

# Global and Local Prompts Cooperation via Optimal Transport for Federated Learning

Hongxia Li<sup>1</sup> Wei Huang<sup>2</sup> Jingya Wang<sup>1</sup> Ye Shi<sup>1,\*</sup>

<sup>1</sup>ShanghaiTech University, Shanghai, China

<sup>2</sup>RIKEN Center for Advanced Intelligence Project, Japan

{lihx2, wangjingya, shiye}@shanghaitech.edu.cn, wei.huang.vr@riken.jp

<https://github.com/HongxiaLee/FedOTP>

## Abstract

*Prompt learning in pretrained visual-language models has shown remarkable flexibility across various downstream tasks. Leveraging its inherent lightweight nature, recent research attempted to integrate the powerful pretrained models into federated learning frameworks to simultaneously reduce communication costs and promote local training on insufficient data. Despite these efforts, current federated prompt learning methods lack specialized designs to systematically address severe data heterogeneities, e.g., data distribution with both label and feature shifts involved. To address this challenge, we present Federated Prompts Cooperation via Optimal Transport (FedOTP), which introduces efficient collaborative prompt learning strategies to capture diverse category traits on a per-client basis. Specifically, for each client, we learn a global prompt to extract consensus knowledge among clients, and a local prompt to capture client-specific category characteristics. Unbalanced Optimal Transport is then employed to align local visual features with these prompts, striking a balance between global consensus and local personalization. By relaxing one of the equality constraints, FedOTP enables prompts to focus solely on the core regions of image patches. Extensive experiments on datasets with various types of heterogeneities have demonstrated that our FedOTP outperforms the state-of-the-art methods.*

## 1. Introduction

Federated learning [52] is a distributed machine learning framework that enables decentralized collaboration among participants without sharing their training data. However, current federated learning methods involve high training and communication costs due to the need to update and share model parameters with the server. This constraint has

typically restricted these methods to modest backbone architectures, hindering their feature capacity and resulting in performance limitations and training instability [74].

Recently, vision-language pre-trained models like Contrastive Language-Image Pretraining (CLIP) [60] have shown potential in learning robust and versatile representations suitable for various image distributions, aligning with the objectives of federated learning. However, the substantial communication overhead between the server and clients renders training CLIP in federated learning frameworks. Besides, overfitting concerns may arise when large-scale models are trained with limited client data. Prompt learning [46, 78] provides a flexible way to adapt pre-trained models to downstream tasks by training only additional parameters. This enables prompts to capture task-specific information while guiding the fixed model’s performance. Leveraging its lightweight nature, prior research [27, 76] has explored the integration of prompt learning into federated learning to overcome the problems outlined above.

In real-world scenarios, client data often exhibits variations in domain discrepancies (feature shift) [44] or imbalanced class distributions (label shift) [40]. Simply applying the FedAvg [52] method on prompts [27] across all clients tends to deviate from their local distribution, leading to unsatisfactory performance. Hence, it’s crucial to develop specialized personalized federated prompt learning approaches to effectively address data heterogeneity. pFedPrompt [26] introduced personalization into federated prompt learning by maintaining personalized attention modules to generate spatial visual features locally while learning user consensus through shared text prompts. However, in the presence of a significant label shift or notable feature shift, merely learning a shared prompt in the language modality is inadequate.

To resolve these limitations, we propose simultaneously learning both a shared global prompt and a personalized local prompt for each client in the local training phase. After local training, the local prompt remains on the client side,

\*Corresponding author.

while the global prompt is transmitted to server to aggregate with prompts from other clients. In this manner, the client owns the capacity to acquire consensus knowledge among clients from the global prompt, while also being able to discern client-specific user traits through the local prompt.

To further achieve a balance between global consensus and local personalization, we introduce Federated Prompts Cooperation via Optimal Transport (FedOTP). FedOTP utilizes Optimal Transport (OT) [33] to align local visual features with both global and local textual features through an adaptive transport plan, promoting fine-grained matching across vision and language modalities and strengthening collaboration between the global and local prompts. The adaptive OT transport plan can provide resilience to visual misalignment and effective adaptation to feature shifts. It's worth noting that the standard OT imposes two hard equality constraints on the transport plan, leading to each image patch being assigned to prompts. This may potentially cause prompts to capture some class-irrelevant information from the image and consequently influence the final results. To mitigate this, we consider employing unbalanced OT by relaxing one of the equality constraints, allowing prompts to focus solely on the most relevant image patches rather than the entire content of the image. For an efficient solution, we apply a fast implementation of Dykstra's algorithm [16] in our FedOTP, enabling swift convergence and focusing on the core area of the image during iterations.

Our main contributions are summarized as follows:

- We are the first to explore the mechanism of prompts' cooperation in federated learning where severe data heterogeneity is present. More precisely, we train both a global prompt for consensus across clients and a local prompt to capture client-specific category traits simultaneously.
- We propose FedOTP, a federated learning framework utilizing unbalanced OT to enhance the cooperation between global and local prompts. Through unbalanced OT, we align local visual features with textual prompts while enable prompts to focus solely on the critical image patches.
- We conducted extensive experiments on widely adopted datasets in various data heterogeneity with feature shifts and label shifts, and significant result improvement verifies the superiority of our FedOTP. In addition, we demonstrated the ability of FedOTP to balance consensus and local personalization through visualizations.

## 2. Related Work

### 2.1. Personalized Federated Learning

Personalized federated learning (PFL) is a highly regarded research field because of its potential to address statistical and systemic heterogeneity across clients. Various approaches have been proposed in prior research to achieve PFL. The most common method involves the inclusion

of regularization terms in the loss function [42, 43, 66], and fine-tuning the global model on clients' local datasets [17, 35, 50, 72]. Additionally, some works focus on explicitly seeking a trade-off between the global model and the local models [8, 28, 45, 51]. To enhance adaptability to diverse data distributions, certain techniques have delved into clustering methods for client grouping [32, 61, 70]. Leveraging the relationships and data distribution among clients, methods like FedPAC [73], and FedDisco [75] introduce novel weighted aggregation techniques to promote intensive collaboration among similar clients. Furthermore, some researchers have explored the decomposition of model parameters into base layers and personalized layers. For instance, FedPer [1], FedRep [12], and FedBABU [55] learn personalized classifier heads locally while sharing the base layers, and FedTP [39] learns personalized self-attention layers for each client. The importance of PFL was pointed out theoretically by [31].

