Dataset Condensation

ATDL Talk

Tong Chen, ML section, DIKU

Illustrative example: SVM

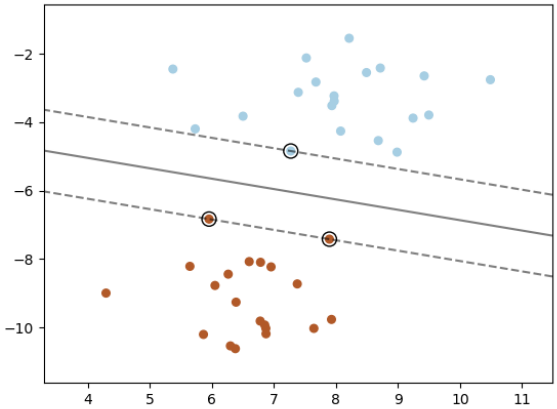

What is Dataset Condensation (DC)?

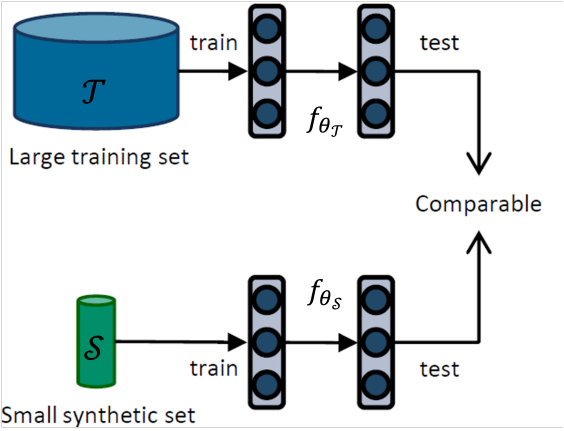

Does DC work?

50 images per class for CNN

98.8 vs 99.6

83.6 vs 93.5

82.3 vs 95.4

53.9 vs 84.8

[Zhao et al. 2021] Dataset Condensation with Gradient Matching

How does DC work?

How does DC work?

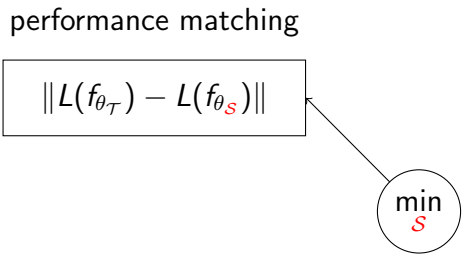

How does DC work?

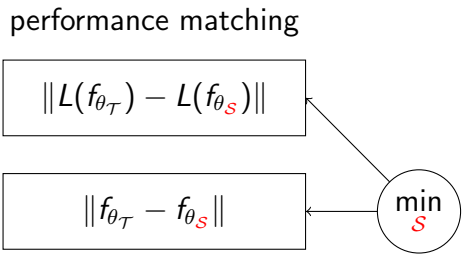

How does DC work?

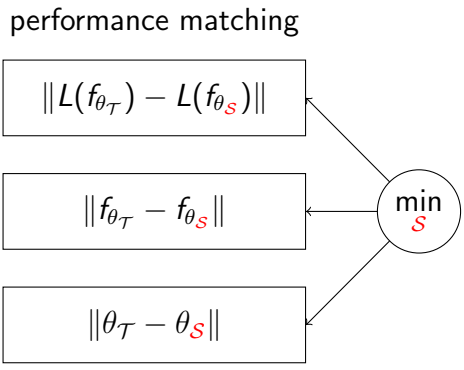

How does DC work?

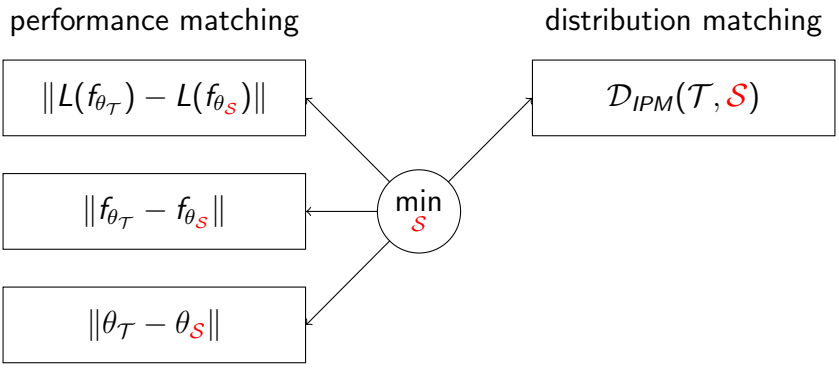

eg. $\mathcal{D}_{IPM}(\mathcal{T}, \mathcal{S}) = \sup_{f \in \mathcal{F}} \|\mathbb{E}_\mathcal{T}[f] - \mathbb{E}_\mathcal{S}[f]\|$

How does DC work?

