

Weakly supervised cross-domain alignment with optimal transport

Siyang Yuan¹
siyang.yuan@duke.edu

Ke Bai¹
ke.bai@duke.edu

Liqun Chen¹
liqun.chen@duke.edu

Yizhe Zhang²
yizzhang@microsoft.com

Chenyang Tao¹
chenyang.tao@duke.edu

Chunyuan Li²
lichunyuan24@gmail.com

Guoyin Wang³
guoyinwang.duke@gmail.com

Ricardo Henao¹
rocardo.henao@duke.edu

Lawrence Carin¹
lcarin@duke.edu

¹ Duke University
Durham, North Carolina, USA

² Microsoft Research
Redmond, Washington, USA

³ Amazon, Alexa AI
Seattle, Washington, USA

Abstract

Cross-domain alignment between image objects and text sequences is key to many visual-language tasks, and it poses a fundamental challenge to both computer vision and natural language processing. This paper investigates a novel approach for the identification and optimization of fine-grained semantic similarities between image and text entities, under a weakly-supervised setup, improving performance over state-of-the-art solutions. Our method builds upon recent advances in optimal transport (OT) to resolve the cross-domain matching problem in a principled manner. Formulated as a drop-in regularizer, the proposed OT solution can be efficiently computed and used in combination with other existing approaches. We present empirical evidence to demonstrate the effectiveness of our approach, showing how it enables simpler model architectures to outperform or be comparable with more sophisticated designs on a range of vision-language tasks.

1 Introduction

The intersection between computer vision (CV) and natural language processing (NLP) has inspired some of the most active research topics in artificial intelligence. Prominent examples of such work includes image-text retrieval [27, 54], image captioning [17, 26, 27, 58],

[62], text-to-image generation [48, 49], phrase localization [12, 47] and visual question answering (VQA) [2, 42]. Core to these applications is the challenge of *cross-domain alignment* (CDA), consisting of accurately associating related entities across different domains in a cost-effective fashion.

Contextualized in image-text applications, the goal of CDA is two-fold: *i*) identify entities in images (*e.g.*, regions or objects) and text sequences (*e.g.*, words or phrases); and then *ii*) quantify the relatedness between identified cross-domain entity pairs. CDA is particularly challenging because it constitutes a *weakly supervised learning task*. More specifically, neither the entities nor their correspondence (*i.e.*, the match between cross-domain entities) is labeled [21]. This means that CDA must learn to identify entities and quantify their correspondence only from the image-text pairs during training.

Given the practical significance of CDA, considerable effort has been devoted to address this challenge in a scalable and flexible fashion. Existing solutions often explore heuristics to design losses that encode cross-domain correspondence. Pioneering investigations, such as [49], considered entity matching via a hinge-based ranking loss applied to shared latent features of image and text, extracted respectively with a convolutional neural network (CNN) [53] and long short-term memory (LSTM) [23] feature encoders. Explicitly modeling the between-entity relations also yields significant improvements [27]. Performance gains can also be expected via exploiting the hardest negatives in a triplet ranking loss specification [16]. More recently, synergies between CDA and attention mechanisms [43] have been explored, further advancing the state of the art with more sophisticated model designs [34].

Despite recent progress, it remains an open question concerning which other (mathematical) principles can be leveraged for scalable automated discovery of cross-domain relations. This study develops a novel solution based on recent developments in *optimal transport* (OT) based learning [11]. Briefly, OT-based learning is a generic framework that seeks to tackle specific problems by recasting them as distribution matching problems, which can then be accurately and efficiently solved by optimizing the transport distance between the distributions. Its recent success in addressing fundamental challenges in artificial intelligence has sparked a surge of interest in extending its reach to other applications [8, 9, 40].

Our work is motivated by the insight that cross-domain alignment can be reformulated as a bipartite matching problem [51], which can be optimized w.r.t. a proper matching score. We show that a solution to the challenge of automated cross-domain alignment can be approached by using the optimal transport distance as the matching score. Notably, our construction is orthogonal to the development of cross-domain attention scores [54, 43, 63], which are essentially advanced feature extractors [4] and cannot be used as an optimization criteria for the purpose of CDA *per se*, necessitating a pre-specified objective function in training for feature alignment. For example, in image captioning, maximum likelihood estimation (MLE) is applied to match generated text sequences to the reference (ground truth), and in image-text retrieval the models are typically optimized w.r.t. their ranking [32]. In this sense, the learning of attention scores is guided by MLE or ranking loss, while our OT objective can be directly optimized during training to learn optimal matching strategies.

The framework developed here makes the following contributions.

- Optimal transport is applied to construct principled matching scores for feature alignment across different domains, in particular, images and text.
- Beyond the functionality as an attention score, OT is also applied as a regularizer on the objective; thus, instead of only being used to match entities within images and text, the proposed OT regularizer can help linking image and text features globally.




