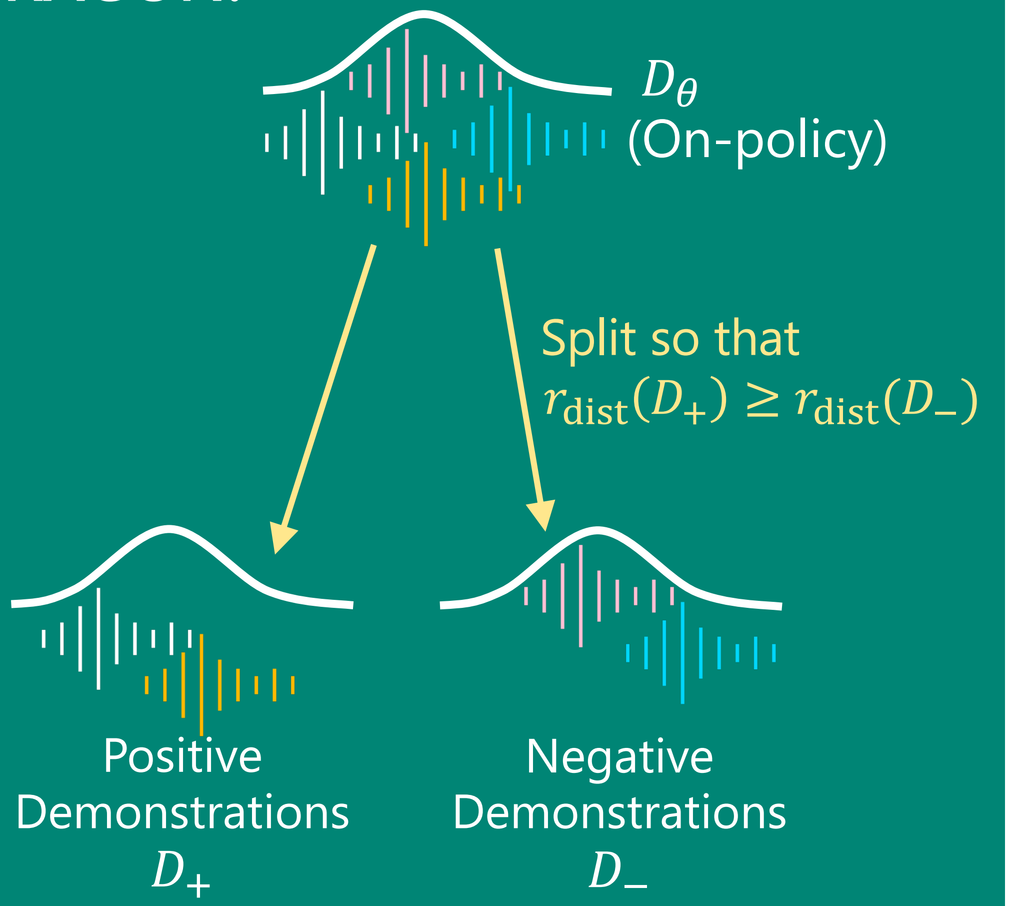

DRAGON Method

- **Goal:**
  - Maximize a reward function
    $$r_{dist}: \mathcal{P} \to \mathbb{R}$$
    that evaluates **distributions**.
  - Per-instance reward special case
    $r_{\text{dist}}(D_\theta) = \mathbb{E}_{X \sim D_\theta}\, r_{\text{instance}}(X).$

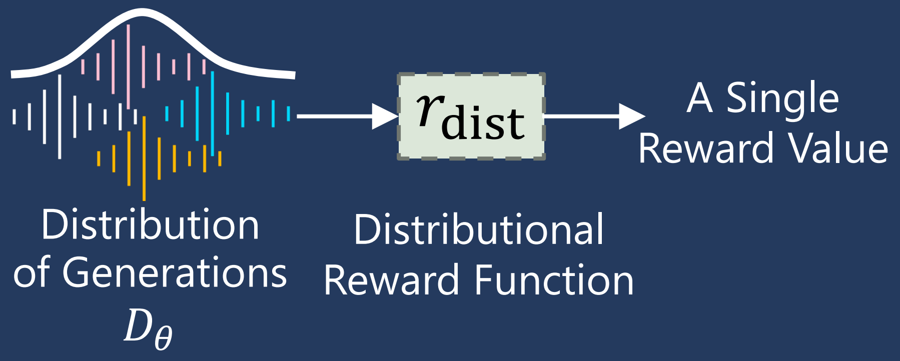

- **DRAGON:**

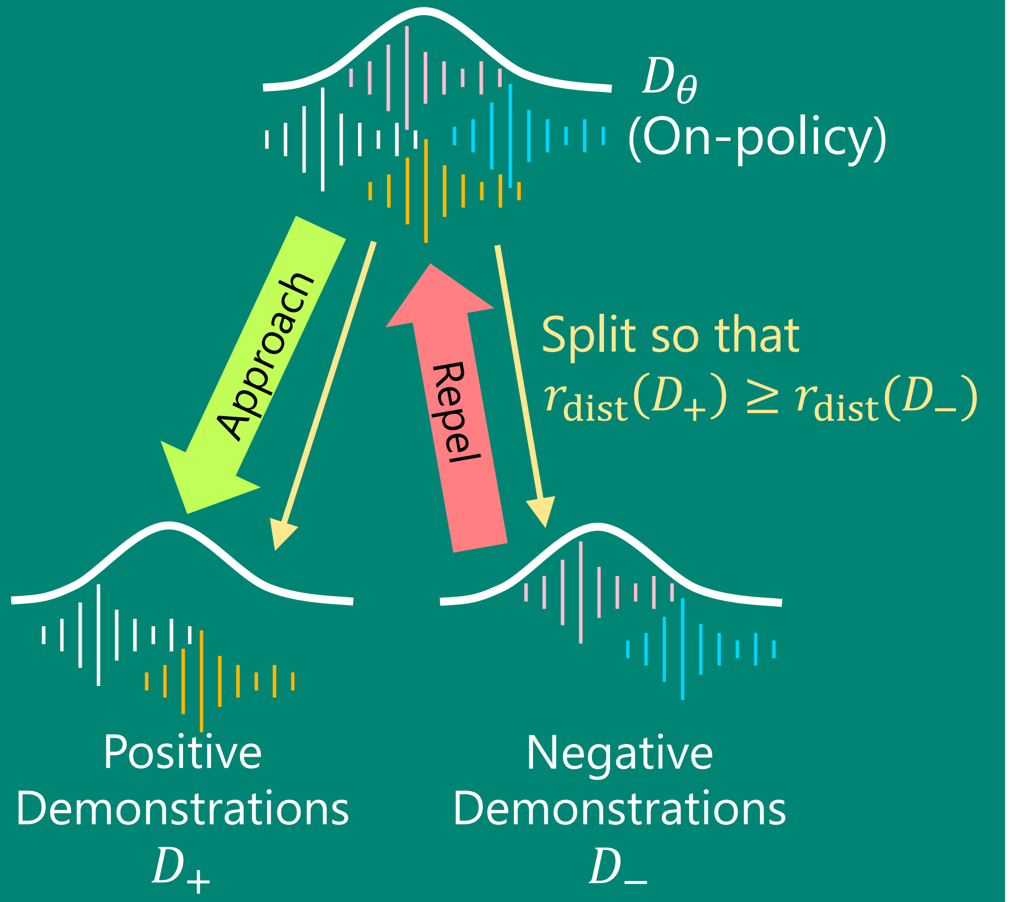

DRAGON Method

- **Goal:**
  - Maximize a reward function
    $$r_{dist}\colon \mathcal{P} \to \mathbb{R}$$
    that evaluates **distributions**.
  - Per-instance reward special case
    $r_{\text{dist}}(D_\theta) = \mathbb{E}_{X \sim D_\theta}\, r_{\text{instance}}(X).$

