

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

11-2022

DAOT: Domain-Agnostically Aligned Optimal Transport for domain-adaptive crowd counting

Huilin ZHU

Jingling YUAN

Xian ZHONG

Zhengwei YANG

Zheng WANG

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Graphics and Human Computer Interfaces Commons](#), and the [Numerical Analysis and Scientific Computing Commons](#)

Citation

ZHU, Huilin; YUAN, Jingling; ZHONG, Xian; YANG, Zhengwei; WANG, Zheng; and HE, Shengfeng. DAOT: Domain-Agnostically Aligned Optimal Transport for domain-adaptive crowd counting. (2022). *MM '23: Proceedings of the 31st ACM International Conference on Multimedia, Ottawa, October 29 - November 3*. 4319-4329.

Available at: https://ink.library.smu.edu.sg/sis_research/8423

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Huilin ZHU, Jingling YUAN, Xian ZHONG, Zhengwei YANG, Zheng WANG, and Shengfeng HE

DAOT: Domain-Agnostically Aligned Optimal Transport for Domain-Adaptive Crowd Counting

Huilin Zhu

School of Computer Science and
Artificial Intelligence,
Wuhan University of Technology
zsj_zhl@whut.edu.cn

Jingling Yuan

School of Computer Science and
Artificial Intelligence,
Wuhan University of Technology
yjl@whut.edu.cn

Xian Zhong*

School of Computer Science and
Artificial Intelligence,
Wuhan University of Technology
zhongx@whut.edu.cn

Zhengwei Yang

National Engineering Research
Center for Multimedia Software,
School of Computer Science,
Wuhan University
yzw_aim@whu.edu.cn

Zheng Wang

National Engineering Research
Center for Multimedia Software,
School of Computer Science,
Wuhan University
wangzwhu@whu.edu.cn

Shengfeng He*

School of Computing and
Information Systems,
Singapore Management University
shengfenghe@smu.edu.sg

ABSTRACT

Domain adaptation is commonly employed in crowd counting to bridge the domain gaps between different datasets. However, existing domain adaptation methods tend to focus on inter-dataset differences while overlooking the intra-differences within the same dataset, leading to additional learning ambiguities. These domain-agnostic factors, *e.g.*, density, surveillance perspective, and scale, can cause significant in-domain variations, and the misalignment of these factors across domains can lead to a drop in performance in cross-domain crowd counting. To address this issue, we propose a Domain-agnostically Aligned Optimal Transport (DAOT) strategy that aligns domain-agnostic factors between domains. The DAOT consists of three steps. First, individual-level differences in domain-agnostic factors are measured using structural similarity (SSIM). Second, the optimal transfer (OT) strategy is employed to smooth out these differences and find the optimal domain-to-domain misalignment, with outlier individuals removed via a virtual “dustbin” column. Third, knowledge is transferred based on the aligned domain-agnostic factors, and the model is retrained for domain adaptation to bridge the gap across domains. We conduct extensive experiments on five standard crowd-counting benchmarks and demonstrate that the proposed method has strong generalizability across diverse datasets. Our code will be available at: <https://github.com/HopooLinZ/DAOT/>.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**; • **Computing methodologies** → **Computer vision**.

*Corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0108-5/23/10...\$15.00

<https://doi.org/10.1145/3581783.3611793>

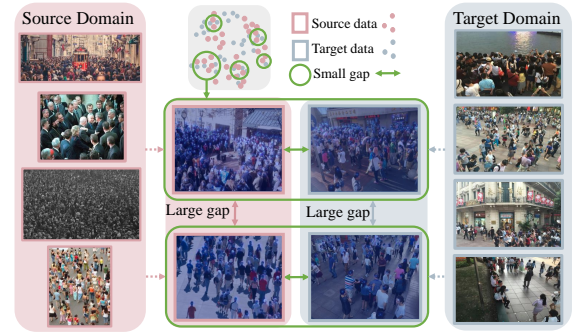


Figure 1: Illustration of intra-domain and inter-domain variations in the source (red) and target (blue) domains for factors, *e.g.*, density, scale, and surveillance perspective (domain-agnostic information). The green boxes highlight the small gap between images from different domains, where the inter-domain variations in these factors are smaller than their intra-domain variations.

KEYWORDS

Crowd Counting, Domain Adaptation, Domain-Agnostic Alignment, Optimal Transport

ACM Reference Format:

Huilin Zhu, Jingling Yuan, Xian Zhong, Zhengwei Yang, Zheng Wang, and Shengfeng He. 2023. DAOT: Domain-Agnostically Aligned Optimal Transport for Domain-Adaptive Crowd Counting. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*, October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3581783.3611793>

1 INTRODUCTION

Crowd counting is a critical task in multimedia and computer vision [12, 14, 42, 45, 47, 49] that aims to estimate the number of people in surveillance images or videos. Three types of methods are commonly used for crowd counting: detection-based [1, 40], regression-based [2], and density map-based [4, 17, 18, 21, 23, 25, 46, 50]. Detection-based methods are unsuitable for estimating the number of people in dense crowds due to occlusion issues, while

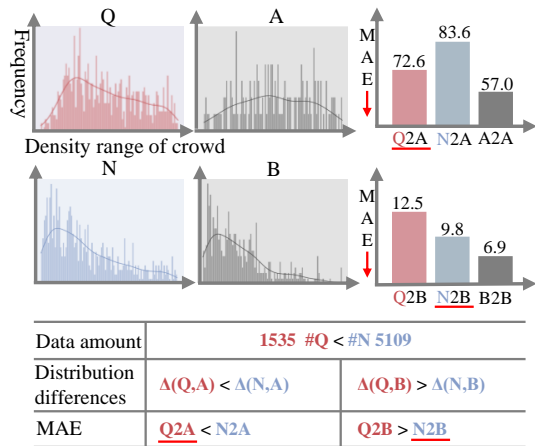


Figure 2: Illustration of density distribution and model performance in cross-domain scenarios. Histograms show density proportions of datasets Q, A, N, and B. Performance under different scenarios are displayed in histograms. Tables provide details on data amounts, distribution differences, and cross-domain performance, highlighting the best performance in red. (Baseline model: FIDT [17].)

regression-based methods are less affected by occlusion but cannot obtain the approximate position distribution of people. The emergence of density map-based methods has addressed the shortcomings of the previous two types of methods and has become the mainstream method for crowd counting. However, most crowd counting methods rely on extensive data annotation and cannot be applied to open scenes. The effect of directly training models with labeled datasets and transferring them to unlabeled scenes is often unsatisfactory due to domain gap issues. As a consequence, unsupervised domain adaptation (UDA) methods [5, 9, 34, 37, 39] have been widely applied in crowd counting to bridge the domain gap. UDA methods in crowd counting cover Synthetic-Real (Syn-Real) or Real-Real scenarios. However, Syn-Real methods often require significant computational resources and cannot capture real crowd characteristics accurately. For the more practical Real-Real cross-domain problem, previous works propose Gaussian process-based iterative learning [31], and combine detection and density estimation with dual knowledge transfer [22]. Nevertheless, these methods prioritize improving their knowledge and fall short in exploring inter-domain relations. To exploit the inter-domain relation, AS-Net [53] use discriminators to evaluate the similarity between the source and target domains at various levels of granularity, while FGFD [51] utilize inter-domain invariant information for model migration. However, as depicted in Fig. 1, many factors within the same domain, *e.g.*, density, perspective, and scale, exhibit more significant intra-domain differences than inter-domain differences, limiting the ability to discriminate or find similarities between single images across domains. Some research also focuses on domain generalization, *e.g.*, C2MoT [41] and DGCC [8]. In C2MoT, only domain-invariant information is explored, but not domain-specific information. In DGCC, although domain information is divided into domain invariant features and domain-specific features, they

are considered mutually exclusive and the factual domain-specific features are not explored.

In this paper, we have three key discoveries related to the cross-domain issue in the context of crowd counting. As depicted in Fig. 1, factors that are domain-agnostic, *e.g.*, density, perception, and scale, can lead to more significant variations within the same dataset compared to those observed across different domains. This because many of these factors are not unique to a particular domain and can vary widely even within the same dataset.

