

# Trustworthy Transfer Learning: A Survey

JUN WU\*, Michigan State University, USA

JINGRUI HE, University of Illinois Urbana-Champaign, USA

Transfer learning aims to transfer knowledge or information from a source domain to a relevant target domain. In this paper, we understand transfer learning from the perspectives of knowledge transferability and trustworthiness. This involves two research questions: How is knowledge transferability quantitatively measured and enhanced across domains? Can we trust the transferred knowledge in the transfer learning process? To answer these questions, this paper provides a comprehensive review of trustworthy transfer learning from various aspects, including problem definitions, theoretical analysis, empirical algorithms, and real-world applications. Specifically, we summarize recent theories and algorithms for understanding knowledge transferability under (within-domain) IID and non-IID assumptions. In addition to knowledge transferability, we review the impact of trustworthiness on transfer learning, e.g., whether the transferred knowledge is adversarially robust or algorithmically fair, how to transfer the knowledge under privacy-preserving constraints, etc. Beyond discussing the current advancements, we highlight the open questions and future directions for understanding transfer learning in a reliable and trustworthy manner.

**JAIR Associate Editor:** Bo Han

## JAIR Reference Format:

Jun Wu and Jingrui He. 2025. Trustworthy Transfer Learning: A Survey. *Journal of Artificial Intelligence Research* 84, Article 20 (November 2025), 59 pages. DOI: [10.1613/jair.1.17602](https://doi.org/10.1613/jair.1.17602)

## 1 Introduction

Standard machine learning assumes that training and testing samples are independently and identically drawn (IID). With this IID assumption, modern machine learning models (e.g., deep neural networks (LeCun et al. 2015)) have achieved promising performance in a variety of high-impact applications. However, this IID assumption is often violated in real-world scenarios, especially when samples are collected from different sources and environments (Pan and Q. Yang 2010; J. Wu, J. He, and Tong 2024). Transfer learning has been introduced to tackle the distribution shifts between training (source domain) and testing (target domain) data sets. In contrast to standard machine learning involving samples from a single domain, transfer learning focuses on modeling heterogeneous data collected from different domains. The intuition behind transfer learning is to bridge the gap between source and target data by discovering and transferring their shared knowledge (Pan and Q. Yang 2010). Compared to learning from the target domain alone, the transferred knowledge could significantly improve the prediction performance on the target domain, especially when the target domain has limited or no labeled data (Ben-David, Blitzer, et al. 2010; Tripuraneni, Jordan, et al. 2020). In recent decades, by instantiating the learning models with modern neural networks, a deep transfer learning paradigm has been introduced with enhanced transferability capabilities (Yosinski et al. 2014).

---

\*This work was mainly completed when Jun Wu was a PhD student at University of Illinois Urbana-Champaign.

---

Authors' Contact Information: Jun Wu, ORCID: [0000-0002-1512-524X](https://orcid.org/0000-0002-1512-524X), [wujun4@msu.edu](mailto:wujun4@msu.edu), Michigan State University, East Lansing, MI, USA; Jingrui He, ORCID: [0000-0002-6429-6272](https://orcid.org/0000-0002-6429-6272), [jingrui@illinois.edu](mailto:jingrui@illinois.edu), University of Illinois Urbana-Champaign, Urbana, IL, USA.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.17602](https://doi.org/10.1613/jair.1.17602)

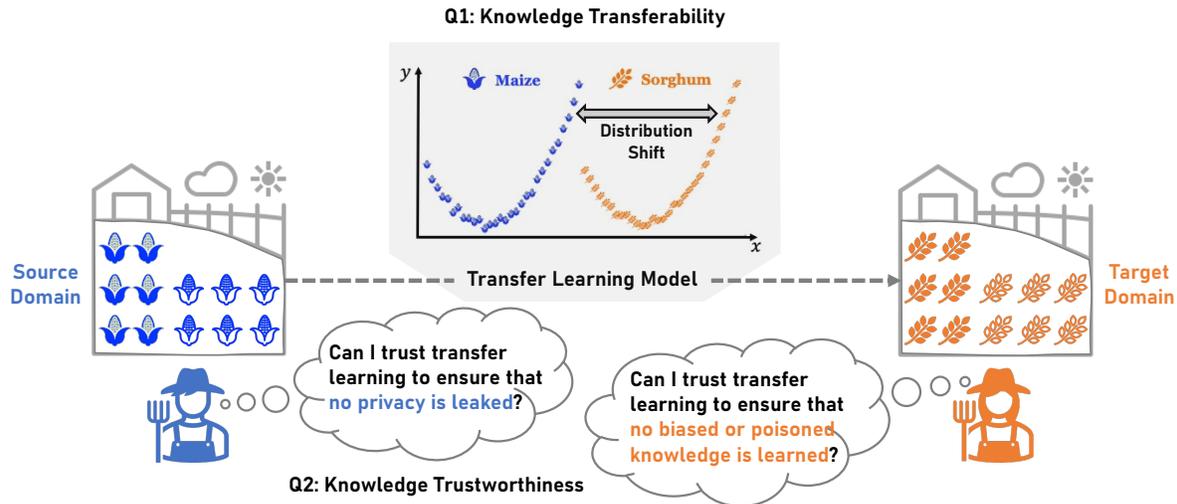


Fig. 1. A motivating example of trustworthy transfer learning

As illustrated in (Pan and Q. Yang 2010), transfer learning is a general term to describe the transfer of knowledge or information from source to target domains. Depending on the data and model assumptions, it can lead to various specific problem settings, such as data-level knowledge transfer (domain adaptation (Ben-David, Blitzer, et al. 2010; Ganin et al. 2016; Mansour, Mohri, and Rostamizadeh 2009a), out-of-distribution generalization (Blanchard et al. 2011; Muandet et al. 2013), and self-taught learning (Raina et al. 2007)) with available source samples and model-level knowledge transfer (fine-tuning (Shachaf et al. 2021), source-free adaptation (Aghbalou and Staerman 2023; J. Liang et al. 2020), knowledge distillation (Hinton et al. 2015)) with a pre-trained source hypothesis. The generalization performance of transfer learning techniques under various data and model assumptions has been studied over the past decades (Minami et al. 2023; Mohri, Sivek, et al. 2019; Tripuraneni, Jordan, et al. 2020; H. Zhao, Combes, et al. 2019). In addition to generalization performance, it is crucial to understand the trustworthiness (Eshete 2021) of the transferred knowledge in the transfer learning process, especially in safety-critical applications such as autonomous driving and medical diagnosis. It is explained (Varshney 2022) that “trust is the relationship between a trustor and a trustee: the trustor trusts the trustee”. In the context of transfer learning, the trustor can be the owners/users/regulators of either the source or the target domain. The trustee can be the transfer learning model itself, or the knowledge transferred from the source domain to the target domain. As summarized in earlier studies (Eshete 2021; Kaur et al. 2023; Varshney 2022), various trustworthiness properties can encourage the “trustor” to trust the “trustee” in real scenarios, including adversarial robustness, privacy, fairness, transparency, etc. Therefore, in this paper, we focus on trustworthy transfer learning (J. Wu and J. He 2023b) that aims to understand transfer learning from the perspective of both knowledge transferability and knowledge trustworthiness.

Fig. 1 provides a motivating example of trustworthy transfer learning in precision agriculture (Adve et al. 2024). In this example, a target farmer aims to train a model over the collected sorghum data. The task is to predict the biochemical traits (e.g., Nitrogen content, chlorophyll, etc.) of sorghum samples using the leaf hyperspectral reflectance (S. Wang et al. 2023; J. Wu, J. He, S. Wang, et al. 2022). Nevertheless, it is expensive and time-consuming to collect the labeled training samples. A feasible solution is to leverage knowledge from a relevant maize data set collected by a source farmer. This transfer learning process might involve several trustworthy concerns from the source and target farmers. To name a few, will the privacy of source data be leaked in transfer learning?

How does the poisoned and biased source knowledge negatively affect the prediction performance on the target domain? What is the fundamental trade-off between transfer performance and trustworthy properties? More generally, from the perspective of data and AI model markets (Pei et al. 2023), this emphasizes the importance of establishing trustworthiness between customers and sellers when purchasing AI models and sharing personal data.

This survey provides a comprehensive review of state-of-the-art theoretical analysis and algorithms for trustworthy transfer learning. More specifically, we summarize recent theories and algorithms for understanding knowledge transferability from two aspects: IID and non-IID transferability. IID transferability assumes that the samples within each domain are independent and identically distributed. In this scenario, we review three major quantitative metrics for evaluating the transferability across domains, including (data-level) *distribution discrepancy*, (task-level) *take diversity*, and (model-level) *transferability estimation*. In contrast, non-IID transferability considers a more relaxed assumption that the samples within each domain can be interdependent, e.g., connected nodes in graphs (Kipf and Welling 2017), word occurrence in texts (J. Y. Lee et al. 2018), temporal observations in time series (Purushotham et al. 2017), etc. We then review how transferability across domains can be quantitatively measured and enhanced in these complex scenarios. In addition to knowledge transferability, we also review the impact of trustworthiness on transfer learning techniques, including privacy, adversarial robustness, fairness, transparency, etc. Finally, we will highlight the open questions and future directions of trustworthy transfer learning.

The rest of this paper is organized as follows. Section 2 presents the main notation and the general problem definition of trustworthy transfer learning. Section 3 and Section 4 summarize the knowledge transferability and trustworthiness in various transfer learning scenarios, respectively. Section 5 provides the applications of transfer learning techniques in real-world applications, and Section 6 summarizes the open questions and future trends of trustworthy transfer learning. Finally, we conclude this survey in Section 7.

## 2 Preliminaries

In this section, we provide the main notation and general problem definition of trustworthy transfer learning.

### 2.1 Notation

In the paper, we let  $\mathcal{X}$  and  $\mathcal{Y}$  denote the input space and output space, respectively. Given a source domain  $\mathcal{D}_S$  and a target domain  $\mathcal{D}_T$ , we denote the probability density (or mass) functions of the source and target domains as  $p_S$  and  $p_T$  (or  $P_S$  and  $P_T$ ) over  $\mathcal{X} \times \mathcal{Y}$ , respectively. In the context of deep transfer learning, a hypothesis function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  can often be decomposed into two components: a feature extraction function  $g : \mathcal{X} \rightarrow \mathbb{R}^d$  and a prediction function  $h : \mathbb{R}^d \rightarrow \mathcal{Y}$ . We let  $\mathcal{F}$  be the class of hypothesis functions (with  $f \in \mathcal{F}$ ). Similarly, we can define  $\mathcal{G}$  (with  $g \in \mathcal{G}$ ) and  $\mathcal{H}$  (with  $h \in \mathcal{H}$ ) for the classes of the feature extraction and prediction functions, respectively. Notice that when the feature extractor is not considered (e.g., in Subsection 3.1.1), we can simply use  $\mathcal{H}$  to represent the class of hypothesis functions with  $h : \mathcal{X} \rightarrow \mathcal{Y}$  for any  $h \in \mathcal{H}$ . In addition, for any hypothesis function  $f \in \mathcal{F}$  and loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ , the expected prediction errors for the source and target domains are denoted by  $\mathcal{E}_S = \mathbb{E}_{p_S}[\ell(f(x), y)]$  and  $\mathcal{E}_T = \mathbb{E}_{p_T}[\ell(f(x), y)]$ , respectively.

### 2.2 Problem Definition

Transfer learning (Pan and Q. Yang 2010) refers to the knowledge or information transfer from the source domain to the target domain such that the prediction performance on the target domain could be significantly improved as compared to learning from the target domain alone. Moreover, in the following definition, we generalize standard transfer learning (Pan and Q. Yang 2010) to trustworthy transfer learning (J. Wu and J. He 2023b).

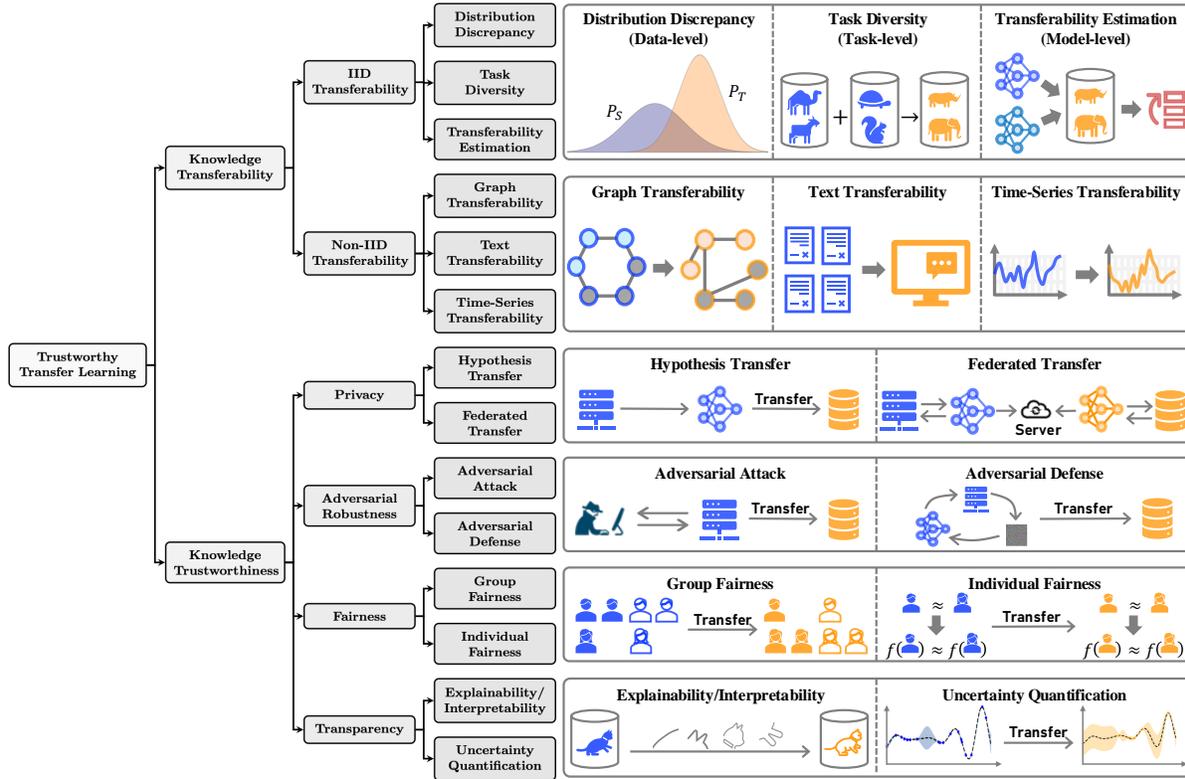


Fig. 2. Overview of trustworthy transfer learning (best viewed in color)

*Definition 2.1 (Trustworthy Transfer Learning).* Given a source domain  $\mathcal{D}_S$  and a target domain  $\mathcal{D}_T$ , trustworthy transfer learning aims at improving the generalization and trustworthiness of a learning algorithm  $f(\cdot)$  on the target domain, by leveraging latent knowledge from the source domain.

The source and target domains might involve different learning tasks (Pan and Q. Yang 2010; Tripuraneni, Jordan, et al. 2020) or data modalities (Bugliarello et al. 2022; J. Shen, L. Li, et al. 2023). There are two key components in trustworthy transfer learning: knowledge transferability and trustworthiness. Specifically, knowledge transferability measures how the source knowledge can be successfully transferred to the target domain. In contrast, knowledge trustworthiness aims to answer whether transfer learning techniques provide reliable and trustworthy results in the target domain. Fig. 2 provides a brief summarization of trustworthy transfer learning regarding knowledge transferability and trustworthiness (discussed in Section 3 and Section 4).

### 3 Knowledge Transferability

This section summarizes the knowledge transferability in different scenarios.

#### 3.1 IID Transferability

Here, we summarize different transferability indicators, including distribution discrepancy, task diversity, and transferability measures.

**3.1.1 Distribution Discrepancy.** Distribution discrepancy quantitatively measures the distribution shifts between two domains in the distribution space, when the source and target domains share the same input and output spaces (this scenario is also known as domain adaptation (Pan and Q. Yang 2010)). There are different types of distribution shifts (Wiles et al. 2022), including covariate shift (Shimodaira 2000) (feedback covariate shift (Fannjiang et al. 2022; Prinster, Saria, et al. 2023)), label/target shift (Lipton et al. 2018; K. Zhang et al. 2013), concept shift (Redko, Morvant, et al. 2019), etc. The covariate shift holds that the conditional probability  $p(y|x)$  is shared across domains, but the marginal  $p(x)$  is different. Label shift assumes that the conditional probability  $p(x|y)$  is shared across domains, while the marginal label distribution  $p(y)$  changes. Concept shift involves changes in the conditional probability  $p(x|y)$  (or  $p(y|x)$ ), while the marginal distribution  $p(y)$  (or  $p(x)$ ) is fixed. The integral probability metric (IPM) (Müller 1997; Sriperumbudur et al. 2010; C. Zhang et al. 2012) is a general framework for quantifying the difference between two distributions, and it can be instantiated by various statistical discrepancy measures (Sriperumbudur et al. 2010), e.g., total variation distance, Wasserstein distance, maximum mean discrepancy (Gretton et al. 2012), etc.

*Definition 3.1 (Integral Probability Metric (Müller 1997)).* Let  $P_S$  and  $P_T$  be the probability distributions of the source  $\mathcal{D}_S$  and target  $\mathcal{D}_T$  domains, respectively. The integral probability metric between  $P_S$  and  $P_T$  is defined as:

$$d_{\text{IPM}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h \in \mathcal{H}} \left| \int_M h dP_S - \int_M h dP_T \right| \quad (1)$$

where  $M$  is a measurable space and  $\mathcal{H}$  is a class of real-valued bounded measurable functions on  $M$ .

The concept of distribution discrepancy is the key to theoretically understanding how knowledge can be transferred from source to target domains. For example, the seminal work of Ben-David, Blitzer, et al. (2010) derives a generalization bound for domain adaptation using a tractable  $\mathcal{H}\Delta\mathcal{H}$ -divergence. Many follow-up works have developed refined generalization bounds by introducing various discrepancy measures. The following theorem provides a unified view of such generalization bounds based on a notion of discrepancy  $d(\mathcal{D}_S, \mathcal{D}_T)$ .

**THEOREM 3.2 (UNIFIED GENERALIZATION BOUND).** *Let  $\mathcal{H}$  denote the hypothesis space, and  $\mathcal{E}_S(h), \mathcal{E}_T(h)$  be the expected prediction error of a hypothesis  $h \in \mathcal{H}$  on the source and target domains, respectively.  $d(\cdot, \cdot)$  measures the difference between source and target distribution probabilities (see more instantiations below). Then for any hypothesis  $h \in \mathcal{H}$ , we can have a unified view of generalization error in the target domain:*

$$\mathcal{E}_T(h) \leq \mathcal{E}_S(h) + d(\mathcal{D}_S, \mathcal{D}_T) + \Omega$$

where  $\Omega$  represents the redundant terms (depending on how  $d(\mathcal{D}_S, \mathcal{D}_T)$  is instantiated), e.g., the difference of labeling functions across domains (Acuna et al. 2021; Ben-David, Blitzer, et al. 2010), the complexity of hypothesis space  $\mathcal{H}$  (Mansour, Mohri, and Rostamizadeh 2009a; C. Zhang et al. 2012), number of training samples (Ben-David, Blitzer, et al. 2010; Redko, Habrard, et al. 2017), etc.

We have the following observations regarding this unified view of generalization error. (1) The complexity of the class of hypothesis functions  $\mathcal{H}$  plays a crucial role in deriving tight generalization error bounds. Various metrics have been applied to quantify this complexity (Redko, Morvant, et al. 2019), including the Vapnik-Chervonenkis (VC) dimension (Ben-David, Blitzer, et al. 2010; Ben-David et al. 2006; Blitzer et al. 2007; X. Peng, Q. Bai, et al. 2019), Rademacher complexities (Acuna et al. 2021; Ghifary et al. 2016; Mansour, Mohri, and Rostamizadeh 2009a; Mohri and Muñoz Medina 2012; Y. Zhang, T. Liu, et al. 2019), covering number (C. Zhang et al. 2012; Y. Zhang, T. Liu, et al. 2019), etc. We refer the reader to the survey (Redko, Morvant, et al. 2019) for more discussion. (2) Generally, the discrepancy  $d(\mathcal{D}_S, \mathcal{D}_T)$  measures the difference between source and target distributions over the joint space  $\mathcal{X} \times \mathcal{Y}$ , when the distribution shifts occur across domains. In practice, it is commonly seen that the discrepancy  $d(\mathcal{D}_S, \mathcal{D}_T)$  is defined over the input space  $\mathcal{X}$  (when no label information is available in

unsupervised domain adaptation), or over the joint space  $\mathcal{X} \times \mathcal{Y}$  (when (pseudo)labels for target samples are available). The first type of  $d(\mathcal{D}_S, \mathcal{D}_T)$ , defined over the input space  $\mathcal{X}$ , is often associated with the covariate shift assumption (Shimodaira 2000) or the redundant term indicating the difference of labeling functions across domains, when deriving the generalization error bound. In the following, we summarize several commonly used discrepancy metrics  $d(\mathcal{D}_S, \mathcal{D}_T)$ .

- **Total Variation Distance:** The total variation distance (also referred to as  $L^1$  divergence) (Ben-David et al. 2006) between source and target domains can be defined as

$$d_{\text{TV}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{B \in \mathcal{B}} |\mathbb{P}_S[B] - \mathbb{P}_T[B]| \quad (2)$$

where  $\mathcal{B}$  is the set of measurable subsets under  $\mathbb{P}_S$  and  $\mathbb{P}_T$ . It is shown (Sriperumbudur et al. 2010) that the total variation distance can be considered as a special case of the integral probability metric.

- **$\mathcal{H}\Delta\mathcal{H}$ -divergence:** It is illustrated (Ben-David, Blitzer, et al. 2010; Ben-David et al. 2006) that the empirical estimate of the total variation distance in Eq. (2) has two limitations. First, it cannot be accurately estimated from finite samples of arbitrary distributions in practice. Second, it results in loose generalization bounds due to involving a supremum over all measurable subsets. To address these limitations, Ben-David et al. (2006) and Blitzer et al. (2007) introduce the following  $\mathcal{H}\Delta\mathcal{H}$ -divergence:

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h, h' \in \mathcal{H}} |\mathbb{P}_S[h(x) \neq h'(x)] - \mathbb{P}_T[h(x) \neq h'(x)]| \quad (3)$$

It can be seen that this distribution difference is defined over the hypothesis-dependent subsets  $\{\mathbb{I}[h(x) \neq h'(x)] | h, h' \in \mathcal{H}\}$ .

- **Discrepancy Distance:** Mansour, Mohri, and Rostamizadeh (2009a) extends the  $\mathcal{H}\Delta\mathcal{H}$ -divergence to a more general discrepancy distance for measuring distribution differences.

$$d_{\text{disc}}(\mathcal{D}_S, \mathcal{D}_T) = \max_{h, h' \in \mathcal{H}} |\mathbb{E}_{x \sim \mathbb{P}_S} [\ell(h(x), h'(x))] - \mathbb{E}_{x \sim \mathbb{P}_T} [\ell(h(x), h'(x))]| \quad (4)$$

where  $\ell(\cdot, \cdot)$  denotes a general loss function (though the derived generalization bounds require that  $\ell(\cdot, \cdot)$  is symmetric and obeys the triangle inequality). When using 0-1 classification loss, this discrepancy distance exactly recovers the  $\mathcal{H}\Delta\mathcal{H}$ -divergence. The discrepancy distance can be flexibly applied to compare distributions across various tasks, e.g., regression (Cortes and Mohri 2011; Mansour, Mohri, and Rostamizadeh 2009a).

- **$\mathcal{Y}$ -discrepancy:** It is notable that the discrepancy distance in Eq. (4) quantifies the difference between two marginal distributions over  $\mathcal{X}$ , when the ground-truth labeling function is unknown in the target domain. Later, Mohri and Muñoz Medina (2012) further extend the discrepancy distance to the  $\mathcal{Y}$ -discrepancy, which is defined over  $\mathcal{X} \times \mathcal{Y}$  as follows.

$$d_{\mathcal{Y}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h \in \mathcal{H}} |\mathbb{E}_{(x,y) \sim \mathbb{P}_S} [\ell(h(x), y)] - \mathbb{E}_{(x,y) \sim \mathbb{P}_T} [\ell(h(x), y)]| \quad (5)$$

In practice, this discrepancy can be estimated using pseudo labels of the target data, when there are no labeled data in the target domain (Courty, Flamary, Habrard, et al. 2017; Long, Z. Cao, et al. 2018).

- **Margin Disparity Discrepancy:** Y. Zhang, T. Liu, et al. (2019) extend the notion of discrepancy distance in Eq. (4) to margin disparity discrepancy (MDD) in multi-class classification settings. Specifically, MDD involves two key refinements: (1) the use of a margin-based loss function, and (2) the formulation of the discrepancy over both a hypothesis space  $\mathcal{H}$  and a specific classifier  $h$ .

$$d_{\text{MDD}}(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h' \in \mathcal{H}} |\mathbb{E}_{x \sim \mathbb{P}_S} [\Phi_\rho(\rho_{h'}(x, h(x)))] - \mathbb{E}_{x \sim \mathbb{P}_T} [\Phi_\rho(\rho_{h'}(x, h(x)))]| \quad (6)$$

where the function  $\rho_{h'}(\cdot, \cdot)$  defines the margin of a hypothesis  $h'$ , and the function  $\Phi_\rho(\cdot)$  defines the margin-based loss over a threshold  $\rho > 0$ .

- $f$ -divergence: Building on the margin disparity discrepancy (MDD) (Y. Zhang, T. Liu, et al. 2019), Acuna et al. (2021) further develop a generic notion of the discrepancy based on the variational characterization of  $f$ -divergence (X. Nguyen et al. 2010). Specifically, the  $f$ -divergence is bounded by

$$d_f(\mathcal{D}_S, \mathcal{D}_T) = \int p_T(x) \phi\left(\frac{p_S(x)}{p_T(x)}\right) dx \geq \sup_{T \in \mathcal{T}} \mathbb{E}_{x \sim P_S} [T(x)] - \mathbb{E}_{x \sim P_T} [\phi^*(T(x))]$$

where  $\phi(\cdot)$  is a convex lower semi-continuous function that satisfies  $\phi(1) = 0$ .  $\phi^*$  is the conjugate function of  $\phi$ .  $\mathcal{T}$  is a set of measurable functions. Based on this observation, (Acuna et al. 2021) define a notion of  $f$ -divergence guided discrepancy as follows.

$$d_\phi(\mathcal{D}_S, \mathcal{D}_T) = \sup_{h' \in \mathcal{H}} \left| \mathbb{E}_{x \sim P_S} [\ell(h(x), h'(x))] - \mathbb{E}_{x \sim P_T} [\phi^*(\ell(h(x), h'(x)))] \right| \quad (7)$$

The flexibility in choosing  $\phi^*$  enables  $f$ -divergence to recover many popular statistical divergences, e.g., Jensen-Shannon (JS) divergence, Kullback-Leibler (KL) divergence, Reverse KL (KL-rev) divergence, Pearson  $\chi^2$  divergence, etc. As a result, different choices of  $\phi^*$  can define various discrepancies from Eq. (7). Besides, it is shown that the notion of discrepancy in Eq. (7) can also recover MDD.

- Generalized Discrepancy: In contrast, Cortes, Mohri, and Medina (2015) and Cortes, Mohri, and Medina (2019) generalize the discrepancy distance in Eq. (4) using reweighting techniques. That is, the difference between two distributions can be adjusted by multiplying the loss for each training example by a non-negative weight (Cortes, Mohri, Riley, et al. 2008; J. Huang et al. 2006; K. Zhang et al. 2013). Formally, for any hypothesis-dependent reweighting function  $U_h$ , the generalized discrepancy is defined as follows.

$$d_{\text{DISC}}(\mathcal{D}_S, \mathcal{D}_T) = \max_{h \in \mathcal{H}, h'' \in \mathcal{H}''} \left| \mathbb{E}_{x \sim \hat{P}_S} [U_h(x) \cdot \ell(h(x), f_S(x))] - \mathbb{E}_{x \sim \hat{P}_T} [\ell(h(x), h''(x))] \right| \quad (8)$$

where  $\hat{P}_S, \hat{P}_T$  denote the empirical distributions of source and target domains, respectively, and  $f_S$  denotes the source labeling function.

- Rényi Divergence: Furthermore, Cortes, Mansour, et al. (2010) and Mansour, Mohri, and Rostamizadeh (2009b) derive the generalization bounds for adaptation approaches based on importance reweighting (e.g., sample reweighting for single-source adaptation (Cortes, Mansour, et al. 2010) and domain reweighting for multi-source adaptation (Mansour, Mohri, and Rostamizadeh 2009b)) based on the following Rényi divergence (Rényi 1961).

$$d_\alpha(\mathcal{D}_T || \mathcal{D}_S) = \frac{1}{\alpha - 1} \log \sum_{x \in \mathcal{X}} P_T(x) \left( \frac{P_T(x)}{P_S(x)} \right)^{\alpha-1} \quad (9)$$

where  $\alpha \geq 0$ .

- Wasserstein Distance (J. Shen, Qu, et al. 2018): In general, for any  $p \geq 1$ , the  $p$ -Wasserstein distance between two distributions can be defined as follows.

$$d_{W_p}(\mathcal{D}_S, \mathcal{D}_T) = \left( \inf_{\pi \in \Pi(P_S, P_T)} \int c(x, x')^p d\pi(x, x') \right)^{1/p} \quad (10)$$

where  $\Pi(P_S, P_T)$  is the set of all measures over  $X \times X$  with marginals  $P_S$  and  $P_T$ , and  $c(\cdot, \cdot)$  is a distance function. When  $p = 1$ , the 1-Wasserstein distance (also known as earth mover's distance) is one special

case of the integral probability metric with  $\mathcal{H} = \{h : \|h\|_L \leq 1\}$ . Specifically, based on the Kantorovich-Rubinstein duality (Dudley 2002), it holds that

$$d_{W_1}(\mathcal{D}_S, \mathcal{D}_T) = \inf_{\pi \in \Pi(\mathcal{P}_S, \mathcal{P}_T)} \int c(x, x') d\pi(x, x') = \sup_{\|h\|_L \leq 1} \mathbb{E}_{\mathcal{P}_S}[h(x)] - \mathbb{E}_{\mathcal{P}_T}[h(x)] \quad (11)$$

where  $\|h\|_L = \sup_{x \neq x'} |h(x) - h(x')|/c(x, x')$ . This enables a practical empirical estimation of the 1-Wasserstein distance using gradient descent optimization (Arjovsky et al. 2017). Therefore, Redko, Habrard, et al. (2017) and J. Shen, Qu, et al. (2018) apply the 1-Wasserstein distance to analyze distribution shifts between source and target domains.

- Maximum Mean Discrepancy: Long, Y. Cao, et al. (2015) and Tzeng, Hoffman, N. Zhang, et al. (2014) leverage the maximum discrepancy discrepancy (MMD) (Gretton et al. 2012) to measure the distribution difference between source and target domains. MMD can be considered as another special case of the integral probability metric by instantiating the hypothesis space with a unit ball in a reproducing kernel Hilbert space associated with kernel  $k(\cdot, \cdot)$ . Given a kernel function  $k(\cdot, \cdot)$ , the MMD between source and target distribution can be defined as:

$$d_{\text{MMD}}(\mathcal{D}_S, \mathcal{D}_T) = \mathbb{E}_{x_S, x'_S \sim \mathcal{P}_S} [k(x_S, x'_S)] - 2\mathbb{E}_{x_S \sim \mathcal{P}_S, x_T \sim \mathcal{P}_T} [k(x_S, x_T)] + \mathbb{E}_{x_T, x'_T \sim \mathcal{P}_T} [k(x_T, x'_T)] \quad (12)$$

Redko, Morvant, et al. (2019) further show the generalization error bounds based on MMD.

- Cauchy-Schwarz Divergence (Yin et al. 2024): Recently, Yin et al. (2024) use the Cauchy-Schwarz (CS) divergence (Principe 2010) to theoretically understand the knowledge transferability across domains.

$$d_{\text{CS}}(\mathcal{D}_S, \mathcal{D}_T) = -\log \left( \frac{(\int p_S(x)p_T(x)dx)^2}{\int p_S^2(x)dx \cdot \int p_T^2(x)dx} \right) \quad (13)$$

It is shown (Yin et al. 2024) that this CS divergence can lead to tighter generalization error bounds than the KL divergence (A. T. Nguyen, T. Tran, et al. 2022). Besides, the empirical estimate of the CS divergence is closely related to MMD (Gretton et al. 2012).

From the perspective of empirical estimation, the discrepancy measures can be broadly categorized into two groups. The first group includes statistical discrepancy measures (X. Chen, S. Wang, et al. 2021; A. T. Nguyen, T. Tran, et al. 2022; B. Sun and Saenko 2016), such as Maximum Mean Discrepancy (MMD) (Long, Y. Cao, et al. 2015; Tzeng, Hoffman, N. Zhang, et al. 2014) and Wasserstein distance (Courty, Flamary, Tuia, et al. 2016; Fatras et al. 2021; Redko, Habrard, et al. 2017), which can be directly estimated from finite samples. The second group is based on adversarial learning (Acuna et al. 2021; Ganin et al. 2016; Hoffman, Tzeng, et al. 2018; Saito et al. 2018; Tzeng, Hoffman, Saenko, et al. 2017; Y. Zhang, T. Liu, et al. 2019), which requires an additional neural network to optimize an adversarial objective. More recently, Kashyap et al. (2021) and Z. Yuan et al. (2022) provide empirical comparisons of various discrepancy measures in natural language processing and computer vision tasks. It is noted that when there are no labeled samples in the target domain, one common strategy in designing practical domain adaptation algorithms is to minimize the discrepancy across domains over  $\mathcal{X}$ . However, it has been shown (Ben-David, T. Lu, et al. 2010; Johansson et al. 2019; Y. Wu et al. 2019; H. Zhao, Combes, et al. 2019) that exact marginal distribution matching might lead to negative transfer in practice.

