

---

# Automatic Domain Adaptation by Transformers in In-Context Learning

---

Ryuichiro Hataya<sup>1</sup> Kota Matsui<sup>2</sup> Masaaki Imaizumi<sup>3</sup><sup>1</sup>

## Abstract

Selecting or designing an appropriate domain adaptation algorithm for a given problem remains challenging. This paper presents a Transformer model that can provably approximate and opt for unsupervised domain adaptation (UDA) methods for a given dataset in the in-context learning framework, where a foundation model performs new tasks without updating its parameters at test time. Specifically, we prove that i) Transformers can approximate instance-based and feature-based unsupervised domain adaptation algorithms, and ii) automatically select the approximated algorithms suited for a given dataset. Numerical results indicate that in-context learning demonstrates an adaptive domain adaptation surpassing existing methods.

## 1. Introduction

Domain adaptation provides a methodology for transferring “knowledge” obtained in one domain to another related domain (Ben-David et al., 2010). One of the challenges associated with domain adaptation is selecting and differentiating effective methods. There are many existing approaches in domain adaptation, such as the instance-based methods (Dai et al., 2007; Sugiyama et al., 2007; Kanamori et al., 2009) and the feature-based methods (Daumé III, 2009; Ganin et al., 2016). However, when each method is effective differs based on data. Specifically, the instance-based methods (Sugiyama et al., 2007; Kanamori et al., 2009) are effective when there exists a valid density ratio between covariates of the source and target domains, whilst the feature-based methods are effective if we can find a domain-invariant representation on a common feature space between the domains. To properly select an appropriate approach to the given data, it is essential to assess whether they meet the specific conditions of each method.

<sup>\*</sup>Equal contribution <sup>1</sup>RIKEN AIP, Tokyo, Japan <sup>2</sup>Nagoya University, Aichi, Japan <sup>3</sup>The University of Tokyo, Tokyo, Japan. Correspondence to: Ryuichiro Hataya <ryuichiro.hataya@riken.jp>.

*Proceedings of the 1<sup>st</sup> Workshop on In-Context Learning at the 41<sup>st</sup> International Conference on Machine Learning*, Vienna, Austria, 2024. Copyright 2024 by the author(s).

As a method for developing versatile algorithms, the in-context learning ability of foundation models has garnered significant attention. In-context learning, a form of meta-learning in foundational models, allows models to adapt to new tasks without updating their parameters. In particular, the capabilities of Transformers in in-context learning have been intensively surveyed: for example, (Bai et al., 2023) showed that Transformers can learn algorithms such as gradient descent, as well as their optimal selection methods. Such in-context learning abilities of Transformers have been validated from both experimental and theoretical aspects (Garg et al., 2022; Li et al., 2023; Von Oswald et al., 2023; Akyürek et al., 2022; Xie et al., 2021; Zhang et al., 2023; Bai et al., 2023; Lin et al., 2023; Ahn et al., 2024; Raventós et al., 2024).

In this study, we demonstrate that Transformer models effectively address the challenge in in-context learning. Specifically, our analysis reveals that Transformers in the framework i) solve domain adaptation problems by approximating the main UDA algorithms, and ii) automatically select suitable methods adaptively to the dataset’s characteristics. The results indicate that Transformers can, in context, not only implement various transfer learning methods but also possess the capability to appropriately select among them. These findings suggest that by accurately choosing a combination of these methods based on the dataset, performance can be enhanced beyond what is achievable by applying the methods individually. The full version of this manuscript can be found at <https://arxiv.org/abs/2405.16819>.

## 2. Preliminary

### 2.1. Unsupervised Domain Adaptation

**Setup** We consider the *Unsupervised domain adaptation* problem with two domains. Let  $\mathcal{X} \subset \mathbb{R}^d$  be a compact input space and  $\mathcal{Y}$  be an output space. Then, the source and the target domains,  $P_S$  and  $P_T$ , are distributions on  $\mathcal{X} \times \mathcal{Y}$ , let their density functions be  $p_S$  and  $p_T$ , and their marginal distributions on  $\mathcal{X}$  be  $P_S^X$  and  $P_T^X$ . Denote  $\mathcal{D}_S = \{(\mathbf{x}_i^S, y_i^S)\}_{i=1}^n$  and  $\mathcal{D}_T = \{\mathbf{x}_i^T\}_{i=1}^{n'}$  for the source labeled data and target unlabeled data, respectively and let  $N := n + n'$ . Given a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$  and a hypothesis space  $\mathcal{F}$ , UDA aims to minimize the target risk  $\operatorname{argmin}_{f \in \mathcal{F}} R_T(f)$ ,

where  $R_T(f) := \mathbb{E}_{(\mathbf{x}, y) \sim P_T}[\ell(f(\mathbf{x}), y)]$  is the target risk, without observing labels of the data from the target domain. To this end, UDA methods aim to match the source and target distributions in the feature space, categorized into instance-based and feature-based approaches.

**Instance-based Methods** Instance-based approaches reweight instances from the source domain to minimize the target risk, based on the following transferability assumption called *covariate shift*.

**Assumption 1** (Covariate Shift, (Shimodaira, 2000)). *Suppose that the distribution of  $y$  conditioned on the input  $\mathbf{x}$  is the same in the source and target domains, i.e.,  $p_S(y|\mathbf{x}) = p_T(y|\mathbf{x})$ .*

Under this assumption,  $R_T(f) = \mathbb{E}_{(\mathbf{x}, y) \sim P_S}(q(\mathbf{x})\ell(f(\mathbf{x}), y))$  holds, where  $q(\mathbf{x}) = p_T(\mathbf{x})/p_S(\mathbf{x})$  is the density ratio between the marginal distributions of source and target. Then, *Importance-weighted learning (IWL)* (Sugiyama et al., 2012; Kimura & Hino, 2024) obtains  $\operatorname{argmin}_{f \in \mathcal{F}} R_T(f)$  in two steps: 1. Empirically learn an estimator of the density ratio  $\hat{q}$  using  $\{\mathbf{x}_i^S\}_{i=1}^n$  and  $\mathcal{D}_T$ , and then, 2. minimize the weighted empirical risk  $\hat{R}_T(f) = \frac{1}{n} \sum_{i=1}^n \hat{q}(\mathbf{x}_i^S)\ell(f(\mathbf{x}_i^S), y_i^S)$  consists of data from the source domain.

