

1 Appendix Contents

- 2 A. Derivations
- 3 B. Algorithm
- 4 C. More Implementation Details
- 5 D. Additional Experimental Results
- 6 E. Discussion on Related Uncertainty Estimation Works
- 7 F. Limitations and Broader Impacts

8 A Derivations

9 A.1 Derivation of the Labeled Source Objective

10 Given a labeled source sample $(\mathbf{x}^s, \mathbf{y}_0^s)$, our goal is to inference the latent data representation \mathbf{z}^s and
 11 a sequence of latent class representations $\mathbf{y}_{1:T}^s$, which controls the generation process of data points
 12 and predictions, respectively. The log-likelihood of the labeled source data can thus be expressed by:

$$\log p(\mathbf{x}^s, \mathbf{y}_0^s) = \log \int p(\mathbf{x}^s, \mathbf{y}_{0:T}^s, \mathbf{z}^s) d\mathbf{y}_{1:T}^s d\mathbf{z}^s. \quad (1)$$

13 Since Eq. (1) is intractable to compute in practice, we then leverage variational inference to ap-
 14 proximate the posterior distribution of unknown variables $(\mathbf{z}^s, \mathbf{y}_{1:T}^s)$ and solve it by optimizing the
 15 following ELBO:

$$\log \int p(\mathbf{x}^s, \mathbf{y}_{0:T}^s, \mathbf{z}^s) d\mathbf{y}_{1:T}^s d\mathbf{z}^s \geq \mathbb{E}_{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p(\mathbf{y}_{0:T}^s, \mathbf{x}^s, \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \right], \quad (2)$$

16 where $q(\mathbf{z}^s, \mathbf{y}_{1:T}^s | \mathbf{x}^s, \mathbf{y}_0^s)$ is the approximation of the ground-truth joint posterior $p(\mathbf{z}^s, \mathbf{y}_{1:T}^s |$
 17 $\mathbf{x}^s, \mathbf{y}_0^s)$, which can be further factorized as:

$$q(\mathbf{z}^s, \mathbf{y}_{1:T}^s | \mathbf{x}^s, \mathbf{y}_0^s) = q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s) q(\mathbf{y}_{1:T}^s | \mathbf{y}_0, \mathbf{z}^s, \mathbf{x}^s). \quad (3)$$

18 With Eq. (3) and the generative process assumed in Eq. (??), we have the following derivation:

$$\begin{aligned} \log p(\mathbf{x}^s, \mathbf{y}_0^s) &= \log \int p(\mathbf{x}^s, \mathbf{y}_{0:T}^s, \mathbf{z}^s) d\mathbf{y}_{1:T}^s d\mathbf{z}^s \\ &\geq \mathbb{E}_{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p(\mathbf{y}_{0:T}^s, \mathbf{x}^s, \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \right] \\ &= \mathbb{E}_{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p(\mathbf{z}^s) p_\phi(\mathbf{x}^s | \mathbf{z}^s) p_\theta(\mathbf{y}_{0:T}^s | \mathbf{x}^s, \mathbf{z}^s)}{q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s) q(\mathbf{y}_{1:T}^s | \mathbf{y}_0, \mathbf{z}^s, \mathbf{x}^s)} \right] \\ &= \mathbb{E}_{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p(\mathbf{z}^s)}{q_\rho(\mathbf{z}^s | \mathbf{y}_0, \mathbf{x}^s)} + \log p_\phi(\mathbf{x}^s | \mathbf{z}^s) + \log \frac{p_\theta(\mathbf{y}_{0:T}^s | \mathbf{x}^s, \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s | \mathbf{y}_0^s, \mathbf{z}^s, \mathbf{x}^s)} \right] \\ &= \mathbb{E}_{q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p(\mathbf{z}^s)}{q_\rho(\mathbf{z}^s | \mathbf{y}_0^s, \mathbf{x}^s)} + \log p_\phi(\mathbf{x}^s | \mathbf{z}^s) \right] + \\ &\quad \mathbb{E}_{q(\mathbf{y}_{1:T}^s, \mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \left[\log \frac{p_\theta(\mathbf{y}_{0:T}^s | \mathbf{x}^s, \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s | \mathbf{y}_0^s, \mathbf{z}^s, \mathbf{x}^s)} \right] \\ &= \mathbb{E}_{q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} [\log p_\phi(\mathbf{x}^s | \mathbf{z}^s)] - D_{KL}(q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s) \| p(\mathbf{z}^s)) + \\ &\quad \underbrace{\mathbb{E}_{\mathbf{z}^s \sim q_\rho(\mathbf{z}^s | \mathbf{x}^s, \mathbf{y}_0^s)} \mathbb{E}_{q(\mathbf{y}_{1:T}^s | \mathbf{x}^s, \mathbf{z}^s, \mathbf{y}_0^s)} \left[\frac{p_\theta(\mathbf{y}_{0:T}^s | \mathbf{x}^s, \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s | \mathbf{y}_0^s, \mathbf{z}^s, \mathbf{x}^s)} \right]}_{\textcircled{1}}. \end{aligned} \quad (4)$$

19 In $\textcircled{1}$, the forward diffusion process $q(\mathbf{y}_{1:T}^s | \mathbf{x}^s, \mathbf{z}^s, \mathbf{y}_0^s)$ and the reverse diffusion process
 20 $p_\theta(\mathbf{y}_{0:T}^s | \mathbf{x}^s, \mathbf{z}^s)$ are still based on the original input \mathbf{x}^s . In this work, we make a simplifica-
 21 tion design to assume that the observed class variable \mathbf{y}_0^s and latent ones $\mathbf{y}_{1:T}^s$ are only conditioned

22 on the latent variable \mathbf{z}^s . We demonstrate the reasonability of this simplification in Appendix A.3.
 23 Consequently, we have

$$\textcircled{1} = \mathbb{E}_{q(\mathbf{y}_{1:T}^s | \mathbf{z}^s, \mathbf{y}_0^s)} \left[\frac{p_\theta(\mathbf{y}_{0:T}^s | \mathbf{z}^s)}{q(\mathbf{y}_{1:T}^s | \mathbf{y}_0^s, \mathbf{z}^s)} \right], \quad (5)$$

24 which has the same form of ELBO as the diffusion classifier in Eq. (??).

