

Graph Matching with Bi-level Noisy Correspondence

Yijie Lin¹, Mouxing Yang¹, Jun Yu², Peng Hu¹, Changqing Zhang^{3*}, Xi Peng^{1*}

¹ Sichuan University, ² Hangzhou Dianzi University, ³ Tianjin University

{linyijie.gm, yangmouxing, penghu.ml, pengx.gm}@gmail.com,

yu jun@hdu.edu.cn, zhangchangqing@tjtu.edu.cn

Abstract

In this paper, we study a novel and widely existing problem in graph matching (GM), namely, Bi-level Noisy Correspondence (BNC), which refers to node-level noisy correspondence (NNC) and edge-level noisy correspondence (ENC). In brief, on the one hand, due to the poor recognizability and viewpoint differences between images, it is inevitable to inaccurately annotate some keypoints with offset and confusion, leading to the mismatch between two associated nodes, *i.e.*, NNC. On the other hand, the noisy node-to-node correspondence will further contaminate the edge-to-edge correspondence, thus leading to ENC. For the BNC challenge, we propose a novel method termed Contrastive Matching with Momentum Distillation. Specifically, the proposed method is with a robust quadratic contrastive loss which enjoys the following merits: i) better exploring the node-to-node and edge-to-edge correlations through a GM customized quadratic contrastive learning paradigm; ii) adaptively penalizing the noisy assignments based on the confidence estimated by the momentum teacher. Extensive experiments on three real-world datasets show the robustness of our model compared with 12 competitive baselines. The code is available at <https://github.com/XLearning-SCU/2023-ICCV-COMMON>.

1. Introduction

Graph Matching (GM) [9, 60] aims to establish correspondences between keypoints of different graphs, which plays an important role in various applications such as object tracking [56, 45], scene graph discovery [5], Simultaneous Localization and Mapping [4], and Structure-from-Motion [41]. The key of GM is to explore and exploit the bi-level affinities between graphs, *i.e.*, node-to-node (*linear*) and edge-to-edge (*quadratic*) affinity. For this purpose, existing methods have been devoted to integrating the bi-level information by designing advanced graph neu-

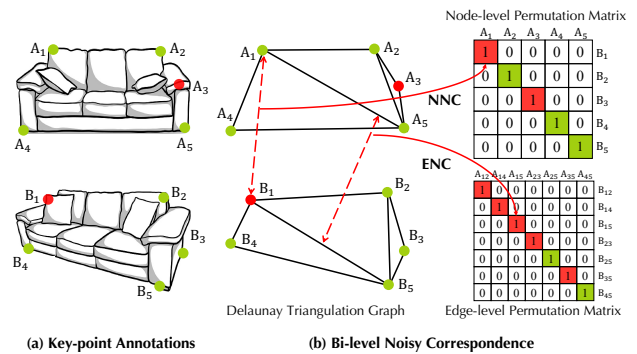


Figure 1. **An illustrative example of Bi-level Noisy Correspondence (BNC).** The green and red dots denote the correctly and wrongly annotated keypoints. The green and red squares denote the true and false assignments. (a) Keypoint A_3 and B_1 are wrongly annotated due to the occlusion caused by the view-point difference, and the low recognizability, respectively. (b) Given the keypoints, GM methods construct the graph structure for further matching. The matching procedure will inevitably encounter BNC, which refers to Node-level Noisy Correspondence (NNC) and Edge-level Noisy Correspondence (ENC). Specifically, NNC denotes the false matching between nodes, *e.g.*, nodes A_1 and B_1 are wrongly matched. ENC is accompanied with NNC as the edge weight is derived from the feature and position of nodes. Once the node is imperfectly annotated, the edges derived from it and the edges derived from its counterpart will be wrongly associated. For instance, the edge between A_1 and A_5 (marked as A_{15}) is wrongly associated with that between B_1 and B_5 (marked as B_{15}).

ral networks [47, 48, 41, 58] or differentiable quadratic loss [13, 40]. Based on the encoded high-order geometrical information, graph matching achieves promising results in correspondence estimation.

However, the success of existing GM methods highly depends on an implicit assumption, *i.e.*, the annotated keypoint pairs are correctly associated. Unfortunately, in practice, the manual annotation process is susceptible to poor recognizability [3] and viewpoint differences between images [34], which probably results in offset and even false keypoint annotations (Please refer to Fig. 4 for examples).

*Corresponding author

As shown in Fig. 1, the inaccurate annotations will inevitably lead to **Bi-level Noisy Correspondence** (BNC) problem. Specifically, BNC refers to node-level noisy correspondence (NNC) and accompanied edge-level noisy correspondence (ENC). As shown in Fig. 3(a), BNC leads to serious performance degradation due to two potential issues. On the one hand, BNC would hinder both node-level and edge-level representation learning. Specifically, representation learning on graphs requires information propagation and aggregation between the nodes and edges, therefore, BNC will cause accumulative errors and mislead the model optimization. On the other hand, GM is subject to one-to-one mapping constraints, *i.e.*, each keypoint in the first graph must have a unique correspondence in the second graph. Hence, “*a slight move in one part may affect the situation as a whole*”, *i.e.*, one single noisy correspondence pair might result in global alignment failure and even affect the optimization of correctly annotated pairs.

To the best of our knowledge, such an essential problem has not been touched so far and it is quite challenging to solve it due to the following two reasons: i) without prior engineering like verification labels [24] for noisy correspondence, it is nearly impossible to pre-process them before training, *i.e.*, correctly distinguish and discard the noisy correspondence in advance. Hence, it is highly expected to develop a robust matching method. ii) however, GM is not a simple one-to-many optimization problem (*e.g.*, classification) but a many-to-many combinatorial optimization problem. One cannot solve the BNC challenge by trivially resorting to label noise learning methods [19, 26, 55] as they only enjoy robustness on one-to-many classification tasks.

To explore an effective solution to this challenge, we propose a novel method termed **CONtrastive Matching with MOMentum distillatiON** (COMMON). Specifically, COMMON is equipped with a robust quadratic contrastive loss that incorporates both linear and quadratic geometrical information into the contrastive learning paradigm [6, 16, 50]. More specifically, this study is motivated by [35] that the vanilla contrastive loss is equivalent to the linear assignment from the combinatorial optimization theory. In other words, the quadratic structure information could be incorporated into contrastive learning. Based on this motivation, we endow the contrastive loss with quadratic information through two novel graph-geometric consistency regularizers. To enhance the robustness against BNC, the proposed loss adaptively penalizes the noisy correspondence through a momentum distillation strategy based on the *memorization effect* of deep neural networks [1, 53, 18], *i.e.*, deep networks are apt to learn the simple patterns before fitting the noise. Motivated by this empirical observation, we keep a momentum version [25, 16] of the GM model by taking the moving average of its parameters during training. The momentum teacher model could generate high-quality pseudo-

targets as additional supervision on the quadratic contrastive loss without resorting to extra verification labels. With bi-level distillation on both node and edge alignment, the proposed robust quadratic contrastive loss effectively mitigates the negative affect of BNC.