Here, we introduce *unconstrained Least-Squares Importance Fitting (uLSIF)* (Kanamori et al., 2009), a method for density-ratio estimation. Firstly, the density ratio is estimated by the linear basis function model  $\hat{q}_\alpha(\mathbf{x}) = \alpha^\top \phi(\mathbf{x}) = \sum_{j=1}^J \alpha_j \phi_j(\mathbf{x})$ , where  $\alpha = [\alpha_1, \dots, \alpha_J]$  and  $\phi = [\phi_1, \dots, \phi_J]$  are parameter and feature vectors respectively, where  $\phi_j : \mathcal{X} \mapsto \mathbb{R}$  are feature maps for an input  $\mathbf{x}$ . uLSIF estimates the weight  $\alpha$  to directly minimize the following squared error:  $L(\alpha) = \frac{1}{2} \int_{\mathcal{X}} (\hat{q}_\alpha(\mathbf{x}) - q(\mathbf{x}))^2 p_S(\mathbf{x}) d\mathbf{x}$ . By some calculation, a minimizer of (2.1) is estimated  $\hat{\alpha} = \operatorname{argmin}_{\alpha \geq 0} \hat{L}(\alpha)$  with

$$\hat{L}(\alpha) = \frac{1}{2n} \sum_{i=1}^n (\hat{q}_\alpha(\mathbf{x}_i^S))^2 - \frac{1}{n'} \sum_{i=1}^{n'} \hat{q}_\alpha(\mathbf{x}_i^T) + \frac{\lambda}{2} \|\alpha\|_2^2 \quad (1)$$

where  $\lambda > 0$  is a regularization parameter,  $\Psi_{jj'} := \frac{1}{n} \sum_{i=1}^n \phi_j(\mathbf{x}_i^S)\phi_{j'}(\mathbf{x}_i^S)$  and  $\psi_j := \frac{1}{n'} \sum_{i=1}^{n'} \phi_j(\mathbf{x}_i^T)$ . Then, if  $\mathcal{F}$  is a space of linear model, we can obtain a classifier with an empirical importance-weighted problem  $\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^J} \hat{R}_T(\mathbf{w})$  with  $\hat{R}_T(\mathbf{w}) := \sum_{(\mathbf{x}, y) \in \mathcal{D}_S} \hat{q}_\alpha(\mathbf{x})\ell(\mathbf{w}^\top \phi(\mathbf{x}), y)$ . Then, we obtain a classifier defined as  $\hat{f}^{\text{IWL}}(\mathbf{x}) := (\mathbf{w}^*)^\top \phi(\mathbf{x})$ .

**Feature-based Methods** A feature-based method learns a domain-invariant feature map  $\phi : \mathcal{X} \rightarrow \mathcal{X}'$ , where  $\mathcal{X}'$  is a feature space, so that  $R_T(f') \approx \mathbb{E}_{(\mathbf{x}, y) \sim P_S} \ell(f'(\phi(\mathbf{x})), y)$ , where  $f' : \mathcal{X}' \rightarrow \mathcal{Y}$ . The *domain adversarial neural network (DANN)* (Ganin et al., 2016) is a typical method to

achieve this by using adversarial learning. DANN consists of three modules, a feature extractor  $f_F : \mathcal{X} \rightarrow \mathcal{X}'$ , a label classifier  $f_L : \mathcal{X}' \rightarrow [0, 1]$ , and a domain discriminator  $f_D : \mathcal{X}' \rightarrow [0, 1]$ , parameterized by  $\theta_F, \theta_L, \theta_D$ , respectively. Then, a classifier  $f_L$  with an invariant feature extractor  $f_F$  can be obtained by solving the following minimax problem:

$$\min_{\theta_F, \theta_L} \max_{\theta_D} L(\theta_F, \theta_L) - \lambda \Omega(\theta_F, \theta_D). \quad (2)$$

where  $\lambda > 0$  is a regularization parameter,  $L(\theta_F, \theta_L) := \sum_{(\mathbf{x}, y) \in \mathcal{D}_S} \gamma(f_L \circ f_F(\mathbf{x}), y)$  is label classification loss and  $\Omega(\theta_F, \theta_D) := \frac{1}{n} \sum_{(\mathbf{x}, y) \in \mathcal{D}_S} \gamma(f_D \circ f_F(\mathbf{x}), 0) + \frac{1}{n'} \sum_{(\mathbf{x}, y) \in \mathcal{D}_T} \gamma(f_D \circ f_F(\mathbf{x}), 1)$  is domain classification loss. Here,  $\gamma(p, q) := -q \log p - (1 - q) \log(1 - p)$  for  $p, q \in [0, 1]$  is a sigmoid cross entropy function. The model parameters are updated by gradient descent with a learning rate of  $\eta$  as follows:

$$\theta_F \leftarrow \theta_F - \eta \nabla_{\theta_F} (L(\theta_F, \theta_L) - \lambda \Omega(\theta_F, \theta_D)), \quad (3)$$

$$\theta_L \leftarrow \theta_L - \eta \nabla_{\theta_L} L(\theta_F, \theta_L), \quad (4)$$

$$\theta_D \leftarrow \theta_D - \eta \lambda \nabla_{\theta_D} \Omega(\theta_F, \theta_D). \quad (5)$$

## 2.2. In-context Learning

**Setup.** In in-context learning, a fixed Transformer observes a dataset  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N \sim (P)^N$  with pairs of an input  $\mathbf{x}_i$  and its label  $y_i$  from a joint distribution  $P$  and a new query input  $\mathbf{x}_*$  then predicts  $y_*$  corresponding to  $\mathbf{x}_*$ . Different from the standard supervised learning, in in-context learning, the Transformer is pre-trained on other datasets  $\mathcal{D}'$  from different distributions to learn an algorithm to predict  $y_*$ . Unlike the standard meta learning, at the inference time of in-context learning, the parameters of a Transformer are fixed. Our interest is to study the expressive power of a Transformer model for algorithms on a given dataset  $\mathcal{D}$ .

**Transformer.** Define an  $L$ -layer Transformer consisting of  $L$  Transformer layers as follows. The  $l$ th Transformer layer maps an input matrix  $\mathbf{H}^{(l)} \in \mathbb{R}^{D \times N}$  to  $\widetilde{\mathbf{H}}^{(l)} \in \mathbb{R}^{D \times N}$  and is composed of a self-attention block and a feed-forward block. The self-attention block  $\text{Attn}^{(l)} : \mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$  is parameterized by  $D \times D$  matrices  $\{(\mathbf{K}_m^{(l)}, \mathbf{Q}_m^{(l)}, \mathbf{V}_m^{(l)})\}_{m=1}^M$ , where  $M$  is the number of heads, and defined as  $\text{Attn}^{(l)}(\mathbf{X}) = \mathbf{X} + \frac{1}{N} \sum_{m=1}^M \mathbf{V}_m^{(l)} \mathbf{X} \sigma((\mathbf{Q}_m^{(l)} \mathbf{X})^\top \mathbf{K}_m^{(l)} \mathbf{X})$ .  $\sigma$  denotes an activation function applied elementwisely.