Meanwhile, we observe that domain-agnostic factors are misaligned across domains, and the cross-domain performance tends to improve as the different domains become more distributionally consistent. We visualize the distribution of density from UCF-QNRF (Q), NWPU-Crowd (N), ShanghaiTech PartB (A), and ShanghaiTech PartB (B) [48] at the top of Fig. 2. We discover that the frequency distribution of A and Q are more similar, with a higher frequency of samples with a medium crowd density range. Similarly, the frequency distribution of N and B are also more similar, with a higher frequency of samples with a low crowd density range. Moreover, we notice that the performance of the model [17] in different cross-domain settings varies significantly. Specifically, the N2B setting outperforms the Q2B setting, and Q2A outperforms N2B, with Q2A even approaching the A2A performance. This discovery indicates that misalignment of domain-agnostic factors at the domain level significantly impacts the model’s cross-domain performance [17].

Furthermore, we discover that the effect of domain misalignment on cross-domain performance can outweigh the impact of the quantity of data available in the source domain. As shown in the bottom of Fig. 2, the Q [13] contains 1,535 images, while the N [38] contains 5,109 images. However, when the same settings are applied to the target domain A [48], which consists of 182 images, N performs worse than Q. This discovery underscores the criticality of aligning domain-agnostic factors, including compensating for insufficient data quantity, in optimizing cross-domain model performance.

Based on the above three discoveries, we posit that the domain gap is not solely dependent on inter-domain factors at the image level, but is strongly influenced by the misalignment of domain-agnostic factors at the domain level. In this paper, the domain factors can be concretely represented as the crowd distribution, as this encompasses information such as density, perspective, scale, and other related factors. The objective of this paper is to address domain adaptation by resolving the misalignment of domain-agnostic factors that contribute to the domain gap. Therefore, we propose a three-step domain-adaption approach in an align-and-transfer manner. In the first step, we measure the differences in domain-agnostic factors at the individual level using structural similarity (SSIM) as an indicator. In the second step, we smooth out these domain-agnostic differences. For this purpose, we introduce the optimal transport (OT) strategy to find the optimal domain-to-domain transformation. To handle outlier individuals in the misalignment measurement process, we augment each domain with a virtual column representing a “dustbin” individual, where outlier individuals are removed. Finally, we can transfer knowledge based on the aligned domain-agnostic factors. By aligning the domain-agnostic factors across domains through the optimal domain-to-domain transformation, we retrain the model [17] for domain adaptation and bridge the gap across domains.

To summarize, the main contributions are threefold:

- We propose a new perspective on the impact of domain-agnostic factors in crowd counting, highlighting their significant influence on in-domain variations and inferior cross-domain performance due to misalignment at the domain level. This discovery provides a better understanding of the domain gap and the importance of domain alignment for improving cross-domain performance.
- We design a novel domain adaptation method, called Domain-agnostically Aligned Optimal Transport (DAOT), to bridge the domain gap by aligning domain-agnostic factor misalignments between the source and target domains.
- We demonstrate the effectiveness and strong generalizability of the proposed DAOT through extensive experiments on five crowd counting datasets. Experimental results indicate that DAOT outperforms state-of-the-art methods.

2 RELATED WORK

2.1 Optimal Transport

It is well-known that optimal transport (OT) [28, 29] has been widely applied in domain adaptation tasks, particularly in generating target domain data by finding the optimal distance between the source and target domain feature distributions. However, most of these methods [3, 7, 44] focus on generating images similar to those in the target domain, which requires finding the optimal distance between the generated and target images. In contrast, the proposed DAOT differs from traditional methods of OT in domain adaptation, as we establish correspondence between source and target domain distributions without generating new images or altering source domain images, similar to a feature-matching task [15].

Compared to point-to-point matching methods [11, 30, 36], the proposed DAOT has a broader scope as it involves region-to-region matching. Moreover, we use a distance measurement method that emphasizes spatial structure, namely SSIM. Note that we do not perform person-level matching, as opposed to DR.VIC [11]. Our focus is not on the existence of a matching relation between the query and gallery, but on specific matching and transfer relations.

In crowd counting, OT has been used as a loss function, with DM-Count [35] being one of the first proposed methods. Several methods [19, 33] have adopted this type of loss. Other than loss functions, we use OT to find the best match rather than simply measuring the difference.

2.2 Unsupervised Domain Adaptation

Most UDA crowd counting methods [9, 24, 37, 39, 42] focus on classifying individual images into either the source or target domain, and therefore aim to bridge the gap between the source and target domains on a per-image basis. However, in Real-Real domain adaptation tasks, many images may have greater intra-domain differences than inter-domain ones. Addressing only the differences within individual images may not be sufficient to fundamentally bridge the gap between domains. Instead, we propose a domain adaptation method based on aligning the overall distribution of data in the source and target domains, rather than focusing solely on individual images. Several crowd counting methods [51, 53] consider that there are similarities between source and target domains

and thus adapt between domains by finding similarities. However, individual similarity retrieval has shifted. Perhaps the found source and target distributions do have strong similarities, but due to the existence of retrieval bias, much of the information in the source distribution is not fully utilized, and only the distribution with the highest similarity is left. In essence, this approach is unable to fully simulate the entire data domain, while the proposed DAOT seeks to maximize the search for transfer relation between the two distributions. Apart from the methods that seek to find the similarities between domains, other methods [8, 41] attempt to find common knowledge between them. However, in our opinion, many images in different domains differ significantly in density, scale, and perception, making it difficult to find common knowledge. Even within the same domain, it is challenging to find common knowledge. We believe that rich data is essential for the model to learn knowledge and to have stronger generalization abilities. Therefore, instead of seeking common knowledge, we bridge the domain gap by aligning domain-agnostic factor misalignment between the source and target domains.

3 PROPOSED METHOD

3.1 Problem Definition

In this paper, we consider domain adaptation as a problem of aligning domain-agnostic factors between domains, and we employ a region-level distribution-perception method, where the distribution is used to describe the domain-agnostic factors. As shown in Fig. 3, the source domain region is set to $S = x_i, i \in (1, n)$, and the target domain region is set to $T = x_j, j \in (1, m)$. The model trained with labeled data in the source domain is defined as M_S , the initial model in the target domain is defined as M_{T0} , with the training labels using the pseudo-labels predicted by the source domain model on the target domain data set. And the regions in the two domains are distribution-perception to obtain the source domain distribution $D^S = d_i^S$ and the target domain distribution $D^T = d_j^T$, respectively. In Sec. 3.2, the distance matrix C is calculated using SSIM to measure the distance between each distribution and extended to form the cost matrix \bar{C} , and the final obtained simulated distribution in Sec. 3.3 is defined as $D^M = D_{r,j}^S, \{r\} \subseteq \{i\}, i \in (1, n), j \in (1, m)$. The main method for domain alignment is optimal transport (OT).

The principle behind OT is to minimize the cost of transporting one distribution to another, where the cost is usually defined as the distance between the two distributions. The OT problem can be formulated as a linear programming problem, where the objective function is to minimize the cost of transportation subject to some constraints. Therefore, the first step is to obtain a probability transition matrix that expresses the transfer between two distributions, called the cost matrix \bar{C} . Then, based on the cost matrix, the Sinkhorn algorithm [6, 29, 32] to find the optimal solution \bar{P} solved to find the optimal solution for transporting between the distributions. In our paper, local distribution matching can be achieved by solving an augmented version of Kantorovich’s OT problem [29]. And $U(A, B)$ denotes the discrete of A and B .

$$L_{\bar{C}}(\bar{A}, \bar{B}) = \min_{P \in U(\bar{A}, \bar{B})} \sum_{i,j}^{n,m} (\bar{C}_{i,j} \bar{P}_{i,j}), \quad (1)$$