The notion of distribution discrepancy has been applied to understand knowledge transferability in various realistic adaptation scenarios, including single-source adaptation (Acuna et al. 2021; Ben-David, Blitzer, et al. 2010; Cortes, Mansour, et al. 2010; A. T. Nguyen, T. Tran, et al. 2022; Y. Zhang, T. Liu, et al. 2019), multi-source adaptation (Hoffman, Mohri, et al. 2018; Mansour, Mohri, Ro, et al. 2021; J. Wu, J. He, and Tong 2024), open-set adaptation (Fang et al. 2020; H. He et al. 2023), domain generalization (Blanchard et al. 2011; Muandet et al. 2013)

(also known as out-of-distribution generalization), privacy-preserving federated adaptation (X. Peng, Z. Huang, et al. 2020), dynamic adaptation (Kumar et al. 2020; J. Wu and J. He 2022), etc.

**3.1.2 Task Diversity.** Task diversity (Tripuraneni, Jordan, et al. 2020; Watkins et al. 2023) is another tool for theoretically understanding the performance of transfer learning. It enables a relaxed data assumption that the source and target domains can have different output spaces, i.e., each domain can be associated with a different learning task (Pan and Q. Yang 2010). In the context of transfer learning, it assumes that a generic nonlinear feature representation function is shared across all tasks. Then each task is associated with a shared representation learning function and a task-specific prediction function. In (Tripuraneni, Jordan, et al. 2020), task diversity is defined to characterize how the worst-case representation difference can be controlled when the task-averaged representation difference is small. In this case, the worst-case representation difference is the distance between two representation functions with the worst-case task-specific prediction function, while the task-averaged representation difference indicates the distance between two representation functions over all the training tasks.

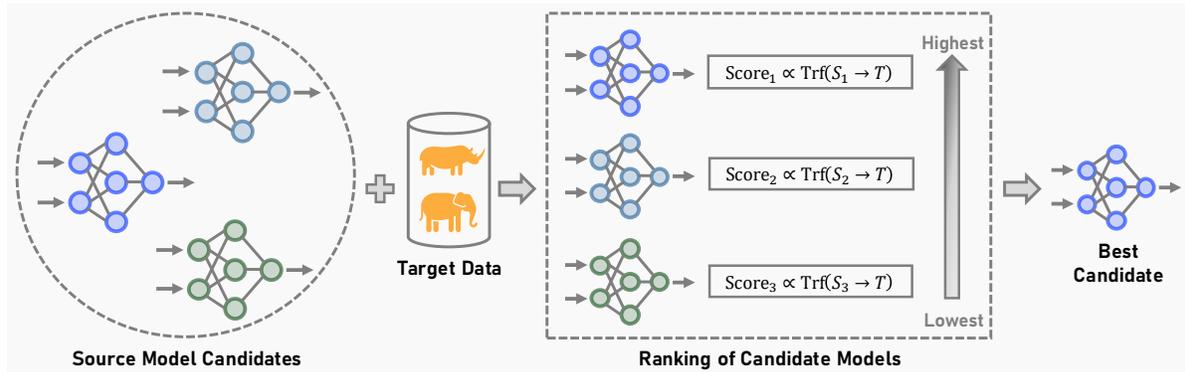
*Definition 3.3 (Task Diversity (Tripuraneni, Jordan, et al. 2020)).* Given  $N$  source tasks associated with a representation function class  $\mathcal{H}$  and a prediction function class  $\mathcal{H}$ , let  $h_i \in \mathcal{H}$  represent the task-specific prediction function for the  $i^{\text{th}}$  source task ( $i = 1, 2, \dots, N$ ), we say that source tasks with the functions  $\{h_1, h_2, \dots, h_N\}$  are  $(\nu, \epsilon)$ -diverse over the function class  $\mathcal{H}_0$  for a representation function  $g \in \mathcal{G}$ , if uniformly for all  $g' \in \mathcal{G}$ ,

$$\underbrace{\sup_{h_0 \in \mathcal{H}_0} \inf_{h' \in \mathcal{H}} \{\mathcal{E}_T(h' \circ g') - \mathcal{E}_T(h_0 \circ g)\}}_{\text{Worst-case representation difference}} \leq \frac{1}{\nu} \cdot \underbrace{\left( \frac{1}{N} \sum_{i=1}^N \inf_{h' \in \mathcal{H}} \{\mathcal{E}_{S_i}(h' \circ g') - \mathcal{E}_{S_i}(h_i \circ g)\} \right)}_{\text{Task-averaged representation difference}} + \epsilon \quad (14)$$

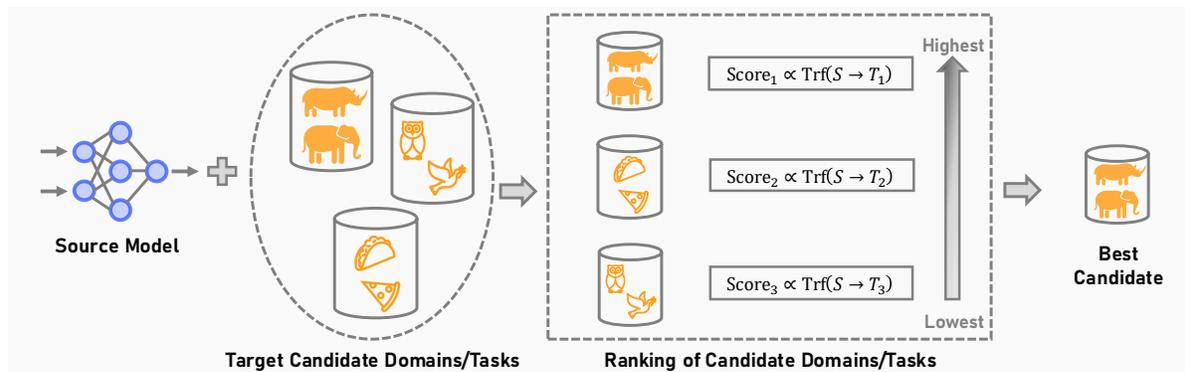
where  $\mathcal{E}_T(h \circ g)$  represents the expected error in the target task using a representation learning function  $g$  and a prediction function  $h$ , and  $\mathcal{E}_{S_i}(h \circ g)$  represents the expected error in the  $i^{\text{th}}$  source task.

Based on the task diversity, Tripuraneni, Jordan, et al. (2020) derive the excess risk bounds of transfer learning for the target task in terms of the complexity of the shared representation function class  $\mathcal{G}$ , the complexity of the prediction function class  $\mathcal{H}$ , the number of tasks  $N$ , and the number of training samples for each task (both source and target). Furthermore, Watkins et al. (2023) show that under the Lipschitz assumption for the loss function, the excess risk in the target task only achieves the standard rate of  $O(n_T^{-1/2})$ , where  $n_T$  is the number of training samples in the target task. By using the smoothness assumption for the loss function (Srebro et al. 2010), they derive optimistic rates that interpolate between the standard rate of  $O(n_T^{-1/2})$  and the fast rate of  $O(n_T^{-1})$  for the excess risk in the target task.

In addition, S. S. Du, W. Hu, et al. (2021) and Tripuraneni, C. Jin, et al. (2021) consider a simplified version of task diversity in the cases of linear prediction functions and quadratic loss. They also theoretically show the benefits (i.e., reduced sample complexity in the target task induced by all available source samples) of representation learning from source tasks. Based on the task diversity, Z. Xu and Tewari (2021) analyze more realistic learning scenarios in which the source and target tasks use different prediction function spaces. However, all the aforementioned theoretical analyses assume uniform sampling from each source task, i.e., all source tasks are equally important for learning a representation function. Instead, Y. Chen, Y. Huang, et al. (2023), Y. Chen, Jamieson, et al. (2022), and Y. Wang, Y. Chen, et al. (2023) study active transfer learning by quantifying the task relatedness and selecting the source tasks that are most relevant to the target task. Similarly, Z. Xu, Z. Shi, et al. (2024) explore the selection of source tasks for multi-task fine-tuning of foundation models, e.g., fine-tuning the foundation model on auxiliary source tasks before adapting it to the target task with limited labeled samples. More recently, Y. Zhao et al. (2023) show that pre-training on a single source task with a high diversity of classes



(a) Transferability estimation for selecting the source model



(b) Transferability estimation for selecting the target domain/task

Fig. 3. Evaluation of transferability between the pre-trained source model and the target data: (a) Transferability scores select the best source model for the target data given a large pool of pre-trained source models. (b) Transferability scores identify the most suitable application domains/tasks for a source model.

can provably improve the sample efficiency of the downstream tasks. In contrast, [Cole et al. \(2024\)](#) leverage task diversity to understand the in-context learning behavior of foundation models.

**3.1.3 Transferability Estimation.** In contrast to the data-centric transferability analyses in Subsection 3.1.1 and Subsection 3.1.2, this subsection explores the knowledge transferability of pre-trained source models. This is driven by the rapidly expanding open-source model repositories such as HuggingFace ([Wolf et al. 2020](#)) and PyTorch Hub ([Paszke et al. 2019](#)). Fine-tuning a pre-trained source model on downstream target data sets with limited sample sizes improves model accuracy and robustness ([Hendrycks et al. 2019](#)). A natural question arises in this scenario: *Given a large pool of pre-trained source models, how can we efficiently select the best one for a target data set?* As shown in Fig. 3, another relevant question is how to identify the most suitable domains/tasks for a given pre-trained source model. One trivial solution is brute-force fine-tuning, where all source models are fine-tuned individually and then ranked based on their transfer accuracy. However, this method is highly time-consuming and computationally expensive. To solve this problem, transferability estimation has been studied to quantitatively measure how effectively the knowledge can be transferred from a pre-trained source model to a target domain ([Agostinelli, Pándy, et al. 2022](#); [Y. Bao et al. 2019](#); [Ibrahim et al. 2022](#); [C. V. Nguyen et al. 2020](#);

A. T. Tran et al. 2019). Following (A. T. Tran et al. 2019), given a pre-trained source model  $f_S(\cdot)$ , the transferability from  $f_S(\cdot)$  to a target domain associated with sampling distribution  $P_T$  can be defined below.

*Definition 3.4 (Transferability Measure (A. T. Tran et al. 2019)).* The transferability from a pre-trained source model  $\mathcal{D}_S$  to a target domain  $\mathcal{D}_T$  with sampling distribution  $P_T$  is measured by the expected accuracy of the fine-tuned model on the target domain:

$$\text{Trf}(S \rightarrow T) = \mathbb{E}_{(x,y) \sim P_T} [\text{acc}(x, y; f_T)] \quad (15)$$

where  $f_T(\cdot)$  is the fine-tuned model from a pre-trained source model  $f_S(\cdot)$ , and  $\text{acc}(\cdot)$  indicates the prediction accuracy.

Thus, a good transferability measurement should have two key properties: (1) the learned transferability score correlates well with the transfer accuracy of the fine-tuned model on the target domain, and (2) it should be significantly more efficient than the fine-tuning approach. Notably, the transferability score does not exactly predict the accuracy of the fine-tuned model on the target domain in practice. Instead, it only needs to correlate with the ranking of fine-tuning accuracy among a pool of source models, i.e., a high transferability score indicates a better source model resulting in higher transfer accuracy.

The estimation of the transferability score is related to the architectures of pre-trained source models. In scenarios where the source models are trained on supervised classification tasks, every source model  $f_S(\cdot)$  is associated with a feature extractor and a predictor. In this case, NCE (A. T. Tran et al. 2019) leverages conditional entropy to define the transferability score, assuming that source and target domains share the same input samples but different labels. It is motivated by the observation that the optimal average log-likelihood on target training samples is lower bounded by the negative conditional entropy. Similarly, LEEP (C. V. Nguyen et al. 2020) estimates the average log-likelihood of target samples using the dummy label distributions generated from the pre-trained source model. Notably, it is shown that LEEP is an upper bound of the NCE (A. T. Tran et al. 2019) plus the average log-likelihood of the dummy labels. It is computationally efficient by using only a single forward pass of the source model through the target data. Nevertheless, LEEP can not handle unsupervised and self-supervised pre-trained models with only a feature extractor.

To address this problem, recent works (Bolya et al. 2021; C. N. Nguyen et al. 2023; K. You et al. 2021) have utilized only the feature extractor  $f_S(\cdot)$  to define transferability scores. They can be divided into two frameworks. One is to measure the class separability of the target samples in the feature space induced by  $f_S(\cdot)$ . For example, H-score (Y. Bao et al. 2019) is defined based on the inter-class variance and feature redundancy of target samples learned from the pre-trained source model. It is inspired by the connection between the optimal prediction error and the modal decomposition of the divergence transition matrix (S.-L. Huang et al. 2024). The follow-up work (Ibrahim et al. 2022) introduces a shrinkage-based H-score to improve the covariance estimation of H-score (Y. Bao et al. 2019) in high-dimensional feature spaces.  $\mathcal{N}$ LEEP (Y. Li, Jia, et al. 2021) replaces the dummy label generation module of LEEP (C. V. Nguyen et al. 2020) with a new Gaussian Mixture Model (GMM). Furthermore, GBC (Pándy et al. 2022) maps each target class as a Gaussian distribution and estimates the pair-wise class separability (i.e., the amount of overlap between two class-wise Gaussian distributions) using the Bhattacharyya coefficient (Bhattacharyya 1946). The other one is to add a probabilistic linear transformation that maps the feature space of  $f_S(\cdot)$  to the target output space in a Bayesian framework. For example, LogME (K. You et al. 2021) is defined over marginalized likelihood  $p(y_i|x_i; f_S) = \int p(w)p(y_i|f_S(x_i), w)dw$ , assuming that the prior distribution of the newly added linear transformation  $w$  is an isotropic multivariate Gaussian  $w \sim \mathcal{N}(0, \alpha^{-1}\mathbb{I})$ . The follow-up work PACTran (N. Ding, X. Chen, et al. 2022) further defines a theoretically grounded family of transferability scores based on the optimal PAC-Bayesian error bound (Germain et al. 2016), taking into consideration various instantiations of the prior distribution for the linear transformation  $w$ , such as Dirichlet, Gamma, and Gaussian priors. Additionally, TransRate (L. Huang et al. 2022) exploits the coding rate to estimate

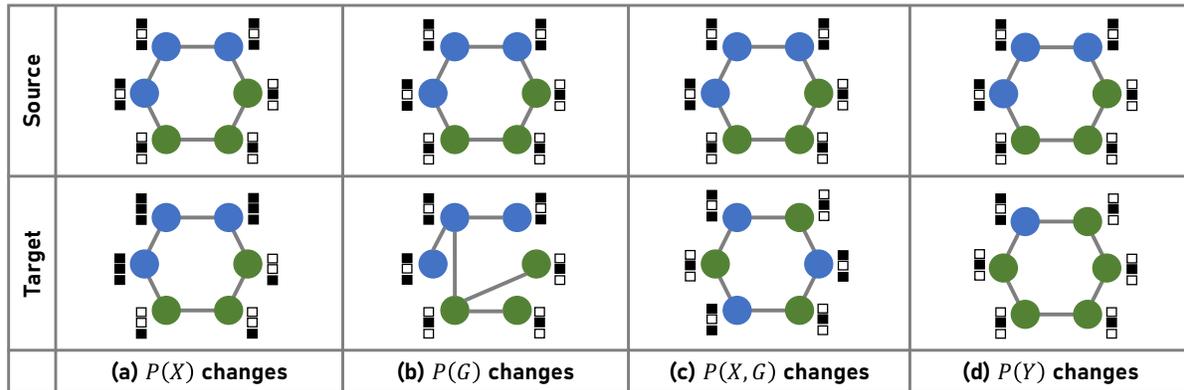


Fig. 4. Illustration of distribution shifts in characterizing graph transferability. The node distribution within the graph can be represented by  $P(X, G, Y)$  where  $X, G, Y$  denote the input node attributes, topology structure, and output class labels, respectively. The color of nodes indicates the class labels  $Y$  (blue or green).

transferability scores of any intermediate layer within the pre-trained model. More recently, inspired by neural scaling laws for fine-tuned LLMs (Tay et al. 2022), Lin et al. (2024) investigated the transferability estimation of large language models (LLMs) based on a rectified scaling law that characterizes the connection between the fine-tuned test loss and the number of target samples.

The transferability metrics mentioned above enable the selection of the best source model from a large pool of open-sourced pre-trained models. Recent studies (B. et al. 2023; Shao et al. 2022) take one step further by studying source model ensemble selections, which transfer knowledge from multiple pre-trained source models to target training samples. This line of research is inspired by the success of ensemble machine learning models (Lakshminarayanan et al. 2017) in improving model performance. Specifically, Agostinelli, Uijlings, et al. (2022) extend LEEP (C. V. Nguyen et al. 2020) to select source model ensembles under the assumption that the source models operate independently. In contrast, OSBORN (B. et al. 2023) relaxes this assumption and explores the inter-model cohesion among source models to estimate the transferability of an ensemble of models to a target domain.

### 3.2 Non-IID Transferability

The IID assumption is often violated in real scenarios, e.g., connected nodes in graphs (Kipf and Welling 2017), word occurrence in texts (J. Y. Lee et al. 2018), temporal observations in time series (Purushotham et al. 2017), etc. To bridge this gap, non-IID transferability explores the knowledge transfer across domains, assuming that samples within each domain can be interdependent.

**3.2.1 Transferability on Graph Data.** Graph data is being generated across a variety of application domains, ranging from bioinformatics (Gilmer et al. 2017) to e-commerce (J. Wu, J. He, and E. A. Ainsworth 2023), from protein-protein interaction prediction (Hamilton et al. 2017) to social network analysis (K. Xu et al. 2019). To capture the complex structure of graph data, graph neural networks (GNNs) (Defferrard et al. 2016; Scarselli et al. 2009) have been introduced to encode the nodes within the graphs into low-dimensional vector representations. It is shown (S. Zhang et al. 2019) that there are two major learning paradigms for GNNs: spectral-based (Defferrard et al. 2016; Kipf and Welling 2017) and spatial-based GNNs (Gilmer et al. 2017; Hamilton et al. 2017; K. Xu et al. 2019). Recently, the transferability of spectral and spatial GNNs has been studied by exploring whether GNNs are transferable across graphs of varying sizes and topologies (Ruiz, Chamon, et al. 2020; J. Wu, J. He, and E. A.

Ainsworth 2023). In this survey, we focus on understanding the transferability of GNNs in node-level graph learning tasks. Note that the key challenge to theoretically understand the transferability of GNNs is to measure the distribution shifts of two graphs. As illustrated in Fig. 4, the distribution shifts between source and target graphs are generally induced by the joint probabilities  $P_S(X, G, Y)$  and  $P_T(X, G, Y)$ , where  $X, G, Y$  denote the input node attributes, topology structure, and output class labels, respectively. Specifically, the transferability of spectral GNNs leverages the graph limits, e.g., graphon (Maskey et al. 2023; Ruiz, Chamon, et al. 2020) and graphop (Le and Jegelka 2023), to determine if two graphs represent the same underlying structure as the number of nodes goes to infinity. In contrast, the transferability of spatial GNNs typically relies on the empirical distribution differences of node representations in a latent embedding space learned by GNNs (J. Wu, J. He, and E. A. Ainsworth 2023).

Spectral GNNs define the graph convolutions in the spectral domain using the graph Fourier transform from the perspective of graph signal processing (Defferrard et al. 2016; Kipf and Welling 2017). Recent efforts (Le and Jegelka 2023; Levie, W. Huang, et al. 2021; Maskey et al. 2023; Ruiz, Chamon, et al. 2020; Ruiz, Gama, et al. 2021) have been dedicated to understanding the transferability of spectral GNNs by answering the following question: *Can spectral GNNs trained on a source graph perform well on a target graph of different sizes?* This question is also known as size generalization (Bevilacqua et al. 2021; Yehudai et al. 2021). The intuition behind the transferability of spectral GNNs is that if two graphs represent the same underlying phenomenon, their GNN outputs will be similar. Thus, the transferability of spectral GNNs can be derived from various aspects, including generic topological space (Levie, W. Huang, et al. 2021), graphon (Ruiz, Chamon, et al. 2020), graphop (Le and Jegelka 2023), and  $k$ -hop ego-graph (Q. Zhu et al. 2021). To be more specific, Levie, W. Huang, et al. (2021) and Levie, Isufi, et al. (2019) study the transferability of spectral graph filters on different discretizations of the same underlying continuous topological space. Later, graphon theory (Lovász 2012) is used to analyze the transferability of spectral GNNs. Formally, a graphon is defined by a bounded symmetric kernel and can be viewed as a graph with an uncountable number of nodes. In particular, Ruiz, Chamon, et al. (2020) leverage graphon to study the asymptotic behavior of GNNs (Defferrard et al. 2016), showing that GNNs converge to graphon neural networks (WNNs) as the number of nodes increases to infinity. This convergence implies that under mild assumptions, GNNs are transferable across graphs with performance guarantees if both graphs are drawn from the same graphon (Maskey et al. 2023; Ruiz, Chamon, et al. 2020; Ruiz, Gama, et al. 2021). Following this observation, recent works (Cerviño et al. 2023; Krishnagopal and Ruiz 2023) further demonstrate the transferability of the gradients of spectral-based GNNs across graphs under similar conditions. Furthermore, using the graphop operator (Backhausz and B. Szegedy 2022), Le and Jegelka (2023) extend the transferability analysis of GNNs to both dense and sparse graphs. Besides, assuming that the  $k$ -hop ego-graphs are independent and identically drawn, Q. Zhu et al. (2021) derive the transferability of a well-designed GNN based on the differences of  $k$ -hop ego-graph Laplacian across graphs.

Spatial GNNs generally follow a recursive message-passing scheme (Gilmer et al. 2017), where each node updates its feature vector by aggregating the message from its local neighborhood. As discussed in (S. Liu, T. Li, et al. 2023; J. Wu, J. He, and E. A. Ainsworth 2023), the marginal distribution shifts  $P(X, G)$  between source and target graph domains can be induced by graph structure and individual node attributes (see Fig. 4(a)-(c)). Notably, three frameworks have been developed to enhance the transferability of spatial GNNs: invariant node representation (H. Wang et al. 2024; J. Wu, J. He, and E. A. Ainsworth 2023; Y. You et al. 2023), structure reweighting (S. Liu, T. Li, et al. 2023; S. Liu, D. Zou, et al. 2024), and graph Gaussian process (J. Wu, L. Ainsworth, et al. 2023).

- (1) **Invariant Node Representation:** Inspired by the domain adaptation theory (Redko, Morvant, et al. 2019), it is theoretically shown (J. Wu, J. He, and E. A. Ainsworth 2023; Y. You et al. 2023) that the target error can be bounded in terms of the source error and the graph domain discrepancy. The crucial idea of invariant node representation learning is to explicitly minimize the graph domain discrepancy in a latent feature space, thereby enhancing the transferability of spatial GNNs. For example, AdaGCN (Q. Dai et al. 2023) and UDA-GCN (M. Wu et al. 2020) leverage a domain discriminator to learn the domain-invariant node

representation learned by the output layer of GNNs. Inspired by the connection between spatial GNNs and Weisfeiler-Lehman graph kernels (Shervashidze et al. 2011; Weisfeiler and Lehman 1968), GRADE (J. Wu, J. He, and E. A. Ainsworth 2023) is proposed based on a graph subtree discrepancy measuring the subtree representation induced distribution shifts across graphs. More recently, SpecReg (Y. You et al. 2023) and A2GNN (M. Liu et al. 2024) further discuss the impact of spectral regularization and asymmetric model architectures on the transferability of GNNs, respectively.

- (2) **Structure Reweighting:** It is noticed (S. Liu, T. Li, et al. 2023) that invariant node representation might lead to sub-optimal solutions under conditional structure shifts. To solve this problem, StruRW (S. Liu, T. Li, et al. 2023) and Pair-Align (S. Liu, D. Zou, et al. 2024) are proposed to reweigh the edges of the source graph based on the label-oriented node connections of source and target graphs.
- (3) **Graph Gaussian Process:** Spatial GNNs are equivalent to graph Gaussian processes in the limit as the width of graph neural layers approaches infinity (Z. Niu et al. 2023). Based on this observation, GraphGP (J. Wu, L. Ainsworth, et al. 2023) is derived from a graph structure-aware neural network in the limit on the layer width, in order to characterize the relationships between nodes across different graph domains. The generalization analysis of GraphGP further reveals the positive correlation between knowledge transferability and graph domain similarity.

**3.2.2 Transferability on Textual Data.** Transfer learning has been widely studied in various natural language processing (NLP) tasks, e.g., text classification (Howard and Ruder 2018), question answering (Wiese et al. 2017), neural machine translation (H. Zhao, J. Hu, et al. 2020), etc. A key challenge in understanding textual transferability is the non-IID nature of words/tokens, as they might co-occur within sequences or documents. Thus, recent theoretical analyses of textual transferability often consider an alternative assumption (Lotfi et al. 2024), i.e., sequences or documents are independently drawn from the same distribution. This assumption enables theoretically deriving the transferability and generalization of transfer learning in sequence-level and document-level NLP tasks. Take multilingual machine translation as an example, the goal is to train a single neural machine translation model to translate between multiple source and target languages (Zoph and Knight 2016). To achieve this, language-invariant representation learning has been introduced to align the sentence distributions of different languages within a shared latent space (Arivazhagan et al. 2019). Nevertheless, H. Zhao, J. Hu, et al. (2020) theoretically analyze the fundamental limits of language-invariant representation learning by deriving a lower bound (w.r.t. marginal sentence distributions from different languages) on the translation error in the many-to-many language translation setting.

More recently, large language models (LLMs) have revolutionized the field of NLP (Brown et al. 2020; OpenAI 2023; Raffel et al. 2020; Touvron et al. 2023). The transferability of LLMs has been studied, as fine-tuning LLMs on downstream tasks has become the *de facto* learning paradigm. However, it is computationally expensive and resource-intensive to fine-tune the entire LLM model weights with billions of parameters via gradient-based optimization (Devlin et al. 2019). To solve this problem, parameter-efficient fine-tuning (PEFT) has been investigated from the perspectives of model tuning (Houlsby et al. 2019; E. J. Hu et al. 2022; Zaken et al. 2022) and prompt tuning (Lester et al. 2021; X. L. Li and P. Liang 2021; T. Shin et al. 2020). The goal is to adapt LLMs to various downstream tasks by adjusting as few parameters as possible.

Model tuning based approaches explore model architectures or parameters of LLMs for parameter-efficient fine-tuning (Aghajanyan et al. 2021). A variety of parameter-efficient model tuning frameworks have been proposed, including adapters (Houlsby et al. 2019), low-rank decomposition (E. J. Hu et al. 2022; Mahabadi et al. 2021), and selective masking (D. Guo et al. 2021; Zaken et al. 2022). The key ideas behind these frameworks are illustrated in Fig. 5. In general, these approaches aim to update only a few parameters by inserting new trainable modules, adding low-rank parameter matrices, or modifying specific parameters (e.g., bias terms).

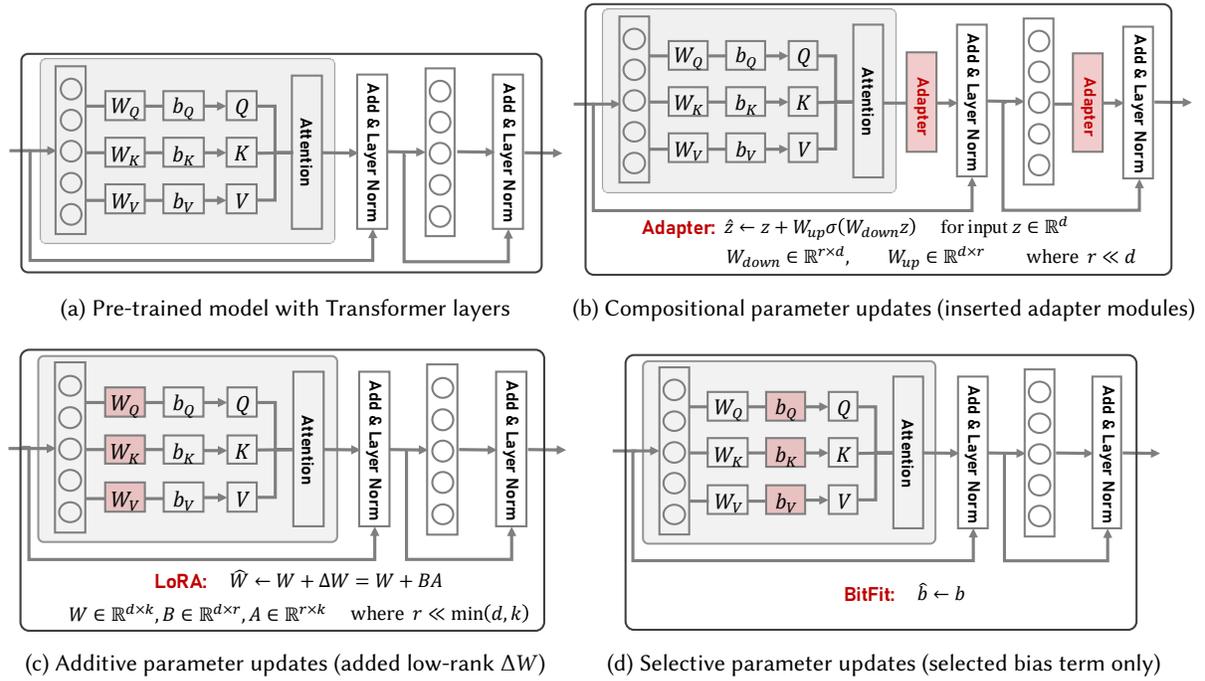


Fig. 5. Illustration of parameter-efficient fine-tuning with (b) Adapters (Houlsby et al. 2019) where new modules are inserted, (c) LoRA (E. J. Hu et al. 2022) where low-rank parameter matrices are added, and (d) BitFit (Zaken et al. 2022) where only bias terms will be updated.

- (1) **Adapters:** Adapter-based approaches add new learnable modules with a small number of parameters to LLMs, e.g., maximizing the likelihood  $p(y|x; \theta_{\text{adapter}} \circ \theta_{\text{LLM}})$  with added modules  $\theta_{\text{adapter}}$ . Initially, inspired by visual adapter modules (Rebuffi et al. 2017), Houlsby et al. (2019) study the adapter-based fine-tuning mechanism in NLP tasks. This method inserts two adapters sequentially within each Transformer block (Vaswani et al. 2017): one following the self-attention layer and another after the feed-forward layer. Nevertheless, follow-up research (Bapna and Firat 2019; Pfeiffer et al. 2021) demonstrates that inserting a single adapter after the feed-forward layer can achieve competitive performance while adding fewer parameters. Furthermore, by highlighting the connections between (model-based) adapters and (prompt-based) Prefix Tuning (X. L. Li and P. Liang 2021), J. He et al. (2022) introduce a family of parallel adapters that directly condition the adapters at different Transformer layers on the input text. On top of adapters, AdapterDrop (Rücklé et al. 2021) and CoDA (T. Lei et al. 2023) improve both fine-tuning and inference efficiency by removing adapters from lower Transformer layers and querying only a small subset of input tokens against the pre-trained LLMs, respectively.
- (2) **Low-rank Decomposition:** Low-rank decomposition injects trainable low-rank decomposition matrices into pre-trained model parameters without changing the model architectures, e.g., maximizing the likelihood  $p(y|x; \theta_{\text{LLM}} + \Delta\theta_{\text{LLM}})$  with low-rank  $\Delta\theta_{\text{LLM}}$ . This line of research is motivated by the phenomenon (Aghajanyan et al. 2021; C. Li et al. 2018) that pre-trained models tend to have a low intrinsic dimension. Here, the intrinsic dimension indicates the lowest dimensional parameter subspace in which satisfactory fine-tuned accuracy on downstream tasks can be achieved. Moreover, by assuming that the parameter change in LLMs during fine-tuning also has a low "intrinsic rank", LoRA (E. J. Hu et al. 2022) is introduced by optimizing low-rank

decomposition matrices of the parameter change. Empirically, LoRA has been further improved in various aspects, including rank selection/optimization (N. Ding, Lv, et al. 2023; Valipour et al. 2023; Q. Zhang et al. 2023), advanced optimizer (Hayou et al. 2024; F. Zhang and Pilanci 2024), etc. Theoretically, the expressiveness and generalization of LoRA have also been analyzed. Notably, Zeng and K. Lee (2024) prove that under mild conditions regarding the rank of LoRA, it can adapt a pre-trained (or randomly initialized) model to approximate any target model of equal or smaller size. Malladi et al. (2023) find that LoRA fine-tuning is approximately equivalent to full fine-tuning in the Neural Tangent Kernel (NTK) regime (Jacot et al. 2018), if  $r \geq \Theta(\log n^t / \epsilon^2)$  where  $r$  is the rank of LoRA and  $\epsilon$  is an approximation tolerance. (Jang et al. 2024) show that LoRA fine-tuning has no spurious local minima in the NTK regime, if  $r(r + 1) > 2Kn^t$  where  $K$  is the output dimension, and  $n^t$  is the number of training samples in the downstream target task. Furthermore, the generalization bounds of LoRA fine-tuning are theoretically derived in recent works (Jang et al. 2024; J. Zhu et al. 2024).