The feed-forward block  $\text{MLP}^{(l)} : \mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$  is a multi-layer perceptron with a skip connection, parameterized by  $(\mathbf{W}_1^{(l)}, \mathbf{W}_2^{(l)}) \in \mathbb{R}^{D' \times D} \times \mathbb{R}^{D \times D'}$ , such that  $\text{MLP}^{(l)}(\mathbf{X}) = \mathbf{X} + \mathbf{W}_2^{(l)} \zeta(\mathbf{W}_1^{(l)} \mathbf{X})$ , where  $\zeta$  is an activation function applied elementwisely. In this paper, we let

both  $\sigma$  and  $\zeta$  be the ReLU function in this paper, following (Bai et al., 2023).

An  $L$ -layer Transformer  $\text{TF}_\theta$ , parameterized by  $\theta = \{(\mathbf{K}_1^{(l)}, \mathbf{Q}_1^{(l)}, \mathbf{V}_1^{(l)}, \dots, \mathbf{K}_M^{(l)}, \mathbf{Q}_M^{(l)}, \mathbf{V}_M^{(l)}, \mathbf{W}_1^{(l)}, \mathbf{W}_2^{(l)})\}_{l=1}^L$ , is a composition of the abovementioned layers as  $\text{TF}_\theta(\mathbf{X}) = \text{MLP}^{(L)} \circ \text{Attn}^{(L)} \circ \dots \circ \text{MLP}^{(1)} \circ \text{Attn}^{(1)}(\mathbf{X})$ .

In the remaining text, the superscript to indicate the layer number  $^{(l)}$  is sometimes omitted for brevity, and the  $n$ th columns of  $\mathbf{H}^{(l)}$ ,  $\widehat{\mathbf{H}}^{(l)}$  are denoted as  $\mathbf{h}_n^{(l)}$ ,  $\widehat{\mathbf{h}}_n^{(l)}$ . We use  $\|\theta\|_{\text{TF}}$  as a norm of the parameter matrices; see Section A in the appendix.

### 3. Approximating UDA Algorithms by Transformers

This section demonstrates that transformers in in-context learning can approximate existing UDA algorithms. Specifically, we show that transformers can approximate uLSIF-based IWL, an instance-based method, and DANN, a feature-based method.

#### 3.1. Setup and Notion

We consider in-context domain adaptation, where a fixed Transformer model is given a tuple  $(\mathcal{D}_S, \mathcal{D}_T, \mathbf{x}_*)$ , where  $\mathbf{x}_* \sim p_T^X(\mathbf{x})$ , and predicts  $y_*$  without updating model parameters.

We construct the input matrix using both source and target data. Specifically, we suppose the input data, namely  $(\mathbf{x}_i^S, y_i^S) \in \mathcal{D}_S$  and  $\mathbf{x}_i^T \in \mathcal{D}_T$ , are encoded into  $\mathbf{H}^{(1)} \in \mathbb{R}^{D \times (N+1)}$  as follows:

$$\mathbf{H}^{(1)} = \begin{bmatrix} \mathbf{x}_1 & \dots & \mathbf{x}_n & \mathbf{x}_{n+1} & \dots & \mathbf{x}_N & \mathbf{x}_{N+1} \\ y_1 & \dots & y_n & 0 & \dots & 0 & 0 \\ t_1 & \dots & t_n & t_{n+1} & \dots & t_N & t_{N+1} \\ s_1 & \dots & s_n & s_{n+1} & \dots & s_N & s_{N+1} \\ 1 & \dots & 1 & 1 & \dots & 1 & 1 \\ \mathbf{0}_{D-(d+4)} & \dots & \mathbf{0}_{D-(d+4)} & \mathbf{0}_{D-(d+4)} & \dots & \mathbf{0}_{D-(d+4)} & \mathbf{0}_{D-(d+4)} \end{bmatrix}, \quad (6)$$

where  $\mathbf{x}_i = \mathbf{x}_i^S$  and  $y_i = y_i^S$  for  $1 \leq i \leq n$ ,  $\mathbf{x}_i = \mathbf{x}_{i-n}^T$  for  $n+1 \leq i \leq N$ ,  $\mathbf{x}_{N+1} = \mathbf{x}_*$ .  $t_i = 1$  for  $1 \leq i \leq n$  otherwise 0 is used to indicate which data point is from the source domain, and  $s_i = 1$  for  $1 \leq i \leq N$  otherwise 0 is used to mark training data. We further define an output of a transformer for the prediction corresponding to  $\mathbf{x}_{N+1} = \mathbf{x}_*$ . For an input  $\mathbf{H}^{(1)}$  and its corresponding output matrix  $\text{TF}_\theta(\mathbf{H}^{(1)})$ , we write its  $(2, N+1)$ -th element of  $\text{TF}_\theta(\mathbf{H}^{(1)})$  as  $\text{TF}_\theta^*(\mathbf{H}^{(1)})$ .

Let  $B_x, B_y > 0$ . To use  $(\epsilon, R, M, C)$ -approximability in Appendix A, we suppose  $\|\mathbf{x}_i\| \leq B_x$  and  $\|y_i\| \leq B_y$  for any  $i \in \{1, \dots, N+1\}$ , and  $\|\mathbf{w}\| \leq B_w$  for a model weight  $\mathbf{w}$  in in-context learning. Additionally, for the feature map of IWL, we suppose  $\|\phi\| \leq 1$  so that  $\|\phi(\mathbf{x})\| \leq B_x$ .

#### 3.2. In-context IWL with uLSIF

We show a transformer in the ICL scheme can approximate the uLSIF estimator-based IWL algorithm, which we refer to as IWL in the following.

In this section, we consider the following setup: Fix any  $B_w > 0, B_\alpha > 0, L_1, L_2 > 0, \eta_1$ , and  $\eta_2$ . Given a loss function  $\ell$  that is convex in the first argument, and  $\nabla_1 \ell$  is  $(\epsilon, R, M, C)$ -approximable by the sum of ReLUs with  $R = \max(B_w B_x, B_y, 1)$ . We first state the main result of this section. This theorem shows that the transformer has the performance to function as an algorithm almost equivalent to IWL for any input  $\mathbf{x}_*$ . Note that the query  $\mathbf{x}_*$  is encoded in the input matrix  $\mathbf{H}^{(1)}$  defined in (6).