Figure 1: Illustration of CDA. Left: OT matching scheme for bipartite matching. Strong signals are marked as blue lines. Upper right: manually labeled correspondence; image regions are matched with words of the same color. Lower right: the automatically learned alignment matrix using optimal transport; darker shades indicate stronger OT matching.

- The effectiveness of our framework is demonstrated on various vision-language tasks, (e.g., image-text matching and phrase localization). Experimental results show that the proposed OT-based CDA module provides consistent performance gains on all tasks.

2 Background

Optimal transport (OT). We consider the problem of transporting mass between two discrete distributions supported on some latent feature space \mathcal{X} . Let $\boldsymbol{\mu} = \{\mathbf{x}_i, \mu_i\}_{i=1}^n$ and $\boldsymbol{\nu} = \{\mathbf{y}_j, \nu_j\}_{j=1}^m$ be the discrete distributions of interest, where $\mathbf{x}_i, \mathbf{y}_j \in \mathcal{X}$ denotes the spatial locations and μ_i, ν_j , respectively, denote the non-negative masses. Without loss of generality, we assume $\sum_i \mu_i = \sum_j \nu_j = 1$. $\pi \in \mathbb{R}_+^{n \times m}$ is called a valid transport plan if its row and column marginals match $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$, respectively, that is to say $\sum_i \pi_{ij} = \nu_j$ and $\sum_j \pi_{ij} = \mu_i$. Intuitively, π transports π_{ij} units of mass at location \mathbf{x}_i to new location \mathbf{y}_j . It is known that such transport plans are not unique, and as such, one often seeks a solution $\pi^* \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})$ that is most preferable in other ways, where $\Pi(\boldsymbol{\mu}, \boldsymbol{\nu})$ denotes the set of all viable transport plans. OT finds a solution that is most cost effective w.r.t. some function $C(\mathbf{x}, \mathbf{y})$, in the sense that [16]

$$\mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\nu}) = \sum_{ij} \pi_{ij}^* C(\mathbf{x}_i, \mathbf{y}_j) = \inf_{\pi \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \sum_{ij} \pi_{ij} C(\mathbf{x}_i, \mathbf{y}_j), \quad (1)$$

where $\mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\nu})$ is known as the *optimal transport distance*. Hence, $\mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\nu})$ minimizes the transport cost from $\boldsymbol{\mu}$ to $\boldsymbol{\nu}$ w.r.t. $C(\mathbf{x}, \mathbf{y})$. Of particular interest is the case for which $C(\mathbf{x}, \mathbf{y})$ defines a distance metric on \mathcal{X} , and then $\mathcal{D}(\boldsymbol{\mu}, \boldsymbol{\nu})$ induces a distance metric on the space of probability distributions supported on \mathcal{X} , commonly known as the Wasserstein distance [57]. The use of OT allows the flexibility to choose task-specific costs for optimal performance, with examples of Euclidean cost $\|\mathbf{x} - \mathbf{y}\|_2^2$ for general probabilistic learning [49] and cosine similarity cost $\cos(\mathbf{x}, \mathbf{y})$ for semantic matching tasks [8].

Image representation. We represent an image as a collection (bag) of feature vectors $\mathbf{V} = \{\mathbf{v}_k\}_{k=1}^K$, where each $\mathbf{v}_k \in \mathbb{R}^d$ represents an image entity in feature space, and K is the number of entities. To simplify our discussion, we identify each entity as a region of interest (RoI), i.e., a bounding box, hereafter referred to as a region. We seek for these features to encode diverse visual concepts, e.g., object class, attributes, etc.

To this end, we follow [11], where $\mathbf{F} = \{\mathbf{f}_k\}_{k=1}^K$, $\mathbf{f} \in \mathbb{R}^{2048}$ is obtained from a pre-trained ResNet-101 [27] concatenated to *faster R-CNN* [50] (fR-CNN) on the heavily annotated Visual Genome dataset [50]. fR-CNN first employs a region proposal network with non-maximum suppression [42] mechanism to propose image regions, then leverages RoI pooling to construct a 2048-dimensional image feature representation, which is then used for object

classification. To project the image features into a feature space shared by sentence features (discussed below), we further apply an affine transformation to \mathbf{f}_k :

$$\mathbf{v}_k = \mathbf{W}_v \mathbf{f}_k + \mathbf{b}_v, \quad (2)$$

where $\mathbf{W}_v \in \mathbb{R}^{d \times 2048}$ and $\mathbf{b}_v \in \mathbb{R}^d$ are learnable parameters.