The main contributions of this work are:

- We reveal a new challenge for graph matching, termed bi-level noisy correspondence (BNC). BNC refers to noisy correspondence on both node and edge levels, which would lead to serious performance degradation.
- To tackle the BNC challenge, we propose a robust quadratic contrastive learning loss that utilizes the momentum distillation strategy to rebalance the noisy assignment. The proposed contrastive loss explores the quadratic information through two simple graph consistency regularizers.
- Extensive experiments on real-world data verify the effectiveness of our method against noisy correspondence. On Pascal VOC, Spair-71k and Willow, we achieve absolute improvements of 1.6%, 1.4%, and 1.4% compared to the state-of-the-art methods.

2. Related work

In this section, we briefly introduce some recent developments in deep graph matching and contrastive learning.

2.1. Deep Graph Matching

Deep GM [60, 11] aims at aligning the associated keypoints from different graphs according to node-to-node and edge-to-edge correlations. To achieve better matching effects, existing methods mainly focus on utilizing the high-order information in the graph structures. According to the choice of exploiting high-order information, most existing methods could be divided into the following two groups. i) network-designed based methods [51, 58, 31, 20] which implicitly aggregate the high-order information through GM-customized networks. For instance, PCA [47] employs graph convolutional networks to aggregate inner-graph and cross-graph structure information, and NGM [49] proposes a matching-aware graph convolution scheme with Sinkhorn iteration [10]. ii) loss-designed based methods [33, 13, 40] that explicitly learn the high-order information through different differentiable quadratic loss or optimization strategies. For instance, QCDGM [13] modifies the Frank-Wolfe algorithm into a differentiable optimization scheme for quadratic constraint, and BBGM [40] proposes a differentiable combinatorial solver for quadratic assignment.

Although deep GM has achieved promising performance, almost all existing methods assume that the node-to-node and edge-to-edge correspondence is faultless and correctly associated. However, in practice, the assumption

is daunting and nearly impossible to satisfy as aforementioned in Introduction, *i.e.*, BNC is inevitable due to poor annotations. Note that, although some efforts have been devoted to achieving robust GM, these methods mainly focus on the robustness against outliers [46, 41, 37, 40] and adversarial attack [39], which is remarkably different from the revealed BNC challenge. To achieve robustness against BNC, this study employs a momentum distillation strategy that rebalances the matching loss in a data-driven way. To the best of our knowledge, this study could be the first work on BNC-robust graph matching.

2.2. Contrastive Learning

As one of the most effective representation learning paradigms, contrastive learning [28, 52, 32, 54, 27, 15, 14] attempts to align representations of the data from different views [29, 57, 21, 36] by maximizing the similarity between positive pairs and minimizing that of negative pairs. The major difference between these works lies in the strategy of constructing positive and negative pairs. For example, SimCLR [6] constructs positive samples through data augmentation and takes different mini-batch images as negatives. Moco [16] proposes a memory bank module that stores a large number of negatives to increase the contrasting effect.

The differences between this study and existing works are two-fold. On the one hand, most existing methods assume that positive pairs are closely associated, which is hard to satisfy in practical applications. In contrast, our momentum distillation solution improves the robustness of contrastive learning against imperfect correspondence. On the other hand, as pointed by [35], almost all contrastive methods only consider instance discrimination problem, *i.e.*, measuring the alignment between individual pairs of objects while ignoring the high-order correlation among these objects. To tackle this problem, we propose a novel quadratic contrastive learning loss that not only aligns the objects (node) but also aligns the correlation between objects (edge). Different from the quadratic loss proposed in [35], our quadratic term is implemented with the stable mean square error instead of singular value decomposition. In addition, our quadratic term incorporates the extra cross-graph quadratic information, whereas [35] does not. What is more distinct, our loss is designed to achieve the robustness against the BNC problem. Experiments demonstrate that the proposed quadratic loss improves the performance of contrastive learning.

3. Method

In this section, we first introduce the problem definition (Section 3.1). Then we delineate how the proposed quadratic contrastive learning loss better explores the node-to-node and edge-to-edge correlations (Section 3.2), and

how the designed loss adaptively penalizes the noisy correspondence through momentum distillation (Section 3.3).

3.1. Problem Definition

For two images with n and m keypoints each ($n \leq m$), graph matching aims to build the node-to-node correspondence between two given graphs \mathcal{G}_A and \mathcal{G}_B built based on these keypoints. Let $\mathbf{V}_A \in \mathbb{R}^{n \times d}$ and $\mathbf{V}_B \in \mathbb{R}^{m \times d}$ denote feature matrix of keypoints in \mathcal{G}_A and \mathcal{G}_B , respectively, where each row of \mathbf{V}_A and \mathbf{V}_B is a feature vector of a keypoint. $\mathbf{F}_A = \mathbf{V}_A \mathbf{V}_A^\top \in \mathbb{R}^{n \times n}$, $\mathbf{F}_B = \mathbf{V}_B \mathbf{V}_B^\top \in \mathbb{R}^{m \times m}$ indicate the adjacency matrices which encode the edge information in graphs \mathcal{G}_A and \mathcal{G}_B . Formally, graph matching seeks to solve,

$$\min L_Y(\mathbf{Y}_{gt}, \mathbf{Y}), \quad (1)$$

where L_Y measures the discrepancy between the ground-truth assignment \mathbf{Y}_{gt} and the matching result \mathbf{Y} , *e.g.*, cross-entropy loss [49, 39] and hamming distance loss [40, 30]. The matching result \mathbf{Y} is generally obtained by,

$$\begin{aligned} \operatorname{argmax}_{\mathbf{Y}} \quad & \underbrace{\operatorname{tr}(\mathbf{Y}^\top \mathbf{F}_A \mathbf{Y} \mathbf{F}_B)}_{\text{Quadratic edge affinity}} + \underbrace{\operatorname{tr}(\mathbf{S}^\top \mathbf{Y})}_{\text{Linear node affinity}} \\ \text{s.t.} \quad & \mathbf{Y} \in \{0, 1\}^{n \times m}, \mathbf{Y} \mathbf{1}_m = \mathbf{1}_n, \mathbf{Y}^\top \mathbf{1}_n \leq \mathbf{1}_m \end{aligned} \quad (2)$$

where $\mathbf{S} = \mathbf{V}_A \mathbf{V}_B^\top \in \mathbb{R}^{n \times m}$ denotes node-to-node similarity matrix. By combining Eqs. (1) and (2), it shows that GM encourages to learn better node and edge features that leads to correct assignment.

