Vicinity-Guided Discriminative Latent Diffusion for Privacy-Preserving Domain Adaptation

Jing Wang[1], Wonho Bae[1], Jiahong Chen[3], Wenxu Wang[4], Junhyug Noh[2]

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

[1]University of British Columbia, [2]Ewha Womans University,
[3]Amazon Inc., [4]Ocean University of China

Can we explicitly transfer knowledge without sharing source data?

Privacy-Sensitive Source Domain

Expert Model

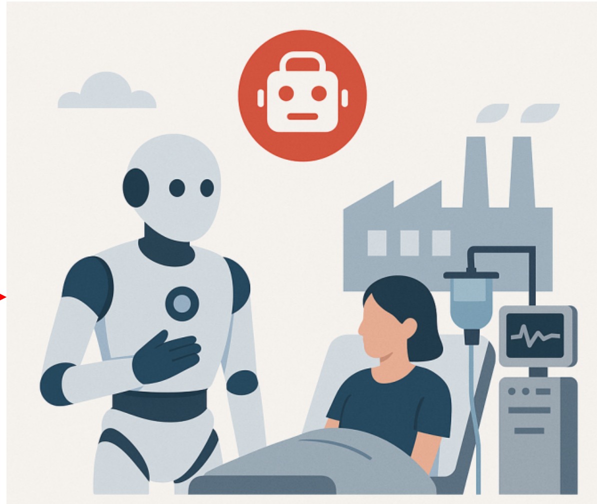

Target Domain

Source-Free Domain Adaptation (SFDA):

- Modern AI models are trained on confidential or proprietary data (e.g., hospitals, industries).
- Regulations and ownership prohibit sharing the original datasets.
- How can we transfer knowledge without transferring data?

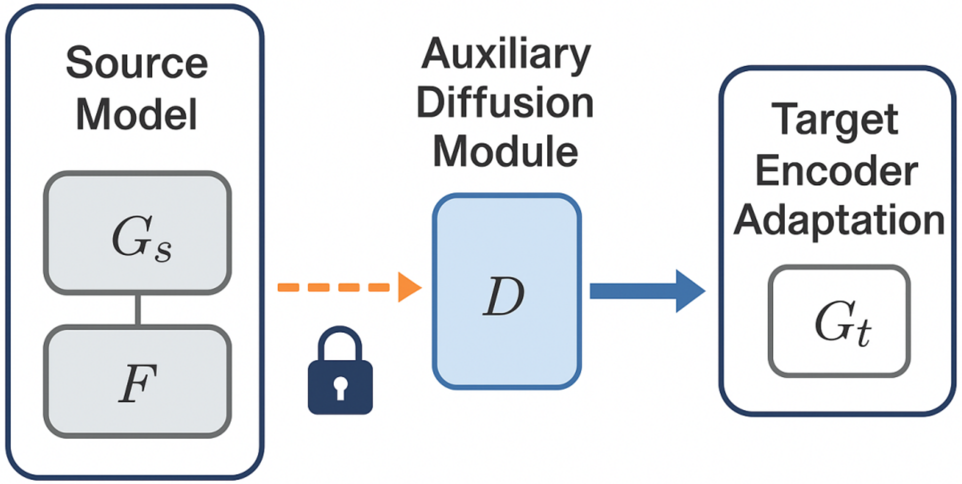

Limitation of SFDA: SFDA usually assumes we can only share a pre-trained model. This limits explicit knowledge transfer. Adaptation becomes implicit and hard to interpret.

Relaxation: Can we have a more practical variant of SFDA: the source provider can release a lightweight auxiliary module, trained once offline, but still never exposes raw data?

Existing SFDA Paradigms — What is Missing?

Category	Intuition	Example Methods	Limitation
Entropy / Pseudo-Labeling	Self-train the target model by minimizing entropy or updating pseudo-labels iteratively.	SHOT (ICML 2020), NRC++ (TPAMI 2023)	Implicit alignment — cannot explicitly transfer source decision geometry.
Contrastive / Prototype-Based	Use contrastive losses or class prototypes to cluster target features by prediction consistency.	DaC (NeurIPS 2022), CPD (PR 2024)	Clusters form implicitly; lacks a controllable mechanism to guide feature clustering.
Generative / Augmentation-Based	Synthesize pseudo-source samples using GANs or diffusion in pixel space.	3C-GAN (CVPR 2020), SFADA (PR 2024)	Requires heavy pixel-level generation; no explicit mapping of discriminative knowledge.

Research Gap: A principled and lightweight way to **explicitly** transfer discriminative knowledge without raw source data.

Question: Can we repurpose **latent diffusion** to bridge this gap: transferring decision boundaries directly in latent space?

Liang et al., "Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation." ICML, 2020.
Li et al., “Model Adaptation: Unsupervised Domain Adaptation without Source Data.” CVPR, 2020
Yang et al., "Trust your good friends: Source-free domain adaptation by reciprocal neighborhood clustering." TPAMI, 2023.
Zhang et al., "Divide and contrast: Source-free domain adaptation via adaptive contrastive learning." NeurIPS, 2022.
Zhou et al., "Source-free domain adaptation with class prototype discovery." PR, 2024.
He et al., "Source-free domain adaptation with unrestricted source hypothesis.” PR, 2024.