- (3) **Selective Masking:** The crucial idea of selective masking is to update only a small subset of model parameters during fine-tuning, e.g., maximizing the likelihood  $p(y|x; \theta_{\text{LLM}} + \Delta\theta_{\text{LLM}})$  with extremely sparse  $\Delta\theta_{\text{LLM}}$ . Intuitively, it aims to find a binary mask that automatically selects a small subset of parameters for fine-tuning (M. Zhao et al. 2020). There are three main frameworks for learning this mask. The first one is random masking (J. Xu and J. Zhang 2024; R. Xu et al. 2021), where all the elements of a mask are sampled independently from a Bernoulli distribution. J. Xu and J. Zhang (2024) demonstrate the effectiveness of random masking under a larger-than-expected learning rate, by theoretically building the connection between random masking and flat loss landscape. The second approach involves heuristically-motivated masks, such as bias terms (Zaken et al. 2022) and cross-attention layers (Gheini et al. 2021). The third approach optimizes masks over model parameters from various perspectives, e.g.,  $L_0$ -norm penalty (D. Guo et al. 2021), Fisher information (Das et al. 2023; Sung et al. 2021; R. Xu et al. 2021), Lottery Ticket Hypothesis (Ansell et al. 2022; Ploner and Akbik 2024), etc.

Unlike model tuning, prompt tuning keeps the pre-trained LLMs fixed and pretends a sequence of virtual token embeddings (referred to as a trainable prompt) to the input text, e.g., maximizing the likelihood  $p(y|[x, z]; \theta_{\text{LLM}})$  for a labeled sample  $(x, y)$  where  $\theta_{\text{LLM}}$  denotes LLM model parameters and  $z$  represents a prompt. Generally, there are three main frameworks for optimizing newly added prompts: soft (continuous) prompt tuning (Lester et al. 2021; X. L. Li and P. Liang 2021), hard (discrete) prompt tuning (T. Shin et al. 2020), and transferable prompt tuning (Vu et al. 2022).

- (1) **Soft Prompt Tuning:** The key idea of soft prompt tuning is to represent the virtual prompt as continuous-valued token embeddings. It will update only the continuous-valued embeddings of prompts  $z$ , either in the input embedding layer (Lester et al. 2021; Razdaibiedina et al. 2023) or in different layers of the LLMs (Hambarzumyan et al. 2021; X. L. Li and P. Liang 2021; X. Liu et al. 2022; G. Qin and Eisner 2021). Theoretically, C. Wei, Xie, et al. (2021) studies the connection between prompt tuning and downstream tasks using an underlying latent variable generative model of text. By assuming that input texts are generated by a Hidden Markov Model (HMM), this work models the downstream task as a function of the posterior distribution of the latent variables. It is then shown that prompt tuning enhances the recovery of the ground-truth labeling function in the downstream classification task. Later, (Y. Wang, Chauhan, et al. 2023) further show that a carefully constructed pre-trained Transformer can leverage prompt tuning to approximate any sequence-to-sequence function in a Lipschitz function space. They also analyze the restricted expressiveness of prompt tuning compared to model fine-tuning (e.g., LoRA (E. J. Hu et al. 2022)).
- (2) **Hard Prompt Tuning:** Although soft prompts can be optimized via gradient-based optimization, Khashabi et al. (2022) reveal that the learned embeddings of soft prompts do not correspond to any human-readable tokens, thus lacking semantic interpretations. An alternative solution is hard prompt optimization (Prasad

et al. 2023; T. Shin et al. 2020), which aims to find human-readable prompts from a pre-defined vocabulary. Specifically, AutoPrompt (T. Shin et al. 2020) greedily selects the optimal token for each location in the prompt based on the gradient of the loss w.r.t. the embeddings over labeled training samples. However, the greedy search strategy can result in disfluent and unnatural prompts. To solve this problem, FluentPrompt (W. Shi et al. 2023) and PEZ (Y. Wen et al. 2023) utilize projected gradient descent optimization to update all the tokens in the prompt. The crucial idea is to project the learned continuous-valued embeddings to their nearest neighbors in a pre-defined discrete token space, and then use the mapped tokens to calculate the gradient of the loss. Besides, Choi et al. (2024) and M. Deng et al. (2022) employ gradient-free reinforcement learning based optimization to discover discrete prompts, especially when LLMs are accessible only via APIs (i.e., model gradients and weights are not accessible).

- (3) **Transferable Prompt Tuning:** Recent studies have also investigated the transferability of prompts (Vu et al. 2022), where soft prompts are first learned from one or more source tasks and then used as the prompt initialization for the target task. This is motivated by the findings (Y. Gu et al. 2022; Vu et al. 2022) that a good prompt initialization is crucial for prompt tuning to achieve competitive performance on the target task, compared to model tuning, especially when the model sizes of LLMs are small. Follow-up research (Su et al. 2022) further analyzes the correlation between soft prompt transferability and the overlapping rate of activated neurons. Inspired by multi-task learning (Misra et al. 2016), MPT (Z. Wang, Panda, et al. 2023) decomposes the soft prompts for source tasks into a shared matrix and low-rank task-specific matrices, and then transfers the shared matrix to the target tasks. Additionally, studies (Su et al. 2022; Z. Wu et al. 2024) have explored the transferability of soft prompts across different language models in zero/few-shot learning settings.

**3.2.3 Transferability on Time Series Data.** A time series is a sequence of observations collected at even intervals of time and ordered chronologically (Chatfield 2004). Time series has been extensively applied to model non-stationary data in various high-impact domains, such as weather monitoring (W. Fan et al. 2023), financial forecasting (D. Zhou et al. 2020), and healthcare (Ragab et al. 2023). The key challenge in time series analysis lies in characterizing the temporal dependencies and non-stationary (i.e., rapidly changing data distribution over time) of time series data. Generally, time series transfer learning involves the following two tasks: time series forecasting (Y. Liu, H. Wu, et al. 2022; Passalis et al. 2020) and classification (Purushotham et al. 2017).

*Definition 3.5 (Time Series Transferability for Forecasting (Passalis et al. 2020)).* Given a target time-series data set with historical observations, time series transferability for forecasting aims to predict future events by utilizing its own historical observations or relevant knowledge from another source domain under temporal distribution shifts.

Time series forecasting leverages historical observations to predict future events. As illustrated in Fig. 6(a), there are two types of distribution shifts within time series forecasting. One is the sample-level temporal distribution shift (Kim et al. 2022; Y. Liu, H. Wu, et al. 2022) of non-stationary time series where the data distribution of time series samples changes over time. The other one is the domain-level distribution shifts that occur between source and target time series domains (X. Jin et al. 2022). To address the first type of distribution shifts, AdaRNN (Y. Du et al. 2021) characterizes the distribution information by splitting the training sequences into diverse periods with the largest distribution gap, and then dynamically reduces the distribution discrepancy across these identified periods. RevIN (Kim et al. 2022) is a symmetrical normalization-and-denormalization method using instance normalization (Ulyanov et al. 2016). It first normalizes the input sequences to mitigate distribution shifts among input sequences and then denormalizes the model outputs to restore the statistical information of input sequences. Follow-up approaches such as SAN (Z. Liu, Cheng, et al. 2023) and Dish-TS (W. Fan et al. 2023) build upon RevIN to further address temporal distribution shifts between input and horizon sequences by adaptively learning

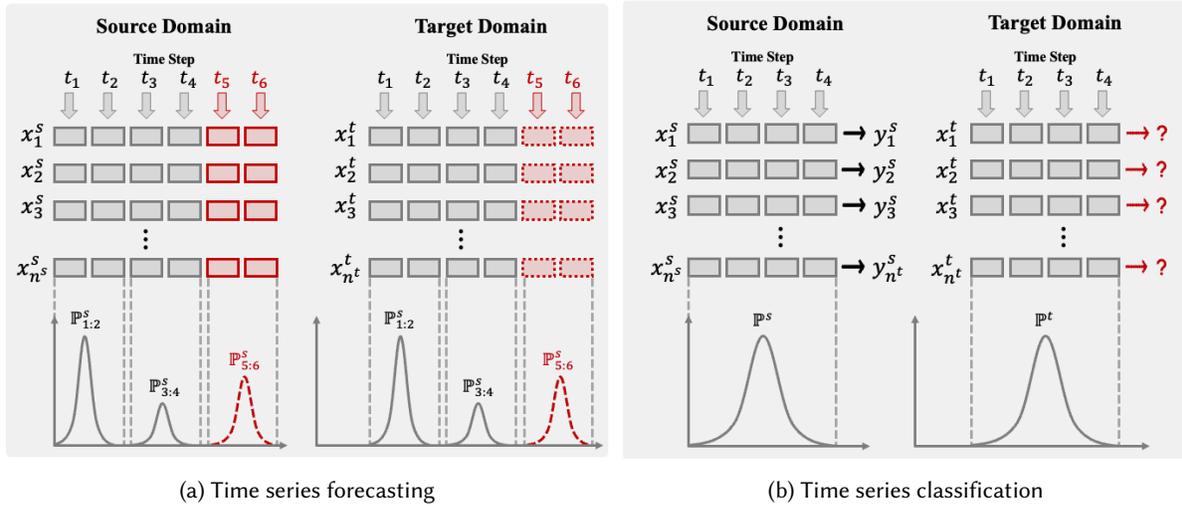


Fig. 6. Illustration of time series analysis under distribution shifts: (a) time series forecasting, and (b) time series classification

normalization coefficients for fine-grained temporal slices. More recently, based on Koopman theory (Brunton et al. 2022; Koopman 1931), KNF (R. Wang et al. 2023) and Koopa (Y. Liu, C. Li, et al. 2023) exploit linear Koopman operators to model the nonlinear dynamics of time series data on the measurement function space. Both methods design a global Koopman operator to learn time-invariant characteristics and a local Koopman operator to capture time-variant dynamics. To address the second type of distribution shifts, DAF (X. Jin et al. 2022) uses attention modules to learn complex temporal patterns within time series data and enforces the time-dependent query-key distribution alignment. Particularly, the queries and keys of attention modules are assumed to be domain-invariant, while the values capture domain-specific information for learning domain-dependent time series forecasters.

*Definition 3.6 (Time Series Transferability for Classification (Purushotham et al. 2017)).* Given a source domain with labeled time series samples and a target domain with limited or no label information, time series transferability for classification aims to improve the prediction performance of the time series classification model in the target domain by leveraging knowledge from the source domain.

Time series classification focuses on identifying time series data as a specific category (shown in Fig. 6(b)). There are two main transfer learning frameworks for time series classification. The first framework involves pre-training a model on a source domain and then fine-tuning it on a target domain. For example, pre-training techniques for time series modeling have been developed using convolutional neural networks (Fawaz et al. 2018; Kashiparekh et al. 2019), recurrent neural networks (Malhotra et al. 2017), ResNet (J. Dong et al. 2023; X. Zhang, Z. Zhao, et al. 2022), and Transformers (Zerveas et al. 2021). The second framework is to learn domain-invariant time series representations using adversarial learning (J. Lu and S. Sun 2024; Özyurt et al. 2023; Purushotham et al. 2017; Wilson et al. 2023, 2020) or statistical divergence metrics (R. Cai et al. 2021; H. He et al. 2023; Q. Liu and Xue 2021; Ott et al. 2022) using both source and target data. Specifically, VRADA (Purushotham et al. 2017) captures the domain-invariant temporal latent dependencies of multivariate time-series data using variational recurrent neural networks (Chung et al. 2015), followed by a gradient reversal layer (Ganin et al. 2016). Similarly, CoDATS (Wilson et al. 2020) is developed based on 1D convolutional neural networks and a gradient reversal layer. CLUDA (Özyurt et al. 2023) and CALDA (Wilson et al. 2023) further leverage contrastive learning losses to enhance the time series representations. In addition, AdvSKM (Q. Liu and Xue 2021) designs a hybrid spectral

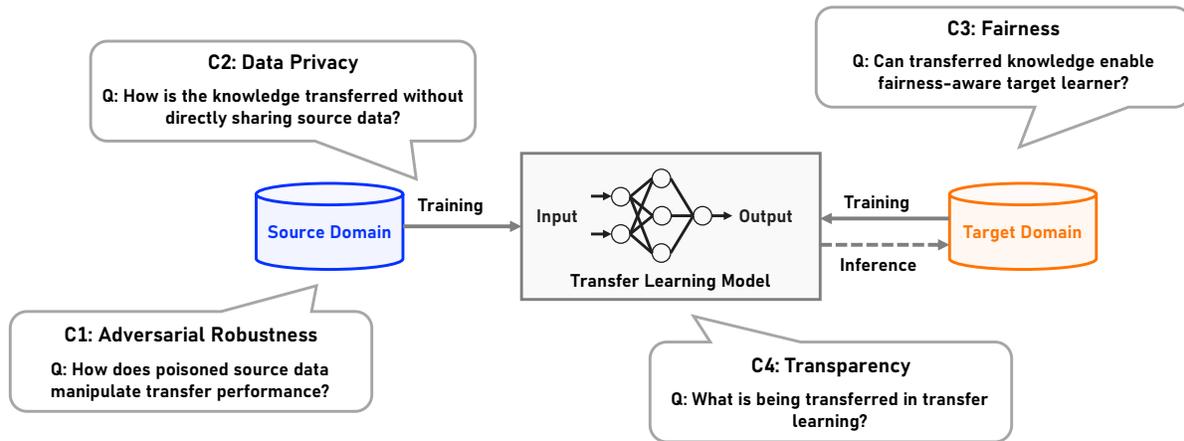


Fig. 7. Illustration of trustworthiness concerns in the knowledge transfer process

kernel network based on maximum mean discrepancy (Gretton et al. 2012) to align source and target time series representations. Assuming that source and target time series domains share the same causal structure, SASA (R. Cai et al. 2021) uses long short-term memory (LSTM) networks to learn sparse associative structures from both domains, and then aligns them via maximum mean discrepancy (Gretton et al. 2012). More recently, CauDiTS (J. Lu and S. Sun 2024) furthermore employs an adaptive causal rationale disentanglement to learn domain-invariant causal rationales and domain-specific correlations from variable interrelationships. RAINCOAT (H. He et al. 2023) is proposed to address open-world adaptation scenarios, where source and target domains might have domain-specific private classes. It extracts both time features via a 1-dimensional convolutional neural network and frequency features via the discrete Fourier transform, and then aligns the time-frequency features across domains using Sinkhorn divergence (Cuturi 2013).

## 4 Knowledge Trustworthiness

In this section, we review the knowledge trustworthiness of transfer learning. Compared to standard trustworthy machine learning over a single domain, this survey discusses whether the source and target users can trust the transferred knowledge, whereas trustworthy machine learning investigates how a user can trust a model trained on private data. As illustrated in Fig. 7, in the context of transfer learning, both the source domain owner and the target domain owner may have trustworthiness concerns about the transfer learning techniques. When considering the owners of the source domain as the “trustor”, do they trust that the transferred knowledge will not leak their data privacy (*C1: Privacy*)? Conversely, if the “trustor” indicates the owners of the target domain, do they trust that the transferred knowledge is not poisoned (*C2: Adversarial Robustness*) or biased (*C3: Fairness*), and how well can the transferred knowledge be explained (*C4: Transparency*)?

### 4.1 Privacy

Privacy protection aims to prevent the unauthorized access or misuse of data that can directly or indirectly reveal sensitive private information, e.g., age, gender, login credential, fingerprint, medical records, etc. In recent years, privacy concerns in understanding the trustworthiness of artificial intelligence (AI) systems have been emphasized in released AI ethics guidelines (Commission et al. 2019; Jobin et al. 2019) and legal laws (e.g., General Data Protection Regulation (GDPR) (Goodman and Flaxman 2017) and California Consumer Privacy Act (CCPA) (Harding et al. 2019)). Maintaining privacy is critical in privacy-sensitive applications, such as patient

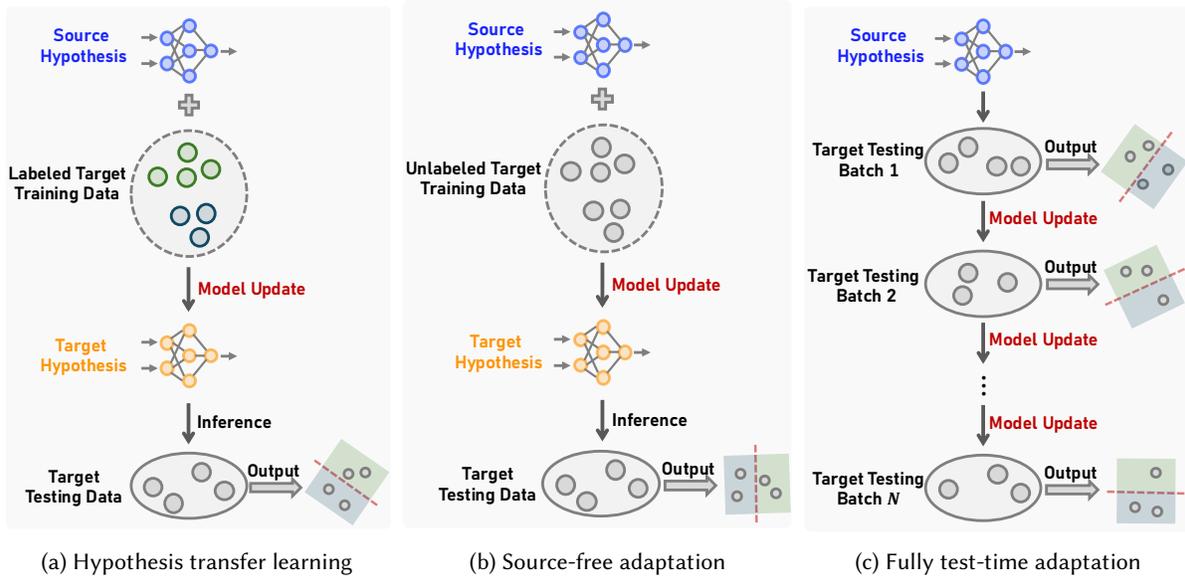


Fig. 8. Illustration of source hypothesis transfer, including (a) hypothesis transfer learning with labeled target training data, (b) source-free adaptation with unlabeled target training data, and (c) test-time adaptation with only target testing data.

clinical data analytics (Dayan et al. 2021) and mobile keyboard prediction (Hard et al. 2018). Particularly, privacy protection in transfer learning frameworks focuses on preventing the leakage of private source data during the knowledge transfer process. This concern has inspired privacy-preserving transfer learning frameworks designed to transfer knowledge from a private source domain to a specific target domain while ensuring data privacy. One key principle of these frameworks is that all source data remains stored locally, with only the updated source models/hypotheses being shared securely.

**4.1.1 Hypothesis Transfer.** Hypothesis transfer involves leveraging the source hypothesis pre-trained from the source data set to solve a learning task on the target domain. It assumes that the target learner has no access to the raw source data or the relatedness between the source and target domains, thereby protecting the data privacy of the source domain. Formally, given a source hypothesis, the problem of hypothesis transfer can be defined as follows.

*Definition 4.1 (Hypothesis Transfer (Kuzborskij and Orabona 2013)).* Given a source hypothesis  $f_S \in \mathcal{F}_S$  and a target data set  $D_T$  with  $n_T$  samples, hypothesis transfer algorithms aim to map the source hypothesis  $f_S \in \mathcal{F}_S$  and  $D_T$  onto a target hypothesis  $f_T \in \mathcal{F}_T$ :

$$A^{\text{htl}} : (\mathcal{X} \times \mathcal{Y})^{n_T} \times \mathcal{F}_S \rightarrow \mathcal{F}_T \quad (16)$$

where  $\mathcal{F}_S$  and  $\mathcal{F}_T$  denote the hypothesis spaces of source and target domains, respectively.

As illustrated in Fig. 8, there are three major learning scenarios for hypothesis transferability: (1) hypothesis transfer learning (Kuzborskij and Orabona 2013) with *labeled target training data*, (2) source-free adaptation (J. Liang et al. 2020) with *unlabeled target training data*, and (3) test-time adaptation (D. Wang et al. 2021) with *only target testing data*.

(1) **Hypothesis Transfer Learning:** The goal of hypothesis transfer learning is to optimize the learning function on the target domain using the basis of hypotheses from the source domain (Kuzborskij and Orabona 2013).

It assumes that both the source hypothesis and a few labeled target samples are accessible during the training of the target model. Earlier works (Fei-Fei et al. 2006; X. Li and Bilmes 2007) utilized the source hypothesis as prior knowledge to guide the target learner in Bayesian learning frameworks. Kuzborskij and Orabona (2013) theoretically analyze the generalization error (instantiated with the leave-one-out error) of regularized empirical risk minimization algorithms for hypothesis transfer learning. It is shown that the generalization error is positively correlated with a quantity  $\mathbb{E}_{(x,y) \sim \mathcal{P}_T} [\ell(f_S(x), y)]$  measuring how the source hypothesis performs on the target domain, and hypothesis transfer learning enjoys faster convergence rates of generalization errors when a good source hypothesis is provided. Later, Kuzborskij and Orabona (2017) extend generalization error bounds of regularized empirical risk minimization with (i) any non-negative smooth loss function, (ii) any strongly convex regularizer, and (iii) a combination of multiple source hypotheses. They further highlight the impact of the quantity  $\mathbb{E}_{(x,y) \sim \mathcal{P}_T} [\ell(f_S(x), y)]$  on the transfer performance and propose a principled approach to optimizing the combination of source hypotheses. Instead, S. S. Du, Koushik, et al. (2017) introduce a notion of transformation function to characterize the relatedness between the source and the target domains. Using this transformation function, they establish excess risk bounds for Kernel Smoothing and Kernel Ridge Regression. Minami et al. (2023) further theoretically derive the optimal form of transformation functions under the squared loss scenario. Aghbalou and Staerman (2023) analyze hypothesis transfer learning through regularized empirical risk minimization under reproducing kernel Hilbert space (RKHS) with surrogate classification losses (e.g., exponential loss, logistical loss, softplus loss, mean squared error, and squared hinge) in the context of binary classification. They establish the generalization error bounds and excess risk bounds based on hypothesis stability and pointwise hypothesis stability, highlighting the connections between surrogate classification losses and the quality of the source hypothesis. In addition, Chi et al. (2021) and R. Dong et al. (2023) study a more challenging few-shot hypothesis adaptation problem, where only a few labeled target samples (e.g., one sample per class) are available. Motivated by the learnability of semi-supervised learning, they propose generating highly-compatible unlabeled data to improve the training of the target learner.

It is noteworthy that gradient-based fine-tuning has become the predominant hypothesis transfer approach in the era of large foundation models (Gouk et al. 2021; Yosinski et al. 2014). Given a source hypothesis pre-trained on a source domain with adequate labeled samples, gradient-based fine-tuning aims to update the hypothesis through gradient descent optimization using a small amount of labeled target samples. The generalization performance of fine-tuning has been theoretically studied recently (Ju et al. 2022; Shachaf et al. 2021). Shachaf et al. (2021) show that the generalization error of fine-tuning under certain architectures (e.g., deep linear networks (Z. Ji and Telgarsky 2019), shallow ReLU networks (S. Arora et al. 2019)) can be affected by the difference between optimal (normalized) source and target hypotheses, the covariance structure of the target data, and the depth of the network. Furthermore, recent works (Gouk et al. 2021; Ju et al. 2022; D. Li and H. R. Zhang 2021) show the generalization error of fine-tuning techniques in terms of the distance between fine-tuned and initialized model parameters.

- (2) **Source-free Adaptation:** In contrast to hypothesis transfer learning, source-free adaptation (Kundu, Venkat, et al. 2020; J. Liang et al. 2020) enables hypothesis transfer from the source to the target domains using only unlabeled target data. To solve this problem, SHOT (J. Liang et al. 2020) is proposed to update the feature extractor of the pre-trained source model for the target domain while keeping the source classifier fixed. It maximizes the mutual information between intermediate feature representations and the output of the classifier, and also minimizes the prediction error using self-supervised pseudo labels. The follow-up works have developed source-free adaptation frameworks from various perspectives: clustering (S. Yang et al. 2022), pseudo-labeling (Boudiaf et al. 2023; J. Lee et al. 2022), data augmentation (Hwang et al. 2024; Kundu, Kulkarni, et al. 2022), etc. As discussed in (Mitsuzumi et al. 2024), most existing source-free adaptation approaches focus on understanding the discriminability-diversity trade-off: the former improves the discriminability of

unlabeled target samples in the latent feature space while the latter ensures prediction diversity for all classes. In particular, Mitsuzumi et al. (2024) establish a theoretical connection between source-free adaptation and self-training (C. Wei, K. Shen, et al. 2021) in terms of discriminability and diversity losses. This connection enables improved training of source-free adaptation by incorporating an auto-adjusting diversity constraint and teacher-student augmentation learning. In contrast, Han et al. (2023) and Kundu, Kulkarni, et al. (2022) study the discriminability-transferability trade-off in the context of source-free adaptation. Theoretically, M. Shen et al. (2023) derive an information-theoretic generalization error bound for multi-source-free adaptation based on a bias-variance trade-off. Here, bias is triggered by the label and feature misalignments across domains, and variance is triggered by the number of pseudo-labeled target samples. Yi et al. (2023) establish the connections between source-free adaptation and learning with noisy labels, given the findings that the pseudo-labels of target samples generated by the source model can be noisy due to domain shift. They theoretically justify the existence of the early-time training phenomenon (ETP) in source-free adaptation scenarios and propose using early learning regularization (S. Liu, Niles-Weed, et al. 2020) to prevent the model from memorizing label noise during training. Empirically, in addition to standard vision tasks, Boudiaf et al. (2023) reevaluate existing source-free adaptation methods in a more challenging set of naturally occurring distribution shifts in bioacoustics. Their findings indicate that these existing methods often lack generalizability and perform worse than no adaptation in some cases. This highlights the necessity of evaluating source-free adaptation methods across a range of tasks, data modalities, and degrees of distribution shifts.

- (3) **Test-time Adaptation:** Fully test-time adaptation (D. Wang et al. 2021) aims to adapt the source hypothesis to the target testing data, where data batches arrive sequentially and each batch can only be observed once. To this end, Test-Time Training (TTT) (Y. Sun et al. 2020) and its modified version (TTT++) (Y. Liu, Kothari, et al. 2021) incrementally update the feature extractor by minimizing the auxiliary task loss. Notably, this approach requires optimizing both this auxiliary task loss and the standard supervised loss. In contrast, without changing the training phase, Tent (D. Wang et al. 2021) is proposed to minimize the entropy of model predictions by only updating the normalization statistics and channel-wise affine transformations in an online manner. The follow-up works further improve this framework from two perspectives. One is to enhance the stability and robustness of test-time adaptation by minimizing the entropy of the average prediction across different augmentations (M. Zhang et al. 2022) or by minimizing sharpness-aware entropy (Gong et al. 2023; S. Niu, J. Wu, Y. Zhang, Z. Wen, et al. 2023). The other addresses catastrophic forgetting by regularizing the updated parameters (S. Niu, J. Wu, Y. Zhang, Y. Chen, et al. 2022; Q. Wang et al. 2022) or adaptive resetting the model parameters (Niloy et al. 2024). Furthermore, Goyal et al. (2022) analyze test-time adaptation through the lens of convex conjugate loss functions and proposes a principled self-training approach based on conjugate pseudo labels for test-time adaptation. Later, J.-K. Wang and Wibisono (2023) theoretically justify the advantages of conjugate labels over hard labels in test-time adaptation by showing the performance gap between gradient descent with conjugate labels and gradient descent with hard labels in a binary classification problem. Empirically, H. Zhao, Y. Liu, et al. (2023) identify the commonly seen pitfalls when evaluating test-time adaptation algorithms, including sensitive hyperparameter selection, inconsistent source hypothesis, and insufficient consideration of various types of distribution shifts. W. Bao, T. Wei, et al. (2023) further demonstrate that the modules (e.g., batch normalization layers (D. Wang et al. 2021), feature extractor layers (Y. Sun et al. 2020), classifier layers (Iwasawa and Matsuo 2021)) selected for test-time adaptation are strongly correlated with the types of distribution shifts.

**4.1.2 Federated Transfer.** In contrast to the unidirectional hypothesis transfer discussed in the previous subsection, federated transfer emphasizes bidirectional knowledge sharing that allows source and target domains to communicate and exchange information while maintaining privacy protection (B. McMahan et al. 2017). This is largely inspired by recent personalized federated learning frameworks (Kairouz et al. 2021; Y. Liu, Y. Kang,

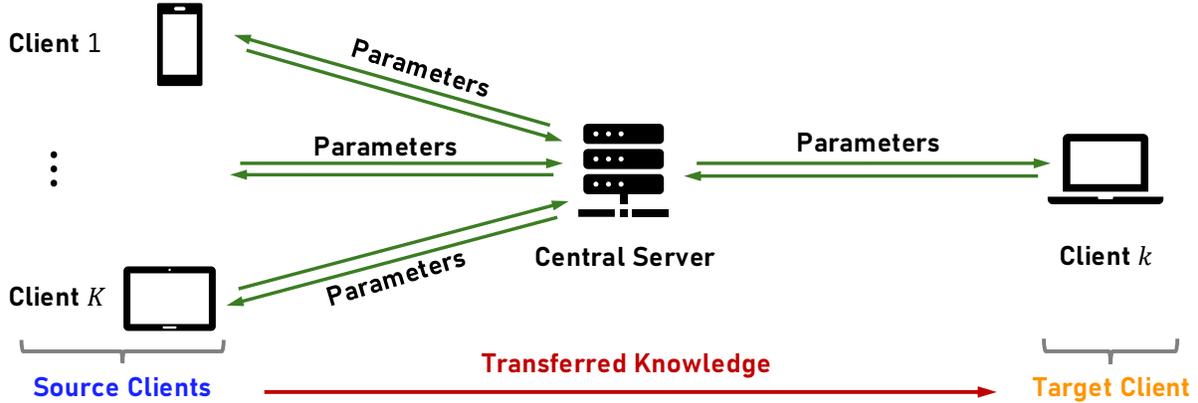


Fig. 9. Illustration of personalized federated learning in two scenarios: (1) Generalization performance is evaluated across all clients, where each client (e.g., client  $k$  ( $k = 1, 2, \dots$ )) is considered as a target client and others as source clients for knowledge transfer. (2) Generalization performance is improved only on a specific target client (i.e., only client  $k$ ).

et al. 2020) which allow private clients to collaborate in training personalized models under the coordination of a central server. As illustrated in (Kairouz et al. 2021; B. McMahan et al. 2017), during each communication round, private clients upload their model updates to the central server, which then securely aggregates these updates and broadcasts the updated model back to each client. In this process, each client exclusively owns their data, which will not be shared with the central server or with other clients. From the perspective of knowledge transferability, the intuition behind personalized federated learning is to transfer knowledge across private clients in a privacy-preserving manner (Y. Chen, X. Qin, et al. 2020; Y. Liu, Y. Kang, et al. 2020; J. Wu, W. Bao, et al. 2023). In other words, each private (target) client participates in federated collaboration to receive knowledge from other (source) clients, and its uploaded parameters can be used as indicators to select only the most relevant source knowledge (e.g., related clients with similar parameters within a coalition (W. Bao, H. Wang, et al. 2023; Donahue and Kleinberg 2021)) under distribution shifts across clients. Formally, given a set of private clients, each with access to a private training set, the problem of federated transfer can be defined as follows.

*Definition 4.2 (Federated Transfer).* Given a central server and  $K$  private clients each with training samples  $D_k$  ( $k = 1, \dots, K$ ), federated transfer aims to learn a personalized model  $f_k \in \mathcal{F}_k$  on the  $k^{\text{th}}$  client ( $k = 1, \dots, K$ ) by leveraging useful knowledge from the other clients  $\{k' | k' \neq k\}$ .

$$A^{\text{fl}} : (\mathcal{X} \times \mathcal{Y})^{n_k} \times (\mathcal{F}_1 \times \dots \times \mathcal{F}_{k-1} \times \mathcal{F}_{k+1} \times \dots \times \mathcal{F}_K) \rightarrow \mathcal{F}_k \quad (17)$$

where  $\mathcal{F}_k$  denotes the hypothesis space of the  $k^{\text{th}}$  client. The hypotheses from the other clients are often aggregated at the central server and then transferred to the  $k^{\text{th}}$  client, i.e.,  $(\mathcal{F}_1 \times \dots \times \mathcal{F}_{k-1} \times \mathcal{F}_{k+1} \times \dots \times \mathcal{F}_K) \rightarrow \mathcal{F}_k$ .

Note that  $\mathcal{F}_1 = \dots = \mathcal{F}_K$  implies that all clients will share the same hypothesis space. It can be seen from Fig. 9 that there are two different scenarios: one where the generalization performance of all clients is important, and another where only the generalization performance of a target client matters.

To address the first scenario, various personalized federated learning frameworks have been proposed, including parameter decoupling (Collins et al. 2021), model interpolation (C. T. Dinh et al. 2020; T. Li et al. 2021), clustering (Ghosh et al. 2020; Sattler et al. 2020), multi-task learning (Smith et al. 2017), meta-learning (Fallah et al. 2020), knowledge distillation (J. Zhang et al. 2021; Z. Zhu, Hong, et al. 2021), Bayesian learning (Achituve et al. 2021; X. Zhang, Y. Li, et al. 2022), etc. Despite the impressive performance of these personalized federated learning frameworks across various applications, it is shown (W. Bao, H. Wang, et al. 2023; J. Wu, W. Bao, et al.