However, due to poor recognizability and viewpoint difference of images, it is extremely hard and time-consuming to precisely annotate keypoints, leading to **incorrect** assignment \mathbf{Y}_{gt} . As shown in Fig. 1 and Eq. (2), the incorrect \mathbf{Y}_{gt} induces false matching on both node and edge levels, *i.e.*, bi-level noisy correspondence. Accordingly, the BNC problem tends to cause false optimization on both node and edge levels, leading to serious performance degradation. As aforementioned, it is hard to preprocess the noisy correspondence before training. Therefore, our method attempts to explore a robust training strategy which could effectively mitigate the negative impact of BNC.

3.2. Quadratic Contrastive Learning for Graph Matching

In this section, we propose the quadratic contrastive learning loss for graph matching. Specifically, based on the combinatorial optimization theory [35] that contrastive learning is equivalent to linear assignment problem, we accordingly introduce two graph consistency regularizers that endow contrastive learning with quadratic information.

As pointed by [35], the linear assignment problem (the second term in Eq. (2)) could be formulated by structured

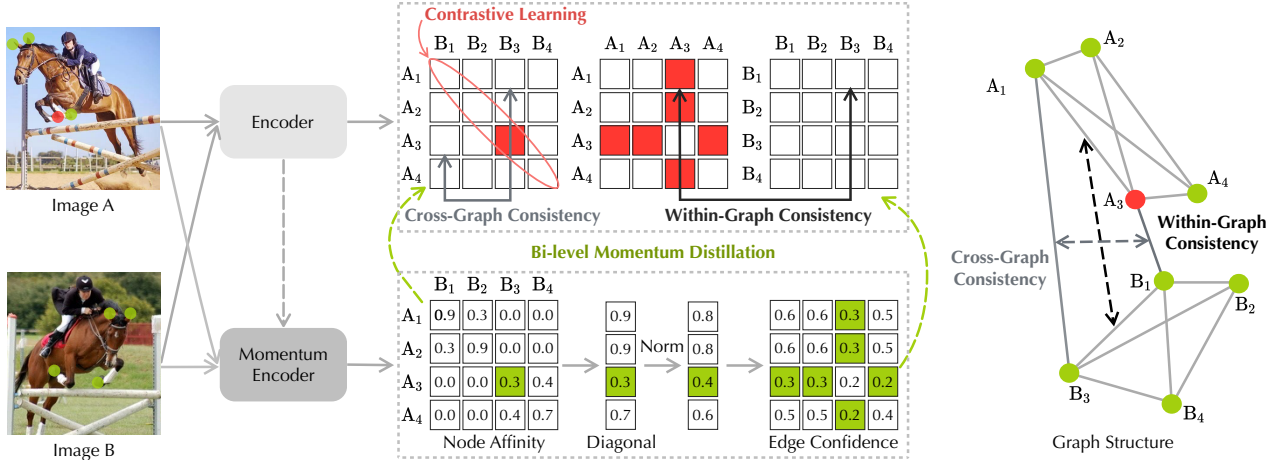


Figure 2. **Overview of COMMON.** In the figure, the green and red dots denote the correctly and wrongly annotated keypoints. The red and green squares indicate the noisy correspondence and corresponding estimated confidence. By feeding a pair of images into the encoder, we obtain the node similarity matrix and two edge adjacency matrices. To better explore the bi-level correlations, we propose a quadratic contrastive loss which consists of three jointly learning objectives, *i.e.*, contrastive learning loss, cross-graph consistency loss, and within-graph consistency loss. To mitigate the influence of BNC, we further feed the input into the momentum encoder that adaptively rebalances the noisy correspondence based on the estimated node and edge confidences.

linear loss [44]:

$$L = \max_{\mathbf{Y} \in \Pi} \text{tr}(\mathbf{S}\mathbf{Y}^\top) - \text{tr}(\mathbf{S}\mathbf{Y}_{gt}^\top), \quad (3)$$

where Π denotes a set containing all $n \times m$ permutation matrices that follows the constraint in Eq. (2). Note that $L \geq 0$ and $L = 0$ i.f.f. the node similarities produced by the encoder lead to the correct assignment. Through minimizing the objective in Eq. (3), the encoder expected to correctly assign keypoints from one image to another.

To make the structure linear loss calculable, one could relax the constraint Π with binary row-stochastic relaxation ($[\mathbf{Y}]_{ij} \in \{0, 1\}$ and $\sum_j \mathbf{Y}_{ij} = 1 \forall i$) and utilize the common log-sum-exp approximation [2] on the max function of Eq. (3). Then the linear assignment loss could be reformulated as the InfoNCE contrastive loss as proved by [35]:

$$L = - \sum_{(i,j) \in \mathbf{Y}_{gt}} [\mathbf{S}]_{ij} + \tau \sum_i \log \left(\sum_j \exp \left(\frac{1}{\tau} [\mathbf{S}]_{ij} \right) \right), \quad (4)$$

where τ controls the degree of smoothness. For our case, we reformulate the equivalence theory between the linear assignment and contrastive learning as follows. The detailed proof please refer to the supplementary material.

Theorem 1 *The log-sum-exp [2] smoothed structured linear assignment loss L with row-stochastic relaxation is equivalent to the InfoNCE contrastive loss [16, 7].*

Notably, the above theorem indicates that contrastive learning Eq. (4) is a fast and differentiable approximation to the linear assignment problem thus could be efficiently

used to align the keypoints in deep graph matching directly. To simply conduct contrastive learning for graph matching, we align the keypoints \mathbf{V}_A and \mathbf{V}_B according to \mathbf{Y}_{gt} and only remain the keypoints which have the counterpart for training, obtaining aligned keypoints $\mathbf{P}_A, \mathbf{P}_B \in \mathbb{R}^{n \times d}$ for graph \mathcal{G}_A and \mathcal{G}_B , respectively. Then we conduct contrastive learning on both rows and columns of the node similarity matrix \mathbf{S} as [38]:

$$L_{\text{InfoNCE}} = \mathcal{H}(\mathbf{I}_n, \rho(\mathbf{P}_A \mathbf{P}_B^\top)) + \mathcal{H}(\mathbf{I}_n, \rho(\mathbf{P}_B \mathbf{P}_A^\top)), \quad (5)$$

where \mathbf{I}_n is the identity matrix, \mathcal{H} is the row-wise cross-entropy function with mean reduction and ρ is the softmax function applied row-wise such that each row sums to one, *i.e.*,

$$[\rho(\mathbf{P}_A \mathbf{P}_B^\top)]_{ij} = \frac{\exp([\mathbf{P}_A]_i [\mathbf{P}_B]_j^\top / \tau)}{\sum_{k=1}^n \exp([\mathbf{P}_A]_i [\mathbf{P}_B]_k^\top / \tau)}, \quad (6)$$

where $[\mathbf{P}_A]_i \in \mathbb{R}^{1 \times d}$ is the i -row of \mathbf{P}_A , and τ is the softmax temperature fixed as 0.07.