From Generative Diffusion to Discriminative Transfer

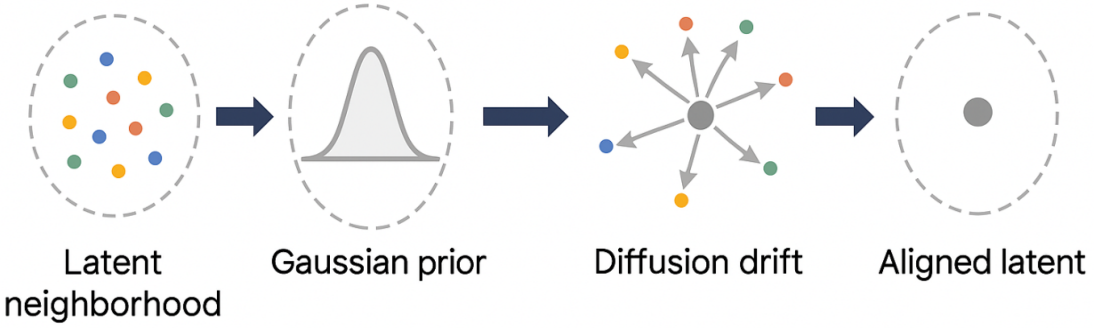

- **From Denoising to Drifting:** Diffusion models denoise samples toward realistic data. We reinterpret this denoising as “drifting” noisy latent points toward label-consistent regions.

- **Encoding Local Semantics:** Each source sample’s k-NN vicinity defines a Gaussian prior, capturing the local geometry and class-level semantics.

- **Learning Discriminative Drift:** The diffusion drift network learns to pull latent features back toward their Gaussian priors — effectively storing discriminative knowledge.

Instead of generating pixels, we use latent diffusion to generate knowledge flow within the latent space across domains.

