
WeaveNet for Approximating Assignment Problems

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Assignment, a task to match a limited number of elements, is a fundamental
2 problem in informatics. Many assignment problems have no exact solvers due
3 to their NP-hardness or incomplete input, and their approximation algorithms
4 have been studied for a long time. However, individual practical applications
5 have various objective functions and prior assumptions, which usually differ from
6 academic studies. This gap hinders applying the algorithms to real problems
7 despite their theoretically ensured performance. In contrast, a learning-based
8 method can be a promising solution to fill the gap. To open a new vista for
9 real-world assignment problems, we propose a novel neural network architecture,
10 *WeaveNet*. Its core module, *feature weaving layer*, is stacked to model frequent
11 communication between elements in a parameter-efficient way for solving the
12 combinatorial problem of assignment. To evaluate the model, we approximated
13 one of the most popular non-linear assignment problems, stable matching with two
14 different *strongly NP-hard* settings. The experimental results showed its impressive
15 performance among the learning-based baselines. Furthermore, we achieved better
16 or comparative performance to the state-of-the-art algorithmic method, depending
17 on the size of problem instances.

18 1 Introduction

19 From multiple object tracking to job matching, assignment problems can represent a wide variety of
20 applications. An assignment problem is typically defined on a bipartite graph, a graph with two sets
21 of nodes A and B with edges $E = A \times B$ ($N = |A|$, $M = |B|$, $N \geq M$). On the graph, the task
22 is to find a matching $m \in \{0, 1\}^{A \times B}$ (a set of edges represented as a binary matrix) that satisfies
23 constraints and/or maximizes objectives. Depending on real-world scenes, there must be various
24 objectives and constraints for m . A typical constraint is a one-to-one correspondence (i.e., every node
25 has at most one matched partner in m) and, for simplicity, we always assume it in this paper.

26 Matching stability is another example of such constraints. It is a non-linear constraint first introduced
27 for a hospital-student assignment problem (Gale and Shapley, 1962) based on the preferences of
28 hospitals among students and vice versa. We say a matching m is unstable when there exist $a \in A$
29 and $b \in B$ which are unmatched in m ($m_{ab} = 0$) but both prefer each other more than their partner
30 in m . We can obtain a stable matching m in $O(N^2)$ by the Gale-Shapley (GS) algorithm (Gale
31 and Shapley, 1962). However, when m is expected to have the minimum difference in the total
32 satisfactions between sides A and B (known as sex-equal stable matching), the problem becomes
33 *strongly NP-hard*¹ (Kato, 1993; McDermid and Irving, 2014).

34 In addition to the NP-hardness, we also face difficulties to obtain the best assignment when assignment
35 candidates may randomly disappear (e.g., multiple object tracking with occlusions (Emami et al.,
36 2020) or joint matching in multi-person pose estimation (Cao et al., 2017)). In such cases, we need

¹*strongly NP-hard* is a subclass of NP-hard and considered more complex than general NP-hard problems

37 to compensate for the inputs of incomplete information by its stochastic properties. The traditional
38 methods often use sub-optimal approximations to avoid solving complex assignment problems.
39 A differential assignment model can be a future option that enables end-to-end training for such
40 applications.

41 Toward such future applications, this paper aims to propose an effective and promising differential
42 solver for assignment problems. The contribution of this paper is four-fold:

- 43 1. We proposed *WeaveNet*, a novel neural network architecture for assignment problems and
44 *set-encoder*, a novel local structure.
- 45 2. We proposed a novel technique, *split batch normalization*, to deal with a strong asymmetry
46 in input distributions for sides A and B .
- 47 3. We focused on stable matching, a classical non-linear assignment problem actively studied
48 even in recent years, and proposed a novel evaluation protocol² with *pseudo costs*, which
49 enables us to compare learning-based solvers and algorithmic solvers directly.
- 50 4. We achieved a better performance with the state-of-the-art algorithmic baseline when
51 $N = 20$, and a comparative performance when $N = 30$. We also outperformed any
52 learning-based baselines with a large margin.

