

SHIFT15M: Fashion-specific dataset for set-to-set matching with several distribution shifts

Masanari Kimura
ZOZO Research
masanari.kimura@zozo.com

Takuma Nakamura
ZOZO Research
takuma.nakamura@zozo.com

Yuki Saito
ZOZO Research
yuki.saito@zozo.com



Figure 1: t-SNE [130] visualization for the SHIFT15M.




Figure 2: Overview of SHIFT15M dataset.

```
from shift15m.datasets import NumLikesRegression

dataset = NumLikesRegression(root="./data", download=True)
(x_tr, y_tr), (x_te, y_te) = dataset.load_dataset(target_shift=True)
```

Figure 3: Minimum sample code using SHIFT15M data loader.

Abstract

This paper addresses the problem of set-to-set matching, which involves matching two different sets of items based on some criteria, especially in the case of high-dimensional items like images. Although neural networks have been applied to solve this problem, most machine learning-based approaches assume that the training and test data follow the same distribution, which is not always true in real-world scenarios. To address this limitation, we introduce SHIFT15M, a dataset that can be used to evaluate set-to-set matching models when the distribution of data changes between training and testing. We conduct benchmark experiments that demonstrate the performance drop of naive methods due to distribution shift. Additionally, we provide software to handle the SHIFT15M dataset in a simple manner: <https://github.com/st-tech/zozo-shift15m>. We believe proposed SHIFT15M dataset provide a valuable resource for evaluating set-to-set matching models under the distribution shift.

1. Introduction

One of the key problems for fashion data analysis is set-to-set matching [110, 7, 3]. For example, we can consider a task that measures the degree of completion of an outfit by matching sets of clothing items (i.e., for two sets $A = \{\text{hat, shirt, skirt}\}$ and $B = \{\text{jacket, shoes}\}$, the matching score of A and B corresponds to the goodness of the outfit $A \cup B$). To solve this, we need to investigate neural networks that handle sets [119, 73, 155, 60, 125, 136, 159, 137]. We summarize neural networks that deal with sets in Section 4.

Another common phenomenon in the domain of fashion is a trend change. These phenomena are observed at various scales, ranging from annual trend changes such as fashionable colors to seasonal trend changes such as summer to winter clothing. In the field of machine learning, such an assumption can be defined as a distribution shift (or dataset shift) [102, 88, 121, 117, 48, 117, 148, 87]. We assume that training examples $\{(\mathbf{x}_i^{tr}, y_i^{tr})\}_{i=1}^{n_{tr}}$ are independently and identically distributed (i.i.d.) according to some fixed but unknown distribution $p_{tr}(\mathbf{x}, y)$, which can be decomposed into the marginal distribution and the conditional probability

distribution, i.e., $p_{tr}(\mathbf{x}, y) = p_{tr}(\mathbf{x})p_{tr}(y|\mathbf{x})$. We also denote the test examples by $\{(\mathbf{x}_i^{te}, y_i^{te})\}_{i=1}^{n_{te}}$ drawn from a test distribution $p_{te}(\mathbf{x}, y) = p_{te}(\mathbf{x})p_{te}(y|\mathbf{x})$.

Definition 1.1. (Covariate shift [118]) We consider that the two distributions $p_{tr}(\mathbf{x}, y)$ and $p_{te}(\mathbf{x}, y)$ satisfy the covariate shift assumption if the following conditions hold:

$$p_{tr}(\mathbf{x}) \neq p_{te}(\mathbf{x}), \quad p(y|\mathbf{x}) = p_{tr}(y|\mathbf{x}) = p_{te}(y|\mathbf{x}).$$

Definition 1.2. (Target shift [156]) We consider that the two distributions $p_{tr}(\mathbf{x}, y)$ and $p_{te}(\mathbf{x}, y)$ satisfy the target shift assumption if the following conditions hold:

$$p_{tr}(y) \neq p_{te}(y), \quad p(\mathbf{x}|y) = p_{tr}(\mathbf{x}|y) = p_{te}(\mathbf{x}|y).$$

Definition 1.3. (General distribution shift [121]) Let $\mathcal{Z} \subset \{\mathcal{X}, \mathcal{Y}\}$ be a set of immutable variables whose marginal distribution should remain fixed, $\mathcal{W} \subset \{\mathcal{X}, \mathcal{Y}\} \setminus \mathcal{Z}$ be a set of mutable variables whose distribution can be shifted, and $\mathcal{V} = \{\mathcal{X}, \mathcal{Y}\} \setminus \{\mathcal{W} \cup \mathcal{Z}\}$ be the remaining dependent variables. This partition of the variables defines a factorization of p_{tr} into

$$p_{tr}(\mathbf{v}|\mathbf{w}, \mathbf{z})p_{tr}(\mathbf{w}|\mathbf{z})p_{tr}(\mathbf{z}), \quad (1)$$

where $\mathbf{z} \in \mathcal{Z}$, $\mathbf{w} \in \mathcal{W}$ and $\mathbf{v} \in \mathcal{V}$. We consider that the two distributions $p_{tr}(\mathbf{x}, y)$ and $p_{te}(\mathbf{x}, y)$ satisfy the general dataset shift assumption if the following hold:

$$p_{tr}(\mathbf{w}|\mathbf{z}) \neq p_{te}(\mathbf{w}|\mathbf{z}). \quad (2)$$

Notably, this formulation generalizes other dataset shifts. For example, if we let $\mathcal{Z} = \emptyset$ and $\mathcal{W} = \mathcal{X}$, then this corresponds to a covariate shift.

To address these problems, we provide SHIFT15M, a real-world dataset that can handle the above two problem settings, that is, the set-to-set matching dataset with distribution shift. Our SHIFT15M dataset is built on data accumulated over the past 10 years in our fashion SNS (see Figure 1). In this SNS, users could post combinations of their clothing items and other users could bookmark them as favorites. The data accumulated by this service, which has been in operation for a decade from 2010 to 2020, is very useful for dealing with distribution shifts in the fashion sector. Figure 2 shows an overview of the SHIFT15M dataset. Each column is a set of posted fashion items, with information such as the user who posted, the date of publication, and the price of each item. We hope that our SHIFT dataset will encourage research on set-to-set matching tasks under the distribution shift.

1.1. Contribution

Our contributions are summarized as follows:

- We propose SHIFT15M, a fashion-specific dataset that can properly evaluate models for set-to-set matching under the distribution shift assumptions. SHIFT15M

also enables the performance evaluation of the model under various magnitudes of dataset shifts by switching the magnitude. Figure 4 shows several sample images from the SHIFT15M dataset. In Section 2, we introduce overall statistics on the SHIFT15M dataset.

- We provide open-source software to handle the SHIFT15M dataset in a very simple way. Figure 3 shows the minimum sample code of our software;
- We propose first-step benchmark methods for set-to-set matching under distribution shift, numerical experiments show the usefulness of these methods. Section 3 presents the proposed benchmark methods and the results of comparative experiments.

2. Statistics on the SHIFT15M dataset

In this section, we present some statistics for our SHIFT15M dataset. First, Table 1 shows the overview of statistics on the SHIFT15M dataset. Since our fashion SNS was launched in 2010, the number of users and posts gradually increased from 2010, reaching a peak around 2014~2015, and slowly decreasing until 2020, the year the service was terminated. Also, the number of items in a set tends to increase over the years, indicating that users tend to construct outfits with more and more items.

The top panel of Figure 5 shows the trend of price for items included in the SHIFT15M dataset. It can be seen that the fashion items posted by users are becoming more expensive every year. The bottom panel of Figure 5 shows the trend of the number of likes for posted sets.

Figure 6 plots the trend of the number of posted sets by year. This figure shows that our fashion SNS, the source of the SHIFT15M dataset, was most active around 2014~2015.

Also, each item from the SHIFT15M dataset has two categories specifying the types of the item. Figure 11 show the distributions of the number of items belonging to each category. The figures show that there are categories in which the number of items changes from year to year and categories in which the number of items does not change much throughout the entire period.

Finally, we confirm the covariate shift of the image features included in SHIFT15M. If covariate shift assumption 1.1 holds, we should be able to construct a classifier $f : \mathbf{x} \mapsto y = \{0, 1\}$, where \mathbf{x} is the image feature of the item and y is the binary classification output for two years. Figure 7 shows the experimental results. The results show that classification between distant years (e.g., acc. of 2010 vs. 2020 is 0.85) is easier, while classification between close years (e.g., acc. of 2010 vs. 2011 is 0.62) is more difficult, indicating a gradual shift in image features. Figure 8 also shows the experimental results of item categorization when the training and test data were generated from different years.



Figure 4: Several sample images from SHIFT15M dataset. See Appendix A for more sample items.

Table 1: Statistics on the SHIFT15M dataset.

Property	Total	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
#sets	2,555,147	1,423	4,813	131,611	466,583	730,443	617,844	299,502	137,510	92,944	59,412	13,062
#items	15,218,721	8,327	29,140	756,532	2,644,564	4,305,802	3,731,864	1,853,647	855,036	576,022	373,549	84,238
mean set size	6.03	5.85	6.05	5.74	5.66	5.89	6.04	6.18	6.21	6.19	6.28	6.44
median set size	6.00	6.00	6.00	5.00	5.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00
mean #likes	26.98	0.94	2.00	15.74	16.84	23.24	37.37	35.67	32.41	24.89	21.34	16.01
median #likes	9.00	0.00	1.00	8.00	6.00	6.00	13.00	18.00	23.00	19.00	17.00	12.00
#unique users	193,574	289	571	16,922	52,283	80,290	49,441	18,854	7,511	4,442	2,739	853




Figure 5: Top panel: the trend of price for items included in SHIFT15M. Bottom panel: the trend of the number of likes for posted sets.

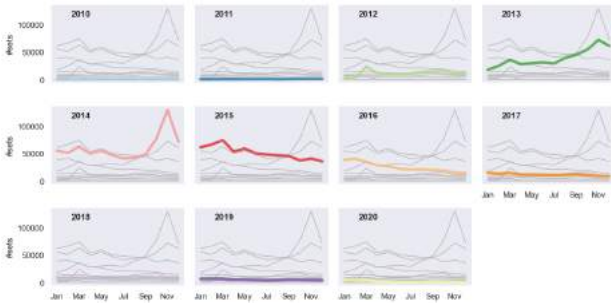


Figure 6: Trends of the number of posted sets by year.

This figure shows that the closer the years of the training and test data are, the higher the classification accuracy.

As we will see in the following sections, density ratio




Figure 7: Covariate shift of image features.




Figure 8: Category classification results under the covariate shift.

$p_{te}(\mathbf{x})/p_{tr}(\mathbf{x})$ is essential in distribution shift adaptation. Let $p(\text{train})$ and $p(\text{test})$ be the probability that some data are generated from train and test distributions, respectively. If we assume that $p(\text{train}) = p(\text{test}) = 0.5$, we can estimate density ratio as $p(\text{test}|\mathbf{x})/p(\text{train}|\mathbf{x})$ via Bayes' theorem:

$$\frac{p_{te}(\mathbf{x})}{p_{tr}(\mathbf{x})} = \frac{p(\mathbf{x}|\text{test})}{p(\mathbf{x}|\text{train})} = \frac{p(\text{train}) p(\text{test}|\mathbf{x})}{p(\text{test}) p(\text{train}|\mathbf{x})} = \frac{p(\text{test}|\mathbf{x})}{p(\text{train}|\mathbf{x})}.$$

Conversely, the Bayes optimal classifier $g^*(\mathbf{x})$ can be written




Figure 9: Density ratio estimation by using binary classifiers.

as a function of the density ratio $r(\mathbf{x}) = \frac{p_{te}(\mathbf{x})}{p_{tr}(\mathbf{x})}$,

$$r(\mathbf{x}) = \frac{p_{te}(\mathbf{x})}{p_{tr}(\mathbf{x})} = \frac{g^*(\mathbf{x})}{1 - g^*(\mathbf{x})}, \quad g^*(\mathbf{x}) = \frac{p_{te}(\mathbf{x})}{p_{te}(\mathbf{x}) + p_{tr}(\mathbf{x})}.$$

Figure 9 shows the density ratio estimation by using the above binary classifiers. These density ratios induce the importance weighted set-to-set matching algorithm in the following section. See Appendix E for more details.

3. Benchmarks

In this section, we introduce several numerical experiments on the SHIFT15M dataset.

3.1. Importance weighted set-to-set matching

As the benchmark strategy for the distribution shift adaptation on the set-to-set matching, we propose importance weighted set-to-set matching which is based on IWERM.

Definition 3.1. (Importance weighted ERM [118]) Importance Weighted Empirical Risk Minimization (IWERM) uses the density ratio $p_{te}(\mathbf{x})/p_{tr}(\mathbf{x})$ as the weighting function:

$$\hat{h} = \arg \min_{h \in \mathcal{H}} \frac{1}{n_{tr}} \sum_{i=1}^{n_{tr}} \frac{p_{te}(\mathbf{x}_i^{tr})}{p_{tr}(\mathbf{x}_i^{tr})} \ell(h(\mathbf{x}_i^{tr}), y_i^{tr}). \quad (3)$$

Adopting the density ratio as the weighting function, as in Definition 3.1, leads to the following statistically important property.

Theorem 3.1. (Consistency of IWERM[118]) If we set $w(\mathbf{x}) = p_{te}(\mathbf{x})/p_{tr}(\mathbf{x})$ as the weighting function, the empirical error computed by the weighted ERM is a consistent estimator of the expected error in the test distribution.

Using the above ideas, we propose a novel covariate shift adaptation method for set-to-set matching. Let $\mathcal{L}(\mathcal{V}, \mathcal{W}, f)$




Figure 10: Overall distribution of the number of items in each category.