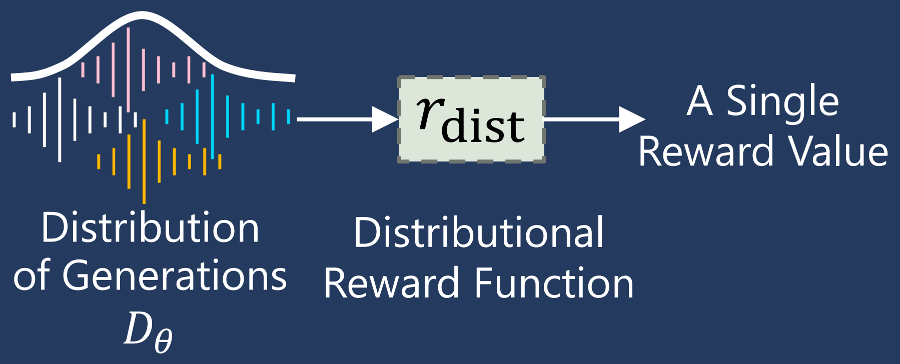

- **DRAGON:**

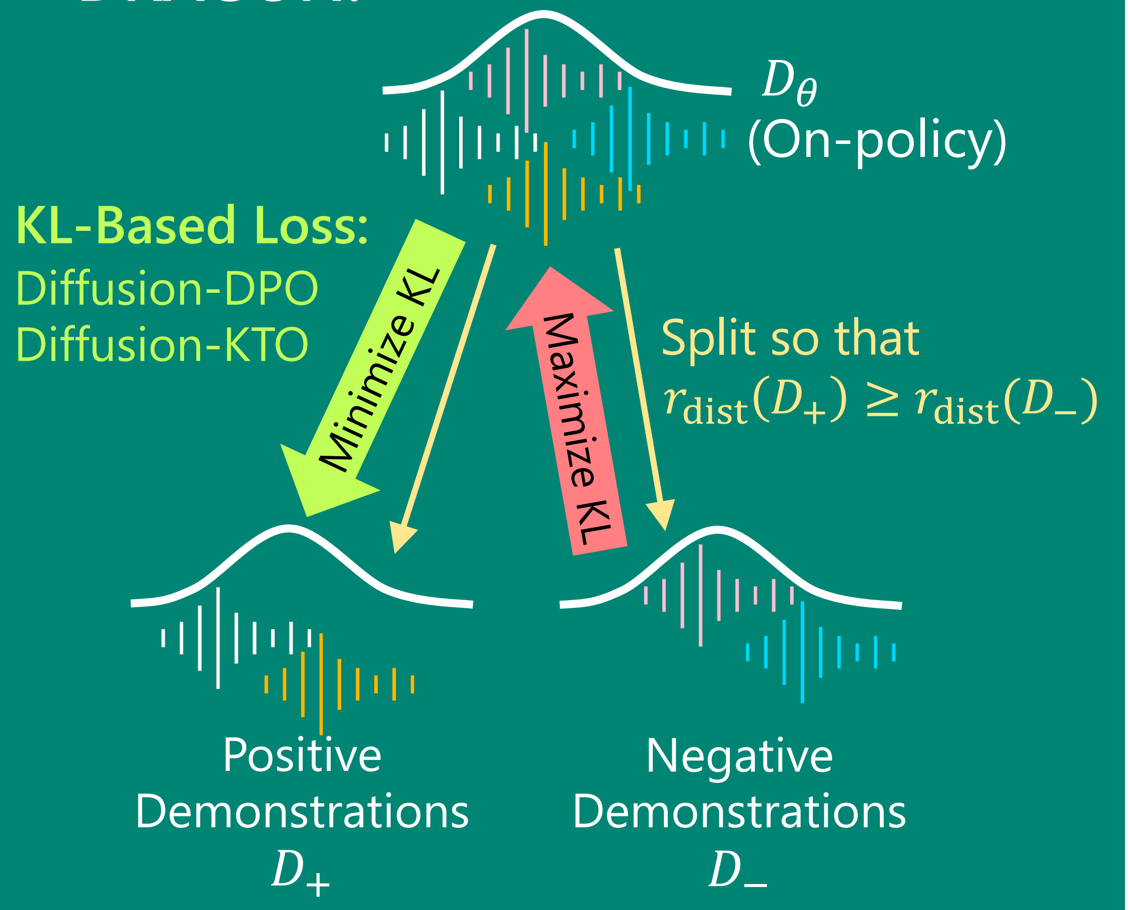

DRAGON Method

- **Goal:**
  - Maximize a reward function
    $$r_{dist}: \mathcal{P} \to \mathbb{R}$$
    that evaluates **distributions**.
  - Per-instance reward special case
    $r_{\text{dist}}(D_\theta) = \mathbb{E}_{X \sim D_\theta}\, r_{\text{instance}}(X)$.