25 A.2 Expression of the Labeled Target Objective

26 For the labeled target data $(\mathbf{x}^t, \mathbf{y}_0^t)$, the unknown latent variables are the same as the labeled source
 27 data, and therefore the ELBO is analogous to that of the labeled source data. The only difference is
 28 that we use the target-specific decoder $p_\psi(\mathbf{x}^t | \mathbf{z}^t)$. We give the full expression as follows:

$$\begin{aligned} \log p(\mathbf{x}^t, \mathbf{y}_0^t) &= \log \int p(\mathbf{x}^t, \mathbf{y}_{0:T}^t, \mathbf{z}^t) d\mathbf{y}_{1:T}^t d\mathbf{z}^t \\ &\geq \mathbb{E}_{q(\mathbf{y}_{1:T}^t, \mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log \frac{p(\mathbf{y}_{0:T}^t, \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_{1:T}^t, \mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \right] \\ &= \mathbb{E}_{q(\mathbf{y}_{1:T}^t, \mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log \frac{p(\mathbf{z}^t) p_\psi(\mathbf{x}^t | \mathbf{z}^t) p_\theta(\mathbf{y}_{0:T}^t | \mathbf{x}^t, \mathbf{z}^t)}{q_\rho(\mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t) q(\mathbf{y}_{1:T}^t | \mathbf{y}_0^t, \mathbf{z}^t, \mathbf{x}^t)} \right] \\ &= \mathbb{E}_{q(\mathbf{y}_{1:T}^t, \mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log \frac{p(\mathbf{z}^t)}{q_\rho(\mathbf{z}^t | \mathbf{y}_0^t, \mathbf{x}^t)} + \log p_\psi(\mathbf{x}^t | \mathbf{z}^t) + \log \frac{p_\theta(\mathbf{y}_{0:T}^t | \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_{1:T}^t | \mathbf{y}_0^t, \mathbf{z}^t, \mathbf{x}^t)} \right] \\ &= \mathbb{E}_{q_\rho(\mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log \frac{p(\mathbf{z}^t)}{q_\rho(\mathbf{z}^t | \mathbf{y}_0^t, \mathbf{x}^t)} + \log p_\psi(\mathbf{x}^t | \mathbf{z}^t) \right] + \\ &\quad \mathbb{E}_{q(\mathbf{y}_{1:T}^t, \mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log \frac{p_\theta(\mathbf{y}_{0:T}^t | \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_{1:T}^t | \mathbf{y}_0^t, \mathbf{z}^t, \mathbf{x}^t)} \right] \\ &= \mathbb{E}_{q_\rho(\mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \left[\log p_\psi(\mathbf{x}^t | \mathbf{z}^t) \right] - D_{KL}(q_\rho(\mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t) \| p(\mathbf{z}^t)) + \\ &\quad \mathbb{E}_{\mathbf{z}^t \sim q_{rho}(\mathbf{z}^t | \mathbf{x}^t, \mathbf{y}_0^t)} \mathbb{E}_{q(\mathbf{y}_{1:T}^t | \mathbf{z}^t, \mathbf{y}_0^t)} \left[\frac{p_\theta(\mathbf{y}_{0:T}^t | \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_{1:T}^t | \mathbf{y}_0^t, \mathbf{z}^t, \mathbf{x}^t)} \right] := \mathcal{L}_{ELBO}^l. \end{aligned} \quad (6)$$

29 Analogously, we additionally impose the classifier $p_\omega(\mathbf{y}_0^t | \mathbf{z}^t)$ in the latent space to jointly train the
 30 source and target labeled data. The final training objective \mathcal{L}_t^l for labeled target data is therefore:

$$\mathcal{L}_t^l = -\mathcal{L}_{ELBO}^l + \mathbb{E}_{(\mathbf{x}^t, \mathbf{y}_0^t) \sim \mathcal{T}_t} \mathbb{E}_{q_\rho(\mathbf{z}^t | \mathbf{x}^t)} [-\log p_\omega(\mathbf{y}_0^t | \mathbf{z}^t)]. \quad (7)$$

31 A.3 Derivation of the Unlabeled Target Objective

32 For the unlabeled target data \mathbf{x}^t , our goal is to inference the low-dimensional latent embedding \mathbf{z}^t that
 33 induces a domain-invariant latent space \mathcal{Z} and the class label \mathbf{y}_0^t based on \mathbf{z}^t . Besides, we assume a
 34 meanfield distribution on $q(\mathbf{z}^t, \mathbf{y}_0^t | \mathbf{x}^t)$, which can then be factorized as:

$$q(\mathbf{z}^t, \mathbf{y}_0^t | \mathbf{x}^t) = q(\mathbf{z}^t | \mathbf{x}^t) q(\mathbf{y}_0^t | \mathbf{x}^t). \quad (8)$$

35 Therefore, we optimize the following ELBO by regarding \mathbf{z}^t and \mathbf{y}_0^t as unknowns:

$$\begin{aligned}
\log p(\mathbf{x}^t) &= \log \int p(\mathbf{x}^t, \mathbf{y}_0^t, \mathbf{z}^t) d\mathbf{y}_0^t d\mathbf{z}^t \\
&\geq \mathbb{E}_{q(\mathbf{y}_0^t, \mathbf{z}^t | \mathbf{x}^t)} \left[\log \frac{p(\mathbf{x}^t, \mathbf{y}_0^t, \mathbf{z}^t)}{q(\mathbf{y}_0^t, \mathbf{z}^t | \mathbf{x}^t)} \right] \\
&= \mathbb{E}_{q(\mathbf{y}_0^t, \mathbf{z}^t | \mathbf{x}^t)} \left[\log \frac{p(\mathbf{z}^t) p_\psi(\mathbf{x}^t | \mathbf{z}^t) p_\omega(\mathbf{y}_0^t | \mathbf{x}^t, \mathbf{z}^t)}{q_\rho(\mathbf{z}^t | \mathbf{x}^t) q(\mathbf{y}_0^t | \mathbf{x}^t)} \right] \\
&= \mathbb{E}_{q(\mathbf{y}_0^t, \mathbf{z}^t | \mathbf{x}^t)} \left[\log \frac{p(\mathbf{z}^t)}{q_\rho(\mathbf{z}^t | \mathbf{x}^t)} + \log p_\psi(\mathbf{x}^t | \mathbf{z}^t) + \log \frac{p_\omega(\mathbf{y}_0^t | \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_0^t | \mathbf{x}^t)} \right] \quad (9) \\
&= \mathbb{E}_{q_\rho(\mathbf{z}^t | \mathbf{x}^t)} [\log p_\psi(\mathbf{x}^t | \mathbf{z}^t)] - D_{KL}(q_\rho(\mathbf{z}^t | \mathbf{x}^t) \| p(\mathbf{z}^t)) + \\
&\mathbb{E}_{\mathbf{z}^t \sim q_\rho(\mathbf{z}^t | \mathbf{x}^t)} \mathbb{E}_{q(\mathbf{y}_0^t | \mathbf{x}^t)} \left[\log \frac{p_\omega(\mathbf{y}_0^t | \mathbf{x}^t, \mathbf{z}^t)}{q(\mathbf{y}_0^t | \mathbf{x}^t)} \right] \\
&= \mathbb{E}_{q_\rho(\mathbf{z}^t | \mathbf{x}^t)} [\log p_\psi(\mathbf{x}^t | \mathbf{z}^t)] - D_{KL}(q_\rho(\mathbf{z}^t | \mathbf{x}^t) \| p(\mathbf{z}^t)) - \\
&\mathbb{E}_{\mathbf{z}^t \sim q(\mathbf{z}^t | \mathbf{x}^t)} \underbrace{[D_{KL}(q(\mathbf{y}_0^t | \mathbf{x}^t) \| p_\omega(\mathbf{y}_0^t | \mathbf{x}^t, \mathbf{z}^t))] }_{\textcircled{2}} := \mathcal{L}_{ELBO}^u.
\end{aligned}$$