53 2 Related work

54 Despite the recent research interest in deep learning technology, we hardly have a fully differential
55 assignment solver. As long as authors know, there are two past attempts to solve assignment problems
56 by a fully differential model. Li (2019) has tried to solve stable matching by multiple layer perceptrons
57 (MLP). Their contribution is in the proposed relaxation of the non-linear stability constraint to a
58 differential loss function. However, the MLP is too redundant to learn the assignment strategy without
59 overfitting. In addition, the proposed auxiliary loss to maintain the output to be one-to-one matching
60 (symmetric doubly stochastic function) overly constrains the solution search space. In this study, we
61 propose a parameter-efficient differential model and a weaker but sufficient constraint to output a
62 one-to-one matching.

63 The second attempt is made by Gibbons *et al.* (2019), where Deep Bipartite Matching (DBM) is
64 proposed. They tested their model with the weapon-target assignment (WTA) problem. WTA is a
65 classical NP-hard problem whose state-of-the-art algorithm (Ahuja *et al.*, 2007) could find optimal
66 solution when $N \leq 20$ in the experiment although there is no theoretical guarantee. In this sense,
67 we can consider WTA is empirically easier than sex-equal stable matching, for which we have no
68 such efficient solvers even for $N = 5$. In addition, DBM is trained in a supervised manner or with
69 reinforcement learning, which is hard to apply to a larger N . Furthermore, the implementation details
70 are not completely explained, and their dataset and source codes are not publicly available. Finally,
71 the architecture of DBM is still parameter-redundant, and their local structure is sub-optimal. In this
72 study, we propose a more parameter-efficient two-stream architecture, *WeaveNet*, with a novel local
73 structure, *set-encoder*, both of which have significant impacts on the performance.

74 In addition to the above methods, it is natural to consider using graph convolutional networks (GCNs).
75 However, there are no GCN methods for assignment problems due to the over-smoothing problem (Li
76 *et al.*, 2018; Oono and Suzuki, 2020). Because any graph-convolutional layer summarizes the output
77 with neighboring nodes, its smoothing effect eliminates expressive power for node classification. To
78 avoid such elimination, GIN (Xu *et al.*, 2019), the state-of-the-art GCN method, stacks only two
79 layers for a node classification task. Such elimination is critical for an assignment-problem solver
80 because it needs to identify any slight difference through frequent communication among nodes.
81 Unlike GIN, our model retains edge-wise features rather than node-wise summaries, which does not
82 cause the smoothing problem. Therefore, we can make the model very deep, which any traditional
83 graph convolutional networks cannot.

84 3 Stable matching problem as a benchmark task

85 To evaluate learning-based assignment solvers, we adopt two *strongly NP-hard* variants of stable
86 matching. They have been actively studied for a long time (Kato, 1993; Iwama *et al.*, 2010; Dworzak,

²The source code and datasets are included in this submission and will be publicly available.

87 [2016; Gupta *et al.*, 2019) and their state-of-the-art algorithm by Tziavelis *et al.* (2019) must be a
 88 strong baseline against learning-based methods. Hence, we set these two variants as the benchmark
 89 task for learning-based assignment problems.

90 An instance I of a stable matching problem consists of two sets of agents A and B on a bipartite
 91 graph. Fig. 1 illustrates an example of I . Each agent a_i in A ($0 < i \leq N$) has a preference list p_i^A ,
 92 which is an ordered set of elements in B and $p_{ij}^A = \text{rank}(b_j; p_i^A)$ is the index of b_j in the list p_i^A . a_i
 93 prefers b_j to $b_{j'}$ if $p_{ij}^A < p_{ij'}^A$. Similarly, each agent b_j in B ($0 < j \leq M$) has a preference list p_j^B .




Figure 1: An example of assignment, where m (black edges) is not stable due to the blocking pair (the orange edge), while m' (green edges) is stable.

94 For a matching m , we say that an unmatched pair $\{a_v, b_w\}$ ($m_{vw} = 0$) blocks m if a_v 's partner
 95 b_j ($m_{vj} = 1$) and b_w 's partner a_i ($m_{iw} = 1$) satisfy the conditions $p_{vw}^A < p_{vj}^A$ and $p_{vw}^B < p_{wi}^B$. Here,
 96 $\{a_v, b_w\}$ is called a blocking pair (the orange edge blocks a matching of black edges in the figure).