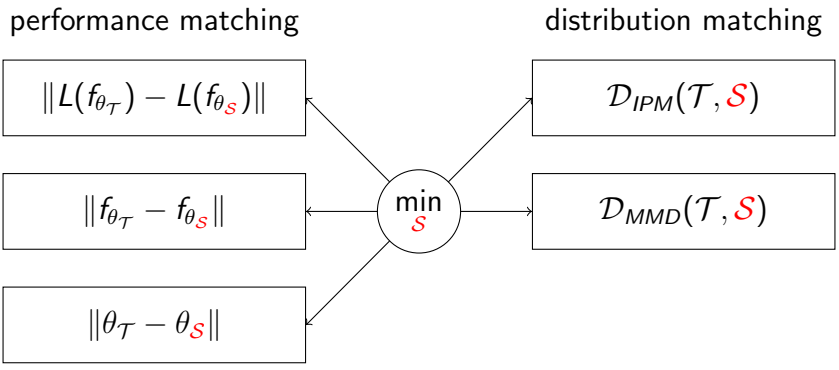

eg. $\mathcal{D}_{IPM}(\mathcal{T}, \mathcal{S}) = \sup_{f \in \mathcal{F}} \|\mathbb{E}_\mathcal{T}[f] - \mathbb{E}_\mathcal{S}[f]\|$

How does DC work?

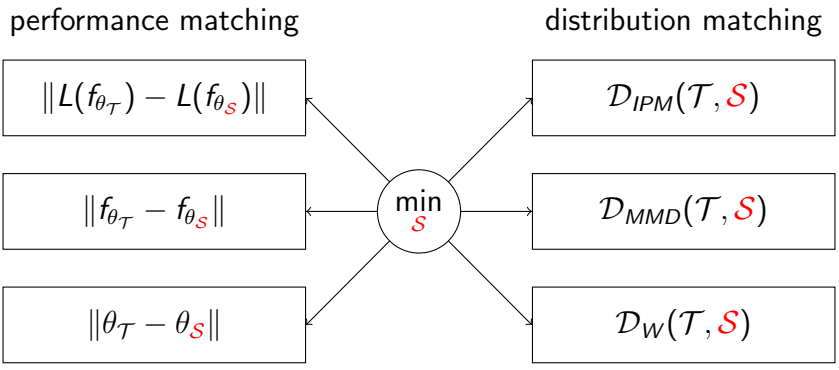

eg. $\mathcal{D}_{IPM}(\mathcal{T}, \mathcal{S}) = \sup_{f \in \mathcal{F}} \|\mathbb{E}_\mathcal{T}[f] - \mathbb{E}_\mathcal{S}[f]\|$