36 Since our deterministic classifier encodes the covariate-dependence between \mathbf{y}_0^t and \mathbf{z}^t , therefore, \mathbf{y}_0^t
37 is not depended on \mathbf{x}^t in our formulation, i.e., $p_\omega(\mathbf{y}_0^t | \mathbf{x}^t, \mathbf{z}^t) = p_\omega(\mathbf{y}_0^t | \mathbf{z}^t)$. $\textcircled{2}$ demonstrates that,
38 to maximize \mathcal{L}_{ELBO}^u , $D_{KL}(q(\mathbf{y}_0^t | \mathbf{x}^t) \| p_\omega(\mathbf{y}_0^t | \mathbf{z}^t)) \equiv 0$ should always be satisfied. On the other
39 hand, we empirically find that $p_\omega(\mathbf{y}_0^t | \mathbf{z}^t)$ can be a good approximation of $q(\mathbf{y}_0^t | \mathbf{x}^t)$ even when it
40 is solely based on latent \mathbf{z}^t . Therefore, we assume that the model output \mathbf{y}_0^t is only depended on the
41 latent embedding \mathbf{z}^t , which supports the derivation in Eq. (4).

42 B Algorithm

43 B.1 Algorithm of DAPM-TT for Conventional Active Domain Adaptation

44 The overall training and selection procedure of DAPM-TT for ADA is summarized in Algorithm 1.

Algorithm 1 Pseudo code of DAPM-TT for ADA

Require: Labeled source dataset \mathcal{S} , whole target dataset \mathcal{T} , unlabeled target dataset \mathcal{T}_u , labeled target dataset \mathcal{T}_l , total training rounds R , total annotation budget B , per round annotation budget b , step number per adaptation stage N_a , step number per diffusion stage N_d .

Ensure: Optimal model parameters $\{\theta, \rho, \phi, \psi, \omega, \tau\}$.

- 1: Initialize student model parameters $\{\rho, \omega\}$ and other parameters $\{\theta, \phi, \psi, \tau\}, \mathcal{T}^l = \emptyset, \mathcal{T}^u = \mathcal{T}$
 - 2: Initialize teacher model parameters $\Omega' = \{\rho, \omega\}$
 - 3: **for** $t = 1$ **to** R **do**
 - 4: **for** $i = 1$ **to** N_a **do**
 - 5: Update parameters $\{\rho, \phi, \psi, \omega, \tau\}$ via optimizing Eq. (??). % Adaptation Stage
 - 6: Update teacher model parameters Ω' with updated $\{\rho, \omega\}$ based on EMA.
 - 7: **end for**
 - 8: **for** $j = 1$ **to** N_d **do**
 - 9: Update diffusion classifier parameters θ via optimizing Eq. (??). % Diffusion Stage
 - 10: **end for**
 - 11: **if** $t \leq \frac{B}{b}$ **then**
 - 12: For each $\mathbf{x}^t \in \mathcal{T}_u$, generate N predictions $\{\tilde{\mathbf{y}}_n^t\}_{n=1}^N$ % Selection Stage
 - 13: Identify the two most predicted classes a, b for each \mathbf{x}^t .
 - 14: Conduct t-test between $\{\tilde{\mathbf{y}}_n^t[a]\}_{n=1}^N$ and $\{\tilde{\mathbf{y}}_n^t[b]\}_{n=1}^N$ and obtain the p-value for each \mathbf{x}^t .
 - 15: $Selected \leftarrow$ Select samples with top- b p-values from \mathcal{T}_u .
 - 16: $\mathcal{T}_u = \mathcal{T}_u \setminus Selected, \mathcal{T}_l = \mathcal{T}_l \cup Selected$.
 - 17: **end if**
 - 18: **end for**
 - 19: **return** Final model parameters $\{\theta, \rho, \phi, \psi, \omega, \tau\}$.
-

45 B.2 Algorithm of DAPM-TT for Source-Free Active Domain Adaptation

46 In SFADA, We do not use the domain classifier since the source domain data and target domain data
 47 cannot co-exist. In addition, we also use confident unlabeled target samples that are pseudo-labeled
 48 by the teacher model to substitute the source-labeled samples in \mathcal{L}_s in the diffusion stage. We denote
 49 the corresponding training loss by \mathcal{L}_t^p .

50 The overall training and selection procedure of DAPM-TT for SFADA is summarized in Algorithm 2.

Algorithm 2 Pseudo code of DAPM-TT for SFADA

Require: Labeled source dataset \mathcal{S} for pre-training, whole target dataset \mathcal{T} , unlabeled target dataset \mathcal{T}_u , labeled target dataset \mathcal{T}_l , total training rounds R , total annotation budget B , per round annotation budget b , step number per adaptation stage N_a , step number per diffusion stage N_d , step number of source pre-training N_s .

Ensure: Optimal model parameters $\{\theta, \rho, \phi, \psi, \omega\}$.

```

1: Initialize model parameters  $\{\rho, \omega\}$  and other parameters  $\{\theta, \phi, \psi\}, \mathcal{T}^l = \emptyset, \mathcal{T}^u = \mathcal{T}$ 
2: for  $i = 1$  to  $N_s$  do
3:   Update source model parameters  $\rho, \phi, \omega$  via optimizing  $\mathcal{L}_s$ .    % Source Pre-training
4: end for
5: Initialize teacher model parameters  $\Omega' = \{\rho, \omega\}$ 
6: for  $t = 1$  to  $R$  do
7:   for  $j = 1$  to  $N_a$  do
8:     Update parameters  $\{\rho, \psi, \omega\}$  via optimizing  $\mathcal{L}_t^u + \mathcal{L}_t^l$ .    % Adaptation Stage
9:     Update teacher model parameters  $\Omega'$  with updated  $\{\rho, \omega\}$  based on EMA.
10:  end for
11:  for  $k = 1$  to  $N_d$  do
12:    Update diffusion classifier parameters  $\theta$  via optimizing  $\mathcal{L}_t^p + \mathcal{L}_t^l$ .    % Diffusion Stage
13:  end for
14:  if  $t \leq \frac{B}{b}$  then
15:    For each  $\mathbf{x}^t \in \mathcal{T}_u$ , generate  $N$  predictions  $\{\tilde{\mathbf{y}}_n^t\}_{n=1}^N$     % Selection Stage
16:    Identify the two most predicted classes  $a, b$  for each  $\mathbf{x}^t$ .
17:    Conduct t-test between  $\{\tilde{\mathbf{y}}_n^t[a]\}_{n=1}^N$  and  $\{\tilde{\mathbf{y}}_n^t[b]\}_{n=1}^N$  and obtain the p-value for each  $\mathbf{x}^t$ .
18:     $Selected \leftarrow$  Select samples with top- $b$  p-values from  $\mathcal{T}_u$ .
19:     $\mathcal{T}_u = \mathcal{T}_u \setminus Selected, \mathcal{T}_l = \mathcal{T}_l \cup Selected$ .
20:  end if
21: end for
22: return Final model parameters  $\{\theta, \rho, \phi, \psi, \omega\}$ .