97 A matching is stable if (and only if) it includes no blocking pair (the green edges in the figure). Note
 98 that I always has at least one stable matching, and the Gale-Shapley (GS) algorithm can find it in
 99 $O(N^2)$. However, the GS algorithm has a biased nature, where one side is prioritized and the other
 100 side only gets the least preferable result among all the possibilities of stable matching.

101 To compensate for the unfairness, we can introduce diverse objectives to maintain a stable matching
 102 fair. Among them, the following two objectives make the stable matching problem *strongly NP-hard*.
 103 The first one is **Sex equality cost** (Seq) (Gusfield and Irving, 1989). It focuses on the unfairness
 104 brought by the gap between the two sides' satisfaction and defined by

$$Seq(m; I) = |P(m; A) - P(m; B)|, \quad P(m; A) = \sum_{\{a_i, b_j\} \in m} p_{ij}^A, \quad P(m; B) = \sum_{\{a_i, b_j\} \in m} p_{ji}^B. \quad (1)$$

105 The other is **Balance cost** (Bal) (Feder, 1995; Gupta *et al.*, 2019), which is a compromise between
 106 side-equality and overall satisfaction. It is defined by

$$Bal(m; I) = \max(P(m; A), P(m; B)). \quad (2)$$

107 In the proposed evaluation protocol, we minimize either cost while maintaining stable one-to-one
 108 matching.

109 **Input and output data format for stable matching** Learning-based approximation is realized by
 110 a trainable function F that outputs a matching $\hat{m} \in [0, 1]^{N \times M}$, which is an $N \times M$ matrix. As
 111 for the input, the value range of the preference rank depends on the problem size, which causes a
 112 range shift of the input distribution. To avoid such shift, we linearly re-scale³ the rank of preference
 113 p_{ij}^* ($* \in \{A, B\}$) from $[1, N]$ to a normalized score s_{ij}^* , ranged in $(0, 1]$ to make it invariant to N ,
 114 where 1 for the highest rank. Then, we obtain the input as matrices S^A and S^B , where s_{ij}^A is the
 115 ij -element of S^A .

116 4 Deep-learning-based fair stable matching with WeaveNet

117 4.1 WeaveNet

118 One of the required properties of $F : (S^A, S^B) \rightarrow \hat{m}$ is to take all the agents' preference into
 119 account when determining the presence of each edge in the output \hat{m} . Li (2019) implemented this by

³The details of this linear re-scaling are based on Li (2019) and described in A.1. Note that sections numbered with capital letters appear in the supplementary material.




Figure 2: WeaveNet architecture. L feature weaving layers are stacked with shortcut paths to be a deep network. The encoded features are fed into $\text{Conv}(1 \times 1)$ layer to obtain logits (\hat{m}^A, \hat{m}^B) . The output \hat{m} will be binarized in prediction phase to represent a matching.

120 MLPs, where S^A and S^B are destructured and concatenated into a single flat vector (with the length
 121 of $2NM$) and fed to the MLP. Its output (a flat vector with the length of NM) is restructured into a
 122 matrix \hat{m} . The MLP model, however, would face difficulties due to the following four problems.

- 123 (a) Preference lists of multiple agents are encoded by independent parameters, though they share a
 124 format so that we could efficiently process them in the same manner.
 125 (b) MLP only supports a fixed-size input, so training different models for different cases of N
 126 becomes mandatory.
 127 (c) F should be permutation invariant, which means the matching result should be unchanged even if
 128 we shuffle the order of agents in S^A and S^B , but MLP does not satisfy.
 129 (d) A shallow MLP model may be insufficient to approximate an exact solver for the NP-hard problem
 130 when N is large.

131 To address the above weaknesses of MLP, we propose the feature weaving network (**WeaveNet**) which
 132 has the properties of (a) **shared encoder**, (b) **variable-size input**, (c) **permutation invariance**, and
 133 (d) **residual structure**. The WeaveNet, as shown in Fig. 2, consists of L feature weaving (FW)
 134 layers. It has two streams of A and B . In a symmetric manner, each stream models the agent’s act of
 135 selecting the one on the opposite side while sharing weights to enhance the parameter efficiency. The
 136 shortcut paths at every two FW layers make them residual blocks, which allows the model to be as
 137 deep as possible. We explain its details as follows.




