The power of two samples for Generative Adversarial Networks

Sewoong Oh

Department of Industrial and Enterprise Systems Engineering
University of Illinois at Urbana-Champaign

joint work with Giulia Fanti, Ashish Khetan, Zinan Lin

Generative models

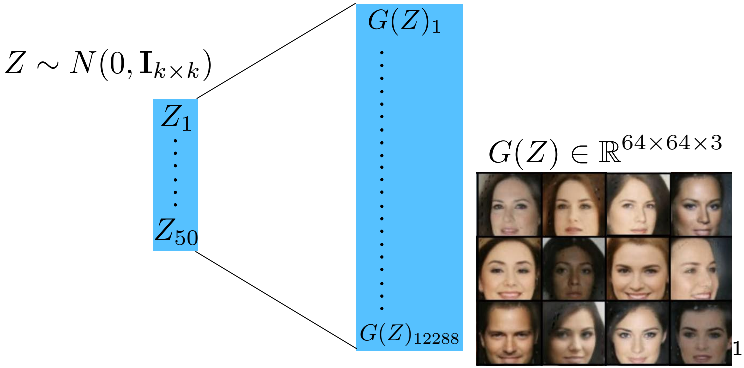

[1]

- A generative model is a black box that takes a random vector $Z \in \mathbb{R}^k$ and produces a sample vector $G(Z) \in \mathbb{R}^n$
- It is a differentiable function and stochastic gradient descent is used to train G

[1] ["Compressed Sensing using Generative Models", A Bora, A Jalal, E Price, AG Dimakis, 2017]

Generative models learn representation[2]

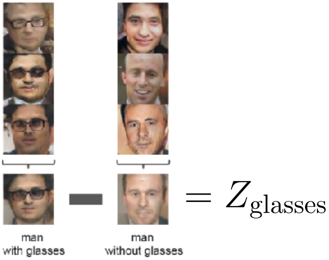

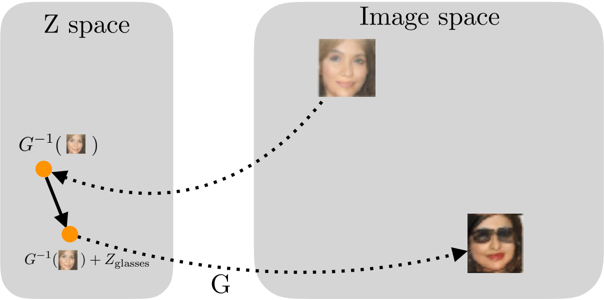

[2] [DCGAN, Radford et al. 2015]

Generative Adversarial Networks (GAN)

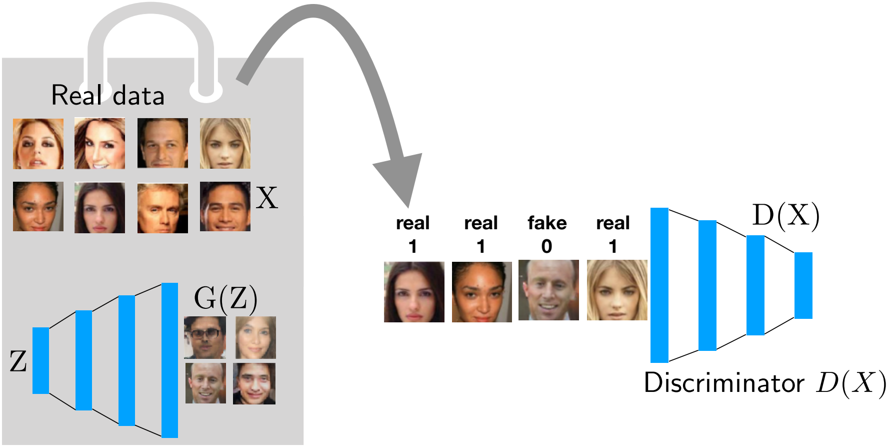

Generator $G(Z)$

$$\min_G \max_D V(G, D)$$

Mode collapse

Target distribution

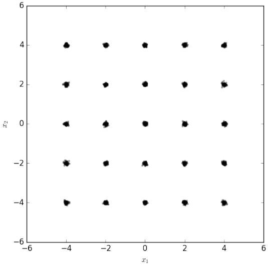

GAN

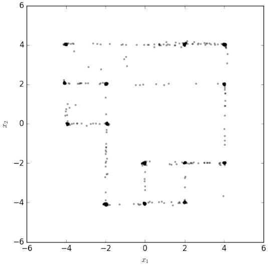

Mode collapse

- “A man in a orange jacket with sunglasses and a hat ski down a hill.”[3]

- “This guy is in black trunks and swimming underwater.”

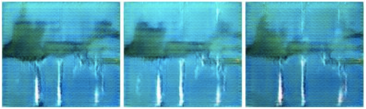

- “A tennis player in a blue polo shirt is looking down at the green court.”

[3] [“Generating interpretable images with controllable structure”, by Reed et al., 2016]

New framework: PacGAN

- lightweight overhead
- experimental results
- principled

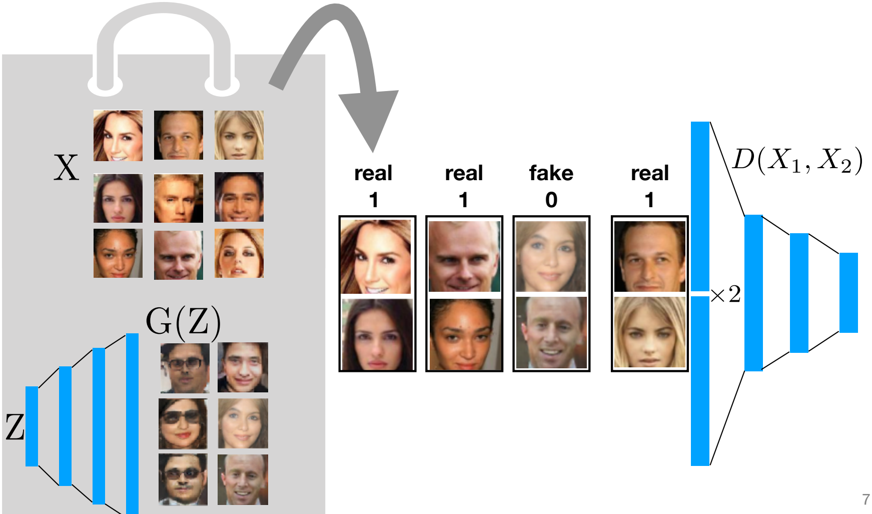

Benchmark datasetes from VEEGAN paper

Target GAN PacGAN2

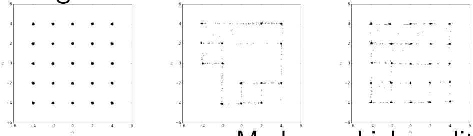

	Modes (Max 25)	high quality samples
GAN	3.3	0.5 %
ALI	15.8	1.6 %
Unrolled GAN	23.6	16.0 %
VEEGAN	24.6	40.0 %
PacGAN2	24.6±0.9	65.8±13.4 %
PacGAN3	24.9±0.3	71.4±13.8 %
PacGAN4	25.0±0.0	76.0±7.1 %