2023) that some clients might suffer from negative transfer in the context of personalized federated learning. It implies that their performance can be worse compared to when they train a model solely on their local data without communicating information with other clients. To mitigate negative transfer issues, INCFE (Cho et al. 2022) is proposed to maximize the incentivized client participation by dynamically adjusting the aggregation weight assigned to each client. FedCollab (W. Bao, H. Wang, et al. 2023) optimizes the collaboration structure by clustering clients into non-overlapping coalitions based on their distribution distances and data quantities. Similarly, FEDORA (J. Wu, W. Bao, et al. 2023) adaptively aggregates relevant source knowledge by considering distribution similarities among clients and regularizes local models when the received knowledge has a positive impact on the generalization performance. DisentAFL (J. Chen and A. Zhang 2024) uses a two-stage knowledge disentanglement and gating mechanism to enhance positive transfer under complex client heterogeneity, e.g., modality heterogeneity, task heterogeneity, and domain heterogeneity among clients.

To address the second scenario, federated domain adaptation (Z. Fan et al. 2023; E. Jiang et al. 2024; X. Peng, Z. Huang, et al. 2020) has been studied to transfer knowledge from multiple source clients with sufficient labeled samples to a target client with limited or no labeled samples. Unlike standard personalized federated learning, it focuses only on the generalization performance of the target clients. Specifically, inspired by domain adaptation theory (Ben-David, Blitzer, et al. 2010), X. Peng, Z. Huang, et al. (2020) derive a weighted error bound for federated domain adaptation. Based on this, the FADA algorithm is proposed to disentangle domain-invariant and domain-specific features for each client and then align the domain-invariant features between source and target clients. Similarly, Feng et al. (2021) leverage knowledge distillation and BatchNorm Maximum Mean Discrepancy (MMD) to address the distribution gaps between source and target clients. More recently, E. Jiang et al. (2024) theoretically analyze the connections between the generalization performance and aggregation rules of federated domain adaptation. This finding also results in an auto-weighting scheme for optimal combinations of the source and target gradients. In addition to federated domain adaptation, federated domain generalization aims to train models using source clients and then apply these models to previously unseen target clients (A. T. Nguyen, Torr, et al. 2022; H. Yuan et al. 2022). Notably, R. Bai et al. (2024) propose a federated domain generalization benchmark, highlighting the necessity of evaluation scenarios that involve a large number of private clients, high client heterogeneity, and more realistic data sets.

## 4.2 Adversarial Robustness

It has been observed (Goodfellow et al. 2015; C. Szegedy et al. 2014) that modern machine learning models can be easily fooled by adversarial examples that are perceptibly indistinguishable with respect to clean inputs. This survey focuses on exploring the adversarial robustness of knowledge transfer models under assumptions where distribution shifts occur across domains.

**4.2.1 Attacks.** Recent efforts have been devoted to understanding the adversarial vulnerability of deep transfer learning techniques (S. Rezaei and X. Liu 2020; B. Wang et al. 2018; Y. Zhang, Song, et al. 2020). In the context of transfer learning, evasion attacks aim to generate adversarial examples to fool the learned transfer learning models on the target domain. Initially, by minimizing the feature representation dissimilarity between adversarial and clean target samples from different classes using only the pre-trained model, B. Wang et al. (2018) demonstrate the vulnerability of fine-tuned models in the transfer learning framework. Based on the observations that the neurons of the activation vector within the pre-trained model correlate with target classes, S. Rezaei and X. Liu (2020) design a simple brute-force attacking mechanism. This approach crafts input data to trigger those neurons individually, thereby exploring which one is highly associated with each target class.

In contrast, as shown in Fig. 10, poisoning attacks allow crafting source samples to control the prediction behavior of transfer learning models on the target domain during model training. Generally, poisoning attacks can occur in two transfer learning scenarios. The first scenario is the joint training of source and target data, assuming

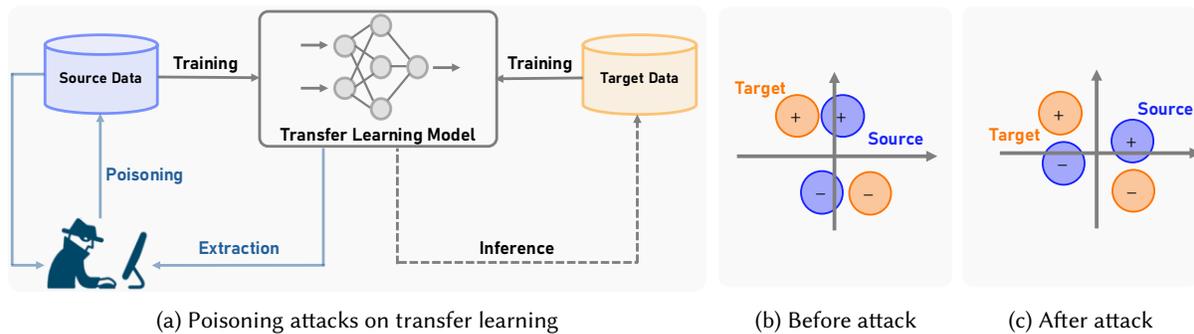


Fig. 10. Illustration of poisoning attacks on transfer learning. (a) By injecting adversarial noise into the source data, the adversary can control the prediction behavior of transfer learning models on the target domain. One intuitive explanation (Mehra et al. 2021; J. Wu and J. He 2021) is that (b) initially the source and target distributions can be correctly aligned, but (c) they become misaligned after applying the attack.

that source and target domains have the same labeling space, and the target domain has only unlabeled training samples (i.e., unsupervised domain adaptation (Pan and Q. Yang 2010)). It is motivated by the findings (H. Zhao, Combes, et al. 2019) that the feature-based marginal distribution matching can result in negative transfer when the target domain has no label information. Notably, I2Attack (J. Wu and J. He 2021) and AdaptAttack (J. Wu and J. He 2023a) maximize the label-informed joint distribution discrepancy between the raw source domain and the poisoned source domain with the following constraints. (1) *Perceptibly Unnoticeable*: All the poisoned input images are natural-looking. (2) *Adversarially Indirect*: Only source samples are maliciously manipulated. (3) *Algorithmically Invisible*: Neither source classification error nor marginal domain discrepancy between source and target domains increases. These constraints imply that in the context of transfer learning, an adversary could potentially manipulate the source data to gain control over the prediction function on the target domain. Similarly, Mehra et al. (2021) propose to generate poisoned source samples with clean labels or mislabeled source samples to fool the discrepancy-based adaptation approaches.

The second scenario is the pre-training and fine-tuning framework, where a model is first pre-trained on the source domain and then fine-tuned on the downstream target domain. In this scenario, backdoor attacks are intended to manipulate pre-trained model weights, thus resulting in malicious prediction behavior of fine-tuned models on the target domain (Y. Ji et al. 2018; L. Shen et al. 2021). The intuition behind backdoor attacks is that when the triggers (e.g., keywords) are activated on target samples, the fine-tuned model will predict predefined class labels (T. Gu et al. 2019). Specifically, backdoor attacks in pre-trained models satisfy the following conditions (Kurita et al. 2020; Y. Yao et al. 2019). (1) Only pre-trained model weights are manipulated, and the infection should be done on the target data through transfer learning. (2) With poisoned pre-trained models, fine-tuned models will behave normally on clean target data, but misclassify any sample with the trigger into a specific class. (3) The designed attacks should be unnoticeable from the viewpoint of the target learner, i.e., the attacker does not alter the fine-tuning process and the training data on the target domain. With these conditions in mind, BadNets (T. Gu et al. 2019) directly uses a poisoned data set to adjust the parameters of pre-trained models. RIPPLE (Kurita et al. 2020) poisons the weights of pre-trained models using a bi-level optimization objective over both the poisoning and fine-tuning losses. However, it is noticed (Kurita et al. 2020) that the fine-tuned models can mitigate the impact of backdoor attacks during the fine-tuning process on clean target samples. Thus, L. Li et al. (2021) and Y. Yao et al. (2019) focus on poisoning only the lower layers of pre-trained models. Furthermore, Y. Li, T. Li, et al. (2024) reformulate backdoor injection as a lightweight knowledge editing problem and adjusts only a subset of model parameters (e.g., key-value pairs) with a minimal amount of poisoned data.

Rather than manipulating pre-trained model weights, recent works explore backdoor attacks on large language models (LLMs) in the fine-tuning phase by inserting triggers into instructions (Shu et al. 2023; J. Xu, M. Ma, et al. 2024; Yan et al. 2024) or prompts (X. Cai et al. 2022; S. Jiang et al. 2023; S. Zhao et al. 2023). Besides, it is empirically observed (Bowen et al. 2025; A. Wan et al. 2023) that larger LLMs are more susceptible to poisoning attacks than smaller ones. All these backdoor attacks above highlight security and ethical concerns in developing and deploying pre-trained models (Hubinger et al. 2024).

**4.2.2 Defenses.** In the context of transfer learning, the adversarial robustness of the prediction function on the target domain can be improved in various scenarios. (1) Given an adversarially pre-trained source model, the adversarial robustness of the source model can be transferred to the target domain (Hendrycks et al. 2019; Shafahi et al. 2020). (2) Given a standard pre-trained model, the adversarial robustness of the target learner can be enhanced via robust fine-tuning (X. Xu, J. Zhang, et al. 2024). (3) Given an attacked pre-trained model, the defense mechanism can be developed to mitigate the negative impact of source knowledge on the target domain during fine-tuning (Chin et al. 2021; Xi et al. 2023).

Recent works (T. Chen et al. 2020; Davchev et al. 2019; Hendrycks et al. 2019) empirically demonstrate the transferability of adversarial robustness across domains, e.g., the robustness of an adversarially pre-trained source model can be transferred to the target domain. Specifically, based on the Learning without Forgetting (LoF) approach (Z. Li and Hoiem 2017), Shafahi et al. (2020) and Vaishnavi et al. (2022) use the distillation regularization to preserve the robust feature representations of the source model during fine-tuning. The intuition is that the lower layers of an adversarially pre-trained source model can capture robust features from input samples. Similarly, to enhance the transferability of adversarial robustness across domains, Awais et al. (2021) utilize knowledge distillation to preserve the feature correlations of the robust source model on the target domain. D. Chen et al. (2021) propose enforcing feature similarity between natural samples and their corresponding adversarial counterparts during pre-training and regularizing the Lipschitz constant of neural networks during fine-tuning. Z. Liu, Y. Xu, et al. (2023) propose a TWINS structure to incorporate the means and variances of batch normalization layers over both pre-training and target data during adversarial fine-tuning. Notably, the studies mentioned above focus on the transferability of empirical adversarial robustness through adversarial training techniques (Goodfellow et al. 2015; Madry et al. 2018), which minimize the adversarial objective against pre-determined strong attacks. In addition to empirical adversarial robustness, Alhamoud et al. (2023) and Vaishnavi et al. (2024) further investigate certified/provable adversarial robustness (Cohen et al. 2019; Jeong and J. Shin 2020) in the context of transfer learning, which seeks to maximize the radius around inputs within which the model output remains consistent. Theoretically, Nern et al. (2023) show that the transferability of adversarial robustness can be guaranteed if the feature extractor of the pre-trained source model is robust and only the newly added linear predictor is updated during fine-tuning. This analysis is consistent with the empirical observations (Hua et al. 2024; Shafahi et al. 2020) that feature extractors from an adversarially pre-trained source model contribute to robustness transfer across domains, where only the last layer is re-trained on the target data.

The second line of research is robust fine-tuning (X. Dong et al. 2021; X. Xu, J. Zhang, et al. 2024), where a standard pre-trained model is fine-tuned using adversarial training (Goodfellow et al. 2015; Madry et al. 2018). Specifically, RIFT (X. Dong et al. 2021) maximizes the mutual information between the feature extracted by the adversarially fine-tuned model and the class label plus the feature extracted by the pre-trained model. AutoLoRa (X. Xu, J. Zhang, et al. 2024) disentangles robust fine-tuning via a low-rank branch to mitigate gradient conflicts between adversarial and natural objectives. It optimizes adversarial objective w.r.t. the standard feature extractor and standard objective w.r.t. the auxiliary LoRA branch. More recently, Y. Wang and R. Arora (2024) examine the adversarial robustness of hypothesis transfer learning (Kuzborskij and Orabona 2013), which involves transferring knowledge from source domains to a target domain using a set of pre-trained auxiliary hypotheses. They derive

generalization error bounds for adversarial robustness in the target domain based on two specific algorithms: adversarial regularized empirical risk minimization and proximal stochastic adversarial training.

The previous two scenarios assume the availability of a clean pre-trained source model for transfer learning. Their goal is to improve the adversarial robustness of fine-tuned models against adversarial perturbations in the target samples during inference. As discussed in Subsection 4.2.1, a more challenging yet realistic scenario occurs when a poisoned source model (X. Cai et al. 2022; S. Rezaei and X. Liu 2020) is deployed for transfer learning. In this scenario, defense mechanisms should handle the negative impact of poisoned source knowledge during fine-tuning. To this end, Chin et al. (2021) design a defense mechanism to counter the attack proposed by S. Rezaei and X. Liu (2020). The key idea is to reduce the similarity between the pre-trained and the fine-tuned models via noisy feature distillation. More recently, in the context of backdoored large language models (LLMs) (X. Cai et al. 2022; S. Zhao et al. 2023), Xi et al. (2023) propose to detect poisoned target samples associated with triggers during inference by leveraging the different masking-sensitivity of poisoned and clean samples. The intuition is that poisoned samples are more sensitive to random masking than clean samples, as fine-tuned LLMs might exhibit significant changes in predictions when the trigger and normal content are masked within poisoned samples. Similarly, Qi et al. (2021) and W. Yang et al. (2021) detect poisoned samples using the perplexity changes of samples under word deletion or different robustness properties of clean and poisoned samples against triggers, respectively.

**4.2.3 Transferability vs. Robustness.** In addition to highlighting the adversarial vulnerability and robustness of transfer learning frameworks, recent studies have also explored the connection between knowledge transferability and adversarial robustness (Salman et al. 2020; Terzi et al. 2021). To be specific, it is empirically demonstrated (Salman et al. 2020) that adversarially robust models can transfer better (i.e., higher transfer accuracy in the target domain) than their standard-trained counterparts. That is, though robustness may be at odds with accuracy within the same domain (Tsipras et al. 2019), the adversarial robustness achieved in a source domain can improve the transfer accuracy in a related target domain. Utrera et al. (2021) further explain that in image classification tasks, adversarial training in the source domain biases the learned representations towards retaining shapes, thereby improving transferability in the target domain. Theoretically, Terzi et al. (2021) provide an information-theoretic justification for adversarial training, implying the trade-off between accuracy on the source domain and transferability on a related target domain. More rigorously, Z. Deng et al. (2021) demonstrate that, for a learning function based on a two-layer linear neural network, adversarially robust representation learning over multiple source domains leads to much tighter transfer error bounds on the target domain than standard representation learning. Alternatively, X. Xu, J. Y. Zhang, et al. (2022) show that adversarial training regularizes the function class of feature representation learning, thus improving knowledge transferability across domains.

### 4.3 Fairness

Fairness involves eliminating discrimination when training machine learning models (Castelnovo et al. 2022; Eshete 2021). In the legal domain, potential discrimination is defined as disparate treatment (triggered by intentionally treating an individual differently) and disparate impact (triggered by negatively affecting members of a protected group) (Pessach and Shmueli 2023). Motivated by this definition, different measures of algorithmic fairness have been proposed in machine learning communities, e.g., individual fairness (Dwork et al. 2012), group fairness (Feldman et al. 2015), etc. To be specific, individual fairness (Dwork et al. 2012) maintains that similar individuals should be treated similarly. Group fairness (Feldman et al. 2015; Hardt et al. 2016) ensures statistical parity among groups with sensitive attributes (e.g., race, gender, age). In the context of transfer learning, a fundamental concern is whether the fairness of a machine learning model can be transferred across domains

under distribution shifts. Following (Y. Chen, Raab, et al. 2022; Schumann et al. 2019), the problem of fairness transfer can be formulated as follows.

*Definition 4.3 (Fairness Transfer (Y. Chen, Raab, et al. 2022; Schumann et al. 2019)).* Given a source domain and a target domain, we denote the fairness violation measures as  $\Delta_S^*(\cdot)$  for the source domain and  $\Delta_T^*(\cdot)$  for the target domain. For any hypothesis  $f \in \mathcal{F}$ , algorithmic fairness can be transferred if the following condition is satisfied:

$$\Delta_T^*(f) \leq \Delta_S^*(f) + \delta \quad (18)$$

where  $\delta$  quantifies the distribution shifts between the source and target domains.

**4.3.1 Group Fairness.** Generally, group fairness (Castelnuovo et al. 2022; Feldman et al. 2015; Hardt et al. 2016) requires that different groups are treated equally. There are several commonly used group fairness metrics: demographic parity (Feldman et al. 2015), equality of opportunity (Hardt et al. 2016), and equalized odds (Hardt et al. 2016). Following (Madras et al. 2018; Schumann et al. 2019), we formally define these metrics in a binary classification problem where  $\mathcal{Y} = \{0, 1\}$ . Assuming there are two groups defined by binary sensitive attributes  $A \in \{0, 1\}$ , fair machine learning seeks to ensure accurate predictions without bias against any particular group.

- *Demographic Parity (Feldman et al. 2015):* Demographic parity, also known as statistical parity, requires the same positive prediction ratio across groups with different sensitive attributes.

$$\Pr(\hat{Y} = 1|A = 0) = \Pr(\hat{Y} = 1|A = 1) \quad (19)$$

where  $\hat{Y}$  denotes the random variable of the predicted class label and  $\Pr(\cdot)$  represents the probability. This criterion implies that the decisions made by machine learning models should be independent of any sensitive attributes. However, it may be limited in scenarios where the base rates of the two groups differ, i.e.,  $\Pr(Y = 1|A = 0) \neq \Pr(Y = 1|A = 1)$  where  $Y$  is the ground-truth class variable. In such cases, it is unrealistic to expect both model accuracy and demographic parity to be achieved simultaneously. Notably, H. Zhao and Gordon (2019) theoretically characterize the inherent trade-off between statistical parity and prediction accuracy, by providing a lower bound on group-wise prediction error for any fair predictor under demographic parity.

- *Equality of Opportunity (Hardt et al. 2016):* A machine learning model is considered fair under equality of opportunity if the false positive rates across groups are equal.

$$\Pr(\hat{Y} = 1|A = 0, Y = 0) = \Pr(\hat{Y} = 1|A = 1, Y = 0) \quad (20)$$

In contrast to demographic parity, equality of opportunity considers the ground-truth class variable  $Y$ . It enables the base rates for the two groups to be different. Similarly, a symmetric definition can be formulated using the false negative rates, i.e.,  $\Pr(\hat{Y} = 0|A = 0, Y = 1) = \Pr(\hat{Y} = 0|A = 1, Y = 1)$ .

- *Equalized Odds (Hardt et al. 2016):* A machine learning model is considered fair under equalized odds if both the false positive rates and false negative rates across groups are equal.

$$\begin{aligned} \Pr(\hat{Y} = 1|A = 0, Y = 0) &= \Pr(\hat{Y} = 1|A = 1, Y = 0) \\ \Pr(\hat{Y} = 0|A = 0, Y = 1) &= \Pr(\hat{Y} = 0|A = 1, Y = 1) \end{aligned} \quad (21)$$

Fair transfer learning that integrates the aforementioned group fairness criteria has been studied in recent years (Dutt et al. 2024; Giguere et al. 2022). For example, Madras et al. (2018) and Zemel et al. (2013) propose learning fair intermediate representations by encoding the data as accurately as possible while obscuring information about sensitive attributes. They demonstrate the transferability of these fair representations across different tasks. Later, Schrouff et al. (2022) empirically investigate the connections between compound distribution shifts (e.g., the co-occurrence of demographic, covariate, and label shifts) and fairness transfer in real-world medical applications via a joint causal framework (Mooij et al. 2020). Furthermore, Y. Chen, Raab, et al. (2022) provide a

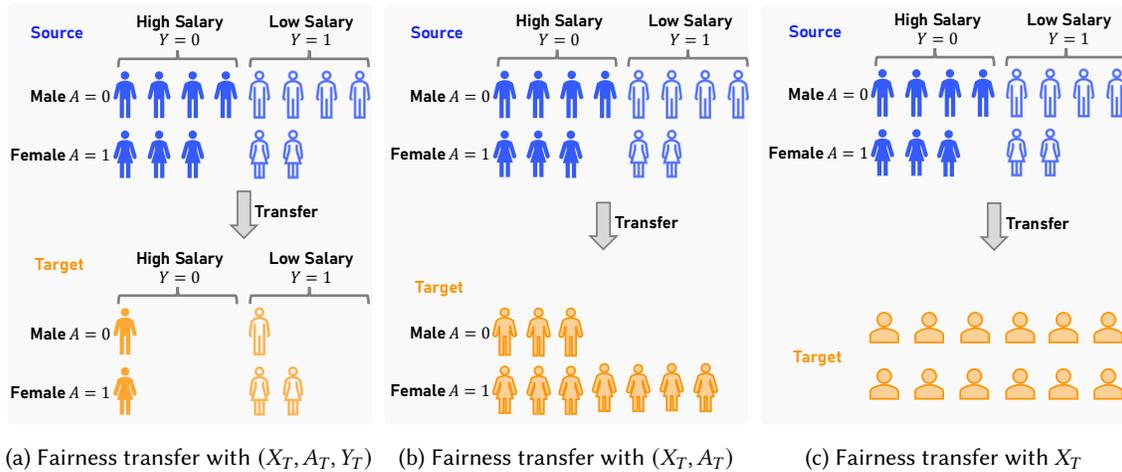


Fig. 11. Transferability of group fairness across domains. (a) The target domain has labeled samples with sensitive attributes. (b) The target domain has unlabeled samples with sensitive attributes. (c) The target domain has only unlabeled samples without sensitive attributes.

generic Lipschitz upper bound for group fairness when the underlying distribution shifts (e.g., covariate shift or label shift between source and target domains) are constrained. Specifically, most existing works (Biswas and S. Mukherjee 2021; Coston et al. 2019; C. Zhao, K. Jiang, et al. 2024) dive into understanding the transferability of fairness, by considering various learning scenarios based on the availability of class labels and sensitive attribute information in the target domain. Fig. 11 illustrates three scenarios for group fairness transfer when training data are available in the target domain. Another related scenario is the domain generalization (Pham et al. 2023) where no target samples are available during training.

- (1) *Labeled target samples with sensitive attributes*: Assuming that the target domain contains a few labeled samples with sensitive attributes, Schumann et al. (2019) provide the generalization error bounds of group fairness (e.g., equality of opportunity, and equalized odds) in the target domain in terms of fairness-aware distribution discrepancy between source and target domains. Oneto, Donini, Luise, et al. (2020) and Oneto, Donini, Pontil, et al. (2020) theoretically show the generalization error bound of group fairness (e.g., demographic parity) across domains from the perspective of multi-task learning via low-rank matrix factorization or parameter decoupling. Similarly, Slack et al. (2020) propose a fair meta-learning algorithm to transfer the fairness across domains.
- (2) *Unlabeled target samples with sensitive attributes*: When the target domain has only unlabeled samples with sensitive attributes, A. Rezaei et al. (2021) propose minimizing both the expected log-loss and the pseudo-label aware fairness penalty over the worst-case approximation of the target distribution to mitigate covariate shifts across domains and ensure fairness (e.g., demographic parity, equality of opportunity, and equalized odds) in the target domain. Havaldar et al. (2024) leverage representation matching across sensitive groups to enforce fairness and sample reweighting to mitigate covariate shifts across domains. Inspired by the theory of self-training (T. Cai et al. 2021; C. Wei, K. Shen, et al. 2021), An et al. (2022) theoretically analyze the transferability of group fairness across domains based on the consistency loss of a machine learning model under input transformations. Then they propose a self-training algorithm with fair consistency regularization to improve fairness transfer in the presence of subpopulation shifts. In contrast, Roh et al. (2023) formalize the notion of correlation shift over labels and sensitive attributes and employs a weighted sampling strategy in data preprocessing to mitigate correlation shifts across domains.

- (3) *Only unlabeled target samples with missing sensitive attributes*: [Coston et al. \(2019\)](#) study a more general learning scenario where the target domain is associated with only unlabeled samples with missing sensitive attributes. To improve group fairness (e.g., demographic parity) in the target domain, they develop fairness-guided sample reweighting approaches by enforcing the similarity of group-wise weighting scores across all pairs of groups.
- (4) *No target samples*: An extreme situation occurs when no target samples are available, commonly referred to as domain generalization or out-of-distribution generalization ([Blanchard et al. 2011](#); [Gulrajani and Lopez-Paz 2021](#)). In this scenario, only source domain data is provided to learn a fair predictor for unseen target domains. To solve this problem, [Singh et al. \(2021\)](#) develop a causal inference framework to minimize the worst-case prediction error under group fairness constraints. Similarly, [Mandal et al. \(2020\)](#) focus on optimizing a fair predictor by minimizing the worst-case error across weighted combinations of the training data. Later, [Pham et al. \(2023\)](#) derive the theoretical upper bounds on generalization error and unfairness in the target domain in terms of source error/unfairness, the domain discrepancy among source domains, and the domain discrepancy between source and unseen target domains. Motivated by this theoretical analysis, they propose an invariant representation learning algorithm to improve the transfer of fairness and accuracy via density matching.

**4.3.2 Individual Fairness.** Individual fairness requires that similar individuals (in the input space) should receive similar decision outcomes (in the output space) ([Dwork et al. 2012](#); [Zemel et al. 2013](#)). Individuals are similar if their only differences lie in protected attributes or features related to those attributes. Mathematically, [Dwork et al. \(2012\)](#) formalize this notion using  $L$ -Lipschitz continuity of a function  $f : \mathcal{X} \rightarrow \mathcal{Y}$ . For all  $x_1, x_2 \in \mathcal{X}$ , the following holds

$$d_{\mathcal{Y}}(f(x_1), f(x_2)) \leq L \cdot d_{\mathcal{X}}(x_1, x_2) \quad (22)$$

where  $L$  is a constant. Here,  $d_{\mathcal{X}}$  and  $d_{\mathcal{Y}}$  represent the distance metrics in the input space and output space, respectively. Recently, [D. Mukherjee et al. \(2022\)](#) investigate the connections between individual fairness and knowledge transferability in unsupervised domain adaptation/generalization scenarios. They show that (i) enforcing individual fairness (e.g., graph Laplacian regularizer ([J. Kang et al. 2020](#))) can theoretically improve the generalization performance of a learning function under the covariate shift assumption, and (ii) invariant representation learning commonly used in existing domain adaptation algorithms ([Ganin et al. 2016](#)) can improve individual fairness. Besides, [Ruoss et al. \(2020\)](#) propose an end-to-end framework to learn individually fair representations with provable certification and demonstrates the transferability of individual fairness using the learned representation. [Wicker et al. \(2023\)](#) further study the certification of distributional individual fairness ([Yurochkin et al. 2020](#)), which enforces the individual fairness within a  $\gamma$ -Wassertein ball of the empirical distribution over a finite set of observed individuals. The proposed distributional individual fairness regularization explicitly enables the transferability of individual fairness under in-the-wild distribution shifts.

#### 4.4 Transparency

Transparency helps non-experts understand the decision-making process of a machine learning model and the confidence level of the model in making decisions ([Varshney 2022](#)). For example, interpretability and explainability have recently been studied to enhance transparency, by designing a simpler and more interpretable model ([Koh and P. Liang 2017](#); [Ribeiro et al. 2016](#)) or providing post-hoc explanations for existing black-box models ([Selvaraju et al. 2017](#)). As a complementary metric of transparency, uncertainty quantification ([Bhatt et al. 2021](#)) illustrates the prediction confidence of a trained model. As a result, we study two major questions behind transparent transfer learning: what knowledge is being transferred in transfer learning, and how to quantify the uncertainty of transfer learning models.

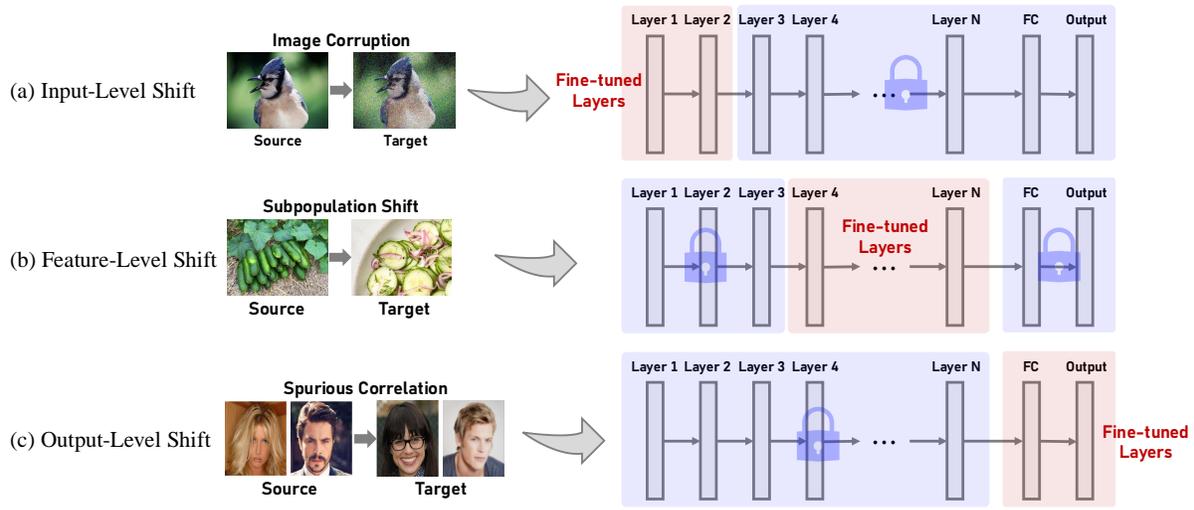


Fig. 12. Illustration of surgical fine-tuning (adapted from (Y. Lee et al. 2023)), where the selected fine-tuning blocks are correlated with the types of distribution shifts between source and target domains.

**4.4.1 Interpretability/Explanability.** Despite the promising performance of transfer learning techniques in a range of applications, limited effort has been devoted to understanding which data and model architecture components contribute to successful knowledge transfer across domains. To bridge this gap, Yosinski et al. (2014) demonstrate that for neural networks pre-trained on ImageNet data set (J. Deng et al. 2009), modules in the lower layers are responsible for capturing general features (e.g., Gabor and color blob features in images), while the higher-layer modules tend to encode task-specific semantic features. Neyshabur et al. (2020) further support this finding from the perspective of module criticality (Chatterji et al. 2020). Besides, they also reveal that both feature reuse and low-level data statistics are crucial for successful knowledge transfer. More recently, Y. Lee et al. (2023) establish the connections between fine-tuned neural layers and types of distribution shifts (shown in Fig. 12). They find that fine-tuning the first block is most effective for input-level shifts (such as image corruption), intermediate blocks excel at feature-level shifts (like shifts in entity subgroups), and tuning the last layer is best for output-level shifts (such as spurious correlations between gender and hair color). Raghu et al. (2019) also investigate transfer learning for medical imaging. They show that using a larger pre-trained ImageNet model does not significantly improve performance compared to smaller lightweight convolutional networks. Additionally, it is observed that transfer learning provides feature-independent benefits, such as improved weight scaling and faster convergence. This is consistent with observations from (K. He et al. 2019; Kornblith et al. 2019).

In addition to understanding the transferability of pre-trained models, recent efforts have been devoted to exploring the explanations of distribution shifts across domains in the distribution space. There are two major frameworks for distribution shift explanations, including interpretable transportation mapping (Kulinski and Inouye 2023; Stein et al. 2023) and natural languages (Dunlap et al. 2024; R. Zhong, P. Zhang, et al. 2023; Z. Zhu, W. Liang, et al. 2022). Specifically, on one hand, Kulinski and Inouye (2022, 2023) explain distribution shifts using interpretable transportation maps indicating how the source distribution can move to the target distribution in the distribution space. The crucial idea is to leverage optimal transport to find the optimal transportation map from user-defined interpretable candidates. Stein et al. (2023) further propose a group-aware shift explanation framework to rectify the group irregularities when explaining distribution shifts. On the other hand, Z. Zhu, W. Liang, et al. (2022) develop a GSCLIP system to explain distribution shifts of different image data sets in natural

language. This system generates human-understandable natural language descriptions of distribution shifts as candidate explanations, and then quantitatively evaluates these candidates to identify the most reasonable ones. R. Zhong, Snell, et al. (2022) study the explanations for text distribution shifts through natural languages. They prompt GPT-3 (Brown et al. 2020) to generate candidate explanations and then employ a verifier neural network to re-rank these explanations. Similarly, Dunlap et al. (2024) leverage visual-language models to generate candidate difference descriptions from image sets and then re-rank these candidates based on their effectiveness in distinguishing the two sets. In contrast, based on graphical causal models, Budhathoki et al. (2021) propose a Shapley value framework to quantify the attribution for each causal mechanism for distribution shifts. The follow-up work (H. Zhang, Singh, et al. 2023) further explores connections between model performance changes across domains and interpretable distribution shifts via Shapley values.

**4.4.2 Uncertainty Quantification.** Uncertainty quantification is essential for decision-making and optimization in machine learning and artificial intelligence (Naeini et al. 2015). For example, high-stakes applications such as medical diagnostics (Begoli et al. 2019) and autonomous driving (Michelmore et al. 2020) require both accurate class predictions and quantification of prediction uncertainty. Generally, there are two types of prediction uncertainty (Hüllermeier and Waegeman 2021): aleatoric (data) uncertainty involving the inherent randomness and variability in the data, and epistemic (model) uncertainty caused by a lack of knowledge about the optimal model parameters. These uncertainties can be formally explained using the Bayesian posterior distribution (Chan et al. 2020).