Text sequence representation. We follow the setup in [54] to extract feature vectors from the text sequences. Every word (token) is first embedded as a feature vector, and we apply a bi-directional Gated Recurrent Unit (Bi-GRU) [9, 52] to account for context. Specifically, let $\mathbf{S} = \{\mathbf{w}_1, \dots, \mathbf{w}_M\}$ be a text sequence, where M is the sequence length and \mathbf{w}_m denotes the p -dimensional word embedding vector for the m -th word in the sequence. Then the m -th feature vector \mathbf{e}_m is constructed by averaging the left and right context of the GRU embedding, *i.e.* $\mathbf{e}_m = (\overrightarrow{\mathbf{h}}_m + \overleftarrow{\mathbf{h}}_m)/2$, where $\overrightarrow{\mathbf{h}}_m = \overrightarrow{\text{GRU}}(\mathbf{w}_m)$, $\overleftarrow{\mathbf{h}}_m = \overleftarrow{\text{GRU}}(\mathbf{w}_m)$, here $\overrightarrow{\mathbf{h}}_m, \overleftarrow{\mathbf{h}}_m, \mathbf{e}_m \in \mathbb{R}^d$. Similar to the image features discussed in the last section, we collectively denote these text sequence features as $\mathbf{E} = \{\mathbf{e}_m\}_{m=1}^M$.

3 Cross-Domain Feature Alignment with OT

To motivate our model, we first review some of the favorable properties of Optimal Transport (OT) that appeal to CDA applications.

- *Sparsity.* It is well known that when solved exactly, OT yields a sparse solution of transportation plan π^* [6], which eliminates matching ambiguity and facilitates model interpretation [13].
- *Mass conservation.* The solution is self-normalized in the sense that π^* 's row-sum and column-sum match the desired marginals [10].
- *Efficient computation.* OT solutions can be readily approximated using iterative procedures known as Sinkhorn iterations, requiring only matrix-vector products [10, 60].

Contextualized in a CDA setup, we can regard image and text sequence embeddings as two discrete distributions supported on the same feature representation space. Solving an OT transport plan between the two naturally constitutes a matching scheme to relate cross-domain entities. Alternatively, this allows OT-matching to be viewed as an attention mechanism, as the model attends to the units with high transportation pairing. The OT distance can further serve as a proxy for assessing the global “relatedness” between the image and text sequence, *i.e.*, a summary of the degree to which the image and text are aligned, justifying its use as a principled regularizer to be incorporated into the training objective.

To evaluate the OT distance, we first choose a pairwise similarity between \mathbf{V} and \mathbf{E} using a cost function $C(\cdot, \cdot)$. In our setup, we choose cosine distance $\mathbf{C}_{km} = C(\mathbf{e}_k, \mathbf{v}_m) = 1 - \frac{\mathbf{e}_k^T \mathbf{v}_m}{\|\mathbf{e}_k\| \|\mathbf{v}_m\|}$ as our cost, so that (1) can be reformulated as:

$$\mathcal{L}_{\text{OT}}(\mathbf{V}, \mathbf{E}) = \min_{\mathbf{T}} \sum_{k=1}^K \sum_{m=1}^M \mathbf{T}_{km} \mathbf{C}_{km} \quad (3)$$

where $\sum_m \mathbf{T}_{km} = \mu_k$, $\sum_k \mathbf{T}_{km} = \nu_m$, $\forall k \in [1, K]$, $m \in [1, M]$. Here, $\mathbf{T} \in \mathbb{R}_+^{K \times M}$ is the transport matrix, d_k and d_m are the weight of \mathbf{v}_k and \mathbf{e}_m in a given image and text sequence, respectively. We assume the weight for different features to be uniform, *i.e.*, $\mu_k = \frac{1}{K}$, $\nu_m = \frac{1}{M}$. We leverage the inexact proximal point method optimal transport algorithm (IPOT) [56] to efficiently solve the linear program (3). More details including pseudo-code implementation are summarized in the Supplementary Material (Supp). Below we elaborate on the use of OT-based image-text cross domain alignment in three tasks.




Figure 2: Illustration of the proposed retrieval model. Image and text sequences features are represented as bag of feature vectors (in blue). Cosine similarity matrix \mathbf{S} is computed (in yellow). Two types of similarity measures are considered: (1) the traditional sum-max text-image aggregation S_{cos} and (2) optimal transport S_{OT} (in green circle). The final score S is obtained as the weighted sum of the two similarity scores.

Image-Text Matching. We start our discussion with image-text matching, a building block of cross-modal retrieval tasks required by many downstream applications. In image-text matching, a model searches for a matching image in an image library based on a text description, or searches for a matching caption in a caption library based on an image. Figure 2 presents a diagram of the proposed OT-based CDA. The feature vectors \mathbf{V} and \mathbf{E} are extracted from images and text sequences using the fR-CNN and Bi-GRU models, respectively. The similarity between an image and a text sequence is obtained in terms of two types of similarity measures, computed over all possible entities (regions and words) within an image and text sequence. Specifically, we consider *i*) a sum-max text-image aggregated cosine similarity, and *ii*) a weighted OT-based similarity that explicitly accounts for all similarities between pairs of entities, as detailed below.