The methods mentioned above primarily target label shift data heterogeneity. However, they may not perform well when substantial domain differences exist among clients. In dealing with these feature shifts, FedBN [44] employs local batch normalization to mitigate the feature shift before model averaging, while PartialFed [65] extends this strategy by selecting personalized parameters according to distinct feature traits of different clients. Besides, FedPCL [67] enhances each client's ability to utilize pre-trained models by extracting client-specific and class-relevant information. Our FedOTP explores the cooperation between global and local prompts to effectively address both label shift and feature shift data heterogeneity.

### 2.2. Prompt-based Federated Learning

Prompt learning, originating from NLP models, has expanded to Vision Language Models. Initial methods like CLIP [60] involved manually crafted templates, while recent approaches concentrate on learning prompts in a continuous embedding space. CoOp [78] fine-tunes CLIP with continuous prompt vectors. Based on this, plenty of studies [7, 30, 34, 47, 49, 77] have been introduced to enhance the effectiveness of prompt learning. To accelerate the global aggregation and handle situations with insufficient user data, FedPrompt [76] and PromptFL [27] have introduced prompt learning into Federated Learning. Based on these two works, several methods have made substantial progress in various domains. For instance, FedPR [22] focuses on learning federated visual prompts within the null space of the global prompt for MRI reconstruction. Based on CLIP, FedAPT [64] introduces a federated adaptive prompt tuning algorithm for cross-domain federated image classification, and FedCLIP [48] utilizes an attention-based adapter to optimize the utilization of pre-trained model information. To tackle statistical heterogeneity among clients,

pFedprompt [26] maintains a non-parametric personalized attention module for each client to generate locally personalized spatial visual features, and pFedPG [74] designs a client-specific prompt generator at the server to create personalized prompts. While these works show the potential of prompt learning in Federated Learning, there remains a deficiency in technical enhancements tailored to PFL scenarios. Compared with these methods, our FedOTP employs OT to balance the global consensus and local personalization from the collaboration of global and local prompts.

### 2.3. Optimal Transport

Initially developed as a solution to optimize the cost of moving multiple items concurrently, Optimal Transport (OT) [33] has gained significant attention in the machine learning and computer vision community. To accelerate the convergence and efficiently deal with large-scale problems, [13] introduced Sinkhorn’s algorithm [63] for computing an approximate transport coupling with entropic regularization. Unbalanced OT relaxes the equality constraint of classical OT, allowing for partial displacement in the transport plan. It can be efficiently solved using the generalized Sinkhorn’s algorithm [10] by incorporating soft penalties based on Kullback-Leibler divergence [23] or parallel algorithm [37]. Given its remarkable ability in distribution matching, OT has been applied in various theoretical and practical tasks, including domain adaptation [5, 14, 19, 38], learning with noisy labels [6, 21], causal discovery [41, 68, 69], federated learning [9, 18], outlier detection [59] and so on. In prompt learning field, PLOT [7] proposes to learn multiple prompt sets for diverse contextual representations and use OT to align the features of vision and language modalities. Different from PLOT, we employ unbalanced OT to enhance the cooperation between global and local prompts by relaxing one of the equality constraints, which allows prompts to concentrate exclusively on the most relevant image patches.

## 3. Preliminaries

### 3.1. Prompt Learning

To adapt pre-trained models like CLIP [60] to downstream tasks, prompt learning methods [77, 78] provide an efficient way by training a few parameters in the prompt. Within the CLIP model, the textual prompts are manually crafted using class labels  $y \in \{1, 2, \dots, K\}$  (e.g., “a photo of a  $\langle \text{classname} \rangle$ ”) representing  $K$  classes. These textual prompts are then tokenized and projected into word embeddings  $W = \{w_1, w_2, \dots, w_L\} \in \mathbb{R}^{L \times d_l}$  where  $L$  is the number of word embeddings and  $d_l$  denotes its dimension. To learn the context prompts, we interpose  $s (\leq L)$  learnable vectors  $\{p_i \in \mathbb{R}^{d_l}\}_{i=1}^s$  in the language branch. Consequently, the textual prompt can be formulated as  $P_k = \{w_1, p_1, \dots, p_s, w_{s+2}, \dots, w_L\} \in \mathbb{R}^{L \times d_l}$ , where we use

$\{p_1, \dots, p_s\}$  in place of  $\{w_2, \dots, w_{s+1}\}$  to be consistent with previous works. Denote the fixed text encoder as  $h(\cdot)$  and image encoder as  $g(\cdot)$ , and the prediction probabilities for each category are computed with the input prompt  $P_k$  of class  $k$  and image  $x$  through matching scores:

$$q(y = k | \mathbf{x}) = \frac{\exp(\text{sim}(g(x), h(P_k))/\tau)}{\sum_{c=1}^K \exp(\text{sim}(g(x), h(P_c))/\tau)}, \quad (1)$$

where  $\text{sim}(\cdot, \cdot)$  denotes a metric function such as cosine similarity, and  $\tau$  represents the temperature of Softmax. Next, we optimize the learnable parameters  $\{p_i\}_{i=1}^s$  by minimizing the cross-entropy loss

$$\ell_{CE} = -\frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} \sum_{k=1}^K y_{\mathbf{x},k} q(y = k | \mathbf{x}), \quad (2)$$

where  $y_{\mathbf{x}}$  is a one-hot label vector.

### 3.2. Optimal Transport

Optimal Transport is a constrained optimization problem aiming to efficiently transfer probability mass between two distributions. Here we briefly recall its formulation of the discrete situation. Given two probability simplex vectors  $\alpha$  and  $\beta$  and a cost matrix  $C \in \mathbb{R}^{|\alpha| \times |\beta|}$ , OT aims to find the optimal transport plan  $T$  by minimizing the following objective:

$$d_C(\alpha, \beta) = \min_{T \in U(\alpha, \beta)} \langle C, T \rangle, \quad (3)$$

$$U(\alpha, \beta) = \left\{ T \in \mathbb{R}_+^{|\alpha| \times |\beta|} \mid T \mathbf{1}_{|\beta|} = \alpha, T^\top \mathbf{1}_{|\alpha|} = \beta \right\},$$

where  $\langle \cdot, \cdot \rangle$  is Frobenius dot-product,  $U(\alpha, \beta)$  denotes the solution space of  $T$ , and  $\mathbf{1}_d$  is a  $d$ -dimensional vector of ones. Directly optimizing the OT problem would be time-consuming. Sinkhorn algorithm [13] introduces an entropic regularization term for fast optimization. The regularized OT formulation can be expressed as:  $\min_{T \in U(\alpha, \beta)} \langle C, T \rangle + \lambda \langle T, \log T \rangle$ , where  $\lambda \geq 0$  is a hyper-parameter. In light of this, the optimal transport plan  $T^*$  has been shown to be unique with the form  $T^* = \text{diag}(u^{(\tilde{t})}) \exp(-C/\lambda) \text{diag}(v^{(\tilde{t})})$ , where  $\tilde{t}$  represents the iteration and in each iteration  $u^{(\tilde{t})} = u / (\exp(-C/\lambda) v^{(\tilde{t}-1)})$  and  $v^{(\tilde{t})} = v / (\exp(-C/\lambda)^\top u^{(\tilde{t})})$ .