*Definition 4.4 (Aleatoric and Epistemic Uncertainty (Chan et al. 2020)).* Given a model  $f$  with parameter  $\theta$  and a test sample  $x^*$ , the Bayesian posterior distribution over  $x^*$  can be formulated as:

$$\underbrace{p(y^*|x^*, D)}_{\text{Total uncertainty}} = \int \underbrace{p(y^*|x^*, f)}_{\text{Aleatoric uncertainty}} \cdot \underbrace{p(f|D)}_{\text{Epistemic uncertainty}} df \quad (23)$$

where  $D$  denotes the set of training samples.

The calibration of uncertainty estimates is vital in determining the trustworthiness of model outputs. A well-calibrated model should provide accurate predictions when it is confident and indicate high uncertainty when it is likely to be incorrect. Thus, calibration can be considered as an orthogonal metric for accuracy when evaluating machine learning systems. In particular, Snoek et al. (2019) conduct a systematic evaluation of traditional uncertainty quantification models under distribution shifts. They observe that the quality of uncertainty consistently degrades with increasing distribution shifts between source and target domains. To solve this problem, various frameworks have been proposed to improve uncertainty quantification under distribution shifts across domains.

- (1) *Temperature Scaling:* Park, Bastani, et al. (2020) derive an upper bound on the expected calibration error in the target domain in terms of the importance-weighted classification error and the error of a domain discriminator. Building on the idea of temperature scaling (C. Guo et al. 2017), they propose a calibration algorithm by minimizing the upper bound over source and target samples. Similarly, X. Wang et al. (2020) develop an adaptive importance weighting approach with lower bias and variance of the estimated calibration errors to improve the uncertainty quantification under the covariate shift assumption. Instead, Y. Zou et al. (2023) focus on learning two calibration functions based on a real in-distribution calibration set and a synthetic out-of-distribution calibration set respectively, and then adaptively combines the two calibrators. D. Hu et al. (2024) optimize the calibration objective function (i.e., temperature scaling optimization (C. Guo et al. 2017)) using a labeled pseudo-target set created via mixup (H. Zhang, Cisse, et al. 2018) over pseudo-labeled target samples.

- (2) *Conformal Prediction*: The model predicts a set of labels instead of a single label (Angelopoulos, Bates, et al. 2023; J. Lei et al. 2018; Romano et al. 2020). Assuming the true importance weights (e.g.,  $w^*(x, y) = p_T(x, y)/p_S(x, y)$ ) are known, Podkopaev and Ramdas (2021) and Tibshirani et al. (2019) study the weighted conformal predictions under covariate shifts and label shifts, respectively. Based on jackknife+ (Barber et al. 2021), Prinster, A. Liu, et al. (2022) and Prinster, Saria, et al. (2023) formulate the sampling weighted jackknife+ prediction interval to handle covariate shifts with a finite-sample coverage guarantee. Cauchois et al. (2024) design prediction sets that are robust against all distribution shifts with bounded  $f$ -divergence. Gibbs and Candès (2021, 2024) further investigate prediction sets in an online setting where the data distribution can shift continuously over time. In addition, Park, Dobriban, et al. (2022) construct probably approximately correct (PAC) prediction sets under bounded covariate shifts in scenarios with known importance weights and an uncertainty set of possible importance weights. The follow-up work (Si et al. 2024) constructs prediction sets with PAC guarantees in the presence of label shifts. The crucial idea is to compute confidence intervals of importance weights (Lipton et al. 2018) through Gaussian elimination.
- (3) *Bayesian Learning*: Chan et al. (2020) develop an approximate Bayesian inference approach based on posterior regularization that captures the distribution difference between source and target domains. A. Zhou and Levine (2021) study the uncertainty quantification problem in the context of test-time adaptation. They develop a probabilistic graphical model for covariate shift scenarios, followed by an instantiated ensemble approach to estimate the uncertainty of trained models over test samples. In addition, transferable Gaussian processes (Bonilla et al. 2007; B. Cao et al. 2010; Maddox et al. 2021; J. Wu, J. He, S. Wang, et al. 2022; K. Yu and Chu 2007) can be applied to model uncertainty in the target domain by leveraging knowledge from the source domain.

## 4.5 Other Trustworthiness Concerns

**4.5.1 Accountability and Auditability.** Accountability is crucial to evaluate the trustworthiness of AI outcomes, as it identifies the organizations and individuals responsible for these results. More specifically, Bovens (2007) defines accountability as “a relationship between an actor and a forum, in which the actor has an obligation to explain and to justify his or her conduct, the forum can pose questions and pass judgment, and the actor may face consequences”. Wieringa (2020) further analyze five key aspects of this definition, including the actor, the forum, the relationship between them, the content and criteria of the account, and the potential consequences resulting from the account. In this case, auditability refers to systematic evaluations to guarantee accountability (Raji et al. 2020). Given the remarkable performance of fine-tuned large language models (LLMs), auditing LLMs has been studied through different principled assessments (Amirizani et al. 2024; Mökander et al. 2024). Recently, Pei et al. (2023) discuss data and AI model markets that facilitate the sharing, discovery, and integration of data and AI models among multiple parties. These markets can enhance knowledge transfer between pre-trained AI models and user-specific tasks, but they raise fundamental concerns regarding accountability in these systems. Further research can be conducted to guarantee accountability for transfer learning systems in supporting model and data knowledge sharing.

**4.5.2 Sustainability and Environmental Well-being.** To establish the trustworthiness of machine learning and artificial intelligence systems, it is crucial to evaluate resource usage and energy consumption within their entire supply chain (Budenny et al. 2022; Nikolinakos 2023). Notably, Schwartz et al. (2020) introduce a simple notion of computational cost in producing AI results.

*Definition 4.5 (Cost of an AI Result (Schwartz et al. 2020)).* The total cost of producing a (R)esult in AI increases with the following quantities.

$$\text{Cost}(R) \propto E \cdot D \cdot H \quad (24)$$

where  $E$  is the cost of executing the model on a single (E)xample,  $D$  is the size of the training (D)ataset, and  $H$  is the number of (H)yperparameter experiments.

To reduce the computational cost, green AI (K. Huang et al. 2024; Memmel et al. 2024; Schwartz et al. 2020) has been promoted by improving the efficiency of AI models with positive impacts on the environment. Several efficiency metrics have been introduced, including carbon emission, electricity usage, floating-point operations (FLOPs), elapsed runtime, and the number of parameters. Transfer learning techniques demonstrated significant improvements in training efficiency by leveraging knowledge from pre-trained models (K. He et al. 2019; Yosinski et al. 2014). This is because these approaches reduce (1) the size of training data  $D$  and hyperparameters  $H$ , and (2) the number of trainable model parameters (via parameter-efficient fine-tuning (Houlsby et al. 2019; E. J. Hu et al. 2022)). Furthermore, K. Huang et al. (2024) recently propose a GreenTrainer method to minimize the FLOPs of LLM fine-tuning via adaptive backpropagation. Qiu et al. (2023) take a first look into the carbon footprint of federated learning, by quantifying carbon emissions from hardware training and communication between server and clients. In real-world applications, transfer learning has been applied to lower energy consumption and reduce carbon emissions by reusing pre-trained models. For example, Ahmed et al. (2024) and Kunwar (2024) analyze transfer learning techniques for garbage classification and flower classification in terms of both prediction accuracy and carbon emissions.

## 5 Applications

Trustworthy transfer learning has been widely applied to artificial intelligence and machine learning fields, including computer vision (Neyshabur et al. 2020), natural language processing (N. Ding, Y. Qin, et al. 2023), and graph learning (Ruiz, Chamon, et al. 2020). In addition, this section highlights real-world applications of trustworthy transfer learning in scientific discovery.

### 5.1 Agriculture

Transfer learning techniques have been applied to various precision agriculture applications (Adve et al. 2024; Y. Ma et al. 2024). Specifically, to improve the management of agricultural stakeholders, S. Wang et al. (2023) and Y. Zhang, Hui, et al. (2021) propose process-guided machine learning frameworks, which transfer knowledge from simulated data generated by soil-vegetation radiative transfer modeling to real-world field data for precise monitoring of cover crop traits. L. Wan et al. (2022) analyze the transferability of support vector regression models for estimating leaf nitrogen concentration across different plant species. Besides, pre-trained vision models have been fine-tuned for crop mapping (Jo et al. 2022), crop pest classification (Thenmozhi and Reddy 2019), and plant phenotyping (Sama et al. 2023).

### 5.2 Bioinformatics

Notably, Theodoris et al. (2023) introduce an attention-based foundation model Geneformer pre-trained on over 30 million single-cell transcriptomes to capture network dynamics (e.g., gene interactions). They also demonstrate the effectiveness of Geneformer in various downstream tasks with limited data through fine-tuning. Later, Hou and Z. Ji (2024) illustrate the efficacy of the pre-trained large language model GPT-4 in cell type annotation of single-cell RNA-seq data. Besides, J. Hu, X. Li, et al. (2020) develop a unified transfer learning framework for open-world single-cell classification across different species and tissues. Similarly, Mieth et al. (2019) study the clustering of single-cell RNA-seq data on the small disease- or tissue-specific data sets by leveraging prior knowledge from large reference data sets. Hetzel et al. (2022) and Lotfollahi et al. (2022) leverage architecture surgery based transfer learning techniques to understand cellular heterogeneity.

In addition, recent efforts (Detlefsen et al. 2022; Heinzinger et al. 2019; Rao et al. 2019) have been devoted to protein representation learning for downstream tasks using language models pre-trained on a large protein

corpus. Typically, [Rao et al. \(2019\)](#) introduce a protein transfer learning benchmark TAPE for learning transferable protein representation, and [Detlefsen et al. \(2022\)](#) further improve the quality of protein representation by considering the geometry of representation space. [Dieckhaus et al. \(2024\)](#) also exploit the pre-trained ProteinMPNN model ([Dauparas et al. 2022](#)) to extract embeddings of input proteins, which are then used to predict stability changes for protein point mutations.

### 5.3 Healthcare

Transfer learning advances the development of effective and efficient health care services ([Jayaraman et al. 2020](#)). For example, [Y. Chen, X. Qin, et al. \(2020\)](#) develop a federated transfer learning framework for privacy-preserving wearable healthcare systems (e.g., Parkinson’s disease auxiliary diagnosis). [Matsoukas et al. \(2022\)](#) and [Raghu et al. \(2019\)](#) further understand the impact of the source domain/model on the downstream medical imaging tasks in the context of transfer learning. To enforce health equity across ethnic groups, [Gao and Cui \(2020\)](#), [T. Lee et al. \(2023\)](#), and [Toseef et al. \(2022\)](#) propose transferring knowledge from majority groups with sufficient data to minority groups with limited data. In addition, transfer learning techniques have been applied to drug discovery ([Chenjing et al. 2020](#)). Specifically, [H. Yao et al. \(2021\)](#) propose a functional rationalized meta-learning algorithm to enable knowledge transfer across assays for virtual screening and ADMET prediction. [Dalkıran et al. \(2023\)](#) and [Goh et al. \(2018\)](#) adopt pre-training and fine-tuning strategies for molecular property prediction and drug-target interaction prediction, respectively.

### 5.4 Education

Transfer learning has been studied in Educational Data Mining (EDM) for predicting student performance in higher education ([Hunt et al. 2017](#)). Over the past decade, Massive Open Online Courses (MOOCs) have supported millions of learners around the world. Early predictions of student performance are crucial for enabling timely interventions in these courses. Transfer learning has been explored to predict student performance in ongoing courses by leveraging knowledge from previous courses. To be specific, [Boyer and Veeramachaneni \(2015\)](#) leverage knowledge from both previous courses and previous weeks of the same course to make real-time predictions for learners in MOOCs. Instead of relying on handcrafted features, [M. Ding et al. \(2019\)](#) aim to learn domain-invariant representation by using an auto-encoder and correlation alignment ([B. Sun and Saenko 2016](#)) between source and target courses. Similarly, [Swamy et al. \(2022\)](#) study the transferability of early success prediction models across MOOCs from different domains and topics. Besides, [Schmucker and Mitchell \(2022\)](#) explore the transferability of student performance in addressing the cold-start problem for new courses in intelligent tutoring systems. More recently, large language models such as GPT-4 and ChatGPT have gained significant attention for improving instructional efficiency and student engagement (e.g., by creating interactive homework with feedback and follow-up questions) ([Kasneji et al. 2023](#); [Vanzo et al. 2025](#)).

### 5.5 Robotics

Sim-to-real transfer aims to transfer knowledge from simulation to real-world environments when training reinforcement learning models for robotic learning ([T. Dai et al. 2024](#); [X. B. Peng et al. 2018](#)). Recently, this framework has been studied in various robotic learning tasks, including Rubik’s cube ([OpenAI et al. 2019](#)), human pose estimation ([Doersch and Zisserman 2019](#)), vision-and-language navigation ([Anderson et al. 2020](#)), biped locomotion ([W. Yu et al. 2019](#)), etc. Specifically, various strategies have been proposed to address the domain shift between simulation and real-world environments ([X. B. Peng et al. 2018](#); [Pinto et al. 2017](#); [Rusu et al. 2017](#); [Tzeng, Devin, et al. 2020](#)). One is to use distribution alignment regularization to learn domain-invariant representation ([Tanwani 2020](#); [Tzeng, Devin, et al. 2020](#)). Another strategy is domain randomization ([Andrychowicz et al. 2020](#); [Tobin et al. 2017](#)), which aims to train the model using a diverse set of randomized simulated environments,

rather than relying on a single simulated environment. X. Chen, J. Hu, et al. (2022) and J. Hu, H. Zhong, et al. (2023) further theoretically highlight the benefits of domain randomization for sim-to-real transfer.

## 5.6 E-commerce

Cross-domain recommendation aims to generate reliable recommendations in a target domain by exploiting knowledge from source recommender systems. It has been studied from various perspectives, e.g., matrix factorization (Man et al. 2017; Samra et al. 2024), neural collaborative filtering (G. Hu et al. 2018; Kanagawa et al. 2019; P. Li and Tuzhilin 2020), graph neural network (J. Wu, J. He, and E. A. Ainsworth 2023; C. Zhao, C. Li, et al. 2019), large language models (Petruzzelli et al. 2024), etc. In addition to prediction accuracy, recent works also investigate the trustworthiness properties of cross-domain recommender systems, such as adversarial vulnerability (H. Chen and J. Li 2019) and user privacy (Z. Yang et al. 2024).

## 6 Open Questions and Future Trends

Despite the rapidly increasing research interest and applications of trustworthy transfer learning in both academia and industry, there remain many open questions, especially in the theoretical understanding of trustworthy transfer learning.

### 6.1 Benchmarking Negative Transfer

Negative transfer can roughly be defined as the phenomenon (Pan and Q. Yang 2010) where “*transferring knowledge from the source can have a negative impact on the target learner*”.

*Definition 6.1 (Negative Transfer (Pan and Q. Yang 2010; Z. Wang, Z. Dai, et al. 2019)).* Given a learning algorithm  $A^{\text{tl}}$ , source data  $D_S$ , and target data  $D_T$ , negative transfer occurs if the following condition holds:

$$\mathcal{E}_T(A^{\text{tl}}(D_S, D_T)) > \mathcal{E}_T(A^{\text{tl}}(\emptyset, D_T)) \quad (25)$$

where  $\mathcal{E}_T(A^{\text{tl}}(S))$  represents the expected error on the target distribution  $P_T$  when the learning algorithm  $A^{\text{tl}}$  is trained on data  $S$ .

This definition reveals (Ben-David, Blitzer, et al. 2010; Kuzborskij and Orabona 2013; Z. Wang, Z. Dai, et al. 2019) that given a learning algorithm, there are two major factors determining if negative transfer occurs: the *distribution discrepancy* between source and target domains and the *size of the labeled target data*. Negative transfer has been observed theoretically and empirically in various applications (Ben-David, T. Lu, et al. 2010; Rosenstein et al. 2005; Z. Wang, Z. Dai, et al. 2019; H. Zhao, Combes, et al. 2019). Recent work (Cohen-Wang et al. 2025) also explores identifying and characterizing the failure modes that pre-training can and cannot address. Despite the extensive work on transfer learning techniques, up until now, little effort (if any) has been devoted to rigorously understanding the boundary between positive and negative transfer given a learning algorithm. It remains an open question to determine when the negative transfer will occur given finite source and target samples (or a source model and finite target samples). Therefore, rather than focusing on performance improvement, more efforts can be dedicated to benchmarking the negative transfer of transfer learning models, e.g., the change from positive to negative transfer can be affected by the magnitude of distribution shifts and the number of target samples. This could provide valuable insights into when a transfer learning model can work well for real-world applications.

### 6.2 Cross-modal Transferability

Cross-modal transfer learning (T. Dinh et al. 2022; J. Shen, L. Li, et al. 2023; Socher et al. 2013) aims at understanding knowledge transferability when the source and target domains have different types of data modalities, e.g.,

transferring knowledge from a text-based source domain to an image-based target domain. This differs from multi-modal learning (Y. Huang et al. 2021; Radford et al. 2021), which maps different data modalities in a unified latent feature space over pair-wise training samples. In contrast, cross-modal transfer learning focuses on investigating what knowledge can be transferred across data modalities. Although large language models (LLMs) have been applied to various scientific discovery tasks (T. Dinh et al. 2022), it is unclear what knowledge is being transferred in this process. Additionally, there is a lack of theoretical understanding regarding how LLMs generalize to downstream tasks with different data modalities.

### 6.3 Physics-informed Transfer Learning

Physics-informed machine learning (Karniadakis et al. 2021; Raissi et al. 2019) aims to improve the training of machine learning models by incorporating physical domain knowledge as soft constraints on an empirical loss function. This alleviates the need for a large amount of high-quality data when training deep neural networks to solve scientific problems. Recent studies have introduced transfer learning to understand the knowledge transferability of physics-informed neural networks across tasks. For example, Desai et al. (2022) study the transferability of physics-informed neural networks across differential equations. Goswami et al. (2022) and W. Xu et al. (2023) study the transfer learning performance of deep operator networks across partial differential equations. Subramanian et al. (2023) further analyze the transfer behavior of neural operators pre-trained on a mixture of different physics problems. However, the theoretical explanation regarding the generalization performance of physics-informed neural networks under distribution shifts is unclear. Besides, it can be seen that in the context of transfer learning, the source knowledge can be provided from multiple aspects, including labeled source samples (Wiles et al. 2022) (e.g., Subsection 3.1.1), pre-trained source models (OpenAI 2023) (e.g., Subsection 4.1.1), synthetic data generated by physics-based simulators (Andrychowicz et al. 2020) (e.g., Subsection 5.5), and fundamental physical rules (Karniadakis et al. 2021). This can motivate a generic physics-informed transfer learning problem involving multi-faceted knowledge transfer from the source to the target domains.

### 6.4 Trade-off between Transferability and Trustworthiness

In standard machine learning, the trade-off between prediction accuracy and trustworthiness has been theoretically studied, e.g., accuracy vs. group fairness (Dutta et al. 2020; H. Zhao and Gordon 2019), accuracy vs. adversarial robustness (Tsipras et al. 2019; Y. Yang et al. 2020; H. Zhang, Y. Yu, et al. 2019), accuracy vs. privacy (Bietti et al. 2022), accuracy vs. explainability (Zarlenga et al. 2022), etc. It has been observed that trustworthiness may be at odds with the prediction accuracy in a single domain. In contrast, recent work (Davchev et al. 2019; Salman et al. 2020) reveals that both trustworthiness (e.g., adversarial robustness) and prediction accuracy can be improved in the target domain by leveraging relevant knowledge from source domains. This motivates us to rethink the fundamental trade-off between knowledge transferability and trustworthiness in the context of transfer learning. Specifically, there are several open questions: (1) Can source knowledge consistently enhance trustworthiness and transfer accuracy in the target domain under various distribution shifts and data modalities? (2) Is there an inherent trade-off between trustworthiness and transfer accuracy in the target domain when discovering and transferring knowledge from the source data/model? These studies will significantly expand the application of transfer learning techniques by clarifying when trained models can be trusted and how well they perform.

### 6.5 Trustworthy Transfer Learning of Foundation Models

Although foundation models have achieved surprising performance across high-impact domains, the underlying properties governing their behavior are not yet fully understood (Bengio et al. 2024; Kapoor et al. 2024), especially when adapting them to downstream tasks. In this work, we summarize the key open questions from the perspective of trustworthy transfer learning. On one hand, knowledge transferability involves whether foundation models

can be successfully adapted to downstream tasks with guaranteed performance. This post-training process leads to the following open questions:

- *Scaling Behaviors of Fine-tuning*: Recent work (Tay et al. 2022; B. Zhang et al. 2024) has investigated the power-law scaling relationship between LLM fine-tuning performance and model/data size. In contrast, Ren and Sutherland (2025) provide a step-wise analysis of the learning dynamics for LLM fine-tuning under gradient descent. However, theoretical understanding still lags behind empirical findings, e.g., how generalization performance of fine-tuning is bounded in terms of model architecture, model and data scale, how learning dynamics vary across different fine-tuning (or other post-training) strategies, and how learning dynamics of fine-tuning explain the potential negative transfer behaviors.
- *Transferability Metrics for Foundation Model Selection*: Recent work (Lin et al. 2024) leverages rectified scaling laws to predict fine-tuning performance, thereby facilitating model selection for LLMs. However, it is computationally expensive due to the need for fine-tuning in low-data regimes, and the theoretical foundations connecting scaling laws to fine-tuning performance are still underexplored. To enable model selection and customization in the AI model market (Pei et al. 2023), there is a growing need for theory-grounded and computationally efficient transferability metrics. In particular, it is feasible and desirable to establish principled guidelines for selecting optimal foundation models for specific target tasks, as well as identifying the application domains in which a given foundation model is best suited for fine-tuning.
- *AI Model Collapse*: Generally, model collapse is defined as “a degenerative process affecting generations of learned generative models, in which the data they generate end up polluting the training set of the next generation” (Shumailov et al. 2024). This phenomenon motivates us to rethink the impact of synthetic data (maybe indistinguishable from human-generated content) on both the pre-training and fine-tuning of foundation models. From the perspective of transfer learning, a key challenge for mitigating model collapse lies in handling open-world distribution shifts between real and AI-generated data.

On the other hand, the knowledge trustworthiness of foundation models in the transfer learning process might involve the following open questions:

- *Data Poisoning and Model Stealing Attacks*: It is shown (Bowen et al. 2025; A. Wan et al. 2023) that larger foundation models are significantly more susceptible to data poisoning attacks than smaller ones during fine-tuning. However, limited efforts have been devoted to theoretically understanding the connections between model size and adversarial vulnerability of foundation models in the context of transfer learning. In addition, recent work (Carlini et al. 2024) discusses the model stealing attacks to extract precise information from black-box production language models via APIs. These emerging threats highlight significant security challenges for the development and deployment of foundation models in both academia and industry. As a result, advancing certified robust techniques is essential to improve the reliability and trustworthiness of foundation models.
- *Transparency and Mechanistic Interpretability*. Mechanistic interpretability aims at understanding the internal mechanisms of foundation models (Dunefsky et al. 2024; Ferrando et al. 2025), thereby providing insights into their decision-making processes and enhancing model transparency. Nevertheless, due to the architectural complexity and scale of foundation models, it is challenging to systematically explain their emergent behaviors and key mechanisms behind their advanced capabilities, when they are adapted to open-world environments. Although fragmented efforts have been made, a systematic framework is still needed to advance the mechanistic interpretability of foundation models.
- *Ethical and Societal Impact in Foundation Model Customization*: The ethical and societal implications, such as misinformation and privacy concerns, of open foundation models have been examined from the perspectives of AI developers (Kapoor et al. 2024; Nikolinakos 2023). As these models are increasingly fine-tuned and customized by millions of non-expert users, a critical gap emerges between expert understanding and

everyday use. This gap highlights the urgent need for intuitive risk assessments and protective measures that empower non-AI users to identify and mitigate intentional or unintentional misuse of foundation models.

## 7 Conclusion

In this survey, we provide a comprehensive review of trustworthy transfer learning from the perspective of knowledge transferability and trustworthiness. With different data and model assumptions, much effort has been devoted to understanding the generalization performance of trustworthy transfer learning and designing advanced techniques in quantifying and enhancing knowledge transfer in a variety of real-world applications. Besides, we also summarize several open questions for trustworthy transfer learning, including benchmarking positive and negative transfer, enabling unified knowledge transfer across different data modalities and physical rules, and achieving the inherent trade-off between transferability and trustworthiness. Ultimately, trustworthy transfer learning could lead to a unified machine learning and artificial intelligence framework that facilitates positive knowledge reuse and transfer in the presence of distribution shifts and across data modalities, while maintaining rigorous standards of trustworthiness.

## Acknowledgments

This work is supported by National Science Foundation under Award No. IIS-2117902, and Agriculture and Food Research Initiative (AFRI) grant no. 2020-67021-32799/project accession no.1024178 from the USDA National Institute of Food and Agriculture. The views and conclusions are those of the authors and should not be interpreted as representing the official policies of the funding agencies or the government.

## References

- I. Achituve, A. Shamsian, A. Navon, G. Chechik, and E. Fetaya. 2021. “Personalized Federated Learning with Gaussian Processes.” *Advances in Neural Information Processing Systems*, 34, 8392–8406.
- D. Acuna, G. Zhang, M. T. Law, and S. Fidler. 2021. “ $f$ -Domain Adversarial Learning: Theory and Algorithms.” In: *International Conference on Machine Learning*. Vol. 139. PMLR, 66–75.
- V. S. Adve, J. M. Wedou, E. A. Ainsworth, G. Chowdhary, A. Green-Miller, and C. Tucker. 2024. “AIFARMS: Artificial Intelligence for Future Agricultural Resilience, Management, and Sustainability.” *AI Magazine*, 45, 1, 83–88.
- A. Aghajanyan, S. Gupta, and L. Zettlemoyer. 2021. “Intrinsic Dimensionality Explains the Effectiveness of Language Model Fine-Tuning.” In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 7319–7328.
- A. Aghbalou and G. Staerman. 2023. “Hypothesis Transfer Learning with Surrogate Classification Losses: Generalization Bounds through Algorithmic Stability.” In: *International Conference on Machine Learning*. PMLR, 280–303.
- A. Agostinelli, M. Pándy, J. R. R. Uijlings, T. Mensink, and V. Ferrari. 2022. “How Stable Are Transferability Metrics Evaluations?” In: *European Conference on Computer Vision*. Springer, 303–321.
- A. Agostinelli, J. R. R. Uijlings, T. Mensink, and V. Ferrari. 2022. “Transferability Metrics for Selecting Source Model Ensembles.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 7926–7936.
- E. Ahmed, A. Darwish, and A. E. Hassanien. 2024. “Flowers Classification with Low Carbon Footprint Using Deep Learning Pretrained Models.” In: *Artificial Intelligence for Environmental Sustainability and Green Initiatives*. Springer, 51–69.
- K. Alhamoud, H. A. A. K. Hammoud, M. Alfarra, and B. Ghanem. 2023. “Generalizability of Adversarial Robustness Under Distribution Shifts.” *Transactions on Machine Learning Research*, 2023.
- M. Amirizani, E. Martin, T. Roosta, A. Chadha, and C. Shah. 2024. “AuditLLM: A Tool for Auditing Large Language Models Using Multiprobe Approach.” In: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 5174–5179.
- B. An, Z. Che, M. Ding, and F. Huang. 2022. “Transferring Fairness under Distribution Shifts via Fair Consistency Regularization.” In: *Advances in Neural Information Processing Systems* 35.
- P. Anderson, A. Shrivastava, J. Truong, A. Majumdar, D. Parikh, D. Batra, and S. Lee. 2020. “Sim-to-Real Transfer for Vision-and-Language Navigation.” In: *4th Conference on Robot Learning*. PMLR, 671–681.
- M. Andrychowicz et al.. 2020. “Learning Dexterous In-hand Manipulation.” *The International Journal of Robotics Research*, 39, 1.

- A. N. Angelopoulos, S. Bates, et al.. 2023. “Conformal prediction: A gentle introduction.” *Foundations and Trends® in Machine Learning*, 16, 4, 494–591.
- A. Ansell, E. M. Ponti, A. Korhonen, and I. Vulic. 2022. “Composable Sparse Fine-Tuning for Cross-Lingual Transfer.” In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1778–1796.
- N. Arivazhagan, A. Bapna, O. Firat, R. Aharoni, M. Johnson, and W. Macherey. 2019. “The Missing Ingredient in Zero-Shot Neural Machine Translation.” *CoRR*, abs/1903.07091.
- M. Arjovsky, S. Chintala, and L. Bottou. 2017. “Wasserstein Generative Adversarial Networks.” In: *International Conference on Machine Learning*. PMLR, 214–223.
- S. Arora, S. S. Du, W. Hu, Z. Li, and R. Wang. 2019. “Fine-Grained Analysis of Optimization and Generalization for Overparameterized Two-Layer Neural Networks.” In: *International Conference on Machine Learning*. PMLR, 322–332.
- M. Awais, F. Zhou, H. Xu, L. Hong, P. Luo, S. Bae, and Z. Li. 2021. “Adversarial Robustness for Unsupervised Domain Adaptation.” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 8548–8557.
- V. K. B., S. Bachu, T. Garg, N. L. Narasimhan, R. Konuru, and V. N. Balasubramanian. 2023. “Building a Winning Team: Selecting Source Model Ensembles using a Submodular Transferability Estimation Approach.” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 11575–11586.
- Á. Backhausz and B. Szegedy. 2022. “Action Convergence of Operators and Graphs.” *Canadian Journal of Mathematics*, 74, 1, 72–121.
- R. Bai, S. Bagchi, and D. I. Inouye. 2024. “Benchmarking Algorithms for Federated Domain Generalization.” In: *The Twelfth International Conference on Learning Representations*.
- W. Bao, H. Wang, J. Wu, and J. He. 2023. “Optimizing the Collaboration Structure in Cross-Silo Federated Learning.” In: *International Conference on Machine Learning*. Vol. 202. PMLR, 1718–1736.
- W. Bao, T. Wei, H. Wang, and J. He. 2023. “Adaptive Test-Time Personalization for Federated Learning.” *Advances in Neural Information Processing Systems*, 36.
- Y. Bao, Y. Li, S. Huang, L. Zhang, L. Zheng, A. Zamir, and L. J. Guibas. 2019. “An Information-Theoretic Approach to Transferability in Task Transfer Learning.” In: *2019 IEEE International Conference on Image Processing*. IEEE, 2309–2313.
- A. Bapna and O. Firat. 2019. “Simple, Scalable Adaptation for Neural Machine Translation.” In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, 1538–1548.
- R. F. Barber, E. J. Candès, A. Ramdas, and R. J. Tibshirani. 2021. “Predictive inference with the jackknife+.” *The Annals of Statistics*, 49, 1, 486–507.
- E. Begoli, T. Bhattacharya, and D. Kusnezov. 2019. “The Need for Uncertainty Quantification in Machine-Assisted Medical Decision Making.” *Nature Machine Intelligence*, 1, 1, 20–23.
- S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. 2010. “A Theory of Learning from Different Domains.” *Machine Learning*, 79, 1-2, 151–175.
- S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira. 2006. “Analysis of Representations for Domain Adaptation.” *Advances in Neural Information Processing Systems*, 19.
- S. Ben-David, T. Lu, T. Luu, and D. Pál. 2010. “Impossibility Theorems for Domain Adaptation.” In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 129–136.
- Y. Bengio et al.. 2024. “Managing extreme AI risks amid rapid progress.” *Science*, 384, 6698, 842–845.
- B. Bevilacqua, Y. Zhou, and B. Ribeiro. 2021. “Size-Invariant Graph Representations for Graph Classification Extrapolations.” In: *International Conference on Machine Learning*. PMLR, 837–851.
- U. Bhatt et al.. 2021. “Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty.” In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 401–413.
- A. Bhattacharyya. July 1946. “On a Measure of Divergence between Two Multinomial Populations.” *Sankhyā: The Indian Journal of Statistics*, 7, 4, (July 1946), 401–406.
- A. Bietti, C. Wei, M. Dudík, J. Langford, and Z. S. Wu. 2022. “Personalization Improves Privacy-Accuracy Tradeoffs in Federated Learning.” In: *International Conference on Machine Learning*. PMLR, 1945–1962.
- A. Biswas and S. Mukherjee. 2021. “Ensuring Fairness under Prior Probability Shifts.” In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 414–424.
- G. Blanchard, G. Lee, and C. Scott. 2011. “Generalizing from Several Related Classification Tasks to a New Unlabeled Sample.” In: *Advances in Neural Information Processing Systems 24*, 2178–2186.
- J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. Wortman. 2007. “Learning Bounds for Domain Adaptation.” *Advances in Neural Information Processing Systems*, 20.
- D. Bolya, R. Mittapalli, and J. Hoffman. 2021. “Scalable Diverse Model Selection for Accessible Transfer Learning.” In: *Advances in Neural Information Processing Systems 34*, 19301–19312.
- E. V. Bonilla, K. Chai, and C. Williams. 2007. “Multi-Task Gaussian Process Prediction.” *Advances in Neural Information Processing Systems*, 20.