**Theorem 1.** *Consider  $\widehat{f}^{\text{IWL}}$  with  $\phi$  which is approximable by a sum of ReLUs. Fix  $\epsilon > 0$  arbitrarily and set  $L_2$  as satisfying  $0 < \epsilon \leq B_w/2L_2$ . Suppose that an input  $(\mathcal{D}_S, \mathcal{D}_T, \mathbf{x}_*)$  satisfies that  $\sup_{\mathbf{w}: \|\mathbf{w}\|_2 \leq B_w} \lambda_{\max}(\nabla^2 \widehat{R}_T(\mathbf{w}; \mathcal{D}_S)) \leq \eta_2/2$  and the minimizer  $\mathbf{w}^* \in \text{argmin} \widehat{R}_T(\mathbf{w}; \mathcal{D}_S)$  satisfies  $\|\mathbf{w}^*\| \leq B_w/2$ . Then, there exists a transformer  $\text{TF}_\theta$  with  $L_1 + L_2 + 1$  layers and  $M$  heads which satisfies the following:  $\|\text{TF}_\theta^*(\mathbf{H}^{(1)}) - \widehat{f}^{\text{IWL}}(\mathbf{x}_*)\| \leq \epsilon$ .*

This result has several implications. First, while transformers with ICL have been shown to approximate algorithms for i.i.d. data (Bai et al., 2023), this theorem extends it to show the capability of handling domain shifts. This extension is nontrivial since it requires showing that the Transformer can express the algorithm's ability to correct the domain shifts. Second, this theorem shows that Transformers do not need to specify the feature map  $\phi(\cdot)$ . Specifically, whatever feature map  $\phi(\cdot)$  is used for uLSIF, the transformer can approximate the uLSIF based on it. In other words, pre-training a transformer induces a situation-specific feature map  $\phi(\cdot)$ .

#### 3.3. In-context DANN

We show the existence of a Transformer that can approximate DANN, presented in Section 2.1.

At the beginning, we identify the structure of the DANN to be approximated. First, we assume that a label classifier  $f_L \circ f_F(\mathbf{x})$  is represented by a two-layer neural network model  $\Lambda(\mathbf{x}; \mathbf{U}, \mathbf{w}) := \sum_{k=1}^K w_k r(\mathbf{u}_k^\top \mathbf{x})$ , and the domain classifier  $f_D \circ f_F(\mathbf{x})$  also follows a similar model  $\Delta(\mathbf{x}; \mathbf{U}, \mathbf{v}) := \sum_{k=1}^K v_k r(\mathbf{u}_k^\top \mathbf{x})$ , with parameters  $\mathbf{w}, \mathbf{v} \in \mathbb{R}^K$  and  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_K] \in \mathbb{R}^{K \times d}$  and an activation function  $r(\cdot)$ . Let  $\mathbf{u} = \text{vec}(\mathbf{U})$  and  $\mathbf{w}, \mathbf{v}, \mathbf{u}$  correspond to  $\theta_L, \theta_D, \theta_F$ , respectively. This simplification using a two-layer neural network was included to simplify the discussion and is easy to generalize.

We further define a sequence of parameters

$\{(\mathbf{u}^{(l)}, \mathbf{v}^{(l)}, \mathbf{w}^{(l)})\}_{l=1}^L$  that can be obtained by  $L$  updates by the DANN. In preparation, we define a closed set  $\mathcal{W} = \{(\mathbf{w}, \mathbf{v}, \mathbf{u}) : \|\mathbf{w}\| \leq B_w, \|\mathbf{v}\| \leq B_v, \max_k \|\mathbf{u}_k\| \leq B_u\} \subset \mathbb{R}^{K \times K \times Kd}$  with some  $B_u, B_w, B_v > 0$  and also define a projection  $\Pi_{\mathcal{W}}(\cdot)$  onto  $\mathcal{W}$ . Let  $(\mathbf{w}^{(1)}, \mathbf{v}^{(1)}, \mathbf{u}^{(1)}) \in \mathcal{W}$  be a tuple of initial values. Then, the subsequent parameters are recursively defined by the following updates:

$$\mathbf{u}^{(l)} = \Pi_{\mathcal{W}} \left( \mathbf{u}^{(l-1)} - \eta \nabla_{\mathbf{u}} \{L(\mathbf{u}^{(l-1)}, \mathbf{w}^{(l-1)}) - \lambda \Omega(\mathbf{u}^{(l-1)}, \mathbf{v}^{(l-1)})\} \right), \quad (7)$$

$$\mathbf{w}^{(l)} = \Pi_{\mathcal{W}} \left( \mathbf{w}^{(l-1)} - \eta \nabla_{\mathbf{w}} L(\mathbf{u}^{(l-1)}, \mathbf{w}^{(l-1)}) \right), \quad (8)$$

$$\mathbf{v}^{(l)} = \Pi_{\mathcal{W}} \left( \mathbf{v}^{(l-1)} - \eta \lambda \nabla_{\mathbf{v}} \Omega(\mathbf{u}^{(l-1)}, \mathbf{v}^{(l-1)}) \right). \quad (9)$$

These updates are an analogy of the original DANN updates Equations (3) to (5) to optimize the minimax loss function Equation (2) of DANN depending on  $\mathcal{D}_S$  and  $\mathcal{D}_T$ . Here, we introduce the projection to explain a more realistic setup, for example, (Shen et al., 2018). Finally, the classifier obtained by  $L$  updates of DANN is defined as  $\hat{f}^{\text{DANN}}(\mathbf{x}) := \Lambda(\mathbf{x}; \mathbf{u}^{(L)}, \mathbf{w}^{(L)})$ .

We show the existence of a Transformer that successively approximates the parameter sequence by DANN defined above. The results are given in the following statement.