Discriminative Vicinity Diffusion (DVD) – Method Overview

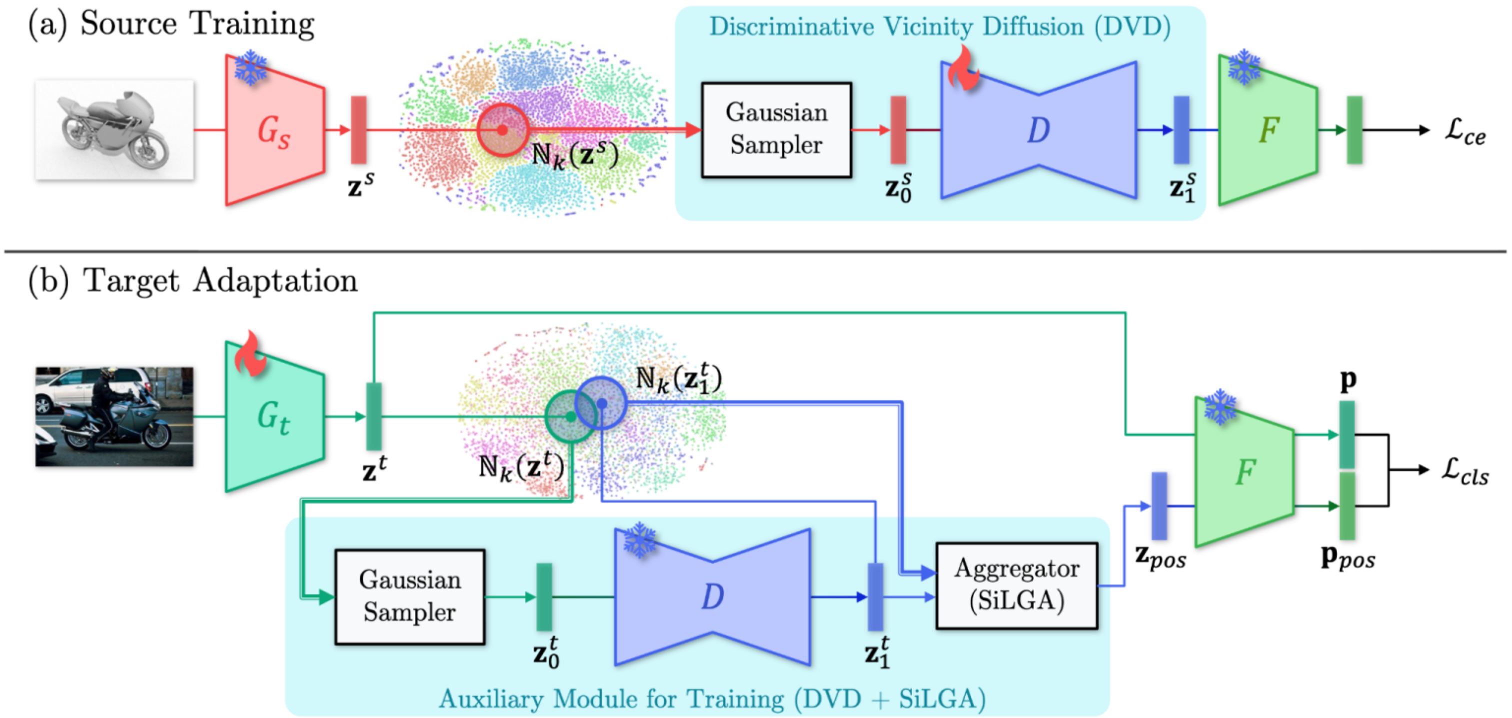

- **Source Training:**
  - Fit Gaussian priors on k-NN latent vicinities.
  - Train drift function D to pull noisy samples back to label-consistent manifold.
- **Target Adaptation:**
  - Use frozen D to guide target features toward nearest source vicinities.
  - Simple contrastive alignment ensures effective transfer of decision boundary.
- **SiLGA:**
  - Blend of two clusters (source-informed and target neighborhoods).
  - D alone may overfit under large domain gaps.
  - SiLGA fuses source-informed cues with local target geometry for stability.
  - Keeps adaptation both label-consistent and geometry-aware.