## 4. Methodology

In this section, we present the design of our FedOTP framework, illustrated in Figure 1. To achieve a balance between global consensus and local personalization, FedOTP utilizes unbalanced Optimal Transport to strengthen the collaboration between global and local prompts, effectively addressing both label shift and feature shift data heterogeneity.

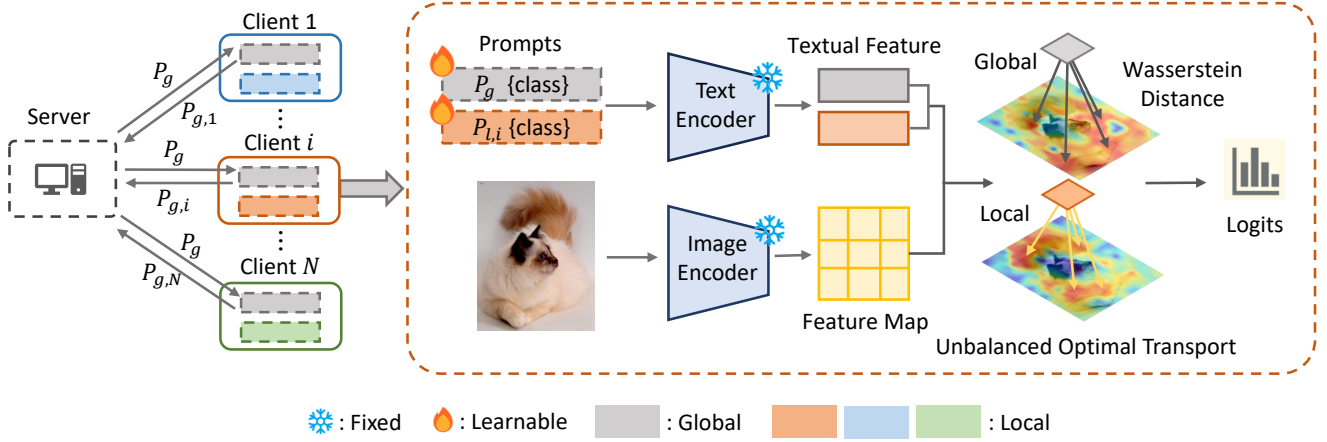


Figure 1. Overview of our FedOTP. On the left, clients transmit global prompts to the server for aggregation while retaining local prompts locally. The right shows the workflow of Global-Local prompt cooperation mechanism, which employs unbalanced Optimal Transport to align visual feature maps with each prompt.

#### 4.1. Federated Learning with Global and Local Prompts

Consider a federated learning scenario involving  $N$  clients and a central server, and each client  $i$  holds a local dataset  $D_i = \{(x_i^j, y_i^j)\}_{j=1}^{m_i}$  ( $i = 1, \dots, N$ ) containing  $m_i$  samples. Let  $D = \{D_1, D_2, \dots, D_N\}$  represent the total datasets where each dataset is derived from a distinct data distribution  $\mathcal{D}_i$ , and let  $C_t$  represent the set of selected clients participating at communication round  $t$ .

Aiming to reduce communication costs and address data heterogeneity, each client is equipped with a pre-trained CLIP model and a prompt learner in our federated learning setup. Through the prompt learner, every client learns both a shared global prompt and a personalized local prompt, allowing clients to extract more individualized insights while maintaining a degree of consensus among them. Specifically, for each client  $i$ , the prompt  $P_i$  comprises a global prompt  $P_g$  and a personalized prompt  $P_{l,i}$ , denoted as  $P_i = [P_g, P_{l,i}]$ . During each communication round  $t$ , client  $i$  initializes the prompt with  $P_i^{t,0} = [P_g^{t-1}, P_{l,i}^{t-1}]$ . Then the global and local prompts are jointly updated through gradient descent  $P_i^{t,r} = P_i^{t,r-1} - \eta \nabla \mathcal{L}_{\mathcal{D}_i}(P_i^{t,r-1})$  for  $R$  iterations locally. After local training, only the updated global prompt  $P_{g,i}^{t,R}$  is transmitted to the server for aggregation to learn global consensus among clients, while the personalized prompt is retained locally to capture client-specific category characteristics. The process of aggregation can be expressed as:

$$P_g^t = \sum_{i \in C_t} \frac{m_i}{\sum_{j \in C_t} m_j} P_{g,i}^{t,R}. \quad (4)$$

Then the objective function of our FedOTP can be formu-

lated as:

$$\min_{P_g, \{P_{l,i}\}_{i=1}^N} \sum_{i=1}^N \frac{m_i}{\sum_{j \in C_t} m_j} \mathcal{L}_{\mathcal{D}_i}(P_g, P_{l,i}), \quad (5)$$

with  $\mathcal{L}_{\mathcal{D}_i}(P_g, P_{l,i}) = \mathbb{E}_{(x_i^j, y_i^j) \in \mathcal{D}_i} \ell(f(P_g, P_{l,i}; x_i^j), y_i^j)$ , where  $f(P_g, P_{l,i}; \cdot)$  represents the personalized model for client  $i$ , and  $\ell(\cdot, \cdot)$  denotes the cross-entropy loss function as introduced in Eq. (2).

#### 4.2. Global-Local Prompt Cooperation by Unbalanced Optimal Transport

In this subsection, we introduce the details of the prompt learning process for each client, which leverages unbalanced OT to integrate insights learned from both global and local prompts. To be specific, as shown in Figure 1, we initialize prompts  $P_g$  and  $P_{l,i}$  as  $\{w_1, p_1, \dots, p_s, \dots, w_L\}$  where  $w_i$  represents the word embedding and  $p_i$  signifies learnable vectors. With the text encoder  $h(\cdot)$ , we obtain a global textual feature  $H_{k,g} = h(P_{g,k}) \in \mathbb{R}^{d_f}$  and a local textual feature  $H_{k,l} = h(P_{l,i,k}) \in \mathbb{R}^{d_f}$  for each class  $k$ , and the combination of these two features is represented as  $H_k = [H_{k,g}, H_{k,l}]$  for convenience. Uploading an image  $x$  to the image encoder  $g(\cdot)$ , we derive a set of visual features  $G = g(x) \in \mathbb{R}^{(V+1) \times d_f}$ , which consists of a class token  $G_c \in \mathbb{R}^{d_f}$  and a feature map  $G_m \in \mathbb{R}^{V \times d_f}$ .