**Theorem 2.** Fix any  $B_u, B_w, B_v > 0$ ,  $L > 0$ ,  $\eta > 0$ , and  $\epsilon > 0$ . Suppose activation function  $r$  is  $C^4$ -smooth. Then, there exists a  $2L$ -layer Transformer  $\text{TF}_{\theta}$  with  $\max_{l \in \{1, \dots, 2L\}} M^{(l)} \leq \tilde{O}(\epsilon^{-2})$ ,  $\max_{l \in \{1, \dots, 2L\}} D^{(l)} \leq \tilde{O}(\epsilon^{-2}) + D_{\text{MLP}}$ ,  $\|\theta\|_{\text{TF}} \leq \mathcal{O}(1 + \eta) + C_{\text{MLP}}$ , with some existing constants  $D_{\text{MLP}}, C_{\text{MLP}} > 0$ , which satisfies the following: for any input  $(\mathcal{D}_S, \mathcal{D}_T, \mathbf{x}_*)$ , an  $2l$ -th layer of the transformer maps  $\tilde{\mathbf{h}}_i^{(2l-1)} := [z_i, \tilde{\mathbf{u}}^{(l-1)}, \tilde{\mathbf{w}}^{(l-1)}, \tilde{\mathbf{v}}^{(l-1)}]$  with any  $(\tilde{\mathbf{w}}^{(l-1)}, \tilde{\mathbf{v}}^{(l-1)}, \tilde{\mathbf{u}}^{(l-1)}) \in \mathcal{W}$  to  $\tilde{\mathbf{h}}_i^{(2l)} := [z_i, \tilde{\mathbf{u}}^{(l)}, \tilde{\mathbf{w}}^{(l)}, \tilde{\mathbf{v}}^{(l)}]$  for each  $i \in \{1, \dots, N + 1\}$  to satisfy Equations (7) to (9) where  $\epsilon_u^{(l-1)}, \epsilon_w^{(l-1)}, \epsilon_v^{(l-1)} \in \mathbb{R}^k$  are some vectors satisfying  $\max\{\|\epsilon_u^{(l-1)}\|, \|\epsilon_w^{(l-1)}\|, \|\epsilon_v^{(l-1)}\|\} \leq \epsilon$ .

The results show that the existing Transformer approximates the original DANN updates Equations (7) to (9), at each step. The terms  $\epsilon_u^{(l-1)}, \epsilon_w^{(l-1)}, \epsilon_v^{(l-1)}$  expresses approximation errors, whose norm are no more than the fixed  $\epsilon$ . Consequently, the final output of the transformer approximates the output of the original DANN after  $L$ -iterations. The results are given as follows:

**Corollary 1.** Consider the setup as in Theorem 2. Then, for any input  $(\mathcal{D}_S, \mathcal{D}_T, \mathbf{x}_*)$ , the transformer  $\text{TF}_{\theta}$  in Theorem 2 outputs a corresponding tuple  $(\tilde{\mathbf{w}}^{(L)}, \tilde{\mathbf{v}}^{(L)}, \tilde{\mathbf{u}}^{(L)}) \in \mathcal{W}$  which satisfies  $\|\text{TF}_{\theta}^*(\mathbf{H}^1) - \hat{f}^{\text{DANN}}(\mathbf{x}_*)\| \leq L\epsilon$ .

## 4. Automatic Algorithm Selection by In-Context UDA

We demonstrate that Transformers can automatically select UDA algorithms based on data in context. Specifically, we consider the selection of methods determined by whether the supports of the source distribution  $p_S$  and the target distribution  $p_T$  sufficiently overlap. In cases where the support of the target distribution  $p_T$  overlaps with that of the source distribution  $p_S$ , the IWL algorithm, which operates using the density ratio  $q(\mathbf{x}) = p_T(\mathbf{x})/p_S(\mathbf{x})$ , is employed. In other words, IWL should be employed when we have

$$p_S(\mathbf{x}) > 0 \text{ holds for all } \mathbf{x} \text{ such that } p_T(\mathbf{x}) > 0. \quad (10)$$

Conversely, when there is no such overlap, i.e.,  $p_S(\mathbf{x}) = 0$  for some  $\mathbf{x}$  with  $p_T(\mathbf{x}) > 0$ , one should select DANN, which does not rely on the density ratio.

We show that Transformers are capable of realizing the above design automatically.

$$\hat{f}^{\text{ICUDA}}(\mathbf{x}) := \begin{cases} \hat{f}^{\text{IWL}}(\mathbf{x}_*) & \text{if } \min_{\mathbf{x}: p_T(\mathbf{x}) > 0} p_S(\mathbf{x}) > 0, \\ \hat{f}^{\text{DANN}}(\mathbf{x}_*) & \text{otherwise,} \end{cases} \quad (11)$$

The statement is as follows:

**Theorem 3.** Fix any  $\epsilon > 0$ . Suppose that  $n' \geq (1/\epsilon)^3 \log n'$  holds and  $p_T(\cdot)$  is Lipschitz continuous. Then, there exists a Transformer  $\text{TF}_{\theta}$  with three layers and  $M$  heads satisfies the following with probability at least  $1 - 1/n' - O(\epsilon)$ : for any input  $(\mathcal{D}_S, \mathcal{D}_T, \mathbf{x}_*)$ , a transformer  $\text{TF}_{\theta}$  satisfies

$$\|\text{TF}_{\theta}^*(\mathbf{H}^1) - \hat{f}^{\text{ICUDA}}(\mathbf{x}_*)\| \leq \epsilon. \quad (12)$$

The results show that the Transformer can automatically check the condition (10) and use the appropriate algorithm without the user having to make a choice. Note that the density ratio condition (10) is one of the options that are automatically learned, and in practice, Transformers can learn more complicated conditions. Importantly, in this case, cross-validation cannot be used to make a selection since true labels of the target domain data cannot be observed.

## 5. Conclusion and Discussion

This paper proved that Transformers can approximate instance-based and feature-based domain adaptation algorithms and select an appropriate one for the given dataset in the in-context learning framework. Technically, our results revealed that Transformers can approximately solve some types of linear equations and minimax problems. Numerical experiments in Figure 1 demonstrated that domain adaptation in in-context learning outperforms UDA methods tailored for specific assumptions.