Benchmark datasetes from VEEGAN paper

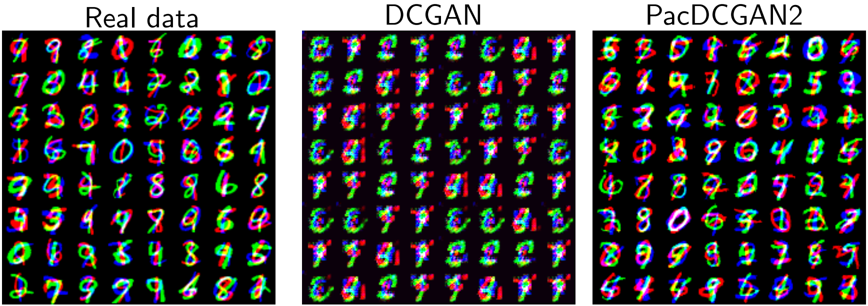

	Modes (Max 1000)	KL
DCGAN	99.0	3.40
ALI	16.0	5.40
Unrolled GAN	48.7	4.32
VEEGAN	150.0	2.95
PacDCGAN2	1000.0$\pm$0.0	0.06$\pm$0.01
PacDCGAN3	1000.0$\pm$0.0	0.06$\pm$0.01
PacDCGAN4	1000.0$\pm$0.0	0.07$\pm$0.01

Benchmark from Unrolled GAN (small discriminator)

Real data

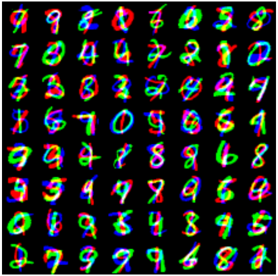

DCGAN

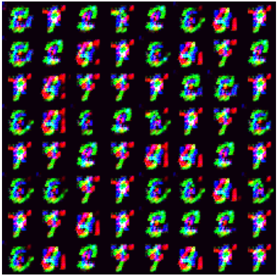

PacDCGAN2

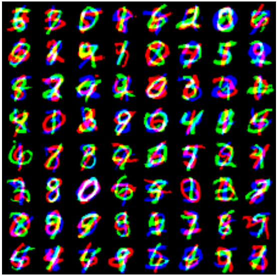

	Modes (Max 1000)	KL
DCGAN	30.6±20.73	5.99±0.42
Unrolled GAN, 1 step	65.4±34.75	5.91±0.14
Unrolled GAN, 5 steps	236.4±63.30	4.67±0.43
Unrolled GAN, 10 steps	327.2±74.67	4.66±0.46
PacDCGAN2	370.8±244.34	3.33±1.02
PacDCGAN3	534.3±103.68	2.11±0.52
PacDCGAN4	557.7±101.37	2.06±0.61