We consider learning an optimal transport plan  $T$  that aligns both global and local textual features  $H_k$  with visual feature map  $G_m$ . By representing features as samples from discrete distributions, the cost matrix can be represented by the cosine distance between  $H_k$  and  $G_m$  as  $C = 1 - G_m^T H_k \in \mathbb{R}^{V \times 2}$ , then the optimization objective

of unbalanced optimal transport is formulated as:

$$d_{C,k}(\alpha, \beta) = \min_{T \in U(\alpha, \beta)} \langle C, T \rangle, \quad (6)$$

$$U(\alpha, \beta) = \{T \in \mathbb{R}_+^{V \times 2} \mid T \mathbb{1}_2 \leq \alpha, T^\top \mathbb{1}_V = \beta\},$$

where  $\alpha \in \mathbb{R}^V$  and  $\beta \in \mathbb{R}^2$  are essentially marginal probability vectors which satisfy  $\|\alpha\|_1 \geq \|\beta\|_1 = \gamma$  ( $\gamma \in [0, 1]$ ). The difference between Eq. (6) and formulation in PLOT [7] lies in their use of classical OT with two hard equality constraints as Eq. (3). This forces prompts to map to each image patch, potentially causing them to capture some class-irrelevant information from the image and thereby influencing the final results. In contrast, our FedOTP relaxes one of the equality constraints, allowing prompts to concentrate solely on the most relevant image patches rather than the entire content of the image. Additionally, by controlling  $\gamma$ , FedOTP owns the ability to regulate the mapping size of prompts on the feature map.

For fast optimization, we add an entropic regularization term to Eq. (6), and the objective function is formulated as follows:

$$d_{C,k}(\alpha, \beta) = \min_{T \in U(\alpha, \beta)} \langle C, T \rangle + \lambda \langle T, \log T \rangle. \quad (7)$$

In line with [2], we can further reformulate Eq. (7) as a Kullback-Leibler (KL) projection, and the solution space  $U(\alpha, \beta)$  is then defined as the intersection of two convex but not affine sets:

$$d_{C,k}(\alpha, \beta) = \min_{T \in U(\alpha, \beta)} \lambda \text{KL}(T \mid e^{-C/\lambda}), \quad (8)$$

$$\mathcal{C}_1 \triangleq \{T \in \mathbb{R}_+^{V \times 2} \mid T \mathbb{1}_2 \leq \alpha\},$$

$$\mathcal{C}_2 \triangleq \{T \in \mathbb{R}_+^{V \times 2} \mid T^\top \mathbb{1}_V = \beta\}.$$

To solve Eq. (8), we employ a rapid implementation of Dijkstra's algorithm [16] as introduced in [6], which efficient scales iterative KL projection between  $\mathcal{C}_1$  and  $\mathcal{C}_2$  by leveraging matrix-vector multiplications exclusively. Initializing  $Q = \exp(-C/\lambda)$  and  $v^{(0)} = \mathbb{1}_2$ , a fast optimization solution is achieved within a few iterations as:

$$T^* = \text{diag}(u^{(\tilde{t})}) Q \text{diag}(v^{(\tilde{t})}), \quad (9)$$

where  $\tilde{t}$  is the iteration, and in each iteration  $u^{(\tilde{t})} = \min(\mathbb{1}_V / Q_\alpha v^{(\tilde{t}-1)}, \mathbb{1}_V)$  and  $v^{(\tilde{t})} = \mathbb{1}_2 / Q_\beta^\top u^{(\tilde{t})}$  with  $Q_\alpha = Q / \text{diag}(\alpha) \mathbb{1}_{V \times 2}$  and  $Q_\beta^\top = Q^\top / \text{diag}(\beta) \mathbb{1}_{V \times 2}$ . The details of this algorithm are shown in Appendix Section A.1.

By Eq. (9), we obtain the optimal transport plan  $T^*$  and the final Wasserstein distance  $d_{C,k}$ , then the matching scores in Eq. (1) is replaced by the following prediction probability:

$$q(y = k \mid \mathbf{x}) = \frac{\exp((1 - d_{C,k})/\tau)}{\sum_{c=1}^K \exp((1 - d_{C,c})/\tau)}. \quad (10)$$

After obtaining  $q(y = k \mid \mathbf{x})$ , we fix the transport plan  $T^*$  and optimize learnable vectors  $\{p_a\}_{a=1}^s$  in both global and local prompts simultaneously for client  $i$  through cross entropy as described in Eq. (2). Then the global prompt  $P_{g,i}$  is sent to the server for aggregation utilizing Eq. (4) with the local prompt retained locally. During local training via OT, the final prediction probability of FedOTP is a synthesis of information derived from both the global and the local prompts. This avoids a straightforward addition of the outcomes from the two prompts, fostering a comprehensive and collaborative learning process. Due to page limitation, the algorithm box is deferred to the Appendix Section A.2.

### 4.3. Generalization Bound

We analyze the generalization bound of our FedOTP in this section. Before starting the analysis, we first introduce some assumptions as follows.

**Assumption 1 (Lipschitz Conditions)** Let  $\mathcal{D}_1, \dots, \mathcal{D}_N$  denote the real data distribution of each client and  $\mathcal{L}_{\mathcal{D}_i}(P_g, P_{l,i}) = \mathbb{E}_{(x_i^j, y_i^j) \in \mathcal{D}_i} \ell(f(P_g, P_{l,i}; x_i^j), y_i^j)$  be the expected loss. We assume the following Lipschitz conditions hold:

$$|\ell(f((P; x), y) - \ell(f((P'; x), y))| \leq L \|f((P; x), y) - f((P'; x), y)\|, \quad (11a)$$

$$\|f(P_g, P_{l,i}) - f(P'_g, P_{l,i})\| \leq L_g \|P_g - P'_g\|, \quad (11b)$$

$$\|f(P_g, P_{l,i}) - f(P_g, P'_{l,i})\| \leq L_{l,i} \|P_{l,i} - P'_{l,i}\|. \quad (11c)$$

**Assumption 2** Since the convergence of global prompt has been proved in [27], we assume  $\|\hat{P}_g - P_g^*\|_2 \leq A_g$  for convenience. And we assume local prompts  $P_{l,i}$  are bounded in a ball of radius  $A_{l,i}$ , which can be denoted as  $\|\hat{P}_{l,i} - P_{l,i}^*\|_2 \leq A_{l,i}$ .