Discriminative Vicinity Diffusion (DVD) – Method Overview

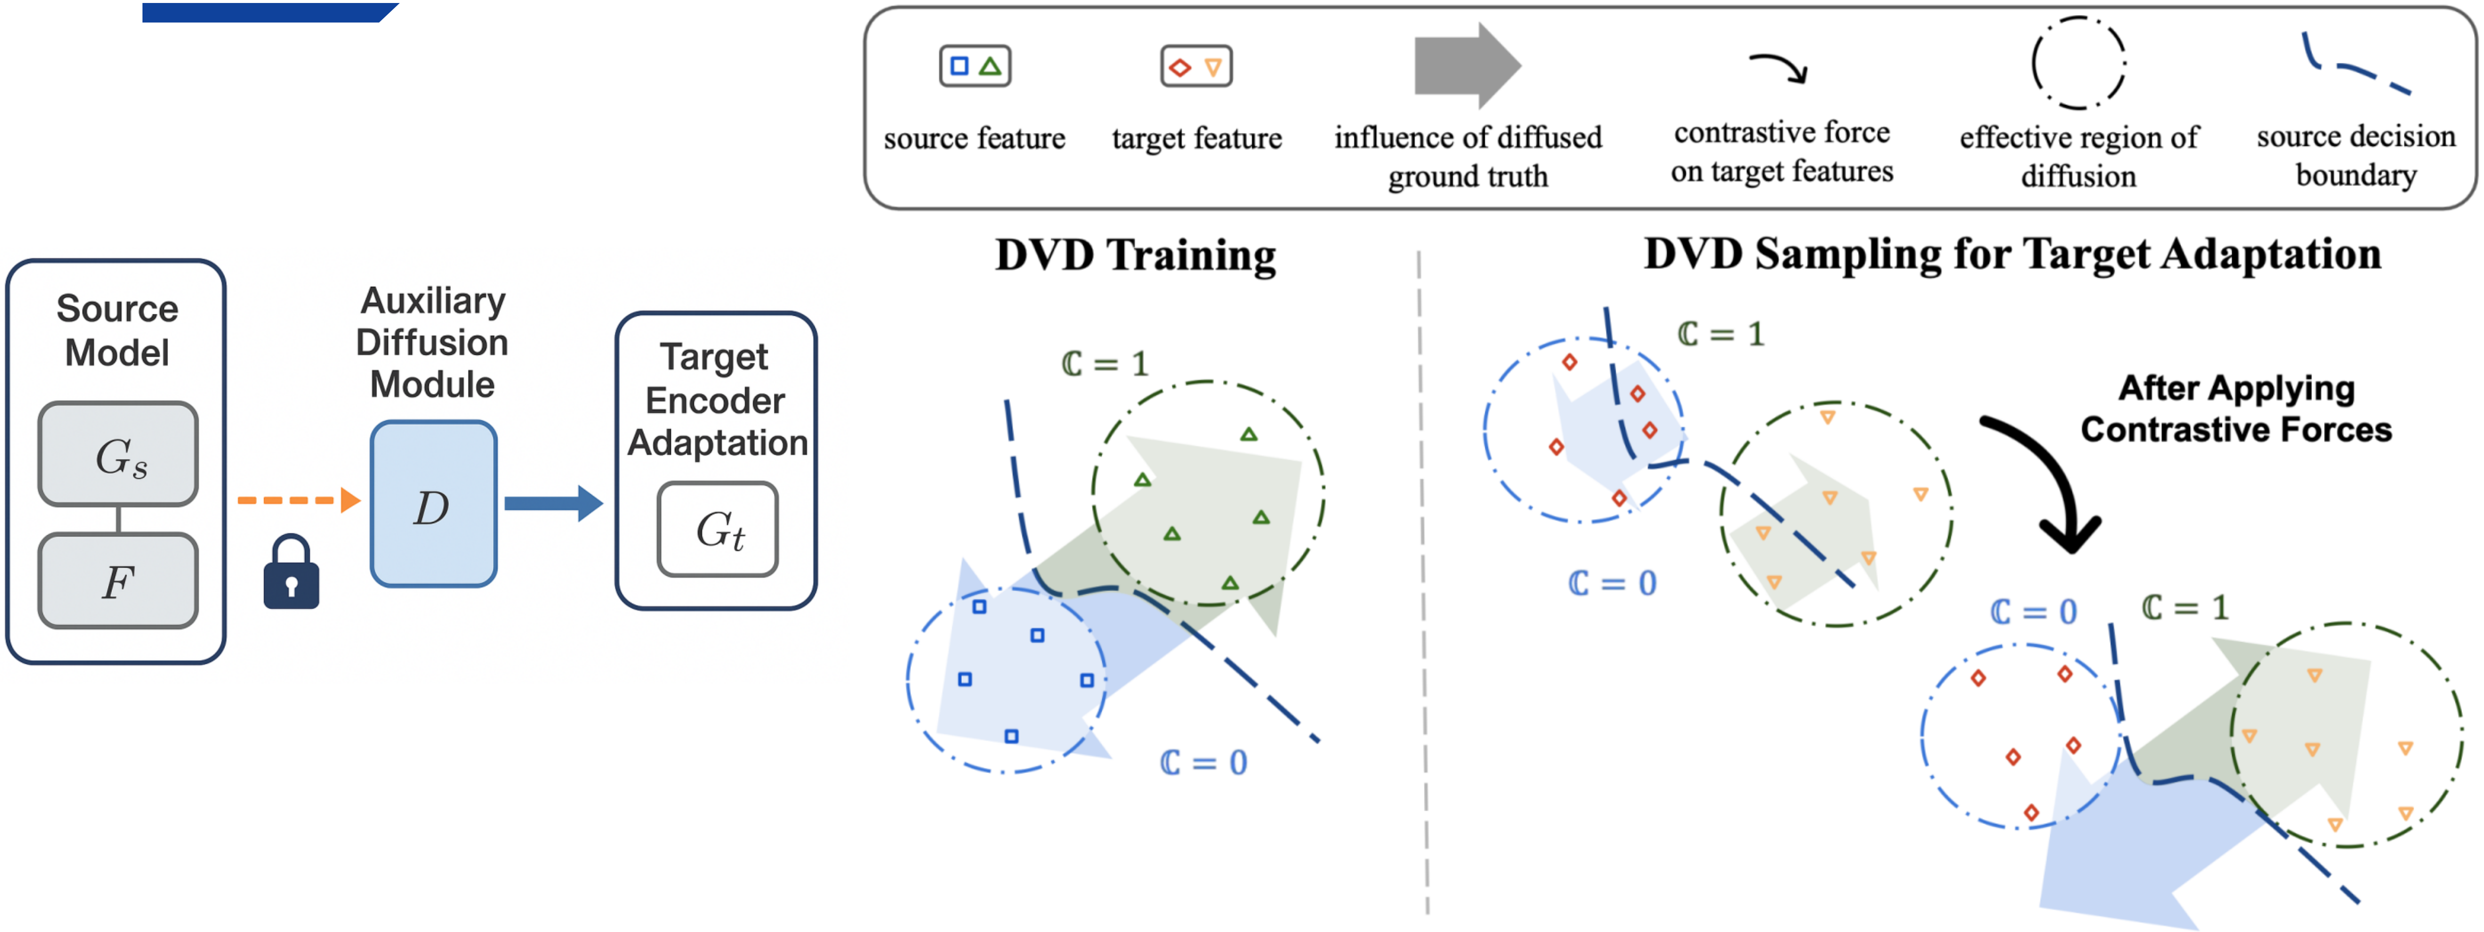

The diffusion module (D) learns to "*drift*" noisy **target latent features** back to their label-consistent **source manifolds**.

- During adaptation, D generates **source-aligned cues** from each **target latent vicinity**.
- **Explicit knowledge transfer** with no source data required.

Why Deterministic Diffusion?