Intuition behind packing via toy example

Target distribution P

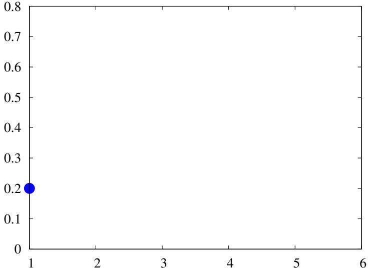

Generator Q_1
with mode collapse

Generator Q_2
without mode collapse

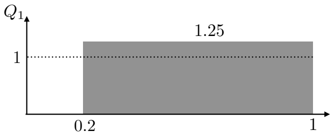

$d_{\mathrm{TV}}(P, Q_1) = 0.2$

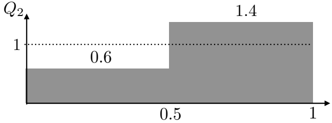

$d_{\mathrm{TV}}(P, Q_2) = 0.2$

Intuition behind packing via toy example

Target distribution P

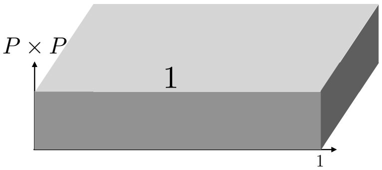

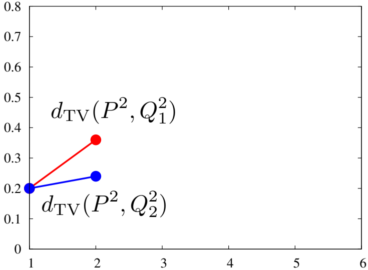

Generator Q_1
with mode collapse

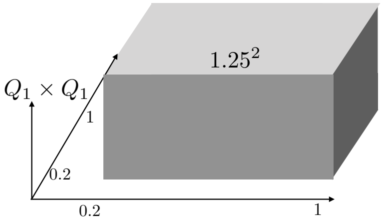

$d_{\mathrm{TV}}(P \times P, Q_1 \times Q_1) = 0.36$

Generator Q_2
without mode collapse

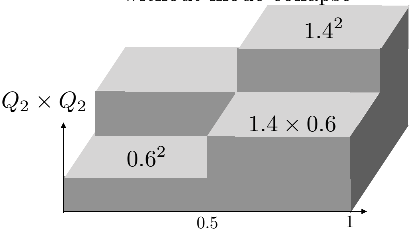

$d_{\mathrm{TV}}(P \times P, Q_2 \times Q_2) = 0.24$

Intuition behind packing via toy example

Target distribution P

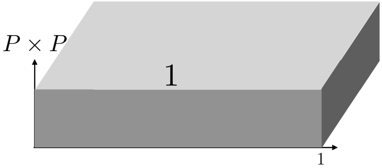

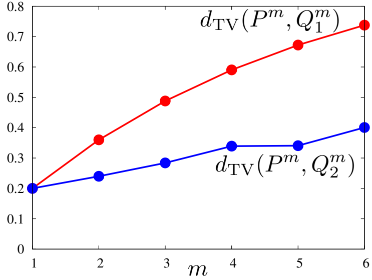

Generator Q_1
with mode collapse

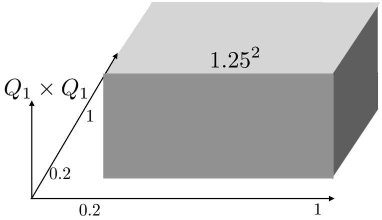

$d_{\mathrm{TV}}(P \times P, Q_1 \times Q_1) = 0.36$

Generator Q_2
without mode collapse

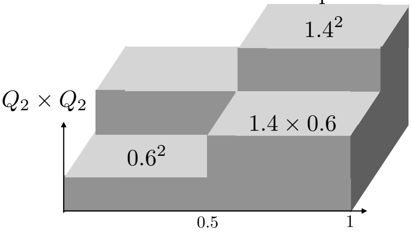

$d_{\mathrm{TV}}(P \times P, Q_2 \times Q_2) = 0.24$

Evolution of TV distances

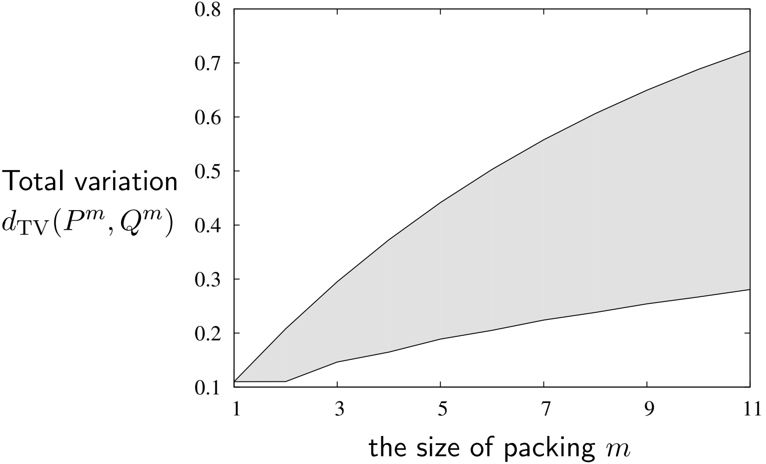

Evolution of TV distances through the prism of packing

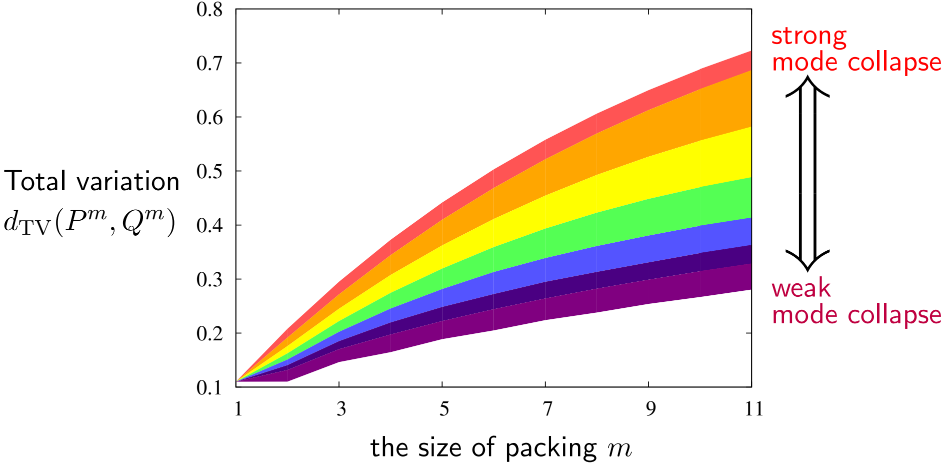

Through packing, the target-generator pairs are expanded over the strengths of the mode collapse

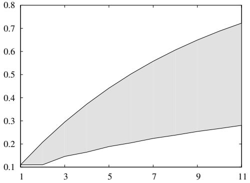

$$\max_{P,Q} / \min_{P,Q} \quad d_{\mathrm{TV}}(P^2, Q^2)$$
$$\text{subject to} \quad d_{\mathrm{TV}}(P, Q) = \tau$$

- we focus on $m = 2$ for this talk

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

Target distribution P

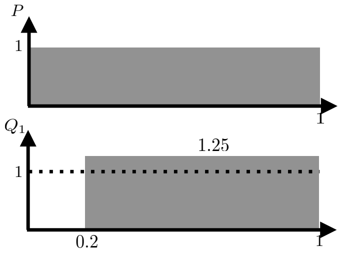

Generator Q_1
with mode collapse

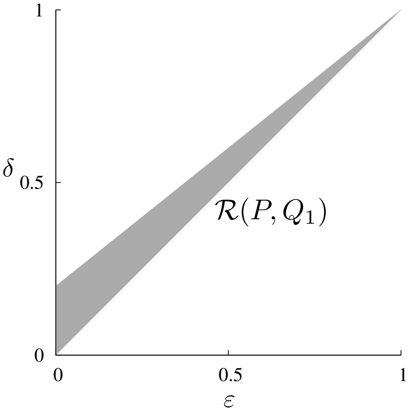

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

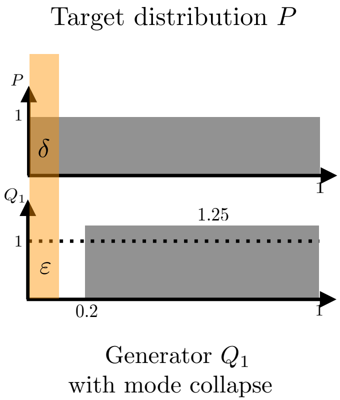

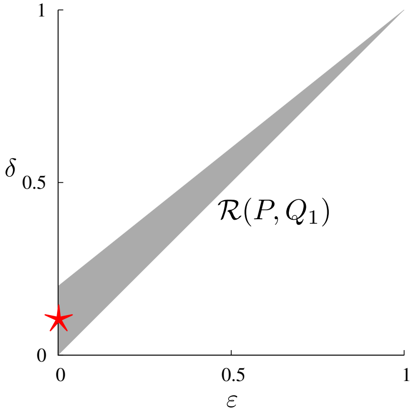

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \qquad \text{and} \qquad Q(S) \leq \varepsilon \ .$$