Figure 11: Yearly distribution of the number of items in each category.

be the K -pair-set loss [110] function for the set matching, which is defined as follows:

$$\mathcal{L}(\mathcal{V}, \mathcal{W}, f) = -\frac{1}{K} \sum_{i=1}^K \sum_{j=1}^K \delta_{ij} \log \frac{\exp(f(\mathcal{V}_i, \mathcal{W}_j))}{\sum_{k=1}^K \exp(f(\mathcal{V}_i, \mathcal{W}_k))},$$

where δ is Kronecker's delta, and we can modify $\mathcal{L}(\mathcal{V}, \mathcal{W}, f)$ as follows:

$$\mathcal{L}_w(\mathcal{V}, \mathcal{W}, f) = -\frac{1}{K} \sum_{i=1}^K \sum_{j=1}^K \delta_{ij} \Gamma_{i,j}^p \log \frac{\Gamma_{i,j}^f}{\sum_{k=1}^K \Gamma_{i,k}^f}, \quad (4)$$

where $\Gamma_{i,j}^p = e^{p(test|\mathcal{V}_i \cup \mathcal{W}_j)}$ and $\Gamma_{i,j}^f = e^{f(\mathcal{V}_i, \mathcal{W}_j)}$. This

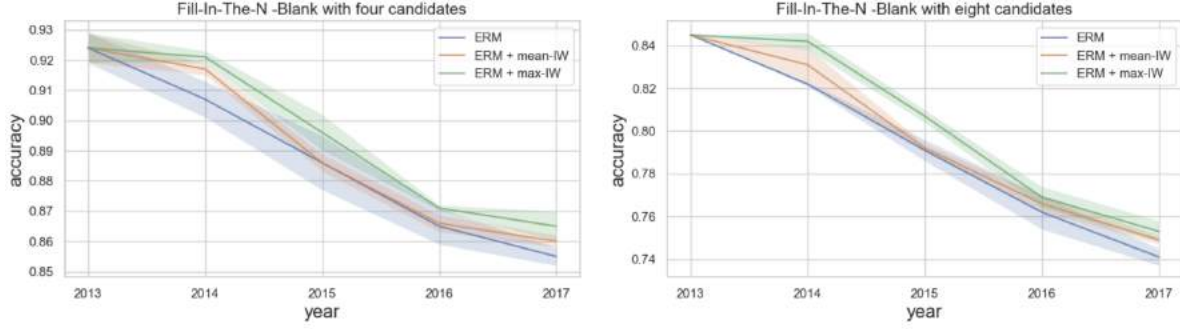


Figure 12: Plots for Fill-In-The-N-Blank experiments.

Table 2: Experimental results of the Fill-In-The- N -Blank with four candidates. Evaluation metrics are the accuracy [%]

Models	2013	2014	2015	2016	2017
ERM [110]	0.924(± 0.005)	0.907(± 0.006)	0.886(± 0.009)	0.865(± 0.006)	0.855(± 0.003)
ERM + mean-IW	0.924(± 0.005)	0.917(± 0.002)	0.886(± 0.003)	0.866(± 0.003)	0.860(± 0.002)
ERM + max-IW	0.924(± 0.005)	0.921(± 0.002)	0.896(± 0.006)	0.871(± 0.001)	0.865(± 0.005)

modification can be regarded as a weighting based on the probability that the pair is included in the test set.

Here, we propose two weighting strategies:

$$\text{max-IW} : p(\text{test}|\mathcal{V}_i \cup \mathcal{W}_j) = \max_{\mathbf{x} \in \mathcal{V}_i \cup \mathcal{W}_j} w(\mathbf{x}),$$

$$\text{mean-IW} : p(\text{test}|\mathcal{V}_i \cup \mathcal{W}_j) = \frac{1}{|\mathcal{V}_i \cup \mathcal{W}_j|} \sum_{\mathbf{x} \in \mathcal{V}_i \cup \mathcal{W}_j} w(\mathbf{x}),$$

where $w(\mathbf{x})$ is the weighting function. Next, we approximate $w(\mathbf{x})$ by using unlabeled data from both p_{tr} and p_{te} . In IWERM, the squared error can be decomposed as follows:

$$\begin{aligned} \mathbb{E}_{p_{te}} [\Delta^2] &= \mathbb{E}_{p_{tr}} [w(\mathbf{x})\Delta^2] \\ &= \mathbb{E}_{p_{tr}} [\hat{w}(\mathbf{x})\Delta^2] + \mathbb{E}_{p_{tr}} [(w(\mathbf{x}) - \hat{w}(\mathbf{x}))\Delta^2], \end{aligned}$$

where $\Delta^2 = \|f(\mathbf{x}) - y\|^2$ and $\hat{w}(\mathbf{x})$ is the approximator of the weighting function $w(\mathbf{x})$. The second term is bounded as

$$\begin{aligned} &\mathbb{E}_{p_{tr}} [(w(\mathbf{x}) - \hat{w}(\mathbf{x}))\Delta^2] \\ &\leq \frac{1}{2} \left(\mathbb{E}_{p_{tr}} [\Delta^2] + \mathbb{E}_{p_{tr}} [(w(\mathbf{x}) - \hat{w}(\mathbf{x}))^2] \right). \quad (5) \end{aligned}$$

Let s be the indicator of the distributions, where $s = 1$ corresponds to the train distribution and $s = 0$ corresponds to the test distribution, and we assume that $p(s) = 0.5$. Then, we also assume that

$$p(\mathbf{x}|s) = \begin{cases} p_{tr}(\mathbf{x}) & (s = 1), \\ p_{te}(\mathbf{x}) & (s = 0). \end{cases} \quad (6)$$

Then, we have $w(\mathbf{x}) = \frac{p(\mathbf{x}|s=0)}{p(\mathbf{x}|s=1)}$. Let $g(\mathbf{x})$ be the optimal source discriminator which identifies whether \mathbf{x} is generated p_{tr} or p_{te} . Then, we can write as $g(\mathbf{x}) = p(s = 1|\mathbf{x}) = \frac{1}{1+w(\mathbf{x})}$. Suppose that the density ratio $p_{te}(\mathbf{x})/p_{tr}(\mathbf{x})$ is bounded by $\beta > 0$, we have $\frac{1}{1+\beta} \leq g(\mathbf{x}) \leq 1$ for all \mathbf{x} . From the unlabeled data generated from p_{tr} and p_{te} , we can learn the estimator \hat{g} of g . Then, we can write the weight estimation term as

$$\begin{aligned} \mathbb{E}_{p_{tr}} [(w(\mathbf{x}) - \hat{w}(\mathbf{x}))^2] &= \mathbb{E}_{p_{tr}} \left[\left(\frac{g(\mathbf{x}) - \hat{g}(\mathbf{x})}{g(\mathbf{x})\hat{g}(\mathbf{x})} \right)^2 \right] \\ &\leq (1 + \beta)^4 \mathbb{E}_{p_{tr}} [(g(\mathbf{x}) - \hat{g}(\mathbf{x}))^2] \\ &= (1 + \beta)^4 \mathbb{E}_{p_{te}} \left[(g(\mathbf{x}) - \hat{g}(\mathbf{x}))^2 \frac{p_{tr}(\mathbf{x})}{p_{te}(\mathbf{x})} \right] \\ &\leq 2(1 + \beta)^4 \mathbb{E}_{p_{te}} [(g(\mathbf{x}) - \hat{g}(\mathbf{x}))^2] \\ &= 2(1 + \beta)^4 \left\{ \mathbb{E}_{p_{te}} [(s - g(\mathbf{x}))^2] - \mathbb{E}_{p_{te}} [(g(\mathbf{x}) - \hat{g}(\mathbf{x}))^2] \right\}. \end{aligned}$$

This indicates that the weighting function is approximated by the function $g(\mathbf{x})$.

3.2. Experimental results on set-to-set matching problem under the covariate shift assumption

We introduce benchmark results for a set-to-set matching under the covariate shift. The model architecture is the same as the previous work [110], which is based on the architecture of Transformer [134, 73, 91]. Our task can be considered an extended version of a standard task, Fill-In-The-

Table 3: Experimental results of the Fill-In-The- N -Blank with eight candidates. Evaluation metrics are the accuracy [%]

Models	2013	2014	2015	2016	2017
ERM [110]	0.845(± 0.000)	0.822(± 0.001)	0.791(± 0.005)	0.762(± 0.008)	0.741(± 0.004)
ERM + mean-IW	0.845(± 0.000)	0.831(± 0.008)	0.792(± 0.002)	0.766(± 0.004)	0.749(± 0.002)
ERM + max-IW	0.845(± 0.000)	0.842(± 0.004)	0.807(± 0.003)	0.769(± 0.005)	0.753(± 0.005)

Blank [21], which requires us to select an item that best extends an outfit from among four candidates. Because selecting a set corresponds to filling multiple blanks, we consider the set matching problem as Fill-In-The- N -Blank [110]. To construct the correct pair of sets to be matched, we randomly halve the given outfit \mathcal{O} into two non-empty proper subsets \mathcal{V} and \mathcal{W} , as follows: $\mathcal{O} \rightarrow \{\mathcal{V}, \mathcal{W}\}$, where $\mathcal{V} \cap \mathcal{W} = \emptyset$. Tables 2, 3 and Figure 12 show the experimental results of the Fill-In-The- N -Blank with four and eight candidates. In these experiments, data from 2013 are used as training data, and data from 2013~2017 are used as test data. ERM refers to empirical risk minimization [131, 132, 17], which assumes that $p_{tr}(\mathbf{x}) = p_{te}(\mathbf{x})$. From these results, we can see that the covariate shift adaptive set-to-set matching methods can achieve better performances than the ordinal ERM. This means that for set-to-set matching on the SHIFT15M dataset, we need to apply some distribution shift adaptation methods.

Table 4 and 5 show the experimental results for the various models. The models used in the experiments are the same as [110]. To quote,

- Cross Attention and Cross Affinity [110]: Set-to-Set matching models which proposed by [110], with the attention-based and affinity-based functions, respectively.
- Set Transformer [73]: Set Transformer which introduced by applying a self-attention based Transformer to a set of data. Set Transformer is trained through supervised or unsupervised learning and transforms a set of data into a vector representation to recognize set features. By using Set Transformer f_{ST} , we perform the extension by calculating the matching score between the two sets \mathcal{V} and \mathcal{W} via the inner product $f_{ST}(\mathcal{V})^\top f_{ST}(\mathcal{W})$, sharing the weights between the two f_{ST} .
- We consider a union of two sets as a set-input for the extension of BERT [27] and omit the individual token embedding. We use the segment embedding to designate items of X and Y. We use three variants: BERT_{BASE} is the same model as described in [27]; BERT_{BASE-AP} uses the average pooling in the last layer; and BERT_{SMALL} is a four-layered version of BERT_{BASE} with eight heads, and the hidden size is 512.

- GNN [21]: We combine two sets as one input for the extension of GNN [21]. Because this model is not presented to train in an end-to-end with the feature extractor, we do not finetune the CNN in fashion set matching, where pre-trained CNNs are used, but train it in an end-to-end manner for the group re-id task. Note that we omit the context provided from the external graphs in the evaluation stage to apply this model in the same scenarios of our tasks. We set the training epoch to 256 in the group re-id to enhance the training results of the GNN.
- HAP2S [154]: A conventional CNN trained by Hard-Aware Point-to-Set loss.

See Appendix B.2 for the additional figures.

3.3. Additional experimental results on regression problem under the target shift assumption

Additionally, we present benchmark results for a regression problem with the target shift. Although SHIFT15M is a dataset for set-to-set matching, it can also be used for simple regression problems by setting the input to image features and the output to attributes associated with the items.

The target variable is the number of likes that each instance possesses, and the input variables are the user ID and the prices of the items. In this experiment, we evaluate the robustness of the model for different shift magnitudes, and the magnitudes of the shift are measured by the Wasserstein distance. We use the simple linear regression as the ordinal ERM, and we compare this with two other covariate shift adaptation methods, IWERM and AIWERM.

Definition 3.2. (Adaptive importance weighted ERM [118]) AIWERM uses $(p_{te}(\mathbf{x})/p_{tr}(\mathbf{x}))^\alpha$ for $\alpha \in [0, 1]$ as the weighting function:

$$\hat{h} = \arg \min_{h \in \mathcal{H}} \frac{1}{n_{tr}} \sum_{i=1}^{n_{tr}} \left(\frac{p_{te}(\mathbf{x}_i^{tr})}{p_{tr}(\mathbf{x}_i^{tr})} \right)^\alpha \ell(h(\mathbf{x}_i^{tr}), y_i^{tr}). \quad (7)$$

Definition 3.3. (Relative importance weighted ERM [146]) RIWERM uses $p_{te}(\mathbf{x})/((1-\alpha)p_{tr}(\mathbf{x}) + \alpha p_{te}(\mathbf{x}))$ for $\alpha \in [0, 1]$ as the weighting function:

$$\hat{h} = \arg \min_{h \in \mathcal{H}} \frac{1}{n_{tr}} \sum_{i=1}^{n_{tr}} \frac{p_{te}(\mathbf{x}_i^{tr})}{m_\alpha(\mathbf{x}_i^{tr})} \ell(h(\mathbf{x}_i^{tr}), y_i^{tr}), \quad (8)$$

Table 4: Experimental results of the Fill-In-The- N -Blank with four candidates with several models. Evaluation metrics are the accuracy [%].