```

51 C More Implementation Details

52 C.1 Network Architecture

53 **Variational Autoencoder** The architecture of the VAE and the deterministic classifier is presented in
 54 detail in Fig. 1. The encoder comprises a pre-trained ResNet-50 backbone and three initialized linear
 55 layers with a ReLU activation following the first linear layer. We assume a Gaussian distribution for
 56 the latent embedding, and its mean and covariance are estimated by two separate linear layers based
 57 on the first linear layer. The decoder is a two-layer MLP that has the same output dimension as the
 58 backbone’s output. This encourages the decoder to reconstruct the feature generated by the backbone.

59 **Deterministic Classifier** The deterministic is simply a single layer linear classifier. We adopt the
 60 weight normalization technique on the classifier to stabilize the training.

61 **Diffusion Classifier** The diffusion classifier is conditioned on the latent embedding \mathbf{z} , the ground-
 62 truth \mathbf{y} , the guided information f_Ω and the time step t . We adopt the same architecture as [1] for class
 63 variable diffusion, except for the dimension of the input variable. For clarity, we describe the detailed
 64 model structure in Fig. 2 (a).

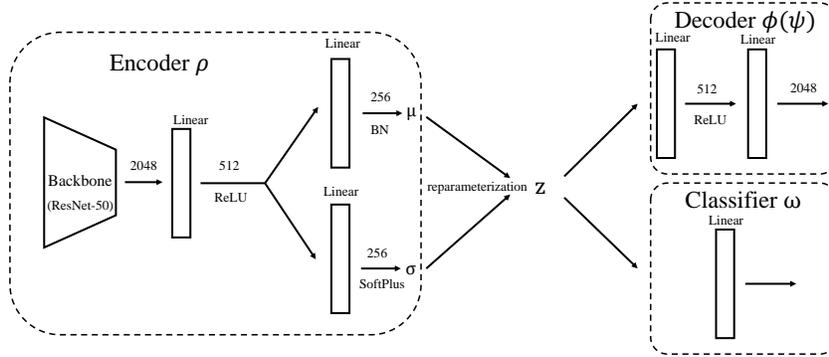


Figure 1: Architecture of the variational autoencoder and the deterministic classifier used in this work. The numbers on the data flow indicate the dimensions of the model output.

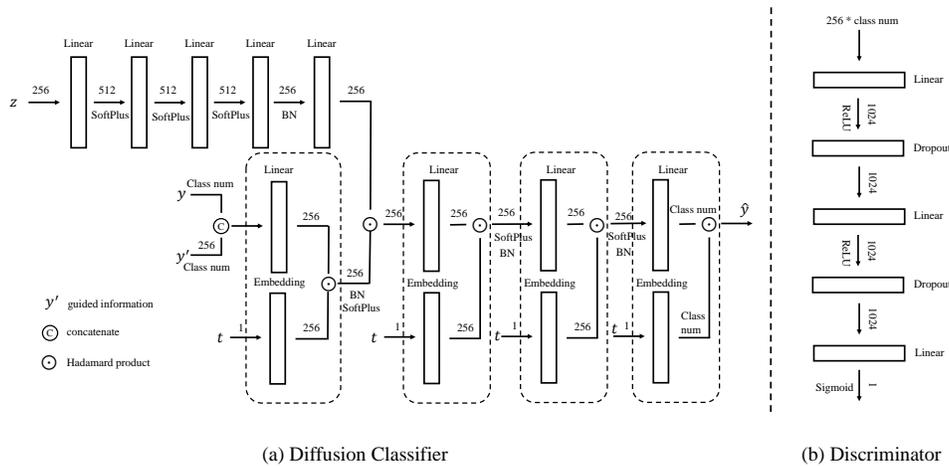


Figure 2: Architecture of (a) the diffusion classifier and (b) the domain discriminator.

65 **Domain Discriminator.** As shown in Fig. 2 (b). The domain discriminator we used is a three-layer
 66 MLP with a Dropout layer after the first and the second layers. And the output is a one-dimensional
 67 value with Sigmoid activation, which indicates the domainness of the sample.

68 C.2 Training Details

69 **Conventional Active Domain Adaptation.** In the adaptation stage, we utilize the SGD optimizer
 70 with a learning rate of 0.01, momentum of 0.9, and weight decay of 0.001. We set the EMA rate
 71 for the teacher model to 0.99. In the VAE objective, we assume a standard Gaussian distribution,
 72 $\mathcal{N}(\mathbf{0}, I)$, for the prior distribution of the latent variable $z^s(z^t)$. For Office-31 and VisDA, we use
 73 the features generated by a pre-trained ResNet-50 and freeze the backbone parameters to accelerate
 74 training and conserve memory. However, for Office-Home, which has more diverse categories, we
 75 jointly train the backbone with other modules to learn more specific category knowledge, and we
 76 set the learning rate to 0.001, which is 10 times lower than that of other models. For Office-31 and
 77 Office-Home, we conduct adaptation for 5 epochs and train the diffusion classifier for 10 epochs in
 78 each training round. The total number of training rounds is 20. For VisDA, we set the epoch number
 79 in each stage to 1, and the total number of training rounds is 10. To train the diffusion classifier, we
 80 use the Adam optimizer with a learning rate of 0.001 and epsilon of $1e-8$. The batch size is the same
 81 as that in the adaptation stage, and we use an EMA strategy with a rate of 0.9999 to update the model
 82 parameters. All experiments are conducted on a single RTX 3090 GPU.

83 **Source-Free Active Domain Adaptation.** In SFADA, we use the SGD optimizer without momentum
 84 and weight decay for adaptation. The learning rate is set to 0.01 for Office-31 and Office-Home, and
 85 0.001 for VisDA. As with ADA, we freeze the backbone for Office-31 and VisDA and open it for

86 Office-Home. We pre-train the source model for 10 epochs for VisDA and 30 epochs for the other
87 datasets. In each active learning round, we set the epoch numbers for the adaptation stage and the
88 diffusion stage to 5 and 10, respectively, for Office-31 and Office-Home, and both to one for VisDA.
89 The training details for the VAE, teacher model, and diffusion classifier are the same as in ADA.

90 C.3 Hyperparameter Choices of the Diffusion Classifier

91 Following [1], the hyperparameters of the diffusion classifier are set in a standard DDPM [2] manner.
92 Specifically, the number of diffusion timesteps T is set to 1000, and a linear noise schedule with
93 $\beta_1 = 0.0001$ and $\beta = 0.02$ is adopted for the forward diffusion process.

94 C.4 Implementation of Compared Baseline Methods

95 Note that for conventional ADA, we cite the results of previous AL methods and ADA methods
96 reproduced in [3] if the experimental settings are the same. For DUC [4] that does not report the
97 result on Office-31 dataset, we report the results by our own runs based on the code from the official
98 repository at <https://github.com/BIT-DA/DUC>. We have tuned some hyperparameters to ensure
99 the best results we can achieve.