Target distribution P

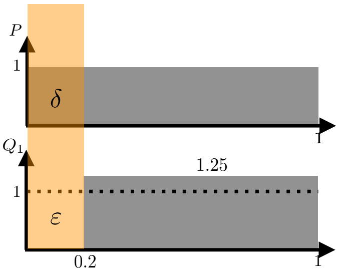

Generator Q_1
with mode collapse

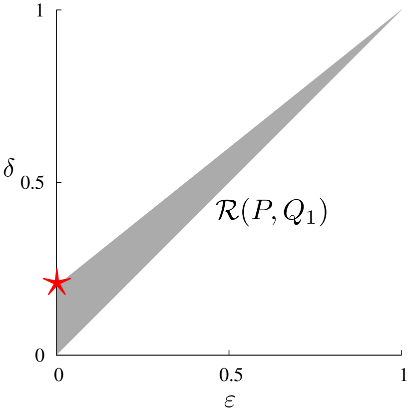

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

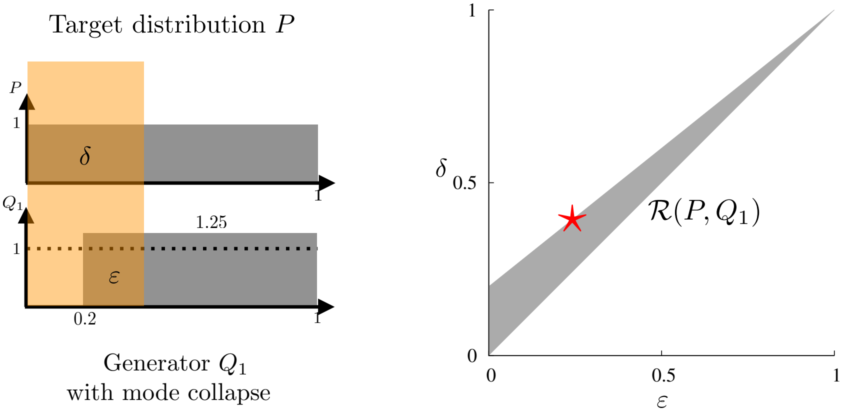

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

Target distribution P

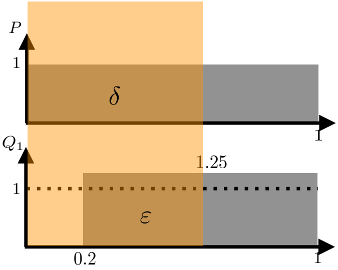

Generator Q_1
with mode collapse

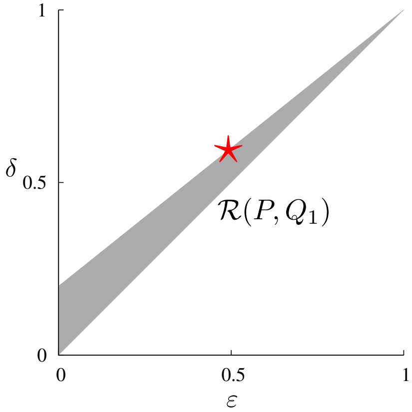

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \;, \quad \text{and} \quad Q(S) \leq \varepsilon \;.$$

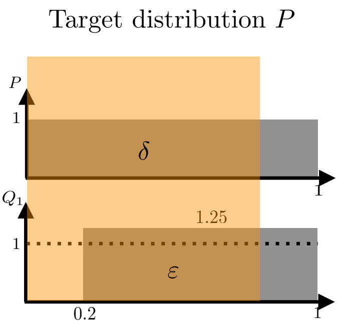

Generator Q_1
with mode collapse

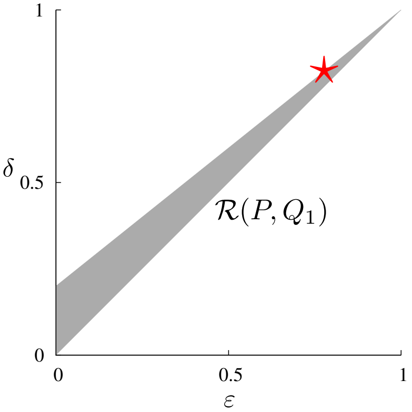

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \qquad \text{and} \qquad Q(S) \leq \varepsilon \ .$$

Target distribution P

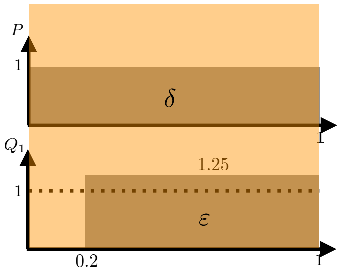

Generator Q_1
with mode collapse

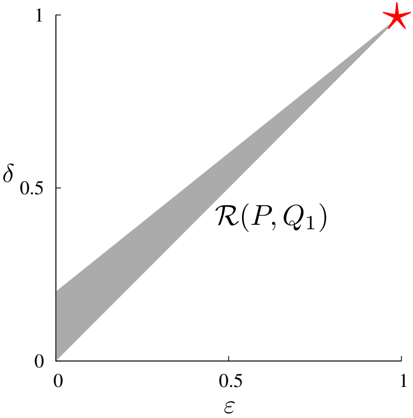

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

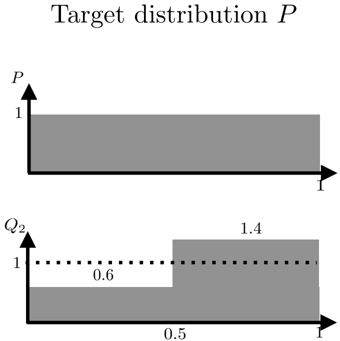

Generator Q_2
without mode collapse

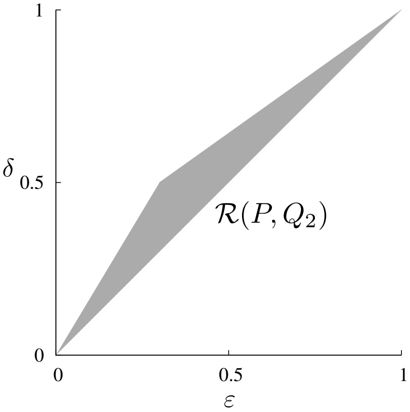

Intuition from Blackwell

Definition [mode collapse region]

We say a pair (P, Q) of a target distribution P and a generator distribution Q has (ε, δ)-**mode collapse** if there exists a set S such that

$$P(S) \geq \delta \ , \quad \text{and} \quad Q(S) \leq \varepsilon \ .$$

Target distribution P

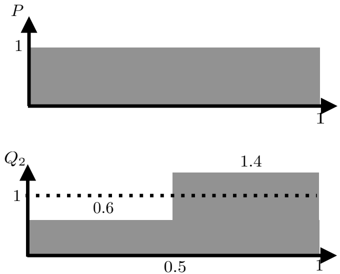

Generator Q_2
without mode collapse

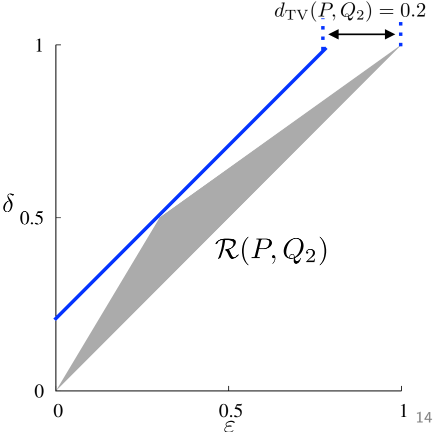

Upper bound

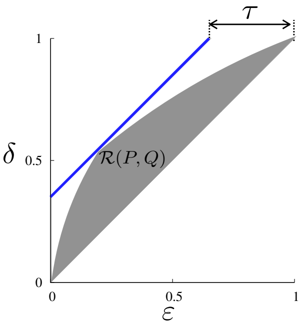

$$\begin{aligned} \max_{P,Q} \quad & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} \quad & d_{\mathrm{TV}}(P, Q) = \tau \end{aligned}$$