Models	2013	2014	2015	2016	2017
Set Transformer [73]	0.791(± 0.036)	0.743(± 0.041)	0.710(± 0.047)	0.698(± 0.050)	0.675(± 0.051)
Set Transformer + mean-IW	0.791(± 0.036)	0.755(± 0.037)	0.732(± 0.041)	0.710(± 0.048)	0.696(± 0.049)
Set Transformer + max-IW	0.791(± 0.036)	0.760(± 0.037)	0.731(± 0.038)	0.714(± 0.045)	0.700(± 0.045)
BERT _{SMALL} [27]	0.898(± 0.008)	0.882(± 0.005)	0.860(± 0.005)	0.842(± 0.007)	0.830(± 0.011)
BERT _{SMALL} + mean-IW	0.898(± 0.008)	0.893(± 0.003)	0.866(± 0.005)	0.844(± 0.007)	0.839(± 0.008)
BERT _{SMALL} + max-IW	0.898(± 0.008)	0.895(± 0.002)	0.878(± 0.003)	0.859(± 0.002)	0.851(± 0.004)
BERT _{BASE} [27]	0.880(± 0.011)	0.875(± 0.009)	0.844(± 0.004)	0.827(± 0.007)	0.817(± 0.010)
BERT _{BASE} + mean-IW	0.880(± 0.011)	0.878(± 0.010)	0.853(± 0.007)	0.840(± 0.013)	0.822(± 0.004)
BERT _{BASE} + max-IW	0.880(± 0.011)	0.877(± 0.003)	0.870(± 0.012)	0.852(± 0.003)	0.830(± 0.003)
BERT _{BASE-AP} [27]	0.869(± 0.008)	0.860(± 0.007)	0.825(± 0.013)	0.802(± 0.010)	0.785(± 0.009)
BERT _{BASE-AP} + mean-IW	0.869(± 0.008)	0.864(± 0.009)	0.837(± 0.009)	0.829(± 0.010)	0.797(± 0.012)
BERT _{BASE-AP} + max-IW	0.869(± 0.008)	0.866(± 0.005)	0.844(± 0.004)	0.833(± 0.008)	0.810(± 0.002)
GNN [21]	0.370(± 0.032)	0.361(± 0.035)	0.355(± 0.035)	0.351(± 0.035)	0.349(± 0.033)
GNN + mean-IW	0.370(± 0.032)	0.366(± 0.033)	0.360(± 0.032)	0.357(± 0.033)	0.354(± 0.033)
GNN + max-IW	0.370(± 0.032)	0.367(± 0.032)	0.362(± 0.030)	0.362(± 0.031)	0.359(± 0.030)
HAP2S [154]	0.433(± 0.041)	0.420(± 0.044)	0.409(± 0.053)	0.385(± 0.053)	0.370(± 0.062)
HAP2S + mean-IW	0.433(± 0.041)	0.428(± 0.042)	0.415(± 0.045)	0.408(± 0.050)	0.395(± 0.057)
HAP2S + max-IW	0.433(± 0.041)	0.427(± 0.042)	0.418(± 0.042)	0.413(± 0.047)	0.400(± 0.055)
Cross Attention [110]	0.920(± 0.007)	0.892(± 0.009)	0.867(± 0.012)	0.844(± 0.012)	0.836(± 0.014)
Cross Attention + mean-IW	0.920(± 0.007)	0.901(± 0.008)	0.873(± 0.012)	0.850(± 0.013)	0.843(± 0.014)
Cross Attention + max-IW	0.920(± 0.007)	0.913(± 0.009)	0.877(± 0.010)	0.849(± 0.011)	0.841(± 0.013)
Cross Affinity [110]	0.924(± 0.005)	0.907(± 0.006)	0.886(± 0.009)	0.865(± 0.006)	0.855(± 0.003)
Cross Affinity + mean-IW	0.924(± 0.005)	0.917(± 0.002)	0.886(± 0.003)	0.866(± 0.003)	0.860(± 0.002)
Cross Affinity + max-IW	0.924(± 0.005)	0.921(± 0.002)	0.896(± 0.006)	0.871(± 0.001)	0.865(± 0.005)

Table 5: Experimental results of the Fill-In-The- N -Blank with eight candidates with several models. Evaluation metrics are the accuracy [%]

Models	2013	2014	2015	2016	2017
Set Transformer [73]	0.711(± 0.042)	0.690(± 0.040)	0.663(± 0.037)	0.637(± 0.044)	0.605(± 0.041)
Set Transformer + mean-IW	0.711(± 0.042)	0.702(± 0.038)	0.685(± 0.043)	0.662(± 0.040)	0.630(± 0.040)
Set Transformer + max-IW	0.711(± 0.042)	0.708(± 0.037)	0.699(± 0.044)	0.675(± 0.040)	0.641(± 0.045)
BERT _{SMALL} [27]	0.824(± 0.015)	0.799(± 0.023)	0.772(± 0.025)	0.740(± 0.039)	0.716(± 0.039)
BERT _{SMALL} + mean-IW	0.824(± 0.015)	0.815(± 0.016)	0.787(± 0.023)	0.759(± 0.031)	0.720(± 0.030)
BERT _{SMALL} + max-IW	0.824(± 0.015)	0.813(± 0.015)	0.805(± 0.020)	0.782(± 0.035)	0.732(± 0.030)
BERT _{BASE} [27]	0.810(± 0.017)	0.780(± 0.027)	0.764(± 0.035)	0.733(± 0.043)	0.700(± 0.044)
BERT _{BASE} + mean-IW	0.810(± 0.017)	0.799(± 0.020)	0.778(± 0.028)	0.750(± 0.022)	0.714(± 0.033)
BERT _{BASE} + max-IW	0.810(± 0.017)	0.805(± 0.023)	0.794(± 0.021)	0.764(± 0.020)	0.738(± 0.029)
BERT _{BASE-AP} [27]	0.801(± 0.012)	0.765(± 0.028)	0.741(± 0.030)	0.719(± 0.042)	0.694(± 0.048)
BERT _{BASE-AP} + mean-IW	0.801(± 0.012)	0.788(± 0.025)	0.763(± 0.026)	0.734(± 0.026)	0.709(± 0.030)
BERT _{BASE-AP} + max-IW	0.801(± 0.012)	0.795(± 0.025)	0.777(± 0.020)	0.748(± 0.028)	0.722(± 0.021)
GNN [21]	0.346(± 0.030)	0.329(± 0.033)	0.319(± 0.036)	0.301(± 0.045)	0.287(± 0.050)
GNN + mean-IW	0.346(± 0.030)	0.337(± 0.031)	0.325(± 0.040)	0.310(± 0.040)	0.299(± 0.044)
GNN + max-IW	0.346(± 0.030)	0.342(± 0.035)	0.338(± 0.036)	0.317(± 0.038)	0.306(± 0.038)
HAP2S [154]	0.382(± 0.044)	0.367(± 0.046)	0.340(± 0.048)	0.325(± 0.045)	0.310(± 0.053)
HAP2S + mean-IW	0.382(± 0.044)	0.371(± 0.040)	0.343(± 0.048)	0.331(± 0.046)	0.312(± 0.051)
HAP2S + max-IW	0.382(± 0.044)	0.380(± 0.043)	0.359(± 0.046)	0.342(± 0.046)	0.324(± 0.048)
Cross Attention [110]	0.839(± 0.002)	0.808(± 0.005)	0.772(± 0.005)	0.740(± 0.009)	0.717(± 0.010)
Cross Attention + mean-IW	0.839(± 0.002)	0.810(± 0.009)	0.788(± 0.007)	0.752(± 0.006)	0.720(± 0.008)
Cross Attention + max-IW	0.839(± 0.002)	0.836(± 0.003)	0.800(± 0.005)	0.783(± 0.007)	0.743(± 0.003)
Cross Affinity [110]	0.845(± 0.000)	0.822(± 0.001)	0.791(± 0.005)	0.762(± 0.008)	0.741(± 0.004)
Cross Affinity + mean-IW	0.845(± 0.000)	0.831(± 0.008)	0.792(± 0.002)	0.766(± 0.004)	0.749(± 0.002)
Cross Affinity + max-IW	0.845(± 0.000)	0.842(± 0.004)	0.807(± 0.003)	0.769(± 0.005)	0.753(± 0.005)

Table 6: Experimental results for the regression problem with distribution shift adaptation. Evaluation metrics are the MSE.

Models	$W = 0$	$W = 10$	$W = 20$	$W = 30$	$W = 40$	$W = 50$
ERM	9.36(± 0.02)	10.44(± 0.04)	17.10(± 0.06)	28.80(± 0.05)	39.56(± 0.05)	48.84(± 0.05)
IWERM (optimal)	9.36(± 0.02)	25.67(± 0.12)	32.58(± 0.12)	26.83(± 0.11)	20.19(± 0.10)	14.52(± 0.10)
RIWERM ($\alpha = 0.25$)	9.36(± 0.02)	9.34(± 0.04)	9.53(± 0.03)	11.37(± 0.04)	14.89(± 0.09)	17.00(± 0.14)
RIWERM ($\alpha = 0.50$)	9.36(± 0.02)	9.73(± 0.04)	9.57(± 0.03)	9.69(± 0.04)	12.37(± 0.10)	14.68(± 0.15)
RIWERM ($\alpha = 0.75$)	9.36(± 0.02)	11.59(± 0.05)	11.35(± 0.05)	9.39(± 0.03)	10.46(± 0.09)	12.60(± 0.14)




Figure 13: Experimental results of the regression problem for the number of likes

where $m_\alpha(\mathbf{x}_i^{tr}) = (1 - \alpha)p_{tr}(\mathbf{x}_i^{tr}) + \alpha p_{te}(\mathbf{x}_i^{tr})$.

Note that there are other variants of IWERM, as in [65].

Figure 13 and Table 6 show the experimental results of the regression task. These results show that the performance of ERM decreases with increasing shift magnitude, while IWERM and its variants allow the robustness to target shifts.

4. Related works and conclusion

There are many studies for handling set data with neural networks. DeepSets [155] proposes a model that satisfies the key concepts of permutation invariant and permutation equivariant for approximating functions that deal with sets.

Definition 4.1 (Permutation invariant). A set function f is said to be permutation invariant if $f(\mathcal{V}, \mathcal{W}) = f(\pi_w \mathcal{W}, \pi_v \mathcal{V})$ for permutations π_w and π_v .

Definition 4.2 (Permutation equivariance). A set function f is said to be permutation equivariant if $f(\pi_w \mathcal{W}, \pi_v \mathcal{V}) = \pi_w f(\mathcal{W}, \mathcal{V})$ for permutations π_w and π_v . Note that f is permutation invariant for permutations within \mathcal{V} .

Theorem 4.1 (Sum-decomposable[155]). A function f on set \mathcal{X} from countable particle space is invariant if and only if there exists a decomposition,

$$f(\mathcal{X}) = \rho \left(\sum_{\mathbf{x} \in \mathcal{X}} \phi(\mathbf{x}) \right), \quad f = \rho \circ \sum \circ \phi,$$

with appropriate functions ϕ and ρ .

It is pointed out that the necessary and sufficient of sum-decomposable is guaranteed only for countable sets [136].

Theorem 4.2 ([136]). A continuous function f on finite sets \mathcal{X} , $|\mathcal{X}| < p$, is invariant if and only if it is sum-decomposable via \mathbb{R}^p .

That is, for an arbitrary continuous function f , the image space of ϕ has to have at least dimension p , which is both necessary and sufficient. For more details, see Appendix C.

SetTransformer [73] is an attention-based neural network module that allows us to handle sets as inputs. SetVAE [63], an extension of VAE to set data, has also been proposed.

Several distribution shift datasets exist for general classification and regression tasks where the input is a vector. WILDS [68] is the collection of benchmark datasets [10, 8, 126, 52, 24, 16, 20, 153, 93, 85] under the distribution shift, including histopathological images, satellite images or sequence of source code tokens. PACS [74] and Office-Home [135] adopt the image style to differentiate distributions, and VLCS [31] takes data collected independently from four sources as environments. Also, DomainNet [160] extends PACS to a far larger scale.

There are also a number of studies that evaluate robustness to distribution shifts by introducing artificial distribution shifts, such as noise corruptions [42, 49, 90, 108, 145], spatial transformations [30, 32], ImageNet [26] variants (e.g. ImageNet-A [50], ImageNet-C [49], ImageNet-R [48]), and adversarial examples [14, 123, 72, 45, 143, 18]. However, a recent study [124] has indicated that there is no correlation between the robustness of such artificial distribution shifts and the robustness of natural distribution shifts.

We believe that our SHIFT15M is a very useful dataset for evaluating the still underdeveloped task of set-to-set matching under natural distribution shifts. See Appendix D and F for more related literature including domain adaptation, out-of-domain generalization, concept drift adaptation, or other fashion datasets. We also provide the datasheet for the SHIFT15M in Appendix H, which is based on [41].

4.1. Future works

- More ablation studies: more model architectures and parameter influences need to be investigated.
- Experiments on model calibration: model calibration and distribution shift are known to be closely related.

lated [138, 95, 70]. We expect that observing the metrics for evaluating model calibration in experiments on the SHIFT15M will provide meaningful insights.

- Additional API development: currently, our API is based on PyTorch [97]. In the future, we would like to expand the APIs for other machine learning libraries (e.g., TensorFlow [1] or Keras [19]).

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016. 9
- [2] Cesare Alippi and Manuel Roveri. Just-in-time adaptive classifiers—part i: Detecting nonstationary changes. *IEEE Transactions on Neural Networks*, 19(7):1145–1153, 2008. 19
- [3] Ognjen Arandjelović. Discriminative extended canonical correlation analysis for pattern set matching. *Machine Learning*, 94:353–370, 2014. 1
- [4] Martin Arjovsky. *Out of distribution generalization in machine learning*. PhD thesis, New York University, 2020. 20
- [5] Stephen H Bach and Marcus A Maloof. Paired learners for concept drift. In *2008 Eighth IEEE International Conference on Data Mining*, pages 23–32. IEEE, 2008. 20
- [6] Manuel Baena-García, José del Campo-Ávila, Raul Fidalgo, Albert Bifet, Ricard Gavaldà, and Rafael Morales-Bueno. Early drift detection method. In *Fourth international workshop on knowledge discovery from data streams*, volume 6, pages 77–86. Citeseer, 2006. 19
- [7] Yunsheng Bai, Hao Ding, Yizhou Sun, and Wei Wang. Convolutional set matching for graph similarity. *arXiv preprint arXiv:1810.10866*, 2018. 1
- [8] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018. 8
- [9] Michele Basseville, Igor V Nikiforov, et al. *Detection of abrupt changes: theory and application*, volume 104. prentice Hall Englewood Cliffs, 1993. 19
- [10] Sara Beery, Elijah Cole, and Arvi Gjoka. The iwildcam 2020 competition dataset. *arXiv preprint arXiv:2004.10340*, 2020. 8
- [11] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79:151–175, 2010. 18
- [12] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19, 2006. 17
- [13] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013. 20
- [14] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*, pages 387–402. Springer, 2013. 8
- [15] François Bolley, Arnaud Guillin, and Cédric Villani. Quantitative concentration inequalities for empirical measures on non-compact spaces. *Probability Theory and Related Fields*, 137:541–593, 2007. 18
- [16] Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*, 2019. 8
- [17] Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In *Summer School on Machine Learning*, pages 169–207. Springer, 2003. 6
- [18] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 8
- [19] Francois Chollet et al. Keras, 2015. 9
- [20] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 8
- [21] Guillem Cucurull, Perouz Taslakian, and David Vazquez. Context-aware visual compatibility prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12617–12626, 2019. 6, 7
- [22] Hal Daumé III. Frustratingly easy domain adaptation. In *Proc. of the 45th Annual Meeting of the Association of Computational Linguistics*, 2007, pages 256–263, 2007. 20
- [23] Hal Daumé III and Daniel Marcu. Domain adaptation for statistical classifiers. *Journal of artificial Intelligence research*, 26:101–126, 2006. 20
- [24] Etienne David, Simon Madec, Pouria Sadeghi-Tehran, Helge Aasen, Bangyou Zheng, Shouyang Liu, Norbert Kirchgessner, Goro Ishikawa, Koichi Nagasawa, Minhajul A Badhon, Curtis Pozniak, Benoit de Solan, Andreas Hund, Scott C. Chapman, Frederic Baret, Ian Stavness, and Wei Guo. Global wheat head detection (gwhd) dataset: a large and diverse dataset of high-resolution rgb-labelled images to develop and benchmark wheat head detection methods. *Plant Phenomics*, 2020, 2020. 8
- [25] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010. 20