Leveraging above assumptions, we can derive the following theorem:

#### Theorem 1 (Generalization Bound of FedOTP)

Suppose  $\hat{\mathcal{D}}_1, \dots, \hat{\mathcal{D}}_N$  denote empirical data distribution of  $N$  clients with learned parameters  $\hat{P}_g$  and  $\hat{P}_{l,i}$ , and  $P_g^*$  and  $P_{l,i}^*$  are optimal parameters for the real distribution  $\mathcal{D}_1, \dots, \mathcal{D}_N$ . Let  $\mathcal{H}$  represent the personalized hypothesis and  $d$  denote the VC-dimension of  $\mathcal{H}$ . Suppose all the clients participate at every communication round and Assumptions 1 and 2 hold, with probability at least  $1 - \delta$ , we have

$$\left| \sum_{i=1}^N \frac{m_i}{M} \left( \mathcal{L}_{\hat{\mathcal{D}}_i}(\hat{P}_g, \hat{P}_{l,i}) - \mathcal{L}_{\mathcal{D}_i}(P_g^*, P_{l,i}^*) \right) \right| \leq \sqrt{\frac{M}{2} \log \frac{N}{\delta}} + \sqrt{\frac{dN}{M} \log \frac{eM}{d}} + L(L_g A_g + L_l A_l), \quad (12)$$

Table 1. The results of our FedOTP and the benchmark methods on the Pathological Non-IID setting with non-overlapping over 10 clients.

Methods	Food101	DTD	Caltech101	Flowers102	OxfordPets
<b>Local Training</b>					
Zero-Shot CLIP [60]	75.27±0.05	40.21±0.12	85.14±0.24	62.17±0.12	84.47±0.10
CoOp [78]	82.54±2.42	82.69±0.63	90.41±0.44	88.23±0.76	94.52±1.30
<b>Prompt-based Federated Learning</b>					
PromptFL [27]	74.81±0.64	50.46±0.54	87.90±0.54	73.68±1.58	88.17±1.18
PromptFL+FT [24]	77.16±1.56	53.74±1.36	89.70±0.25	72.31±0.91	91.23±0.50
PromptFL+FedProx [42]	73.96±0.75	50.89±0.71	87.80±1.10	74.14±0.65	87.25±1.48
PromptFL+FedPer [1]	71.29±1.87	50.23±0.82	86.72±1.45	72.11±1.35	89.50±1.62
PromptFL+FedAMP [32]	74.48±1.71	47.16±0.92	87.31±1.60	69.10±0.13	80.21±0.44
pFedPrompt [26]	92.26±1.34	77.14±0.09	96.54±1.31	86.46±0.15	91.84±0.41
FedOTP (Ours)	<b>92.73±0.15</b>	<b>87.67±0.70</b>	<b>97.02±0.36</b>	<b>96.23±0.44</b>	<b>98.82±0.11</b>

where  $M = \sum_{i=1}^N m_i$ , and we denote  $L_l = \sqrt{\sum_{i=1}^N L_{l,i}^2}$  and  $A_l = \sqrt{\sum_{i=1}^N A_{l,i}^2}$  for simplicity. Theorem 1 indicates that the performance of FedOTP trained on the empirical distribution relates to the model complexity and Lipschitz assumptions. More details and proof of Theorem 1 are provided in the Appendix Section E.

## 5. Experiments

In this section, we conducted comprehensive experiments to numerically evaluate our FedOTP in the scenarios of heterogeneous data distribution.

### 5.1. Experimental Setup

**Datasets and Data Heterogeneity.** We evaluated the performance of FedOTP on nine public benchmark datasets with different types of heterogeneity, including label shift and feature shift. To investigate label shift, we selected two types of datasets. Following previous research [26, 27], we utilized five visual classification datasets to simulate datasets with limited samples: Food101 [3], DTD [11], Caltech101 [20], Flowers102 [54], and OxfordPets [56]. Referring to these datasets as the CLIP dataset for convenience, we utilized a Pathological setting by randomly allocating a distinct number of non-overlapping classes to each client. We also employed two image benchmark datasets: CIFAR-10, and CIFAR-100 [36]. We considered the Dirichlet distribution as introduced in [4, 39, 62] where datasets are partitioned randomly among clients using a symmetric Dirichlet distribution with  $\alpha = 0.3$ . For feature shift, we utilized two datasets with multiple domains: DomainNet [58] with 6 domains, and Office-Caltech10 [25] with 4 domains. In line with prior studies [44, 67], each client participating in the federated learning system is assigned data from one of these distinct domains. The details of these dataset setup can be found in the Appendix Section B.1.

**Baselines.** We compared FedOTP with three kinds of base-

lines: (1) Local training methods: (i) Zero-shot CLIP [60] with hand-crafted text prompt templates; (ii) CoOp [78] with learnable prompt vectors trained on each client locally. (2) Existing prompt-based federated learning methods: (i) PromptFL [27] learning a unified prompt across clients; (ii) pFedPrompt [26] learning a shared prompt with personalized visual attention modules. (3) Four adapted methods derived from traditional PFL techniques, including PromptFL + FT [24], PromptFL+FedProx [42], PromptFL+FedPer [1] and PromptFL+FedAMP [32], as introduced in [26].

**Implementation Details.** To simulate federated learning in various scenarios, we consider the following two settings: (1)  $n = 10$  clients with a full 100% partition, (2)  $n = 100$  clients with a 10% partition. We employ SGD optimizer with a learning rate  $lr = 0.001$  and local epoch  $R = 5$  for CLIP datasets while  $R = 1$  for other cases. The communication round is set to  $T = 10$  for CLIP datasets with 10 clients and  $T = 150$  for CIFAR-10/CIFAR-100 datasets with 100 clients. We present the results using two representative backbones, ResNet50 [29] and ViT\_B16 [15], defaulting to ViT\_B16 if not explicitly specified. More implementation details can be found in the Appendix Section B.2.

### 5.2. Performance Evaluation

**Evaluation Protocol.** We evaluated the models on each client’s private test data whose distribution is consistent with its training set. The reported results are the average test accuracy across all clients from three different seeds.