Upper bound

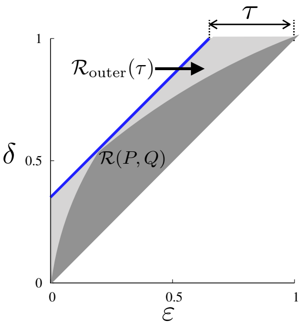

$$\begin{array}{ll} \max\limits_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\mathrm{TV}}(P,Q) = \tau \end{array}$$

$$\mathcal{R}(P,Q) \subseteq \mathcal{R}_{\mathrm{outer}}(\tau)$$

Upper bound

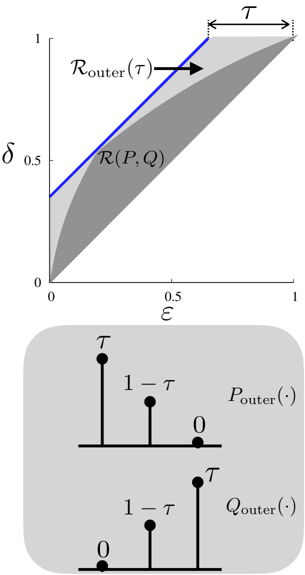

$$\max_{P,Q} \quad d_{\mathrm{TV}}(P^2, Q^2)$$

$$\text{subject to} \quad d_{\mathrm{TV}}(P, Q) = \tau$$

$$\mathcal{R}(P,Q) \subseteq \mathcal{R}_{\text{outer}}(\tau)$$

Upper bound

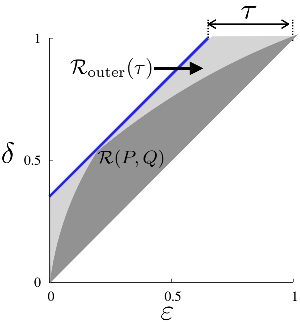

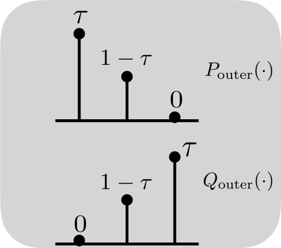

$$\begin{array}{ll} \max\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P,Q) = \tau \end{array}$$

$$\begin{array}{rcl} \mathcal{R}(P,Q) & \subseteq & \mathcal{R}_{\text{outer}}(\tau) \\ \mathcal{R}(P^2,Q^2) & \subseteq & \mathcal{R}(P_{\text{outer}}^2, Q_{\text{outer}}^2) \end{array}$$

Blackwell's theorem

$$\begin{array}{rl} & \mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q') \\ \Rightarrow & \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2) \end{array}$$

Upper bound

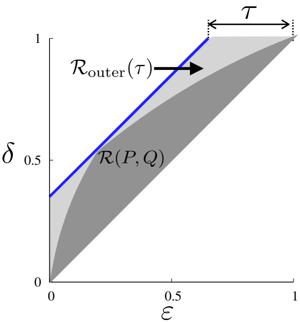

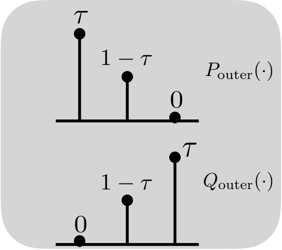

$$\max_{P,Q} \quad d_{\mathrm{TV}}(P^2, Q^2)$$

$$\text{subject to} \quad d_{\mathrm{TV}}(P, Q) = \tau$$

$$\begin{aligned}
\mathcal{R}(P,Q) &\subseteq \mathcal{R}_{\mathrm{outer}}(\tau) \\
\mathcal{R}(P^2,Q^2) &\subseteq \mathcal{R}(P^2_{\mathrm{outer}}, Q^2_{\mathrm{outer}}) \\
d_{\mathrm{TV}}(P^2,Q^2) &\leq \underbrace{d_{\mathrm{TV}}(P^2_{\mathrm{outer}}, Q^2_{\mathrm{outer}})}_{1-(1-\tau)^2}
\end{aligned}$$

Blackwell's theorem

$$\begin{aligned}
& \mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q') \\
\Rightarrow \quad & \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2)
\end{aligned}$$

Lower bound

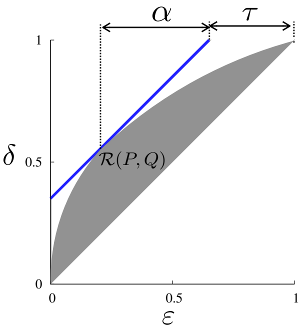

$$\begin{array}{ll} \min_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\mathrm{TV}}(P,Q) = \tau \end{array}$$

Lower bound

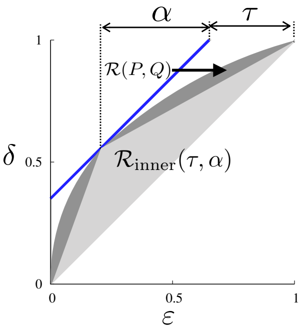

$$\begin{array}{ll} \min\limits_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\mathrm{TV}}(P,Q) = \tau \end{array}$$

$$\mathcal{R}_{\text{inner}}(\tau,\alpha) \quad \subseteq \quad \mathcal{R}(P,Q)$$

Lower bound

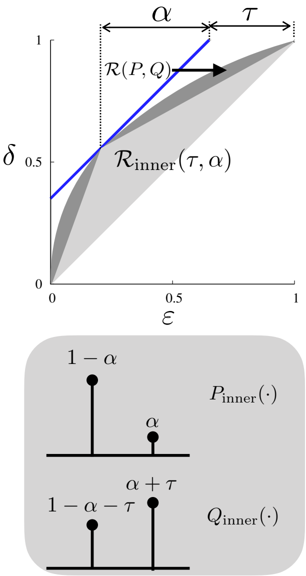

$$\begin{array}{ll} \min\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P,Q) = \tau \end{array}$$

$$\mathcal{R}_{\text{inner}}(\tau,\alpha) \quad \subseteq \quad \mathcal{R}(P,Q)$$