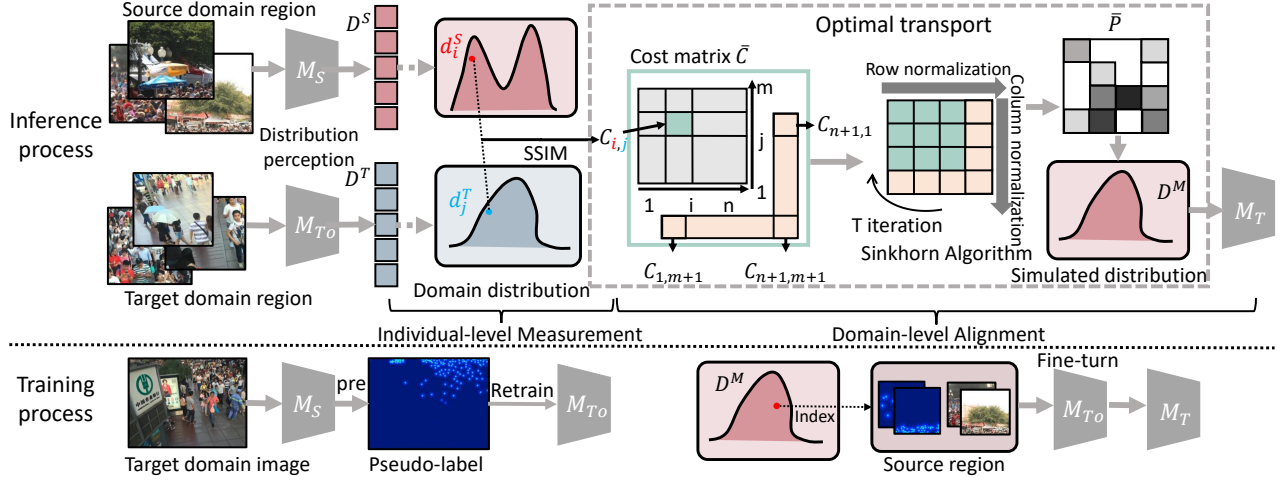


Figure 3: Structure of DAOT. The figure illustrates the inference and training processes, where the training process describes the acquisition of M_{T_0} and M_T , respectively. The inference process is divided into two stages. In the individual-level measurement stage (stage 1), region information from both the source and target domains is collected. The source domain model M_S and target domain model M_{T_0} trained with pseudo-labels from M_S are used for distribution perception, yielding the source domain distribution D^S and target domain distribution D^T . The distance matrix C is calculated by SSIM to measure the distance between each distribution and extended to form the cost matrix \bar{C} . In the domain-level alignment stage (stage 2), we use the Sinkhorn algorithm with iterative updates to obtain the optimal transfer matrix solution \bar{P} and the final simulated distribution. We fine-tune the initial model M_{T_0} using the simulated distribution to obtain the target domain model M_T .

where \bar{P} is a transportation matrix. $\bar{C}_{i,j}$ represents the distance between each element in the source and target domains, in Eq. (2).

3.2 Individual-Level Measurement

We believe that domain-agnostic factors are variable in both inter-domains and intra-domain, which makes individual images not exclusively belong to a particular domain, while misalignment of these factors at the domain level lead to the domain gap. Therefore, this paper aims to bridge the domain gap by aligning domain-agnostic factors. To achieve domain-level alignment, we measure individual-level differences using the following algorithm.

3.2.1 Describing Domain-Agnostic Factors. Due to the diversity of distributions increasing with scale, it is challenging to align at the individual level without altering the image itself. Thus, we align the regions in the image to address this issue. Therefore, we divide the whole image in the source and target domains into different regions and forms the sets $S = x_i$ and $T = y_j$. Subsequently, the distributions D^S and D^T of the source and target domains are obtained after distribution perception of the source domain model M_S and the preliminary target domain model M_{T_0} . We utilize FIDT [17] as our location-sensitive distribution-perception model.

3.2.2 Structural Similarity and Cost Matrix. The solution of the transfer matrix between the distributions heavily relies on the distance metric function between the distributions. In the crowd counting distribution distance metric, due to the huge and dense crowd, it possesses rich brightness, contrast, and structural variation. The SSIM similarity can fully combine this information, and the calculated similarity is not only the value of the vector but the structural properties in the whole distribution.

For a given source domain distribution set D^S and target domain distribution set D^T , we can calculate an SSIM distance, $C_{i,j}$, between each element d_i^S and d_j^T . The formula is defined as Eq. (2).

$$C_{i,j} = 1 - \text{SSIM}(d_i^S, d_j^T) = 1 - \frac{(2\mu_{d_i^S} \mu_{d_j^T} + Q_1)(2\sigma_{d_i^S, d_j^T} + Q_2)}{(\mu_{d_i^S}^2 + \mu_{d_j^T}^2 + Q_1)(\sigma_{d_i^S}^2 + \sigma_{d_j^T}^2 + Q_2)}, \quad (2)$$

$$C = \begin{bmatrix} C_{1,1} & \cdots & C_{i,1} & \cdots & C_{n,1} \\ C_{1,j} & \cdots & C_{i,j} & \cdots & C_{n,j} \\ C_{1,m} & \cdots & C_{i,m} & \cdots & C_{n,m} \end{bmatrix}, \quad (3)$$

where $\mu_{d_i^S}$ and $\sigma_{d_j^T}$ separately denote the mean of d_i^S and standard deviation of d_j^T . $\sigma_{d_i^S, d_j^T}$ is covariance of d_i^S and d_j^T . We set Q_1 and Q_2 as e^{-4} and $9e^{-4}$. \bar{C} as the final cost matrix, the shape of $(n+1, m+1)$.

$$\bar{C} = \begin{bmatrix} C_{i,j} & C_{n+1,1} \\ C_{1,m+1} & C_{n+1,m+1} \end{bmatrix}, \quad (4)$$

where \bar{C} is used to represent the distance between source domain region distribution and target domain region distribution. As shown in Fig. 4, to ensure the reliability of the distance scores, we set $C_{i,m+1}$, $C_{n+1,j}$, and $C_{n+1,m+1}$ as ‘‘dustbin’’ column to remove unreliable rows and columns, where $C_{i,m+1}$ represents SSIM value between the distribution of each source domain sample and a distribution consisting of all ones, while $C_{n+1,j}$ represents SSIM value between the distribution of each target domain sample and the same distribution consisting of all ones. $C_{n+1,m+1}$ represents the SSIM value between a matrix consisting of all zeros and the same distribution consisting of all ones, which serves as a threshold for determining the reliability of each element in the distance matrix \bar{C} .

If the minimum value in any row or column of C exceeds $C_{n+1,m+1}$, the corresponding row or column is zeroed out to filter out unreliable elements and enhance the Sinkhorn algorithm's robustness.

3.3 Domain-Level Alignment

In this paper, the inter-domain adaptation task is transformed into an inter-domain domain-agnostic factor alignment problem, and further, the alignment of domain-agnostic factors is considered as finding the optimal mapping between domain-level distributions, which are used to describe domain-agnostic factors at the domain level. It is also the solution of Eq. (1). The function we use to find the optimal transport is Sinkhorn [32]. The ultimate goal of the OT problem is to find a matrix \bar{P} that optimally transfers the source distribution to the target distribution so that the product of $\sum (\bar{P} \times \bar{C})$. This is essentially a linear programming problem Eq. (1) with $n + m + 2$ constraints Eq. (4), the constraints ensure that the probabilities of mapping between the source and target distributions sum up to one for each element, and vice versa.

3.3.1 Sinkhorn Algorithm. Sinkhorn Algorithm [29, 32] is applied to solve for \bar{P} , which represents the cost of transporting each element from the source domain to the target domain. The algorithm computes the optimal transport plan from the source domain to the target domain and returns a mapping between the indices of the source and target elements that minimizes the overall cost of the transport plan. In Sec. 3.3, the cost matrix C , between the domains has been obtained.

This algorithm is a differentiable version of the Hungarian algorithm [27], which is typically used for bipartite matching. It involves iterative normalizing \bar{C} along rows and columns, similar to row and column Softmax, and the DAOT for calculating discrete is SSIM. This means minimizing Eq. (1). Specifically, subject to the constraints that the marginals of \bar{P} match the given weights:

$$\sum_j \bar{P}_{i,j} = \bar{A}_i, \quad \sum_i \bar{P}_{i,j} = \bar{B}_j, \quad (5)$$

where the algorithm updates the rows of \bar{P} as follows:

$$\bar{A}_i \frac{\bar{B}_j}{\sum_k \bar{P}_{i,k} \bar{C}_{i,k} + \epsilon} \rightarrow \bar{P}_{i,j}, \quad (6)$$

where k is an index variable used for summation, it represents that for the i -th source sample and the j -th target sample, we need to sum over all k , which represents the k -th sample in the source domain. Therefore, the formula can be understood as computing the transition probability between the i -th source sample and the j -th target sample and using information from all source domain samples to calculate it. And then updates the columns of P by applying the same operation to the transpose of P . The regularization parameter $\epsilon > 0$ prevents division by zero. This row and column scaling process is repeated iteratively until the resulting transport plan \bar{P} converges to a stable solution. The symbol " \rightarrow " in Eq. 6 indicates the final output of the Sinkhorn Algorithm. The algorithm stops at $T_{\max} = 1000$, following the convention in [6].