Figure 3: Feature weaving layer orthogonally cross-concatenates the weftwise and warpwise components (Z_ℓ^A and Z_ℓ^B) in a symmetric way (cross-concatenation). Then, the concatenated tensors are separated into $z_\ell^{a_i}$ (or $z_\ell^{b_j}$), which represents a set of outgoing edges from agent a_i (or b_j), and independently fed to E_ℓ .




Figure 4: Illustration of the process in set encoder E_ℓ , where $z_\ell^{a_i}$ (colored in white) is once encoded to D' channel features (colored in pale blue), then max-pooled to obtain statistics in the feature set (colored in blue). The statistics information is concatenated to each input feature and further encoded (color in a gradation).

138 Fig. 3 illustrates the detail of a single FW layer, which is the core architecture of the proposed network.
 139 FW layer is a two-stream layer whose inputs consist of a *weftwise* component Z_ℓ^A and a *warpwise*
 140 component Z_ℓ^B , which are the output of $(l - 1)$ -th layer and $Z_0^A = S^A$ and $Z_0^B = S^B$ for the first
 141 layer. The two components are symmetrically concatenated in each stream (**cross-concatenation**).
 142 Then these concatenations are separated into agent-wise features, each of which is a set of outgoing-
 143 edge features of an agent (indicating the preference from that agent to every matching candidate).
 144 These features are processed by the encoder E_ℓ **shared by every agent in both A and B** . As for
 145 an encoder that can embed **variable-size** input in a **permutation invariant** manner, we adopted the
 146 structure inspired by DeepSet (Zaheer et al., 2017) and PointNet (Qi et al., 2017) (Fig. 4), which

147 consists of two convolutional layers with kernel size 1 and a set-wise max-pooling layer, followed by
 148 batch-normalization and PReLU activation. We refer to this structure as *set encoder*.

149 **Mathematical formulations** Z_ℓ^A in Fig. 3 is a third-order tensor whose dimensions, in sequence,
 150 corresponding to the agent, candidate, and feature dimension, with a size of (N, M, D) . Similarly,
 151 Z_ℓ^B has a size of (M, N, D) . The **cross-concatenation** is defined as

$$Z_\ell'^A = \text{cat}(Z_\ell^A, P_{A \leftrightarrow B}(Z_\ell^B)), \quad (3)$$

152 where $P_{A \leftrightarrow B}$ swaps the first and second dimensions of the tensor, and $\text{cat}(\{Z_1, Z_2, \dots\})$ concatenates
 153 the features of two tensors Z_1, Z_2 . $Z_\ell'^A$ is then sliced into agent-wise features $z_\ell'^{a_i}$ and we obtain
 154 $Z_{\ell+1}^A = (E_\ell(z_\ell'^{a_i}) | 0 < i \leq N)$, which is also a third-order tensor (and fed to the next layer). We can
 155 calculate $Z_{\ell+1}^B$ in a symmetric manner (with the same encoder E_ℓ).

156 After the process of L FW layers, Z_L^A and Z_L^B are further cross-concatenated and fed to the matching
 157 estimator (in Fig. 2). It outputs a non-deterministic edge assignment \hat{m} . In the training phase, \hat{m}
 158 is input to an objective function, and the loss is minimized. In the prediction phase, the matching
 159 is obtained by binarizing \hat{m} . In this sense, matching estimation through a neural network can be
 160 considered as an approximation by relaxing the binary assignment space $\{0, 1\}^{N \times M}$ into a continuous
 161 assignment space $[0, 1]^{N \times M}$.

162 **Asymmetric variant with split batch normalization** WeaveNet is designed to be fully symmetric
 163 for S^A and S^B . Hence, it satisfies the equation $F(S^A, S^B) = F(S^B, S^A)^\top$. This condition ensures
 164 that the model architecture cannot distinguish the two sides A and B innately. This property is
 165 beneficial when mathematically fair treatment between A and B is desirable. However, when inputs
 166 from A and B are differently biased (e.g., the two sides have different trends of preference or the
 167 objective is asymmetric for A and B), this symmetric treatment degrades the performance. To
 168 eliminate the bias difference without losing the parameter-efficiency, we further propose to **a**) apply
 169 batch normalization independently for each stream (*split batch normalization*), and **b**) adding a
 170 side-identifiable code (e.g., 1 for A and 0 for B) to Z_0^A and Z_0^B as a $(D+1)$ -th element of the feature.
 171 We call this variant ‘‘asymmetric’’.