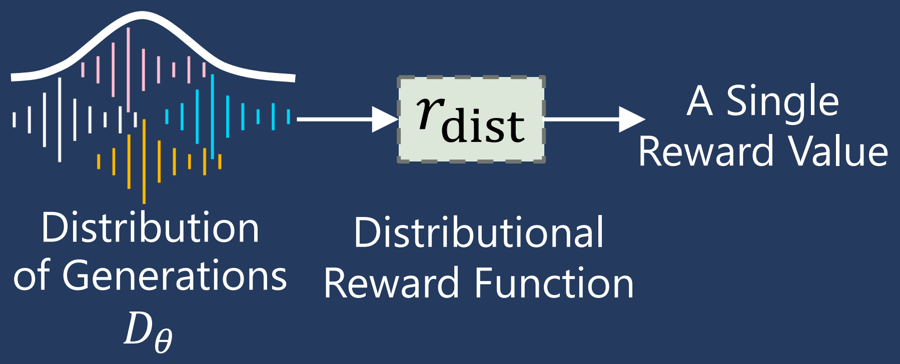

- **DRAGON:**

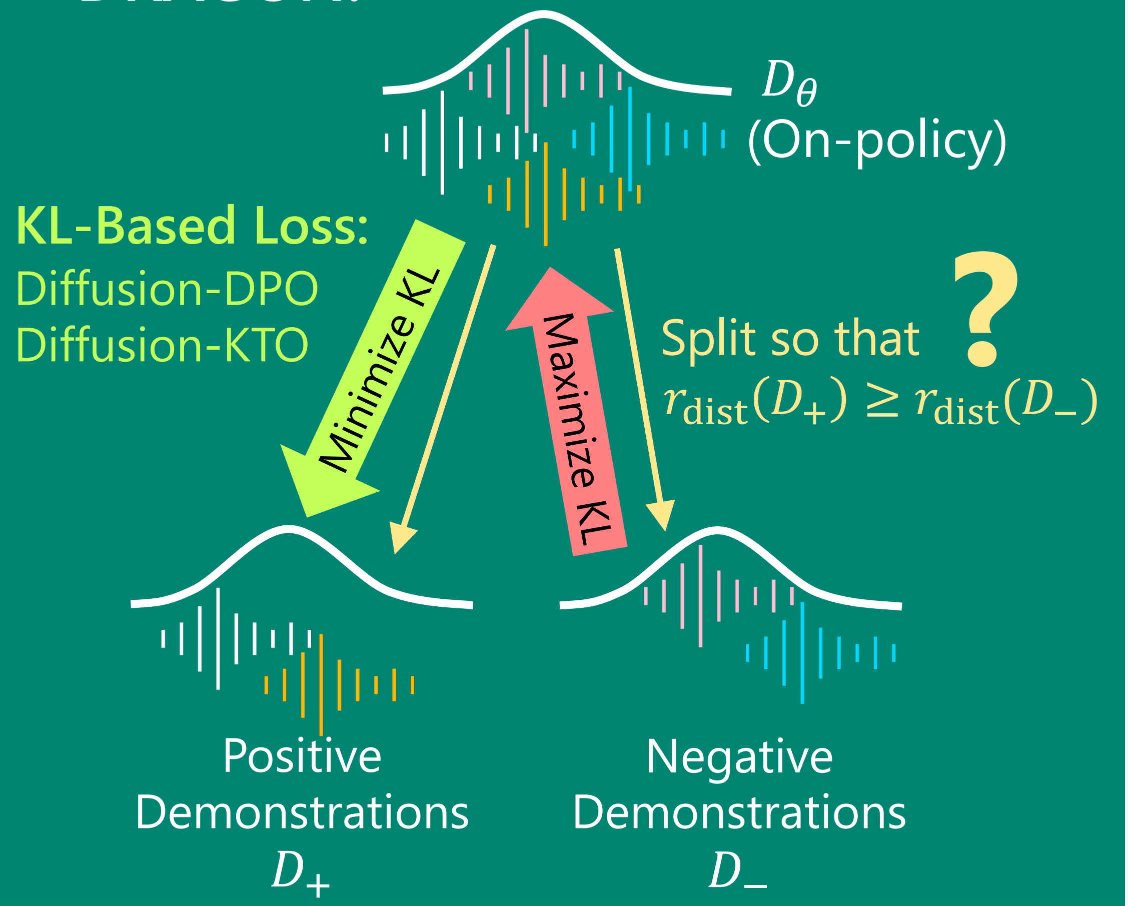

Splitting into D_+ and D_-

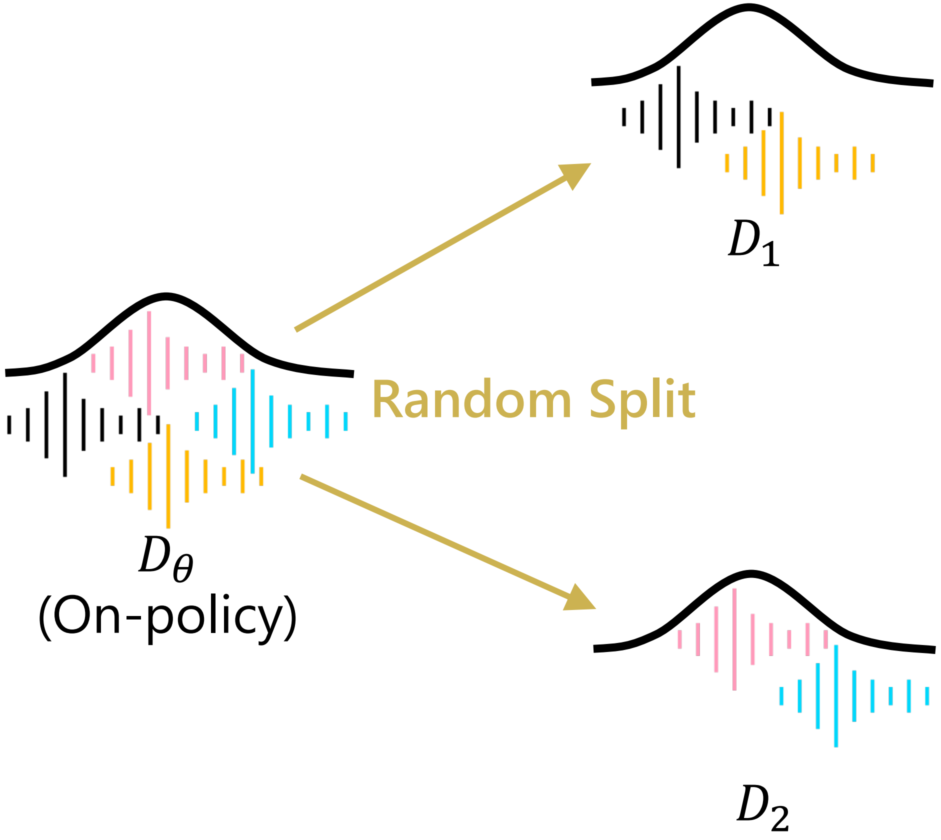

Splitting into D_+ and D_-

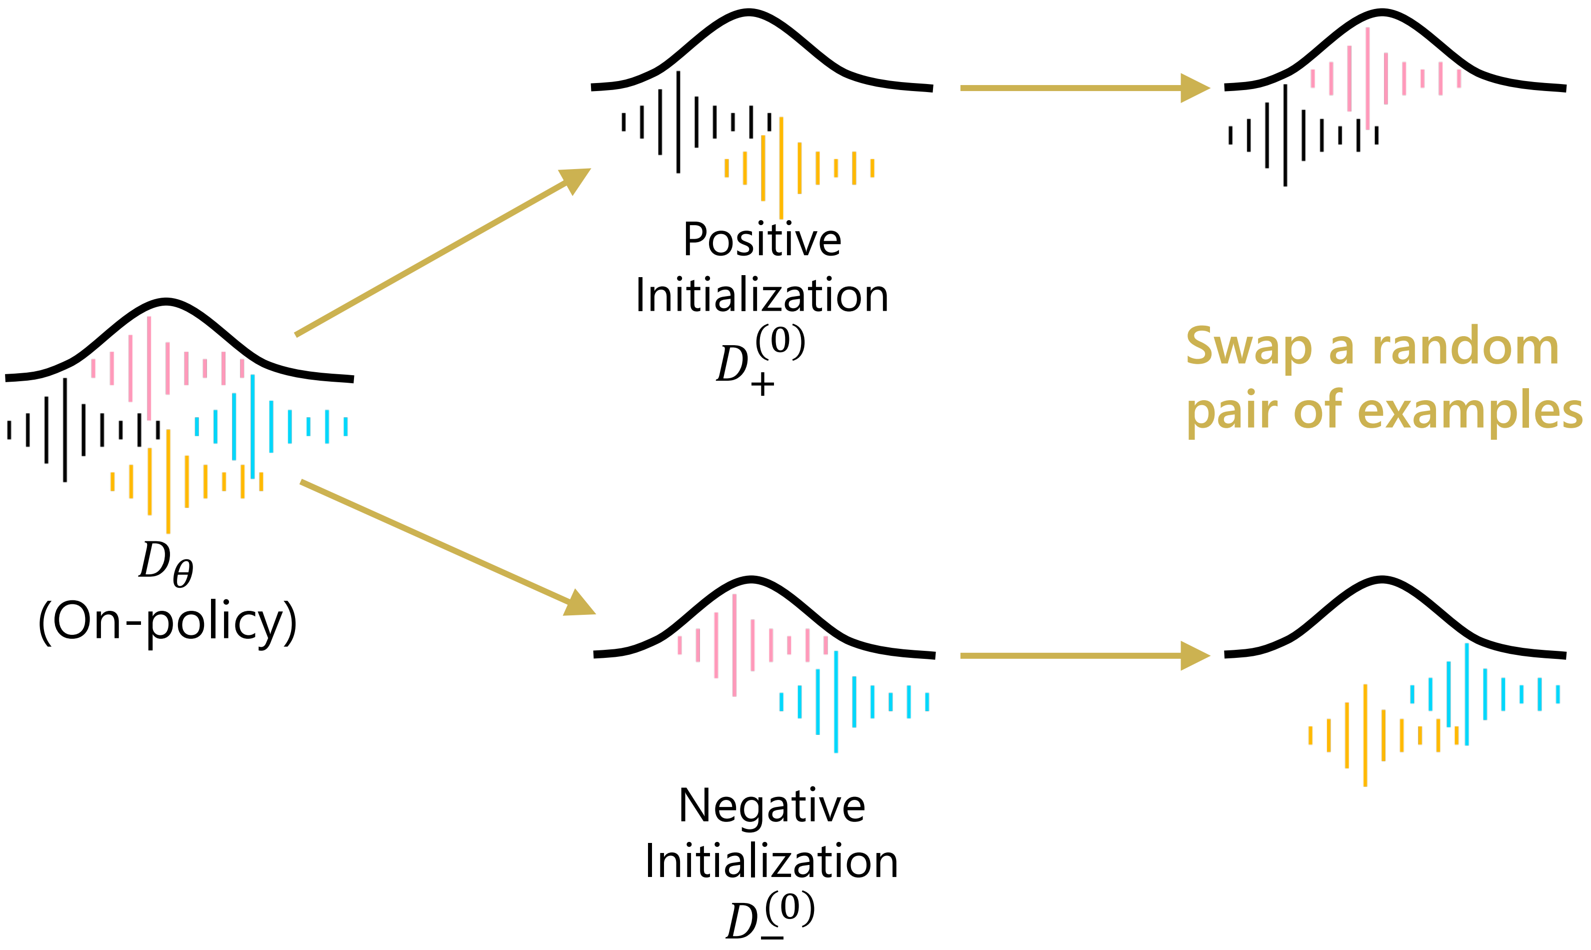

Splitting into D_+ and D_-

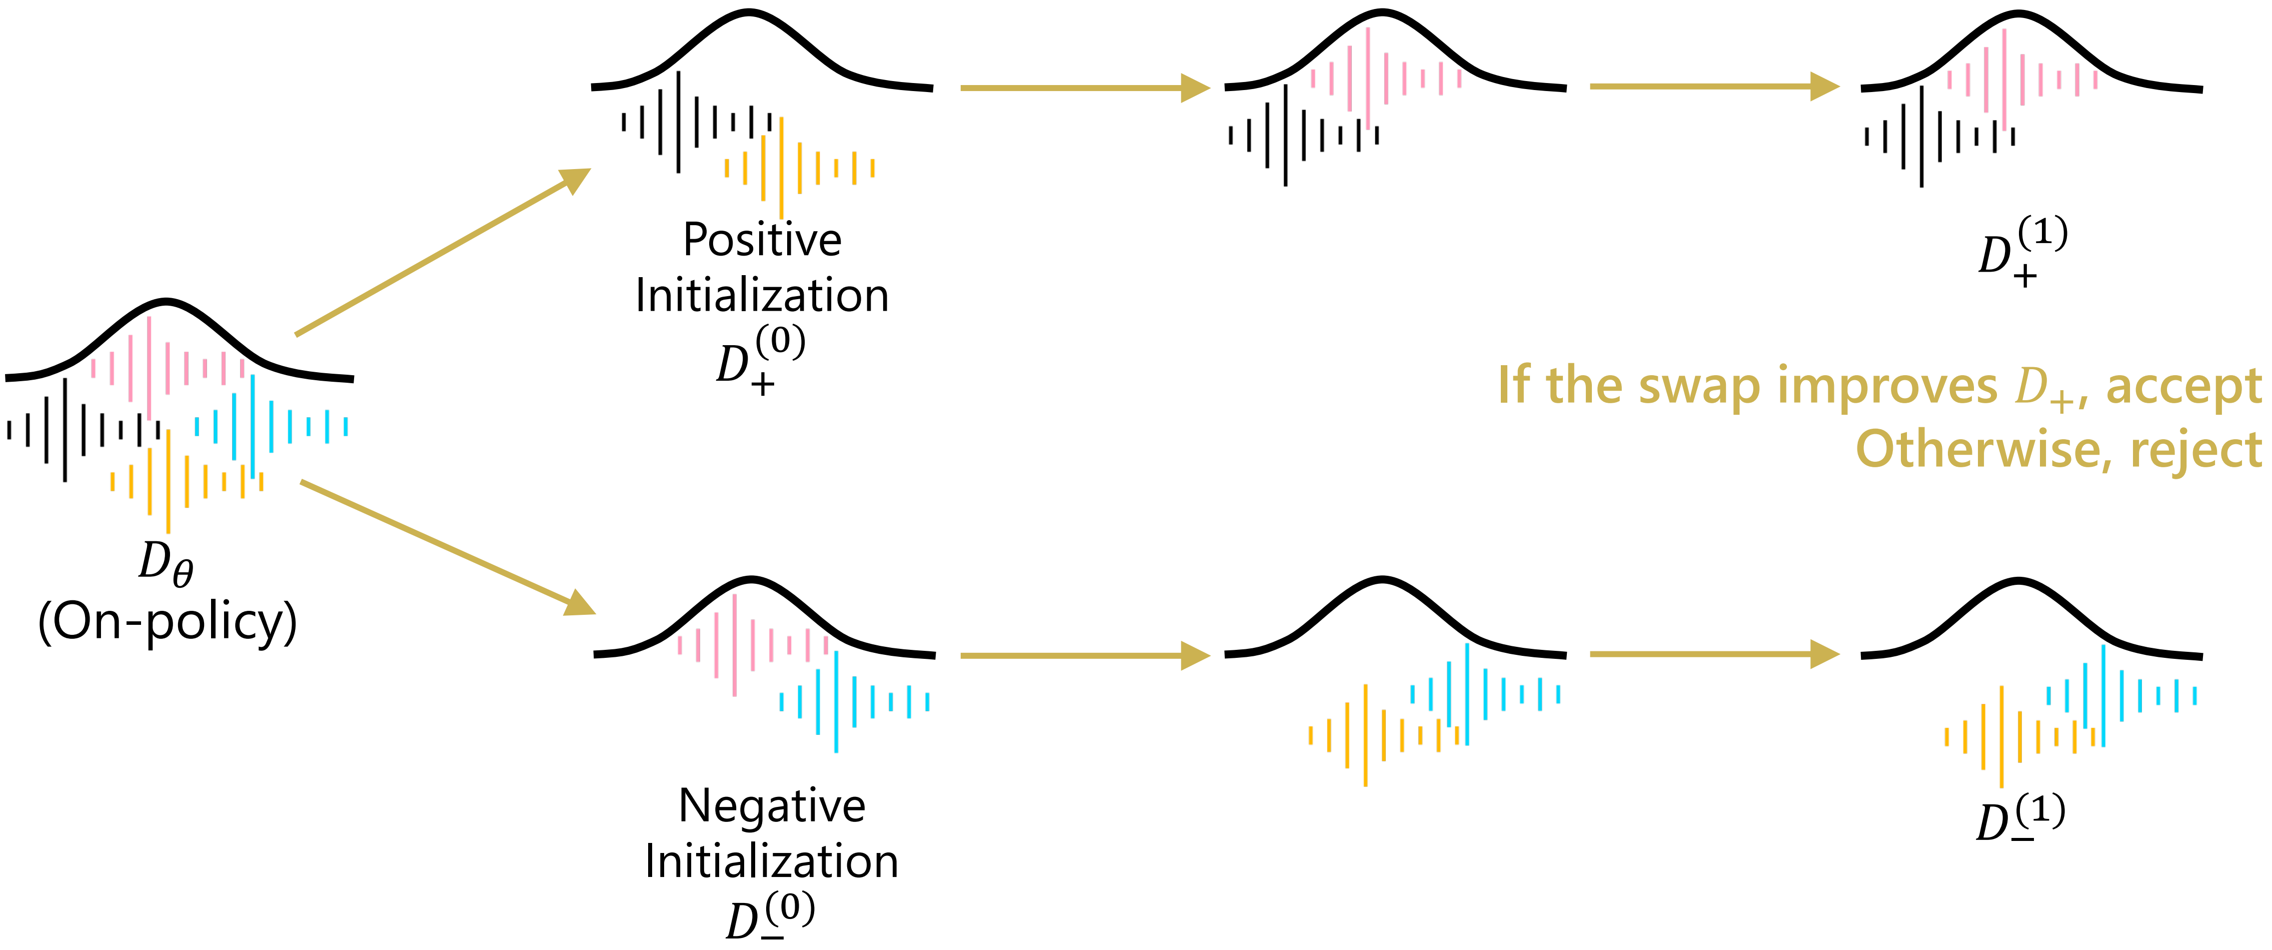

Splitting into D_+ and D_-

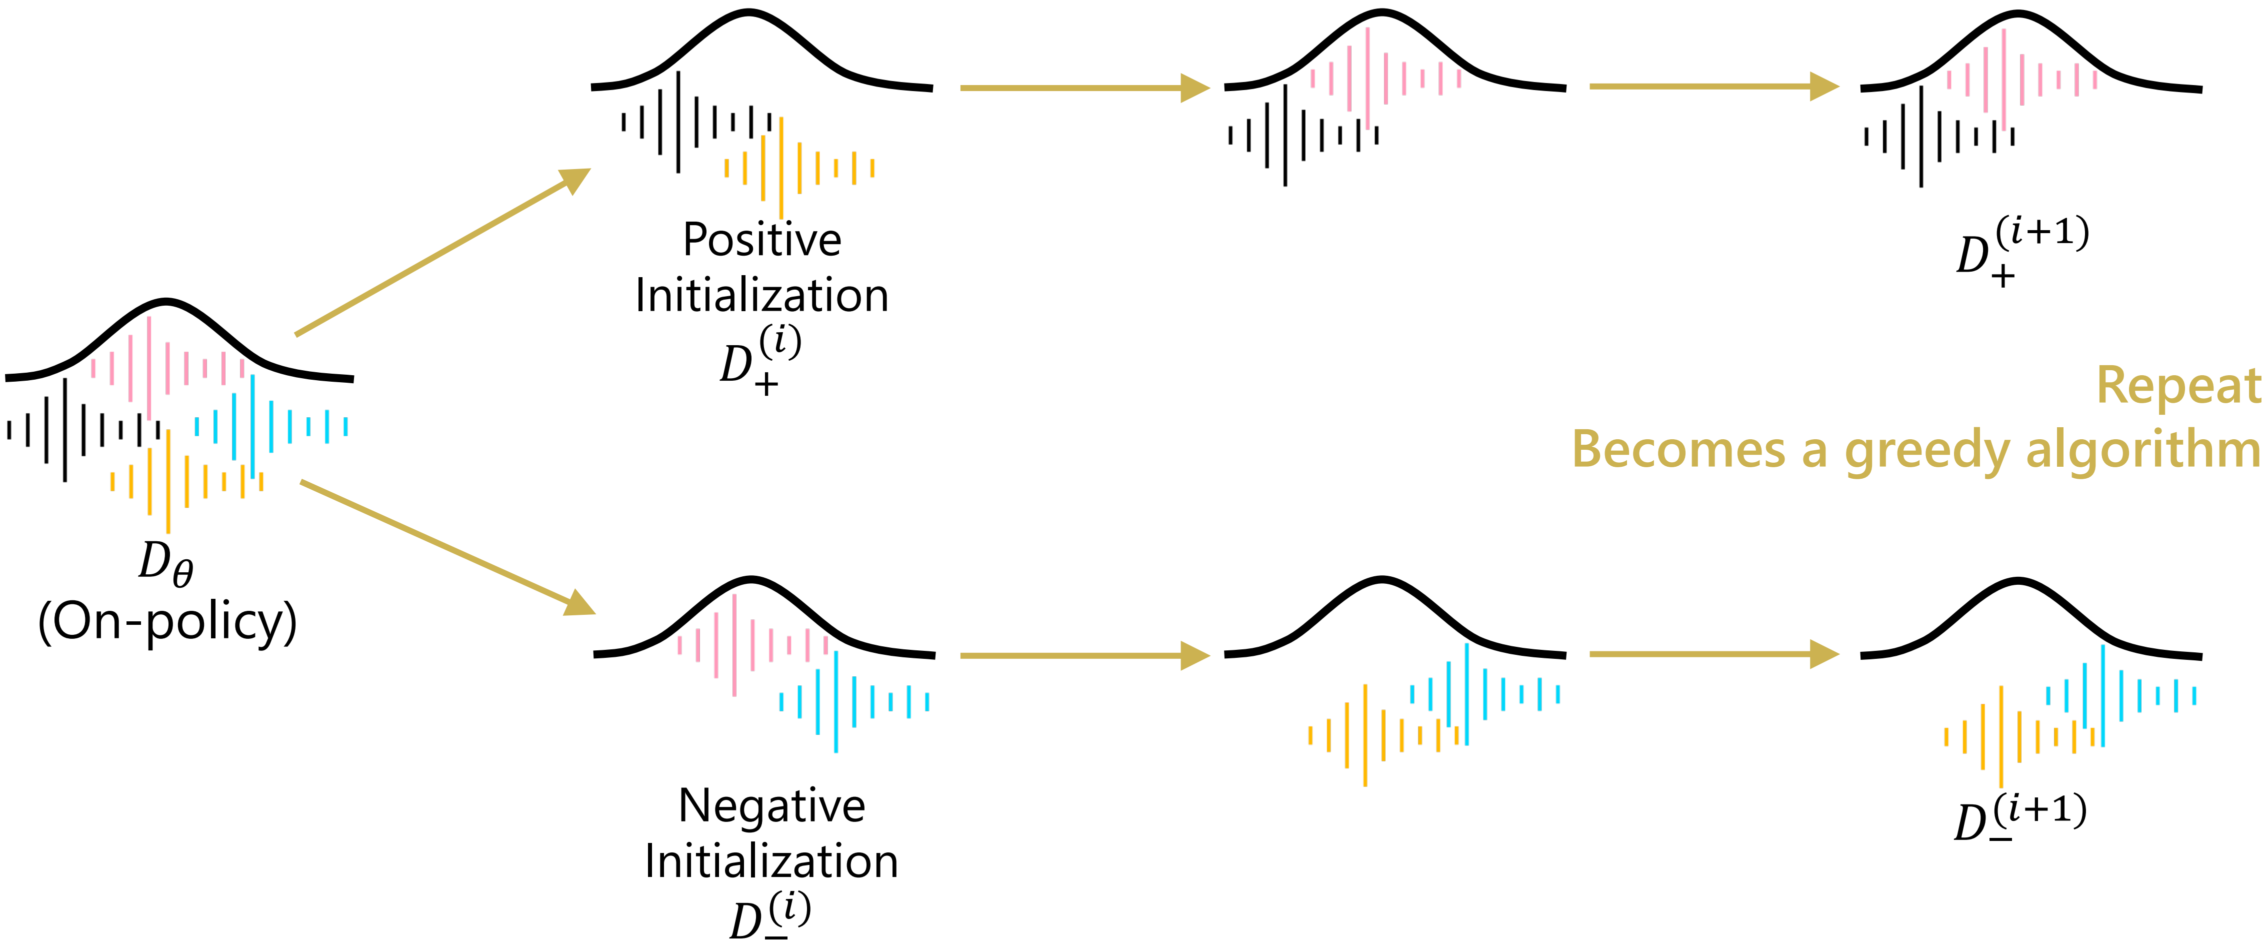

DRAGON Workflow

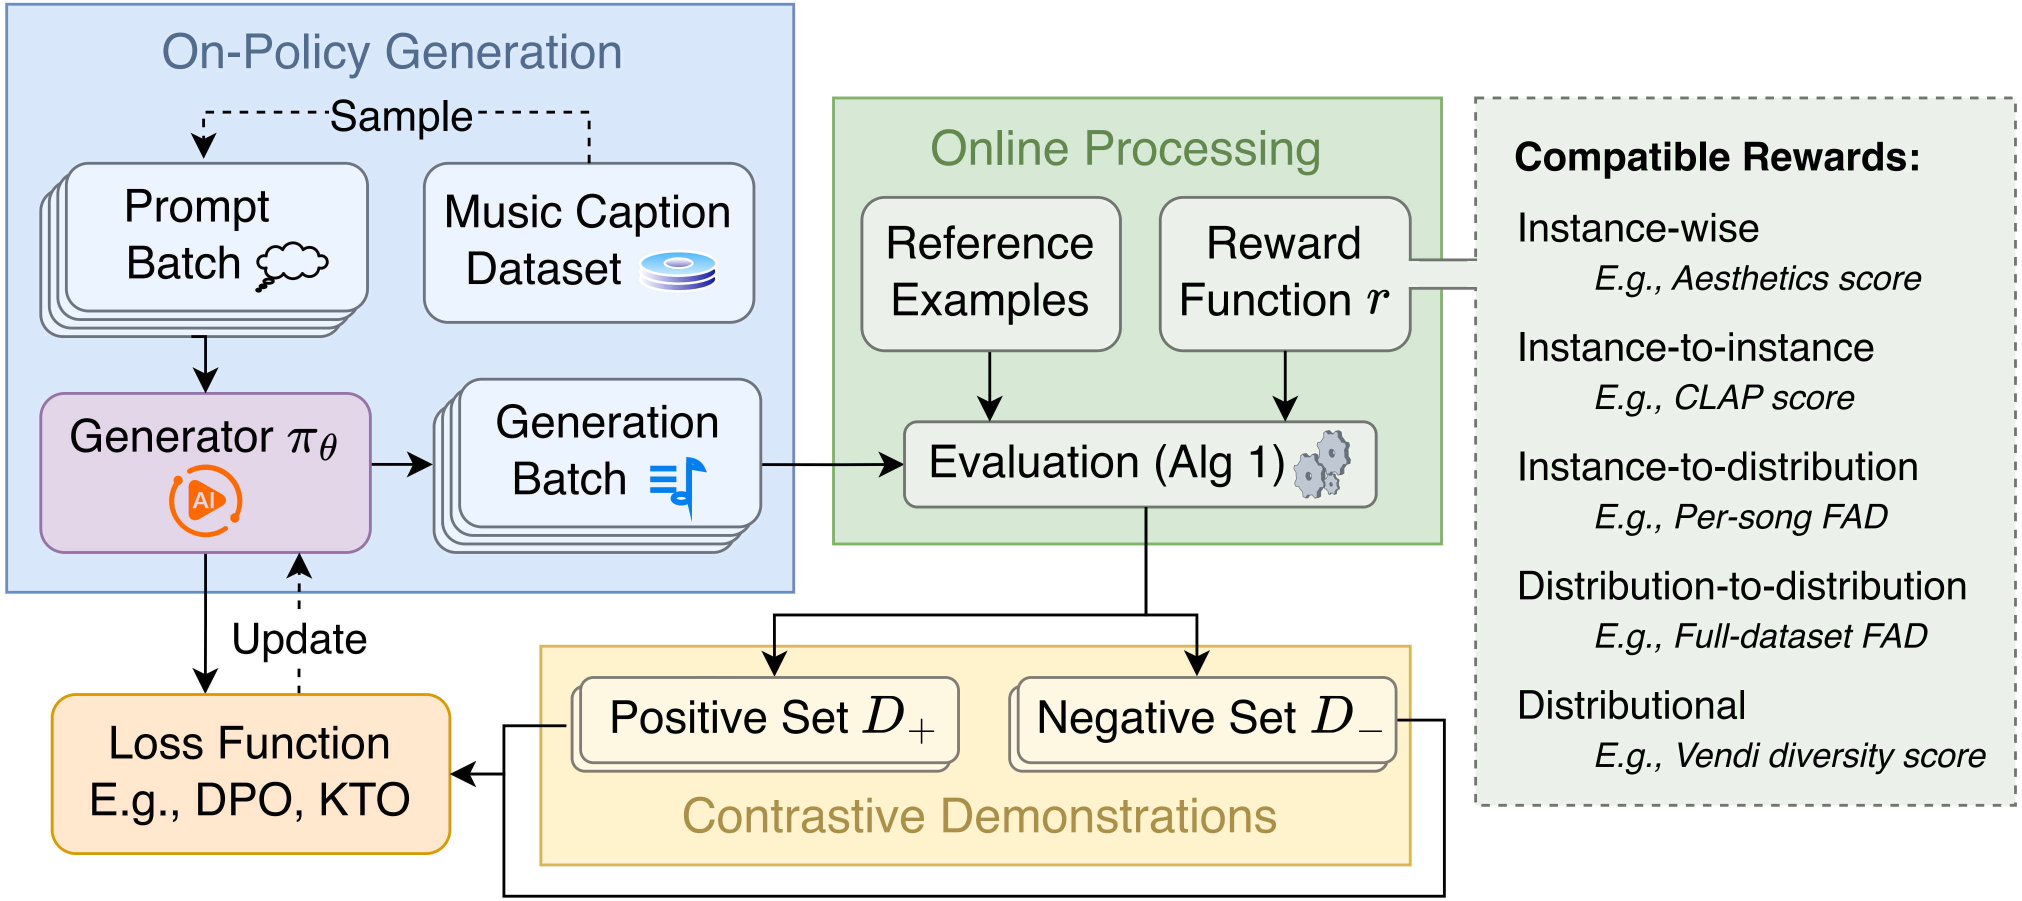

A New Way to Construct Rewards

- **Reward construction has been challenging for media generation.**