- M. Boudiaf, T. Denton, B. Van Merriënboer, V. Dumoulin, and E. Triantafillou. 2023. “In Search for a Generalizable Method for Source Free Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 2914–2931.
- M. Bovens. 2007. “Analysing and Assessing Accountability: A Conceptual Framework.” *European Law Journal*, 13, 4, 447–468.
- D. Bowen, B. Murphy, W. Cai, D. Khachaturov, A. Gleave, and K. Pelrine. 2025. “Scaling Trends for Data Poisoning in LLMs.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 26. Vol. 39, 27206–27214.
- S. Boyer and K. Veeramachaneni. 2015. “Transfer Learning for Predictive Models in Massive Open Online Courses.” In: *Artificial Intelligence in Education*. Springer, 54–63.
- T. B. Brown et al.. 2020. “Language Models are Few-Shot Learners.” In: *Advances in Neural Information Processing Systems* 33.
- S. L. Brunton, M. Budišić, E. Kaiser, and J. N. Kutz. 2022. “Modern Koopman Theory for Dynamical Systems.” *SIAM Review*, 64, 2, 229–340.
- S. A. Budenny et al.. 2022. “Eco2AI: Carbon Emissions Tracking of Machine Learning Models as the First Step Towards Sustainable AI.” In: *Doklady Mathematics Suppl* 1. Vol. 106. Springer, S118–S128.
- K. Budhathoki, D. Janzing, P. Bloebaum, and H. Ng. 2021. “Why Did the Distribution Change?” In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 1666–1674.
- E. Bugliarello, F. Liu, J. Pfeiffer, S. Reddy, D. Elliott, E. M. Ponti, and I. Vulic. 2022. “IGLUE: A Benchmark for Transfer Learning across Modalities, Tasks, and Languages.” In: *International Conference on Machine Learning*. PMLR, 2370–2392.
- R. Cai, J. Chen, Z. Li, W. Chen, K. Zhang, J. Ye, Z. Li, X. Yang, and Z. Zhang. 2021. “Time Series Domain Adaptation via Sparse Associative Structure Alignment.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 8. Vol. 35, 6859–6867.
- T. Cai, R. Gao, J. Lee, and Q. Lei. 2021. “A Theory of Label Propagation for Subpopulation Shift.” In: *International Conference on Machine Learning*. PMLR, 1170–1182.
- X. Cai, H. Xu, S. Xu, Y. Zhang, and X. Yuan. 2022. “BadPrompt: Backdoor Attacks on Continuous Prompts.” In: *Advances in Neural Information Processing Systems* 35.
- B. Cao, S. J. Pan, Y. Zhang, D.-Y. Yeung, and Q. Yang. 2010. “Adaptive Transfer Learning.” In: *proceedings of the AAAI Conference on Artificial Intelligence* 1. Vol. 24, 407–412.
- N. Carlini et al.. 2024. “Stealing part of a production language model.” In: *Forty-first International Conference on Machine Learning*.
- A. Castelnovo, R. Crupi, G. Greco, D. Regoli, I. G. Penco, and A. C. Cosentini. 2022. “A Clarification of the Nuances in the Fairness Metrics Landscape.” *Scientific Reports*, 12, 1, 4209.
- M. Cauchois, S. Gupta, A. Ali, and J. C. Duchi. 2024. “Robust Validation: Confident Predictions Even When Distributions Shift.” *Journal of the American Statistical Association*, 0, 0, 1–66.
- J. Cerviño, L. Ruiz, and A. Ribeiro. 2023. “Learning by Transference: Training Graph Neural Networks on Growing Graphs.” *IEEE Transactions on Signal Processing*, 71, 233–247.
- A. J. Chan, A. M. Alaa, Z. Qian, and M. van der Schaar. 2020. “Unlabelled Data Improves Bayesian Uncertainty Calibration under Covariate Shift.” In: *International Conference on Machine Learning*. PMLR, 1392–1402.
- C. Chatfield. 2004. *The Analysis of Time Series: An Introduction*. CRC Press.
- N. Chatterji, B. Neyshabur, and H. Sedghi. 2020. “The Intriguing Role of Module Criticality in the Generalization of Deep Networks.” In: *International Conference on Learning Representations*.
- D. Chen, H. Hu, Q. Wang, Y. Li, C. Wang, C. Shen, and Q. Li. 2021. “CARTL: Cooperative Adversarially-Robust Transfer Learning.” In: *International Conference on Machine Learning*. PMLR, 1640–1650.
- H. Chen and J. Li. 2019. “Data Poisoning Attacks on Cross-domain Recommendation.” In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2177–2180.
- J. Chen and A. Zhang. 2024. “On Disentanglement of Asymmetrical Knowledge Transfer for Modality-Task Agnostic Federated Learning.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 10. Vol. 38, 11311–11319.
- T. Chen, S. Liu, S. Chang, Y. Cheng, L. Amini, and Z. Wang. 2020. “Adversarial Robustness: From Self-Supervised Pre-Training to Fine-Tuning.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 696–705.
- X. Chen, J. Hu, C. Jin, L. Li, and L. Wang. 2022. “Understanding Domain Randomization for Sim-to-real Transfer.” In: *International Conference on Learning Representations*.
- X. Chen, S. Wang, J. Wang, and M. Long. 2021. “Representation Subspace Distance for Domain Adaptation Regression.” In: *International Conference on Machine Learning*. PMLR, 1749–1759.
- Y. Chen, R. Raab, J. Wang, and Y. Liu. 2022. “Fairness Transferability Subject to Bounded Distribution Shift.” *Advances in Neural Information Processing Systems*, 35, 11266–11278.
- Y. Chen, Y. Huang, S. S. Du, K. G. Jamieson, and G. Shi. 2023. “Active Representation Learning for General Task Space with Applications in Robotics.” *Advances in Neural Information Processing Systems*, 36.
- Y. Chen, K. Jamieson, and S. Du. 2022. “Active Multi-Task Representation Learning.” In: *International Conference on Machine Learning*. PMLR, 3271–3298.
- Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao. 2020. “FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare.” *IEEE Intelligent Systems*, 35, 4, 83–93.

- C. Chenjing, S. Wang, Y. Xu, W. Zhang, K. Tang, Q. Ouyang, L. Lai, and J. Pei. 2020. “Transfer Learning for Drug Discovery.” *Journal of Medicinal Chemistry*, 63, 16, 8683–8694.
- H. Chi, F. Liu, W. Yang, L. Lan, T. Liu, B. Han, W. K. Cheung, and J. T. Kwok. 2021. “TOHAN: A One-step Approach towards Few-shot Hypothesis Adaptation.” In: *Advances in Neural Information Processing Systems 34*, 20970–20982.
- T. Chin, C. Zhang, and D. Marculescu. 2021. “Renofeation: A Simple Transfer Learning Method for Improved Adversarial Robustness.” In: *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 3243–3252.
- Y. J. Cho, D. Jhunjhunwala, T. Li, V. Smith, and G. Joshi. 2022. “To Federate or Not To Federate: Incentivizing Client Participation in Federated Learning.” In: *Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022)*.
- Y. Choi, S. Bae, S. Ban, M. Jeong, C. Zhang, L. Song, L. Zhao, J. Bian, and K.-E. Kim. 2024. “Hard Prompts Made Interpretable: Sparse Entropy Regularization for Prompt Tuning with RL.” In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8252–8271.
- J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio. 2015. “A Recurrent Latent Variable Model for Sequential Data.” In: *Advances in Neural Information Processing Systems 28*, 2980–2988.
- J. Cohen, E. Rosenfeld, and J. Z. Kolter. 2019. “Certified Adversarial Robustness via Randomized Smoothing.” In: *International Conference on Machine Learning*. PMLR, 1310–1320.
- B. Cohen-Wang, J. Vendrow, and A. Madry. 2025. “Ask Your Distribution Shift if Pre-Training is Right for You.” *Transactions on Machine Learning Research*.
- F. Cole, Y. Lu, T. Zhang, and R. C. W. O’Neill. 2024. “Provable in-context learning of linear systems and linear elliptic PDEs with transformers.” In: *Neurips 2024 Workshop Foundation Models for Science: Progress, Opportunities, and Challenges*.
- L. Collins, H. Hassani, A. Mokhtari, and S. Shakkottai. 2021. “Exploiting Shared Representations for Personalized Federated Learning.” In: *International Conference on Machine Learning*. PMLR, 2089–2099.
- E. Commission, C. Directorate-General for Communications Networks, and Technology. 2019. *Ethics Guidelines for Trustworthy AI*. Publications Office.
- C. Cortes, Y. Mansour, and M. Mohri. 2010. “Learning Bounds for Importance Weighting.” *Advances in neural information processing systems*, 23.
- C. Cortes and M. Mohri. 2011. “Domain Adaptation in Regression.” In: *International Conference on Algorithmic Learning Theory*. Springer, 308–323.
- C. Cortes, M. Mohri, and A. M. Medina. 2015. “Adaptation Algorithm and Theory Based on Generalized Discrepancy.” In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 169–178.
- C. Cortes, M. Mohri, and A. M. Medina. 2019. “Adaptation Based on Generalized Discrepancy.” *Journal of Machine Learning Research*, 20, 1, 1–30.
- C. Cortes, M. Mohri, M. Riley, and A. Rostamizadeh. 2008. “Sample Selection Bias Correction Theory.” In: *International Conference on Algorithmic Learning Theory*. Springer, 38–53.
- A. Coston, K. N. Ramamurthy, D. Wei, K. R. Varshney, S. Speakman, Z. Mustahsan, and S. Chakraborty. 2019. “Fair Transfer Learning with Missing Protected Attributes.” In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 91–98.
- N. Courty, R. Flamary, A. Habrard, and A. Rakotomamonjy. 2017. “Joint Distribution Optimal Transportation for Domain Adaptation.” *Advances in Neural Information Processing Systems*, 30.
- N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. 2016. “Optimal Transport for Domain Adaptation.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39, 9, 1853–1865.
- M. Cuturi. 2013. “Sinkhorn Distances: Lightspeed Computation of Optimal Transport.” In: *Advances in Neural Information Processing Systems 26*. Ed. by C. J. C. Burges, L. Bottou, Z. Ghahramani, and K. Q. Weinberger, 2292–2300.
- Q. Dai, X. Wu, J. Xiao, X. Shen, and D. Wang. 2023. “Graph Transfer Learning via Adversarial Domain Adaptation With Graph Convolution.” *IEEE Transactions on Knowledge and Data Engineering*, 35, 5, 4908–4922.
- T. Dai, J. Wong, Y. Jiang, C. Wang, C. Gokmen, R. Zhang, J. Wu, and L. Fei-Fei. 2024. “ACDC: Automated Creation of Digital Cousins for Robust Policy Learning.” In: *8th Annual Conference on Robot Learning*.
- A. Dalkiran, A. Atakan, A. S. Rifaioğlu, M. J. Martin, R. Ç. Atalay, A. C. Acar, T. Doğan, and V. Atalay. 2023. “Transfer Learning for Drug-Target Interaction Prediction.” *Bioinformatics*, 39, i103–i110.
- S. S. S. Das, H. Zhang, P. Shi, W. Yin, and R. Zhang. 2023. “Unified Low-Resource Sequence Labeling by Sample-Aware Dynamic Sparse Finetuning.” In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 6998–7010.
- J. Dauparas et al.. 2022. “Robust Deep Learning–based Protein Sequence Design Using ProteinMPNN.” *Science*, 378, 6615, 49–56.
- T. Davchev, T. Korres, S. Fotiadis, N. Antonopoulos, and S. Ramamoorthy. 2019. “An Empirical Evaluation of Adversarial Robustness under Transfer Learning.” In: *ICML 2019 Workshop on Understanding and Improving Generalization in Deep Learning*.
- I. Dayan et al.. 2021. “Federated Learning for Predicting Clinical Outcomes in Patients with COVID-19.” *Nature Medicine*, 27, 10, 1735–1743.
- M. Defferrard, X. Bresson, and P. Vandergheynst. 2016. “Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering.” In: *Advances in Neural Information Processing Systems 29*, 3837–3845.

- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. "ImageNet: A Large-Scale Hierarchical Image Database." In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 248–255.
- M. Deng, J. Wang, C. Hsieh, Y. Wang, H. Guo, T. Shu, M. Song, E. P. Xing, and Z. Hu. 2022. "RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning." In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 3369–3391.
- Z. Deng, L. Zhang, K. Vodrahalli, K. Kawaguchi, and J. Y. Zou. 2021. "Adversarial Training Helps Transfer Learning via Better Representations." In: *Advances in Neural Information Processing Systems 34*, 25179–25191.
- S. Desai, M. Mattheakis, H. Joy, P. Protopapas, and S. J. Roberts. 2022. "One-Shot Transfer Learning of Physics-Informed Neural Networks." In: *ICML 2022 2nd AI for Science Workshop*.
- N. S. Dettlarsen, S. Hauberg, and W. Boomsma. 2022. "Learning Meaningful Representations of Protein Sequences." *Nature Communications*, 13, 1, 1914.
- J. Devlin, M. Chang, K. Lee, and K. Toutanova. 2019. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186.
- H. Dieckhaus, M. Brocidiaco, N. Z. Randolph, and B. Kuhlman. 2024. "Transfer Learning to Leverage Larger Datasets for Improved Prediction of Protein Stability Changes." *Proceedings of the National Academy of Sciences*, 121, 6, e2314853121.
- M. Ding, Y. Wang, E. Hemberg, and U.-M. O'reilly. 2019. "Transfer Learning Using Representation Learning in Massive Open Online Courses." In: *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*, 145–154.
- N. Ding, X. Chen, T. Levinboim, S. Changpinyo, and R. Soricut. 2022. "PACTran: PAC-Bayesian Metrics for Estimating the Transferability of Pretrained Models to Classification Tasks." In: *European Conference on Computer Vision*. Springer, 252–268.
- N. Ding, X. Lv, Q. Wang, Y. Chen, B. Zhou, Z. Liu, and M. Sun. 2023. "Sparse Low-rank Adaptation of Pre-trained Language Models." In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 4133–4145.
- N. Ding, Y. Qin, et al.. 2023. "Parameter-Efficient Fine-Tuning of Large-scale Pre-trained Language Models." *Nature Machine Intelligence*, 5, 3, 220–235.
- C. T. Dinh, N. H. Tran, and T. D. Nguyen. 2020. "Personalized Federated Learning with Moreau Envelopes." In: *Advances in Neural Information Processing Systems 33*.
- T. Dinh, Y. Zeng, R. Zhang, Z. Lin, M. Gira, S. Rajput, J. Sohn, D. S. Papailiopoulos, and K. Lee. 2022. "LIFT: Language-Interfaced Fine-Tuning for Non-language Machine Learning Tasks." In: *Advances in Neural Information Processing Systems 35*.
- C. Doersch and A. Zisserman. 2019. "Sim2real Transfer Learning for 3D Human Pose Estimation: Motion to the Rescue." In: *Advances in Neural Information Processing Systems 32*, 12929–12941.
- K. Donahue and J. Kleinberg. 2021. "Model-sharing Games: Analyzing Federated Learning Under Voluntary Participation." In: *Proceedings of the AAAI Conference on Artificial Intelligence 6*. Vol. 35, 5303–5311.
- J. Dong, H. Wu, H. Zhang, L. Zhang, J. Wang, and M. Long. 2023. "SimMTM: A Simple Pre-Training Framework for Masked Time-Series Modeling." In: *Advances in Neural Information Processing Systems 36*.
- R. Dong, F. Liu, H. Chi, T. Liu, M. Gong, G. Niu, M. Sugiyama, and B. Han. 2023. "Diversity-enhancing Generative Network for Few-shot Hypothesis Adaptation." In: *International Conference on Machine Learning*. PMLR, 8260–8275.
- X. Dong, A. T. Luu, M. Lin, S. Yan, and H. Zhang. 2021. "How Should Pre-Trained Language Models Be Fine-Tuned Towards Adversarial Robustness?" In: *Advances in Neural Information Processing Systems 34*, 4356–4369.
- S. S. Du, J. Koushik, A. Singh, and B. Póczos. 2017. "Hypothesis Transfer Learning via Transformation Functions." In: *Advances in Neural Information Processing Systems 30*, 574–584.
- S. S. Du, W. Hu, S. M. Kakade, J. D. Lee, and Q. Lei. 2021. "Few-Shot Learning via Learning the Representation, Provably." In: *International Conference on Learning Representations*.
- Y. Du, J. Wang, W. Feng, S. J. Pan, T. Qin, R. Xu, and C. Wang. 2021. "AdaRNN: Adaptive Learning and Forecasting of Time Series." In: *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*. ACM, 402–411.
- R. Dudley. 2002. *Real Analysis and Probability*. Cambridge University Press.
- J. Dunefsky, P. Chlenski, and N. Nanda. 2024. "Transcoders find interpretable llm feature circuits." *Advances in Neural Information Processing Systems*, 37, 24375–24410.
- L. Dunlap, Y. Zhang, X. Wang, R. Zhong, T. Darrell, J. Steinhardt, J. E. Gonzalez, and S. Yeung-Levy. 2024. "Describing Differences in Image Sets with Natural Language." In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24199–24208.
- R. Dutt, O. Bohdal, S. A. Tsafaris, and T. Hospedales. 2024. "FairTune: Optimizing Parameter Efficient Fine Tuning for Fairness in Medical Image Analysis." In: *International Conference on Learning Representations*.
- S. Dutta, D. Wei, H. Yueksel, P. Chen, S. Liu, and K. R. Varshney. 2020. "Is There a Trade-Off Between Fairness and Accuracy? A Perspective Using Mismatched Hypothesis Testing." In: *International Conference on Machine Learning*. PMLR, 2803–2813.
- C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. S. Zemel. 2012. "Fairness through Awareness." In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM, 214–226.
- B. Eshete. 2021. "Making Machine Learning Trustworthy." *Science*, 373, 6556, 743–744.

- A. Fallah, A. Mokhtari, and A. E. Ozdaglar. 2020. “Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach.” In: *Advances in Neural Information Processing Systems* 33.
- W. Fan, P. Wang, D. Wang, D. Wang, Y. Zhou, and Y. Fu. 2023. “Dish-TS: A General Paradigm for Alleviating Distribution Shift in Time Series Forecasting.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 6. Vol. 37, 7522–7529.
- Z. Fan, H. Ding, A. Deoras, and T. N. Hoang. 2023. “Personalized Federated Domain Adaptation for Item-to-Item Recommendation.” In: *Uncertainty in Artificial Intelligence*. PMLR, 560–570.
- Z. Fang, J. Lu, F. Liu, J. Xuan, and G. Zhang. 2020. “Open Set Domain Adaptation: Theoretical Bound and Algorithm.” *IEEE Transactions on Neural Networks and Learning Systems*, 32, 10, 4309–4322.
- C. Fannjiang, S. Bates, A. N. Angelopoulos, J. Listgarten, and M. I. Jordan. 2022. “Conformal Prediction Under Feedback Covariate Shift for Biomolecular Design.” *Proceedings of the National Academy of Sciences*, 119, 43, e2204569119.
- K. Fatras, T. Séjourné, R. Flamary, and N. Courty. 2021. “Unbalanced minibatch Optimal Transport; applications to Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 3186–3197.
- H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P. Muller. 2018. “Transfer Learning for Time Series Classification.” In: *IEEE International Conference on Big Data*. IEEE, 1367–1376.
- L. Fei-Fei, R. Fergus, and P. Perona. 2006. “One-Shot Learning of Object Categories.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28, 4, 594–611.
- M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. 2015. “Certifying and Removing Disparate Impact.” In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 259–268.
- H. Feng, Z. You, M. Chen, T. Zhang, M. Zhu, F. Wu, C. Wu, and W. Chen. 2021. “KD3A: Unsupervised Multi-Source Decentralized Domain Adaptation via Knowledge Distillation.” In: *International Conference on Machine Learning*. PMLR, 3274–3283.
- J. Ferrando, O. B. Obeso, S. Rajamanoharan, and N. Nanda. 2025. “Do I Know This Entity? Knowledge Awareness and Hallucinations in Language Models.” In: *The Thirteenth International Conference on Learning Representations*.
- Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. March, and V. Lempitsky. 2016. “Domain-Adversarial Training of Neural Networks.” *Journal of Machine Learning Research*, 17, 59, 1–35.
- Y. Gao and Y. Cui. 2020. “Deep Transfer Learning for Reducing Health Care Disparities Arising from Biomedical Data Inequality.” *Nature Communications*, 11, 1, 5131.
- P. Germain, F. R. Bach, A. Lacoste, and S. Lacoste-Julien. 2016. “PAC-Bayesian Theory Meets Bayesian Inference.” In: *Advances in Neural Information Processing Systems* 29, 1876–1884.
- M. Gheini, X. Ren, and J. May. 2021. “Cross-Attention is All You Need: Adapting Pretrained Transformers for Machine Translation.” In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 1754–1765.
- M. Ghifary, D. Balduzzi, W. B. Kleijn, and M. Zhang. 2016. “Scatter Component Analysis: A Unified Framework for Domain Adaptation and Domain Generalization.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39, 7, 1414–1430.
- A. Ghosh, J. Chung, D. Yin, and K. Ramchandran. 2020. “An Efficient Framework for Clustered Federated Learning.” In: *Advances in Neural Information Processing Systems* 33.
- I. Gibbs and E. J. Candès. 2021. “Adaptive Conformal Inference Under Distribution Shift.” *Advances in Neural Information Processing Systems*, 34, 1660–1672.
- I. Gibbs and E. J. Candès. 2024. “Conformal Inference for Online Prediction with Arbitrary Distribution Shifts.” *Journal of Machine Learning Research*, 25, 162, 1–36.
- S. Giguere, B. Metevier, Y. Brun, P. S. Thomas, S. Niekum, and B. C. da Silva. 2022. “Fairness Guarantees under Demographic Shift.” In: *International Conference on Learning Representations*.
- J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. 2017. “Neural Message Passing for Quantum Chemistry.” In: *International Conference on Machine Learning*. PMLR, 1263–1272.
- G. B. Goh, C. Siegel, A. Vishnu, and N. Hodas. 2018. “Using Rule-Based Labels for Weak Supervised Learning: A ChemNet for Transferable Chemical Property Prediction.” In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 302–310.
- T. Gong, Y. Kim, T. Lee, S. Chottananurak, and S.-J. Lee. 2023. “SoTTA: Robust Test-Time Adaptation on Noisy Data Streams.” *Advances in Neural Information Processing Systems*, 36.
- I. J. Goodfellow, J. Shlens, and C. Szegedy. 2015. “Explaining and Harnessing Adversarial Examples.” In: *3rd International Conference on Learning Representations*.
- B. Goodman and S. Flaxman. 2017. “European Union Regulations on Algorithmic Decision-Making and a “Right to Explanation.”” *AI Magazine*, 38, 3, 50–57.
- S. Goswami, K. Kontolati, M. D. Shields, and G. E. Karniadakis. 2022. “Deep Transfer Operator Learning for Partial Differential Equations under Conditional Shift.” *Nature Machine Intelligence*, 4, 12, 1155–1164.
- H. Gouk, T. M. Hospedales, and M. Pontil. 2021. “Distance-Based Regularisation of Deep Networks for Fine-Tuning.” In: *9th International Conference on Learning Representations*.

- S. Goyal, M. Sun, A. Raghunathan, and J. Z. Kolter. 2022. “Test Time Adaptation via Conjugate Pseudo-labels.” In: *Advances in Neural Information Processing Systems* 35.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola. 2012. “A Kernel Two-Sample Test.” *Journal of Machine Learning Research*, 13, 723–773.
- T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg. 2019. “BadNets: Evaluating Backdooring Attacks on Deep Neural Networks.” *IEEE Access*, 7, 47230–47244.
- Y. Gu, X. Han, Z. Liu, and M. Huang. 2022. “PPT: Pre-trained Prompt Tuning for Few-shot Learning.” In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8410–8423.
- I. Gulrajani and D. Lopez-Paz. 2021. “In Search of Lost Domain Generalization.” In: *International Conference on Learning Representations*.
- C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. 2017. “On Calibration of Modern Neural Networks.” In: *International Conference on Machine Learning*. PMLR, 1321–1330.
- D. Guo, A. M. Rush, and Y. Kim. 2021. “Parameter-Efficient Transfer Learning with Diff Pruning.” In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4884–4896.
- K. Hambardzumyan, H. Khachatrian, and J. May. 2021. “WARP: Word-level Adversarial ReProgramming.” In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 4921–4933.
- W. L. Hamilton, Z. Ying, and J. Leskovec. 2017. “Inductive Representation Learning on Large Graphs.” In: *Advances in Neural Information Processing Systems* 30, 1024–1034.
- Z. Han, Z. Zhang, F. Wang, R. He, W. Su, X. Xi, and Y. Yin. 2023. “Discriminability and Transferability Estimation: A Bayesian Source Importance Estimation Approach for Multi-Source-Free Domain Adaptation.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 6. Vol. 37, 7811–7820.
- A. Hard, K. Rao, R. Mathews, S. Ramaswamy, F. Beaufays, S. Augenstein, H. Eichner, C. Kiddon, and D. Ramage. 2018. “Federated Learning for Mobile Keyboard Prediction.” *arXiv preprint arXiv:1811.03604*.
- E. L. Harding, J. J. Vanto, R. Clark, L. Hannah Ji, and S. C. Ainsworth. 2019. “Understanding the Scope and Impact of the California Consumer Privacy Act of 2018.” *Journal of Data Protection & Privacy*, 2, 3, 234–253.
- M. Hardt, E. Price, and N. Srebro. 2016. “Equality of Opportunity in Supervised Learning.” In: *Advances in Neural Information Processing Systems* 29, 3315–3323.
- S. Havaldar, J. Chauhan, K. Shanmugam, J. Nandy, and A. Raghuvver. 2024. “Fairness under Covariate Shift: Improving Fairness-Accuracy Tradeoff with Few Unlabeled Test Samples.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 11. Vol. 38, 12331–12339.
- S. Hayou, N. Ghosh, and B. Yu. 2024. “LoRA+: Efficient Low Rank Adaptation of Large Models.” In: *International Conference on Machine Learning*.
- H. He, O. Queen, T. Koker, C. Cuevas, T. Tsiligkaridis, and M. Zitnik. 2023. “Domain Adaptation for Time Series Under Feature and Label Shifts.” In: *International Conference on Machine Learning*. Vol. 202. PMLR, 12746–12774.
- J. He, C. Zhou, X. Ma, T. Berg-Kirkpatrick, and G. Neubig. 2022. “Towards a Unified View of Parameter-Efficient Transfer Learning.” In: *The Tenth International Conference on Learning Representations*.
- K. He, R. Girshick, and P. Dollár. 2019. “Rethinking ImageNet Pre-training.” In: *Proceedings of the IEEE/CVF international conference on computer vision*, 4918–4927.
- M. Heinzinger, A. Elnaggar, Y. Wang, C. Dallago, D. Nechaev, F. Matthes, and B. Rost. 2019. “Modeling Aspects of the Language of Life through Transfer-Learning Protein Sequences.” *BMC Bioinformatics*, 20, 1–17.
- D. Hendrycks, K. Lee, and M. Mazeika. 2019. “Using Pre-Training Can Improve Model Robustness and Uncertainty.” In: *International Conference on Machine Learning*. PMLR, 2712–2721.
- L. Hetzel, S. Böhm, N. Kilbertus, S. Günemann, M. Lotfollahi, and F. J. Theis. 2022. “Predicting Cellular Responses to Novel Drug Perturbations at a Single-Cell Resolution.” *Advances in Neural Information Processing Systems*, 35, 26711–26722.
- G. E. Hinton, O. Vinyals, and J. Dean. 2015. “Distilling the Knowledge in a Neural Network.” *arXiv preprint arXiv:1503.02531*.
- J. Hoffman, M. Mohri, and N. Zhang. 2018. “Algorithms and Theory for Multiple-Source Adaptation.” *Advances in Neural Information Processing Systems*, 31.
- J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell. 2018. “CyCADA: Cycle-Consistent Adversarial Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 1989–1998.
- W. Hou and Z. Ji. 2024. “Assessing GPT-4 for Cell Type Annotation in Single-cell RNA-seq Analysis.” *Nature Methods*, 21, 1–4.
- N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly. 2019. “Parameter-Efficient Transfer Learning for NLP.” In: *International Conference on Machine Learning*. PMLR, 2790–2799.
- J. Howard and S. Ruder. 2018. “Universal Language Model Fine-tuning for Text Classification.” In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 328–339.

- D. Hu, J. Liang, X. Wang, and C.-S. Foo. 2024. “Pseudo-Calibration: Improving Predictive Uncertainty Estimation in Unsupervised Domain Adaptation.” In: *International Conference on Machine Learning*.
- E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. 2022. “LoRA: Low-Rank Adaptation of Large Language Models.” In: *The Tenth International Conference on Learning Representations*.
- G. Hu, Y. Zhang, and Q. Yang. 2018. “CoNet: Collaborative Cross Networks for Cross-Domain Recommendation.” In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. ACM, 667–676.
- J. Hu, H. Zhong, C. Jin, and L. Wang. 2023. “Provable Sim-to-real Transfer in Continuous Domain with Partial Observations.” In: *International Conference on Learning Representations*.
- J. Hu, X. Li, G. Hu, Y. Lyu, K. Susztak, and M. Li. 2020. “Iterative Transfer Learning with Neural Network for Clustering and Cell Type Classification in Single-cell RNA-seq Analysis.” *Nature Machine Intelligence*, 2, 10, 607–618.
- A. Hua, J. Gu, Z. Xue, N. Carlini, E. Wong, and Y. Qin. 2024. “Initialization Matters for Adversarial Transfer Learning.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24831–24840.
- J. Huang, A. Gretton, K. Borgwardt, B. Schölkopf, and A. Smola. 2006. “Correcting Sample Selection Bias by Unlabeled Data.” *Advances in Neural Information Processing Systems*, 19.
- K. Huang, H. Yin, H. Huang, and W. Gao. 2024. “Towards Green AI in Fine-tuning Large Language Models via Adaptive Backpropagation.” In: *International Conference on Learning Representations*.
- L. Huang, J. Huang, Y. Rong, Q. Yang, and Y. Wei. 2022. “Frustratingly Easy Transferability Estimation.” In: *International Conference on Machine Learning*. PMLR, 9201–9225.
- S.-L. Huang, A. Makur, G. W. Wornell, L. Zheng, et al. 2024. “Universal Features for High-Dimensional Learning and Inference.” *Foundations and Trends® in Communications and Information Theory*, 21, 1-2, 1–299.
- Y. Huang, C. Du, Z. Xue, X. Chen, H. Zhao, and L. Huang. 2021. “What Makes Multi-modal Learning Better than Single (Probably).” *Advances in Neural Information Processing Systems*, 34, 10944–10956.
- E. Hubinger et al. 2024. “Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training.” *arXiv preprint arXiv:2401.05566*.
- E. Hüllermeier and W. Waegeman. 2021. “Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods.” *Machine Learning*, 110, 3, 457–506.
- X. J. Hunt, I. K. Kabul, and J. Silva. 2017. “Transfer Learning for Education Data.” In: *Proceedings of the ACM SIGKDD Conference*. Vol. 17.
- U. Hwang, J. Lee, J. Shin, and S. Yoon. 2024. “SF(DA)<sup>2</sup>: Source-free Domain Adaptation Through the Lens of Data Augmentation.” In: *The Twelfth International Conference on Learning Representations*.
- S. Ibrahim, N. Ponomareva, and R. Mazumder. 2022. “Newer is Not Always Better: Rethinking Transferability Metrics, Their Peculiarities, Stability and Performance.” In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 693–709.
- Y. Iwasawa and Y. Matsuo. 2021. “Test-Time Classifier Adjustment Module for Model-Agnostic Domain Generalization.” *Advances in Neural Information Processing Systems*, 34, 2427–2440.
- A. Jacot, C. Hongler, and F. Gabriel. 2018. “Neural Tangent Kernel: Convergence and Generalization in Neural Networks.” In: *Advances in Neural Information Processing Systems* 31, 8580–8589.
- U. Jang, J. D. Lee, and E. K. Ryu. 2024. “LoRA Training in the NTK Regime has No Spurious Local Minima.” In: *International Conference on Machine Learning*.
- P. P. Jayaraman, A. R. M. Forkan, A. Morshed, P. D. Haghighi, and Y.-B. Kang. 2020. “Healthcare 4.0: A Review of Frontiers in Digital Health.” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10, 2, e1350.
- J. Jeong and J. Shin. 2020. “Consistency Regularization for Certified Robustness of Smoothed Classifiers.” In: *Advances in Neural Information Processing Systems* 33.
- Y. Ji, X. Zhang, S. Ji, X. Luo, and T. Wang. 2018. “Model-Reuse Attacks on Deep Learning Systems.” In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 349–363.
- Z. Ji and M. Telgarsky. 2019. “Gradient Descent Aligns the Layers of Deep Linear Networks.” In: *7th International Conference on Learning Representations*.
- E. Jiang, Y. J. Zhang, and S. Koyejo. 2024. “Principled Federated Domain Adaptation: Gradient Projection and Auto-Weighting.” In: *The Twelfth International Conference on Learning Representations*.
- S. Jiang, S. Kadhe, Y. Zhou, L. Cai, and N. Baracaldo. 2023. “Forcing Generative Models to Degenerate Ones: The Power of Data Poisoning Attacks.” In: *NeurIPS 2023 Workshop on Backdoors in Deep Learning-The Good, the Bad, and the Ugly*.
- X. Jin, Y. Park, D. C. Maddix, H. Wang, and Y. Wang. 2022. “Domain Adaptation for Time Series Forecasting via Attention Sharing.” In: *International Conference on Machine Learning*. Vol. 162. PMLR, 10280–10297.
- H.-W. Jo, A. M. Koukos, V. Sitokostantinou, W.-K. Lee, and C. Kontoes. 2022. “Towards Global Crop Maps with Transfer Learning.” In: *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*.
- A. Jobin, M. Ienca, and E. Vayena. 2019. “The Global Landscape of AI Ethics Guidelines.” *Nature Machine Intelligence*, 1, 9, 389–399.
- F. D. Johansson, D. A. Sontag, and R. Ranganath. 2019. “Support and Invertibility in Domain-Invariant Representations.” In: *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 527–536.