172 4.2 Relaxed continuous optimization for fair stable matching

173 Generally, a combinatorial optimization problem has discrete objective functions and conditions,
 174 which are not differentiable. To optimize the model in an end-to-end manner without inaccessible
 175 ground truth, we optimize the model by relaxing such discrete loss functions into continuous ones.

176 Assume we target to obtain a fair stable matching that has the minimum *SEq*, for example. Then, we
 177 have the following three loss functions.

178 \mathcal{L}_m conditions the binarization of \hat{m} to represent a matching.

179 \mathcal{L}_s conditions the matching to be stable.

180 \mathcal{L}_f minimizing the fairness cost *SEq* of the matching

181 The overall loss function is defined as

$$\mathcal{L}_{\text{fsm}}(\hat{m}) = \lambda_m \mathcal{L}_m + \frac{1}{2} \sum_{m \in \{\hat{m}^A, \hat{m}^B\}} (\lambda_s \mathcal{L}_s(m) + \lambda_f \mathcal{L}_f(m)), \quad (4)$$

182 where $\hat{m}^A = \text{softmax}(\hat{m})$ and $\hat{m}^B = \text{softmax}(\hat{m}^\top)$.

183 An important advantage of learning-based approximation is its flexibility. We can modify the above
 184 loss functions to easily obtain other variants. For example, removing \mathcal{L}_f in Eq. (4) leads to standard
 185 stable matching, and replacing \mathcal{L}_f with \mathcal{L}_b (which minimizes *Bal*) leads to balanced stable matching,
 186 as follows:

$$\mathcal{L}_{\text{sm}}(\hat{m}) = \lambda_m \mathcal{L}_m + \frac{1}{2} \sum_{m \in \{\hat{m}^A, \hat{m}^B\}} \lambda_s \mathcal{L}_s(m), \quad (5)$$

$$\mathcal{L}_{\text{bsm}}(\hat{m}) = \lambda_m \mathcal{L}_m + \frac{1}{2} \sum_{m \in \{\hat{m}^A, \hat{m}^B\}} (\lambda_s \mathcal{L}_s(m) + \lambda_b \mathcal{L}_b(m)). \quad (6)$$

187 **One-to-one matching constraint** \hat{m} can be safely converted into a binarized matching by column-
 188 wise or row-wise argmax operation when it is a symmetric doubly stochastic matrix (Li, 2019). To
 189 satisfy this condition, we defined \mathcal{L}_m with an average of the cosine distance as

$$\begin{aligned}\mathcal{L}_m(\hat{m}^A, \hat{m}^B) &= 1 - \frac{1}{2}(C(\hat{m}^A, \hat{m}^B) + C(\hat{m}^B, \hat{m}^A)), \\ C(\hat{m}^A, \hat{m}^B) &= \frac{1}{N} \sum_{i=0}^N \frac{\hat{m}_{i*}^A \cdot \hat{m}_{*i}^B}{\|\hat{m}_{i*}^A\|_2 \|\hat{m}_{*i}^B\|_2},\end{aligned}\tag{7}$$

190 where \hat{m}_{i*}^A means the i -th row of \hat{m}^A . This formulation binds \hat{m} to be a symmetric⁴ doubly stochastic
 191 matrix when $\mathcal{L}_m(\hat{m}^A, \hat{m}^B) = 0$. The advantage of this implementation against the original one in Li
 192 (2019) is described in B.1 with additional experimental results.