Although the popular InfoNCE loss might be an effective solution to the linear assignment problem, it ignores another essential perspective in graph matching, *i.e.*, edge alignment. In fact, it is generally accepted that considering the edge information in graphs makes the matching more robust [47]. Therefore, contrastive learning can not sufficiently utilize the graph structure and might obtain sub-optimal results. Hence, we augment the contrastive loss with two novel graph-geometric consistency regularizers, namely, within-graph consistency and cross-graph consistency, to further exploit the edge correlation.

Within-graph consistency aims to encourage the alignment between edges within each image,

$$L_{\text{within}} = \|\mathbf{P}_A \mathbf{P}_A^\top - \mathbf{P}_B \mathbf{P}_B^\top\|_F^2, \quad (7)$$

where $\|\cdot\|_F$ is the Frobenius norm. The within-graph consistency is a direct quadratic assignment regularizer.

Cross-graph consistency encourages the alignment of the edges across two images:

$$L_{\text{cross}} = \|\mathbf{P}_A \mathbf{P}_B^\top - \mathbf{P}_B \mathbf{P}_A^\top\|_F^2. \quad (8)$$

For clarity, we provide a typical example in Fig. 2 to illustrate the above formulation. Intuitively, Eq. (7) minimizes the difference between the edge in image A and its counterpart in image B, *e.g.*, A_1A_3 and B_1B_3 . Apart from the alignment between within-graph edges, it is also encouraged to align the cross-graph edges. Specifically, to better align the keypoints between graphs, the semantic difference between objects needs to be eliminated, *e.g.*, the difference between two kinds of horses. Hence, given two associated within-graph edges (*e.g.*, A_1A_3 and B_1B_3), it is highly expected that the corresponding keypoints are equivalent and also exchangeable (*e.g.*, A_3 and B_3).

Therefore, we establish the cross-graph edges from within-graph edges by exchanging the counterpart of each keypoint and then minimize the difference between them (*e.g.*, A_1B_3 and B_1A_3). Mathematically, Eq. (8) makes the matrix \mathbf{S} symmetric to encourage the alignment between cross-graph edges. Although [47, 41] propose the cross-graph neural network to encode the information from both graphs, they only propagate the information from one node in image A (B) to other nodes in image B (A). In contrast, our cross-graph consistency offers a more explicit quadratic supervision to ensure the cross-graph semantic consistency.

Notably, the above two geometric consistency terms are elegantly implemented as simple regularizers on contrastive learning and facilitate the matching procedure. Accordingly, the quadratic contrastive loss for graph matching is given as:

$$L_{\text{quadratic}} = L_{\text{InfoNCE}} + L_{\text{within}} + L_{\text{cross}}. \quad (9)$$

3.3. Momentum Distillation for Robust Matching

As aforementioned, bi-level noisy correspondence would cause failure on both node-to-node and edge-to-edge matching, resulting in serious performance degradation. Different from the traditional noisy label problem, BNC is a many-to-many combinatorial optimization problem rather than a one-to-many classification problem. It is inapplicable to distinguish [24, 22] or rectify [26] the noisy correspondence through existing noisy label classification methods.

Fortunately, Bengio *et al.* [1] have empirically found the memorization effect of the deep neural networks, *i.e.*, networks are apt to fit the simple patterns first. To be specific, precisely annotated correspondence could be regarded as simple patterns while noisy correspondence could be treated as complex ones. Motivated by this effect, we propose to learn from the high-quality pseudo-target generated by the momentum encoder [16, 25] as shown in Fig 2. The momentum encoder is a continuously-evolving teacher, which will keep the memory at the simple patterns to some extent by taking the exponential-moving-average of the parameters from the base encoder. Formally, denoting the parameters of the base encoder as θ_q and those of momentum encoder as θ_k , we update θ_k by:

$$\theta_k \leftarrow t\theta_k + (1-t)\theta_q, \quad (10)$$

where t is the momentum coefficient fixed as 0.995 in our experiments. During training, we utilize the matching prediction of the momentum encoder to rebalance the proposed quadratic contrastive learning loss. Specifically, we propose a bi-level distillation on both node and edge levels with respect to contrastive loss and graph consistency loss.

Distillation on node level. We first compute the node features $\hat{\mathbf{P}}_A$ and $\hat{\mathbf{P}}_B$ from the momentum encoder. Then we introduce a new cross-entropy term into Eq. (5) to match the alignment results between base encoder and the soft target of its momentum teacher [25]. Formally,

$$L_{\text{RInfoNCE}} = (1-\alpha) (\mathcal{H}(\mathbf{I}_n, \rho(\mathbf{P}_A \mathbf{P}_B^\top)) + \mathcal{H}(\mathbf{I}_n, \rho(\mathbf{P}_B \mathbf{P}_A^\top))) + \alpha (\mathcal{H}(\rho(\hat{\mathbf{P}}_A \hat{\mathbf{P}}_B^\top), \rho(\mathbf{P}_A \mathbf{P}_B^\top)) + \mathcal{H}(\rho(\hat{\mathbf{P}}_B \hat{\mathbf{P}}_A^\top), \rho(\mathbf{P}_B \mathbf{P}_A^\top))), \quad (11)$$

where α is the distillation temperature. The second term in Eq. (11) allows the teacher to re-calibrate the keypoint alignment by replacing the impractical perfect alignment target \mathbf{I}_n in Eq. (5) with estimated soft-alignment probability. To keep the stability in the preliminary training stage, α is linearly increasing from 0 to 0.4 in the first epoch and fixed to 0.4 afterward. As shown in Table 6, our method is insensitive to the choice of α .

Distillation on edge level. As the edge-level noisy correspondence is accompanied with node-level noisy correspondence, we first estimate the confidence of each node and then estimate that of the edge based on the confidence of its two vertices. Formally, let $\mathbf{s} \in \mathbb{R}^n$ denote the confidence of nodes where

$$s_i = \frac{1}{2} \left(\frac{[\hat{\mathbf{P}}_A \hat{\mathbf{P}}_B^\top]_{ii}}{\sum_j [\hat{\mathbf{P}}_A \hat{\mathbf{P}}_B^\top]_{ij}} + \frac{[\hat{\mathbf{P}}_B \hat{\mathbf{P}}_A^\top]_{ii}}{\sum_j [\hat{\mathbf{P}}_B \hat{\mathbf{P}}_A^\top]_{ji}} \right). \quad (12)$$

The confidence of each edge is estimated by accumulating the confidence of its two vertexes, *i.e.*, $\mathbf{W} = \sqrt{\mathbf{s}} \otimes \mathbf{s}$ where \otimes is the outer product. Based on the edge confidence \mathbf{W} , we re-weight the graph consistency loss (Eqs. (7) and (8))