Feature transformation: generative models

Input space

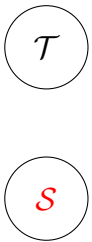

Feature transformation: generative models

Input space

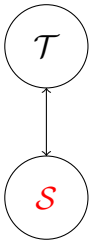

Feature transformation: generative models

Input space

Latent space

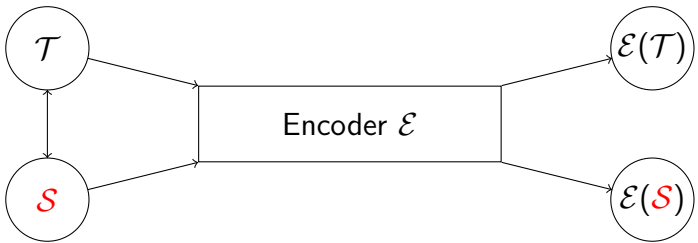

Feature transformation: generative models

Input space

Latent space

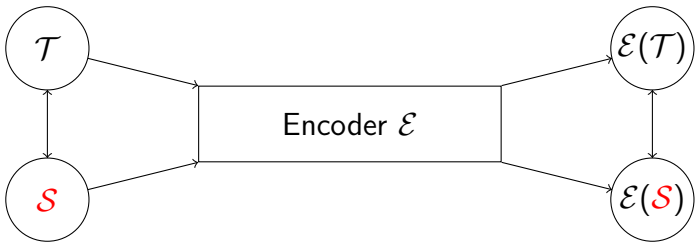

Feature transformation: generative models

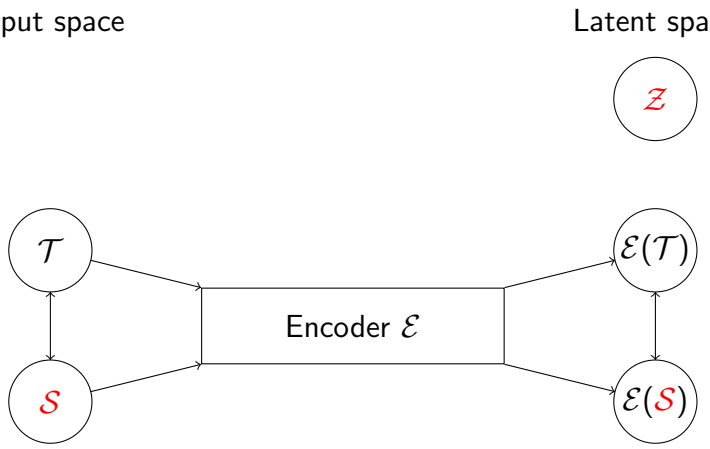

Feature transformation: generative models

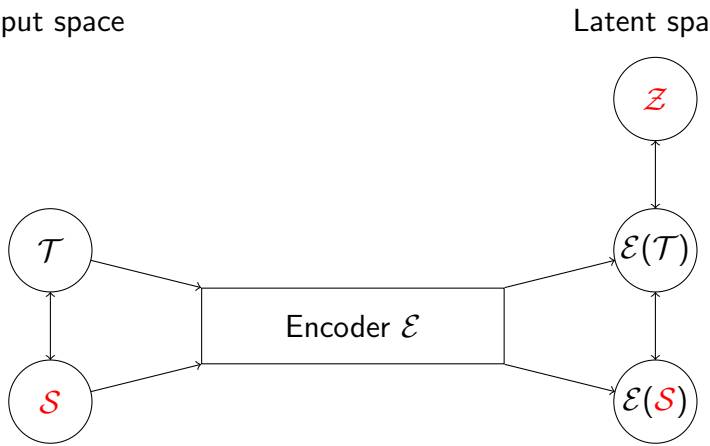

Feature transformation: generative models

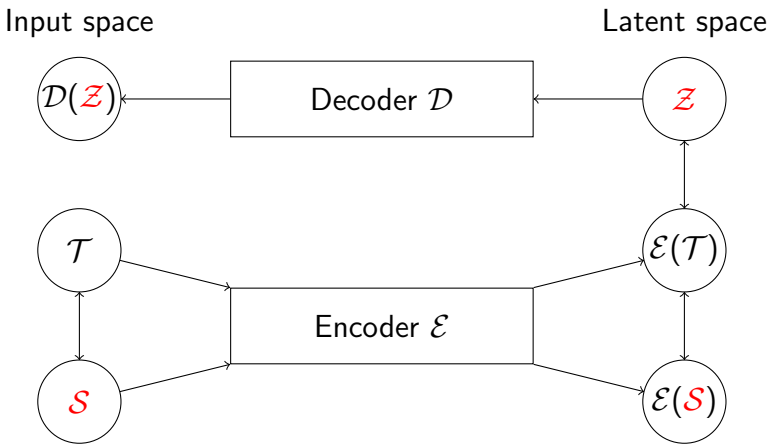

Feature transformation: generative models

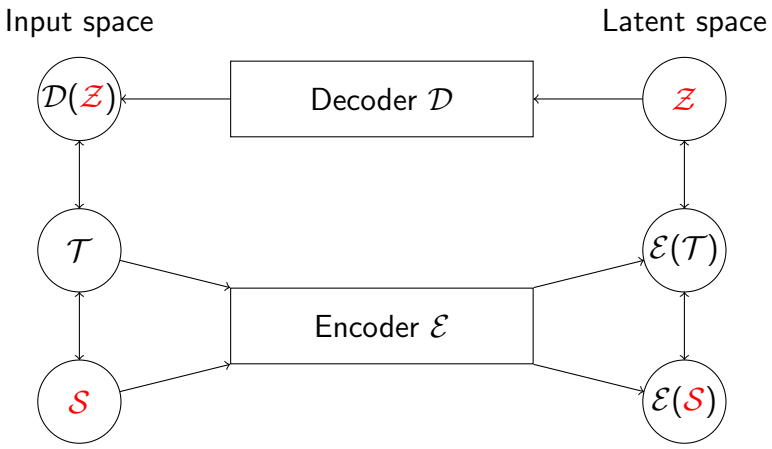

Examples: DC with performance matching

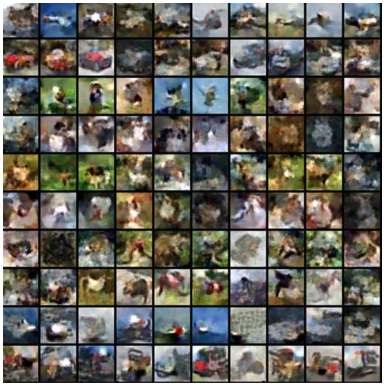

Matching f_θ [Wang et al. 2022]