---

## References

- Ahn, K., Cheng, X., Daneshmand, H., and Sra, S. Transformers learn to implement preconditioned gradient descent for in-context learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Akyürek, E., Schuurmans, D., Andreas, J., Ma, T., and Zhou, D. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*, 2022.
- Bai, Y., Chen, F., Wang, H., Xiong, C., and Mei, S. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=liMSqUuVg9>.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine learning*, 79:151–175, 2010.
- Dai, W., Yang, Q., Xue, G.-R., and Yu, Y. Boosting for transfer learning. In *Proceedings of the 24th international conference on Machine learning*, pp. 193–200, 2007.
- Daumé III, H. Frustratingly easy domain adaptation. *arXiv preprint arXiv:0907.1815*, 2009.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., and Lempitsky, V. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59):1–35, 2016.
- Garg, S., Tsipras, D., Liang, P. S., and Valiant, G. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598, 2022.
- Kanamori, T., Hido, S., and Sugiyama, M. A least-squares approach to direct importance estimation. *The Journal of Machine Learning Research*, 10:1391–1445, 2009.
- Kimura, M. and Hino, H. A short survey on importance weighting for machine learning. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=IhXM3g2gxxg>. Survey Certification.
- Li, Y., Ildiz, M. E., Papailiopoulos, D., and Oymak, S. Transformers as algorithms: Generalization and implicit model selection in in-context learning. *arXiv preprint arXiv:2301.07067*, 2023.
- Lin, L., Bai, Y., and Mei, S. Transformers as decision makers: Provable in-context reinforcement learning via supervised pretraining. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- Raventós, A., Paul, M., Chen, F., and Ganguli, S. Pretraining task diversity and the emergence of non-bayesian in-context learning for regression. *Advances in Neural Information Processing Systems*, 36, 2024.
- Shen, J., Qu, Y., Zhang, W., and Yu, Y. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, 2018.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.
- Sugiyama, M., Krauledat, M., and Müller, K.-R. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(5), 2007.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., and Vladymyrov, M. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.
- Xie, S. M., Raghunathan, A., Liang, P., and Ma, T. An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*, 2021.
- Zhang, R., Frei, S., and Bartlett, P. Trained transformers learn linear models in-context. In *R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*, 2023.

## A. Additional Setup

We define a norm of a parameter  $\|\theta\|_{\text{TF}}$  of Transformers as

$$\|\theta\|_{\text{TF}} = \max_{l \in \{1, \dots, L\}} \left\{ \max_{m \in \{1, \dots, M\}} \left\{ \|\mathbf{Q}_m^{(l)}\|, \|\mathbf{K}_m^{(l)}\| \right\} + \sum_{m=1}^M \left( \|\mathbf{V}_m^{(l)}\| + \|\mathbf{W}_1^{(l)}\| + \|\mathbf{W}_2^{(l)}\| \right) \right\}, \quad (13)$$

where  $\|\cdot\|$  for matrices indicates the operator norm in this equation.

We also define the concepts necessary for our theoretical results. To approximate smooth functions, such as loss functions  $\ell$  and  $\gamma$ , we use the following notion.

**Definition 1** ( $(\varepsilon, R, M, C)$ -approximability by sum of ReLUs, (Bai et al., 2023)). For  $\varepsilon > 0$  and  $R \geq 1$ , a function  $g : \mathbb{R}^k \rightarrow \mathbb{R}$  is  $(\varepsilon, R, M, C)$ -approximable by sum of ReLUs if there exist a function  $f(\mathbf{z}) = \sum_{m=1}^M c_m \sigma(\mathbf{a}_m^\top \mathbf{z} + b_m)$  with  $\sum_{m=1}^M |c_m| \leq C$  with  $\max_{m \in \{1, \dots, M\}} \|\mathbf{a}_m\|_1 + b_m \leq 1$ , where  $\mathbf{a}_m \in \mathbb{R}^k$ ,  $b_m \in \mathbb{R}$ ,  $c_m \in \mathbb{R}$ , and  $\sup_{\mathbf{z} \in [-R, R]^k} |g(\mathbf{z}) - f(\mathbf{z})| \leq \varepsilon$ .

All functions approximated in this paper are included in this class.

## B. Proof Outline of In-context IWL

In our proof, we construct three types of sub-transformers and approximate IWL-uLSIF by combining them. Specifically, the first sub-transformer with one layer to construct a feature map  $\phi(\cdot)$ , the second one calculates  $\alpha$  with  $L_1$  layers, and this one optimizes the model weight  $\mathbf{w} \in \mathbb{R}^J$  with  $L_2$  layers. For simplicity, we define a notation  $\mathbf{z}_i = [\mathbf{x}_i, y_i, t_i, s_i, 1]$ . In the following, we describe each approximation step by step.

**Step 1: Feature map approximation.** First, we construct a transformer that approximates the feature map  $\phi(\cdot)$ . This construction is trivial under the assumption that the feature map  $\phi(\cdot)$  is approximable by the sum of ReLUs. For example, a self-attention block can construct a feature map with the RBF kernel.

**Lemma 1.** *There exists a transformer with one layer which maps  $\mathbf{h}_i^{(1)} = [\mathbf{z}_i, \mathbf{0}_J, \mathbf{0}_{D-(d+4)}]$  to  $\tilde{\mathbf{h}}_i^{(1)} = [\mathbf{z}_i, \hat{\phi}(\mathbf{x}_i), \mathbf{0}_{D-J-(d+4)}]$  for  $i = 1, \dots, N + 1$ , such that  $\|\phi - \hat{\phi}\|_{L^\infty} \leq \varepsilon$  holds.*

**Step 2: Density ratio parameter approximation.** Second, we construct a transformer that approximates  $\hat{\alpha}$  that minimizes the loss  $\hat{L}(\alpha)$  in Equation (1), in which a technical difficulty lies. For the approximation, we define a sequence of parameters  $\{\alpha_{\text{GD}}^{(l)}\}_{l=1,2,\dots}$  that converge to  $\hat{\alpha}$ , using an update equation by the gradient descent algorithm:

$$\alpha_{\text{GD}}^{(l+1)} = \alpha_{\text{GD}}^{(l)} - \eta_1 \nabla \hat{L}(\alpha_{\text{GD}}^{(l)}), \quad l = 1, 2, \dots, \quad \text{and} \quad \alpha_{\text{GD}}^{(1)} = \mathbf{0}_J. \quad (14)$$

Then, we develop a Transformer layer that exactly implements a single update step of this equation. We give the following statement.

**Lemma 2.** *There exists a Transformer with  $L_1$  layer which maps  $\tilde{\mathbf{h}}_i^{(1)} = [\mathbf{z}_i, \phi(\mathbf{x}_i), \mathbf{0}_J, \mathbf{0}_{D-2J-(d+4)}]$  to  $\tilde{\mathbf{h}}_i^{(L_1+1)} = [\mathbf{z}_i, \phi(\mathbf{x}_i), \tilde{\alpha}^{(L_1)}, \mathbf{0}_{D-2J-(d+4)}]$  for  $i = 1, \dots, N + 1$ , such that  $\alpha_{\text{GD}}^{(L_1)} = \tilde{\alpha}^{(L_1)}$ .*

We note that this approach with the gradient descent is more efficient in the sense of layer size than the direct minimization of  $\hat{L}(\alpha)$  using an inverse matrix.