- H. Ju, D. Li, and H. R. Zhang. 2022. “Robust Fine-Tuning of Deep Neural Networks with Hessian-based Generalization Guarantees.” In: *International Conference on Machine Learning*. PMLR, 10431–10461.
- P. Kairouz et al. 2021. “Advances and open problems in federated learning.” *Foundations and Trends® in Machine Learning*, 14, 1–2, 1–210.
- H. Kanagawa, H. Kobayashi, N. Shimizu, Y. Tagami, and T. Suzuki. 2019. “Cross-domain Recommendation via Deep Domain Adaptation.” In: *European Conference on Information Retrieval*. Springer, 20–29.
- J. Kang, J. He, R. Maciejewski, and H. Tong. 2020. “InFoRM: Individual Fairness on Graph Mining.” In: *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 379–389.
- S. Kapoor et al. 2024. “Position: On the Societal Impact of Open Foundation Models.” In: *Forty-first International Conference on Machine Learning*.
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. 2021. “Physics-Informed Machine Learning.” *Nature Reviews Physics*, 3, 6, 422–440.
- K. Kashiparekh, J. Narwariya, P. Malhotra, L. Vig, and G. Shroff. 2019. “ConvTimeNet: A Pre-trained Deep Convolutional Neural Network for Time Series Classification.” In: *International Joint Conference on Neural Networks*. IEEE, 1–8.
- A. R. Kashyap, D. Hazarika, M. Kan, and R. Zimmermann. 2021. “Domain Divergences: A Survey and Empirical Analysis.” In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1830–1849.
- E. Kasneci et al. 2023. “ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education.” *Learning and Individual Differences*, 103.
- D. Kaur, S. Uslu, K. J. Rittichier, and A. Durresi. 2023. “Trustworthy Artificial Intelligence: A Review.” *ACM Computing Surveys*, 55, 2, 39:1–39:38.
- D. Khashabi et al. 2022. “Prompt Waywardness: The Curious Case of Discretized Interpretation of Continuous Prompts.” In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3631–3643.
- T. Kim, J. Kim, Y. Tae, C. Park, J. Choi, and J. Choo. 2022. “Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift.” In: *The Tenth International Conference on Learning Representations*.
- T. N. Kipf and M. Welling. 2017. “Semi-Supervised Classification with Graph Convolutional Networks.” In: *5th International Conference on Learning Representations*.
- P. W. Koh and P. Liang. 2017. “Understanding Black-box Predictions via Influence Functions.” In: *International Conference on Machine Learning*. PMLR, 1885–1894.
- B. O. Koopman. 1931. “Hamiltonian Systems and Transformation in Hilbert Space.” *Proceedings of the National Academy of Sciences*, 17, 5, 315–318.
- S. Kornblith, J. Shlens, and Q. V. Le. 2019. “Do Better ImageNet Models Transfer Better?” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2661–2671.
- S. Krishnagopal and L. Ruiz. 2023. “Graph Neural Tangent Kernel: Convergence on Large Graphs.” In: *International Conference on Machine Learning*. Vol. 202. PMLR, 17827–17841.
- S. Kulinski and D. I. Inouye. 2022. “Towards Explaining Image-Based Distribution Shifts.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4788–4792.
- S. Kulinski and D. I. Inouye. 2023. “Towards Explaining Distribution Shifts.” In: *International Conference on Machine Learning*. PMLR, 17931–17952.
- A. Kumar, T. Ma, and P. Liang. 2020. “Understanding Self-Training for Gradual Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 5468–5479.
- J. N. Kundu, A. R. Kulkarni, S. Bhambri, D. Mehta, S. A. Kulkarni, V. Jampani, and V. B. Radhakrishnan. 2022. “Balancing Discriminability and Transferability for Source-Free Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 11710–11728.
- J. N. Kundu, N. Venkat, R. M. V., and R. V. Babu. 2020. “Universal Source-Free Domain Adaptation.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4543–4552.
- S. Kunwar. 2024. “Managing Household Waste Through Transfer Learning.” *Industrial and Domestic Waste Management*, 4, 1, 14–22.
- K. Kurita, P. Michel, and G. Neubig. 2020. “Weight Poisoning Attacks on Pretrained Models.” In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2793–2806.
- I. Kuzborskij and F. Orabona. 2017. “Fast Rates by Transferring from Auxiliary Hypotheses.” *Machine Learning*, 106, 171–195.
- I. Kuzborskij and F. Orabona. 2013. “Stability and Hypothesis Transfer Learning.” In: *International Conference on Machine Learning*, 942–950.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. 2017. “Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles.” In: *Advances in Neural Information Processing Systems* 30, 6402–6413.
- T. Le and S. Jegelka. 2023. “Limits, Approximation and Size Transferability for GNNs on Sparse Graphs via Graphops.” In: *Advances in Neural Information Processing Systems* 36.
- Y. LeCun, Y. Bengio, and G. Hinton. 2015. “Deep Learning.” *Nature*, 521, 7553, 436–444.
- J. Y. Lee, F. Dernoncourt, and P. Szolovits. 2018. “Transfer Learning for Named-Entity Recognition with Neural Networks.” In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*.

- J. Lee, D. Jung, J. Yim, and S. Yoon. 2022. “Confidence Score for Source-Free Unsupervised Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 12365–12377.
- T. Lee, G. Wollstein, C. T. Madu, A. Wronka, L. Zheng, R. Zambrano, J. S. Schuman, and J. Hu. 2023. “Reducing Ophthalmic Health Disparities Through Transfer Learning: A Novel Application to Overcome Data Inequality.” *Translational Vision Science & Technology*, 12, 12, 2–2.
- Y. Lee, A. S. Chen, F. Tajwar, A. Kumar, H. Yao, P. Liang, and C. Finn. 2023. “Surgical Fine-Tuning Improves Adaptation to Distribution Shifts.” In: *The Eleventh International Conference on Learning Representations*.
- J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman. 2018. “Distribution-Free Predictive Inference For Regression.” *Journal of the American Statistical Association*, 113, 523, 1094–1111.
- T. Lei et al.. 2023. “Conditional Adapters: Parameter-efficient Transfer Learning with Fast Inference.” In: *Advances in Neural Information Processing Systems 36*.
- B. Lester, R. Al-Rfou, and N. Constant. 2021. “The Power of Scale for Parameter-Efficient Prompt Tuning.” In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045–3059.
- R. Levie, W. Huang, L. Bucci, M. Bronstein, and G. Kutyniok. 2021. “Transferability of Spectral Graph Convolutional Neural Networks.” *Journal of Machine Learning Research*, 22, 272, 1–59.
- R. Levie, E. Isufi, and G. Kutyniok. 2019. “On the Transferability of Spectral Graph Filters.” In: *2019 13th International conference on Sampling Theory and Applications (SampTA)*. IEEE, 1–5.
- C. Li, H. Farkhoor, R. Liu, and J. Yosinski. 2018. “Measuring the Intrinsic Dimension of Objective Landscapes.” In: *6th International Conference on Learning Representations*.
- D. Li and H. R. Zhang. 2021. “Improved Regularization and Robustness for Fine-tuning in Neural Networks.” In: *Advances in Neural Information Processing Systems 34*, 27249–27262.
- L. Li, D. Song, X. Li, J. Zeng, R. Ma, and X. Qiu. 2021. “Backdoor Attacks on Pre-trained Models by Layerwise Weight Poisoning.” In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3023–3032.
- P. Li and A. Tuzhilin. 2020. “DDTCDR: Deep Dual Transfer Cross Domain Recommendation.” In: *Proceedings of the 13th International Conference on Web Search and Data Mining*, 331–339.
- T. Li, S. Hu, A. Beirami, and V. Smith. 2021. “Ditto: Fair and Robust Federated Learning Through Personalization.” In: *International Conference on Machine Learning*. PMLR, 6357–6368.
- X. L. Li and P. Liang. 2021. “Prefix-Tuning: Optimizing Continuous Prompts for Generation.” In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 4582–4597.
- X. Li and J. Bilmes. 2007. “A Bayesian Divergence Prior for Classifier Adaptation.” In: *Artificial Intelligence and Statistics*. PMLR, 275–282.
- Y. Li, X. Jia, R. Sang, Y. Zhu, B. Green, L. Wang, and B. Gong. 2021. “Ranking Neural Checkpoints.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2663–2673.
- Y. Li, T. Li, K. Chen, J. Zhang, S. Liu, W. Wang, T. Zhang, and Y. Liu. 2024. “BadEdit: Backdooring Large Language Models by Model Editing.” In: *The Twelfth International Conference on Learning Representations*.
- Z. Li and D. Hoiem. 2017. “Learning without Forgetting.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40, 12, 2935–2947.
- J. Liang, D. Hu, and J. Feng. 2020. “Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 6028–6039.
- H. Lin et al.. 2024. “Selecting Large Language Model to Fine-tune via Rectified Scaling Law.” In: *International Conference on Machine Learning*.
- Z. Lipton, Y.-X. Wang, and A. Smola. 2018. “Detecting and Correcting for Label Shift with Black Box Predictors.” In: *International Conference on Machine Learning*. PMLR, 3122–3130.
- M. Liu, Z. Fang, Z. Zhang, M. Gu, S. Zhou, X. Wang, and J. Bu. 2024. “Rethinking Propagation for Unsupervised Graph Domain Adaptation.” In: *Proceedings of the AAAI Conference on Artificial Intelligence 12*. Vol. 38, 13963–13971.
- Q. Liu and H. Xue. 2021. “Adversarial Spectral Kernel Matching for Unsupervised Time Series Domain Adaptation.” In: *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 2744–2750.
- S. Liu, J. Niles-Weed, N. Razavian, and C. Fernandez-Granda. 2020. “Early-Learning Regularization Prevents Memorization of Noisy Labels.” *Advances in Neural Information Processing Systems*, 33, 20331–20342.
- S. Liu, T. Li, Y. Feng, N. Tran, H. Zhao, Q. Qiu, and P. Li. 2023. “Structural Re-weighting Improves Graph Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 21778–21793.
- S. Liu, D. Zou, H. Zhao, and P. Li. 2024. “Pairwise Alignment Improves Graph Domain Adaptation.” In: *International Conference on Machine Learning*.
- X. Liu, K. Ji, Y. Fu, W. Tam, Z. Du, Z. Yang, and J. Tang. 2022. “P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks.” In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 61–68.
- Y. Liu, Y. Kang, C. Xing, T. Chen, and Q. Yang. 2020. “A Secure Federated Transfer Learning Framework.” *IEEE Intelligent Systems*, 35, 4, 70–82.
- Y. Liu, C. Li, J. Wang, and M. Long. 2023. “Koopa: Learning Non-stationary Time Series Dynamics with Koopman Predictors.” In: *Advances in Neural Information Processing Systems 36*.

- Y. Liu, H. Wu, J. Wang, and M. Long. 2022. “Non-stationary Transformers: Exploring the Stationarity in Time Series Forecasting.” In: *Advances in Neural Information Processing Systems 35*.
- Y. Liu, P. Kothari, B. van Delft, B. Bellot-Gurlet, T. Mordan, and A. Alahi. 2021. “TTT++: When Does Self-Supervised Test-Time Training Fail or Thrive?” In: *Advances in Neural Information Processing Systems 34*, 21808–21820.
- Z. Liu, M. Cheng, Z. Li, Z. Huang, Q. Liu, Y. Xie, and E. Chen. 2023. “Adaptive Normalization for Non-stationary Time Series Forecasting: A Temporal Slice Perspective.” In: *Advances in Neural Information Processing Systems 36*.
- Z. Liu, Y. Xu, X. Ji, and A. B. Chan. 2023. “TWINS: A Fine-Tuning Framework for Improved Transferability of Adversarial Robustness and Generalization.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 16436–16446.
- M. Long, Y. Cao, J. Wang, and M. Jordan. 2015. “Learning Transferable Features with Deep Adaptation Networks.” In: *International Conference on Machine Learning*. PMLR, 97–105.
- M. Long, Z. Cao, J. Wang, and M. I. Jordan. 2018. “Conditional Adversarial Domain Adaptation.” *Advances in Neural Information Processing Systems*, 31.
- S. Lotfi, M. A. Finzi, Y. Kuang, T. G. J. Rudner, M. Goldblum, and A. G. Wilson. 2024. “Non-Vacuous Generalization Bounds for Large Language Models.” In: *International Conference on Machine Learning*.
- M. Lotfollahi et al. 2022. “Mapping Single-cell Data to Reference Atlases by Transfer Learning.” *Nature Biotechnology*, 40, 1, 121–130.
- L. Lovász. 2012. *Large Networks and Graph Limits*. Vol. 60. American Mathematical Soc.
- J. Lu and S. Sun. 2024. “CauDiTS: Causal Disentangled Domain Adaptation of Multivariate Time Series.” In: *International Conference on Machine Learning*.
- Y. Ma, S. Chen, S. Ermon, and D. Lobell. Feb. 2024. “Transfer Learning in Environmental Remote Sensing.” *Remote Sensing of Environment*, 301, (Feb. 2024), 113924.
- W. Maddox, S. Tang, P. Moreno, A. G. Wilson, and A. Damianou. 2021. “Fast Adaptation with Linearized Neural Networks.” In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2737–2745.
- D. Madras, E. Creager, T. Pitassi, and R. S. Zemel. 2018. “Learning Adversarially Fair and Transferable Representations.” In: *International Conference on Machine Learning*. PMLR, 3381–3390.
- A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. 2018. “Towards Deep Learning Models Resistant to Adversarial Attacks.” In: *International Conference on Learning Representations*.
- R. K. Mahabadi, J. Henderson, and S. Ruder. 2021. “Compacter: Efficient Low-Rank Hypercomplex Adapter Layers.” In: *Advances in Neural Information Processing Systems 34*, 1022–1035.
- P. Malhotra, V. TV, L. Vig, P. Agarwal, and G. Shroff. 2017. “TimeNet: Pre-trained deep recurrent neural network for time series classification.” In: *25th European Symposium on Artificial Neural Networks*.
- S. Malladi, A. Wettig, D. Yu, D. Chen, and S. Arora. 2023. “A Kernel-Based View of Language Model Fine-Tuning.” In: *International Conference on Machine Learning*. PMLR, 23610–23641.
- T. Man, H. Shen, X. Jin, and X. Cheng. 2017. “Cross-Domain Recommendation: An Embedding and Mapping Approach.” In: *Proceedings of the 26th International Conference on Artificial Intelligence*, 2464–2470.
- D. Mandal, S. Deng, S. Jana, J. Wing, and D. J. Hsu. 2020. “Ensuring Fairness Beyond the Training Data.” *Advances in Neural Information Processing Systems*, 33, 18445–18456.
- Y. Mansour, M. Mohri, J. Ro, A. T. Suresh, and K. Wu. 2021. “A Theory of Multiple-Source Adaptation with Limited Target Labeled Data.” In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2332–2340.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. 2009a. “Domain Adaptation: Learning Bounds and Algorithms.” In: *The 22nd Conference on Learning Theory*.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. 2009b. “Multiple Source Adaptation and the Rényi Divergence.” In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 367–374.
- S. Maskey, R. Levie, and G. Kutyniok. 2023. “Transferability of Graph Neural Networks: An Extended Graphon Approach.” *Applied and Computational Harmonic Analysis*, 63, 48–83.
- C. Matsoukas, J. F. Haslum, M. Sorkhei, M. Söderberg, and K. Smith. 2022. “What Makes Transfer Learning Work for Medical Images: Feature Reuse & Other Factors.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9225–9234.
- B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. 2017. “Communication-Efficient Learning of Deep Networks from Decentralized Data.” In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR, 1273–1282.
- A. Mehra, B. Kailkhura, P. Chen, and J. Hamm. 2021. “Understanding the Limits of Unsupervised Domain Adaptation via Data Poisoning.” In: *Advances in Neural Information Processing Systems 34*, 17347–17359.
- E. Memmel, C. Menzen, J. Schuurmans, F. Wesel, and kim batselier. 2024. “Position: Tensor Networks are a Valuable Asset for Green AI.” In: *International Conference on Machine Learning*.
- R. Michelmor, M. Wicker, L. Laurenti, L. Cardelli, Y. Gal, and M. Kwiatkowska. 2020. “Uncertainty Quantification with Statistical Guarantees in End-to-End Autonomous Driving Control.” In: *2020 IEEE International Conference on Robotics and Automation*. IEEE, 7344–7350.

- B. Mieth, J. Hockley, N. Görnitz, M. Höhne, K.-R. Müller, A. Gutteridge, and D. Ziemek. Dec. 2019. “Using Transfer Learning from Prior Reference Knowledge to Improve the Clustering of Single-Cell RNA-Seq Data.” *Scientific Reports*, 9, (Dec. 2019), 20353.
- S. Minami, K. Fukumizu, Y. Hayashi, and R. Yoshida. 2023. “Transfer Learning with Affine Model Transformation.” In: *Thirty-seventh Conference on Neural Information Processing Systems*.
- I. Misra, A. Shrivastava, A. Gupta, and M. Hebert. 2016. “Cross-Stitch Networks for Multi-task Learning.” In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3994–4003.
- Y. Mitsuzumi, A. Kimura, and H. Kashima. 2024. “Understanding and Improving Source-free Domain Adaptation from a Theoretical Perspective.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28515–28524.
- M. Mohri and A. Muñoz Medina. 2012. “New Analysis and Algorithm for Learning with Drifting Distributions.” In: *23rd International Conference on Algorithmic Learning Theory*. Springer, 124–138.
- M. Mohri, G. Sivek, and A. T. Suresh. 2019. “Agnostic Federated Learning.” In: *International Conference on Machine Learning*, 4615–4625.
- J. Mökander, J. Schuett, H. R. Kirk, and L. Floridi. 2024. “Auditing large language models: A three-layered approach.” *AI and Ethics*, 4, 4, 1085–1115.
- J. M. Mooij, S. Magliacane, and T. Claassen. 2020. “Joint Causal Inference from Multiple Contexts.” *Journal of Machine Learning Research*, 21, 99, 1–108.
- K. Muandet, D. Balduzzi, and B. Schölkopf. 2013. “Domain Generalization via Invariant Feature Representation.” In: *International Conference on Machine Learning*, 10–18.
- D. Mukherjee, F. Petersen, M. Yurochkin, and Y. Sun. 2022. “Domain Adaptation meets Individual Fairness. And they get along.” In: *Advances in Neural Information Processing Systems 35*.
- A. Müller. 1997. “Integral Probability Metrics and Their Generating Classes of Functions.” *Advances in Applied Probability*, 29, 2, 429–443.
- M. P. Naeini, G. Cooper, and M. Hauskrecht. 2015. “Obtaining Well Calibrated Probabilities Using Bayesian Binning.” In: *Proceedings of the AAAI Conference on Artificial Intelligence 1*. Vol. 29.
- L. F. Nern, H. Raj, M. A. Georgi, and Y. Sharma. 2023. “On Transfer of Adversarial Robustness from Pretraining to Downstream Tasks.” In: *Advances in Neural Information Processing Systems 36*.
- B. Neyshabur, H. Sedghi, and C. Zhang. 2020. “What is Being Transferred in Transfer Learning?” In: *Advances in Neural Information Processing Systems 33*.
- A. T. Nguyen, P. Torr, and S. N. Lim. 2022. “FedSR: A Simple and Effective Domain Generalization Method for Federated Learning.” *Advances in Neural Information Processing Systems*, 35, 38831–38843.
- A. T. Nguyen, T. Tran, Y. Gal, P. Torr, and A. G. Baydin. 2022. “KL Guided Domain Adaptation.” In: *International Conference on Learning Representations*.
- C. N. Nguyen, P. Tran, L. S. T. Ho, V. C. Dinh, A. T. Tran, T. Hassner, and C. V. Nguyen. 2023. “Simple Transferability Estimation for Regression Tasks.” In: *Uncertainty in Artificial Intelligence*. PMLR, 1510–1521.
- C. V. Nguyen, T. Hassner, M. W. Seeger, and C. Archambeau. 2020. “LEEP: A New Measure to Evaluate Transferability of Learned Representations.” In: *International Conference on Machine Learning*. PMLR, 7294–7305.
- X. Nguyen, M. J. Wainwright, and M. I. Jordan. 2010. “Estimating Divergence Functionals and the Likelihood Ratio by Convex Risk Minimization.” *IEEE Transactions on Information Theory*, 56, 11, 5847–5861.
- N. T. Nikolinakos. 2023. “Ethical Principles for Trustworthy AI.” In: *EU Policy and Legal Framework for Artificial Intelligence, Robotics and Related Technologies-The AI Act*. Springer, 101–166.
- F. F. Niloy, S. M. Ahmed, D. S. Raychaudhuri, S. Oymak, and A. K. Roy-Chowdhury. 2024. “Effective Restoration of Source Knowledge in Continual Test Time Adaptation.” In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2091–2100.
- S. Niu, J. Wu, Y. Zhang, Y. Chen, S. Zheng, P. Zhao, and M. Tan. 2022. “Efficient Test-Time Model Adaptation without Forgetting.” In: *International Conference on Machine Learning*. PMLR, 16888–16905.
- S. Niu, J. Wu, Y. Zhang, Z. Wen, Y. Chen, P. Zhao, and M. Tan. 2023. “Towards Stable Test-time Adaptation in Dynamic Wild World.” In: *The Eleventh International Conference on Learning Representations*.
- Z. Niu, M. Anitescu, and J. Chen. 2023. “Graph Neural Network-Inspired Kernels for Gaussian Processes in Semi-Supervised Learning.” In: *The Eleventh International Conference on Learning Representations*.
- L. Oneto, M. Donini, G. Luise, C. Ciliberto, A. Maurer, and M. Pontil. 2020. “Exploiting MMD and Sinkhorn Divergences for Fair and Transferable Representation Learning.” *Advances in Neural Information Processing Systems*, 33, 15360–15370.
- L. Oneto, M. Donini, M. Pontil, and A. Maurer. 2020. “Learning Fair and Transferable Representations with Theoretical Guarantees.” In: *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, 30–39.
- OpenAI. 2023. “GPT-4 Technical Report.” *arXiv preprint arXiv:2303.08774*.
- OpenAI et al.. 2019. “Solving Rubik’s Cube with a Robot Hand.” *arXiv preprint arXiv:1910.07113*.
- F. Ott, D. Rügamer, L. Heublein, B. Bischl, and C. Mutschler. 2022. “Domain Adaptation for Time-Series Classification to Mitigate Covariate Shift.” In: *Proceedings of the 30th ACM International Conference on Multimedia*. ACM, 5934–5943.

- Y. Özyurt, S. Feuerriegel, and C. Zhang. 2023. “Contrastive Learning for Unsupervised Domain Adaptation of Time Series.” In: *The Eleventh International Conference on Learning Representations*.
- S. J. Pan and Q. Yang. 2010. “A Survey on Transfer Learning.” *IEEE Transactions on Knowledge and Data Engineering*, 22, 10, 1345–1359.
- M. Pándy, A. Agostinelli, J. R. R. Uijlings, V. Ferrari, and T. Mensink. 2022. “Transferability Estimation using Bhattacharyya Class Separability.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 9162–9172.
- S. Park, O. Bastani, J. Weimer, and I. Lee. 2020. “Calibrated Prediction with Covariate Shift via Unsupervised Domain Adaptation.” In: *The 23rd International Conference on Artificial Intelligence and Statistics*. PMLR, 3219–3229.
- S. Park, E. Dobriban, I. Lee, and O. Bastani. 2022. “PAC Prediction Sets Under Covariate Shift.” In: *International Conference on Learning Representations*.
- N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis. 2020. “Deep Adaptive Input Normalization for Time Series Forecasting.” *IEEE Transactions on Neural Networks and Learning Systems*, 31, 9, 3760–3765.
- A. Paszke et al. 2019. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” In: *Advances in Neural Information Processing Systems* 32, 8024–8035.
- J. Pei, R. C. Fernandez, and X. Yu. 2023. “Data and AI Model Markets: Opportunities for Data and Model Sharing, Discovery, and Integration.” *Proceedings of the VLDB Endowment*, 16, 12, 3872–3873.
- X. Peng, Q. Bai, X. Xia, Z. Huang, K. Saenko, and B. Wang. 2019. “Moment Matching for Multi-source Domain Adaptation.” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1406–1415.
- X. Peng, Z. Huang, Y. Zhu, and K. Saenko. 2020. “Federated Adversarial Domain Adaptation.” In: *International Conference on Learning Representations*.
- X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel. 2018. “Sim-to-Real Transfer of Robotic Control with Dynamics Randomization.” In: *2018 IEEE International Conference on Robotics and Automation*. IEEE, 3803–3810.
- D. Pessach and E. Shmueli. 2023. “Algorithmic Fairness.” In: *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*. Springer, 867–886.
- A. Petruzzelli, C. Musto, L. Laraspata, I. Rinaldi, M. de Gemmis, P. Lops, and G. Semeraro. 2024. “Instructing and Prompting Large Language Models for Explainable Cross-domain Recommendations.” In: *Proceedings of the 18th ACM Conference on Recommender Systems*, 298–308.
- J. Pfeiffer, A. Kamath, A. Rücklé, K. Cho, and I. Gurevych. 2021. “AdapterFusion: Non-Destructive Task Composition for Transfer Learning.” In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 487–503.
- T.-H. Pham, X. Zhang, and P. Zhang. 2023. “Fairness and Accuracy under Domain Generalization.” In: *International Conference on Learning Representations*.
- L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta. 2017. “Robust Adversarial Reinforcement Learning.” In: *International Conference on Machine Learning*. PMLR, 2817–2826.
- M. Ploner and A. Akbik. 2024. “Parameter-Efficient Fine-Tuning: Is There An Optimal Subset of Parameters to Tune?” In: *Findings of the Association for Computational Linguistics: EAACL 2024*.
- A. Podkopaev and A. Ramdas. 2021. “Distribution-free Uncertainty Quantification for Classification under Label Shift.” In: *Uncertainty in Artificial Intelligence*. PMLR, 844–853.
- A. Prasad, P. Hase, X. Zhou, and M. Bansal. 2023. “GrIPS: Gradient-free, Edit-based Instruction Search for Prompting Large Language Models.” In: *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 3827–3846.
- J. C. Principe. 2010. *Information Theoretic Learning: Renyi’s Entropy and Kernel Perspectives*. Springer Science & Business Media.
- D. Prinster, A. Liu, and S. Saria. 2022. “JAWS: Auditing Predictive Uncertainty Under Covariate Shift.” *Advances in Neural Information Processing Systems*, 35, 35907–35920.
- D. Prinster, S. Saria, and A. Liu. 2023. “JAWS-X: Addressing Efficiency Bottlenecks of Conformal Prediction Under Standard and Feedback Covariate Shift.” In: *International Conference on Machine Learning*. PMLR, 28167–28190.
- S. Purushotham, W. Carvalho, T. Nilanon, and Y. Liu. 2017. “Variational Recurrent Adversarial Deep Domain Adaptation.” In: *International Conference on Learning Representations*.
- F. Qi, Y. Chen, M. Li, Y. Yao, Z. Liu, and M. Sun. 2021. “ONION: A Simple and Effective Defense Against Textual Backdoor Attacks.” In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9558–9566.
- G. Qin and J. Eisner. 2021. “Learning How to Ask: Querying LMs with Mixtures of Soft Prompts.” In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 5203–5212.
- X. Qiu, T. Parcollet, J. Fernandez-Marques, P. P. Gusmao, Y. Gao, D. J. Beutel, T. Topal, A. Mathur, and N. D. Lane. 2023. “A First Look into the Carbon Footprint of Federated Learning.” *Journal of Machine Learning Research*, 24, 129, 1–23.
- A. Radford et al. 2021. “Learning Transferable Visual Models From Natural Language Supervision.” In: *International Conference on Machine Learning*. PMLR, 8748–8763.
- C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. 2020. “Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer.” *Journal of Machine Learning Research*, 21, 140:1–140:67.

- M. Ragab, E. Eldele, W. L. Tan, C.-S. Foo, Z. Chen, M. Wu, C.-K. Kwok, and X. Li. 2023. “ADATIME: A Benchmarking Suite for Domain Adaptation on Time Series Data.” *ACM Transactions on Knowledge Discovery from Data*, 17, 8, 1–18.
- M. Raghu, C. Zhang, J. Kleinberg, and S. Bengio. 2019. “Transfusion: Understanding Transfer Learning for Medical Imaging.” *Advances in Neural Information Processing Systems*, 32.
- R. Raina, A. J. Battle, H. Lee, B. Packer, and A. Y. Ng. 2007. “Self-Taught Learning: Transfer Learning from Unlabeled Data.” In: *International Conference on Machine Learning*. ACM, 759–766.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. 2019. “Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations.” *Journal of Computational Physics*, 378, 686–707.
- I. D. Raji, A. Smart, R. N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, and P. Barnes. 2020. “Closing the AI Accountability Gap: Defining An End-to-End Framework for Internal Algorithmic Auditing.” In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.
- R. Rao, N. Bhattacharya, N. Thomas, Y. Duan, P. Chen, J. Canny, P. Abbeel, and Y. Song. 2019. “Evaluating Protein Transfer Learning with TAPE.” *Advances in Neural Information Processing Systems*, 32.
- A. Razdaibiedina, Y. Mao, M. Khabsa, M. Lewis, R. Hou, J. Ba, and A. Almahairi. 2023. “Residual Prompt Tuning: Improving Prompt Tuning with Residual Reparameterization.” In: *Findings of the Association for Computational Linguistics: ACL 2023*, 6740–6757.
- S. Rebuffi, H. Bilen, and A. Vedaldi. 2017. “Learning Multiple Visual Domains with Residual Adapters.” In: *Advances in Neural Information Processing Systems* 30, 506–516.
- I. Redko, A. Habrard, and M. Sebban. 2017. “Theoretical Analysis of Domain Adaptation with Optimal Transport.” In: *Machine Learning and Knowledge Discovery in Databases: European Conference*. Springer, 737–753.
- I. Redko, E. Morvant, A. Habrard, M. Sebban, and Y. Bennani. 2019. *Advances in domain adaptation theory*. Elsevier.
- Y. Ren and D. J. Sutherland. 2025. “Learning Dynamics of LLM Finetuning.” In: *The Thirteenth International Conference on Learning Representations*.
- A. Rényi. 1961. “On Measures of Entropy and Information.” In: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. Vol. 4. University of California Press, 547–562.
- A. Rezaei, A. Liu, O. Memarrast, and B. D. Ziebart. 2021. “Robust Fairness under Covariate Shift.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 11. Vol. 35, 9419–9427.
- S. Rezaei and X. Liu. 2020. “A Target-Agnostic Attack on Deep Models: Exploiting Security Vulnerabilities of Transfer Learning.” In: *International Conference on Learning Representations*.
- M. T. Ribeiro, S. Singh, and C. Guestrin. 2016. “‘Why Should I Trust You?’: Explaining the Predictions of Any Classifier.” In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1135–1144.
- Y. Roh, K. Lee, S. E. Whang, and C. Suh. 2023. “Improving Fair Training under Correlation Shifts.” In: *International Conference on Machine Learning*. PMLR, 29179–29209.
- Y. Romano, M. Sesia, and E. Candes. 2020. “Classification with Valid and Adaptive Coverage.” *Advances in Neural Information Processing Systems*, 33, 3581–3591.
- M. T. Rosenstein, Z. Marx, L. P. Kaelbling, and T. G. Dietterich. 2005. “To Transfer or Not To Transfer.” In: *NIPS 2005 Workshop on Transfer Learning* 3. Vol. 898, 4.
- A. Rücklé, G. Geigle, M. Glockner, T. Beck, J. Pfeiffer, N. Reimers, and I. Gurevych. 2021. “AdapterDrop: On the Efficiency of Adapters in Transformers.” In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 7930–7946.
- L. Ruiz, L. F. O. Chamon, and A. Ribeiro. 2020. “Graphon Neural Networks and the Transferability of Graph Neural Networks.” In: *Advances in Neural Information Processing Systems* 33.
- L. Ruiz, F. Gama, and A. Ribeiro. 2021. “Graph Neural Networks: Architectures, Stability, and Transferability.” *Proceedings of the IEEE*, 109, 5, 660–682.
- A. Ruoss, M. Balunovic, M. Fischer, and M. Vechev. 2020. “Learning Certified Individually Fair Representations.” *Advances in Neural Information Processing Systems*, 33, 7584–7596.
- A. A. Rusu, M. Večerík, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell. 2017. “Sim-to-Real Robot Learning from Pixels with Progressive Nets.” In: *Conference on Robot Learning*. PMLR, 262–270.
- K. Saito, K. Watanabe, Y. Ushiku, and T. Harada. 2018. “Maximum Classifier Discrepancy for Unsupervised Domain Adaptation.” In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3723–3732.
- H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, and A. Madry. 2020. “Do Adversarially Robust ImageNet Models Transfer Better?” In: *Advances in Neural Information Processing Systems* 33.
- N. Sama, E. David, S. Rossetti, A. Antona, B. Franchetti, and F. Pirri. 2023. “A New Large Dataset and a Transfer Learning Methodology for Plant Phenotyping in Vertical Farms.” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 540–551.
- A. Samra, E. Frolov, A. Vasilev, A. Grigorevskiy, and A. Vakhruhev. 2024. “Cross-Domain Latent Factors Sharing via Implicit Matrix Factorization.” In: *Proceedings of the 18th ACM Conference on Recommender Systems*, 309–317.