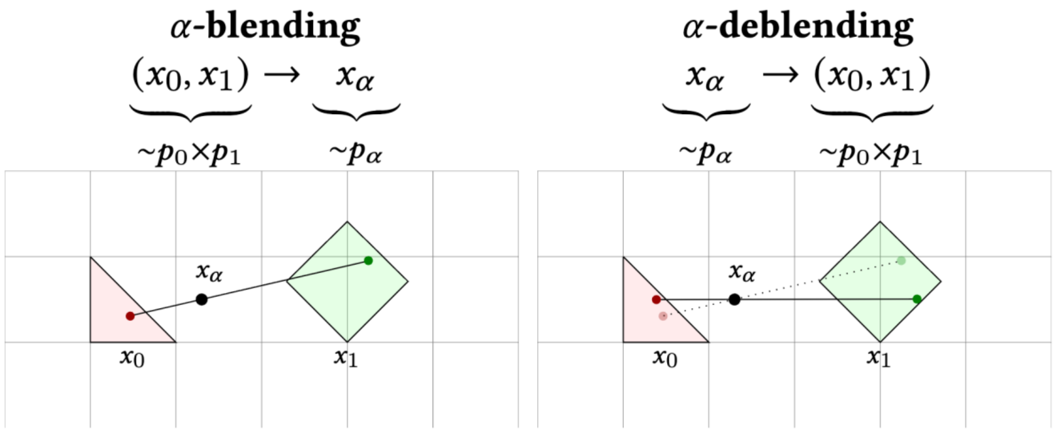

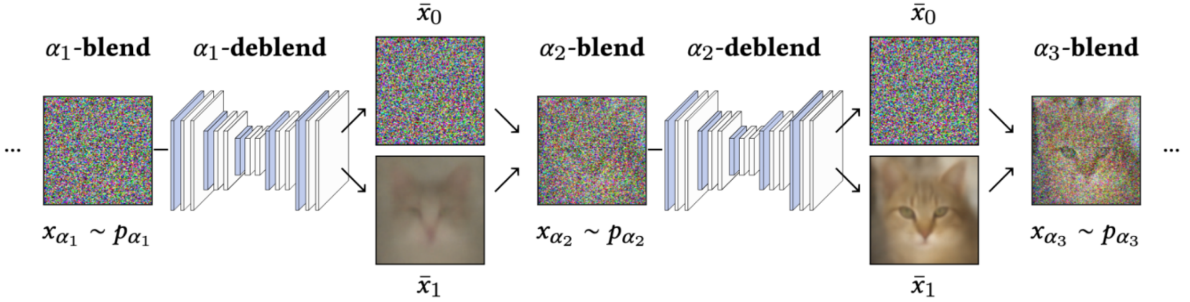

$$\mathbf{z}_{\alpha_{t+1}} = \mathbf{z}_{\alpha_t} + (\alpha_{t+1} - \alpha_t)\, D(\mathbf{z}_{\alpha_t}, \alpha_t)$$

$$\mathcal{L}_{dif} = \mathbb{E}_{\alpha_t, \mathbf{z}_{\alpha_t}}\left[|| D(\mathbf{z}_{\alpha_t}, \alpha_t) - \mathbb{E}_{(\mathbf{z}_0, \mathbf{z}_1)|(\mathbf{z}_{\alpha_t}, \alpha_t)}[\mathbf{z}_1 - \mathbf{z}_0] ||^2 \right]$$

Algorithm 1: Latent Diffusion Training

Input:
$(\mathbf{z}_0, \mathbf{z}_1) \sim (p_0 \times p_1)$
$\alpha \sim \text{Uniform}[0, 1]$
D's parameters ϕ; steps T

for $t = 0$ **to** T **do**
- $\mathbf{z}_{\alpha_t} \leftarrow (1 - \alpha_t)\mathbf{z}_0 + \alpha_t \mathbf{z}_1$;
- Update ϕ using $\mathcal{L}_{\text{dif}}$ (Eq. 6);
- Sample $\hat{\mathbf{z}}_1$ via Algorithm 2;
- If $\mathbf{z}_1$ has label, minimize $\mathcal{L}_{\text{ce}}$ (Eq. 7);

end

Output: Updated parameters ϕ

Algorithm 2: Latent Diffusion Sampling

Input:
$(\mathbf{z}_0, \mathbf{z}_1) \sim (p_0 \times p_1)$
Discrete $\alpha_t = \frac{t}{T}$; steps T

Initialize: $\mathbf{z}_{\alpha_0} \leftarrow \mathbf{z}_0$
for $t = 0$ **to** $T - 1$ **do**
- $\mathbf{z}_{\alpha_{t+1}} \leftarrow \mathbf{z}_{\alpha_t} + (\alpha_{t+1} - \alpha_t)\, D(\mathbf{z}_{\alpha_t}, \alpha_t)$;

end

Output: $\hat{\mathbf{z}}_1 = \mathbf{z}_{\alpha_T} \approx \mathbf{z}_1$

- **Stable feature transport:** The drift field defines a single path between noisy and clean latents, preventing instability or variance introduced by randomness.
- **Preserves label consistency:** Since every trajectory is deterministic, each latent consistently returns to its label-consistent manifold, improving discriminative reliability.
- **No error accumulation:** By predicting a **global displacement vector** (from noisy to clean latent) instead of incremental noise, it avoids iterative drift errors common in stochastic or score-based diffusion.
- **Efficiency:** Deterministic (drift-only) diffusion eliminates random sampling at each step, making the process faster and more lightweight than stochastic diffusion.

Heitz et al., "Iterative α-(de) blending: A minimalist deterministic diffusion model." SIGGRAPH, 2023.

t-SNE Visualization

Before DVD — Scattered and Overlapping

- Target features are disorganized and overlapping across classes, reflecting severe domain shift.
- Features lack geometric alignment with the source classifier's structure, leading to uncertain representations.