**Step 3: Model parameter approximation.** In this step, we employ a similar approach to develop a Transformer that approximates  $\mathbf{w}^*$ , which is the minimizer of  $\hat{R}_T(\mathbf{w})$  in Section 2.1. To the aim, we define a sequence of parameters  $\{\mathbf{w}_{\text{GD}}^{(l)}\}_{l=1,2,\dots}$  that converge to  $\mathbf{w}^*$  by the gradient descent algorithm:

$$\mathbf{w}_{\text{GD}}^{(l+1)} = \mathbf{w}_{\text{GD}}^{(l)} - \eta_2 \nabla \hat{R}_T(\mathbf{w}_{\text{GD}}^{(l)}), \quad l = 1, 2, \dots, \quad \text{and} \quad \mathbf{w}_{\text{GD}}^{(1)} = \mathbf{0}_J. \quad (15)$$

Then, we develop a Transformer layer that approximates this algorithm as follows:

**Lemma 3.** *There exists a Transformer with  $L_2$  layer which maps  $\tilde{\mathbf{h}}_i^{(L_1+1)} = [\mathbf{z}_i, \phi(\mathbf{x}_i), \tilde{\alpha}^{(L_1)}, \mathbf{0}_{D-2J-(d+4)}]$  to  $\tilde{\mathbf{h}}_i^{(L_1+L_2+1)} = [\mathbf{z}_i, \phi(\mathbf{x}_i), \tilde{\alpha}^{(L_1)}, \tilde{\mathbf{w}}^{(L_2)}, 0]$  for  $i = 1, \dots, N + 1$ , such that  $\|\mathbf{w}_{\text{GD}}^{(L_2)} - \tilde{\mathbf{w}}^{(L_2)}\|_{L^\infty} \leq \varepsilon L_2 \eta_2 B_x$  holds.*

Combining these layers constitutes the transformer that achieves the approximation capability of Theorem 1. Importantly, this method requires as many layers as there are iterations, but this can be shortened by using regression coupling.

---

## C. Proof Outline of In-context DANN

Here, we focus on the sub-problem that a pair of Transformer layers approximates a single step of update of DANN consisting of two-layer neural networks. Specifically, the attention block at the first layer approximates the forward passes of networks, the succeeding MLP block computes the partial derivatives of the loss function, then the next attention block implements the minimax optimization step, and the last MLP block projects the weights.

The specific procedure is presented in the following steps. For brevity, we write  $\gamma := (\mathbf{u}, \mathbf{w}, \mathbf{v})$ .

**Step 1: Forward pass approximation.** The first attention block performs the forward passes, that is,  $\mathbf{x}_i \mapsto (\Lambda(\mathbf{x}_i), \Delta(\mathbf{x}_i))$ , by approximating the activation function  $r$  by sum-of-ReLUs.

**Lemma 4.** *For each  $\varepsilon > 0$ , there exists an attention block which maps  $\mathbf{h}_i = [\mathbf{z}_i, \gamma, \mathbf{0}_4]$  to  $\mathbf{h}'_i = [\mathbf{z}_i, \gamma, \tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), \mathbf{0}_2]$  for  $i = 1, \dots, N$  such that  $\|\tilde{\Lambda} - \Lambda\|_{L^\infty} \leq \varepsilon$  and  $\|\tilde{\Delta} - \Delta\|_{L^\infty} \leq \varepsilon$  hold.*

**Step 2: Loss derivative approximation.** The next MLP block obtains loss derivatives, i.e.,  $(\tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), y_i, t_i, s_i) \mapsto (\mathbb{1}(i \leq n)\partial_1\gamma(\tilde{\Lambda}(\mathbf{x}_i), y_i), \mathbb{1}(i \leq N)\partial_1\gamma(\tilde{\Delta}(\mathbf{x}_i), y_i))$ . We also use sum-of-ReLUs to approximate  $\partial_1\gamma$ .

**Lemma 5.** *For any  $\varepsilon > 0$ , there exists an MLP block which maps  $\mathbf{h}'_i = [\mathbf{z}_i, \gamma, \tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), \mathbf{0}_2]$  to  $\mathbf{h}''_i = [\mathbf{z}_i, \gamma, \tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), g_{\Lambda,i}, g_{\Delta,i}]$ , where  $g_{\Lambda,i}$  and  $g_{\Delta,i}$  are scalars satisfying  $|g_{\Lambda,i} - \mathbb{1}(i \leq n)\partial_1\gamma(\tilde{\Lambda}(\mathbf{x}_i), y_i)| \leq \varepsilon$  and  $|g_{\Delta,i} - \mathbb{1}(i \leq N)\partial_1\gamma(\tilde{\Delta}(\mathbf{x}_i), y_i)| \leq \varepsilon$ .*

**Step 3: Gradient descent approximation** The attention block at the second layer approximates the optimization step:  $(\mathbf{u}, \mathbf{w}, \mathbf{v}) \mapsto (\mathbf{u} - \eta\nabla_{\mathbf{u}}(L(\mathbf{u}, \mathbf{w}) - \lambda\Omega(\mathbf{u}, \mathbf{v})), \mathbf{w} - \eta\nabla_{\mathbf{w}}L(\mathbf{u}, \mathbf{w}), \mathbf{v} - \eta\lambda\nabla_{\mathbf{v}}\Omega(\mathbf{u}, \mathbf{v}))$ , which can be obtained by approximating  $(s, t) \mapsto s \cdot r'(t)$ .

**Lemma 6.** *For each  $\varepsilon > 0$ , there exists an attention block that maps  $\mathbf{h}''_i = [\mathbf{z}_i, \gamma, \tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), g_{\Lambda,i}, g_{\Delta,i}]$  to  $\mathbf{h}'''_i = [\mathbf{z}_i, \mathbf{u} - \eta\mathbf{g}_u, \mathbf{w} - \eta\mathbf{g}_w, \mathbf{v} - \eta\mathbf{g}_v, \tilde{\Lambda}(\mathbf{x}_i), \tilde{\Delta}(\mathbf{x}_i), g_{\Lambda,i}, g_{\Delta,i}]$ , where  $\|\mathbf{g}_u - \nabla_{\mathbf{u}}(L(\mathbf{u}, \mathbf{w}) - \lambda\Omega(\mathbf{u}, \mathbf{v}))\|_2, \|\mathbf{g}_w - \nabla_{\mathbf{w}}L(\mathbf{u}, \mathbf{w})\|_2, \|\mathbf{g}_v - \lambda\nabla_{\mathbf{v}}\Omega(\mathbf{u}, \mathbf{v})\|_2 \leq \varepsilon$ .*

**Step 4: Projection approximation.** The last MLP block projects  $\mathbf{u}, \mathbf{w}, \mathbf{v}$  onto  $\mathcal{W}$  appropriately by the assumption.