3.3.2 Final Mapping and Target Model. After obtaining the optimal transport plan \bar{P} using the Sinkhorn algorithm, we can use it to obtain a mapping between the source and target domains. As shown

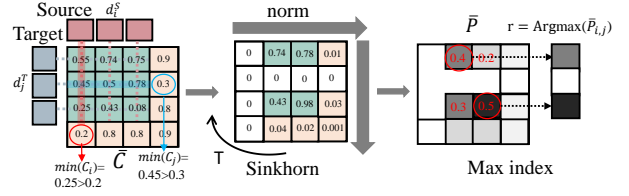


Figure 4: Illustration of OT solving process. Red: source domain. Blue: target domain. Green: SSIM distance. Yellow: garbage bin score. Yellow and green form cost matrix \bar{C} . \bar{C} regularized iteratively to obtain optimal solution \bar{P} , transfer probability matrix. Max value in each row equals final matching distribution index.

in Fig. 4, this mapping represents the correspondence between the source and target samples, and can be used to train a model that can translate samples from the source to the target domain. To obtain the mapping, we find the index i in the source domain that has the highest probability of being paired with each target sample (d_j^T, d_r^S) , which can be computed as $r = \arg \max_i \bar{P}_{i,j}$. We consider the correspondences between the source and target samples only if the corresponding probability value is non-zero. This procedure gives us the mapping between the source and target samples with non-zero correspondence values.

Finally, we obtain the distribution $D^M = D_{r,j}^S$ to fine-tune the initial model M_{T0} and obtain the final target domain model M_T .

4 EXPERIMENTAL RESULT

4.1 Datasets and Implementation Details

Shanghai Tech [48] is a benchmark for crowd counting contains 1,198 scene images and 330,165 labeled head positions, which is subdivided into PartA (A) and PartB (B). A is densely populated and contains 300 training images and 182 testing images; B is relatively sparser and contains 400 training images and 316 testing images.

UCF-QNRF (Q) [13] is a challenging crowd counting dataset with rich scenes and varying viewpoints, densities, and illumination conditions. It contains 1,535 images of dense crowd scenes, with 1,201 training images and 334 testing images.

NWPU-Crowd (N) [38] has 5,109 high resolution images and annotated entities 2,133,238. The large range of annotated entity numbers for a single image is extremely challenging. Additionally, some images have object annotation number of 0.

JHU-CROWD++ (J) [31] contains 4,372 images with a total of 1.51 million annotations, including some based on extreme weather variation and light variation, and introduces negative samples of entities annotated with zero.

Loss Function. We use the loss function proposed in FIDT [17] for training, where η is set to 1 following FIDT.

$$\mathcal{L} = \mathcal{L}_{MSE} + \eta \mathcal{L}_{I-S}. \quad (7)$$

Evaluation Metrics. In this paper, the evaluation metrics used are mean absolute error (MAE) and root mean square error (RMSE) introduced in [48]. MAE assesses model accuracy, while RMSE evaluates model robustness.

Table 1: Evaluation performance comparison of MAE and RMSE with the state-of-the-arts for cross-domain crowd counting. “Syn” and “Real” indicate the synthetic GCC dataset and the real datasets as source domains, respectively. “AD” stands for Domain Adaption. † indicates reproduced results. **Bold and underline numbers denote the best and second-best results, respectively.**

Method	backbone	Type	AD	A2B		A2Q		Q2A		Q2B		N2A		N2B		J2B		J2Q	
				MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Cycle GAN [52]	CycleGAN	Syn	●	25.4	39.7	257.3	400.6	143.3	204.3	25.4	39.7	143.3	204.3	25.4	39.7	25.4	39.7	257.3	400.6
GP [31]	ResNet50	Syn	●	12.8	19.2	210.0	351.0	121.0	181.0	12.8	19.2	121.0	181.0	12.8	19.2	12.8	19.2	210.0	351.0
DACC [9]	CycleGAN	Syn	●	13.1	19.4	203.5	343.0	112.4	176.9	13.1	19.4	112.4	176.9	13.1	19.4	13.1	19.4	203.5	343.0
BLA [10]	VGG16	Syn	●	11.9	18.9	198.9	316.1	99.3	145.0	11.9	18.9	99.3	145.0	11.9	18.9	11.9	18.9	198.9	316.1
CDCC [20]	CycleGAN	Syn	●	11.4	17.3	198.3	332.9	109.2	168.1	11.4	17.3	109.2	168.1	11.4	17.3	11.4	17.3	198.3	332.9
SPN+L2SM [43]	VGG16	Real	○	21.2	38.7	227.2	405.2	73.4	119.4	-	-	-	-	-	-	-	-	-	-
BL [26]	VGG16	Real	○	15.9	25.8	135.8	252.6	69.8	123.8	15.3	26.5	-	-	-	-	-	-	-	-
RBT [22]	VGG16	Real	●	13.4	29.3	175.0	294.8	-	-	-	-	-	-	-	-	-	-	-	-
D2CNet [4]	VGG16	Real	●	21.6	34.6	126.8	245.5	61.2	100.8	10.2	19.1	93.3	167.1	11.1	<u>16.0</u>	27.4	43.2	171.4	302.2
ASNet [53]	VGG16	Real	●	13.6	23.2	-	-	-	-	-	-	-	-	-	-	-	-	-	-
C2MoT [41]	VGG19	Real	●	<u>12.4</u>	<u>21.1</u>	125.7	218.3	-	-	-	-	-	-	-	-	-	-	135.3	251.9
FGFD [51] †	HRNet	Real	●	12.7	23.3	124.1	242.0	70.2	118.4	12.5	20.3	<u>83.2</u>	<u>138.1</u>	10.2	18.6	<u>13.4</u>	<u>20.9</u>	166.7	320.0
DGCC [8]	VGG16	Real	●	12.6	24.6	<u>119.4</u>	<u>216.6</u>	67.4	<u>112.8</u>	12.1	20.9	-	-	-	-	-	-	-	-
FIDT [17] †	HRNet	Real	○	16.0	26.9	129.3	254.1	72.6	132.3	12.5	21.7	83.6	<u>138.1</u>	<u>9.8</u>	17.0	14.2	27.3	166.6	282.4
DAOT (Ours)	HRNet	Real	●	10.9	18.8	113.9	215.6	<u>67.0</u>	128.4	<u>11.3</u>	<u>19.6</u>	79.9	135.0	8.8	14.8	10.5	17.4	<u>146.5</u>	<u>255.2</u>

Optimal Transport. In **A2Q** setting, the regularization parameter ϵ is set to $e-2$, while in **Q2B** setting, it is set to $e-1$. All other settings are set to $e-3$. The algorithm stops at $T_{\max} = 1000$.

Retraining. We use adaptive moment estimation (Adam) [16] to optimize the model with the learning rate of $e-5$, and the weight decay is set as $5e-4$. We set the size of the training batch to 16.

4.2 Comparison with the State-of-the-Arts

To evaluate the effectiveness of DAOT, this subsection compares it with other state-of-the-art cross-domain methods. We conduct quantitative evaluations through five sets of adaptation experiments, reporting the error between the predicted density maps and ground-truth maps on the target set for each dataset pair. We compare DAOT with three types of counting methods: 1) state-of-the-art counting methods without domain adaptation [4, 43], 2) Syn-Real domain adaptation methods [9, 10, 20, 31, 39, 52], and 3) Real-Real domain adaptation methods [8, 41, 51, 53], as well as our baseline method, FIDT [17]. Notably, the source domains for Syn-Real methods are large-scale synthetic GCC [39].

As shown in Table 1, the proposed DAOT achieves the best performance in five groups’ experiments (**A2B**, **A2Q**, **N2A**, **N2B**, and **J2B**) and obtained the second-best performance among the other groups (**Q2A**, **Q2B**, and **J2Q**). Especially with a significant improvement over other methods in the setting of **A2B** and **N2B**. The DAOT outperforms the baseline in all 8 settings and achieves better results than the similarity-based FGFD [51] method in multiple settings. Compared with D2CNet [4], the proposed DAOT is slightly lower than D2CNet, indicating that D2CNet is more suitable for the group with the source domain as **Q**, while the DAOT has a stronger generalization ability. Compared with C2MoT [41], the DAOT is slightly lower than C2MoT, but the amount of data used for retraining is much lower than that of C2MoT, and C2MoT is more dependent on large amounts of data. When the target domain is **B**, DAOT often achieves better performance, which proves that it has a stronger fitting ability for sparse density distribution. In the setting of **A2Q** and **N2B** with aligned domain-agnostic factors, the proposed DAOT method can further improve the excellent performance.