Baseline similarity score: Following the practice of [47], we first derive a global similarity score as the baseline target to optimize for cross-domain alignment. More specifically, we begin by computing the pairwise similarities for the k -th region and the m -th token using *cosine similarity* in the feature space

$$s_{km} = \frac{\mathbf{v}_k^\top \mathbf{e}_m}{\|\mathbf{v}_k\| \|\mathbf{e}_m\|}, \quad k \in [1, K], \quad m \in [1, M]. \quad (4)$$

The *global similarity score* is built via the aggregation:

$$S_{\text{cos}}(\mathbf{V}, \mathbf{E}) = \sum_{m=1}^M \max_k(s_{km}). \quad (5)$$

This strategy is known as *sum-max text-image* aggregation and has been applied successfully in image-text matching tasks. Alternatively, we can use *sum-max image-text*, where $S'_{\text{cos}}(\mathbf{V}, \mathbf{E}) = \sum_{k=1}^K \max_m(s_{km})$. Our choice of *sum-max text-image* is based on the ablation study from [54] and [47], which showed empirical evidence that *sum-max text-image* works better in practice.

OT similarity score: In addition to the above *sum-max text-image* score, we present an OT construction of global similarity, which is the key regularizer in our framework. Specifically, we choose the cost matrix \mathbf{C} to be $C_{km} = 1 - s_{km}$. The OT-based similarity score can be defined as $S_{\text{OT}}(\mathbf{V}, \mathbf{E}) = -\mathcal{L}_{\text{OT}}(\mathbf{V}, \mathbf{E})$ using (3), and the transport plan \mathbf{T} naturally corresponds to the cross-domain alignment strategy.

Composed similarity score: We integrate both similarity scores discussed above using a simple linear combination

$$S(\mathbf{V}, \mathbf{E}) = S_{\cos}(\mathbf{V}, \mathbf{E}) + \lambda S_{\text{OT}}(\mathbf{V}, \mathbf{E}), \quad (6)$$

where λ is a hyper-parameter weighting the relative importance of the OT similarity score. Intuitively, the baseline $S_{\cos}(\mathbf{V}, \mathbf{E})$ provides an *unweighted* account of the aggregated agreement between regions and words of an image and text sequence, while the OT-based S_{OT} provides a *weighted* summary of how well every region in the image matches every word in the text sequence. Consequently, (6) accounts for both aggregated and weighted alignments to assess global similarity. From the perspective of attention models, S_{\cos} can be understood as a *hard* attention whereas S_{OT} can be perceived as a *soft* attention. The hyperparameter λ controls the smoothness and sparsity of the final matching plan, which gives the similarity score $S(\mathbf{V}, \mathbf{E})$.

Final training objective: To derive our final training loss, we consider the construction known as *triplet loss with hardest negatives*, originally proposed in [16, 59]. For each batch of B image and sentence pairs $\{\mathbf{V}_j, \mathbf{E}_j\}_{j=1}^B$, the total loss is given by

$$\mathcal{L} = \sum_{j=1}^B \left\{ \max \left[0, S(\mathbf{V}_j, \mathbf{E}_j^-) - S(\mathbf{V}_j, \mathbf{E}_j) + \eta \right] + \max \left[0, S(\mathbf{V}_j^-, \mathbf{E}_j) - S(\mathbf{V}_j, \mathbf{E}_j) + \eta \right] \right\}, \quad (7)$$

where the hardest negatives are given by $\mathbf{V}_j^- = \arg \max_{\mathbf{v} \in \mathbf{V}_j} S(\mathbf{v}, \mathbf{E}_j)$ and

$\mathbf{E}_j^- = \arg \max_{\mathbf{e} \in \mathbf{E}_j} S(\mathbf{V}_j, \mathbf{e})$, and $\{\setminus j\}$ denotes all indices except for j .

This means that once the score of the positive pair, *i.e.*, $S(\mathbf{V}_j, \mathbf{E}_j)$, is higher by η units over the score for the negative pair with the highest score in a batch, the hinge loss is zero. This training objective encourages separation in similarity score between paired data and unpaired data.

Table 1: Cross-domain matching results with Recall@ K (R@K). Upper panel: Flickr30K, lower panel: MSCOCO.