**Model Evaluation on Label Shifts.** We first measured the performance of FedOTP against baselines on datasets with label shifts. The experimental results on CLIP datasets and CIFAR-10/CIFAR-100 datasets are summarized in Table 1 and Table 3. For easy comparison, Table 1 reports results utilizing ResNet50 as the backbone, maintaining consistency with [26]. As shown in Table 1, our FedOTP outperforms state-of-the-art algorithms by a large

Table 2. Experimental results on DomainNet dataset with feature & label shifts.

Datasets	DomainNet						
	Clipart	Infograph	Painting	Quickdraw	Real	Sketch	Avg.
<b>Local Training</b>							
Zero-Shot CLIP [60]	8.72±1.73	12.48±3.78	8.53±4.32	9.31±0.69	9.13±2.55	11.96±2.80	10.02±2.65
CoOp [78]	44.40±14.89	45.68±16.53	<b>47.21±18.20</b>	<b>41.13±20.62</b>	48.02±24.49	39.47±5.68	44.32±16.74
<b>Prompt-based Federated Learning</b>							
PromptFL [27]	9.31±6.53	12.58±9.91	8.23± 8.47	14.79±12.07	9.37±10.82	7.48±11.32	10.29±10.35
PromptFL+FedProx [42]	9.84±6.60	11.16±11.17	10.64±6.79	13.40±16.09	9.39±7.69	6.78±11.76	10.20±10.99
FedOTP (Ours)	<b>46.14±6.53</b>	<b>60.14±18.23</b>	45.2±16.86	38.66±7.60	<b>49.30±17.80</b>	<b>49.02±24.22</b>	<b>48.08±15.21</b>




Figure 2. Heatmaps of similarity between text features and image feature maps for different methods on 4 categories in OxfordPets dataset. “FedOTP-G” denotes the results from the global prompt and “FedOTP-L” refers to the local prompt.

Table 3. The results of our FedOTP and the benchmark methods on Dirichlet settings in CIFAR-10 and CIFAR-100 over 100 clients.

Methods	CIFAR-10	CIFAR-100
<b>Local Training</b>		
Zero-Shot CLIP [60]	87.71±0.68	64.92±0.53
CoOp [78]	93.11±0.39	74.83±0.45
<b>Prompt-based Federated Learning</b>		
PromptFL [27]	92.30±0.87	73.67±0.56
PromptFL+FedProx [42]	91.83±0.47	71.11±0.91
FedOTP (Ours)	<b>96.05±0.12</b>	<b>78.03±0.08</b>

margin across all datasets, which confirms the effectiveness of our Global-Local prompt cooperation mechanism to handle label shift scenarios. Remarkably, while both PromptFL+FedPer (which splits the learnable prompt vector into “base+personalized” vectors) and pFedPrompt (utilizing a shared prompt with a personalized attention module in the vision modal) experience significant declines when datasets are altered, FedOTP exhibits slight fluctuations. This verifies the robustness of our method across diverse scenarios. Table 3 shows the results of our FedOTP and benchmark methods on CIFAR-10/CIFAR-100 datasets under Dirichlet setting over 100 clients with 10% partition. Even in this scenario with Dirichlet settings and a large number of clients, FedOTP consistently outperforms the baseline methods, further highlighting the superiority of our approach.

**Model Evaluation on Feature & Label Shifts.** In this

set of experiments, we explored scenarios involving both feature shifts and label shifts by partitioning data within a domain into five clients based on the Dirichlet distribution with  $\alpha = 0.1$ . We analyzed the mean and variance of clients in the same domain, and the outcomes for the DomainNet dataset are summarized in Table 2. In the presence of two types of data heterogeneity, our method performs favorably against baselines. We observe that, with significant data heterogeneity across clients, traditional federated learning methods experience a pronounced performance decline compared to local training. In contrast, our FedOTP exhibits superior performance, achieving a 3.7% increase in average accuracy on each domain. Additional experimental results on feature shifts and in the Office-Caltech10 dataset are available in the Appendix Section C.1 and C.2.

**Visualization.** We first investigated the interplay between global and local prompts by representing the similarity between text features and image feature maps as a heatmap on OxfordPets [56] dataset. To be specific, we compared the heatmaps of our FedOTP with PromptFL and Local Training using CoOp, and the original images and corresponding heatmaps are illustrated in Figure 2. We observed that global prompts of FedOTP might concentrate more on common features, like limbs and facial characteristics, while local prompts tended to capture client-specific details such as the special tail of “Birman”, unique patterns of “Egyptian Mau” and “Basset Hound”, and the distinct dent on the head of “Terrier”. This demonstrates the effectiveness of FedOTP in balancing global consensus and local personal-




Figure 3. Performance with the different number of shots.

Table 4. Quantitative comparisons on the Pathological Non-IID setting across different numbers of shots over 10 clients.

Datasets	Food101		DTD		Caltech101		Flowers102		OxfordPets	
	2	8	2	8	2	8	2	8	2	8
FedOTP (Similarity Averaging)	83.38±0.54	87.59±1.05	81.01±0.23	88.17±0.73	92.68±0.44	96.73±0.29	91.73±0.68	97.09±0.18	96.23±0.25	98.34±0.15
FedOTP (Classical OT)	88.07±0.63	89.77±0.62	81.42±0.99	88.43±0.45	93.17±0.68	96.80±0.23	92.84±1.34	97.07±0.25	96.55±0.26	98.51±0.27
FedOTP (Unbalanced OT)	<b>89.12±0.28</b>	<b>92.94±0.18</b>	<b>85.50±0.35</b>	<b>90.25±0.74</b>	<b>95.05±0.49</b>	<b>97.34±0.18</b>	<b>93.96±0.48</b>	<b>98.23±0.32</b>	<b>97.73±0.57</b>	<b>99.02±0.38</b>

ization. More visualization results on transport plans of our FedOTP will be given in the Appendix Section D.1.

### 5.3. Ablation Study

**Impact of Number of Shots.** Following the few-shot evaluation setting adopted in [26, 27], we further investigated the impact of the number of shots in FedOTP. To analyze this, we varied the number of shots during the training process within the range of [1, 2, 4, 8, 16]. Results are summarized in Figure 3, where the horizontal axis denotes the number of shots and the vertical axis represents the average test accuracy. We observe that as the number of shots increases, the corresponding performance of each method gradually improves. However, our FedOTP consistently exhibits a dominant edge over methods with a shared global prompt in all scenarios.