- [26] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 8
- [27] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 6, 7
- [28] Jeff Donahue, Judy Hoffman, Erik Rodner, Kate Saenko, and Trevor Darrell. Semi-supervised domain adaptation with instance constraints. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 668–675, 2013. 20
- [29] Ryan Elwell and Robi Polikar. Incremental learning of concept drift in nonstationary environments. *IEEE Transactions on Neural Networks*, 22(10):1517–1531, 2011. 20
- [30] Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *International conference on machine learning*, pages 1802–1811. PMLR, 2019. 8
- [31] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1657–1664, 2013. 8
- [32] Alhussein Fawzi and Pascal Frossard. Manitest: Are classifiers really invariant? *arXiv preprint arXiv:1507.06535*, 2015. 8
- [33] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of the IEEE international conference on computer vision*, pages 2960–2967, 2013. 20
- [34] Isvani Frias-Blanco, José del Campo-Ávila, Gonzalo Ramos-Jimenez, Rafael Morales-Bueno, Agustin Ortiz-Diaz, and Yailé Caballero-Mota. Online and non-parametric drift detection methods based on hoeffding’s bounds. *IEEE Transactions on Knowledge and Data Engineering*, 27(3):810–823, 2014. 19
- [35] Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. 16
- [36] Joao Gama and Gladys Castillo. Learning with local drift detection. In *Advanced Data Mining and Applications: Second International Conference, ADMA 2006, Xi’an, China, August 14-16, 2006 Proceedings 2*, pages 42–55. Springer, 2006. 19
- [37] Joao Gama, Pedro Medas, Gladys Castillo, and Pedro Rodrigues. Learning with drift detection. In *Advances in Artificial Intelligence—SBIA 2004: 17th Brazilian Symposium on Artificial Intelligence, Sao Luis, Maranhao, Brazil, September 29-October 1, 2004. Proceedings 17*, pages 286–295. Springer, 2004. 19
- [38] Joao Gama, Ricardo Rocha, and Pedro Medas. Accurate decision trees for mining high-speed data streams. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 523–528, 2003. 20
- [39] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014. 19
- [40] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 20
- [41] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021. 8, 29
- [42] Robert Geirhos, Carlos RM Temme, Jonas Rauber, Heiko H Schütt, Matthias Bethge, and Felix A Wichmann. Generalisation in humans and deep neural networks. *Advances in neural information processing systems*, 31, 2018. 8
- [43] Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010. 20
- [44] Heitor M Gomes, Albert Bifet, Jesse Read, Jean Paul Barddal, Fabrício Enembreck, Bernhard Pfharinger, Geoff Holmes, and Talel Abdesslem. Adaptive random forests for evolving data stream classification. *Machine Learning*, 106:1469–1495, 2017. 20
- [45] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 8
- [46] Arthur Gretton, Alex Smola, Jiayuan Huang, Marcel Schmittfull, Karsten Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4):5, 2009. 19
- [47] Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018. 20
- [48] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 1, 8, 20
- [49] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. 8
- [50] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. 8
- [51] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2017. 20
- [52] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine

- learning on graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 8
- [53] Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex Smola. Correcting sample selection bias by unlabeled data. *Advances in neural information processing systems*, 19, 2006. 19
- [54] Geoff Hulten, Laurie Spencer, and Pedro Domingos. Mining time-changing data streams. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 97–106, 2001. 20
- [55] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7610–7619, 2020. 20
- [56] Menglin Jia, Mengyun Shi, Mikhail Sirotenko, Yin Cui, Claire Cardie, Bharath Hariharan, Hartwig Adam, and Serge Belongie. Fashionpedia: Ontology, segmentation, and an attribute localization dataset. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 316–332. Springer, 2020. 21
- [57] Takafumi Kanamori, Shohei Hido, and Masashi Sugiyama. A least-squares approach to direct importance estimation. *The Journal of Machine Learning Research*, 10:1391–1445, 2009. 19
- [58] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4893–4902, 2019. 20
- [59] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 20
- [60] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022. 1
- [61] Jack Kiefer and Jacob Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, pages 462–466, 1952. 16
- [62] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International Conference on Machine Learning*, pages 2649–2658. PMLR, 2018. 20
- [63] Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical composition for generative modeling of set-structured data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15059–15068, 2021. 8
- [64] Masanari Kimura. Generalization bounds for set-to-set matching with negative sampling. *arXiv preprint arXiv:2302.12991*, 2023. 21
- [65] Masanari Kimura and Hideitsu Hino. Information geometrically generalized covariate shift adaptation. *Neural Computation*, 34(9):1944–1977, 2022. 8
- [66] Masanari Kimura, Kei Wakabayashi, and Atsuyuki Morishima. Batch prioritization of data labeling tasks for training classifiers. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 163–167, 2020. 20
- [67] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 20
- [68] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. 8
- [69] J Zico Kolter and Marcus A Maloof. Dynamic weighted majority: An ensemble method for drifting concepts. *The Journal of Machine Learning Research*, 8:2755–2790, 2007. 20
- [70] Ananya Kumar, Tengyu Ma, Percy Liang, and Aditi Raghunathan. Calibrated ensembles can mitigate accuracy trade-offs under distribution shift. In *Uncertainty in Artificial Intelligence*, pages 1041–1051. PMLR, 2022. 9
- [71] Abhishek Kumar, Avishek Saha, and Hal Daume. Co-regularization based semi-supervised domain adaptation. *Advances in neural information processing systems*, 23, 2010. 20
- [72] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pages 99–112. Chapman and Hall/CRC, 2018. 8
- [73] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, pages 3744–3753. PMLR, 2019. 1, 5, 6, 7, 8
- [74] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 8, 20
- [75] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 20
- [76] Haoyang Li, Xin Wang, Ziwei Zhang, and Wenwu Zhu. Out-of-distribution generalization on graphs: A survey. *arXiv preprint arXiv:2202.07987*, 2022. 20
- [77] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9641–9650, 2020. 20
- [78] Anjin Liu, Yiliao Song, Guangquan Zhang, and Jie Lu. Regional concept drift detection and density synchronized drift adaptation. In *IJCAI International Joint Conference on Artificial Intelligence*, 2017. 20

- [79] Anjin Liu, Guangquan Zhang, and Jie Lu. Fuzzy time windowing for gradual concept drift adaptation. In *2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, pages 1–6. IEEE, 2017. 19
- [80] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 21
- [81] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. *Advances in neural information processing systems*, 29, 2016. 20
- [82] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. Learning under concept drift: A review. *IEEE transactions on knowledge and data engineering*, 31(12):2346–2363, 2018. 19
- [83] Ning Lu, Jie Lu, Guangquan Zhang, and Ramon Lopez De Mantaras. A concept drift-tolerant case-base editing technique. *Artificial Intelligence*, 230:108–133, 2016. 19
- [84] Ning Lu, Guangquan Zhang, and Jie Lu. Concept drift detection via competence models. *Artificial Intelligence*, 209:11–28, 2014. 19
- [85] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, et al. Codexglue: A machine learning benchmark dataset for code understanding and generation. *arXiv preprint arXiv:2102.04664*, 2021. 8
- [86] Bryan FJ Manly and Darryl Mackenzie. A cumulative sum type of method for environmental monitoring. *Environmetrics: The official journal of the International Environmetrics Society*, 11(2):151–166, 2000. 20
- [87] John P Miller, Rohan Taori, Aditi Raghunathan, Shiori Sagawa, Pang Wei Koh, Vaishaal Shankar, Percy Liang, Yair Carmon, and Ludwig Schmidt. Accuracy on the line: on the strong correlation between out-of-distribution and in-distribution generalization. In *International Conference on Machine Learning*, pages 7721–7735. PMLR, 2021. 1
- [88] Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern recognition*, 45(1):521–530, 2012. 1
- [89] Saeid Motiian, Marco Piccirilli, Donald A Adjero, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5715–5725, 2017. 20
- [90] Norman Mu and Justin Gilmer. Mnist-c: A robustness benchmark for computer vision. *arXiv preprint arXiv:1906.02337*, 2019. 8
- [91] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 5
- [92] XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010. 19
- [93] Jianmo Ni, Jiacheng Li, and Julian McAuley. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019. 8
- [94] Kyosuke Nishida and Koichiro Yamauchi. Detecting concept drift using statistical testing. In *Discovery science*, volume 4755, pages 264–269. Springer, 2007. 20
- [95] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019. 9
- [96] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019. 20
- [97] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 9
- [98] Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE signal processing magazine*, 32(3):53–69, 2015. 20
- [99] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 20
- [100] Niklas Pfister, Peter Bühlmann, and Jonas Peters. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527):1264–1276, 2019. 20
- [101] Michael Prince. Does active learning work? a review of the research. *Journal of engineering education*, 93(3):223–231, 2004. 20
- [102] Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008. 1
- [103] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019. 20
- [104] Ievgen Redko, Amaury Habrard, and Marc Sebban. Theoretical analysis of domain adaptation with optimal transport. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18–22, 2017, Proceedings, Part II 10*, pages 737–753. Springer, 2017. 18
- [105] Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younes Bennani. *Advances in domain adaptation theory*. Elsevier, 2019. 17

- [106] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. [16](#)
- [107] Dominik Rothenhäusler, Nicolai Meinshausen, Peter Bühlmann, and Jonas Peters. Anchor regression: heterogeneous data meets causality. *arXiv preprint arXiv:1801.06229*, 2018. [20](#)
- [108] Evgenia Rusak, Lukas Schott, Roland S Zimmermann, Julian Bitterwolf, Oliver Bringmann, Matthias Bethge, and Wieland Brendel. A simple way to make neural networks robust against diverse image corruptions. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 53–69. Springer, 2020. [8](#)
- [109] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8050–8058, 2019. [20](#)
- [110] Yuki Saito, Takuma Nakamura, Hirotaka Hachiya, and Kenji Fukumizu. Exchangeable deep neural networks for set-to-set matching and learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII*, pages 626–646. Springer, 2020. [1](#), [4](#), [5](#), [6](#), [7](#), [16](#)
- [111] Danielle Saunders. Domain adaptation and multi-domain adaptation for neural machine translation: A survey. *Journal of Artificial Intelligence Research*, 75:351–424, 2022. [20](#)
- [112] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. [20](#)
- [113] Ozan Sener, Hyun Oh Song, Ashutosh Saxena, and Silvio Savarese. Learning transferrable representations for unsupervised domain adaptation. *Advances in neural information processing systems*, 29, 2016. [20](#)
- [114] Burr Settles. Active learning literature survey. 2009. [20](#)
- [115] Junming Shao, Zahra Ahmadi, and Stefan Kramer. Prototype-based learning on concept-drifting data streams. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 412–421, 2014. [19](#)
- [116] Zheyang Shen, Peng Cui, Jiashuo Liu, Tong Zhang, Bo Li, and Zhitang Chen. Stable learning via differentiated variable decorrelation. In *Proceedings of the 26th acm sigkdd international conference on knowledge discovery & data mining*, pages 2185–2193, 2020. [20](#)
- [117] Zheyang Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*, 2021. [1](#), [20](#)
- [118] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000. [2](#), [4](#), [6](#)
- [119] Maximilian Soelch, Adnan Akhundov, Patrick van der Smagt, and Justin Bayer. On deep set learning and the choice of aggregations. In *Artificial Neural Networks and Machine Learning—ICANN 2019: Theoretical Neural Computation: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part I 28*, pages 444–457. Springer, 2019. [1](#)
- [120] Xiuyao Song, Mingxi Wu, Christopher Jermaine, and Sanjay Ranka. Statistical change detection for multi-dimensional data. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 667–676, 2007. [19](#)
- [121] Adarsh Subbaswamy, Roy Adams, and Suchi Saria. Evaluating model robustness and stability to dataset shift. In *International Conference on Artificial Intelligence and Statistics*, pages 2611–2619. PMLR, 2021. [1](#), [2](#)
- [122] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. *Advances in neural information processing systems*, 20, 2007. [19](#)
- [123] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [8](#)
- [124] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020. [8](#)
- [125] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *ACM Computing Surveys*, 55(6):1–28, 2022. [1](#)
- [126] J. Taylor, B. Earnshaw, B. Mabey, M. Victors, and J. Yosinski. Rrx1: An image set for cellular morphological variation across many experimental batches. In *International Conference on Learning Representations (ICLR)*, 2019. [8](#)
- [127] Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI global, 2010. [20](#)
- [128] Alexey Tsymbal. The problem of concept drift: definitions and related work. *Computer Science Department, Trinity College Dublin*, 106(2):58, 2004. [19](#)
- [129] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE international conference on computer vision*, pages 4068–4076, 2015. [20](#)
- [130] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [1](#)
- [131] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013. [6](#)
- [132] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999. [6](#)
- [133] Mariya I Vasileva, Bryan A Plummer, Krishna Dusad, Shreya Rajpal, Ranjitha Kumar, and David Forsyth. Learning type-aware embeddings for fashion compatibility. In