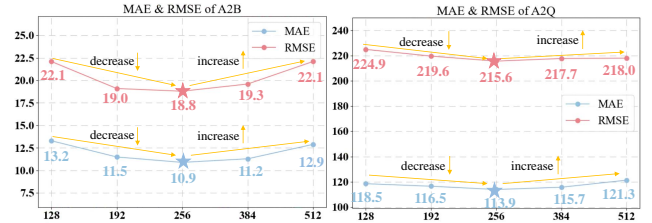


Figure 5: Line plot depicting the effect of region size on cross-domain performance, with stars indicating optimal values.

4.3 Ablation Study

Ablation Study on Region Size The transition relation between entire images across domains is limited, while the transition relation between regions exists extensively across different domains. As the region size decreases, the accuracy of distance measurement between distributions will decrease. As the region size increases, the transition relation between distributions gradually decreases. As shown in Fig. 5, as the region size decreases from the entire image to 256, the MAE and RMSE metrics correspondingly improve. However, as the region size continues to decrease beyond 256, MAE and RMSE begin to exhibit a rising trend.

Ablation Study on Various Factors Performance To achieve domain adaptation in crowd counting, we transform it into a domain-agnostic factors alignment problem and adopt optimal transport (OT) to find the best transfer distance between distributions. In Table 2, using all factors achieves the best performance in all experiments. OT significantly improved performance, especially in **Q2A** experiment. Using only pseudo-labels and SSIM achieves the second-best performance, showing a slight improvement over using only source-domain data. This suggests that adjusting the density distribution misalignment can bridge the domain gap.

Ablation Study on Distance Metric This paper aims to find the optimal transport between distributions, by measuring the distribution distance between the source and target domains before using OT. The OT ultimately obtains a more reasonable distribution

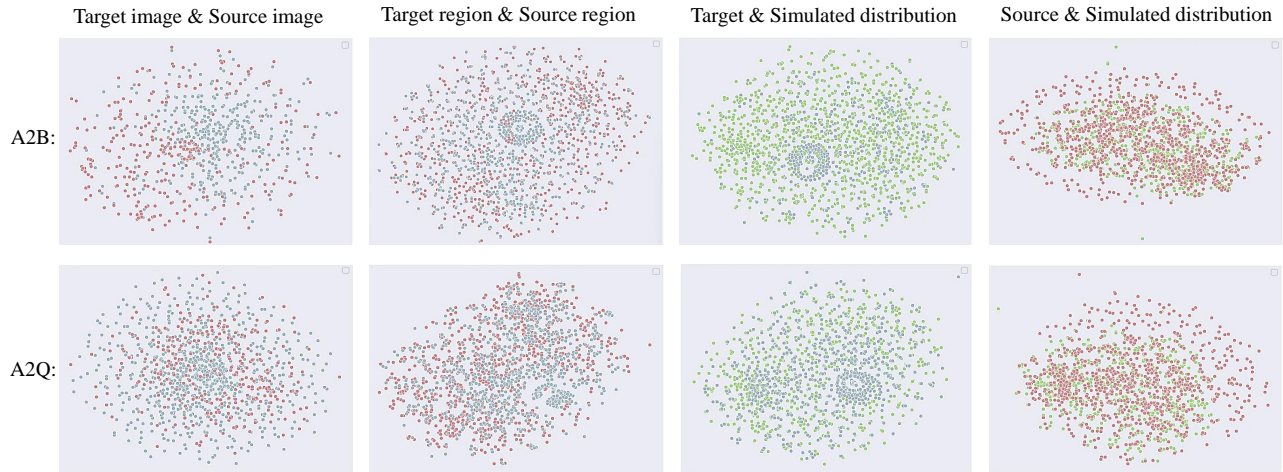


Figure 6: Visualization of t-SNE for different distributions. The two rows respectively showcase the clustering results of A2B and A2Q, where blue represents the target domain distribution, red represents the source domain distribution, and green represents the simulated distribution. The four columns illustrate the image distribution of the target and source domains, the distribution of the regions in the target and source domains, the clustering of the target domain and simulated distribution, and the clustering of the source domain and simulated distribution.

Table 2: Ablation studies of diverse factors on cross-domain performance. Each row exhibits the MAE of models trained with only source domain data, pseudo-labels generated by predicted density maps (Pre.), SSIM, and OT, respectively.

Source	Pre.	SSIM	OT	A2B		A2Q		Q2A		Q2B	
				MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
●	○	○	○	16.0	26.9	129.3	254.1	72.6	132.3	12.5	21.7
○	●	○	○	13.1	27.1	120.5	216.8	76.9	138.4	13.7	22.6
○	●	●	○	11.0	19.6	116.5	219.6	76.4	138.1	12.3	24.3
○	●	●	●	10.9	18.8	113.9	215.6	67.0	128.4	11.3	19.6

Table 3: Ablation studies of distance metrics on cross-domain performance. Each row shows the performance using Euclidean distance (ED), Kullback-Leibler (KL) divergence, and SSIM.

Metric	A2B		A2Q		Q2A		Q2B	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Source	16.0	26.9	129.3	254.1	72.6	132.3	12.5	21.7
ED	14.4	25.3	125.6	218.3	68.1	130.0	14.1	25.9
KL	11.2	21.0	114.7	218.9	80.7	150.3	13.3	27.0
SSIM	10.9	18.8	113.9	215.6	67.0	128.4	11.3	19.6

distance. The distance measurement method is crucial for the final solution, and we conduct ablation studies on different similarity metrics, which are shown in Table 3. The three distance metric methods have all shown some improvement compared to using source domain data alone, especially when using SSIM.

4.4 Qualitative Analysis

Analysis of t-SNE Visualization To more intuitively demonstrate the changes in domain-agnostic factors before and after domain adaptation, we visualize Fig. 6. In A2B, the clustering of the overall image distribution shows the most obvious difference, with

the distribution of B more concentrated in the lower left corner. This indicates that there is more similarity in the crowd distribution between A and B, but the range of A is wider. In the regional clustering, most of the data in the regions of A and B overlap, but the distribution range of A is still wider than that of B. From the third and fourth columns, it can be seen that although the simulated distribution comes entirely from the source domain, it is highly overlapped with the target domain and more concentrated than the distribution of the source domain.

In A2Q, the clustering of the overall image distribution shows higher overlap, but the distribution of Q is more concentrated in the center. In regional clustering, the distribution of Q shows a more obvious clustering phenomenon. In the visualizations of the third and fourth columns, it can be clearly seen that the simulated distribution is closer to the target domain.

Analysis of Cross-Domain Density Map Prediction To visualize the cross-domain performance of different methods in a straightforward manner, we visualize Fig. 7. The figure illustrates that the proposed DAOT method outperforms other methods in various scenarios, ranging from dense to sparse (B2A), from large-scale to small-scale (Q2A), and even from small-scale dense to large-scale (A2Q). Particularly in Q2A setting, the DAOT prediction is very close to the ground truth. In the first two rows (Q2A and A2Q), there are fewer erroneous classifications of the background, and in the third row (A2B), the estimation of the crowd is more accurate. We have highlighted these parts with red boxes.