Method	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	Mbike	Person	Plant	Sheep	Sofa	Train	Tv	Mean
GMN [60]	41.6	59.6	60.3	48.0	79.2	70.2	67.4	64.9	39.2	61.3	66.9	59.8	61.1	59.8	37.2	78.2	68.0	49.9	84.2	91.4	62.4
PCA [47]	49.8	61.9	65.3	57.2	78.8	75.6	64.7	69.7	41.6	63.4	50.7	67.1	66.7	61.6	44.5	81.2	67.8	59.2	78.5	90.4	64.8
NGM [49]	50.1	63.5	57.9	53.4	79.8	77.1	73.6	68.2	41.1	66.4	40.8	60.3	61.9	63.5	45.6	77.1	69.3	65.5	79.2	88.2	64.1
IPCA [48]	53.8	66.2	67.1	61.2	80.4	75.3	72.6	72.5	44.6	65.2	54.3	67.2	67.9	64.2	47.9	84.4	70.8	64.0	83.8	90.8	67.7
LCS [51]	46.9	58.0	63.6	69.9	87.8	79.8	71.8	60.3	44.8	64.3	79.4	57.5	64.4	57.6	52.4	96.1	62.9	65.8	94.4	92.0	68.5
CIE [58]	52.5	68.6	70.2	57.1	82.1	77.0	70.7	73.1	43.8	69.9	62.4	70.2	70.3	66.4	47.6	85.3	71.7	64.0	83.9	91.7	68.9
QC-DGM [13]	49.6	64.6	67.1	62.4	82.1	79.9	74.8	73.5	43.0	68.4	66.5	67.2	71.4	70.1	48.6	92.4	69.2	70.9	90.9	92.0	70.3
DGMC [11]	50.4	67.6	70.7	70.5	87.2	85.2	82.5	74.3	46.2	69.4	69.9	73.9	73.8	65.4	51.6	98.0	73.2	69.6	94.3	89.6	73.2
BBGM [40]	61.9	71.1	<u>79.7</u>	79.0	87.4	<u>94.0</u>	<u>89.5</u>	<u>80.2</u>	56.8	79.1	64.6	78.9	76.2	75.1	<u>65.2</u>	<u>98.2</u>	<u>77.3</u>	77.0	94.9	93.9	79.0
NGM-v2 [49]	61.8	71.2	77.6	78.8	87.3	93.6	87.7	79.8	55.4	77.8	<u>89.5</u>	78.8	<u>80.1</u>	<u>79.2</u>	62.6	97.7	77.7	75.7	96.7	93.2	80.1
SCGM [30]	<u>62.9</u>	72.9	79.6	<u>79.5</u>	89.3	94.1	89.1	79.2	58.4	79.3	<u>80.5</u>	<u>79.9</u>	79.5	76.8	64.8	98.1	<u>78.0</u>	75.9	98.0	93.2	80.5
ASAR [39]	<u>62.9</u>	<u>74.3</u>	79.5	80.1	89.2	<u>94.0</u>	88.9	78.9	<u>58.8</u>	<u>79.8</u>	88.2	78.9	79.5	77.9	64.9	<u>98.2</u>	<u>77.5</u>	<u>77.1</u>	<u>98.6</u>	93.7	<u>81.1</u>
COMMON	65.6	75.2	80.8	<u>79.5</u>	89.3	92.3	90.1	81.8	61.6	80.7	95.0	82.0	81.6	79.5	66.6	98.9	78.9	80.9	99.3	<u>93.8</u>	82.7

Table 1. Keypoint matching accuracy (%) on Pascal VOC with standard intersection filtering. The best and second-best results are **highlighted** and underlined, respectively.

as follows,

$$L_{Rgraph} = \|\mathbf{W} \circ (\mathbf{P}_A \mathbf{P}_A^T - \mathbf{P}_B \mathbf{P}_B^T)\|_F^2 + \|\mathbf{W} \circ (\mathbf{P}_A \mathbf{P}_B^T - \mathbf{P}_B \mathbf{P}_A^T)\|_F^2, \quad (13)$$

where \circ is the Hadamard product.

Different from most existing distillation methods that acquire knowledge from a pre-trained model [17, 43], our momentum distillation strategy owns the following two merits: i) enjoying the robustness against BNC by exploiting the memory effect characteristic of the momentum model, *i.e.*, exponential-moving-average parameters would resist overfitting the noise; ii) bootstrapping the matching performance without resorting to extra knowledge or models.

Finally, by combining Eqs. (11) and (13), the overall robust quadratic contrastive learning loss is induced as,

$$L_{Rquadratic} = L_{RInfoNCE} + L_{Rgraph}. \quad (14)$$

Note that different loss terms are weighted equally in Eqs. (13) and (14), and our method performs well without laborious searching on the balance hyper-parameters.

4. Experiments

In this section, we carry out extensive experiments on three widely-used GM datasets with the comparisons of 12 state-of-the-art deep graph matching approaches.

4.1. Experimental Settings

Datasets. We conduct experiments on Pascal VOC with Berkeley annotation [3], SPair-71K [34], and Willow Object Class [8]. For extensive evaluation, we report the average and per-category performance. The objective of all experiments is to maximize the average matching accuracy.

Implementation details. We implement our method in PyTorch 1.10.0 and conduct all evaluations on the Ubuntu-20.04 OS with an NVIDIA 3090 GPU. For all datasets, we use the exact same set of hyper-parameters. The encoder

network consists of an ImageNet-pretrained VGG16 [42] image encoder, a graph neural network SplineCNN [12], and a two-layer projection head [6]. All details of the network have been presented in the supplementary material. To optimize the networks, we adopt Adam optimizer [23] with the default parameters and set the initial learning rate as $3e^{-4}$. The learning rate for fine-tuning the VGG network is $2e^{-5}$. The batch size is set to 8 image pairs and the distillation temperature α is fixed to 0.4. To obtain the permutation matrix \mathbf{Y} , we perform the Hungarian algorithm on the similarity matrix \mathbf{S} obtained from the base encoder following [49, 31, 47, 48, 39, 59].

Compared methods. We compare our COMMON with the following 12 popular deep graph matching methods: GMN [60], PCA [47], NGM [49], IPCA [48], CIE [58], DGMC [11], LCS [51], BBGM [40], QC-DGM [13], NGM-v2 [49], SCGM [30], and ASAR [39] on the popular open-source graph matching toolbox ThinkMatch¹, for the purpose of better reproducibility and more fair comparison. Note that our code has also been included in ThinkMatch.

4.2. Results on Graph Matching

Pascal VOC [3] consists of 7,020 training images and 1,682 testing images with 20 classes in total. The number of nodes per graph ranges from 6 to 23. Following [40], we test our method by filtering out the non-matched points before matching. As shown in Table 1, our method outperforms the counterparts with +1.6% in terms of accuracy. It is worth pointing out that our method achieves remarkable performance improvements on the classes with prominent noisy correspondence, *e.g.*, table (+6.8%) and sofa (+3.8%).