Lower bound

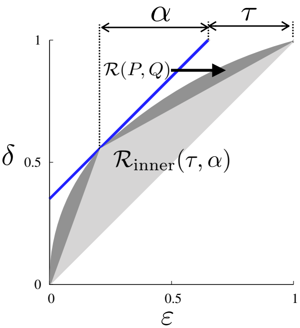

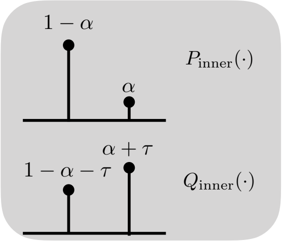

$$\begin{aligned} \min_{P,Q} \quad & d_{\text{TV}}(P^2,Q^2) \\ \text{subject to} \quad & d_{\text{TV}}(P,Q)=\tau \end{aligned}$$

$$\begin{aligned} \mathcal{R}_{\text{inner}}(\tau,\alpha) \quad &\subseteq \quad \mathcal{R}(P,Q) \\ \mathcal{R}(P^2_{\text{inner}},Q^2_{\text{inner}}) \quad &\subseteq \quad \mathcal{R}(P^2,Q^2) \end{aligned}$$

Blackwell's theorem

$$\begin{aligned} & \mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q') \\ \Rightarrow \quad & \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2) \end{aligned}$$

Lower bound

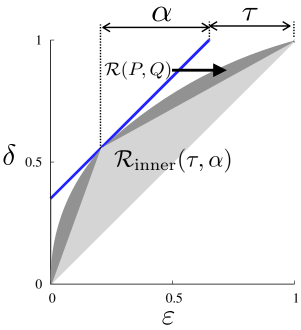

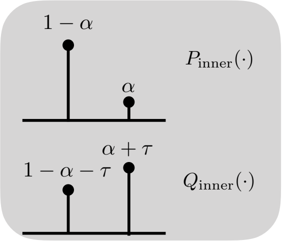

$$\begin{aligned} \min_{P,Q} \quad & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} \quad & d_{\text{TV}}(P, Q) = \tau \end{aligned}$$

$$\begin{aligned} \mathcal{R}_{\text{inner}}(\tau,\alpha) &\subseteq \mathcal{R}(P,Q) \\ \mathcal{R}(P^2_{\text{inner}}, Q^2_{\text{inner}}) &\subseteq \mathcal{R}(P^2,Q^2) \\ \underbrace{d_{\text{TV}}(P^2_{\text{inner}}, Q^2_{\text{inner}})}_{\tau} &\leq d_{\text{TV}}(P^2,Q^2) \end{aligned}$$

Blackwell's theorem

$$\begin{aligned} & \mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q') \\ \Rightarrow \quad & \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2) \end{aligned}$$

$d_{\mathrm{TV}}(P^m, Q^m)$

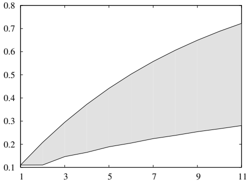

m

$$\begin{array}{ll} \max_{P,Q} / \min_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\mathrm{TV}}(P, Q) = \tau \end{array}$$

$d_{\mathrm{TV}}(P^m, Q^m)$

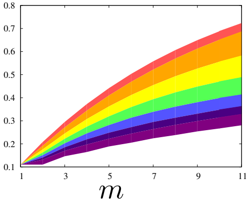

$$\max_{P,Q} \quad d_{\mathrm{TV}}(P^2, Q^2)$$

subject to $\quad d_{\mathrm{TV}}(P, Q) = \tau$

no $(\varepsilon_0, \delta_0)$-mode collapse

Upper bound without $(\varepsilon_0, \delta_0)$-mode collapse

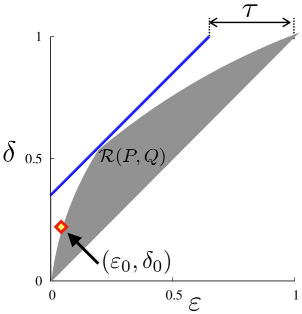

$$\begin{array}{rl} \max\limits_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\mathrm{TV}}(P,Q) = \tau \\ & \text{no } (\varepsilon_0, \delta_0)\text{-mode collapse} \end{array}$$

Upper bound without $(\varepsilon_0, \delta_0)$-mode collapse

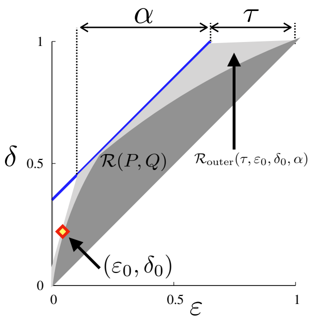

$$\max_{P,Q} \quad d_{\text{TV}}(P^2, Q^2)$$

subject to $\quad d_{\text{TV}}(P, Q) = \tau$

no $(\varepsilon_0, \delta_0)$-mode collapse

$$\mathcal{R}(P,Q) \subseteq \mathcal{R}_{\text{outer}}(\tau, \varepsilon_0, \delta_0, \alpha)$$

Upper bound without $(\varepsilon_0, \delta_0)$-mode collapse

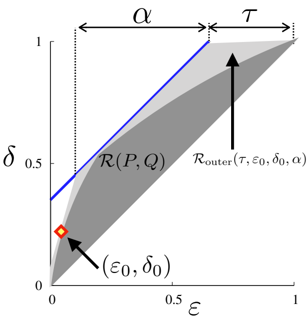

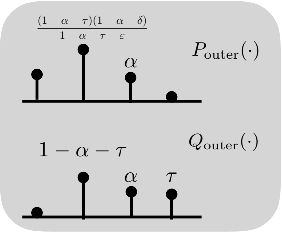

$$\max_{P,Q} \quad d_{\text{TV}}(P^2, Q^2)$$

subject to $\quad d_{\text{TV}}(P, Q) = \tau$

no $(\varepsilon_0, \delta_0)$-mode collapse

$$\mathcal{R}(P,Q) \subseteq \mathcal{R}_{\text{outer}}(\tau, \varepsilon_0, \delta_0, \alpha)$$

Upper bound without $(\varepsilon_0, \delta_0)$-mode collapse

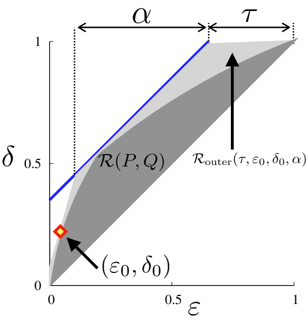

$$
\begin{array}{rl}
\displaystyle\max_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\
\text{subject to} & d_{\mathrm{TV}}(P,Q) = \tau \\
 & \text{no } (\varepsilon_0, \delta_0)\text{-mode collapse}
\end{array}
$$

$$
\begin{array}{rcl}
\mathcal{R}(P,Q) & \subseteq & \mathcal{R}_{\text{outer}}(\tau, \varepsilon_0, \delta_0, \alpha) \\
\mathcal{R}(P^2,Q^2) & \subseteq & \mathcal{R}(P^2_{\text{outer}}, Q^2_{\text{outer}})
\end{array}
$$