- Proceedings of the European conference on computer vision (ECCV)*, pages 390–405, 2018. [21](#)
- [134] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [5](#)
- [135] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. [8](#)
- [136] Edward Wagstaff, Fabian Fuchs, Martin Engelcke, Ingmar Posner, and Michael A Osborne. On the limitations of representing functions on sets. In *International Conference on Machine Learning*, pages 6487–6494. PMLR, 2019. [1](#), [8](#)
- [137] Edward Wagstaff, Fabian B Fuchs, Martin Engelcke, Michael A Osborne, and Ingmar Posner. Universal approximation of functions on sets. *Journal of Machine Learning Research*, 23(151):1–56, 2022. [1](#), [17](#)
- [138] Yoav Wald, Amir Feder, Daniel Greenfeld, and Uri Shalit. On calibration and out-of-domain generalization. *Advances in neural information processing systems*, 34:2215–2227, 2021. [9](#)
- [139] Heng Wang and Zubin Abraham. Concept drift detection for streaming data. In *2015 international joint conference on neural networks (IJCNN)*, pages 1–9. IEEE, 2015. [19](#)
- [140] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. [20](#)
- [141] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016. [20](#)
- [142] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. The fashion iq dataset: Retrieving images by combining side information and relative natural language feedback. *CVPR*, 2021. [21](#)
- [143] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. *arXiv preprint arXiv:1801.02610*, 2018. [8](#)
- [144] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. [21](#)
- [145] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. *arXiv preprint arXiv:2007.13003*, 2020. [8](#)
- [146] Makoto Yamada, Taiji Suzuki, Takafumi Kanamori, Hiro-taka Hachiya, and Masashi Sugiyama. Relative density-ratio estimation for robust distribution comparison. *Neural computation*, 25(5):1324–1370, 2013. [6](#)
- [147] Hang Yang and Simon Fong. Incrementally optimized decision tree for noisy big data. In *Proceedings of the 1st International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications*, pages 36–44, 2012. [20](#)
- [148] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021. [1](#)
- [149] Luyu Yang, Yogesh Balaji, Ser-Nam Lim, and Abhinav Shrivastava. Curriculum manager for source selection in multi-source domain adaptation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 608–624. Springer, 2020. [20](#)
- [150] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9593–9602, 2021. [20](#)
- [151] Ting Yao, Yingwei Pan, Chong-Wah Ngo, Houqiang Li, and Tao Mei. Semi-supervised domain adaptation with subspace learning for visual recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2142–2150, 2015. [20](#)
- [152] Haotian Ye, Chuanlong Xie, Tianle Cai, Ruichen Li, Zhen-guo Li, and Liwei Wang. Towards a theoretical framework of out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34:23519–23531, 2021. [20](#)
- [153] Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and deep learning to understand economic well-being in africa. *Nature Communications*, 2020. [8](#)
- [154] Rui Yu, Zhiyong Dou, Song Bai, Zhaoxiang Zhang, Yongchao Xu, and Xiang Bai. Hard-aware point-to-set deep metric for person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 188–204, 2018. [6](#), [7](#)
- [155] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017. [1](#), [8](#), [16](#)
- [156] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pages 819–827. PMLR, 2013. [2](#)
- [157] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyang Shen. Deep stable learning for out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5372–5382, 2021. [20](#)
- [158] Yuhong Zhang, Guang Chu, Peipei Li, Xuegang Hu, and Xindong Wu. Three-layer concept drifting detection in text data streams. *Neurocomputing*, 260:393–403, 2017. [19](#)
- [159] Yan Zhang, Jonathon Hare, and Adam Prugel-Bennett. Deep set prediction networks. *Advances in Neural Information Processing Systems*, 32, 2019. [1](#)
- [160] Sicheng Zhao, Bo Li, Xiangyu Yue, Yang Gu, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source domain adaptation for semantic segmentation. *Advances in Neural Information Processing Systems*, 32, 2019. [8](#)
- [161] Sicheng Zhao, Guangzhi Wang, Shanghang Zhang, Yang Gu, Yaxian Li, Zhichao Song, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source distilling domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12975–12983, 2020. [20](#)

- [162] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 20
- [163] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. 20
- [164] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020. 20

A. Sample items from the SHIFT15M

Figure 15,16,17,18 show the additional sample items from the SHIFT15M dataset.

In addition, we provide the year-wise visualization for the SHIFT15M dataset with t-SNE in Figure 19,20,21 like as Figure 1. t-SNE is a dimensionality reduction technique that is often used for visualizing high-dimensional data in a lower-dimensional space. In this case, the data being visualized is likely a collection of fashion-related features, such as color, texture, and style, that have been extracted from the dataset. These figures are likely to show the distribution of fashion-related features across different years or time periods, which can help to reveal trends and patterns in the fashion industry over time.

B. Details and additional figures for the numerical experiments

Here, we describe the details for numerical experiments and introduce additional figures.

B.1. Details for numerical experiments

Training Settings We use a stochastic gradient descent method [106, 61] with a learning rate of 0.005, a momentum of 0.5, and a weight decay of 0.00004. We train both the CNN and set-matching model in an end-to-end manner. In each iteration, we randomly swap pairs of sets and items in each set, and randomly flip images horizontally, to learn all the methods stably.

Preparing Set Pairs To construct the correct pair of sets to be matched, we randomly halve the given outfit \mathcal{O} into two non-empty proper subsets \mathcal{X} and \mathcal{Y} as follows: $\mathcal{O} \rightarrow \{\mathcal{X}, \mathcal{Y}\}$, where $\mathcal{X} \cap \mathcal{Y} = \emptyset$. Here, we extend this setting to include more general situations. We select Q outfits $\mathcal{O}^{(1)}, \dots, \mathcal{O}^{(Q)}$ randomly and split the respective outfits in half $\mathcal{O}^{(q)} \rightarrow \{\mathcal{X}^{(q)}, \mathcal{Y}^{(q)}\}$, where $q \in \{1, \dots, Q\}$. We regard the two sets $\{\mathcal{X}^{(1)}, \dots, \mathcal{X}^{(Q)}\}$ and $\{\mathcal{Y}^{(1)}, \dots, \mathcal{Y}^{(Q)}\}$ as the correct pair, which consists of Q fashion styles. In the training phase, we set $Q = 4$. Figure 14 shows the overview of set-to-set matching problem.

B.2. Additional figures for the numerical experiments

Figure 22,23,24,25,26 show the additional figures for the Fill-In-The-N Blank experiments. These figures are based on Table 4 and 5.

Figure 27,28 and 29 shows the PCA [35] dimensionality reduction for items in each year. The color of each point corresponds to the item category.

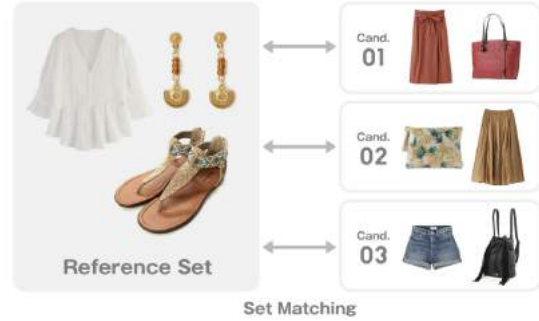


Figure 14: This figure is a citation of Figure 1 from Y.Saito et al. [110]. Set-to-set matching aims to answer a fundamental question: which candidate set is more compatible with the reference set than others? In this process, we match the reference set with each candidate set and select the best pair based on some criteria.

C. Details for the sum-decomposable set function

From Definition 4.1, we say that (ρ, ϕ) is a sum-decomposition of f . Given sum-decomposition (ρ, ϕ) , we write $\Phi(\mathcal{X}) = \sum_{x \in \mathcal{X}} \phi(x)$ and $f = \rho \circ \Phi$. We may also refer to the function $\rho \circ \Phi$ as a sum-decomposition.

Definition C.1. Let (ρ, ϕ) be a sum-decomposition. Write Z for the domain of ρ and the codomain of ϕ . We refer to Z as the latent space of the sum-decomposition (ρ, ϕ) .

Definition C.2. Given a space Z , we say that f is sum-decomposable via Z if f has a sum-decomposition whose latent space is Z .

Definition C.3. We say that f is continuously sum-decomposable when there exists a sum-decomposition (ρ, ϕ) of f such that both ρ and ϕ are continuous. (ρ, ϕ) is then a continuous sum-decomposition of f .

We give a brief reproduction of the statements and proof of two key theorems from [155].

Theorem C.1. Let $f : 2^{\mathcal{U}} \rightarrow \mathbb{R}$ where \mathcal{U} is countable. Then f is sum-decomposable via \mathbb{R} .

Proof. Since \mathcal{U} is countable, each $x \in \mathcal{U}$ can be mapped to a unique element in \mathbb{N} by a bijective function $c : \mathcal{U} \rightarrow \mathbb{N}$. If we can choose ϕ so that Φ is invertible, then we can set $\rho = f \circ \Phi^{-1}$, giving

$$f = \rho \circ \Phi, \quad (9)$$

so f is sum-decomposable via \mathbb{R} .

Considering the mapping $\phi(x) = 4^{-c(x)}$, each set $\mathcal{X} \subset \mathcal{U}$ corresponds to a unique real number $r := \Phi(\mathcal{X})$. The

number r can be decoded to the set \mathcal{X} and the element $c^{-1}(n) \in \mathcal{U}$ belongs to \mathcal{X} if and only if the n -th digit of r is 1. This decoding procedure shows that Φ is invertible, and the conclusion follows. \square

Theorem C.2. Let $M \in \mathbb{N}$, and let $f : [0, 1]^M \rightarrow \mathbb{R}$ be a continuous permutation-invariant function. Then, f is continuously sum-decomposable via \mathbb{R}^{M+1} .

Theorem C.3. Deep Sets can represent any continuous permutation-invariant function of M elements if the dimension of the latent space of the model is at least $M + 1$.

Recent work [137] argue that the guarantee of sum-decomposability via \mathbb{R} given by Theorem C.1 cannot hold in practice, and prove that the guarantee of sum-decomposability via \mathbb{R}^{M+1} is essentially the bestpossible.

Theorem C.4. Let $M, N \in \mathbb{N}$, with $M > N$. Then there exist continuous permutation-invariant functions $f : \mathbb{R}^M \rightarrow \mathbb{R}$ which are not continuously sum-decomposable via \mathbb{R}^N .

Theorem C.5. Let $M \in \mathbb{N}$, and let $f : \mathbb{R}^M \rightarrow \mathbb{R}$ be a continuous permutation-invariant function. Then, f is continuously sum-decomposable via \mathbb{R}^M .

Theorem C.6. Denote the set of subsets $[0, 1]$ containing at most M elements by $[0, 1]^{\leq M}$. Let $f : [0, 1]^{\leq M} \rightarrow \mathbb{R}$ be continuous and permutation-invariant. Then, f is continuously sum-decomposable via \mathbb{R}^M .

Let M be a positive integer, $U \subset \mathbb{R}^M$ be compact, and $f : U \rightarrow \mathbb{R}$.

Definition C.4. Let $\epsilon > 0$. (ϕ, ρ) is a within- ϵ sum-decomposition of f if $|f(\mathcal{U}) - \rho(\Phi(\mathcal{U}))| < \epsilon$ for every $\mathcal{U} \in U$.

Definition C.5. A sequence $(\phi, \rho)_k := \{(\phi_k, \rho_k) : k \in \mathbb{N}\}$ is an approximate sum-decomposition of f if, for any $\epsilon > 0$, there is some $K \in \mathbb{N}$ such that (ϕ_K, ρ_K) is a within- ϵ sum-decomposition of f . We also require that (ϕ_k, ρ_k) is a sequence of ever-closer approximations to f . The existence of an approximate sum-decomposition of f guarantees that f can be approximated arbitrarily closely by sum-decomposition.

Theorem C.7. Let $M, N \in \mathbb{N}$ with $M > N$, and recall that $I_M = [-1, 1]^M \subset \mathbb{R}^M$. Then there exists a continuous permutation-invariant function $f : I_M \rightarrow \mathbb{R}$ which has no continuous approximate sum-decomposition via \mathbb{R}^N .

D. Generalization bounds under the distribution shift

The SHIFT15M dataset is a valuable resource for evaluating the performance of machine learning models in settings

where the underlying distribution of the data may shift over time. One important aspect of this dataset is that it allows researchers to calculate the distance between the distributions in the train and test splits. This distance measure can provide valuable insight into how much the distribution has shifted between the two sets of data, which in turn can help researchers to better understand the behavior of their models under distributional shifts.

Furthermore, researchers have proposed several generalization bounds that are specific to distribution shifts and that depend on the distance between distributions. These bounds provide a way to quantify the relationship between the performance of a model and the extent of the distributional shift, which is a crucial factor in evaluating the robustness of a model. By considering these generalization bounds when analyzing experimental results on the SHIFT15M dataset, researchers can gain a more precise understanding of how their models perform under distribution shifts and how to improve them. Overall, the SHIFT15M dataset and the associated generalization bounds represent important tools for evaluating and improving the robustness of machine learning models. Here, we introduce several known generalization bounds under the distribution shift [105].

Definition D.1 (Total variation distance). Denote by \mathcal{B} be the set of measurable subsets under two probability distributions p_1 and p_2 . Then, the total variation distance between p_1 and p_2 is defined as

$$d_{TV}(p_1, p_2) = 2 \sup_{B \in \mathcal{B}} |p_1(B) - p_2(B)|. \quad (10)$$

Let

$$\mathcal{R}_{\mathcal{D}}^{\ell}(h) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\ell(\mathbf{x}, y)]. \quad (11)$$

Theorem D.1 ([12]). Given two domains \mathfrak{S} and \mathfrak{T} , and a hypothesis class \mathcal{H} over $\mathcal{X} \times \mathcal{Y}$, the following holds for any $h \in \mathcal{H}$.

$$\begin{aligned} \mathcal{R}_{\mathfrak{T}}^{\ell_{01}}(h) &\leq \mathcal{R}_{\mathfrak{S}}^{\ell_{01}}(h) + d_{TV}(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}) \\ &+ \min \{ \mathbb{E}_{\mathbf{x} \sim \mathfrak{S}_{\mathcal{X}}} [J(\mathbf{x}; \mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}})], \mathbb{E}_{\mathbf{x} \sim \mathfrak{T}_{\mathcal{X}}} [J(\mathbf{x}; \mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}})] \}, \end{aligned}$$

where $J(\mathbf{x}; \mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}) = |f_{\mathfrak{S}}(\mathbf{x}) - f_{\mathfrak{T}}(\mathbf{x})|$, and $f_{\mathfrak{S}}(\mathbf{x})$ and $f_{\mathfrak{T}}(\mathbf{x})$ are source and target true labeling functions associated with \mathfrak{S} and \mathfrak{T} , respectively.