Method	Sentence Retrieval			Image Retrieval			Rsum
	R@1	R@5	R@10	R@1	R@5	R@10	
DVSA (R-CNN, AlexNet) [17]	22.2	48.2	61.4	15.2	37.7	50.5	235.2
HM-LSTM (R-CNN, AlexNet) [18]	38.1	—	76.5	27.7	—	68.8	—
2WayNet (VGG) [19]	49.8	67.5	—	36.0	55.6	—	—
SM-LSTM (VGG) [20]	42.5	71.9	81.5	30.2	60.4	72.3	358.8
VSE++ (ResNet) [14]	52.9	—	87.2	39.6	—	79.5	—
DPC (ResNet) [21]	55.6	81.9	89.5	39.1	69.2	80.9	416.2
DAN (ResNet) [22]	55.0	81.8	89.0	39.4	69.2	79.1	413.5
SCO (ResNet) [23]	55.5	82.0	89.3	41.1	70.5	80.1	418.5
SCAN (Faster R-CNN, ResNet) [24]	67.7	88.9	94.0	44.0	74.2	82.6	452.2
BFAN (Faster R-CNN, ResNet) [25]	65.5	89.4	—	47.9	77.6	—	—
PFAN (Faster R-CNN, ResNet) [26]	66	89.6	94.3	49.6	77	84.2	460.7
VSRN (Faster R-CNN, ResNet) [27]	65	89	93.1	49	76	84.4	456.5
Ours (Faster R-CNN, ResNet):							
cos + OT	69	91.8	95.9	50.4	77.6	85.5	470.2
Order-embeddings (VGG) [28]	23.3	—	84.7	31.7	—	74.6	—
VSE++ (ResNet) [14]	41.3	—	81.2	30.3	—	72.4	—
DPC (ResNet) [21]	41.2	70.5	81.1	25.3	53.4	66.4	337.9
GXN (ResNet) [29]	42.0	—	84.7	31.7	—	74.6	—
SCO (ResNet) [23]	42.8	72.3	83.0	33.1	62.9	75.5	369.6
SCAN (Faster R-CNN, ResNet) [24]	46.4	77.4	87.2	34.4	63.7	75.7	384.8
VSRN (Faster R-CNN, ResNet) [27]	48.6	78.9	87.7	37.8	68	77.1	398.1
Ours (Faster R-CNN, ResNet):							
cos + OT	49.9	81.4	89.8	37.8	66.7	78.1	403.6

Weakly supervised phrase localization. The phrase-localization task aims to learn relatedness between text phrases and image regions. Weakly supervised phrase localization



Figure 3: Examples of image-text retrieval results using OT regularization. First row shows text-to-image retrieval results. For each sentence, the top-3 matched images are listed from left to right. Right-bottom corner in each image indicates if this is a ground truth image. Image-to-text retrieval results are shown in the second row, where the top-5 sentences given an image query are provided. The mark at the end of each sentence denotes if this is a ground truth sentence. Throughout the text a green checks indicates ground-truth, while a red cross indicates otherwise.

guided by an image-sentence pair can serve as an evaluation of CDA methods, as the performance of phrase localization shows the model’s ability to capture vision-language interactions. Phrase localization seeks to learn a mapping model $f(k|m, \mathbf{I}, \mathbf{w})$ that evaluates the probability that the m -th token in text sequence \mathbf{w} references the k -th region in image \mathbf{I} . For our model, we define the mapping model as:

$$f(k|m, \mathbf{I}, \mathbf{w}) \propto \mathbf{T}_{km}. \quad (8)$$

For the baseline model, we use the cosine similarity matrix as the mapping model:

$$f(k|m, \mathbf{I}, \mathbf{w}) \propto s_{km}. \quad (9)$$

For each model, we first train with the image-sentence matching task, and then directly apply the model to the phrase localization task without further tuning.

4 Related Work

Optimal transport. Efforts have been made to use OT to find intra-domain similarities. In computer vision, the earth mover’s distance (EMD), also known as the OT distance, is used to match the distribution of the content between two images [60]. OT has also been applied successfully to NLP tasks such as document classification [62], sequence-to-sequence learning [2] and text generation [8]. In these works, OT has been applied to within-domain alignment, either for image regions or text sequences, capturing the intra-domain semantics. This paper constitutes the first work to use OT for cross-domain feature alignment, *e.g.* in image-text retrieval and weakly supervised phrase grounding tasks.

Image-text matching. Many works have investigated embedding image and text sequence features into a joint semantic space for image-text matching. The first attempt was made by [29], where the authors proposed to use CNNs to encode the images and LSTMs to encode text. The model was trained with a hinge-base triplet ranking loss, and it was improved by adding hardest negatives in the triplet ranking loss [47]. To consider the relationship between image regions and text sequences, [27] first computed the similarity matrix for all regions and word pairs via a dot product, and then calculated the similarity score with a sum or max aggregation function (denoted as *dot* model). Recently, SCAN [54] was proposed to use two-step stacked cross attention to measure similarities in image region and text pairs. Further works include VSRN[36], PFAN[60] and BFAN[69]. In our model, we

share the same motivation as SCAN, but we propose OT to obtain the optimal relevance correspondence between entities from the two domains. Recently, large-scale vision-language pre-training [9, 20, 65, 67, 44, 63, 64, 65] has provided more informative representations for image-text pairs, and has achieved state-of-the-art matching performance. The proposed OT is orthogonal to this, and we leave it as future research to combine OT with pre-trained models.