100 For SFADA, we implement compared baseline algorithms on our DAPM baseline with following
101 details:

102 **Random.** We abandon the use of any selection strategy and randomly select samples from the
103 unlabeled target dataset \mathcal{T}_u for annotation.

104 **BvSB.** We compute the best-versus-second-best score based on the output of the deterministic
105 classifier for each unlabeled sample and select b samples with the lowest scores for annotation.

106 **Entropy.** We use the conditional entropy based on the output of the deterministic classifier to measure
107 the prediction confidence. And samples with highest entropy values are selected for annotation.

108 **CoreSet.** We regard the sample selection in each round as a core-set cover problem and solve it with
109 the code at https://github.com/ozansener/active_learning_coreset.

110 **BADGE.** We obtain the gradient vectors based on the pseudo labels generated by the deterministic
111 classifier and utilize K-Means++ on the gradient vectors for diverse sampling. The algorithm is
112 implemented based on the repository at <https://github.com/JordanAsh/badge>.

113 **ELPT.** We cite the results on Office-31 and Office-Home dataset from the original paper [5]. For
114 VisDA, we run this method and report the results on ResNet-50 backbone based on the official code at
115 <https://github.com/TL-UESTC/ELPT>.

116 D Additional Experimental Results

117 D.1 Accuracies of Confident Predictions under Different Thresholds

118 Fig. 3 illustrates the accuracies of the teacher model’s predictions for confident samples across
119 different threshold settings. At each training step, the teacher model is updated with the student
120 model, and we have computed the accuracy of the current batch and presented it as a curve. As
121 expected, increasing the threshold leads to an increase in the teacher model’s accuracy. However,
122 when the threshold is relatively high, only a small number of samples are considered confident at
123 the beginning, leading to a higher accuracy initially and a subsequent drop. It is worth noting that
124 the teacher model provided relatively reliable predictions at a threshold value of 0.9. Raising the
125 threshold further would result in too few confident samples, making 0.9 a more appropriate choice.

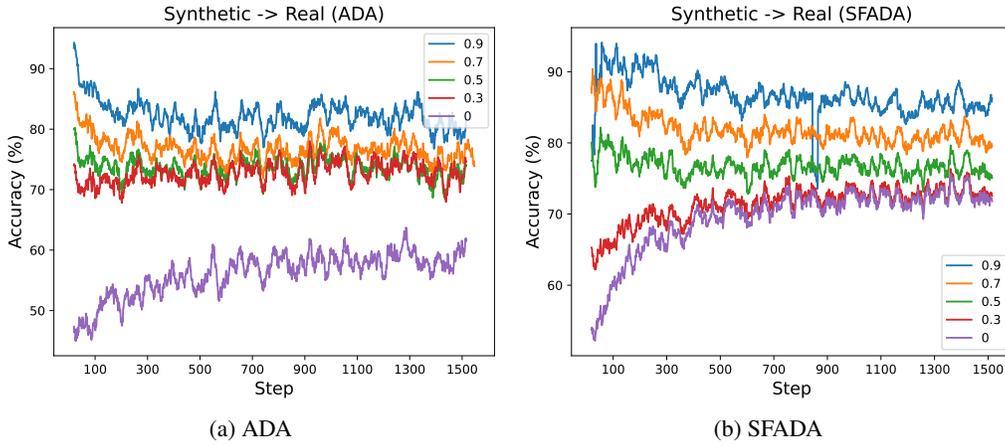
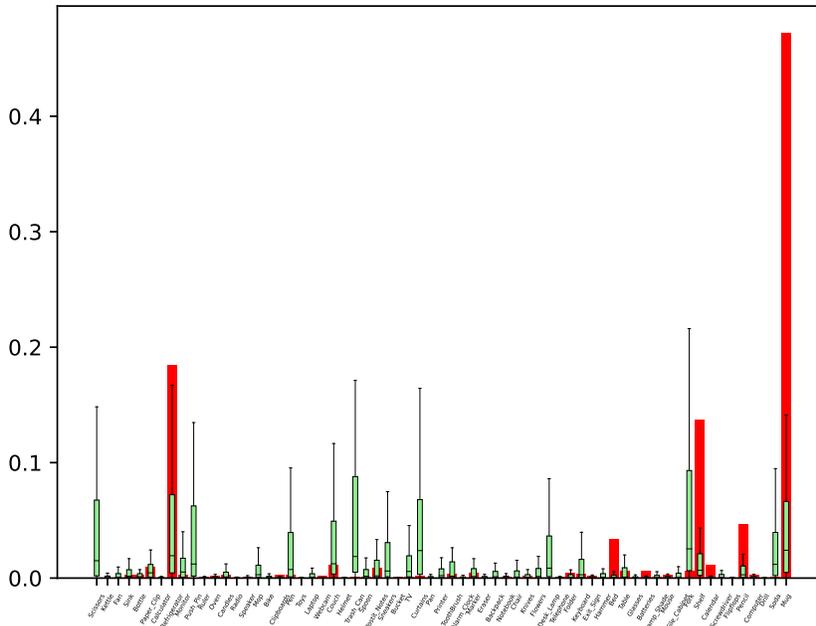


Figure 3: Accuracies of confident predictions made by the teacher model under different threshold values on VisDA dataset.

126 **D.2 Prediction Visualization of Different Classifiers**

127 To investigate the contrasting behaviors of deterministic and diffusive classifiers, we randomly select
 128 two samples on task Ar \rightarrow Cl, and visualize the predictions made by the deterministic classifier and
 129 the diffusion classifier ($N = 100$). As shown in Fig. 4a, the deterministic classifier exhibits high
 130 confidence in predicting a *refrigerator* as a *mug*. This verifies the overconfident issue in traditional
 131 softmax-based deterministic model, making it challenging for the active learning methods to detect
 132 the error and select such hard samples. In contrast, the diffusion classifier produces an uncertain
 133 prediction, indicating confusion in its output with a p-value of 0.873. As depicted in Fig. 3b, the
 134 deterministic classifier displays high uncertainty and misclassified the sample, whereas the diffusion
 135 classifier correctly classifies the sample with a p-value of 0.024, which saves budget and resources
 136 that would have been wasted on correcting the misclassification.