  - Media is perceptive. Hard to use criteria-based rewards like LLM alignment.
  - Hard to gather high-quality large-scale preference annotations.
- Leveraging DRAGON's versatility, **we construct exemplar-based rewards.**
  - Exemplars: A set of high-quality music embeddings.

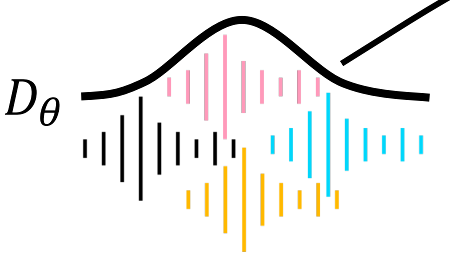

Match D_θ with exemplars

- Instance-to-instance (CLAP score)
- Distribution-to-distribution (FAD)
- Instance-to-distribution (Per-song FAD)

A New Way to Construct Rewards

- **Reward construction has been challenging for media generation.**
  - Media is perceptive. Hard to use criteria-based rewards like LLM alignment.
  - Hard to gather high-quality large-scale preference annotations.
- Leveraging DRAGON's versatility, **we construct exemplar-based rewards.**
  - Exemplars: A set of high-quality ***text*** (e.g., music captions, via cross-modal embedding spaces).

D_θ

Match D_θ with exemplars

- Instance-to-instance (CLAP score)
- Distribution-to-distribution (FAD)
- Instance-to-distribution (Per-song FAD)

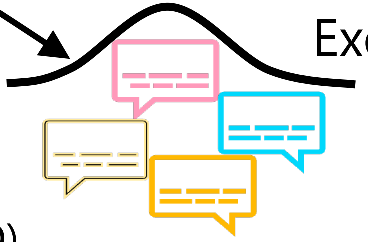

Exemplar Set

Main Experiment Result

Each vertex considers a reward metric and reports the win rate of the DRAGON model optimized for the metric.