SPair-71k [34] consists of 70,958 image pairs collected from Pascal VOC 2012 and Pascal 3D+. Following the data preparation in [49, 47, 48], each object is cropped to

¹<https://github.com/Thinklab-SJTU/ThinkMatch>

Method	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Dog	Horse	Mbike	Person	Plant	Sheep	Train	Tv	Mean
GMN [60]	59.9	51.0	74.3	46.7	63.3	75.5	69.5	64.6	57.5	73.0	58.7	59.1	63.2	51.2	86.9	57.9	70.0	92.4	65.3
PCA [47]	64.7	45.7	78.1	51.3	63.8	72.7	61.2	62.8	62.6	68.2	59.1	61.2	64.9	57.7	87.4	60.4	72.5	92.8	66.0
NGM [49]	66.4	52.6	77.0	49.6	67.7	78.8	67.6	68.3	59.2	73.6	63.9	60.7	70.7	60.9	87.5	63.9	79.8	91.5	68.9
IPCA [48]	69.0	52.9	80.4	54.3	66.5	80.0	68.5	71.4	61.4	74.8	66.3	65.1	69.6	63.9	91.1	65.4	82.9	97.5	71.2
CIE [58]	71.5	57.1	81.7	56.7	67.9	82.5	73.4	74.5	62.6	78.0	68.7	66.3	73.7	66.0	92.5	67.2	82.3	97.5	73.3
NGM-v2 [49]	68.8	63.3	86.8	70.1	69.7	94.7	87.4	77.4	72.1	80.7	74.3	72.5	79.5	73.4	98.9	81.2	94.3	98.7	80.2
BBGM [40]	75.3	65.0	87.6	78.0	69.8	94.0	87.8	78.3	72.8	82.7	76.6	76.3	80.1	75.0	98.7	85.2	96.3	98.0	82.1
ASAR [39]	72.4	61.8	91.8	79.1	71.2	97.4	90.4	78.3	74.2	83.1	77.3	77.0	83.1	76.4	99.5	85.2	97.8	99.5	83.1
COMMON	77.3	68.2	92.0	79.5	70.4	97.5	91.6	82.5	72.2	88.0	80.0	74.1	83.4	82.8	99.9	84.4	98.2	99.8	84.5

Table 2. Keypoint matching accuracy (%) on SPair-71k for all classes.

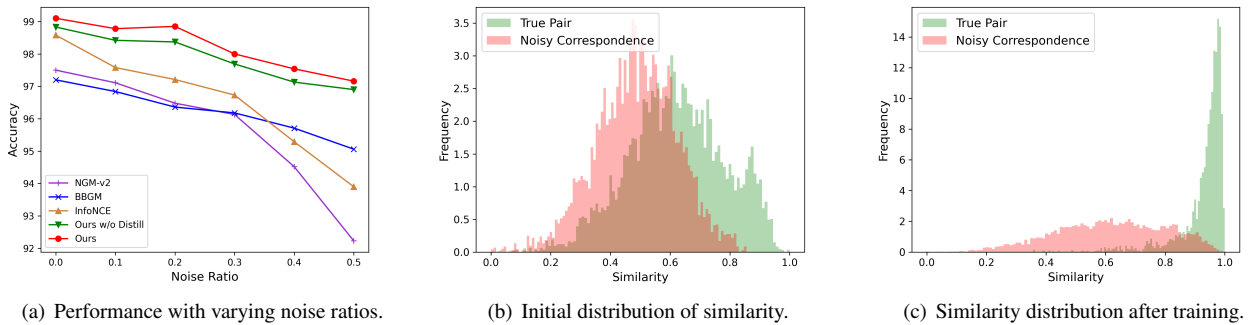


Figure 3. Effectiveness on noisy correspondence.

its bounding box and scaled to 256×256 . As shown in Table 2, our method consistently improves the matching performance (+1.4%).

Willow Object [8] consists of 256 images in 5 categories, where each target object is annotated with 10 distinctive landmarks. Following the protocol in [47, 48, 49], we train our methods on the first 20 images and report testing results on the rest. As shown in Table 3, our method significantly outperforms baselines (+1.4%).

4.3. Evaluation on Noisy Correspondence

In this section, we evaluate the effectiveness of our method against noisy correspondence with synthetic noise experiments, varying viewpoint experiments, ablation studies, and parameter analysis.

4.3.1 Synthetic Noise Experiment

To explicitly demonstrate the effectiveness of our method against noisy correspondence, we evaluate models on Willow Object with synthetic noisy correspondence. Concretely, we randomly select some keypoints in the training set as the noise points by adding displacement to the location coordinates. The displacement (s, θ) is generated from uniform distribution: $s \sim \mathcal{U}(0.1, 0.2)$, $\theta \sim \mathcal{U}(0, 360)$ where s is the displacement value and θ is the angle. The displacement value is further scaled with respect to the size

of the bounding box. The noisy rate η is defined as the percentage of noise points.

Varying noise ratios. As shown in Fig. 3(a), we conduct experiments by varying the noise rate η from 0 to 0.5 with an interval of 0.1. From the results, one could observe that: i) our method significantly outperforms baselines in all noise rate settings; ii) momentum distillation consistently improves the robustness against noisy correspondence.

Distribution of similarity scores. We further explore the similarity scores of the correspondence learned by our method under noise rate $\eta = 0.3$. In Figs. 3(b) and 3(c), we show the histograms of similarity scores for true pairs (in which keypoints are without synthetic noise) and noisy correspondence (at least one of the keypoints is with synthetic noise) separately. The initial similarity of these two kinds of pairs is confused and hardly distinguishable. After training, the similarity of noisy correspondence is 0.6 on average, while that of true pairs is 0.95 on average and ranges into $[0.7, 1.0]$. In other words, our method prevents noisy correspondence from dominating the network optimization, thus eliminating the negative impact of noisy correspondence.

4.3.2 Evaluation with Different Viewpoint Difficulty

SPair-71k contains images with diverse variations in viewpoint and divides image pairs into easy, medium, and hard ones. In practice, we find that noisy correspondence such

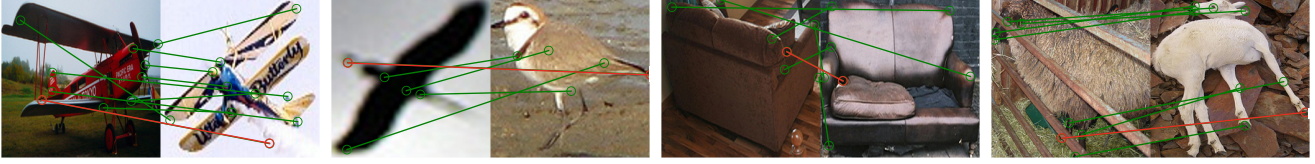


Figure 4. Examples of noisy correspondence in Pascal VOC and SPair-71k. The red line indicates wrongly annotated correspondence.