Weakly supervised phrase localization. Motivated by the large amount of annotation efforts for supervised approaches, some previous works [6, 12, 15, 27, 62] attempted to use matched image-text pairs as supervision to guide phrase localization training. In [12, 27] a local region-phrase similarity score was calculated, followed by aggregating the local scores to calculate the global image-text similarity score.

5 Experiments

Datasets. We evaluate our model on the Flickr30K [47] and MS-COCO [58] datasets. Flickr30K contains 31,000 images, and each photo has five human-annotated captions. We split the data following the same setup as [16, 27]. The dataset has 29,000 training images, 1,000 validation images and 1,000 test images. MS-COCO contains 123,287 images, and each image is annotated with 5 human-generated text descriptions. We use the split in [16], *i.e.*, the dataset is split into 113,287 training images, 5,000 validation images and 5,000 test images.



Figure 4: Visualization of the learned OT alignment. We show attended image regions with matched key words. The brightness reflects the alignment strength. The left-most figure is the original image. Each bounding box is the region with highest OT alignment score w.r.t the matched key word. Our model successfully identifies the correct pairing without seeing any ground-truth (*i.e.*, weak supervision) during training.

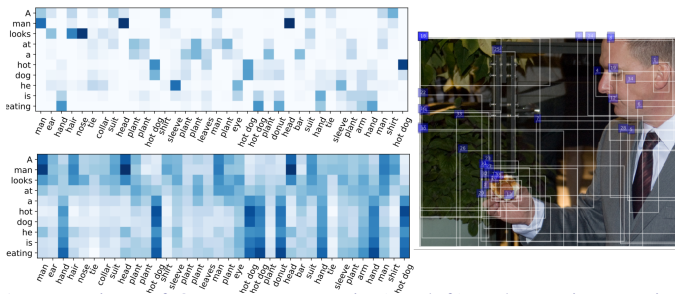


Figure 5: A comparison of OT transport matrix (top left) and attention matrix (bottom left). The horizontal axis represents image regions (annotated here to facilitate understanding), and the vertical axis represents words. Original image on the right.

Implementation details. For image-text matching, we use the Adam optimizer [25] to train the models. For the Flickr30K data, we train the model for 30 epochs. The initial learning rate is set to 0.0002, and decays by a factor of 10 after 15 epochs. For MS-COCO data, we train the model for 20 epochs. The initial learning rate is set to 0.0005, and decays by 10 after 10 epochs. The batch size to 128, and the maximum gradient norm is thresholded to 2.0 for gradient clipping. The dimension of the GRU and joint embedding space is set

to 1800, and the dimension of the word embedding to 300. Twenty iterations of the IPOT algorithm are considered. Since the performance of a single model is not reported in VSRN [36] paper, we ran associated experiments based on their github repository.

5.1 Image-text matching

We evaluate image-text matching on both datasets. The performance of sentence retrieval with image query or image retrieval with sentence query is measured by recall at K ($R@K$) [24], defined as the percentage of queries which retrieve the correct object within those with top K highest similarity scores as determined by the model. For each retrieval task, $K = \{1, 5, 10\}$ is recorded. We use Rsum [24] to evaluate the overall performance, defined as: $Rsum = \sum_K R@K_{I2T} + R@K_{T2I}$, where I2T denotes image-to-text retrieval and T2I denotes text-to-image retrieval.

Table 1 shows the quantitative results on Flickr30K and MS-COCO, with η representing the margin in (7) and λ the weight on the OT regularizer in (6). Hyper-parameters η and λ are determined with a grid search using the validation set, specifically, $\eta = 0.12$, $\lambda = 1.5$ for Flickr30K, and $\eta = 0.05$, $\lambda = 0.1$ for MS-COCO. We see that for a single model, our approach outperforms or is comparable with the current state-of-the-art method VSRN [36]. Similar results are observed under an ensemble setup (see Supp for detailed results).

5.2 Weakly supervised phrase localization

In order to demonstrate the efficiency of our CDA method under weak-supervision, we executed the weakly supervised phrase grounding experiment using pretrained retrieval models described in Section 3. Our implementation is based on Bilinear Attention Network code-base¹. We evaluate the models with the percentage of phrases that are correctly localized with respect to the ground truth bounding box across all images, where correct localization is defined as $IoU \geq 0.5$ [47]. Specifically, K predictions are permitted to find at least one correction, called Recall at K ($R@K$). Table 2 shows the comparison between our model and baseline SCAN model on Flickr30k [47]. When training the retrieval model, we choose the set of hyper-parameter that achieves best performance for both our model and the baseline SCAN model. In particular, OT_T represents the model described in Eq. (8), and OT_S represents the model described in Eq. (9) with image/text encoders trained by our model. Our approach outperforms the baseline model on all three metrics. This indicates that by leveraging OT, not only better alignment is computed, but also better feature encoders are trained.