193 **Blocking pair suppression** As for L_s , we used the function proposed in Li (2019), which is

$$\begin{aligned}\mathcal{L}_s(\hat{m}; I) &= \sum_{(v,w) \in A \times B} g(a_v; b_w, \hat{m}) g(b_w; a_v, \hat{m}) \\ g(a_i; b_w, \hat{m}) &= \sum_{b_j \neq b_w} \hat{m}_{ij} \cdot \max(S_{iw}^A - S_{ij}^A, 0) \\ g(b_j; a_v, \hat{m}) &= \sum_{a_i \neq a_v} \hat{m}_{ji}^\top \cdot \max(S_{jv}^B - S_{ji}^B, 0),\end{aligned}\tag{8}$$

194 where $g(a_i; b_w, \hat{m})$ is a criterion known as ex-ante justified envy, which has a positive value when
 195 a_i prefers b_w more than any b_j in $\{b_j | j \neq w, \hat{m}_{ij} > 0\}$. This is the same for $g(b_j; a_v, \hat{m})$. Hence,
 196 $\{a_v, b_w\}$ becomes a (soft) blocking pair when both $g(a_v; b_w, \hat{m})$ and $g(b_w; a_v, \hat{m})$ are positive.

197 **Fairness measurements** $\mathcal{L}_f, \mathcal{L}_b$ minimize $Seq(m; I), Bal(m; I)$, respectively, and are defined as

$$\mathcal{L}_f(\hat{m}; I) = \frac{1}{N} |S(\hat{m}; A) - S(\hat{m}; B)| \quad \mathcal{L}_b(\hat{m}; I) = -\frac{1}{N} \min(S(\hat{m}; A), S(\hat{m}; B)),\tag{9}$$

198 where

$$S(\hat{m}; A) = \sum_{i=1}^N \sum_{j=1}^M \hat{m}_{ij} \cdot S_{ij}^A, \quad S(\hat{m}; B) = \sum_{j=1}^M \sum_{i=1}^N \hat{m}_{ij} \cdot S_{ji}^B.\tag{10}$$

199 5 Experiments

200 We evaluated WeaveNet with different sizes of N . First, with test samples of $N < 10$, we compared
 201 its performance with learning-based baselines and optimal solutions obtained by a brute-force search.
 202 Second, we compared WeaveNet with algorithmic baselines at $N = 20, 30$, where neither existing
 203 learning-based methods nor brute-force search work. We also demonstrated the generalization ability
 204 of WeaveNet under the mismatched training/test dataset distributions. Third, we demonstrated the
 205 performance of WeaveNet at $N = 100$. Note that we always assume $M = N$ hereafter.

206 **Sample generation protocol** In the experiments, we used the same method as Tziavelis *et al.*
 207 (2019) to generate synthetic datasets that draw preference lists from the following distributions.

208 **Uniform (U)** Each agent’s preference towards any matching candidate is totally random, defined by
 209 a uniform distribution $\mathcal{U}(0, 1)$ (larger value means prior in the preference list).

210 **Discrete (D)** Each agent has a preference of $\mathcal{U}(0.5, 1)$ towards a certain group of $\lfloor 0.4N \rfloor$ popular
 211 candidates, while $\mathcal{U}(0, 0.5)$ towards the rest.

212 **Gauss (G)** Each agent’s preference towards i -th candidate is defined by a Gaussian distribution
 213 $\mathcal{N}(i/N, 0.4)$.

214 **LibimSeTi (Lib)** Simulate real rating activity on the online dating service LibimSeTi (Brozovsky
 215 and Petricek, 2007) based on the 2D distribution of frequency of each rating pair (p_{ij}^A, p_{ji}^B) .

⁴Here a (possibly non-square) matrix \hat{m} ($N \geq M$) is symmetric if and only if $\hat{m}_{i*} = \hat{m}_{*i}$, ($0 < i \leq M$).

216 Choosing the above preference distributions for group A and B respectively, we obtained five different
 217 dataset settings, namely UU, DD, GG, UD, and Lib. We randomly generated 1,000 test samples and
 218 1,000 validation samples for each of the five distribution settings.

219 **Training protocol** We trained any learning-based models 200k total iterations at $N \leq 30$ and 300k
 220 at $N = 100$, with a batch size of 8. We randomly generated training samples at each iteration based
 221 on the distribution of each dataset and used the Adam optimizer (Kingma and Ba, 2015). We set
 222 learning rate 0.0001 and loss weights $\lambda_s = 0.7$, $\lambda_m = 1.0$, $\lambda_f = \lambda_b = 0.01$ based on a preliminary
 223 experiment (see A.4).