Method	Car	Duck	Face	Mbike	Wbottle	Mean
GMN [60]	67.9	76.7	99.8	69.2	83.1	79.3
NGM [49]	84.2	77.6	99.4	76.8	88.3	85.3
PCA [47]	87.6	83.6	100	77.6	88.4	87.4
CIE [58]	85.8	82.1	99.9	88.4	88.7	89.0
IPCA [48]	90.4	88.6	100	83.0	88.3	90.1
SCGM [30]	91.3	73.0	100	95.6	96.6	91.3
ASAR [39]	92.5	84.0	100	95.4	99.0	94.2
LCS [51]	91.2	86.2	100	99.4	97.9	94.9
DGMC [11]	98.3	90.2	100	98.5	98.1	97.0
BBGM [40]	96.8	89.9	100	99.8	99.4	97.2
NGM-v2 [49]	97.4	<u>93.4</u>	100	98.6	98.3	97.5
QC-DGM [13]	<u>98.0</u>	92.8	100	98.8	99.0	<u>97.7</u>
COMMON	97.6	98.2	100	100	99.6	99.1

Table 3. Keypoint matching accuracy (%) across all objects on Willow Object.

Method	Viewpoint difficulty			All
	Easy	Medium	Hard	
BBGM [40]	84.7	78.9	73.6	82.1
ASAR [39]	86.5	79.1	72.5	83.1
COMMON	86.6 (+0.1)	81.4 (+2.3)	76.4 (+2.8)	84.5 (+1.4)

Table 4. Keypoint matching accuracy (%) on SPair-71k grouped by levels of difficulty in the viewpoint of the matching pair.

as occlusion is more likely to occur in hard image pairs. In other words, with the increment of viewpoint difficulty, noisy correspondence appears more frequently. As shown in Table 4, our method consistently improves the matching result, particularly for the image pairs with high viewpoint difficulty (+2.8% for hard pairs). This experiment highlights the ability of our method against noisy correspondence.

4.3.3 Ablation Studies and Parameter Analysis

To evaluate our framework, we conduct comprehensive ablation studies by separately investigating each component. As shown in Table 5, all modules are inseparable and bring substantial performance gains. Notably, the momentum distillation strategy remarkably improves the performance by alleviating the negative impact of bi-level noisy correspondence. Besides, the graph consistency terms play indispensable roles in learning quadratic correlations and our loss outperforms the quadratic regularization [35]. Following, we present the parameter sensitivity analysis on the distil-

Method	Pascal VOC	SPair-71k
COMMON (FULL)	82.67	84.54
- w/o graph consistency	81.95	83.83
- w/o distillation	81.77	83.94
- InfoNCE and quadratic regularization [35]	81.67	83.66
- InfoNCE	81.33	83.38

Table 5. Ablation study of COMMON in terms of accuracy. Default settings are marked in gray.

α	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
ACC	81.90	82.34	82.69	82.62	82.67	82.50	82.59	82.62	82.66	82.41

Table 6. Parameter analysis of COMMON with the increase of the distillation parameter α on Pascal VOC. Default settings are marked in gray.

lation temperature α in Table 6. As shown, the distillation strategy indeed improves the performance and our method is insensitive to the choice of α .

4.4. Visualization on Noisy Correspondence

In this section, we identify some noisy correspondence pairs predicted by our model with the lowest matching similarity in Fig. 4. More visualization results have been presented in the supplementary material.

5. Conclusion

This paper reveals a new problem for graph matching, *i.e.*, bi-level noisy correspondence, which refers to wrongly annotated node-to-node correspondence and accompanied edge-to-edge correspondence. To tackle this challenge, the proposed method introduces the momentum distillation strategy to rebalance the quadratic contrastive loss and mitigate the influence of BNC. This work potentially provides a novel insight into the community and may inspire some further exploration of the noisy correspondence problem. In terms of application scenario, noisy correspondence might appear in not only two-graph matching but also multi-graph and hyper-graph matching. In terms of methodology, it may be feasible to design NC-robust GM networks by controlling the information propagation in graph neural networks.

Acknowledgments

This work was supported by the National Key R&D Program of China under Grant 2020YFB1406702, in part by NSFC under Grant U21B2040 and 62176171, and in part by the 111 Project under grant B21044.

References

- [1] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *ICML*, pages 233–242, 2017. [2](#), [5](#)
- [2] Amir Beck and Marc Teboulle. Smoothing and first order methods: A unified framework. *SIAM Journal on Optimization*, 22(2):557–580, 2012. [4](#)
- [3] Lubomir Bourdev and Jitendra Malik. Poselets: Body part detectors trained using 3d human pose annotations. In *ICCV*, pages 1365–1372, 2009. [1](#), [6](#)
- [4] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, 2016. [1](#)
- [5] Lichang Chen, Guosheng Lin, Shijie Wang, and Qingyao Wu. Graph edit distance reward: Learning to edit scene graph. In *ECCV*, pages 539–554, 2020. [1](#)
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv:2002.05709*, 2020. [2](#), [3](#), [6](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv:2003.04297*, 2020. [4](#)
- [8] Minsu Cho, Karteek Alahari, and Jean Ponce. Learning graphs to match. In *ICCV*, pages 25–32, 2013. [6](#), [7](#)
- [9] Minsu Cho, Jungmin Lee, and Kyoung Mu Lee. Reweighted random walks for graph matching. In *ECCV*, pages 492–505, 2010. [1](#)
- [10] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, volume 26, 2013. [2](#)
- [11] Matthias Fey, Jan E Lenssen, Christopher Morris, Jonathan Masci, and Nils M Kriege. Deep graph matching consensus. In *ICLR*, 2020. [2](#), [6](#), [8](#)
- [12] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *CVPR*, pages 869–877, 2018. [6](#)
- [13] Quankai Gao, Fudong Wang, Nan Xue, Jin-Gang Yu, and Gui-Song Xia. Deep graph matching under quadratic constraint. In *CVPR*, pages 5069–5078, 2021. [1](#), [2](#), [6](#), [8](#)
- [14] Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple contrastive learning of sentence embeddings. In *EMNLP*, 2021. [3](#)
- [15] Shashank Goel, Hritik Bansal, Sumit Bhatia, Ryan Rossi, Vishwa Vinay, and Aditya Grover. Cyclop: Cyclic contrastive language-image pretraining. *Advances in Neural Information Processing Systems*, 35:6704–6719, 2022. [3](#)
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. [2](#), [3](#), [4](#), [5](#)
- [17] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2(7), 2015. [6](#)
- [18] Zhenyu Huang, Guocheng Niu, Xiao Liu, Wenbiao Ding, Xinyan Xiao, Hua Wu, and Xi Peng. Learning with noisy correspondence for cross-modal matching. In *NeurIPS*, pages 29406–29419, 2021. [2](#)
- [19] Ahmet Iscen, Jack Valmadre, Anurag Arnab, and Cordelia Schmid. Learning with neighbor consistency for noisy labels. In *CVPR*, pages 4672–4681, 2022. [2](#)
- [20] Bo Jiang, Pengfei Sun, Ziyang Zhang, Jin Tang, and Bin Luo. Gamnet: Robust feature matching via graph adversarial-matching network. In *ACMMM*, pages 5419–5426, 2021. [2](#)
- [21] Yangbangyan Jiang, Qianqian Xu, Zhiyong Yang, Xiaochun Cao, and Qingming Huang. Dm2c: Deep mixed-modal clustering. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *NeurIPS*, pages 5880–5890, 2019. [3](#)
- [22] Nazmul Karim, Mamshad Nayeem Rizve, Nazanin Rahnavard, Ajmal Mian, and Mubarak Shah. Unicon: Combating label noise through uniform selection and contrastive learning. In *CVPR*, pages 9676–9686, 2022. [5](#)
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. [6](#)
- [24] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pages 5447–5456, 2018. [2](#), [5](#)
- [25] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34:9694–9705, 2021. [2](#), [5](#)
- [26] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020. [2](#), [5](#)
- [27] Yunfan Li, Mouxing Yang, Dezhong Peng, Taihao Li, Jiantao Huang, and Xi Peng. Twin contrastive learning for online clustering. *International Journal of Computer Vision*, 130(9):2205–2221, 2022. [3](#)
- [28] Yijie Lin, Yuanbiao Gou, Xiaotian Liu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Dual contrastive prediction for incomplete multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–14, 2022. [3](#)
- [29] Yijie Lin, Yuanbiao Gou, Zitao Liu, Boyun Li, Jiancheng Lv, and Xi Peng. Completer: Incomplete multi-view clustering via contrastive prediction. In *CVPR*, pages 11174–11183, 2021. [3](#)
- [30] Chang Liu, Shaofeng Zhang, Xiaokang Yang, and Junchi Yan. Self-supervised learning of visual graph matching. In *ECCV*, 2022. [3](#), [6](#), [8](#)