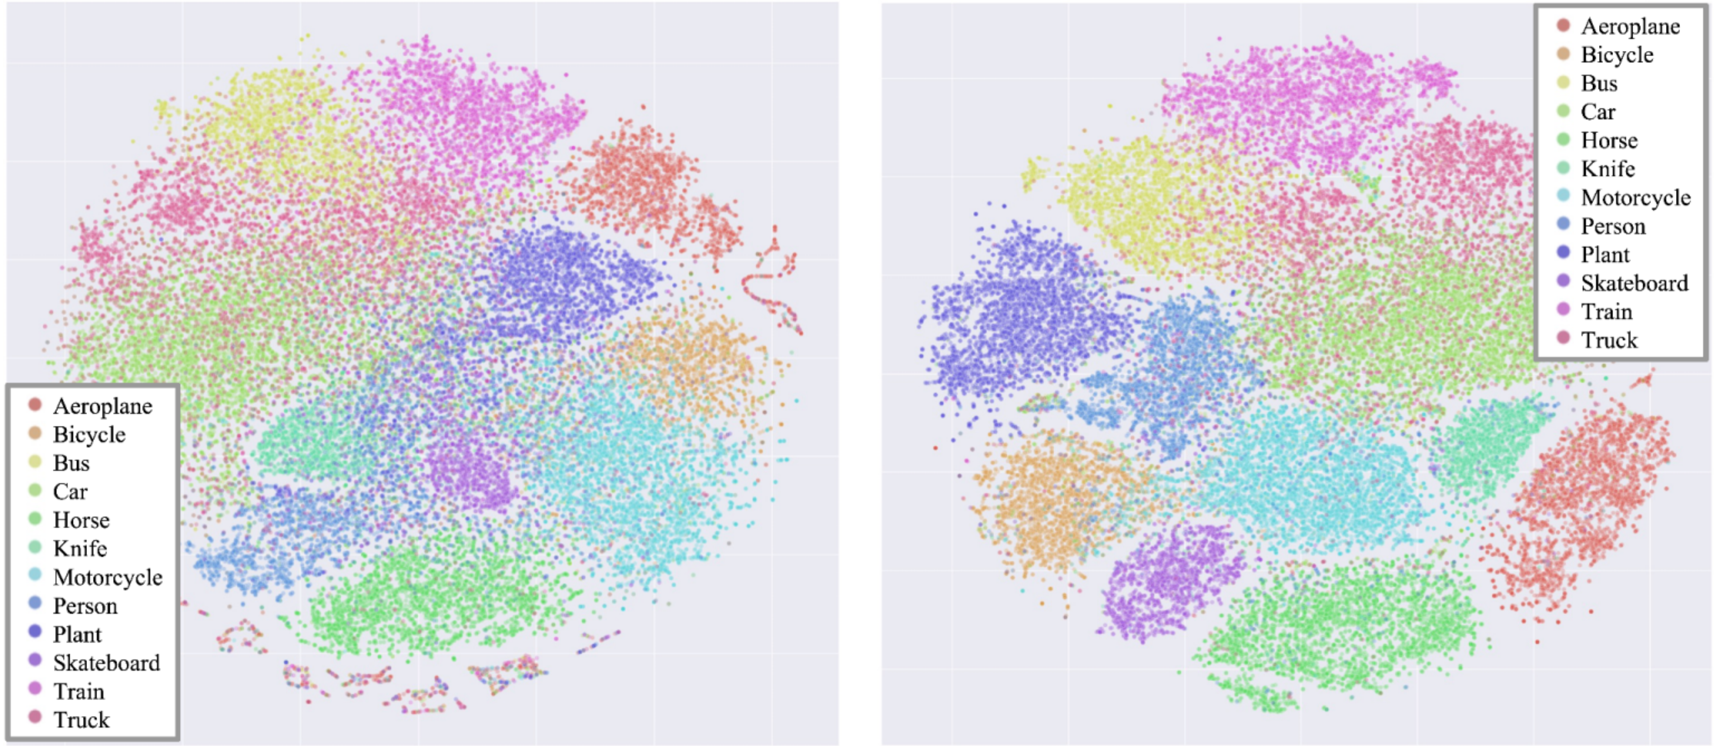

(a) Without DVD (b) With DVD

After DVD — Structured and Aligned Clusters

- DVD produces compact and well-separated clusters in latent space.
- Decision boundaries from the source model can now be directly transferred to the target domain.

✔ **Explicit Knowledge Transfer:** DVD bridges domains through latent space dynamics, not heuristic fine-tuning.
✔ **Discriminative Alignment:** The drift field enforces label-consistent geometry, yielding interpretable cluster formation.
✔ **Privacy-Preserving Adaptation:** Source data never appear, yet their discriminative structure is effectively reproduced in the target domain.

Experiments on SFDA

Comparison of SFDA methods on Office-31 (ResNet-50).

Method	Add.	A→D	A→W	D→W	D→A	W→D	W→A	Avg. ± s.d.
ResNet-50 [He et al., 2016]		68.9	68.4	96.7	62.5	99.3	60.7	76.1
SHOT [Liang et al., 2020]		94.0	90.1	98.4	74.7	99.9	74.3	88.6
AaD [Yang et al., 2022]		95.5	92.1	98.5	74.0	99.4	75.8	89.2
NRC++ [Yang et al., 2023b]		95.9	91.2	**99.1**	75.5	100.0	75.0	89.5
SFADA [He et al., 2024]		94.8	92.0	97.6	76.5	99.8	75.7	89.4
CPD [Zhou et al., 2024]		96.6	94.2	98.2	77.3	100.0	**78.3**	90.8
3C-GAN [Li et al., 2020]	✓	93.4	92.9	97.5	76.5	99.8	77.3	89.6
SFADA [He et al., 2024]	✓	95.2	91.4	98.2	77.8	100.0	76.3	89.8
DVD (Ours)	✓	**96.7**	**95.2**	98.6	**79.2**	**100.0**	77.4	**91.2** ± 0.5

Comparison of SFDA methods on Office-Home (ResNet-50).