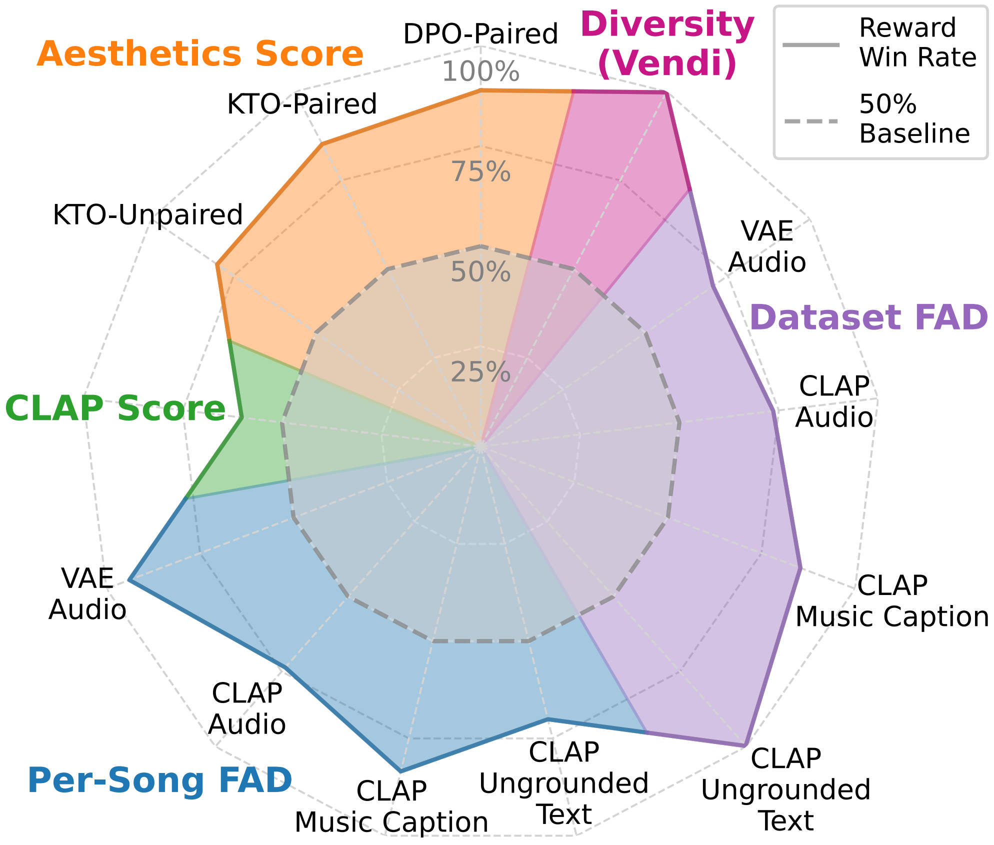

- **Experiment results on optimizing a text-to-music diffusion model.**
  - Over 20 reward functions, DRAGON achieves an ***81.45% win rate*** on average.

Human Listening Test

- **DRAGON-vs-Baseline binary comparison test.**
  - 21 raters, each rate 20 random blinded pairs (420 total).

- Via exemplar sets, DRAGON **improves human-perceived quality without human annotated preference dataset.**

Generation Examples

Electro dance song to play in the pub to cheer up the crowd.