By stacking such pairs of layers, a  $2L$ -layer Transformer can implement  $L$  steps of DANN in context.

## D. Proof Outline of In-context Algorithm Selection

The key idea is to first implement kernel density estimation of  $p_S(\mathbf{x})$  for  $\mathbf{x} \in \mathcal{D}_T$  and then select an algorithm with  $\min_{\mathbf{x} \in \mathcal{D}_T} \hat{p}_S(\mathbf{x}) > \delta$ , where  $\delta > 0$  is a given threshold. Here, we assume that each token is  $\mathbf{h}_i = [\mathbf{z}_i, \gamma, \tilde{f}^{\text{IWL}}(\mathbf{x}_i), \tilde{f}^{\text{DANN}}(\mathbf{x}_i), \mathbf{0}]$ , where  $\gamma$  consists of weights of in-context IWL and DANN and is unused in this selection process.

**Step 1: Approximation of density estimation.** The first step is to approximate the density function  $p_S$  with kernel density estimation  $\hat{p}_S(\cdot) = \frac{1}{n} \sum_{i=1}^n K(\cdot, \mathbf{x}_i)$ , where  $K$  is a kernel, such as the RBF kernel, and evaluate it on each data point.

**Lemma 7.** *For any  $\epsilon > 0$ , there exists an attention block that maps  $\mathbf{h}_i$  to  $\mathbf{h}'_i = [\mathbf{z}_i, \gamma, \tilde{f}^{\text{IWL}}(\mathbf{x}_i), \tilde{f}^{\text{DANN}}(\mathbf{x}_i), p_i, \mathbf{0}]$ , where  $|p_i - \hat{p}_S(\mathbf{x}_i)| \leq \epsilon$ .*

**Step 2: Algorithm selection approximation.** Then, the Transformer decides which algorithm is appropriate based on  $\text{softmax}_{\beta, i} p_i > \delta$ , where  $\text{softmax}_{\beta, i} p_i = -\frac{1}{\beta} \log \sum (-\beta p_i)$  is a relaxation of min with an inverse temperature parameter  $\beta > 0$ .

**Lemma 8.** *For any  $\epsilon > 0$ , there is an MLP block succeeded by a Transformer layer that maps  $\mathbf{h}'_{N+1}$  to  $\mathbf{h}''_{N+1} = [\mathbf{z}_i, \gamma, \tilde{f}^{\text{IWL}}(\mathbf{x}_*), \tilde{f}^{\text{DANN}}(\mathbf{x}_*), p_i, -\frac{1}{\beta} \log \sum_i \exp(-\beta p_i), \tilde{f}(\mathbf{x}_*)]$ , where  $|\tilde{f} - \tilde{f}^{\text{ICUDA}}| \leq \epsilon$ .*

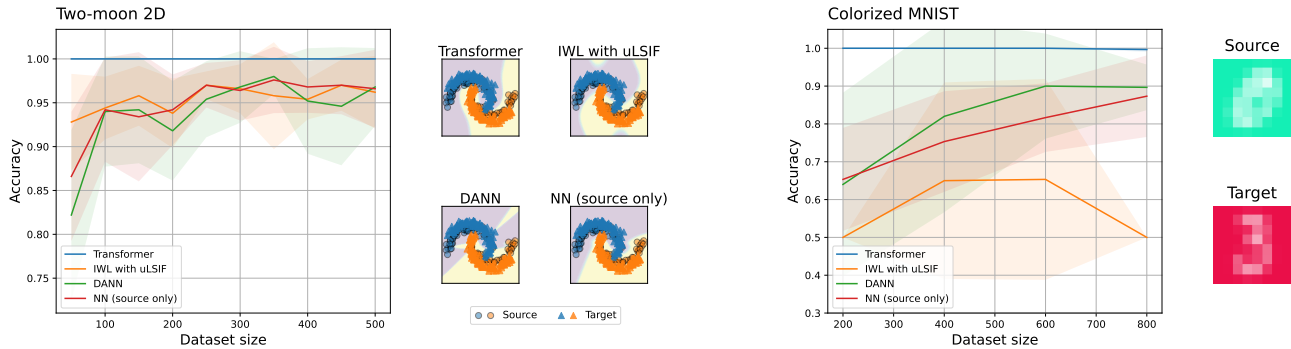


Figure 1. (Left) Test accuracy averaged over five runs of Transformer (ICL) and baseline models on the two-moon 2D problem. Decision boundaries of the models are presented when  $N = 200$ . (Right) Test accuracy averaged over five runs on the colored MNIST.

## E. Experiments

We verify our theory by using the in-context domain adaptation abilities of Transformers using two synthetic problems. The ICL domain adapter is compared with 1. IWL with uLSIF using the RBF kernel; 2. DANN of a two-layer neural network with the ELU activation function; and 3. the same neural network trained on the source domain only. For the in-context learning, we used an eight-layer Transformer with eight heads and pre-trained it to minimize  $\gamma(\text{TF}_\theta(\mathbf{x}'_*; \mathcal{D}'_S, \mathcal{D}'_T), y'_*)$ , where  $(\mathbf{x}'_*, y'_*) \sim \mathcal{D}'_T$ , for randomly synthesized datasets  $(\mathcal{D}'_S, \mathcal{D}'_T) (\neq (\mathcal{D}_T, \mathcal{D}_S))$  for  $10^4$  iterations. Each dataset consists of  $n = n' = N/2$  data points. Note that test data used to report test accuracy are unseen during pre-training.

Figure 1 (Left) presents test accuracy on the two-moon 2D problem, where the target distribution is a rotation of the source one. The decision boundaries are also presented, and we can observe that the Transformer learns a smoother boundary than others. Additionally, Figure 1 (Right) demonstrates the results of the colored MNIST problem, where each dataset consists of images of two digits resized to  $8 \times 8$  pixels, and their background colors alter based on the domains, as presented at the rightmost. On both problems, ICL consistently achieves much better performance than the baselines. These results indicate that Transformer implements adaptive domain adaptation algorithms to given datasets.