Method	Add.	Ar →			Cl →			Pr →			Rw →			Avg. ± s.d.
		Cl	Pr	Rw	Ar	Pr	Rw	Ar	Cl	Rw	Ar	Cl	Pr	
ResNet-50 [He et al., 2016]		34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DaC [Zhang et al., 2022]		59.5	79.5	81.2	**69.3**	78.9	79.2	67.4	56.4	82.4	**74.0**	**61.4**	84.4	72.8
NRC++ [Yang et al., 2023b]		57.8	**80.4**	81.6	69.0	80.3	79.5	65.6	57.0	83.2	72.3	59.6	85.7	72.5
SFADA [He et al., 2024]		56.1	78.0	81.6	68.5	79.5	78.5	67.8	56.0	82.3	73.6	57.8	83.0	71.9
CPD [Zhou et al., 2024]		59.1	79.0	82.4	68.5	79.7	79.5	67.9	57.9	82.8	73.8	61.2	84.6	73.0
3C-GAN [Li et al., 2020]	✓	57.4	79.3	81.8	69.1	79.8	80.0	66.5	56.2	83.9	72.4	60.1	85.4	72.7
SFADA [He et al., 2024]	✓	57.8	78.5	82.3	68.2	79.6	79.2	66.9	57.4	83.8	73.4	58.0	84.1	72.4
DVD (Ours)	✓	**60.1**	79.6	**82.5**	69.1	**80.8**	**80.6**	**67.9**	**58.5**	**84.3**	73.5	59.3	**87.6**	**73.7** ± 0.7

Comparison of the SFDA methods on VisDA-C 2017 (ResNet-101).

Method	Add.	plane	bcycl	bus	car	horse	knife	mcycl	person	plant	sktbrd	train	truck	Avg. ± s.d.
ResNet-101 [He et al., 2016]		55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
SHOT [Liang et al., 2020]		94.3	88.5	80.1	57.3	93.1	94.9	80.7	80.3	91.5	89.1	86.3	58.2	82.9
HCL [Huang et al., 2021]		93.3	85.4	80.7	68.5	91.0	88.1	86.0	78.6	86.6	88.8	80.0	**74.7**	83.5
DaC [Zhang et al., 2022]		96.6	86.8	**86.4**	78.4	96.4	96.2	**93.6**	**83.8**	96.8	95.1	89.6	50.0	87.3
NRC++ [Yang et al., 2023b]		96.8	91.9	88.2	82.8	97.1	96.2	90.0	81.1	95.2	93.8	91.1	49.6	87.8
SFADA [He et al., 2024]		94.2	79.6	79.8	65.7	92.6	94.1	87.3	80.8	88.1	91.4	83.3	55.0	82.7
CPD [Zhou et al., 2024]		96.7	88.5	79.6	69.0	95.9	96.3	87.3	83.3	94.4	92.9	87.0	58.7	85.8
3C-GAN [Li et al., 2020]	✓	95.4	75.8	70.5	73.9	92.4	94.3	89.4	82.5	91.7	90.1	85.2	50.3	82.6
SFADA [He et al., 2024]	✓	94.9	80.4	80.6	68.2	94.3	94.0	86.5	82.1	91.4	92.2	85.0	51.4	83.4
DVD (Ours)	✓	**98.4**	**92.1**	83.9	**83.6**	**98.1**	**96.5**	92.1	82.9	**97.0**	**95.2**	**92.6**	**54.6**	**88.9** ± 0.6

- Across SFDA benchmarks, DVD consistently achieves **state-of-the-art** performance.
- On *VisDA-C,* which presents the largest visual domain gap, DVD achieves notable performance improvements over existing methods.
- The results indicates that DVD effectively reconstructs **source-like discriminative geometry** from the target features alone.

Beyond SFDA

Supervised Classification. Top-1 test error (%).

Method	CIFAR-10	CIFAR-100	ImageNet
ResNet-18	7.07	22.74	31.46
+ DVD	**6.52** ± 0.1	**22.01** ± 0.2	**31.06** ± 0.2
ResNet-50	6.35	22.23	24.68
+ DVD	**5.92** ± 0.2	**21.95** ± 0.3	**24.30** ± 0.3
VGG-16	7.36	28.82	25.82
+ DVD	**6.42** ± 0.1	**26.58** ± 0.3	**25.02** ± 0.3

Classification accuracy (%) for source-domain testing using DVD features.

Method	Office-31			Office-Home				VisDA-C 2017
	A	D	W	Ar	Cl	Pr	Rw	
ResNet	91.5	100.0	98.7	81.5	81.5	93.9	85.1	99.6
+ DVD	**96.9**	**100.0**	**100.0**	**96.9**	**90.6**	**100.0**	**96.9**	**100.0**

Comparison of the SFDA methods on VisDA-C 2017 (ResNet-101).