- [31] He Liu, Tao Wang, Yidong Li, Congyan Lang, Yi Jin, and Haibin Ling. Joint graph learning and matching for semantic feature correspondence. *arXiv:2109.00240*, 2021. 2, 6
- [32] Junhong Liu, Yijie Lin, Liang Jiang, Jia Liu, Zujie Wen, and Xi Peng. Improve interpretability of neural networks via sparse contrastive coding. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 2022. 3
- [33] Linfeng Liu, Michael C Hughes, Soha Hassoun, and Liping Liu. Stochastic iterative graph matching. In *ICML*, pages 6815–6825, 2021. 2
- [34] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Spair-71k: A large-scale benchmark for semantic correspondence. *arXiv:1908.10543*, 2019. 1, 6
- [35] Artem Moskalev, Ivan Sosnovik, Volker Fischer, and Arnold Smeulders. Contrasting quadratic assignments for set-based representation learning. In *ECCV*, pages 88–104, 2022. 2, 3, 4, 8
- [36] Pedro O O Pinheiro, Amjad Almahairi, Ryan Benmalek, Florian Golemo, and Aaron C Courville. Unsupervised learning of dense visual representations. In *NeurIPS*, volume 33, pages 4489–4500, 2020. 3
- [37] Jingwei Qu, Haibin Ling, Chenrui Zhang, Xiaoqing Lyu, and Zhi Tang. Adaptive edge attention for graph matching with outliers. In *IJCAI*, pages 966–972, 2021. 3
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 4
- [39] Qibing Ren, Qingquan Bao, Runzhong Wang, and Junchi Yan. Appearance and structure aware robust deep visual graph matching: Attack, defense and beyond. In *CVPR*, pages 15263–15272, 2022. 3, 6, 7, 8
- [40] Michal Rolínek, Paul Swoboda, Dominik Zietlow, Anselm Paulus, Vít Musil, and Georg Martius. Deep graph matching via blackbox differentiation of combinatorial solvers. In *ECCV*, pages 407–424, 2020. 1, 2, 3, 6, 7, 8
- [41] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, pages 4938–4947, 2020. 1, 3, 5
- [42] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2014. 6
- [43] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357, 2021. 6
- [44] Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, Yasemin Altun, and Yoram Singer. Large margin methods for structured and interdependent output variables. *Journal of machine learning research*, 6(9), 2005. 4
- [45] Nikolai Ufer and Bjorn Ommer. Deep semantic feature matching. In *CVPR*, pages 6914–6923, 2017. 1
- [46] Fudong Wang, Nan Xue, Jin-Gang Yu, and Gui-Song Xia. Zero-assignment constraint for graph matching with outliers. In *CVPR*, pages 3033–3042, 2020. 3
- [47] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Learning combinatorial embedding networks for deep graph matching. In *ICCV*, pages 3056–3065, 2019. 1, 2, 4, 5, 6, 7, 8
- [48] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Combinatorial learning of robust deep graph matching: an embedding based approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 1, 6, 7, 8
- [49] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Neural graph matching network: Learning lawler’s quadratic assignment problem with extension to hypergraph and multiple-graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2, 3, 6, 7, 8
- [50] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, pages 9929–9939, 2020. 2
- [51] Tao Wang, He Liu, Yidong Li, Yi Jin, Xiaohui Hou, and Haibin Ling. Learning combinatorial solver for graph matching. In *CVPR*, pages 7568–7577, 2020. 2, 6, 8
- [52] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, pages 3024–3033, 2021. 3
- [53] Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang. Robust early-learning: Hindering the memorization of noisy labels. In *ICLR*, 2020. 2
- [54] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *CVPR*, pages 16684–16693, 2021. 3
- [55] Mouxing Yang, Zhenyu Huang, Peng Hu, Taihao Li, Jiancheng Lv, and Xi Peng. Learning with twin noisy labels for visible-infrared person re-identification. In *CVPR*, pages 14308–14317, 2022. 2
- [56] Yiding Yang, Zhou Ren, Haoxiang Li, Chunlun Zhou, Xinchao Wang, and Gang Hua. Learning dynamics via graph neural networks for human pose estimation and tracking. In *CVPR*, pages 8074–8084, 2021. 1
- [57] Zhiyong Yang, Qianqian Xu, Weigang Zhang, Xiaochun Cao, and Qingming Huang. Split multiplicative multi-view subspace clustering. *IEEE Trans. Image Process.*, 28(10):5147–5160, 2019. 3
- [58] Tianshu Yu, Runzhong Wang, Junchi Yan, and Baoxin Li. Learning deep graph matching with channel-independent embedding and hungarian attention. In *ICLR*, 2019. 1, 2, 6, 7, 8
- [59] Tianshu Yu, Runzhong Wang, Junchi Yan, and Baoxin Li. Deep latent graph matching. In *ICML*, pages 12187–12197, 2021. 6
- [60] Andrei Zanfir and Cristian Sminchisescu. Deep learning of graph matching. In *CVPR*, pages 2684–2693, 2018. 1, 2, 6, 7, 8