Analysis of Distribution Match The final output distribution has a matching relation with the initial target domain distribution, and we visualize the distribution matching in Fig. 11. We demonstrate the retrieved source domain regions using different methods on the target domain regions ranging from dense to sparse. As shown in the figure, the DAOT achieves superior performance across all density levels, even in scenarios with few people or only

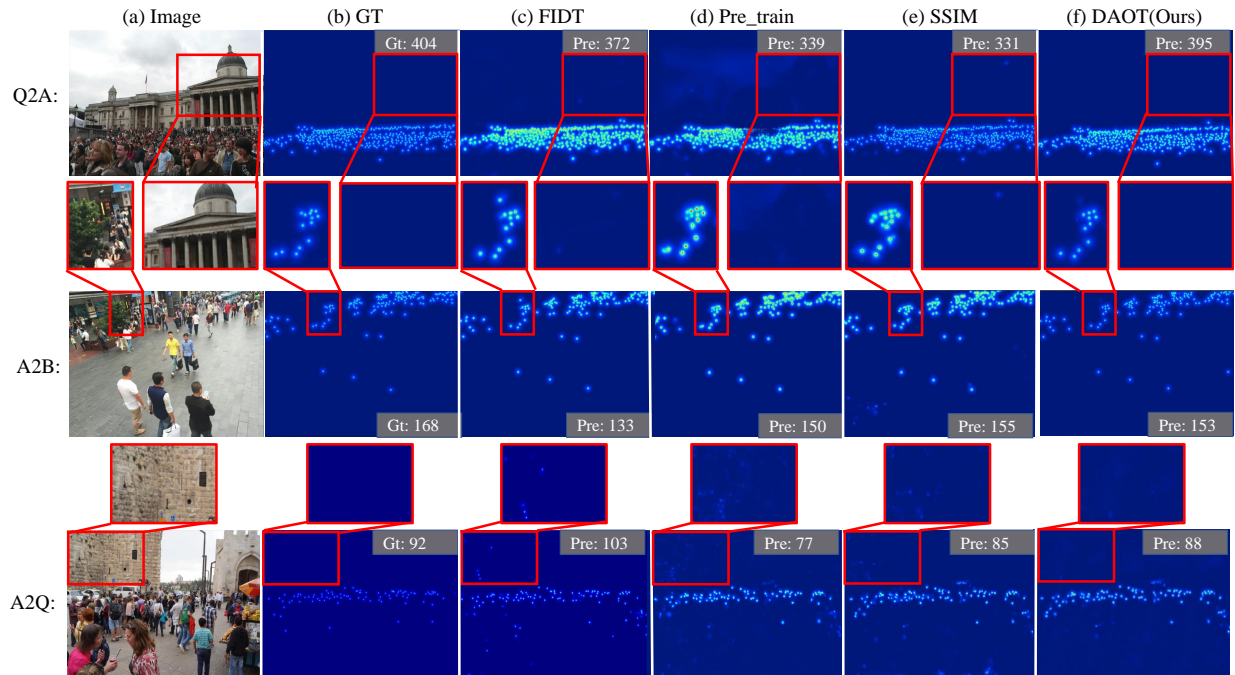


Figure 7: Visualization results in cross-domain setting involving Q2A, A2B, and A2Q. We demonstrate cross-scene results of baseline FIDT [17] (c), the method of using the target domain predicted density map as a pseudo label (d), the method only based on SSIM (e), as well as DAOT (Ours) (f). Our DAOT yields superior visual results and more accurate counting consequences than others.



Figure 8: Visualization of patches retrieved from source domain to target domain of varying densities in Q2B. Each row corresponds to the best-matching region retrieved by Diffusion, SSIM-based, and SSIM + OT (Ours), respectively, with a region size of 256 pixels.

background. Compared with the method using only SSIM, the proposed DAOT achieves better performance in the second column with denser scenarios, and has better perception simulation for the densest scenario. Compared with the diffusion-based FGFD [51], it achieves better results in the densest and sparsest scenarios, and has better simulation for the location of pedestrians.

5 CONCLUSION

In crowd counting domain adaptation task, it has been observed that the intra-domain differences are larger than the inter-domain differences in many data domains. However, even for large-scale source domains, the model’s performance in cross-domain tasks is consistently inferior to that in the source domain. Our research has revealed that this is mainly due to the misalignment of domain-agnostic factors across domains, which strongly influences the cross-domain performance of the model. Therefore, we propose to bridge the inter-domain gap in crowd counting by aligning the misalignment of domain-agnostic factors using optimal transportation in both the source and target domains. Our approach is the first domain adaptation method for crowd counting that considers the entire data domain, unlike previous methods that deal with individual distributions. We treat the overall domain distribution as a whole and achieve the best performance on multiple datasets. This work provides a new perspective on crowd counting domain adaptation, and we believe that solving the problem of domain misalignment will benefit not only cross-domain tasks but also other related applications.

ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China (62271361 and 62171325), and the Department of Science and Technology, Hubei Provincial People’s Government (2021CFB513), Guangdong Distinguished Young Scholar (2023B1515020097), and Singapore MOE Tier 1 Fund (MSS23C002).