Table 2: Phrase localization result on Flickr30K Entities

Method	R@1	R@5	R@10
SCAN	20.79	47.45	55.14
Dot	35.09	64.35	68.48
MATN [34]	33.10	-	-
KAC Net[9]	38.71	-	-
OT_T	35.98	70.33	78.97
OT_S	41.12	70.42	77.48

5.3 Qualitative results

We provide samples of image-text retrieval results from Flickr30K test set in Figure 3. For each sentence query, we present the top-3 images ranked by similarity score, as calculated by

¹<https://github.com/jnhwkim/ban-vqa>

Table 3: Ablation study on Flickr30K. We study the impact of hyper-parameters for OT and the baseline.

Method	Sentence Retrieval			Image Retrieval			Rsum
	R@1	R@5	R@10	R@1	R@5	R@10	
cos, $\eta=0.2$	61.7	87.4	93.5	48.5	76.0	83.7	450.8
cos + OT, $\eta=0.2, \lambda=1$	66.2	89.0	94.1	48.9	77.5	85.4	461.1
cos, $\eta=0.12$	63.1	89.5	94.3	50.5	77.1	84.7	459.2
cos + OT, $\eta=0.12, \lambda=2$	69.3	91.0	95.7	48.4	77.2	84.7	466.3

our model. For each image query, we present the top-5 sentences. From this representative sample we see that our model matches images and sentences with high correlation. Although the query text and retrieved images (and query image and retrieved text) are not the exact pairs, they are still highly correlated and share the same theme. More qualitative results for image-text retrieval, image captioning and VQA are presented in Supp.

5.4 Analysis

Ablation study. We consider several ablation settings to further examine the capabilities of the proposed OT algorithm. To show the effectiveness of OT, we consider an ablation experiment for the image-text retrieval task. In Table 3, we compare our model with the baseline, which only uses cosine similarity to measure the distance between image and text features, *i.e.*, only (5) is applied. Two hyper-parameter combinations are considered. In both cases, the OT-enhanced similarity outperforms the baseline model, demonstrating the effectiveness of optimal transport. The ablation study on network architecture choices and adaptive region numbers are found in the Supp.

Interpretable alignment. One favorable property of OT is the interpretability of the optimal transport plan \mathbf{T} . To illustrate this, we visualize \mathbf{T} in comparison with the attention matrix $(1 - \mathbf{C})$ in Figure 5. The darker shade implies stronger OT matching or attention weights. We see that OT transport mapping is more interpretable, as the alignment is sparse and self-normalized. See Supp for more examples.

6 Conclusions

We have proposed to use optimal transport to provide a principled alignment between features from text-image domains. We take advantage of such alignment when computing similarity scores for image and text entities in matching tasks, and the results outperform the state of the art. Moreover, we show the accuracy of OT-based alignment with phrase localization and achieve better performance than baseline models. As future work, it is of interest to take advantage of OT alignment in other text-image cross-domain tasks, such as visual question answering and text-to-image generation.

7 Acknowledgements

The authors would like to thank the anonymous reviewers for their insightful comments. The research at Duke University was supported in part by DARPA, DOE, NIH, NSF and ONR.

References

- [1] Peter Anderson et al. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, pages 6077–6086, 2018.
- [2] Stanislaw Antol et al. Vqa: Visual question answering. In *ICCV*, pages 2425–2433, 2015.
- [3] Martin Arjovsky et al. Wasserstein generative adversarial networks. In *ICML*, 2017. URL <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- [5] Richard A Brualdi and Herbert J Ryser. *Combinatorial matrix theory*, volume 39. 1991.
- [6] Kan Chen, Jiyang Gao, and Ram Nevatia. Knowledge aided consistency for weakly supervised phrase grounding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4042–4050, 2018.
- [7] Liqun Chen, Yizhe Zhang, et al. Improving sequence-to-sequence learning via optimal transport. In *ICLR*, 2019.
- [8] Liqun Chen et al. Adversarial text generation via feature-mover’s distance. In *NeurIPS*, 2018.
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740*, 2019.
- [10] M Cuturi and G Peyré. Computational optimal transport. 2017.
- [11] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, pages 2292–2300, 2013.
- [12] Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment. In *ICCV*, 03 2019.
- [13] Fernando De Goes et al. An optimal transport approach to robust reconstruction and simplification of 2d shapes. In *Computer Graphics Forum*, volume 30, 2011.
- [14] Aviv Eisenschat and Lior Wolf. Linking image and text with 2-way nets. In *CVPR*, 2017.
- [15] Martin Engilberge, Louis Chevallier, Patrick Pérez, and Matthieu Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *CVPR*, June 2018.
- [16] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improved visual-semantic embeddings. In *BMVC*, volume 2, page 8, 2018.
- [17] Hao Fang, Saurabh Gupta, et al. From captions to visual concepts and back. In *CVPR*, pages 1473–1482, 2015.