**Effectiveness of the Unbalanced OT.** In this subsection, we explored the effectiveness of OT on two variants of FedOTP briefly described below: (1) FedOTP (Similarity Averaging): removing OT in FedOTP and matching global and local prompts with visual feature maps by averaging similarities of each visual-textual pair; (2) FedOTP (Classical OT): employing classical OT during the matching process. The results in Table 4 demonstrate the effectiveness of utilizing OT to align feature maps with global and local prompts compared to FedOTP (Similarity Averaging) in almost all cases, particularly on the Food101 dataset. This is because the absence of OT leads to the feature map’s distance from prompts reverting to the mean distance of each feature-prompt pair, highlighting the crucial role of OT in providing resilience to visual misalignment. In addition, the persistent superiority of unbalanced OT over classical OT across all scenarios serves as a compelling testament to the effectiveness of our approach.

## 6. Conclusion

In this paper, we proposed Federated Prompts Cooperation via Optimal Transport (FedOTP), a novel framework designed to facilitate efficient model personalization across heterogeneous clients. In our approach, each client is equipped with both a global prompt and a local prompt, and then unbalanced Optimal Transport is utilized to align local visual features with these prompts, fostering enhanced collaboration between global and local prompts. With fine-grained matching facilitated by OT, FedOTP effectively addresses data heterogeneity characterized by domain discrepancy and imbalanced class distributions. Our extensive experiments across diverse datasets consistently demonstrate the superior performance of FedOTP in tackling both label shifts and feature shifts, which verifies the effectiveness of our Global-Local prompt cooperation mechanism via OT. Through visualization results, we confirmed that global prompts learned by FedOTP concentrated on common features among all clients, while local prompts captured individual client-specific details. In future work, we aim to investigate the generalization capabilities of our method on novel clients unseen during the training process.

## Acknowledgement

This work was supported by NSFC (No.62303319), Shanghai Sailing Program (22YF1428800, 21YF1429400), Shanghai Local College Capacity Building Program (23010503100), Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University), and Shanghai Engineering Research Center of Intelligent Vision and Imaging.



## References

- [1] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019. [2](#), [6](#)
- [2] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015. [5](#)
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VI 13*, pages 446–461. Springer, 2014. [6](#), [2](#)
- [4] Yu-Tong Cao, Ye Shi, Baosheng Yu, Jingya Wang, and Dacheng Tao. Knowledge-aware federated active learning with non-iid data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22279–22289, 2023. [6](#)
- [5] Wanxing Chang, Ye Shi, Hoang Tuan, and Jingya Wang. Unified optimal transport framework for universal domain adaptation. *Advances in Neural Information Processing Systems*, 35:29512–29524, 2022. [3](#)
- [6] Wanxing Chang, Ye Shi, and Jingya Wang. CSOT: Curriculum and Structure-Aware Optimal Transport for Learning with Noisy Labels. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [3](#), [5](#), [1](#)
- [7] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. Prompt learning with optimal transport for vision-language models. *arXiv preprint arXiv:2210.01253*, 2022. [2](#), [3](#), [5](#)
- [8] Hong-You Chen and Wei-Lun Chao. On bridging generic and personalized federated learning for image classification. *arXiv preprint arXiv:2107.00778*, 2021. [2](#)
- [9] Yi-Han Chiang, Koudai Terai, Tsung-Wei Chiang, Hai Lin, Yusheng Ji, and John CS Lui. Optimal Transport based One-Shot Federated Learning for Artificial Intelligence of Things. *IEEE Internet of Things Journal*, 2023. [3](#)
- [10] Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018. [3](#)
- [11] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. [6](#), [2](#)
- [12] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International Conference on Machine Learning*, pages 2089–2099. PMLR, 2021. [2](#)
- [13] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013. [3](#), [1](#)
- [14] Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 447–463, 2018. [3](#)
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [6](#)
- [16] Richard L Dykstra. An algorithm for restricted least squares regression. *Journal of the American Statistical Association*, 78(384):837–842, 1983. [2](#), [5](#)
- [17] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems*, 33:3557–3568, 2020. [2](#)
- [18] Farzan Farnia, Amirhossein Reisizadeh, Ramtin Pedarsani, and Ali Jadbabaie. An optimal transport approach to personalized federated learning. *IEEE Journal on Selected Areas in Information Theory*, 3(2):162–171, 2022. [3](#)
- [19] Kilian Fatras, Thibault SÉjourné, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *International Conference on Machine Learning*, pages 3186–3197. PMLR, 2021. [3](#)
- [20] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004. [6](#), [2](#)
- [21] Chuanwen Feng, Yilong Ren, and Xike Xie. OT-Filter: An Optimal Transport Filter for Learning With Noisy Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16164–16174, 2023. [3](#)
- [22] Chun-Mei Feng, Bangjun Li, Xinxing Xu, Yong Liu, Huazhu Fu, and Wangmeng Zuo. Learning Federated Visual Prompt in Null Space for MRI Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8064–8073, 2023. [2](#)
- [23] Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a Wasserstein loss. *Advances in neural information processing systems*, 28, 2015. [3](#)
- [24] Karan N. Chadha Gary Cheng and John C. Duchi. Fine-tuning is Fine in Federated Learning. *CoRR*, abs/2108.07313, 2021. [6](#)
- [25] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2066–2073. IEEE, 2012. [6](#), [2](#)
- [26] Tao Guo, Song Guo, and Junxiao Wang. pFedPrompt: Learning Personalized Prompt for Vision-Language Models in Federated Learning. In *Proceedings of the ACM Web Conference 2023*, pages 1364–1374, 2023. [1](#), [3](#), [6](#), [8](#)
- [27] Tao Guo, Song Guo, Junxiao Wang, Xueyang Tang, and Wenchao Xu. PromptFL: Let federated participants cooperatively learn prompts instead of models-federated learning in age of foundation model. *IEEE Transactions on Mobile Computing*, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#), [4](#)