224 **Pseudo fairness costs for comparing learning-based results with algorithmic results** Note that
 225 for learning-based methods, there is a trade-off between fairness scores and stable matching rate.
 226 Hence they may violate the constraints of stable one-to-one matching and yield an SEq or Bal even
 227 lower than the ideal value. To compare the methods fairly with traditional algorithmic methods, we
 228 evaluate our methods using pseudo SEq ($pSEq$) and pseudo Bal ($pBal$) cost, in which the cost of
 229 violation cases is replaced by the worst result of the GS algorithm (prioritizing each side once and
 230 adopting the *worse* one).

231 5.1 Comparison with learning-based methods ($N = 3, 5, 7, 9$)

232 **Baselines and ablations** In this experiment, we show results obtained by following baselines and
 233 WeaveNet variants. **MLP** is the model proposed in Li (2019). **GIN** is the state-of-the-art GCN model
 234 proposed in Xu et al. (2019). We use each (normalized) preference list as a node feature and bipartite
 235 edges as the graph structure. After two graph-convolution calculations, as MLP, we destructed the
 236 node-wise embeddings and concatenated them into a single vector, which is fed to one Linear layer to
 237 output \hat{m} . **DBM** is the model in Gibbons et al. (2019). **SSWN** is the single-stream WeaveNet, which
 238 is equivalent to a DBM adopting the set-encoder of WeaveNet. **WN** is the standard WeaveNet.




Figure 5: Change of the success rates of stable matching (↑) according to N




Figure 6: Change of $pSEq$ - ideal scores (↓) according to N .




Figure 7: Change of $pBal$ - ideal scores (↓) according to N .

239 Fig. 5 shows the success rates of finding a stable matching, where we trained models to minimize
 240 Eq. (5), considering only the stable matching constraints. Since MLP and GIN have size-dependency,
 241 we trained the models independently for $N = 3, 5, 7, 9$. The other models were trained with
 242 $N = 10$ and tested on $N = 3, 5, 7, 9$. We maintained models with $L = 6$ layers (the model names
 243 are noted as XXX-6) to have a similar number of parameters with MLP for $N = 5$ (see A.5), while
 244 WN-18 is prepared to demonstrate the full performance (with the residual blocks).

245 MLP and GIN can hardly find stable matchings when $N \geq 5$. Note that the number of total cases for
 246 size N instances is estimated by $N!^{2(N-1)}$. Hence, when $N = 3$, there are only 1,296 cases at most,
 247 and the test set will fully overlap with the training set. In contrast, when $N = 5$, we have 4.3×10^{16}
 248 cases, and the overlap is negligible. Therefore, we can say that methods working only with $N = 3$,
 249 such as MLP and GIN, have little generalization ability.

250 DBM performs better than MLP but obviously worse than SSWN and WN. The performance gain of
 251 SSWN-6 over DBM-6 represents the advantage of the set-encoder. Similarly, the improvement of
 252 WN-6 over SSWN-6 shows the benefit of the two-stream architecture. Finally, that of WN-18 over
 253 WN-6 demonstrates the impact of stacked layers on the performance. Fig. 15 of the appendix shows

254 some additional baselines, including a performance of our L_m against the original one proposed in Li
 255 (2019).

256 Figs. 6 and 7 show $pSeq$ and $pBal$ (their difference from the ideal values⁵), respectively. XXX-18f/b
 257 are trained to minimize Eqs. (4) and (6), respectively⁶. We omitted MLP and GIN due to their poor
 258 performance in Fig. 5. In the results, both SSWN and WN largely outperformed DBM, which again
 259 proved the advantage of the set-encoder. WN performed better than SSWN for larger N , owing to the
 260 parameter efficiency of the two-stream architecture. Note that the performance gain of XXX-18f/b
 261 from XXX-18 proved the flexibility of general learning-based methods for customized objective
 262 functions.