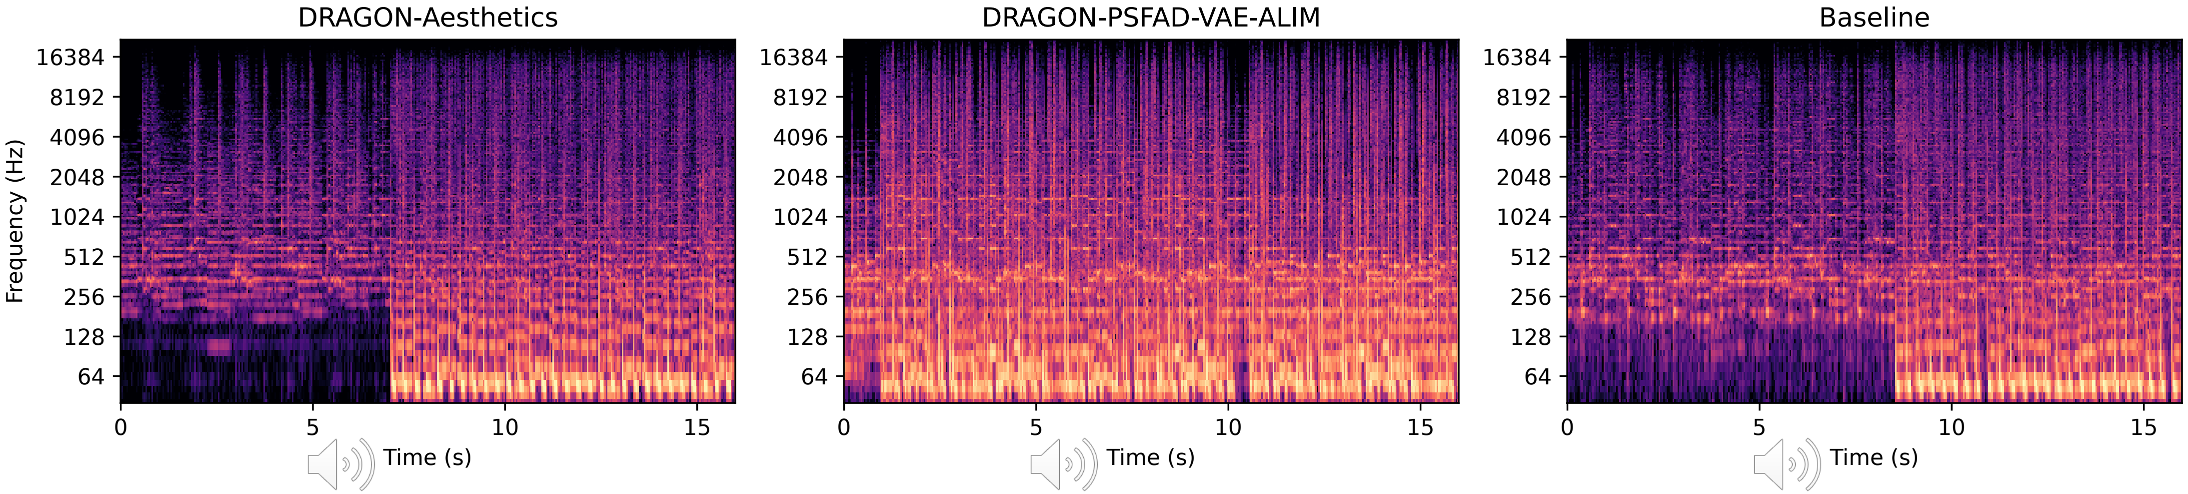

Generation Examples

Electro dance song to play in the pub to cheer up the crowd.