- [28] Filip Hanzely and Peter Richtárik. Federated learning of a mixture of global and local models. *arXiv preprint arXiv:2002.05516*, 2020. **2**
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **6**
- [30] Yun He, Steven Zheng, Yi Tay, Jai Gupta, Yu Du, Vamsi Aribandi, Zhe Zhao, YaGuang Li, Zhao Chen, Donald Metzler, et al. Hyperprompt: Prompt-based task-conditioning of transformers. In *International Conference on Machine Learning*, pages 8678–8690. PMLR, 2022. **2**
- [31] Wei Huang, Ye Shi, Zhongyi Cai, and Taiji Suzuki. Understanding convergence and generalization in federated learning through feature learning theory. In *The Twelfth International Conference on Learning Representations*, 2023. **2**
- [32] Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized Cross-Silo Federated Learning on Non-IID Data. In *AAAI*, pages 7865–7873, 2021. **2, 6**
- [33] Leonid V Kantorovich. On the translocation of masses. *Journal of mathematical sciences*, 133(4):1381–1382, 2006. **2, 3**
- [34] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023. **2**
- [35] Mikhail Khodak, Maria-Florina F Balcan, and Ameet S Talwalkar. Adaptive gradient-based meta-learning methods. *Advances in Neural Information Processing Systems*, 32, 2019. **2**
- [36] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. **6, 2**
- [37] Nathaniel Lahn, Sharath Raghvendra, and Kaiyi Zhang. A Combinatorial Algorithm for Approximating the Optimal Transport in the Parallel and MPC Settings. *Advances in Neural Information Processing Systems*, 36, 2024. **3**
- [38] Bin Li, Ye Shi, Qian Yu, and Jingya Wang. Unsupervised cross-domain image retrieval via prototypical optimal transport. *arXiv preprint arXiv:2402.18411*, 2024. **3**
- [39] Hongxia Li, Zhongyi Cai, Jingya Wang, Jiangnan Tang, Weiping Ding, Chin-Teng Lin, and Ye Shi. FedTP: Federated Learning by Transformer Personalization. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. **2, 6**
- [40] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10713–10722, 2021. **1**
- [41] Qian Li, Zhichao Wang, Shaowu Liu, Gang Li, and Guandong Xu. Causal optimal transport for treatment effect estimation. *IEEE transactions on neural networks and learning systems*, 2021. **3**
- [42] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020. **2, 6, 7, 4**
- [43] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pages 6357–6368. PMLR, 2021. **2**
- [44] Xiaoxiao Li, Meirui JIANG, Xiaofei Zhang, Michael Kamp, and Qi Dou. FedBN: Federated Learning on Non-IID Features via Local Batch Normalization. In *International Conference on Learning Representations*, 2020. **1, 2, 6**
- [45] Paul Pu Liang, Terrance Liu, Liu Ziyin, Nicholas B Allen, Randy P Auerbach, David Brent, Ruslan Salakhutdinov, and Louis-Philippe Morency. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*, 2020. **2**
- [46] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35, 2023. **1**
- [47] Yajing Liu, Yuning Lu, Hao Liu, Yaozu An, Zhuoran Xu, Zhuokun Yao, Baofeng Zhang, Zhiwei Xiong, and Chenguang Gui. Hierarchical Prompt Learning for Multi-Task Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10888–10898, 2023. **2**
- [48] Wang Lu, Xixu Hu, Jindong Wang, and Xing Xie. FedCLIP: Fast Generalization and Personalization for CLIP in Federated Learning. *arXiv preprint arXiv:2302.13485*, 2023. **2**
- [49] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5206–5215, 2022. **2**
- [50] Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020. **2, 10**
- [51] Othmane Marfoq, Giovanni Neglia, Richard Vidal, and Laetitia Kameni. Personalized Federated Learning through Local Memorization. In *International Conference on Machine Learning*, pages 15070–15092. PMLR, 2022. **2**
- [52] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. **1**
- [53] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018. **8**
- [54] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008. **6, 2**
- [55] Jaehoon Oh, Sangmook Kim, and Se-Young Yun. FedBABU: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*, 2021. **2**
- [56] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on*

- computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. [6](#), [7](#), [2](#)
- [57] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [3](#)
- [58] 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. [6](#), [2](#)
- [59] Abhijeet Phatak, Sharath Raghvendra, Chittaranjan Tripathy, and Kaiyi Zhang. Computing all optimal partial transports. In *The Eleventh International Conference on Learning Representations*, 2022. [3](#)
- [60] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [4](#)
- [61] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems*, 32(8):3710–3722, 2020. [2](#)
- [62] Aviv Shamsian, Aviv Navon, Ethan Fetaya, and Gal Chechik. Personalized federated learning using hypernetworks. In *International Conference on Machine Learning*, pages 9489–9502. PMLR, 2021. [6](#)
- [63] Richard Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, 74(4):402–405, 1967. [3](#)
- [64] Shangchao Su, Mingzhao Yang, Bin Li, and Xiangyang Xue. Cross-domain federated adaptive prompt tuning for clip. *arXiv preprint arXiv:2211.07864*, 2022. [2](#)
- [65] Benyuan Sun, Hongxing Huo, Yi Yang, and Bo Bai. Partialfed: Cross-domain personalized federated learning via partial initialization. *Advances in Neural Information Processing Systems*, 34:23309–23320, 2021. [2](#)
- [66] Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020. [2](#)
- [67] Yue Tan, Guodong Long, Jie Ma, Lu Liu, Tianyi Zhou, and Jing Jiang. Federated learning from pre-trained models: A contrastive learning approach. *Advances in Neural Information Processing Systems*, 35:19332–19344, 2022. [2](#), [6](#)
- [68] William Torous, Florian Gunsilius, and Philippe Rigollet. An optimal transport approach to causal inference. *arXiv preprint arXiv:2108.05858*, 2021. [3](#)
- [69] Ruibo Tu, Kun Zhang, Hedvig Kjellström, and Cheng Zhang. Optimal transport for causal discovery. *arXiv preprint arXiv:2201.09366*, 2022. [3](#)
- [70] Saeed Vahidian, Mahdi Morafah, Weijia Wang, Vyacheslav Kungurtsev, Chen Chen, Mubarak Shah, and Bill Lin. Efficient distribution similarity identification in clustered federated learning via principal angles between client data subspaces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10043–10052, 2023. [2](#)
- [71] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [6](#)
- [72] Kangkang Wang, Rajiv Mathews, Chloé Kiddon, Hubert Eichner, Françoise Beaufays, and Daniel Ramage. Federated evaluation of on-device personalization. *arXiv preprint arXiv:1910.10252*, 2019. [2](#)
- [73] Jian Xu, Xinyi Tong, and Shao-Lun Huang. Personalized federated learning with feature alignment and classifier collaboration. *arXiv preprint arXiv:2306.11867*, 2023. [2](#)
- [74] Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient model personalization in federated learning via client-specific prompt generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19159–19168, 2023. [1](#), [3](#)
- [75] Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, and Yanfeng Wang. FedDisco: Federated Learning with Discrepancy-Aware Collaboration. *arXiv preprint arXiv:2305.19229*, 2023. [2](#)
- [76] Haodong Zhao, Wei Du, Fangqi Li, Peixuan Li, and Gongshen Liu. Reduce communication costs and preserve privacy: Prompt tuning method in federated learning. *arXiv preprint arXiv:2208.12268*, 2022. [1](#), [2](#)
- [77] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [2](#), [3](#)
- [78] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [1](#), [2](#), [3](#), [6](#), [7](#), [4](#)