- [18] Jiuxiang Gu, Jianfei Cai, Shafiq R Joty, Li Niu, and Gang Wang. Look, imagine and match: Improving textual-visual cross-modal retrieval with generative models. In *CVPR*, 2018.
- [19] Ishaan Gulrajani et al. Improved training of Wasserstein GANs. In *NeurIPS*, 2017.
- [20] Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. In *CVPR*, 2020.
- [21] David Harwath, Adria Recasens, Dídac Surís, Galen Chuang, Antonio Torralba, and James Glass. Jointly discovering visual objects and spoken words from raw sensory input. In *ECCV*, 2018.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [23] Sepp Hochreiter et al. Long short-term memory. *Neural computation*, 1997.
- [24] Yan Huang, Wei Wang, and Liang Wang. Instance-aware image and sentence matching with selective multimodal lstm. In *CVPR*, pages 2310–2318, 2017.
- [25] Yan Huang, Qi Wu, Chunfeng Song, and Liang Wang. Learning semantic concepts and order for image and sentence matching. In *CVPR*, pages 6163–6171, 2018.
- [26] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *CVPR*, pages 4565–4574, 2016.
- [27] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, pages 3128–3137, 2015.
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [29] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. In *NeurIPS*, 2014.
- [30] Ranjay Krishna et al. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *IJCV*, 2017.
- [31] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 1955.
- [32] Matt Kusner et al. From word embeddings to document distances. In *ICML*, 2015.
- [33] Yann LeCun et al. Object recognition with gradient-based learning. In *Shape, contour and grouping in computer vision*. 1999.
- [34] Kuang-Huei Lee et al. Stacked cross attention for image-text matching. In *ECCV*, 2018.
- [35] Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. *arXiv preprint arXiv:1908.06066*, 2019.

- [36] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4654–4662, 2019.
- [37] Xijun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. *ECCV*, 2020.
- [38] Tsung-Yi Lin et al. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [39] Chunxiao Liu, Zhendong Mao, An-An Liu, Tianzhu Zhang, Bin Wang, and Yongdong Zhang. Focus your attention: A bidirectional focal attention network for image-text matching. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 3–11, 2019.
- [40] Yishu Liu et al. Scene classification using hierarchical wasserstein cnn. *IEEE Transactions on Geoscience and Remote Sensing*, 2018.
- [41] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- [42] Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *NeurIPS*, pages 1682–1690, 2014.
- [43] Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *CVPR*, pages 299–307, 2017.
- [44] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 3, pages 850–855. IEEE, 2006.
- [45] Zhenxing Niu et al. Hierarchical multimodal lstm for dense visual-semantic embedding. In *ICCV*, 2017.
- [46] Gabriel Peyré and Marco Cuturi. Computational optimal transport. Technical report, 2017.
- [47] Bryan A Plummer et al. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, pages 2641–2649, 2015.
- [48] Tingting Qiao et al. Mirrorgan: Learning text-to-image generation by redescription. *CVPR*, 2019.
- [49] Scott Reed et al. Generative adversarial text to image synthesis. *ICML*, 2016.
- [50] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015.
- [51] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.

- [52] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [53] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.
- [54] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A joint model for video and language representation learning. *ICCV*, 2019.
- [55] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. *EMNLP*, 2019.
- [56] Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. Order-embeddings of images and language. In *ICLR*, 2016.
- [57] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [58] Oriol Vinyals et al. Show and tell: A neural image caption generator. In *CVPR*, pages 3156–3164, 2015.
- [59] Liwei Wang et al. Learning deep structure-preserving image-text embeddings. In *CVPR*, 2016.
- [60] Yaxiong Wang, Hao Yang, Xueming Qian, Lin Ma, Jing Lu, Biao Li, and Xin Fan. Position focused attention network for image-text matching. *arXiv preprint arXiv:1907.09748*, 2019.
- [61] Yujia Xie, Xiangfeng Wang, Ruijia Wang, and Hongyuan Zha. A fast proximal point method for computing exact wasserstein distance. *arXiv preprint arXiv:1802.04307*, 2018.
- [62] Kelvin Xu et al. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015.
- [63] Dongfei Yu et al. Multi-level attention networks for visual question answering. In *CVPR*, 2017.
- [64] Fang Zhao, Jianshu Li, Jian Zhao, and Jiashi Feng. Weakly supervised phrase localization with multi-scale anchored transformer network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5696–5705, 2018.
- [65] Zhedong Zheng et al. Dual-path convolutional image-text embedding with instance loss. *arXiv*, 2017.