Definition D.2. Given marginal distributions of two domains $\mathfrak{S}_{\mathcal{X}}$ and $\mathfrak{T}_{\mathcal{X}}$ over the input space \mathcal{X} , let \mathcal{H} be a hypothesis class, and denote $\mathcal{H}\Delta\mathcal{H}$ the symmetric difference hypothesis space defined as

$$h \in \mathcal{H}\Delta\mathcal{H} \iff h(\mathbf{x}) = g(\mathbf{x}) \oplus g'(\mathbf{x}) \quad (12)$$

for some $(g, g') \in \mathcal{H}^2$, where \oplus stands for XOR operation. Let $I(h)$ denote the set for which $h \in \mathcal{H}\Delta\mathcal{H}$ is the characteristic function, that is $\mathbf{x} \in I(h) \iff g(\mathbf{x}) = 1$. The $\mathcal{H}\Delta\mathcal{H}$ -divergence between $\mathfrak{S}_{\mathcal{X}}$ and $\mathfrak{T}_{\mathcal{X}}$ is defined as:

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}) = 2 \sup_{h \in \mathcal{H}\Delta\mathcal{H}} \left| \Pr_{\mathfrak{S}}(I(h)) - \Pr_{\mathfrak{T}}(I(h)) \right|. \quad (13)$$

Lemma D.2. Let \mathcal{H} be a hypothesis space of VC dimension $VC(\mathcal{H})$. If S_u, T_u are unlabeled samples of size m each, drawn independently from $\mathfrak{S}_{\mathcal{X}}$ and $\mathfrak{T}_{\mathcal{X}}$ respectively, then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$ over the random choice of the samples we have

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}) \leq \hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathfrak{S}_u, \mathfrak{T}_u) + 4\sqrt{\frac{2VC(\mathcal{H}) \ln 2m + \ln \frac{2}{\delta}}{m}},$$

where $\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathfrak{S}_u, \mathfrak{T}_u)$ is the empirical $\mathcal{H}\Delta\mathcal{H}$ -divergence estimated on \mathfrak{S}_u and \mathfrak{T}_u .

Lemma D.3 ([11]). Let \mathcal{H} be a hypothesis space. Then, for two unlabeled samples S_u, T_u of size m we have

$$\begin{aligned} & \hat{d}_{\mathcal{H}\Delta\mathcal{H}}(S_u, T_u) \\ &= 2 \left\{ 1 - \min_{h \in \mathcal{H}\Delta\mathcal{H}} \left(\frac{1}{m} \sum_{\mathbf{x} \in H_0} 1_{\mathbf{x} \in S_u} + \frac{1}{m} \sum_{\mathbf{x} \in H_1} 1_{\mathbf{x} \in T_u} \right) \right\}, \end{aligned}$$

where $H_0 = \{\mathbf{x} : h(\mathbf{x}) = 0\}$ and $H_1 = \{\mathbf{x} : h(\mathbf{x}) = 1\}$.

Lemma D.4 ([11]). Let \mathfrak{S} and \mathfrak{T} be two domains on $\mathcal{X} \times \mathcal{Y}$. For any pair of hypothesis $(h, h') \in \mathcal{H}\Delta\mathcal{H}^2$, we have

$$\left| \mathcal{R}_{\mathfrak{T}}^{\ell_{01}}(h, h') - \mathcal{R}_{\mathfrak{S}}^{\ell_{01}}(h, h') \right| \leq \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}), \quad (14)$$

where

$$\mathcal{R}_{\mathfrak{D}}^{\ell}(h, h') = \mathbb{E}_{\mathbf{x} \sim \mathfrak{D}_{\mathcal{X}}} [\ell(h(\mathbf{x}), h'(\mathbf{x}))]. \quad (15)$$

Theorem D.5 ([11]). Let \mathcal{H} be a hypothesis space of VC dimension $VC(\mathcal{H})$. If S_u and T_u are unlabeled samples of size m' each, drawn independently from $\mathfrak{S}_{\mathcal{X}}$ and $\mathfrak{T}_{\mathcal{X}}$, respectively, then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$ over the random choice of the samples, we have that for all $h \in \mathcal{H}$

$$\begin{aligned} \mathcal{R}_{\mathfrak{T}}^{\ell_{01}}(h) &\leq \mathcal{R}_{\mathfrak{S}}^{\ell_{01}}(h) + \frac{1}{2} \hat{d}_{\mathcal{H}\Delta\mathcal{H}}(S_u, T_u) \\ &\quad + 4\sqrt{\frac{2VC(\mathcal{H}) \ln 2m' + \ln \frac{2}{\delta}}{m'}} + \lambda, \end{aligned}$$

where λ is the combined error of the ideal hypothesis h^* .

The trade-off between source risk, divergence and the capability to adapt for distribution shift is a crucial phenomenon that plays a significant role in experiments related to distribution shift adaptation using SHIFT15M. This trade-off refers to the balance that needs to be maintained between the risk of the source, the extent of divergence between the source and target domains, and the ability of a model to adapt to distribution shift. Therefore, it is essential to evaluate the performance of models while keeping this trade-off in mind. By doing so, researchers can ensure that the models they develop are not only accurate but also robust and adaptable

to changes in the underlying distribution of the data. Ultimately, this can lead to better real-world performance of machine learning models and improve their utility in various applications.

Since our numerical experiments use the Wasserstein distance W_p^p as the distance between distributions, we also introduce generalization bounds based on this distance:

$$W_p^p(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}) := \inf_{\gamma \in \Pi(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}})} \int_{\mathcal{X} \times \mathcal{X}} c(\mathbf{x}, \mathbf{x}')^p d\gamma(\mathbf{x}, \mathbf{x}'),$$

where $c : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ is a cost function for transporting one unit of mass \mathbf{x} to \mathbf{x}' , and $p \in [1, +\infty]$.

Lemma D.6 ([104]). Let $\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}} \in \mathcal{P}(\mathcal{X})$ be two probability measures on \mathbb{R}^d . Assume that the cost function $c(\mathbf{x}, \mathbf{x}') = \|\phi(\mathbf{x}) - \phi(\mathbf{x}')\|_{\mathfrak{H}_{k_{\ell}}}$, where \mathfrak{H} is an RKHS equipped with kernel $k_{\ell} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ induced by $\phi : \mathcal{X} \rightarrow \mathfrak{H}_{k_{\ell}}$ and $k_{\ell}(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathfrak{H}_{k_{\ell}}}$. Assume further that the loss function $\ell_{h,f} : \mathbf{x} \mapsto \ell(h(\mathbf{x}), f(\mathbf{x}))$ is convex, symmetric and bounded and obeys the triangular equality and has the parametric form $|h(\mathbf{x}) - f(\mathbf{x})|^q$ for some $q > 0$. Assume also that kernel k_{ℓ} in the RKHS $\mathfrak{H}_{k_{\ell}}$ is square-root integrable with respect to both $\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}$ for all $\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}} \in \mathcal{P}(\mathcal{X})$ where \mathcal{X} is separable and $0 \leq k_{\ell}(\mathbf{x}, \mathbf{x}') \leq K, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$ if $\|\ell\|_{\mathfrak{H}_{k_{\ell}}} \leq 1$, then the following holds:

$$\mathcal{R}_{\mathfrak{T}}^{\ell_q}(h, h') \leq \mathcal{R}_{\mathfrak{S}}^{\ell_q}(h, h') + W_1(\mathfrak{S}_{\mathcal{X}}, \mathfrak{T}_{\mathcal{X}}), \quad (16)$$

for all $(h, h') \in \mathfrak{H}_{k_{\ell}}^2$.

This lemma allows us to relate the source and target errors by using Wasserstein distance. Also, we have

$$\begin{aligned} \|\phi(\mathbf{x}) - \phi(\mathbf{x}')\|_{\mathfrak{H}} &= \sqrt{\langle \phi(\mathbf{x}) - \phi(\mathbf{x}'), \phi(\mathbf{x}) - \phi(\mathbf{x}') \rangle_{\mathfrak{H}}} \\ &= \sqrt{k(\mathbf{x}, \mathbf{x}) - 2k(\mathbf{x}, \mathbf{x}') + k(\mathbf{x}', \mathbf{x}')}. \end{aligned}$$

Theorem D.7 ([15]). Let μ be a probability measure in \mathbb{R}^d so that for some $\alpha > 0$, $\int_{\mathbb{R}^d} \exp\{\alpha \|\mathbf{x}\|^2\} d\mu < \infty$, and $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{x}_i}$ be its associated empirical measure defined on a sample of independent variables $\{\mathbf{x}_i\}_{i=1}^N$ drawn from μ . Then for any $d' > d$ and $\zeta' < \sqrt{2}$ there exists some constant N_0 depending on d' and some square exponential moment of μ such that, for any $\epsilon > 0$ and $N \geq N_0 \max(\epsilon^{-(d'+2)}, 1)$,

$$\Pr \{W_1(\mu, \mu') > \epsilon\} \leq \exp \left\{ -\frac{\zeta'}{2} N \epsilon^2 \right\} \quad (17)$$

where d' and ζ' can be calculated explicitly.

Theorem D.8. Under the assumption of Lemma D.6, let S_u and T_u be two samples of size N_S and N_T drawn i.i.d. from $\mathfrak{S}_{\mathcal{X}}$ and $\mathfrak{T}_{\mathcal{X}}$, respectively. Let $\hat{\mathfrak{S}}_{\mathcal{X}} = \frac{1}{N_S} \sum_{i=1}^{N_S} \delta_{\mathbf{x}_i^S}$ and $\hat{\mathfrak{T}}_{\mathcal{X}} = \frac{1}{N_T} \sum_{i=1}^{N_T} \delta_{\mathbf{x}_i^T}$ be the associated empirical measures. Then for any $d' > d$ and $\zeta' < \sqrt{2}$, there exists some

constant N_0 depending on d' such that for any $\delta > 0$ and $\min(N_S, N_T) \geq N_0 \max(\delta^{-(d'+2)}, 1)$ with probability at least $1 - \delta$ for all h , we have

$$\begin{aligned} \mathcal{R}_{\mathfrak{X}}^{\ell_q}(h) &\leq \mathcal{R}_{\mathfrak{G}}^{\ell_q}(h) + W_1(\hat{\mathfrak{G}}_{\mathcal{X}}, \hat{\mathfrak{X}}_{\mathcal{X}}) \\ &\quad + \sqrt{\frac{2 \ln \frac{1}{\delta}}{\zeta'}} \left(\sqrt{\frac{1}{N_S}} + \sqrt{\frac{1}{N_T}} \right) + \lambda, \end{aligned}$$

where λ is the combined error of the ideal hypothesis h^* that minimizes the combined error of $\mathcal{R}_{\mathfrak{G}}^{\ell_q}(h) + \mathcal{R}_{\mathfrak{X}}^{\ell_q}(h)$.

E. Density ratio estimation for the importance weighted set-to-set matching

In the experiments, we proposed the importance-weighted set-to-set matching as the baseline method. Recall that, under the covariate shift assumption, we have

$$\begin{aligned} \mathbb{E}_{tr} \left[\frac{p_{te}(\mathbf{x})}{p_{tr}(\mathbf{x})} \ell(h(\mathbf{x}), y) \right] &= \int \frac{p_{te}(\mathbf{x})}{p_{tr}(\mathbf{x})} \ell(h(\mathbf{x}), y) p_{tr}(\mathbf{x}, y) d\mathbf{x} dy \\ &= \int \ell(h(\mathbf{x}), y) p_{te}(\mathbf{x}, y) d\mathbf{x} dy \\ &= \mathbb{E}_{te} [\ell(h(\mathbf{x}), y)]. \end{aligned}$$

Density ratio estimation is a technique used in machine learning to quantify the difference between the distributions of two datasets. The goal of density ratio estimation is to estimate the density ratio $r(\mathbf{x})$, which is defined as the ratio of the probability density functions of the test distribution $p_{te}(\mathbf{x})$ and the training distribution $p_{tr}(\mathbf{x})$ as $r(\mathbf{x}) = p_{te}(\mathbf{x})/p_{tr}(\mathbf{x})$.

The density ratio is a powerful tool because it allows us to compare the distributions of the two datasets and measure the extent of the distribution shift. In particular, if we can estimate the density ratio accurately, we can obtain a consistent estimator under the distribution shift. This means that we can use this estimator to accurately predict the performance of our model on the test set, even if the distribution of the test set is different from that of the training set.

To estimate the density ratio, we used a probabilistic classifier. However, there are several other strategies that exist for estimating density ratios. Each of these methods has its own strengths and weaknesses, and the choice of method will depend on the specific application and the characteristics of the data.

The main idea of moment matching is, to match the moments of $\hat{p}_{te}(\mathbf{x}) = \hat{r}(\mathbf{x})p_{tr}(\mathbf{x})$ and $p_{te}(\mathbf{x})$. For example, matching the mean is

$$\int \mathbf{x} \hat{r}(\mathbf{x}) p_{tr}(\mathbf{x}) d\mathbf{x} = \int \mathbf{x} p_{te}(\mathbf{x}) d\mathbf{x}. \quad (18)$$

However, matching a finite number of moments does not necessarily yield the true density ratio even asymptotically.

Kernel mean matching [53, 46] allows that All moments are efficiently matched in Gaussian RKHS \mathfrak{H} :

$$\min_{\hat{r} \in \mathfrak{H}} \left\| \int K(\mathbf{x}, \cdot) \hat{r}(\mathbf{x}) p_{tr}(\mathbf{x}) d\mathbf{x} - \int K(\mathbf{x}, \cdot) p_{tr}(\mathbf{x}) d\mathbf{x} \right\|_{\mathfrak{H}}^2.$$

KLIEP [92, 122] minimize KL-divergence from $p_{te}(\mathbf{x})$ to $\hat{p}_{te}(\mathbf{x}) = \hat{r}(\mathbf{x})p_{tr}(\mathbf{x})$ as

$$\min_{\hat{r}} D_{KL}[p_{te}(\mathbf{x}) \parallel \hat{p}_{te}] = \min_{\hat{r}} \int p_{te}(\mathbf{x}) \frac{p_{te}(\mathbf{x})}{\hat{r}(\mathbf{x})p_{tr}(\mathbf{x})} d\mathbf{x}.$$

Least-Squares Importance Fitting (LSIF) [57] minimize squared loss:

$$\min_{\hat{r}} \int (\hat{r}(\mathbf{x}) - r(\mathbf{x}))^2 p_{tr}(\mathbf{x}) d\mathbf{x}.$$

F. Other related literature

Here we present the rest of the relevant literature.