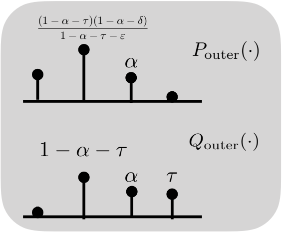

Blackwell's theorem

$$
\begin{array}{rl}
 & \mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q') \\
\Rightarrow & \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2)
\end{array}
$$

Upper bound without $(\varepsilon_0, \delta_0)$-mode collapse

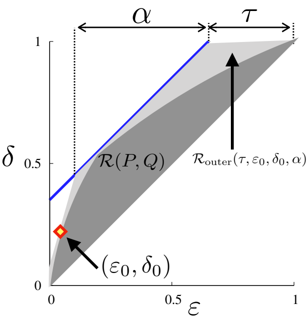

$$\begin{array}{ll} \max\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P,Q) = \tau \\ & \text{no } (\varepsilon_0,\delta_0)\text{-mode collapse} \end{array}$$

$$\begin{array}{rcl} \mathcal{R}(P,Q) & \subseteq & \mathcal{R}_{\text{outer}}(\tau,\varepsilon_0,\delta_0,\alpha) \\ \mathcal{R}(P^2,Q^2) & \subseteq & \mathcal{R}(P^2_{\text{outer}}, Q^2_{\text{outer}}) \\ d_{\text{TV}}(P^2,Q^2) & \leq & \underbrace{d_{\text{TV}}(P^2_{\text{outer}}, Q^2_{\text{outer}})}_{\text{simple to evaluate}} \end{array}$$

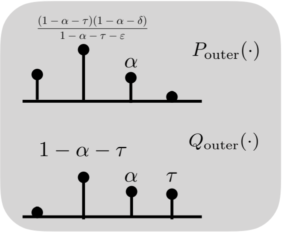

Blackwell's theorem

$$\mathcal{R}(P,Q) \subseteq \mathcal{R}(P',Q')$$
$$\Rightarrow \quad \mathcal{R}(P^2,Q^2) \subseteq \mathcal{R}(P'^2,Q'^2)$$

$d_{\mathrm{TV}}(P^2, Q^2)$

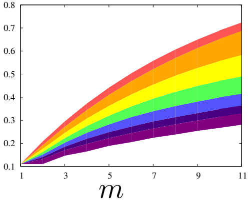

$$\begin{aligned} \min_{P,Q} \quad & d_{\mathrm{TV}}(P^2, Q^2) \\ \text{subject to} \quad & d_{\mathrm{TV}}(P, Q) = \tau \\ & (\varepsilon_1, \delta_1)\text{-mode collapse} \end{aligned}$$

Lower bound with $(\varepsilon_1, \delta_1)$-mode collapse

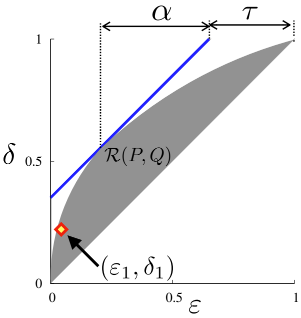

$$
\begin{array}{rl}
\min\limits_{P,Q} & d_{\mathrm{TV}}(P^2, Q^2) \\
\text{subject to} & d_{\mathrm{TV}}(P, Q) = \tau \\
& (\varepsilon_1, \delta_1)\text{-mode collapse}
\end{array}
$$

Lower bound with $(\varepsilon_1, \delta_1)$-mode collapse

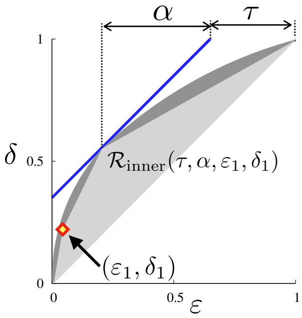

$$
\begin{aligned}
\min_{P,Q} \quad & d_{\text{TV}}(P^2, Q^2) \\
\text{subject to} \quad & d_{\text{TV}}(P, Q) = \tau \\
& (\varepsilon_1, \delta_1)\text{-mode collapse}
\end{aligned}
$$

$$\mathcal{R}_{\text{inner}}(\tau, \alpha, \varepsilon_1, \delta_1) \subseteq \mathcal{R}(P, Q)$$

Lower bound with $(\varepsilon_1, \delta_1)$-mode collapse

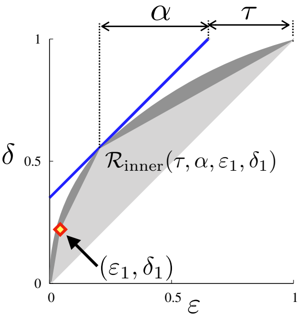

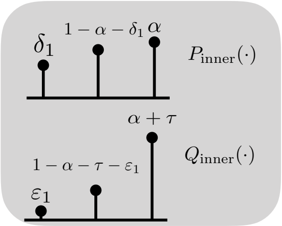

$$\begin{array}{rl} \min\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P, Q) = \tau \\ & (\varepsilon_1, \delta_1)\text{-mode collapse} \end{array}$$

$$\mathcal{R}_{\text{inner}}(\tau, \alpha, \varepsilon_1, \delta_1) \subseteq \mathcal{R}(P, Q)$$

Lower bound with $(\varepsilon_1, \delta_1)$-mode collapse

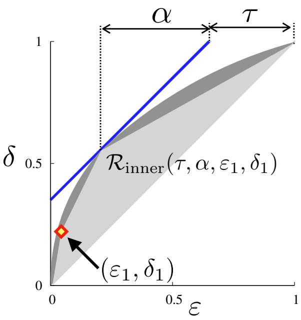

$$\begin{array}{rl} \min\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P, Q) = \tau \\ & (\varepsilon_1, \delta_1)\text{-mode collapse} \end{array}$$

$$\begin{array}{rcl} \mathcal{R}_{\text{inner}}(\tau, \alpha, \varepsilon_1, \delta_1) & \subseteq & \mathcal{R}(P, Q) \\ \mathcal{R}(P^2_{\text{inner}}, Q^2_{\text{inner}}) & \subseteq & \mathcal{R}(P^2, Q^2) \end{array}$$