REFERENCES

- [1] Carlos Arteta, Victor S. Lempitsky, J. Alison Noble, and Andrew Zisserman. 2014. Interactive Object Counting. In *Proc. Eur. Conf. Comput. Vis.* 504–518.
- [2] José M. Bernardo and Adrian F. M. Smith. 2001. Bayesian Theory. *Meas. Sci. Technol.* 12 (2001), 221–222.
- [3] Liangyu Chai, Yongtuo Liu, Wenxi Liu, Guoqiang Han, and Shengfeng He. 2022. CrowdGAN: Identity-Free Interactive Crowd Video Generation and Beyond. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 6 (2022), 2856–2871.
- [4] Jian Cheng, Haipeng Xiong, Zhiguo Cao, and Hao Lu. 2021. Decoupled Two-Stage Crowd Counting and Beyond. *IEEE Trans. Image Process.* 30 (2021), 2862–2875.
- [5] Yoonki Cho, Woo Jae Kim, Seunghoon Hong, and Sung-Eui Yoon. 2022. Part-based Pseudo Label Refinement for Unsupervised Person Re-identification. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 7298–7308.
- [6] Marco Cuturi. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Adv. Neural Inf. Process. Syst.* 2292–2300.
- [7] Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. 2018. DeepJDOT: Deep Joint Distribution Optimal Transport for Unsupervised Domain Adaptation. In *Proc. Eur. Conf. Comput. Vis.* 467–483.
- [8] Zhipeng Du, Jiankang Deng, and Miaojing Shi. 2022. Domain-general Crowd Counting in Unseen Scenarios. *arXiv:2212.02573* (2022).
- [9] Junyu Gao, Tao Han, Yuan Yuan, and Qi Wang. 2021. Domain-adaptive crowd counting via high-quality image translation and density reconstruction. *IEEE Trans. Neural Networks Learn. Syst.* (2021), 1–13.
- [10] Shenjian Gong, Shanshan Zhang, Jian Yang, Dengxin Dai, and Bernt Schiele. 2022. Bi-level Alignment for Cross-Domain Crowd Counting. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 7532–7540.
- [11] Tao Han, Lei Bai, Junyu Gao, Qi Wang, and Wanli Ouyang. 2022. DRVIC: Decomposition and Reasoning for Video Individual Counting. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 3073–3082.
- [12] Zongbo Han, Changqing Zhang, Huazhu Fu, and Joey Tianyi Zhou. 2023. Trusted Multi-View Classification With Dynamic Evidential Fusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 45, 2 (2023), 2551–2566.
- [13] Haroon Idrees, Muhammad Tayyab, Kishan Athrey, Dong Zhang, Somaya Al-Máadeed, Nasir M. Rajpoot, and Mubarak Shah. 2018. Composition Loss for Counting, Density Map Estimation and Localization in Dense Crowds. In *Proc. Eur. Conf. Comput. Vis.* 544–559.
- [14] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Zheng Wang, Xiao Wang, Junjun Jiang, and Chia-Wen Lin. 2021. Rain-Free and Residue Hand-in-Hand: A Progressive Coupled Network for Real-Time Image Deraining. *IEEE Trans. Image Process.* 30 (2021), 7404–7418.
- [15] Yuhe Jin, Dmytro Mishkin, Anastasiia Mishchuk, Jiri Matas, Pascal Fua, Kwang Moo Yi, and Eduard Trulls. 2021. Image Matching Across Wide Baselines: From Paper to Practice. *Int. J. Comput. Vis.* 129, 2 (2021), 517–547.
- [16] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proc. Int. Conf. Learn. Represent.* 1–11.
- [17] Dingkang Liang, Wei Xu, Yingying Zhu, and Yu Zhou. 2021. Focal Inverse Distance Transform Maps for Crowd Localization. *IEEE Trans. Multimedia* (2021), 1–13.
- [18] Hui Lin, Zhiheng Ma, Rongrong Ji, Yaowei Wang, and Xiaopeng Hong. 2022. Boosting Crowd Counting via Multifaceted Attention. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 19596–19605.
- [19] Wei Lin and Antoni B. Chan. 2023. Optimal Transport Minimization: Crowd Localization on Density Maps for Semi-Supervised Counting. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 21663–21673.
- [20] Weizhe Liu, Nikita Durasov, and Pascal Fua. 2022. Leveraging Self-Supervision for Cross-Domain Crowd Counting. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 5331–5342.
- [21] Yongtuo Liu, Sucheng Ren, Liangyu Chai, Hanjie Wu, Dan Xu, Jing Qin, and Shengfeng He. 2023. Reducing Spatial Labeling Redundancy for Active Semi-Supervised Crowd Counting. *IEEE Trans. Pattern Anal. Mach. Intell.* 45, 7 (2023), 9248–9255.
- [22] Yuting Liu, Zheng Wang, Miaojing Shi, Shin'ichi Satoh, Qijun Zhao, and Hongyu Yang. 2020. Towards unsupervised crowd counting via regression-detection bi-knowledge transfer. In *Proc. ACM Int. Conf. Multimedia.* 129–137.
- [23] Yongtuo Liu, Qiang Wen, Haoxin Chen, Wenxi Liu, Jing Qin, Guoqiang Han, and Shengfeng He. 2020. Crowd Counting Via Cross-Stage Refinement Networks. *IEEE Trans. Image Process.* 29 (2020), 6800–6812.
- [24] Yongtuo Liu, Dan Xu, Sucheng Ren, Hanjie Wu, Hongmin Cai, and Shengfeng He. 2023. Fine-grained Domain Adaptive Crowd Counting via Point-derived Segmentation. In *Proc. IEEE Int. Conf. Multimedia and Expo.*
- [25] Yu-Jen Ma, Hong-Han Shuai, and Wen-Huang Cheng. 2022. Spatiotemporal Dilated Convolution With Uncertain Matching for Video-Based Crowd Estimation. *IEEE Trans. Multimedia* 24 (2022), 261–273.
- [26] Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. 2019. Bayesian Loss for Crowd Count Estimation With Point Supervision. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.* 6141–6150.
- [27] James Munkres. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Ind. Appl. Math.* 5, 1 (1957), 32–38.
- [28] Victor M. Panaretos and Yoav Zemel. 2019. Statistical aspects of Wasserstein distances. *Annu. Rev. Stat. Its Appl.* 6 (2019), 405–431.
- [29] Gabriel Peyré, Marco Cuturi, et al. 2019. Computational optimal transport: With applications to data science. *Found. Trends Mach. Learn.* 11, 5–6 (2019), 355–607.
- [30] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. 2020. SuperGlue: Learning Feature Matching With Graph Neural Networks. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 4937–4946.
- [31] Vishwanath A Sindagi, Rajevee Yasarla, and Vishal M Patel. 2020. Jhu-crowd++: Large-scale crowd counting dataset and a benchmark method. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 5 (2020), 2594–2609.
- [32] Richard Sinkhorn and Paul Knopp. 1967. Concerning nonnegative matrices and doubly stochastic matrices. *Pac. J. Math.* 21, 2 (1967), 343–348.
- [33] Qingyu Song, Changan Wang, Zhengkai Jiang, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yang Wu. 2021. Rethinking Counting and Localization in Crowds: A Purely Point-Based Framework. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.* 3345–3354.
- [34] Geyu Tang, Xingyu Gao, Zhenyu Chen, and Huicai Zhong. 2021. Unsupervised adversarial domain adaptation with similarity diffusion for person re-identification. *Neurocomputing* 442 (2021), 337–347.
- [35] Boyu Wang, Huidong Liu, Dimitris Samaras, and Minh Hoai Nguyen. 2020. Distribution Matching for Crowd Counting. In *Adv. Neural Inf. Process. Syst.*
- [36] Haipeng Wang, Yuan Liu, Qingyong Hu, Bing Wang, Jianguo Chen, Zhen Dong, Yulan Guo, Wenping Wang, and Bisheng Yang. 2023. RoReg: Pairwise Point Cloud Registration With Oriented Descriptors and Local Rotations. *IEEE Trans. Pattern Anal. Mach. Intell.* 45, 8 (2023), 10376–10393.
- [37] Li Wang, Yongbo Li, and Xiangyang Xue. 2019. CODA: Counting Objects via Scale-Aware Adversarial Density Adaption. In *Proc. IEEE Int. Conf. Multimedia Expo.* 193–198.
- [38] Qi Wang, Junyu Gao, Wei Lin, and Xuelong Li. 2020. NWPU-crowd: A large-scale benchmark for crowd counting and localization. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 6 (2020), 2141–2149.
- [39] Qi Wang, Junyu Gao, Wei Lin, and Yuan Yuan. 2019. Learning From Synthetic Data for Crowd Counting in the Wild. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 8198–8207.
- [40] Xiao Wang, Zheng Wang, Wu Liu, Xin Xu, Jing Chen, and Chia-Wen Lin. 2021. Consistency-Constasy Bi-Knowledge Learning for Pedestrian Detection in Night Surveillance. In *Proc. ACM Int. Conf. Multimedia.* 4463–4471.
- [41] Qiangqiang Wu, Jia Wan, and Antoni B. Chan. 2021. Dynamic Momentum Adaptation for Zero-Shot Cross-Domain Crowd Counting. In *Proc. ACM Int. Conf. Multimedia.* 658–666.
- [42] Haiyang Xie, Zhengwei Yang, Huilin Zhu, and Zheng Wang. 2023. Striking a Balance: Unsupervised Cross-Domain Crowd Counting via Knowledge Diffusion. In *Proc. ACM Int. Conf. Multimedia.*
- [43] Chenfeng Xu, Kai Qiu, Jianlong Fu, Song Bai, Yongchao Xu, and Xiang Bai. 2019. Learn to Scale: Generating Multipolar Normalized Density Maps for Crowd Counting. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.* 8381–8389.
- [44] Renjun Xu, Pelen Liu, Liyan Wang, Chao Chen, and Jindong Wang. 2020. Reliable Weighted Optimal Transport for Unsupervised Domain Adaptation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 4393–4402.
- [45] Xin Xu, Xin Yuan, Z. Wang, Kai Zhang, and Ruimin Hu. 2022. Rank-in-Rank Loss for Person Re-identification. *ACM Trans. Multimedia Comput. Commun. Appl.* 18, 2s (2022), 1–21.
- [46] Zhaoyi Yan, Ruimao Zhang, Hongzhi Zhang, Qingfu Zhang, and Wangmeng Zuo. 2022. Crowd Counting Via Perspective-Guided Fractional-Dilation Convolution. *IEEE Trans. Multimedia* 24 (2022), 2633–2647.
- [47] Zhengwei Yang, Xian Zhong, Zhun Zhong, Hong Liu, Zheng Wang, and Shin'ichi Satoh. 2023. Win-Win by Competition: Auxiliary-Free Cloth-Changing Person Re-identification. *IEEE Trans. Image Process.* 32 (2023), 2985–2999.
- [48] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. 2016. Single-Image Crowd Counting via Multi-Column Convolutional Neural Network. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* 589–597.
- [49] Xian Zhong, Shidong Tu, Xianzheng Ma, Kui Jiang, Wenxin Huang, and Zheng Wang. 2022. Rainy WCity: A Real Rainfall Dataset with Diverse Conditions for Semantic Driving Scene Understanding. In *Proc. Int. Joint Conf. Artif. Intell.* 1743–1749.
- [50] Joey Tianyi Zhou, Le Zhang, Jiawei Du, Xi Peng, Zhiwen Fang, Zhe Xiao, and Hongyuan Zhu. 2022. Locality-Aware Crowd Counting. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 7 (2022), 3602–3613.
- [51] Huilin Zhu, Jingling Yuan, Zhengwei Yang, Xian Zhong, and Zheng Wang. 2022. Fine-Grained Fragment Diffusion for Cross Domain Crowd Counting. In *Proc. ACM Int. Conf. Multimedia.* 5659–5668.
- [52] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.* 2242–2251.
- [53] Zhikang Zou, Xiaoye Qu, Pan Zhou, Shuangjie Xu, Xiaoqing Ye, Wenhao Wu, and Jin Ye. 2021. Coarse to Fine: Domain Adaptive Crowd Counting via Adversarial

Scoring Network. In *Proc. ACM Int. Conf. Multimedia*. 2185–2194.

APPENDIX

A EXPERIMENT ON VIRTUAL DATA

To demonstrate the significant impact of the misalignment between the source and target domains on the cross-domain performance, we simulated a dataset (from A and Q) similar to the A test set distribution as the training set and evaluated it on the test set.