p-value: 0.873



(a) Refrigerator

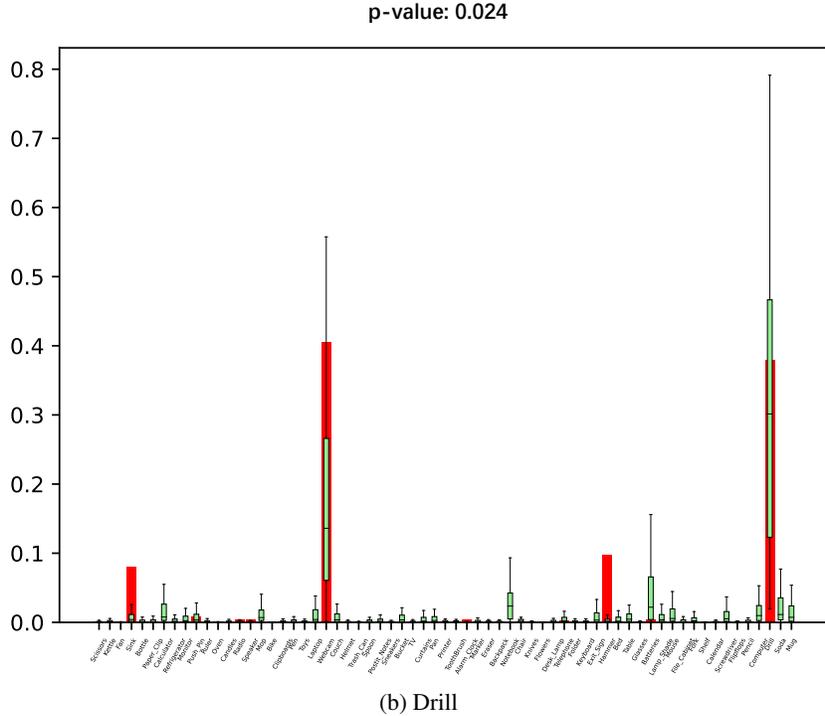


Figure 3: Visualization of predictions made by different classifiers for 2 randomly picked samples from *refrigerator* and *drill*, respectively. Red histograms represent the prediction of the deterministic classifier and green boxes denote the predictions of the diffusion classifier.

137 D.3 Effect of the Modeling of the Latent Feature Distribution

138 In comparison to CARD [1] which utilizes a deterministic feature extraction network and a diffusion
 139 classifier to evaluate uncertainty based on the original image, our method employs an additional VAE
 140 to model data uncertainty in a low-dimensional latent space, and the diffusion classifier is based
 141 on latent variables. To investigate the benefits of this improvement for ADA tasks, we implement
 142 CARD in the ADA task and report the results of it on Office-31 and VisDA. Specifically, we use a
 143 deterministic ResNet-50 network as the feature extractor and train a deterministic classifier on top of
 144 it to guide the diffusion classifier. The diffusion classifier takes the original image x as one of the
 145 inputs and uses the same independent two-sample t-test-based criterion for sample selection. We
 146 denote this implementation by CARD-TT. As shown in Table 1, DAPM-TT significantly outperforms
 147 CARD-TT on both datasets. It is exciting to see that although VAE is mainly designed for modeling
 148 the uncertainty of the data generation process, it results in a significant improvement with respect to
 149 accuracy. We conjecture that the reason for this improvement is two-fold: firstly, domain shift leads
 150 to a significant distribution shift in the image space, and in such case, CARD fails to work as intended
 151 [1]. This effect is mitigated to a certain extent in the low-dimensional and less noisy latent space.
 152 Secondly, in VAE training, we use the same prior distribution, i.e., the standard Gaussian distribution,
 153 for the latent variables of data in both the source and target domains. This design draws the samples in
 154 both domains closer to the standard Gaussian distribution, thereby achieving an indirect distribution
 155 alignment.

156 D.4 Qualitative Analysis on Selected Samples

157 We present in Fig. 3 a list of all the selected samples by our approach to gain insight into which
 158 samples are chosen. Intuitively, our method tends to select samples that are challenging for the model,
 159 such as those with complex backgrounds or different styles from the other images in the dataset.
 160 Labeling these samples can help reduce the ambiguity in the model and enable it to better capture
 161 the semantic aspects of the category. Interestingly, we observe that our approach naturally selects
 162 a diverse range of sample classes, even though we did not explicitly impose a diversity constraint.
 163 Furthermore, we observe that the p-values of the selected samples gradually decrease with each

Table 1: Comparison between probabilistic and deterministic feature extractors, and between different t-test strategies on ADA task (ResNet-50). w/o and w/ are short for without and with, respectively.

Category	Method	Office-Home						VisDA	
		A→D	A→W	D→A	D→W	W→A	W→D	Avg	Synthetic→Real
w/o adaptaion stage	CARD-TT	95.1	94.2	78.5	98.7	78.2	99.1	90.6	84.6
	DAPM-TT	96.1	95.9	79.5	98.7	79.2	99.1	91.4	86.3
w/ adaptaion stage	DAPM-TT*	96.8	97.8	83.3	99.8	81.7	100	93.2	88.6
	DAPM-TT	96.8	98.6	82.3	99.8	83.3	100	93.5	89.1

164 training round, and by the 5th round, the lowest p-value is 0.518. This indicates that selecting
 165 approximately 5% of the samples is sufficient to mitigate much of the ambiguity in the model.

166 D.5 Comparison between Different T-test Strategies

167 Based on the scores of the two most probable classes predicted by the diffusion classifier, we can
 168 use either paired two-sample t-test or independent two-sample t-test for selection, which correspond
 169 to different assumptions for the generation of predictions. The former assumes that the scores of
 170 different classes are generated in pairs, while the latter assumes that they are generated independently.
 171 For a paired t-test, the t-value of a target sample \mathbf{x}^t is calculated as follows:

$$t = (\bar{d} - \mu_d) / (s_d / \sqrt{N}), \quad (10)$$

172 where $\bar{d} = \frac{1}{N} \sum d_i$ is the mean of sample differences $d_i = \tilde{\mathbf{y}}_i^t[a] - \tilde{\mathbf{y}}_i^t[b]$, μ_d is the difference of the
 173 null hypothesis (usually set as 0), and $s_d = \sqrt{\sum (d_i - \bar{d})^2 / (N - 1)}$ is the standard deviation of the
 174 sample difference.

175 We denote our method with paired t-test-based criterion by DAPM-TT*, and report the results on
 176 Office-31 and VisDA in Table 1. Empirically, we find that independent two-sample t-test yields
 177 superior performance on both datasets. We conjecture the reason might be that the independent two-
 178 sample t-test considers the internal variance of each group of samples. Therefore, when evaluating
 179 uncertainty, it considers an additional dimension compared to the paired t-test. In this work, we adopt
 180 independent two-sample t-test for all experiments.

181 D.6 T-SNE Visualization of Latent Representations

182 We visualize the latent representations of unlabeled target data and selected target data in Fig. 4
 183 using t-SNE [6]. In this visualization experiment, we compared our DAPM-TT with BvSB [7] that
 184 is based on the deterministic classifier. It can be observed that BvSB tends to select samples from
 185 relatively ambiguous regions (the center region) since these samples often have ambiguity between
 186 different classes. However, many samples selected by BvSB are in areas where the model is able to
 187 make predictions accurately. Therefore, it will not help to correct the samples with wrong predictions,
 188 resulting in modest improvement on performance. Our DAPM-TT, on the other hand, can select
 189 samples from both the regions where there is ambiguity between classes and the regions where a large
 190 number of samples are misclassified, which are exactly the ones we want to select for annotation.