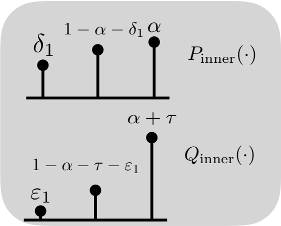

Blackwell's theorem

$$\begin{array}{rl} & \mathcal{R}(P, Q) \subseteq \mathcal{R}(P', Q') \\ \Rightarrow & \mathcal{R}(P^2, Q^2) \subseteq \mathcal{R}(P'^2, Q'^2) \end{array}$$

Lower bound with $(\varepsilon_1, \delta_1)$-mode collapse

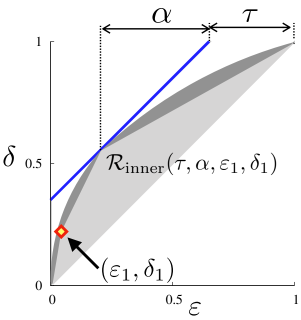

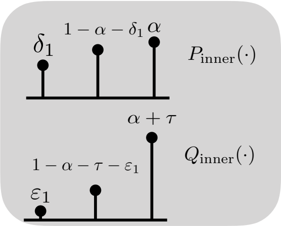

$$\begin{array}{rl} \min\limits_{P,Q} & d_{\text{TV}}(P^2, Q^2) \\ \text{subject to} & d_{\text{TV}}(P, Q) = \tau \\ & (\varepsilon_1, \delta_1)\text{-mode collapse} \end{array}$$

$$\begin{array}{rcl} \mathcal{R}_{\text{inner}}(\tau, \alpha, \varepsilon_1, \delta_1) & \subseteq & \mathcal{R}(P, Q) \\ \mathcal{R}(P^2_{\text{inner}}, Q^2_{\text{inner}}) & \subseteq & \mathcal{R}(P^2, Q^2) \\ \underbrace{d_{\text{TV}}(P^2_{\text{inner}}, Q^2_{\text{inner}})}_{\text{simple to evaluate}} & \leq & d_{\text{TV}}(P^2, Q^2) \end{array}$$

Blackwell's theorem

$$\begin{array}{rl} & \mathcal{R}(P, Q) \subseteq \mathcal{R}(P', Q') \\ \Rightarrow & \mathcal{R}(P^2, Q^2) \subseteq \mathcal{R}(P'^2, Q'^2) \end{array}$$

Achievable TV distances for distributions

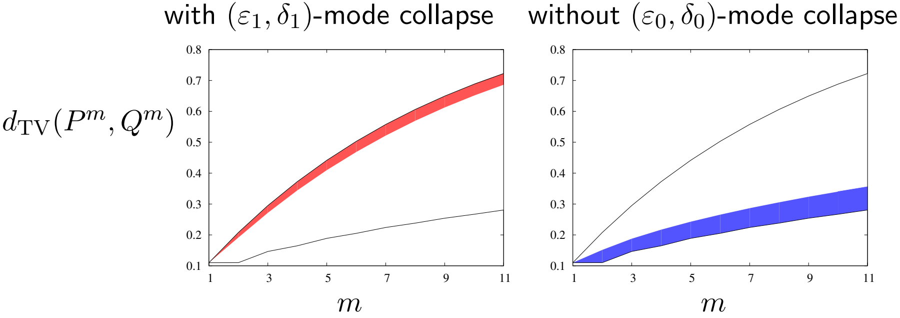

with packing, the discriminator naturally penalizes (P, Q) with severe mode collapses

Does larger m help?

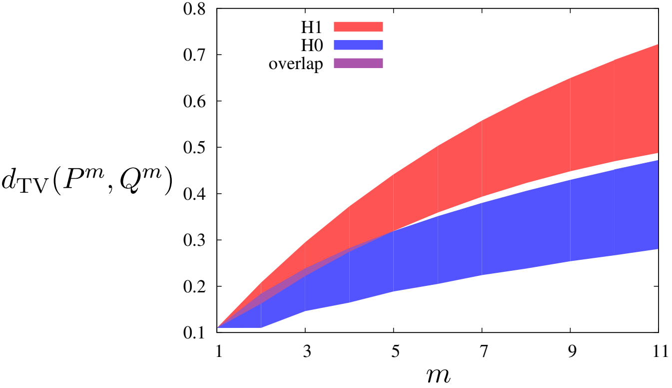

Evolution of TV distances through the prism of packing

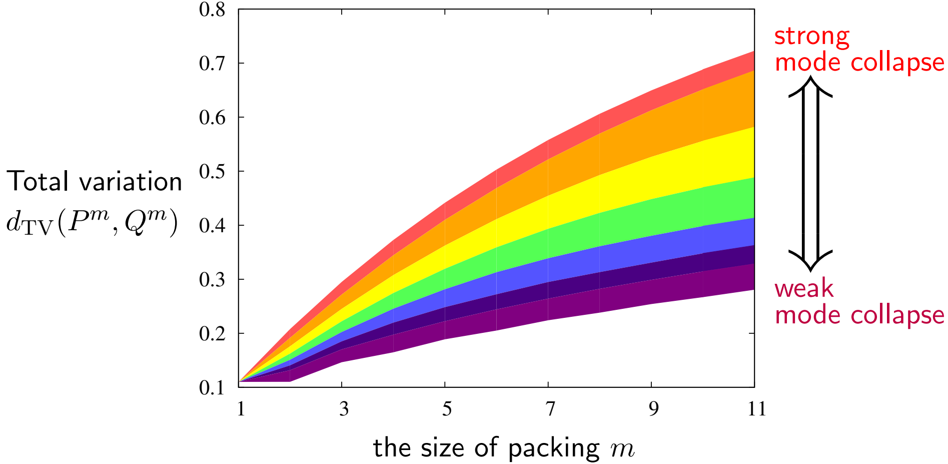

Through the **prism** packing, the target-generator pairs are expanded over the **spectrum** of (the strengths of) mode collapse

Our paper is available at:
https://arxiv.org/abs/1712.04086

- *Packing* is a principled solution to mode collapse
  - Lightweight overhead
  - Excellent numerical results
  - Theoretical analysis
- Exciting time for bridging the theory and practice of GAN
  - "Generalization and Equilibrium in Generative Adversarial Nets", Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, Yi Zhang
  - "Approximation and Convergence Properties of Generative Adversarial Learning", Shuang Liu, Olivier Bousquet, Kamalika Chaudhuri
  - "Compressed Sensing using Generative Models", Ashish Bora, Ajil Jalal, Eric Price, Alexandros G. Dimakis
  - "Towards Understanding the Dynamics of Generative Adversarial Networks", Jerry Li, Aleksander Madry, John Peebles, Ludwig Schmidt


Giulia Fanti (CMU)


Ashish Khetan (UIUC)


Zinan Lin (CMU)