F.1. Concept drift

In addition to covariate shift and target shift, the following concept drift [128, 39, 82] is also well known.

Definition F.1 (Concept drift). We consider that the two distributions $p_{tr}(\mathbf{x}, y)$ and $p_{te}(\mathbf{x}, y)$ satisfy the concept drift assumption if the following conditions hold:

$$\begin{aligned} p_{tr}(\mathbf{x}|y) &\neq p_{te}(\mathbf{x}|y), \\ p_{tr}(y|\mathbf{x}) &\neq p_{tr}(y|\mathbf{x}). \end{aligned}$$

Unlike covariate shift and target shift, concept drift assumes that the conditional probabilities between the two distributions are different. This means that a model that is well-specified in the training distribution will be misspecified in the test distribution, making it the most difficult problem setup. The strategies for addressing concept drift can be broadly categorized as follows

- concept drift detection;
- concept drift understanding;
- concept drift adaptation.

Concept drift detection refers to the strategies that characterize and quantify concept drift via identifying change points or change time intervals [9]. Many algorithms focus on tracking changes in the online error rate of base classifiers, such as Drift Detection Method (DDM) [37], LLDD [36], EDDM [6], HDDM [34] or FW-DDM [79]. Other strategies based on data distribution [84, 120, 83] or multiple hypothesis test [2, 139, 158].

Concept drift understanding refers to retrieving concept drift information about "When" (the time at which the concept drift occurs and how long the drift lasts) [115, 84],

"How" (the severity/degree of concept drift) [94], and "Where" (the drift regions of concept drift) [78].

The main approaches for concept drift adaptation are training new models for global drift [5, 86], model ensemble for recurring drift [44, 69, 29], or adjusting existing models for regional drift [54, 38, 147].

F.2. Domain adaptation

Domain adaptation [23, 22, 40, 111] is often used in a similar context to distribution shift adaptation. It is often referred to as visual domain adaptation [33, 55, 140, 98], especially in the field of computer vision. This concept is often referred to indistinguishably from covariate shift. Depending on the availability of the source and target domain data, the domain adaptation problem can be defined in many different ways.

- **supervised domain adaptation:** In supervised domain adaptation [89, 129], labeled data is available in both the source and target domains. The model is trained on the labeled data from the source domain and then adapted to the target domain using the labeled data from the target domain.
- **semi-supervised domain adaptation:** In semi-supervised domain adaptation [71, 109, 28, 151], a small amount of labeled data is available in the target domain in addition to the labeled data in the source domain. The model is trained on the labeled data from both domains and then adapted to the target domain using both the labeled and unlabeled data from the target domain.
- **unsupervised domain adaptation:** In unsupervised domain adaptation [40, 58, 77, 113, 81], labeled data is only available in the source domain. The model is trained on the labeled data from the source domain and then adapted to the target domain using only the unlabeled data from the target domain.
- **multi-source domain adaptation:** Unlike traditional domain adaptation, which involves adapting from a single source domain, multi-source domain adaptation [99, 149, 161] deals with situations where there are multiple source domains with different but related feature distributions.

F.3. Out-of-distribution generalization

Another similar concept is out-of-distribution generalization [4, 48, 76, 117, 152]. The main difference with distribution shift problem is that in the out-of-distribution generalization problem setting we do not have any access to the test data during training. For example, under the covariate shift hypothesis, we often assume that we are allowed to access unlabeled test data. Major approaches

for out-of-distribution generalization are disentangled representation learning [13, 51, 62], causal representation learning [150, 112], domain generalization [162, 75, 163, 74], invariant learning [100, 107, 47], stable learning [116, 157] and distributionally robust optimization [103, 25, 43].

F.4. Other related topics

Active learning Active learning [114, 101, 66] is a subfield of machine learning that focuses on developing algorithms that can automatically select the most informative data to be labeled by an expert or a human annotator. The goal of active learning is to achieve high accuracy models with minimal labeled data. Active learning can help mitigate the effects of distribution shift by actively selecting the most informative data points to be labeled by an expert or a human annotator. By doing so, the active learning algorithm can ensure that the labeled data used to train the model is representative of the data that the model will encounter in the real world.

Continual learning Continual learning [96], also known as lifelong learning or incremental learning, is a subfield of machine learning that deals with the problem of learning from a continuous stream of data over time, without forgetting previously learned knowledge. One of the key challenges in continual learning is avoiding catastrophic forgetting [67, 59], which occurs when a model overwrites previously learned knowledge with new information. Both catastrophic forgetting and distribution shift can lead to poor model performance and may require the model to be re-trained or updated to address these issues. Continual learning algorithms are designed to address these challenges by enabling models to learn from a continuous stream of data over time, without forgetting previously learned knowledge and without being affected by distribution shift.

Transfer learning Transfer learning [141, 127, 164] is a machine learning technique that involves leveraging knowledge learned from one task to improve performance on a different, but related task. In transfer learning, a model is first trained on a source task, which provides a foundation of knowledge and skills. Then, the model is fine-tuned on a target task, which is related to the source task but may have different characteristics. Transfer learning is motivated by the fact that many tasks in machine learning share common features and patterns. By leveraging knowledge from a related task, a model can learn to recognize patterns and features more effectively, even if the target task has different characteristics.

F.5. Fashion datasets

Several fashion datasets have been introduced in the research community to facilitate the development and eval-

uation of fashion-related machine learning models. These datasets contain various types of fashion-related data, including images, textual descriptions, and attribute labels.

Fashion MNIST [144] Fashion MNIST is a publicly available dataset of Zalando’s article images that is widely used for research and development in computer vision and machine learning. The dataset consists of a training set of 60,000 examples and a test set of 10,000 examples, each of which is a 28×28 grayscale image. The images in the dataset are associated with labels from 10 different classes.

DeepFashion [80] DeepFashion is a dataset containing around 800K diverse fashion images with their rich annotations (46 categories, 1,000 descriptive attributes, bounding boxes and landmark information) ranging from well-posed product images to real-world-like consumer photos.

Fashionpedia [56] Fashionpedia consists of user uploaded 48K street-fashion photos collected from free license websites such as Unsplash, Kaboompics etc. These photos contain people wearing variety of clothes and accessories captured in different background, weather and camera conditions.

Polyvore [133] Polyvore is a crowd-sourced dataset containing outfits or sets of fashion items that complement each other. It consists of manually labeled 68K outfits that are split into 53K, 10K and 5K into training, validation and testing sets, respectively

Polyvore-disjoint [133] Polyvore-disjoint is a subset of the Polyvore dataset created by removing outfits that have common items between training, validation and testing sets. The dataset is challenging compared to Polyvore dataset and consists of 32K outfits. During inference, we use only product images patches and do not use the metadata associated with the products.

Fashion IQ [142] The Fashion IQ dataset is a collection of images and associated metadata designed to facilitate research on natural language-based interactive image retrieval in the fashion domain. It is the fashion dataset that includes human-written relative captions that have been annotated for similar pairs of images, as well as real-world product descriptions and attribute labels as side information. The dataset was created to help researchers develop new methods for retrieving fashion images using natural language queries.

G. Open questions

Learning guarantees for set-to-set matching under the distribution shift As introduced in Appendix D, there

are a number of studies on generalized error analysis under distribution shifts for ordinal classification and regression problems. However, theoretical analysis of set matching under distributional shifts remains unexplored. Moreover, even under the i.i.d. assumption, there is little research on theoretical analysis of set matching [64].

Correlation of performance on SHIFT15M dataset with performance on other different datasets As introduced in Section 4, there are many datasets for classification and regression under distribution shifts. Since SHIFT15M provides data loaders for ordinary classification and regression as well as set matching, it is useful to examine the correlation between performance on these tasks and performance on other data sets.

2014



Figure 15: Additional sample items from 2014.

2015



Figure 16: Additional sample items from 2015.

2016



Figure 17: Additional sample items from 2016.

2017



Figure 18: Additional sample items from 2017.



Figure 19: t-SNE visualization for the SHIFT15M (2015).




Figure 20: t-SNE visualization for the SHIFT15M (2016).

2017



Figure 21: t-SNE visualization for the SHIFT15M (2017).




Figure 22: Plots for Fill-In-The-N-Blank experiments with Cross Attention.




Figure 23: Plots for Fill-In-The-N-Blank experiments with Set Transformer.

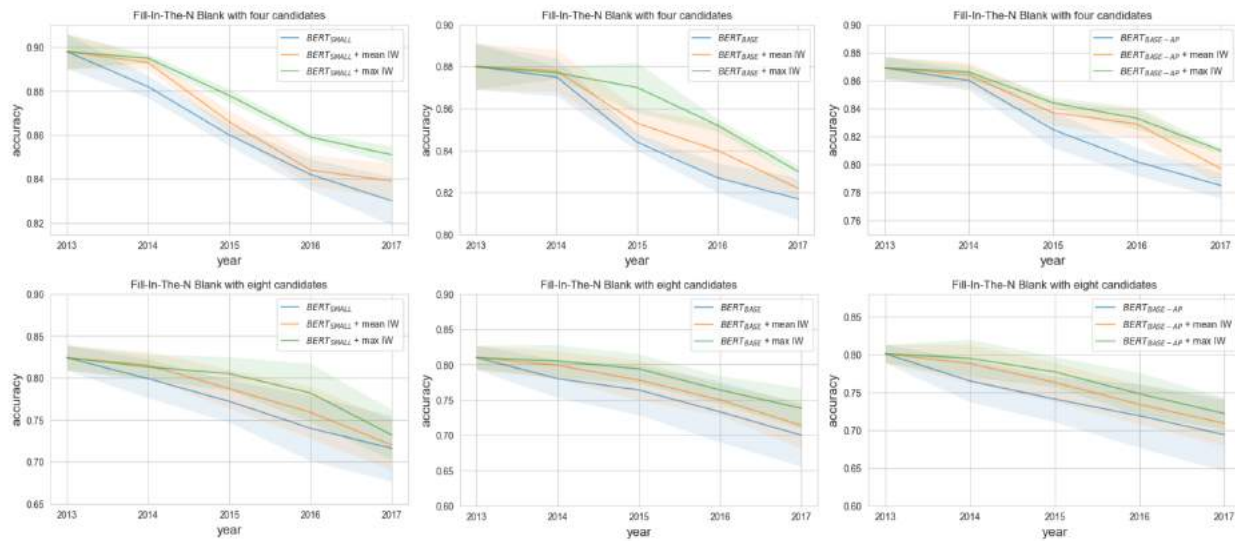


Figure 24: Plots for Fill-In-The-N-Blank experiments with BERT.

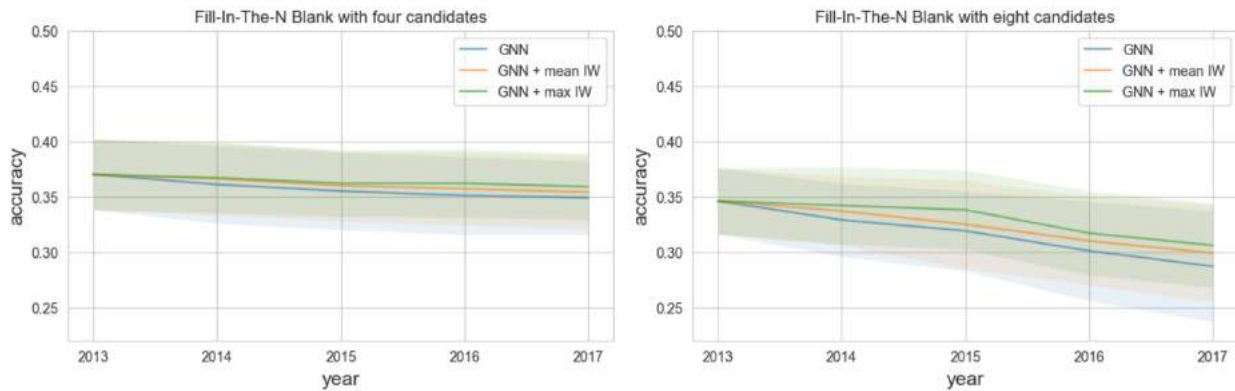


Figure 25: Plots for Fill-In-The-N-Blank experiments with GNN.




Figure 26: Plots for Fill-In-The-N-Blank experiments with HAP2S.

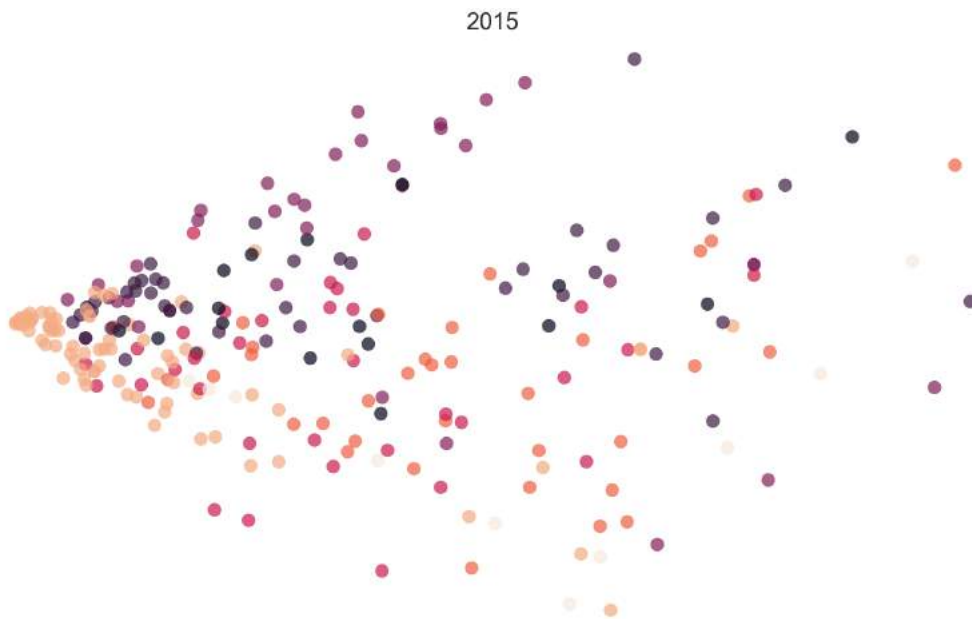


Figure 27: PCA for the SHIFT15M (2015). The color of each point corresponds to the item category.