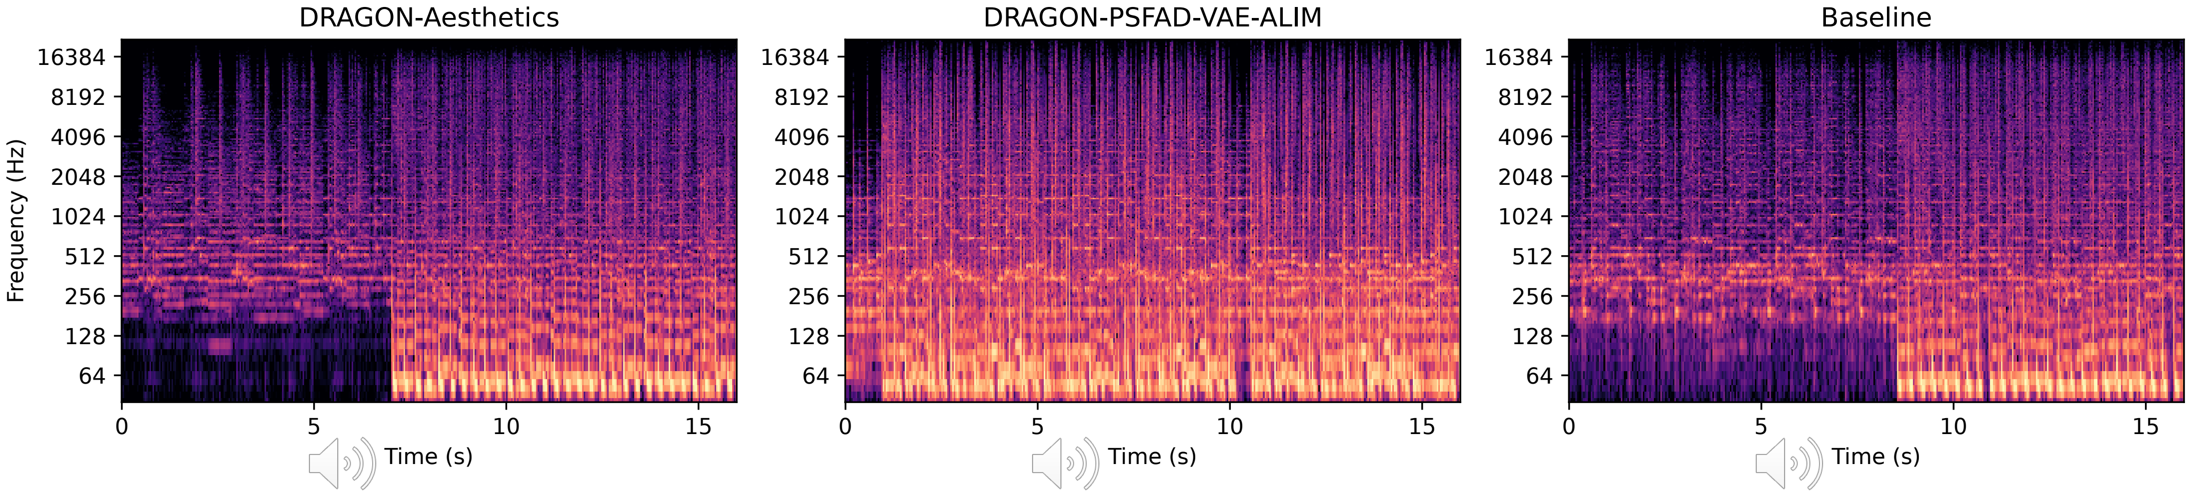

a show stopping broadway musical opening number

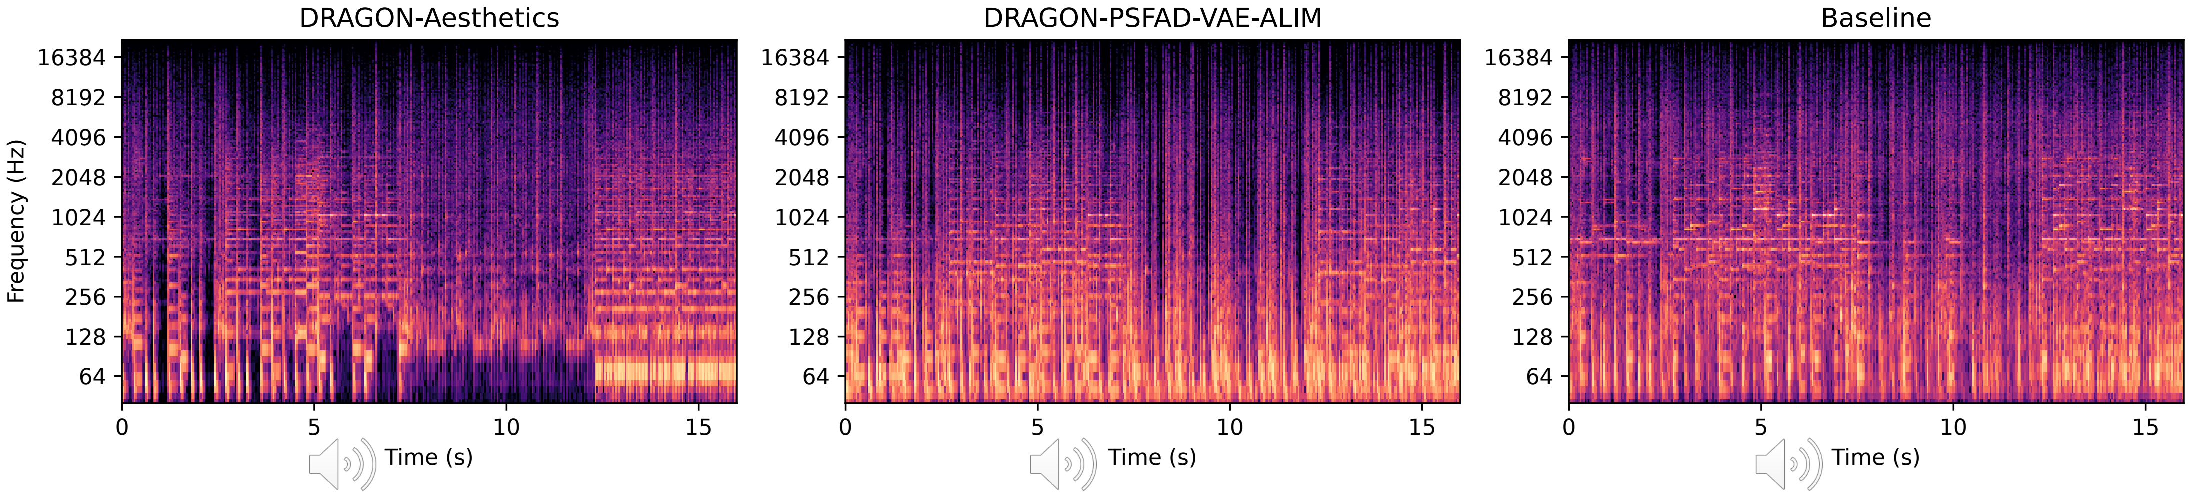

This Presentation

- An overview of my PhD research.
- **Efficient and reliable discriminative models** under input uncertainties.
  - Efficient Convex Optimization for Neural Network (Adversarial) Training.
  - Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
- **Efficient and reliable media generation** aligned with human preference.
  - ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation.
  - DRAGON: Optimizing Distributional Rewards Enhances Diffusion Models.
- **Summary.**

Summary

Diffusion Models – Audio/Music Generation

- Distillation/Acceleration
- Reward Optimization

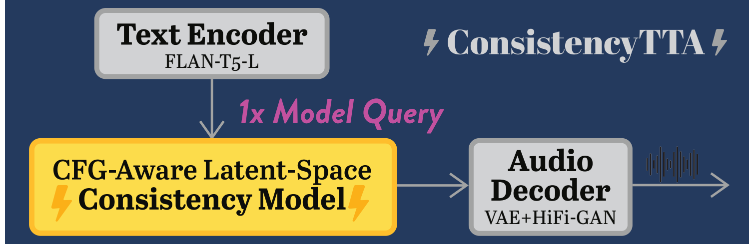

ML Safety – Adversarial Robustness

- Accuracy-Robustness Balance

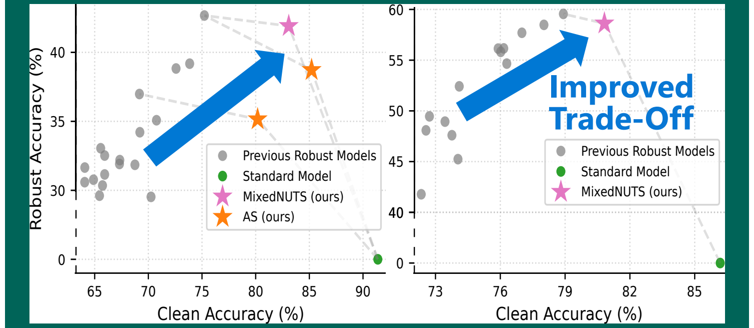

Convex Optimization for Training Neural Nets

- Convex Training
- Convex Adversarial Training

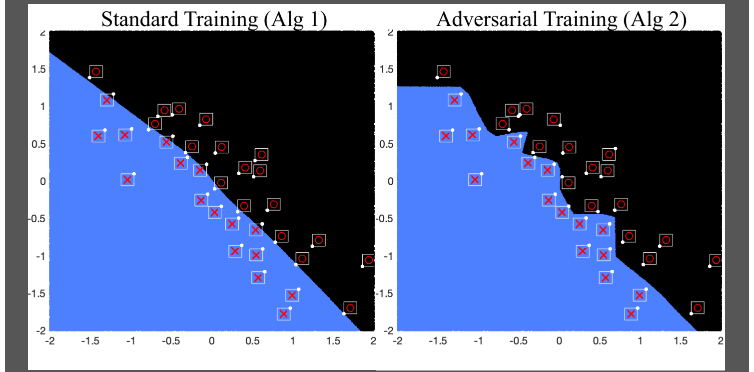

Summary

	Efficiency	Reliability
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers	• Training-free • Plug-and-play	• Interpretable formulation • Robust models are now practical
Diffusion Distillation		
Distributional Reward		

Summary

	Efficiency	Reliability
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers	• Training-free • Plug-and-play	• Interpretable formulation • Robust models are now practical
Diffusion Distillation	• 400x speedup	• End-to-end optimizes reward functions
Distributional Reward		

Summary

	Efficiency	Reliability
Convex Training	• Polynomial-time	• Global optimality guarantee • Robustness guarantees w/ adversarial training
Mixing Classifiers	• Training-free • Plug-and-play	• Interpretable formulation • Robust models are now practical
Diffusion Distillation	• 400x speedup	• End-to-end optimizes reward functions
Distributional Reward	• Exemplar-based reward	• Reward optimization on a distribution level • Address the training objective mismatch

Next Steps

- **Efficient and Reliable Optimization for Deep Learning and Media Generation** in an industry setting.
  - Distillation + reward optimization for diffusion models.
  - Adversarial attack and defense with generative models.
  - Optimizing more fine-grained rewards for media generation (e.g., text adherence).
- Research scientist at the music generation team of ByteDance.

Thanks to my collaborators and peers!

- Somayeh group:


Somayeh Sojoudi


Samuel Pfrommer


Brendon G Anderson


Tanmay Gautam


Jingqi Li


Ziye Ma


Elizabeth Glista


Eli Brock


Hyunin Lee


Yixiao Huang


Jiangyan Ma

- Other research collaborators:


Aerin Kim


Apaar Shanker

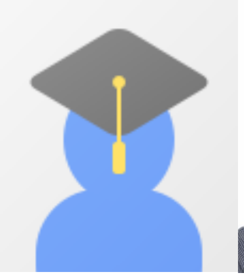
Yu Gai


Kazuhito Koishida


Dung Tran


Trung Dang


Mo Zhou


Vishal M Patel


Nicholas J Bryan


Jonah Casebeer

- Dissertation Committee:

Somayeh Sojoudi


Javad Lavaei


Kameshwar Poolla


- And many others!

Publications Presented

1. Practical Convex Formulation of Robust One-Hidden-Layer Neural Network Training.
   Yatong Bai, Tanmay Gautam, Yu Gai, Somayeh Sojoudi, in American Control Conference (ACC), 2022.

2. Efficient Global Optimization of Two-Layer ReLU Networks: Adversarial Training and Quadratic-time Algorithms.
   Yatong Bai, Tanmay Gautam, Somayeh Sojoudi, in *SIAM Journal on Mathematics of Data Science (SIMODS)* 5 (2), 446-474, 2023.

3. Mixing Classifiers to Alleviate the Accuracy-Robustness Trade-Off.
   Yatong Bai, Brendon G. Anderson, Somayeh Sojoudi, in *Annual Learning for Dynamics & Control Conference (L4DC)*, 2024.

4. Improving the Accuracy-Robustness Trade-Off of Classifiers via Adaptive Smoothing.
   Yatong Bai, Brendon G. Anderson, Aerin Kim, Somayeh Sojoudi, in *SIAM Journal on Mathematics of Data Science (SIMODS)* 6 (3), 2024.

5. MixedNUTS: Training-Free Accuracy-Robustness Balance via Nonlinearly Mixed Classifiers.
   Yatong Bai, Mo Zhou, Vishal M. Patel, Somayeh Sojoudi, in *Transactions on Machine Learning Research (TMLR)*, 2024.

6. ConsistencyTTA: Accelerating Diffusion-Based Text-to-Audio Generation with Consistency Distillation.
   Yatong Bai, Trung Dang, Dung Tran, Kazuhito Koishida, and Somayeh Sojoudi, in *INTERSPEECH*, 2024.

7. DRAGON: Distributional Rewards Optimize Diffusion Generative Models.
   Yatong Bai, Jonah Casebeer, Somayeh Sojoudi, Nicholas J. Bryan, under submission.