As shown in Fig. 9 and Table 4, when the crowd distribution in the training data is similar to that in the test set, the model’s performance is also close to the non-cross-domain performance. Even in cases where the data distribution is closer, less data has achieved better performance.

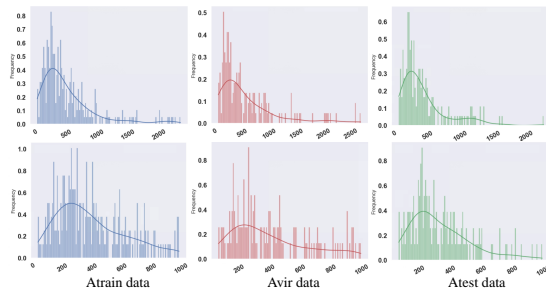


Figure 9: Histograms of crowd density frequency for the training set (blue), test set (green), and virtual dataset (red, simulated data with the same distribution as the test set) of Dataset A, with a group setting of 100. The first row shows all data, while the second row shows data from 0 to 1000.

Table 4: Comparison of performance between virtual data training and target domain training set training.

Setting	A2A	V2A
Amount of Training Data	300	295
MAE	70.9	67.0
RMSE	140.7	134.0

B VISUALIZATION OF DENSITY MAP

We show in Fig. 10 the detailed visualization of the cross-domain experiments on A2Q, J2Q, A2B, and N2B. Each column stands for the crowd image, the corresponding Ground truth, the density maps produced by the Baseline method FIDT and the results predicted by our DAOT method. Our method has been demonstrated to outperform the baseline approach in scenes with varying crowd densities, including dense-to-sparse (A2B), sparse-to-dense (A2Q), and dense-to-dense (J2Q, N2B) scenes. The primary reason is that our method leverages the alignment of domain-agnostic factors between the source and target domains, facilitating effective model migration. Specifically, our method is able to learn to match the density, surveillance perspective, scale, and other relevant factors in a domain-agnostic way, which improves the ability of the model to generalize to new domains with different crowd distributions.

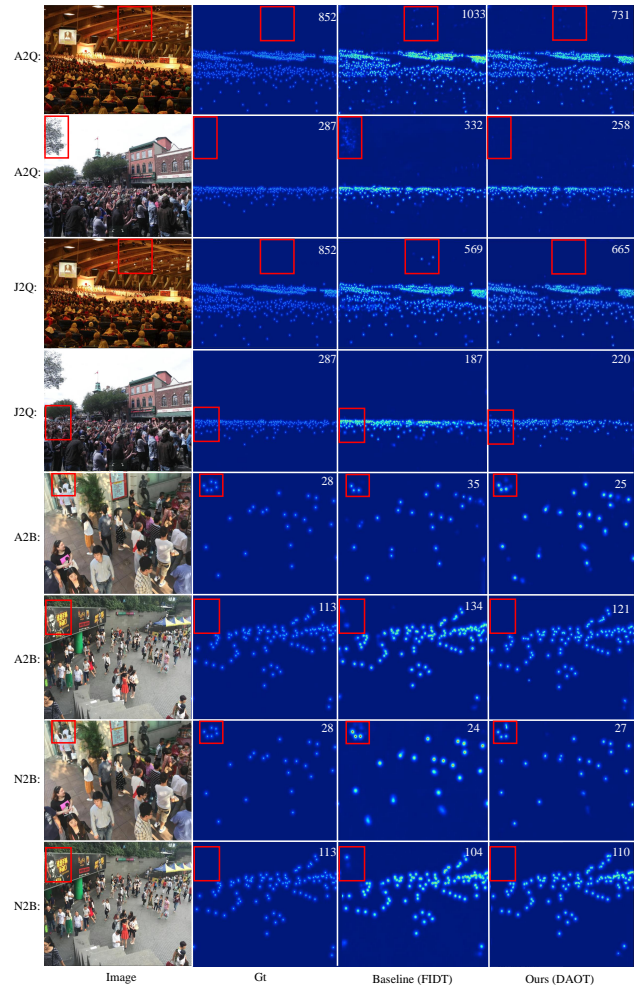


Figure 10: Visualization results of density map. The density prediction maps for A2Q, J2Q, A2B, and N2B are displayed. Each column represents the image, ground truth (GT), baseline (FIDT), and our method (DAOT), respectively.

This lets the model perceive the distribution and background better, leading to more accurate crowd counting.

C VISUALIZATION OF VARIOUS FACTORS ALIGNMENT

We visualize the distribution matching results in Fig. 11 under different settings, arranged from sparse to dense from top to bottom. Our method aligns the distributions in terms of density, head scale, and perception in each group, as demonstrated by the visualizations in Fig. 11, which are arranged from sparse to dense. Notably, in the third row of A2Q, the crowds in the two patches on the left and right are extremely similar in density, scale, perception, and even background factors. Additionally, the first row of A2B and Q2A shows similar background matching in the absence of people. These results demonstrate that our method can align various factors of crowd distributions.



Figure 11: Visualization results in distribution match. In the four sets of cross-domain experiments A2B, A2Q, Q2A, and Q2B. The shown patches come from the source and target domains and are arranged from top to bottom according to different crowd density levels.

Table 5: Ablation studies on regularization term ϵ . Each row represents the MAE and RMSE for different cross-domain settings when ϵ is set to e^{-3} , e^{-2} , e^{-1} , and 1, respectively.

Reg-term	A2B		A2Q		Q2A		Q2B	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
$\epsilon = e^{-3}$	10.9	18.8	118.5	214.9	67.0	128.4	12.4	21.4
$\epsilon = e^{-2}$	11.5	20.3	113.9	215.6	72.4	136.2	12.8	22.9
$\epsilon = e^{-1}$	11.9	20.2	116.3	217.3	73.7	142.1	11.3	19.6
$\epsilon = 1$	11.1	18.4	121.1	219.7	72.1	139.1	13.0	23.1

D REGULARIZATION TERM

The OT regularization term enhances the smoothness and stability of the OT plan by penalizing the cost function. It usually measures the distance or divergence between the transport plan and a reference distribution, often a prior distribution. This regularization prevents overfitting, improves stability and robustness, and allows incorporating prior knowledge or constraints when dealing with noisy or high-dimensional data.

As shown in Table 5, we varied ϵ from e^{-3} to 1. When ϵ was between 1 and 10, the number of mapped distribution data points was less than 100. We also experimented with ϵ set as e^{-4} and 0 in four groups. With $\epsilon = e^{-4}$, the final optimal mapping data points were all less than 50, and with $\epsilon = 0$, no mapping was obtained in any of the four groups, highlighting the importance of the regularization term. Different values for the regularization term slightly affected the performance of the four groups. A2B and Q2A achieved the best performance with $\epsilon = e^{-3}$, A2Q performed best with $\epsilon = e^{-2}$, and Q2B achieved its best performance with $\epsilon = e^{-1}$.

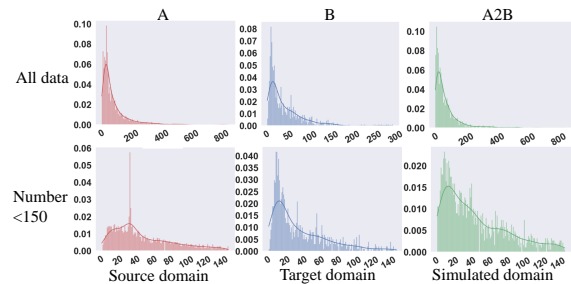


Figure 12: Visualization of frequency histograms depicting the number of people in the source domain distribution, target domain distribution, and simulated distribution (A2B).

E ANALYSIS OF DOMAIN ALIGNMENT

To demonstrate the density frequency before and after the transfer, we visualize frequency distribution histograms of the source and target domain distributions and the simulated distribution histograms after transformation, as shown in Fig. 12. Since our distribution transfer is at the regional level, we show distribution transfer when the region is set to 256. Before distribution transfer, the overall distribution of datasets A and B is concentrated in range of 0 to 150, so we show frequency distribution histograms before and after the transfer for images with 0 to 150 people. From the figure, it can be seen that A has a very high peak around 40, while the distribution of B peaks at around 10 and exhibits a sudden increase in a few specific regions. After distribution transfer, the distribution is closer to the distribution of B, peaking at around 10 and trending similarly to B overall.