191 E Discussion on Related Uncertainty Estimation Works

192 In machine learning, there are primarily two kinds of uncertainties that are studied, i.e., *epistemic*
 193 *uncertainty* which arises due to a lack of knowledge or data and can be reduced with more data or
 194 improved models, and *aleatoric uncertainty* that stems from the inherent randomness in the data [8].
 195 To model these uncertainties, the community has proposed many Bayesian deep learning methods.
 196 From a Bayesian perspective, these two uncertainties can be modeled by the posterior of model
 197 parameters W and outputs \mathbf{y} , respectively, using the following formulation:

$$P(\mathbf{y} | \mathbf{x}, \mathcal{D}) = \int \underbrace{P(\mathbf{y} | \mathbf{x}, W)}_{\text{aleatoric uncertainty}} \underbrace{P(W | \mathcal{D})}_{\text{epistemic uncertainty}} dW. \quad (11)$$

198 The family of Bayesian neural networks (BNNs) [9, 10, 11] is specifically designed to capture
199 epistemic uncertainty by assuming a probability distribution over the network parameters. This
200 involves estimating the posterior distribution over the parameters of the neural network given the
201 observed data. However, due to the intractable form of the posterior, BNNs are often trained using
202 appropriate approximations like Markov Chain Monte Carlo (MCMC) or Variational Inference
203 (VI). Another approximation for BNNs is Monte Carlo Dropout [12], which assumes a Bernoulli
204 distribution over network parameters. During inference, the output of the network is averaged
205 over multiple stochastic forward passes with dropout enabled, resulting in a distribution over the
206 predictions.

207 Evidential deep learning (EDL) [13, 4] is another Bayesian method that models uncertainty associated
208 with the output of the model based on evidential theory. In EDL, this is typically achieved by using a
209 distributional output, such as a Dirichlet distribution over class probabilities for classification tasks,
210 instead of a point estimate.

211 For non-Bayesian methods, ensemble-based methods [14, 15] have been proposed to model predictive
212 uncertainty by combining multiple deterministic neural networks with different initializations. How-
213 ever, all these methods are designed to capture either epistemic uncertainty or aleatoric uncertainty
214 alone by modeling probability distributions over the model parameters and outputs, respectively.
215 Moreover, they still impose a restricted form of distributions, such as Gaussian or Dirichlet, which
216 limits their applicability in practice.

217 To capture both sources of uncertainties in a single model, Kendall et al. [8] propose modeling
218 aleatoric uncertainty in the model outputs beyond model parameters by predicting the noise term for
219 the output variable of each sample as part of the model output. However, the form of the noise is still
220 assumed to be Gaussian.

221 Recently, Han [1] proposed modeling the implicit output distribution by leveraging the generative
222 capability of the diffusion model. However, they only model aleatoric uncertainty in their formulation
223 since the model they use is still a deterministic neural network, and the proposed method, CARD,
224 only enables in-distribution generalization.

225 In addition to modeling aleatoric uncertainty with the diffusion classifier, we also incorporate a VAE
226 to model the underlying data generation process. The VAE learns a probabilistic distribution over the
227 latent space, which represents the model’s uncertainty about the true underlying distribution of the
228 data, given the limited amount of training data. As more data is provided during training, the learned
229 distribution should converge to the true underlying distribution, reducing epistemic uncertainty.
230 Therefore, our DAPM also offers a way to measure epistemic uncertainty. Furthermore, our diffusion
231 classifier is conditioned on the latent variables in \mathcal{Z} rather than the ones in the original image space
232 \mathcal{X} . We argue that the latent space contains less noise and is more suitable for cross-domain tasks.

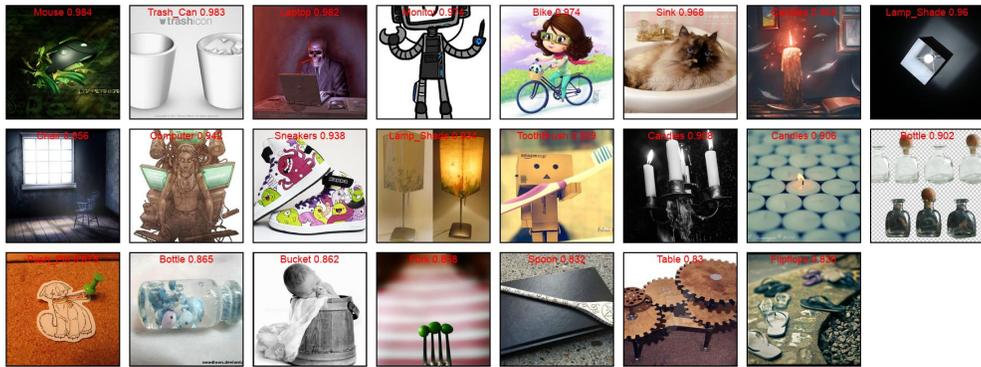
233 F Limitations and Broader Impacts

234 **Limitations.** Our work presents a way to recover the predictive distribution of deep models by
235 combining the power of diffusion models and variational autoencoders. However, like any research,
236 our study may have some limitations that should be acknowledged. Firstly, the task we focus on is
237 limited to image classification in this work. In our future study, we may extend the scope of research
238 to other areas like image segmentation and object detection, etc. Secondly, probabilistic models are
239 often more difficult to interpret compared to deterministic models. It is important to study insights in
240 future research to help interpret probabilistic models.

241 **Broader Impacts.** Indeed, the impact of active domain adaptation is significant, especially in
242 scenarios where labeled data is scarce in the target domain. The ability to adapt to new domains with
243 limited labeled data can potentially reduce the time and cost required to gather labeled data for each
244 specific task, thus making the deployment of machine learning models more accessible and cost-
245 effective. This can also facilitate the development of more robust and generalizable machine learning
246 models that can be used across multiple domains, which is particularly important for applications that
247 operate in dynamic and diverse environments. Overall, our work contributes to advancing the field of
248 machine learning and promoting the development of more efficient and adaptive technologies.