- F. Sattler, K.-R. Müller, and W. Samek. 2020. “Clustered Federated Learning: Model-Agnostic Distributed Multitask Optimization Under Privacy Constraints.” *IEEE Transactions on Neural Networks and Learning Systems*, 32, 8, 3710–3722.
- F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. 2009. “The Graph Neural Network Model.” *IEEE Transactions on Neural Networks*, 20, 1, 61–80.
- R. Schmucker and T. M. Mitchell. 2022. “Transferable Student Performance Modeling for Intelligent Tutoring Systems.” In: *Proceedings of the 30th International Conference on Computers in Education*, 13–23.
- J. Schrouff et al. 2022. “Diagnosing Failures of Fairness Transfer across Distribution Shift in Real-World Medical Settings.” In: *Advances in Neural Information Processing Systems* 35.
- C. Schumann, X. Wang, A. Beutel, J. Chen, H. Qian, and E. H. Chi. 2019. “Transfer of Machine Learning Fairness across Domains.” In: *NeurIPS Joint Workshop on AI for Social Good*.
- R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni. 2020. “Green AI.” *Communications of the ACM*, 63, 12, 54–63.
- R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. 2017. “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization.” In: *Proceedings of the IEEE International Conference on Computer Vision*, 618–626.
- G. Shachaf, A. Brutzkus, and A. Globerson. 2021. “A Theoretical Analysis of Fine-tuning with Linear Teachers.” In: *Advances in Neural Information Processing Systems* 34, 15382–15394.
- A. Shafahi, P. Saadatpanah, C. Zhu, A. Ghiasi, C. Studer, D. W. Jacobs, and T. Goldstein. 2020. “Adversarially Robust Transfer Learning.” In: *8th International Conference on Learning Representations*.
- W. Shao, X. Zhao, Y. Ge, Z. Zhang, L. Yang, X. Wang, Y. Shan, and P. Luo. 2022. “Not All Models Are Equal: Predicting Model Transferability in a Self-challenging Fisher Space.” In: *European Conference on Computer Vision*. Springer, 286–302.
- J. Shen, Y. Qu, W. Zhang, and Y. Yu. 2018. “Wasserstein Distance Guided Representation Learning for Domain Adaptation.” In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. AAAI Press, 4058–4065.
- J. Shen, L. Li, L. M. Dery, C. Staten, M. Khodak, G. Neubig, and A. Talwalkar. 2023. “Cross-Modal Fine-Tuning: Align then Refine.” In: *International Conference on Machine Learning*. PMLR, 31030–31056.
- L. Shen, S. Ji, X. Zhang, J. Li, J. Chen, J. Shi, C. Fang, J. Yin, and T. Wang. 2021. “Backdoor Pre-trained Models Can Transfer to All.” In: *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 3141–3158.
- M. Shen, Y. Bu, and G. W. Wornell. 2023. “On Balancing Bias and Variance in Unsupervised Multi-Source-Free Domain Adaptation.” In: *International Conference on Machine Learning*. PMLR, 30976–30991.
- N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. 2011. “Weisfeiler-Lehman Graph Kernels.” *Journal of Machine Learning Research*, 12, 2539–2561.
- W. Shi, X. Han, H. Gonen, A. Holtzman, Y. Tsvetkov, and L. Zettlemoyer. 2023. “Toward Human Readable Prompt Tuning: Kubrick’s The Shining is a good movie, and a good prompt too?” In: *Findings of the Association for Computational Linguistics: EMNLP 2023*, 10994–11005.
- H. Shimodaira. 2000. “Improving Predictive Inference under Covariate Shift by Weighting the Log-Likelihood Function.” *Journal of Statistical Planning and Inference*, 90, 2, 227–244.
- T. Shin, Y. Razeghi, R. L. L. IV, E. Wallace, and S. Singh. 2020. “AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts.” In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 4222–4235.
- M. Shu, J. Wang, C. Zhu, J. Geiping, C. Xiao, and T. Goldstein. 2023. “On the Exploitability of Instruction Tuning.” In: *Advances in Neural Information Processing Systems* 36.
- I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal. 2024. “AI models collapse when trained on recursively generated data.” *Nature*, 631, 8022, 755–759.
- W. Si, S. Park, I. Lee, E. Dobriban, and O. Bastani. 2024. “PAC Prediction Sets Under Label Shift.” In: *International Conference on Learning Representations*.
- H. Singh, R. Singh, V. Mhasawade, and R. Chunara. 2021. “Fairness Violations and Mitigation under Covariate Shift.” In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 3–13.
- D. Slack, S. A. Friedler, and E. Givental. 2020. “Fairness Warnings and Fair-MAML: Learning Fairly with Minimal Data.” In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 200–209.
- V. Smith, C. Chiang, M. Sanjabi, and A. Talwalkar. 2017. “Federated Multi-Task Learning.” In: *Advances in Neural Information Processing Systems* 30, 4424–4434.
- J. Snoek, Y. Ovadia, E. Fertig, B. Lakshminarayanan, S. Nowozin, D. Sculley, J. V. Dillon, J. Ren, and Z. Nado. 2019. “Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty under Dataset Shift.” In: *Advances in Neural Information Processing Systems* 32, 13969–13980.
- R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. 2013. “Zero-Shot Learning Through Cross-Modal Transfer.” *Advances in Neural Information Processing Systems*, 26.
- N. Srebro, K. Sridharan, and A. Tewari. 2010. “Optimistic Rates for Learning with a Smooth Loss.” *arXiv Preprint*, arXiv:1009.3896.
- B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet. 2010. “Hilbert Space Embeddings and Metrics on Probability Measures.” *The Journal of Machine Learning Research*, 11, 1517–1561.

- A. Stein, Y. Wu, E. Wong, and M. Naik. 2023. "Rectifying Group Irregularities in Explanations for Distribution Shift." In: *XAI in Action: Past, Present, and Future Applications*.
- Y. Su et al. 2022. "On Transferability of Prompt Tuning for Natural Language Processing." In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3949–3969.
- S. Subramanian, P. Harrington, K. Keutzer, W. Bhimji, D. Morozov, M. W. Mahoney, and A. Gholami. 2023. "Towards Foundation Models for Scientific Machine Learning: Characterizing Scaling and Transfer Behavior." *Advances in Neural Information Processing Systems*, 36.
- B. Sun and K. Saenko. 2016. "Deep CORAL: Correlation Alignment for Deep Domain Adaptation." In: *Computer Vision - ECCV 2016 Workshops - Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part III*. Vol. 9915, 443–450.
- Y. Sun, X. Wang, Z. Liu, J. Miller, A. Efros, and M. Hardt. 2020. "Test-Time Training with Self-Supervision for Generalization under Distribution Shifts." In: *International Conference on Machine Learning*. PMLR, 9229–9248.
- Y. Sung, V. Nair, and C. Raffel. 2021. "Training Neural Networks with Fixed Sparse Masks." In: *Advances in Neural Information Processing Systems* 34, 24193–24205.
- V. Swamy, M. Marras, and T. Käser. 2022. "Meta Transfer Learning for Early Success Prediction in MOOCs." In: *Proceedings of the Ninth ACM Conference on Learning @ Scale*. ACM, 121–132.
- C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. 2014. "Intriguing Properties of Neural Networks." In: *2nd International Conference on Learning Representations*.
- A. K. Tanwani. 2020. "DIRL: Domain-Invariant Representation Learning for Sim-to-Real Transfer." In: *4th Conference on Robot Learning*. PMLR, 1558–1571.
- Y. Tay et al. 2022. "Scale Efficiently: Insights from Pretraining and Finetuning Transformers." In: *International Conference on Learning Representations*.
- M. Terzi, A. Achille, M. Maggipinto, and G. A. Susto. 2021. "Adversarial Training Reduces Information and Improves Transferability." In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2674–2682.
- K. Thenmozhi and U. S. Reddy. 2019. "Crop Pest Classification Based on Deep Convolutional Neural Network and Transfer Learning." *Computers and Electronics in Agriculture*, 164, 104906.
- C. V. Theodoris et al. 2023. "Transfer Learning Enables Predictions in Network Biology." *Nature*, 618, 616–624.
- R. J. Tibshirani, R. Foygel Barber, E. Candès, and A. Ramdas. 2019. "Conformal Prediction Under Covariate Shift." *Advances in Neural Information Processing Systems*, 32.
- J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. 2017. "Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World." In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 23–30.
- M. Toseef, X. Li, and K.-C. Wong. 2022. "Reducing Healthcare Disparities using Multiple Multiethnic Data Distributions with Fine-tuning of Transfer Learning." *Briefings in Bioinformatics*, 23, 3, bbac078.
- H. Touvron et al. 2023. "LLaMA: Open and Efficient Foundation Language Models." *arXiv preprint arXiv:2302.13971*.
- A. T. Tran, C. V. Nguyen, and T. Hassner. 2019. "Transferability and Hardness of Supervised Classification Tasks." In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 1395–1405.
- N. Tripuraneni, C. Jin, and M. Jordan. 2021. "Provable Meta-Learning of Linear Representations." In: *International Conference on Machine Learning*. PMLR, 10434–10443.
- N. Tripuraneni, M. I. Jordan, and C. Jin. 2020. "On the Theory of Transfer Learning: The Importance of Task Diversity." In: *Advances in Neural Information Processing Systems* 33.
- D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry. 2019. "Robustness May Be at Odds with Accuracy." In: *7th International Conference on Learning Representations*.
- E. Tzeng, C. Devin, J. Hoffman, C. Finn, P. Abbeel, S. Levine, K. Saenko, and T. Darrell. 2020. "Adapting Deep Visuomotor Representations with Weak Pairwise Constraints." In: *Algorithmic Foundations of Robotics XII: Proceedings of the Twelfth Workshop on the Algorithmic Foundations of Robotics*. Springer, 688–703.
- E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. 2017. "Adversarial Discriminative Domain Adaptation." In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7167–7176.
- E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell. 2014. "Deep Domain Confusion: Maximizing for Domain Invariance." *arXiv preprint arXiv:1412.3474*.
- D. Ulyanov, A. Vedaldi, and V. S. Lempitsky. 2016. "Instance Normalization: The Missing Ingredient for Fast Stylization." *CoRR*, abs/1607.08022.
- F. Utrera, E. Kravitz, N. B. Erichson, R. Khanna, and M. W. Mahoney. 2021. "Adversarially-Trained Deep Nets Transfer Better: Illustration on Image Classification." In: *9th International Conference on Learning Representations*.
- P. Vaishnavi, K. Eykholt, and A. Rahmati. 2024. "A Study of the Effects of Transfer Learning on Adversarial Robustness." *Transactions on Machine Learning Research*, 2024.
- P. Vaishnavi, K. Eykholt, and A. Rahmati. 2022. "Transferring Adversarial Robustness Through Robust Representation Matching." In: *31st USENIX Security Symposium*. USENIX Association, 2083–2098.

- M. Valipour, M. Rezagholizadeh, I. Kobayzev, and A. Ghodsi. 2023. “DyLoRA: Parameter-Efficient Tuning of Pre-trained Models using Dynamic Search-Free Low-Rank Adaptation.” In: *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 3266–3279.
- A. Vanzo, S. Pal Chowdhury, and M. Sachan. 2025. “GPT-4 as a Homework Tutor Can Improve Student Engagement and Learning Outcomes.” In: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 31119–31136.
- K. R. Varshney. 2022. *Trustworthy Machine Learning*. Independently Published, Chappaqua, NY, USA.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. “Attention is All you Need.” In: *Advances in Neural Information Processing Systems 30*, 5998–6008.
- T. Vu, B. Lester, N. Constant, R. Al-Rfou’, and D. Cer. 2022. “SPoT: Better Frozen Model Adaptation through Soft Prompt Transfer.” In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5039–5059.
- A. Wan, E. Wallace, S. Shen, and D. Klein. 2023. “Poisoning Language Models During Instruction Tuning.” In: *International Conference on Machine Learning*. PMLR, 35413–35425.
- L. Wan, W. Zhou, Y. He, T. C. Wanger, and H. Cen. 2022. “Combining Transfer Learning and Hyperspectral Reflectance Analysis to Assess Leaf Nitrogen Concentration across Different Plant Species Datasets.” *Remote Sensing of Environment*, 269, 112826.
- B. Wang, Y. Yao, B. Viswanath, H. Zheng, and B. Y. Zhao. 2018. “With Great Training Comes Great Vulnerability: Practical Attacks against Transfer Learning.” In: *27th USENIX Security Symposium (USENIX Security 18)*, 1281–1297.
- D. Wang, E. Shelhamer, S. Liu, B. A. Olshausen, and T. Darrell. 2021. “Tent: Fully Test-Time Adaptation by Entropy Minimization.” In: *9th International Conference on Learning Representations*.
- H. Wang et al.. 2024. “EvoluNet: Advancing Dynamic Non-IID Transfer Learning on Graphs.” In: *International Conference on Machine Learning*.
- J.-K. Wang and A. Wibisono. 2023. “Towards Understanding GD with Hard and Conjugate Pseudo-labels for Test-Time Adaptation.” In: *The Eleventh International Conference on Learning Representations*.
- Q. Wang, O. Fink, L. Van Gool, and D. Dai. 2022. “Continual Test-Time Domain Adaptation.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7201–7211.
- R. Wang, Y. Dong, S. Ö. Arik, and R. Yu. 2023. “Koopman Neural Operator Forecaster for Time-series with Temporal Distributional Shifts.” In: *The Eleventh International Conference on Learning Representations*.
- S. Wang et al.. Feb. 2023. “Airborne Hyperspectral Imaging of Cover Crops through Radiative Transfer Process-Guided Machine Learning.” *Remote Sensing of Environment*, 285, (Feb. 2023), 113386.
- X. Wang, M. Long, J. Wang, and M. I. Jordan. 2020. “Transferable Calibration with Lower Bias and Variance in Domain Adaptation.” In: *Advances in Neural Information Processing Systems 33*.
- Y. Wang, J. Chauhan, W. Wang, and C. Hsieh. 2023. “Universality and Limitations of Prompt Tuning.” In: *Advances in Neural Information Processing Systems 36*.
- Y. Wang, Y. Chen, K. Jamieson, and S. S. Du. 2023. “Improved Active Multi-Task Representation Learning via Lasso.” In: *International Conference on Machine Learning*. PMLR, 35548–35578.
- Y. Wang and R. Arora. 2024. “Adversarially Robust Hypothesis Transfer Learning.” In: *International Conference on Machine Learning*.
- Z. Wang, R. Panda, L. Karlinsky, R. Feris, H. Sun, and Y. Kim. 2023. “Multitask Prompt Tuning Enables Parameter-Efficient Transfer Learning.” In: *The Eleventh International Conference on Learning Representations*.
- Z. Wang, Z. Dai, B. Póczos, and J. G. Carbonell. 2019. “Characterizing and Avoiding Negative Transfer.” In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11293–11302.
- A. Watkins, E. Ullah, T. Nguyen-Tang, and R. Arora. 2023. “Optimistic Rates for Multi-Task Representation Learning.” In: *Advances in Neural Information Processing Systems 36*.
- C. Wei, K. Shen, Y. Chen, and T. Ma. 2021. “Theoretical Analysis of Self-Training with Deep Networks on Unlabeled Data.” In: *International Conference on Learning Representations*.
- C. Wei, S. M. Xie, and T. Ma. 2021. “Why Do Pretrained Language Models Help in Downstream Tasks? An Analysis of Head and Prompt Tuning.” In: *Advances in Neural Information Processing Systems 34*, 16158–16170.
- B. Weisfeiler and A. Lehman. 1968. “A Reduction of a Graph to a Canonical Form and an Algebra Arising during This Reduction.” *Nauchno-Tekhnicheskaya Informatsia*, 2, 9, 12–16.
- Y. Wen, N. Jain, J. Kirchenbauer, M. Goldblum, J. Geiping, and T. Goldstein. 2023. “Hard Prompts Made Easy: Gradient-Based Discrete Optimization for Prompt Tuning and Discovery.” In: *Advances in Neural Information Processing Systems 36*.
- M. Wicker, V. Piratla, and A. Weller. 2023. “Certification of Distributional Individual Fairness.” In: *Advances in Neural Information Processing Systems 36*.
- M. Wieringa. 2020. “What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability.” In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 1–18.
- G. Wiese, D. Weissenborn, and M. L. Neves. 2017. “Neural Domain Adaptation for Biomedical Question Answering.” In: *Proceedings of the 21st Conference on Computational Natural Language Learning*, 281–289.

- O. Wiles, S. Goyal, F. Stimberg, S.-A. Rebuffi, I. Ktena, K. D. Dvijotham, and A. T. Cemgil. 2022. "A Fine-Grained Analysis on Distribution Shift." In: *International Conference on Learning Representations*.
- G. Wilson, J. R. Doppa, and D. J. Cook. 2023. "CALDA: Improving Multi-Source Time Series Domain Adaptation With Contrastive Adversarial Learning." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 12, 14208–14221.
- G. Wilson, J. R. Doppa, and D. J. Cook. 2020. "Multi-Source Deep Domain Adaptation with Weak Supervision for Time-Series Sensor Data." In: *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 1768–1778.
- T. Wolf et al. 2020. "Transformers: State-of-the-Art Natural Language Processing." In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45.
- J. Wu, L. Ainsworth, A. Leakey, H. Wang, and J. He. 2023. "Graph-Structured Gaussian Processes for Transferable Graph Learning." In: *Advances in Neural Information Processing Systems* 36.
- J. Wu, W. Bao, E. A. Ainsworth, and J. He. 2023. "Personalized Federated Learning with Parameter Propagation." In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2594–2605.
- J. Wu and J. He. 2023a. "A Unified Framework for Adversarial Attacks on Multi-Source Domain Adaptation." *IEEE Transactions on Knowledge and Data Engineering*, 35, 11, 11039–11050.
- J. Wu and J. He. 2022. "Dynamic Transfer Learning with Progressive Meta-Task Scheduler." *Frontiers Big Data*, 5.
- J. Wu and J. He. 2021. "Indirect Invisible Poisoning Attacks on Domain Adaptation." In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM, 1852–1862.
- J. Wu and J. He. 2023b. "Trustworthy transfer learning: Transferability and trustworthiness." In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5829–5830.
- J. Wu, J. He, and E. A. Ainsworth. 2023. "Non-IID Transfer Learning on Graphs." In: *Proceedings of the AAAI Conference on Artificial Intelligence* 9. Vol. 37, 10342–10350.
- J. Wu, J. He, and H. Tong. 2024. "Distributional Network of Networks for Modeling Data Heterogeneity." In: *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3379–3390.
- J. Wu, J. He, S. Wang, K. Guan, and E. A. Ainsworth. 2022. "Distribution-Informed Neural Networks for Domain Adaptation Regression." In: *Advances in Neural Information Processing Systems* 35.
- M. Wu, S. Pan, C. Zhou, X. Chang, and X. Zhu. 2020. "Unsupervised Domain Adaptive Graph Convolutional Networks." In: *Proceedings of the Web Conference 2020*, 1457–1467.
- Y. Wu, E. Winston, D. Kaushik, and Z. C. Lipton. 2019. "Domain Adaptation with Asymmetrically-Relaxed Distribution Alignment." In: *International Conference on Machine Learning*. PMLR, 6872–6881.
- Z. Wu, Y. Wu, and L. Mou. 2024. "Zero-Shot Continuous Prompt Transfer: Generalizing Task Semantics Across Language Models." In: *The Twelfth International Conference on Learning Representations*.
- Z. Xi, T. Du, C. Li, R. Pang, S. Ji, J. Chen, F. Ma, and T. Wang. 2023. "Defending Pre-trained Language Models as Few-shot Learners against Backdoor Attacks." In: *Advances in Neural Information Processing Systems* 36.
- J. Xu, M. Ma, F. Wang, C. Xiao, and M. Chen. 2024. "Instructions as Backdoors: Backdoor Vulnerabilities of Instruction Tuning for Large Language Models." In: *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 3111–3126.
- J. Xu and J. Zhang. 2024. "Random Masking Finds Winning Tickets for Parameter Efficient Fine-tuning." In: *International Conference on Machine Learning*.
- K. Xu, W. Hu, J. Leskovec, and S. Jegelka. 2019. "How Powerful are Graph Neural Networks?" In: *7th International Conference on Learning Representations*.
- R. Xu, F. Luo, Z. Zhang, C. Tan, B. Chang, S. Huang, and F. Huang. 2021. "Raise a Child in Large Language Model: Towards Effective and Generalizable Fine-tuning." In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9514–9528.
- W. Xu, Y. Lu, and L. Wang. 2023. "Transfer Learning Enhanced DeepONet for Long-Time Prediction of Evolution Equations." In: *Proceedings of the AAAI Conference on Artificial Intelligence* 9. Vol. 37, 10629–10636.
- X. Xu, J. Y. Zhang, E. Ma, H. H. Son, S. Koyejo, and B. Li. 2022. "Adversarially Robust Models May not Transfer Better: Sufficient Conditions for Domain Transferability from the View of Regularization." In: *International Conference on Machine Learning*. PMLR, 24770–24802.
- X. Xu, J. Zhang, and M. Kankanhalli. 2024. "AutoLoRa: An Automated Robust Fine-Tuning Framework." In: *The Twelfth International Conference on Learning Representations*.
- Z. Xu, Z. Shi, J. Wei, F. Mu, Y. Li, and Y. Liang. 2024. "Towards Few-Shot Adaptation of Foundation Models via Multitask Finetuning." In: *The Twelfth International Conference on Learning Representations*.
- Z. Xu and A. Tewari. 2021. "Representation Learning Beyond Linear Prediction Functions." *Advances in Neural Information Processing Systems*, 34, 4792–4804.
- J. Yan, V. Yadav, S. Li, L. Chen, Z. Tang, H. Wang, V. Srinivasan, X. Ren, and H. Jin. 2024. "Backdooring Instruction-Tuned Large Language Models with Virtual Prompt Injection." In: *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 6065–6086.

- S. Yang, S. Jui, J. van de Weijer, et al. 2022. "Attracting and Dispersing: A Simple Approach for Source-Free Domain Adaptation." *Advances in Neural Information Processing Systems*, 35, 5802–5815.
- W. Yang, Y. Lin, P. Li, J. Zhou, and X. Sun. 2021. "RAP: Robustness-Aware Perturbations for Defending against Backdoor Attacks on NLP Models." In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8365–8381.
- Y. Yang, C. Rashtchian, H. Zhang, R. Salakhutdinov, and K. Chaudhuri. 2020. "A Closer Look at Accuracy vs. Robustness." In: *Advances in Neural Information Processing Systems* 33.
- Z. Yang, Z. Peng, Z. Wang, J. Qi, C. Chen, W. Pan, C. Wen, C. Wang, and X. Fan. 2024. "Federated Graph Learning for Cross-Domain Recommendation." *Advances in Neural Information Processing Systems*, 37.
- H. Yao, Y. Wei, L.-K. Huang, D. Xue, J. Huang, and Z. J. Li. 2021. "Functionally Regionalized Knowledge Transfer for Low-resource Drug Discovery." *Advances in Neural Information Processing Systems*, 34, 8256–8268.
- Y. Yao, H. Li, H. Zheng, and B. Y. Zhao. 2019. "Latent Backdoor Attacks on Deep Neural Networks." In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2041–2055.
- G. Yehudai, E. Fetaya, E. A. Meiron, G. Chechik, and H. Maron. 2021. "From Local Structures to Size Generalization in Graph Neural Networks." In: *International Conference on Machine Learning*. PMLR, 11975–11986.
- L. Yi, G. Xu, P. Xu, J. Li, R. Pu, C. Ling, I. McLeod, and B. Wang. 2023. "When Source-Free Domain Adaptation Meets Learning with Noisy Labels." In: *The Eleventh International Conference on Learning Representations*.
- W. Yin, S. Yu, Y. Lin, J. Liu, J.-J. Sonke, and S. Gavves. 2024. "Domain Adaptation with Cauchy-Schwarz Divergence." In: *The 40th Conference on Uncertainty in Artificial Intelligence*.
- J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. 2014. "How Transferable are Features in Deep Neural Networks?" In: *Advances in Neural Information Processing Systems* 27, 3320–3328.
- K. You, Y. Liu, J. Wang, and M. Long. 2021. "LogME: Practical Assessment of Pre-trained Models for Transfer Learning." In: *International Conference on Machine Learning*. PMLR, 12133–12143.
- Y. You, T. Chen, Z. Wang, and Y. Shen. 2023. "Graph Domain Adaptation via Theory-Grounded Spectral Regularization." In: *The Eleventh International Conference on Learning Representations*.
- K. Yu and W. Chu. 2007. "Gaussian Process Models for Link Analysis and Transfer Learning." *Advances in Neural Information Processing Systems*, 20.
- W. Yu, V. C. V. Kumar, G. Turk, and C. K. Liu. 2019. "Sim-to-Real Transfer for Biped Locomotion." In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3503–3510.
- H. Yuan, W. R. Morningstar, L. Ning, and K. Singhal. 2022. "What Do We Mean by Generalization in Federated Learning?" In: *International Conference on Learning Representations*.
- Z. Yuan, X. Hu, Q. Wu, S. Ma, C. H. Leung, X. Shen, and Y. Huang. 2022. "A Unified Domain Adaptation Framework with Distinctive Divergence Analysis." *Transactions on Machine Learning Research*, 2022.
- M. Yurochkin, A. Bower, and Y. Sun. 2020. "Training Individually Fair ML Models with Sensitive Subspace Robustness." In: *International Conference on Learning Representations*.
- E. B. Zaken, Y. Goldberg, and S. Ravfogel. 2022. "BitFit: Simple Parameter-efficient Fine-tuning for Transformer-based Masked Language-models." In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 1–9.
- M. E. Zarlenga et al. 2022. "Concept Embedding Models: Beyond the Accuracy-Explainability Trade-Off." In: *Advances in Neural Information Processing Systems* 35.
- R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. 2013. "Learning Fair Representations." In: *International Conference on Machine Learning*. PMLR, 325–333.
- Y. Zeng and K. Lee. 2024. "The Expressive Power of Low-Rank Adaptation." In: *The Twelfth International Conference on Learning Representations*.
- G. Zerveas, S. Jayaraman, D. Patel, A. Bhamidipaty, and C. Eickhoff. 2021. "A Transformer-based Framework for Multivariate Time Series Representation Learning." In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM, 2114–2124.
- B. Zhang, Z. Liu, C. Cherry, and O. Firat. 2024. "When Scaling Meets LLM Finetuning: The Effect of Data, Model and Finetuning Method." In: *The Twelfth International Conference on Learning Representations*.
- C. Zhang, L. Zhang, and J. Ye. 2012. "Generalization Bounds for Domain Adaptation." *Advances in Neural Information Processing Systems*, 25.
- F. Zhang and M. Pilanci. 2024. "Riemannian Preconditioned LoRA for Fine-Tuning Foundation Models." In: *International Conference on Machine Learning*.
- H. Zhang, H. Singh, M. Ghassemi, and S. Joshi. 2023. "'Why did the Model Fail?': Attributing Model Performance Changes to Distribution Shifts." In: *International Conference on Machine Learning*, 41550–41578.
- H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan. 2019. "Theoretically Principled Trade-off between Robustness and Accuracy." In: *International Conference on Machine Learning*. PMLR, 7472–7482.
- H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. 2018. "mixup: Beyond Empirical Risk Minimization." In: *International Conference on Learning Representations*.

- J. Zhang, S. Guo, X. Ma, H. Wang, W. Xu, and F. Wu. 2021. "Parameterized Knowledge Transfer for Personalized Federated Learning." *Advances in Neural Information Processing Systems*, 34, 10092–10104.
- K. Zhang, B. Schölkopf, K. Muandet, and Z. Wang. 2013. "Domain Adaptation under Target and Conditional Shift." In: *International Conference on Machine Learning*. PMLR, 819–827.
- M. Zhang, S. Levine, and C. Finn. 2022. "MEMO: Test Time Robustness via Adaptation and Augmentation." *Advances in Neural Information Processing Systems*, 35, 38629–38642.
- Q. Zhang, M. Chen, A. Bukharin, P. He, Y. Cheng, W. Chen, and T. Zhao. 2023. "Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning." In: *International Conference on Learning Representations*.
- S. Zhang, H. Tong, J. Xu, and R. Maciejewski. Nov. 2019. "Graph Convolutional Networks: A Comprehensive Review." *Computational Social Networks*, 6, (Nov. 2019).
- X. Zhang, Z. Zhao, T. Tsiligkaridis, and M. Zitnik. 2022. "Self-Supervised Contrastive Pre-Training For Time Series via Time-Frequency Consistency." In: *Advances in Neural Information Processing Systems* 35.
- X. Zhang, Y. Li, W. Li, K. Guo, and Y. Shao. 2022. "Personalized Federated Learning via Variational Bayesian Inference." In: *International Conference on Machine Learning*. PMLR, 26293–26310.
- Y. Zhang, J. Hui, Q. Qin, Y. Sun, T. Zhang, H. Sun, and M. Li. 2021. "Transfer-learning-based Approach for Leaf Chlorophyll Content Estimation of Winter Wheat from Hyperspectral Data." *Remote Sensing of Environment*, 267, 112724.
- Y. Zhang, Y. Song, J. Liang, K. Bai, and Q. Yang. 2020. "Two Sides of the Same Coin: White-box and Black-box Attacks for Transfer Learning." In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2989–2997.
- Y. Zhang, T. Liu, M. Long, and M. I. Jordan. 2019. "Bridging Theory and Algorithm for Domain Adaptation." In: *International Conference on Machine Learning*. PMLR, 7404–7413.
- C. Zhao, K. Jiang, X. Wu, H. Wang, L. Khan, C. Grant, and F. Chen. 2024. "Algorithmic Fairness Generalization under Covariate and Dependence Shifts Simultaneously." In: *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4419–4430.
- C. Zhao, C. Li, and C. Fu. 2019. "Cross-Domain Recommendation via Preference Propagation GraphNet." In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2165–2168.
- H. Zhao, R. T. des Combes, K. Zhang, and G. J. Gordon. 2019. "On Learning Invariant Representations for Domain Adaptation." In: *International Conference on Machine Learning*. PMLR, 7523–7532.
- H. Zhao and G. J. Gordon. 2019. "Inherent Tradeoffs in Learning Fair Representations." In: *Advances in Neural Information Processing Systems* 32, 15649–15659.
- H. Zhao, J. Hu, and A. Risteski. 2020. "On Learning Language-Invariant Representations for Universal Machine Translation." In: *International Conference on Machine Learning*. PMLR, 11352–11364.
- H. Zhao, Y. Liu, A. Alahi, and T. Lin. 2023. "On Pitfalls of Test-time Adaptation." In: *International Conference on Machine Learning*.
- M. Zhao, T. Lin, F. Mi, M. Jaggi, and H. Schütze. 2020. "Masking as an Efficient Alternative to Finetuning for Pretrained Language Models." In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2226–2241.
- S. Zhao, J. Wen, A. T. Luu, J. Zhao, and J. Fu. 2023. "Prompt as Triggers for Backdoor Attack: Examining the Vulnerability in Language Models." In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 12303–12317.
- Y. Zhao, J. Chen, and S. Du. 2023. "Blessing of Class Diversity in Pre-training." In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 283–305.
- R. Zhong, C. Snell, D. Klein, and J. Steinhardt. 2022. "Describing Differences between Text Distributions with Natural Language." In: *International Conference on Machine Learning*. PMLR, 27099–27116.
- R. Zhong, P. Zhang, S. Li, J. Ahn, D. Klein, and J. Steinhardt. 2023. "Goal Driven Discovery of Distributional Differences via Language Descriptions." *Advances in Neural Information Processing Systems*, 36, 40204–40237.
- A. Zhou and S. Levine. 2021. "Bayesian Adaptation for Covariate Shift." *Advances in Neural Information Processing Systems*, 34, 914–927.
- D. Zhou, L. Zheng, Y. Zhu, J. Li, and J. He. 2020. "Domain Adaptive Multi-Modality Neural Attention Network for Financial Forecasting." In: *Proceedings of The Web Conference 2020*, 2230–2240.
- J. Zhu, K. Greenewald, K. Nadjahi, H. S. de Ocariz Borde, R. B. Gabrielsson, L. Choshen, M. Ghassemi, M. Yurochkin, and J. Solomon. 2024. "Asymmetry in Low-Rank Adapters of Foundation Models." In: *International Conference on Machine Learning*.
- Q. Zhu, C. Yang, Y. Xu, H. Wang, C. Zhang, and J. Han. 2021. "Transfer Learning of Graph Neural Networks with Ego-graph Information Maximization." In: *Advances in Neural Information Processing Systems* 34, 1766–1779.
- Z. Zhu, W. Liang, and J. Zou. 2022. "GSCLIP: A Framework for Explaining Distribution Shifts in Natural Language." In: *Proceedings of the ICML Workshop on DataPerf: Benchmarking Data for Data-Centric AI*.
- Z. Zhu, J. Hong, and J. Zhou. 2021. "Data-Free Knowledge Distillation for Heterogeneous Federated Learning." In: *International Conference on Machine Learning*. PMLR, 12878–12889.
- B. Zoph and K. Knight. 2016. "Multi-Source Neural Translation." In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 30–34.

Y. Zou, W. Deng, and L. Zheng. 2023. “Adaptive Calibrator Ensemble: Navigating Test Set Difficulty in Out-of-Distribution Scenarios.” In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19333–19342.

Received 12 November 2024; accepted 12 June 2025