Method	PACS	VLCS	Office-H	DomainNet
ResNet-50 [He et al., 2016]	84.5	72.8	61.3	40.2
ZS-CLIP(C) [Radford et al., 2021]	90.6	76.0	68.6	45.6
CAD [Dubois et al., 2021]	90.0	81.2	70.5	45.5
ZS-CLIP(PC) [Radford et al., 2021]	90.7	80.1	72.0	46.3
PromptStyler [Cho et al., 2023]	93.2	**82.3**	73.6	49.5
DVD (Ours)	**93.8**± 0.4	81.9 ± 1.2	**74.5** ± 0.8	**50.8** ± 0.6

- **Benefits for the source provider:**
  - During source training, the diffusion drift acts as a **regularizer** that enforces smooth feature evolution within each class manifold.
  - Effectively reducing intra-class variance and promotes discriminative embeddings.
  - DVD enhances the source classifier itself, improving standard supervised classification and source-domain testing performance through latent-space augmentation.

- **Boosting Domain Generalization (DG):**
  - Beyond adaptation, DVD enhances model generalization in general.
  - By learning a deterministic drift field that maintains semantic consistency across varying conditions, the model acquires inherently domain-invariant representations.

Efficiency and Ablation Studies

Training cost and deployment efficiency on VisDA-C (ResNet-101, Nvidia V100).

Method	Epoch (s)	Convergence (s)	Inference (ms)	FPS	Latency (ms)
DaC [Zhang et al., 2022]	632.8±3.1	12,656.3±55.2	50.6±0.8	19.8±0.5	86.2±1.2
NRC++ [Yang et al., 2023b]	469.2±2.7	4,692.8±41.5	26.0±0.6	38.4±0.4	51.0±0.7
SHOT [Liang et al., 2020]	439.3±2.4	6,589.5±47.2	25.9±0.4	38.6±0.3	50.8±0.6
DVD (Ours)	516.3±2.9	5,163.1±43.6	26.1±0.5	38.2±0.2	51.1±0.8

Ablation on deterministic vs. stochastic drift on VisDA-2017.

Method	Target Acc. (%)	Runtime (second/iter)
Stochastic Drift (+ Noise)	87.2	2.5
Deterministic Drift (Ours)	**88.9**	**1.2**

Adaptation accuracy on VisDA-2017 under different diffusion steps.

Train Steps	Inference Steps	Target Accuracy (%)
8	8	88.5
16	**16**	**88.9**
32	32	88.2
64	64	87.6
100	100	87.1
100	8	85.1
8	100	86.2

Sensitivity analysis of vicinity parameters.

$(k_s^{dif}, k_t^{dif}, k_t)$	VisDA-2017 (%)	Office-Home (%)	DomainNet (%)
$(5, 5, 3)$	87.4	72.7	49.2
$(10, 10, 5)$	88.5	73.4	49.7
$\mathbf{(15, 15, 6)}$	**88.9**	**73.8**	**50.8**
$(20, 20, 10)$	88.3	73.1	50.0
$(25, 25, 12)$	88.0	72.8	49.6
$(30, 30, 15)$	87.8	72.6	49.5
$(40, 40, 20)$	87.5	72.3	49.1

- **No extra cost at inference:**
  - DVD achieves training efficiency similar to NRC++ and SHOT, despite adding a diffusion module.
  - Once adaptation is complete, the diffusion module is **removed**. Thus, DVD retains **real-time inference capability**.
- **Ablations on Design Choices:**
  - Deterministic formulation is not only simpler but a **more principled match** to discriminative goal.
  - Step schedule is important for **the transport mechanism** learned by DVD, which requires **consistent T** to preserve semantic alignment.
  - DVD is robust on different vicinity radii.
  - Flat response across hyperparameters suggests DVD does not depend on method-specific tuning.

Hyperparameter ablation on VisDA-C.

Batch Size		Learning Rate		Temperature	
Value	Acc.	Value	Acc.	Value	Acc.
32	88.2	1×10^{-4}	88.1	0.05	88.2
64	88.6	1×10^{-3}	88.6	0.07	88.5
128	**88.9**	$\mathbf{3 \times 10^{-3}}$	**88.9**	**0.13**	**88.9**
256	88.2	1×10^{-2}	88.3	0.20	88.6
512	87.6	1×10^{-1}	87.2	0.50	87.9
				0.80	87.1
				1.00	86.6

Summary

- **Rethinking Diffusion.** DVD transforms diffusion from pixel generation into latent knowledge transfer, aligning features across domains without source data.

- **Explicit Geometry Transfer.** DVD learns a deterministic drift field that restores label-consistent manifolds, enabling clear and interpretable adaptation.

- **Privacy-Preserving Learning.** DVD transfers decision geometry without sharing any source data, ensuring compliance with privacy constraints.

- **Lightweight Yet Powerful.** Our DVD achieves state-of-the-art SFDA performance with no added inference cost and robustness to hyperparameters.

- **Robust and Stable.** DVD demonstrates robustness across vicinity radii, drift steps, and architectures; deterministic design ensures stability.

- **Benefits both source and target.** Our DVD acts as a latent regularization improves source classifiers, and diffusion dynamics enhance domain generalization.