249 **References**

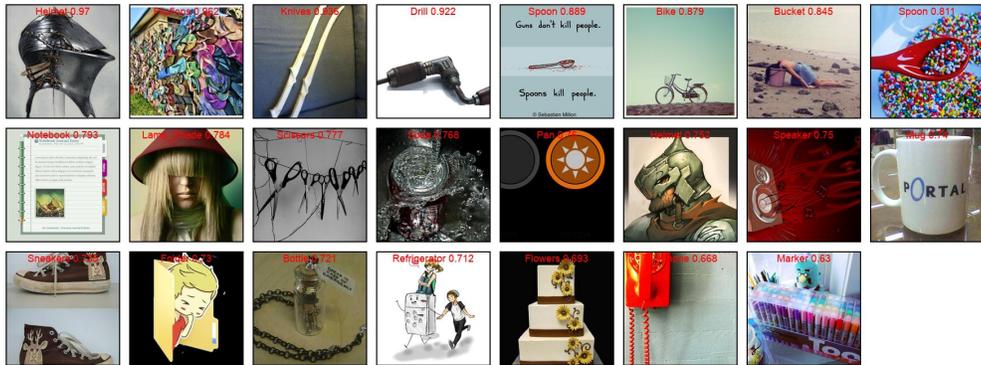
- 250 [1] Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. Card: Classification and regression
251 diffusion models. In *NIPS*, volume 35, pages 18100–18115, 2022.
- 252 [2] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NIPS*,
253 volume 33, pages 6840–6851, 2020.
- 254 [3] Binhui Xie, Longhui Yuan, Shuang Li, Chi Harold Liu, Xinjing Cheng, and Guoren Wang.
255 Active learning for domain adaptation: An energy-based approach. In *AAAI*, volume 36, pages
256 8708–8716, 2022.
- 257 [4] Mixue Xie, Shuang Li, Rui Zhang, and Chi Harold Liu. Dirichlet-based uncertainty calibration
258 for active domain adaptation. In *ICLR*, 2023.
- 259 [5] Xinyao Li, Zhekai Du, Jingjing Li, Lei Zhu, and Ke Lu. Source-free active domain adaptation
260 via energy-based locality preserving transfer. In *MM*, pages 5802–5810, 2022.
- 261 [6] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11), 2008.
- 262 [7] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for
263 image classification. In *CVPR*, pages 2372–2379. IEEE, 2009.
- 264 [8] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for
265 computer vision? volume 30, 2017.
- 266 [9] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty
267 in neural network. In *ICML*, pages 1613–1622. PMLR, 2015.
- 268 [10] José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable
269 learning of bayesian neural networks. In *ICML*, pages 1861–1869. PMLR, 2015.
- 270 [11] Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameter-
271 ization trick. *Advances in neural information processing systems*, 28, 2015.
- 272 [12] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model
273 uncertainty in deep learning. In *ICML*, pages 1050–1059. PMLR, 2016.
- 274 [13] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify
275 classification uncertainty. In *NIPS*, volume 31, 2018.
- 276 [14] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable
277 predictive uncertainty estimation using deep ensembles. In *NIPS*, volume 30, 2017.
- 278 [15] Shiwei Liu, Tianlong Chen, Zahra Atashgahi, Xiaohan Chen, Ghada Sokar, Elena Mocanu,
279 Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Deep ensembling
280 with no overhead for either training or testing: The all-round blessings of dynamic sparsity.
281 *arXiv preprint arXiv:2106.14568*, 2021.



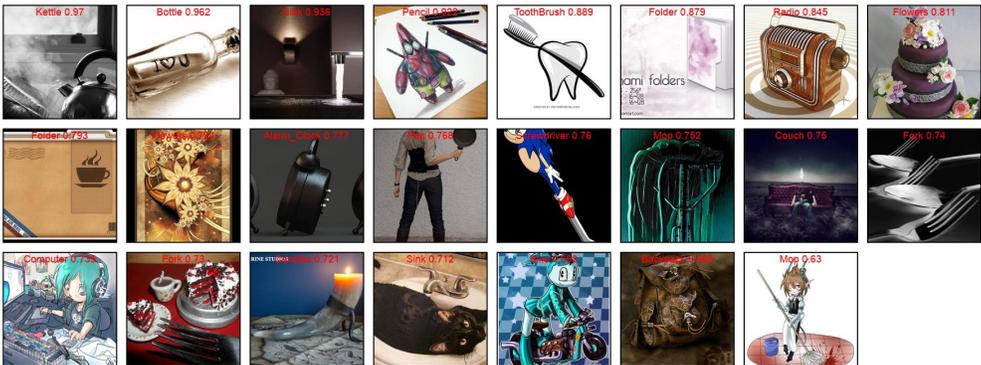
(a) Round 1



(b) Round 2



(c) Round 3



(d) Round 4



(e) Round 5

Figure 3: All selected samples on task $Cl \rightarrow Ar$ (ADA). For samples in each round, the priority (p-value) gradually decreases from top left to bottom right.

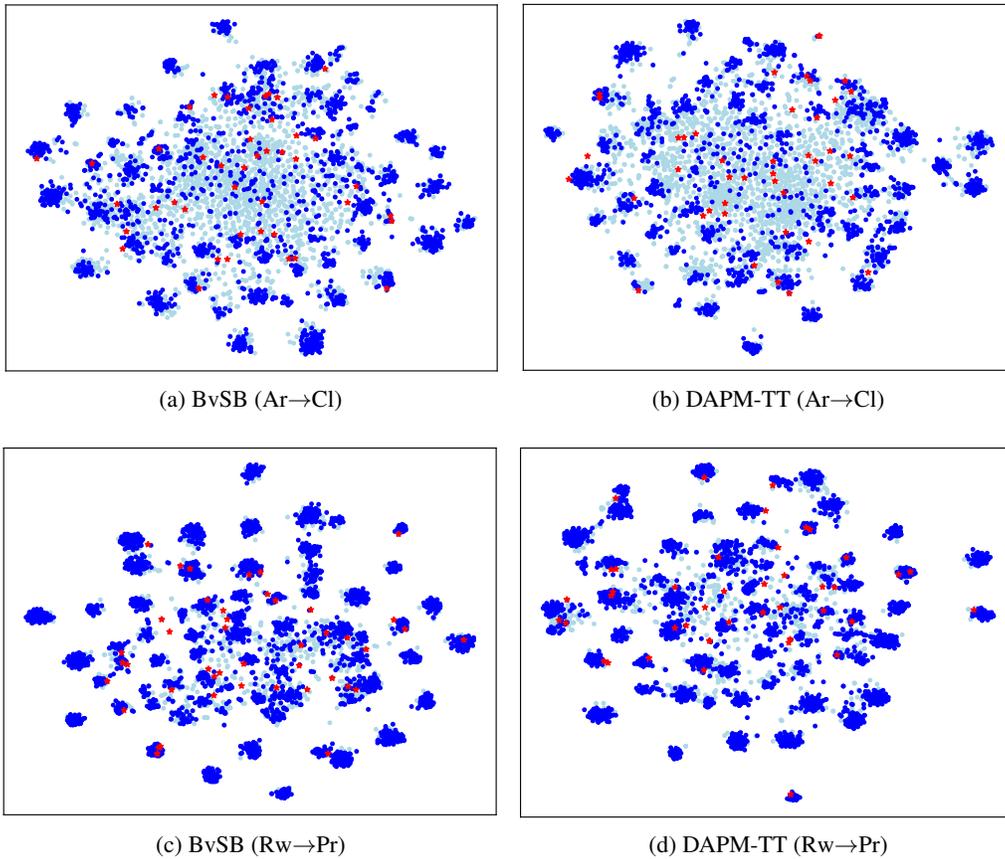


Figure 4: Visualization of latent representations using t-SNE [6] on task $Ar \rightarrow Cl$ (a to b) and $Rw \rightarrow Pr$ (c to d) of ADA. Darkblue points are unlabeled target samples correctly classified by our model. Lightblue points represent unlabeled target samples misclassified by our model. Red stars are the selected target samples.