2016



Figure 28: PCA for the SHIFT15M (2016). The color of each point corresponds to the item category.

2017



Figure 29: PCA for the SHIFT15M (2017). The color of each point corresponds to the item category.

H. Datasheet for SHIFT15M

In accordance with [41], we provide the datasheet for SHIFT15M. This datasheet offers an overview of crucial information about the dataset that can aid users in making informed decisions regarding its use. It encompasses a comprehensive range of details, including but not limited to the dataset's creation purpose, sources of data, and data processing methods. Overall, the datasheet provides a vital resource for those interested in utilizing the SHIFT15M dataset, as it offers a transparent and comprehensive account of the dataset's characteristics and construction.

Motivation The purpose of this section is to emphasize the importance of transparency and clarity in the process of dataset creation, particularly with regards to the motivations behind the creation of the dataset and any potential conflicts of interest that may arise. Dataset creators are encouraged to clearly articulate their reasons for creating the dataset, including the research questions or goals that the dataset is intended to address.

Composition Dataset creators are advised to review these questions thoroughly prior to commencing any data collection, and provide responses after data collection has concluded. The primary purpose of the questions in this section is to equip dataset consumers with the information necessary to make informed decisions about utilizing the dataset for their specific needs. Additionally, some of the questions have been tailored to obtain information about adherence to the General Data Protection Regulation (GDPR) of the European Union or similar regulatory frameworks in other jurisdictions. Notably, questions specific to datasets involving individuals have been consolidated towards the end of the section. It is recommended a broad interpretation of what constitutes a dataset relating to people. For instance, any dataset containing text that was produced by individuals can be considered to relate to people.

Collection Process To ensure potential issues are identified, dataset creators are advised to review the questions in this section before initiating data collection and then furnish responses upon completion of collection, similar to the previous section. Apart from the objectives set out in the preceding section, the questions in this section are aimed at extracting information that could assist researchers and practitioners in producing alternative datasets with comparable attributes. Similarly, inquiries exclusive to datasets pertaining to individuals are categorized towards the end of this section.

Preprocessing/cleaning/labeling Before proceeding with any preprocessing, cleaning, or labeling, it is recommended

that dataset creators review the questions presented in this section and subsequently provide responses upon completing these tasks. The purpose of the questions in this section is to equip dataset consumers with the requisite information to evaluate whether the "raw" data has been processed in a manner that aligns with their intended use.

Uses The questions in this section serve to prompt dataset creators to consider the appropriate and inappropriate uses of their dataset. By explicitly specifying such tasks, dataset creators can assist dataset consumers in making informed decisions, minimizing potential hazards or adverse outcomes.

Distribution Dataset creators should provide answers to these questions prior to distributing the dataset either internally within the entity on behalf of which the dataset was created or externally to third parties.

Maintenance As with the questions in the previous section, dataset creators should provide answers to these questions prior to distributing the dataset. The questions in this section are intended to encourage dataset creators to plan for dataset maintenance and communicate this plan to dataset consumers.

Motivation

For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

This paper addresses the problem of set-to-set matching, which involves matching two different sets of items based on some criteria, especially in the case of high-dimensional items like images. Although neural networks have been applied to solve this problem, most machine learning-based approaches assume that the training and test data follow the same distribution, which is not always true in real-world scenarios. To address this limitation, we introduce SHIFT15M, a dataset that can be used to evaluate set-to-set matching models when the distribution of data changes between training and testing. We conduct benchmark experiments that demonstrate the performance drop of naive methods due to distribution shift. Additionally, we provide software to handle the SHIFT15M dataset in a simple manner, with the URL for the software to be made available after publication of this manuscript. We believe proposed SHIFT15M dataset provide a valuable resource for evaluating set-to-set matching models under the distribution shift.

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

The initial version of the dataset was created by Masanari Kimura, Yuki Saito, Kazuya Morishita, Ryosuke Goto, and Takuma Nakamura most of whom were researchers at the ZOZO Research.

Who funded the creation of the dataset? If there is an associated grant, please provide the name of the grantor and the grant name and number.

Not applicable.

Any other comments?

Composition

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

The SHIFT15M dataset is an extensive collection of outfits that were previously shared on a fashion-oriented social networking service. The service is no longer available, but the dataset continues to be a valuable resource for researchers studying set-to-set matching problems, particularly in the context of fashion. The dataset contains a vast array of information about each outfit, including details about the user who posted it and some meta-information. Each record in the dataset comprises five different fields, making it easy to organize and analyze the data.

- `set_id`: An ID that identifies the outfit that was posted.
- `items`: Provides information about the items that comprise the posted outfit and consists of 4 subfields.
 - `item_id`: An ID that identifies an item.
 - `category_id1`: An ID indicating the item category (e.g., outerwear, tops, ...).
 - `category_id2`: An ID indicating the item subcategory (e.g., T-shirts, blouses, ...).

- `price`: Price of the item. `user`: Provides information about the user who posted the outfit and consists of 2 subfields. An ID that identifies the user who posted the outfit. A list of brands that users have voted for as their favorites. The number is an ID that identifies the brand.

- `like_num`: the number of times this outfit has been favorited by other users.
- `publish_date`: The date the outfit was posted.

How many instances are there in total (of each type, if appropriate)?

The dataset consists of 15,218,721 item images and 2,555,147 outfits which created by users of our fashion SNS.

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set? If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).

The fashion-oriented social networking service from which we collected outfits for the SHIFT15M dataset was a rich source of data that allowed us to obtain insights into the fashion trends and preferences of millions of users. With approximately 2 million users, the website was a bustling hub of activity where fashion enthusiasts could share their outfit ideas, provide feedback to others, and explore the latest styles and trends.

The vast majority of users on the website were women in their 20s and 30s, representing a demographic that is known for their fashion-forward mindset and interest in new trends. This demographic was particularly valuable for our research, as it allowed us to obtain a large amount of data on outfits that were representative of current fashion trends.

Our collection period spanned over a decade, starting on January 1st, 2010 and ending on April 6th, 2020. During this time, we meticulously gathered outfits consisting of multiple items, each of which was carefully categorized into a specific category. Our focus was on outfits that contained four or more items from the main categories, including outerwear, tops, bottoms, shoes, bags, hats, and accessories. This selection criterion was chosen with the aim of creating a dataset that would be useful for set-to-set matching tasks involving fashion items.

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images) or features? In either case, please provide a description.

Each item consists of 4096-dimensional features extracted via the VGG16 model trained using the ILSVRC2012 dataset.

Is there a label or target associated with each instance? If so, please provide a description.

Indeed, each instance in the SHIFT15M dataset contains a wealth of information that can be leveraged for various tasks. Along with the outfit items, each instance also includes several numerical values such as the category ID and number of likes. These values provide additional context that can be used to train models for various set-to-set matching problems.

One of the strengths of the SHIFT15M dataset is its versatility. By choosing one of these numerical values as the target variable, researchers can easily switch between several tasks, each with its own unique set of challenges and opportunities. For example, if the focus is on predicting outfit popularity, the number of likes can be used as the target variable. On the other hand, if the goal is to perform set-to-set matching between outfits, the number of likes can be used as the target variable.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

It is important to note that the SHIFT15M dataset only includes items that belong to the main categories, such as outerwear, tops, bottoms, shoes, bags, hats, and accessories. This means that items outside of these categories, such as underwear or background images for collage, are missing from the dataset.

While this selection criterion may seem limiting, it was chosen to ensure that the dataset is focused on items that are most commonly used in fashion-related set-to-set matching tasks. By excluding items outside of the main categories, we were able to curate a dataset that is more manageable and less noisy, while still providing a diverse range of fashion items for analysis.

Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)? If so, please describe how these relationships are made explicit.

Each instance is assigned the ID of the user who submitted the outfit.

Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.

SHIFT15M is a valuable dataset that simulates various types of dataset shifts that are commonly observed in real-world applications. The collected data spans a decade, from 2010 to 2020, and encompasses various shifts that arise due to factors such as changes in user behavior, fashion trends, and cultural preferences. To make it easy for researchers to evaluate their models on the SHIFT15M dataset, we have developed software that allows them to experiment with different types and sizes of shifts. The software automates the train/val/test splitting process, making it easier for researchers to evaluate the performance of their models under various shift scenarios. With this software, researchers can simulate shifts that arise due to various factors and assess their models' robustness to such shifts. By doing so, they can gain insights into how their models perform in real-world settings where data distributions are constantly changing.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.

No.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

The dataset is self-contained.

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals non-public communications)? If so, please provide a description.

No.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

No.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

Yes. Each instance is a combination of outfits created by an individual and preferred by that individual.

Does the dataset identify any subpopulations (e.g., by age, gender)? If so, please describe how these subpopulations are identified and provide a description of their respective distributions within the dataset.

No.

Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset? If so, please describe how.

It is impossible to identify individuals from the dataset.

Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)? If so, please provide a description.

No.

Any other comments?

Collection Process

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

Except for the item attributes, the data was generated by users. Item attributes (category and price) were collected from e-commerce sites that sell the item. All data was viewable on the website.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)? How were these mechanisms or procedures validated?

The fashion-oriented social networking service that we collected the SHIFT15M dataset from provided its users with an easy-to-use outfit editor that allowed them to create and publish their outfits on the platform. This editor featured a wide range of clothing items, accessories, and other fashion-related items that users could choose from to create their outfits. Once a user had selected the items they wanted to include in their outfit, the editor registered this selection as a new outfit on the platform.

To ensure that this outfit creation function was tested appropriately, we followed general software development procedures. This involved conducting thorough testing to ensure that the editor functioned as intended, with no bugs or glitches that could affect the accuracy or reliability of

the data collected. By following this rigorous testing process, we were able to gather a high-quality dataset that accurately reflects the fashion choices made by users on the social networking service during the ten-year collection period.

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

We collected a complete dataset without sampling to create our dataset, except for data deleted by the user.

Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?

The employees in ZOZO Technologies, Inc. or VASILY, inc. (merged into ZOZO Technologies, Inc.) were involved in collecting data.

Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.

The dataset was collected in the period of 2010 2020. Each outfit includes a timestamp that describes when the outfit created.

Were any ethical review processes conducted (e.g., by an institutional review board)? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

No.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

Yes. Each instance is a combination of outfits created by an individual and preferred by that individual.

Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?

Collected directly through the website.

Were the individuals in question notified about the data collection? If so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.

Notified in the Terms of Service.

Did the individuals in question consent to the collection and use of their data? If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.

The use of the service was deemed as consent.

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses? If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate).

It is possible to contact the company that provided the service.

Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted? If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

No, there had been no potential impact analysis conducted.

Any other comments?

Preprocessing/cleaning/labeling

Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)? If so, please provide a description. If not, you may skip the remainder of the questions in this section.

We extracted the CNN features from images and treated them as input data in our image-based tasks. As a result, our dataset contains the features but does not include raw photos, making them anonymized. The CNN we used is an official pre-trained VGG16, and we adopted the outputs of the 'fc6' layer before applying ReLU as the feature. We exclude the outfits that contain less than four items. Other than that, we did not remove any instances in creating our dataset. However, we excluded some data in each independent task. In detail, please refer to each task description.

Was the "raw" data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the "raw" data.

No.

Is the software used to preprocess/clean/label the instances available? If so, please provide a link or other access point.

All software are provided on the SHIFT15M repository.

Any other comments?

Uses

Has the dataset been used for any tasks already? If so, please provide a description.

Benchmarks using this dataset and the specified evaluation protocol are listed in GitHub page. <https://github.com/st-tech/zozo-shift15m/tree/main/benchmarks>.

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

All benchmarks that use this dataset will be available at GitHub page. <https://github.com/st-tech/zozo-shift15m/tree/main/benchmarks>.

What (other) tasks could the dataset be used for?

Here, we list candidate tasks for which SHIFT15M can be applied as follows:

- set-to-set matching;
- regression (e.g., number of likes or sum of prices);

- classification (e.g., category ids or publish years).

Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

No.

Are there tasks for which the dataset should not be used? If so, please provide a description.

This dataset is distributed in a way that excluding raw images and anonymizing the users/brands. Therefore, it requires the dataset users not to reconstruct raw images from the image features or restore the anonymized parts in a future task.

Any other comments?

Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

Yes. The dataset will be distributed to third parties based on the licence.

How will the dataset will be distributed (e.g., tarball on website, API, GitHub) Does the dataset have a digital object identifier (DOI)?

The dataset will be distributed via a website or the links indicated in our Github repository. We will add DOI for the SHIFT15M dataset.

When will the dataset be distributed?

The dataset will be first released in August 2021.

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

The SHIFT15M dataset will be made available for distribution under the Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) license. This means that users are free to share and adapt the dataset, as long as they provide attribution and do not use it for commercial purposes. For more information about the license, please refer to the following link: <https://creativecommons.org/licenses/by-nc/4.0/>. This license ensures that the dataset can be used by the academic community for research purposes, and that any derivative works or publications based on the dataset will be properly attributed.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

There are no fees or restrictions.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions,

and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

Unknown.

Any other comments?

Maintenance

Who will be supporting/hosting/maintaining the dataset?

ZOZO Research is supporting/maintaining the dataset.

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

All changes to the dataset will be announced through the GitHub Releases.

Is there an erratum? If so, please provide a link or other access point.

To ensure the accuracy and transparency of the SHIFT15M dataset, any changes made to the dataset will be immediately announced through the GitHub Releases page. This will include updates to the dataset's documentation, modifications to the dataset's format or metadata, or any other changes that may impact the dataset's use. Additionally, any errors or issues found in the dataset will be listed in the "Errata" section of the SHIFT15M repository. This will allow users to stay informed of any updates or issues related to the dataset and ensure that they are working with the most accurate and up-to-date version of the data.

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

All changes to the dataset will be announced through the GitHub Releases.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

No.

Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

They will continue to be supported with all information on SHIFT15M repository. We also provide the contribution guides for software that supports the dataset.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

Others may do so and should contact the original authors about incorporating fixes/extensions.

Any other comments?