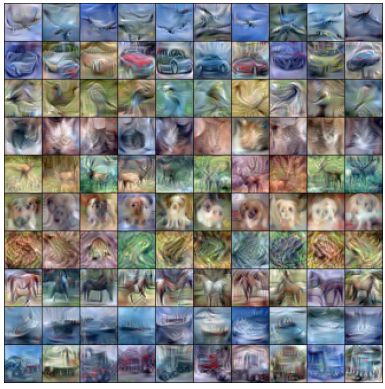

Matching θ [Cazenavette et al. 2022]

Examples: DC with distribution matching

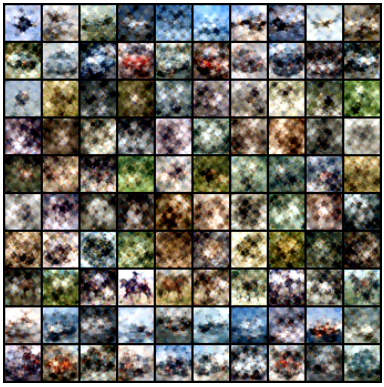

IPM with CNN [Zhao et al. 2021]

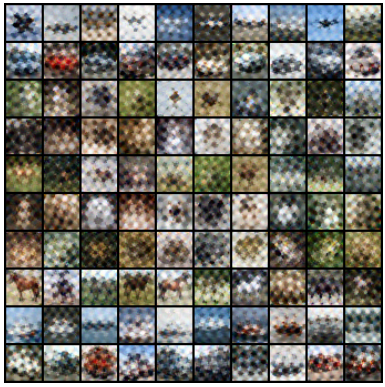

MMD with Gaussian [Zhang et al. 2023]

Examples: DC with generative models

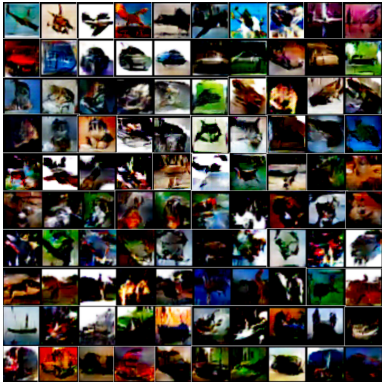

Pre-trained GAN [Zhang et al. 2023]

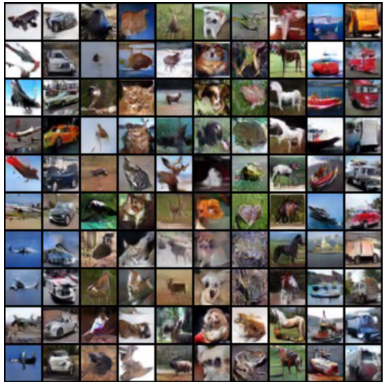

GAN [Wang et al. 2023]

Open questions about DC

- Beyond accuracy: robustness, fairness, privacy, etc.

Open questions about DC

- Beyond accuracy: robustness, fairness, privacy, etc.
- Does synthetic dataset preserves privacy, robustness or fairness?

Open questions about DC

- Beyond accuracy: robustness, fairness, privacy, etc.
- Does synthetic dataset preserves privacy, robustness or fairness?
- Adversarial loss:

$$L^{adv}(f) = \mathbb{E}_{\mathcal{T}}\left[\max_{\|\delta\| \leq \varepsilon} l(f(\mathbf{x}+\delta), y)\right]$$

Open questions about DC

- Beyond accuracy: robustness, fairness, privacy, etc.
- Does synthetic dataset preserves privacy, robustness or fairness?
- Adversarial loss:

$$L^{adv}(f) = \mathbb{E}_{\mathcal{T}}\left[\max_{\|\delta\|\leq\varepsilon} l(f(\mathbf{x}+\delta), y)\right]$$

Performance matching:

$$\min_{\mathcal{S}} \|L^{adv}(f^{adv}_{\theta_{\mathcal{T}}}) - L^{adv}(f^{adv}_{\theta_{\mathcal{S}}})\|$$

Open questions about DC

- Beyond accuracy: robustness, fairness, privacy, etc.
- Does synthetic dataset preserves privacy, robustness or fairness?
- Adversarial loss:

$$L^{adv}(f) = \mathbb{E}_{\mathcal{T}}\left[\max_{\|\delta\|\leq\varepsilon} l(f(\mathbf{x}+\delta), y)\right]$$

Performance matching:

$$\min_{\mathcal{S}} \|L^{adv}(f^{adv}_{\theta_{\mathcal{T}}}) - L^{adv}(f^{adv}_{\theta_{\mathcal{S}}})\|$$

- How does the distribution of synthetic dataset looks like?