263 5.2 Comparison with algorithmic methods ($N = 20, 30$)

264 As the algorithmic methods, we prepared four baselines. **GS** is the *better* result of applying the GS
 265 algorithm to prioritize each side once, which runs in $O(N^2)$. **PolyMin** minimizes some alternative
 266 fairness costs (the regret and egalitarian costs, which can be solved in $O(N^2)$ and $O(N^3)$, respectively
 267 (Gusfield, 1987; Irving *et al.*, 1987; Feder, 1992)). **DACC** by Dworzak (2016) is an approximate
 268 algorithm that runs in $O(N^4)$. **PowerBalance** is the state-of-the-art method that runs in $O(N^2)$.

269 **WN-60f/b(20/30)** is WeaveNet with $L = 60$ layers trained with samples of $N = 20$ and $N = 30$.
 270 Note that we used the asymmetric variant for UD and Lib. Moreover, we do not involve any traditional
 271 learning-based methods in this part since they scored clear performance drops with increasing N (see
 272 Fig. 5) and the problem size of $N = 20, 30$ is clearly beyond their capabilities, but an ablation with
 273 WeaveNet variants is reported in B.2.

Table 1: Average Seq (\downarrow) and success rate of stable matching (\uparrow). Bold and underlined scores shows the **best** and **second best** ones, respectively.

Agents ($N \times M$) Datasets (Dist. Type)	20 × 20					30 × 30				
	UU	DD	GG	UD	Lib	UU	DD	GG	UD	Lib
GS	41.89	18.81	19.52	70.97	19.66	94.03	43.46	36.56	163.77	39.78
PolyMin	19.93	11.83	20.57	87.08	18.47	35.52	21.21	37.37	209.62	31.85
DACC	24.34	20.13	23.07	101.75	20.40	40.87	34.35	40.59	240.48	33.88
Power Balance	16.28	8.93	17.07	<u>71.09</u>	15.40	<u>18.45</u>	<u>11.05</u>	27.22	<u>163.90</u>	21.57
WN-60f(20) ($pSeq$)	<u>12.23</u>	6.37	15.50	71.31	<u>14.59</u>	25.21	11.38	29.36	172.63	23.53
Stably Matched (%)	98.90	99.50	99.40	99.60	99.30	94.60	97.30	95.70	91.30	97.70
WN-60f(30) ($pSeq$)	12.16	<u>6.53</u>	<u>15.56</u>	71.34	14.53	18.30	10.52	<u>27.39</u>	170.35	<u>22.17</u>
Stably Matched (%)	99.10	99.40	99.40	99.50	99.80	98.10	99.00	98.00	93.90	98.60

Table 2: Average Bal (\downarrow) and success rate of stable matching (\uparrow).

Agents ($N \times M$) Datasets (Dist. Type)	20 × 20					30 × 30				
	UU	DD	GG	UD	Lib	UU	DD	GG	UD	Lib
GS	89.14	146.16	108.36	140.53	68.62	184.05	322.05	225.49	312.12	137.59
PolyMin	74.19	140.99	108.04	145.28	66.94	144.48	306.28	224.13	324.54	130.79
DACC	78.49	146.71	110.06	151.34	68.75	150.71	316.18	227.52	337.43	133.59
Power Balance	73.28	140.12	106.92	<u>140.55</u>	65.89	138.04	<u>302.30</u>	220.26	312.12	126.96
WN-60b(20) ($pBal$)	71.89	<u>138.79</u>	106.20	140.84	<u>65.85</u>	141.49	302.73	<u>221.92</u>	317.60	130.58
Stably Matched (%)	98.50	98.80	99.50	99.70	98.80	96.10	96.70	95.00	88.90	93.80
WN-60b(30) ($pBal$)	<u>72.33</u>	138.75	<u>106.65</u>	140.79	65.84	<u>140.40</u>	301.59	223.02	313.59	<u>127.93</u>
Stably Matched (%)	98.00	99.10	98.60	99.80	99.10	97.00	98.60	93.70	98.80	98.00

274 We show the results in Tables 1 and 2. When $N = 20$, except for UD, the proposed method constantly
 275 performed better than any algorithmic methods for both Seq and Bal . When $N = 30$, they are
 276 comparative. For UD, GS performed even better than PowerBalance. That means that the ideal
 277 solution constantly prioritizes one side (a kind of the strongest bias). Since we designed the WeaveNet
 278 architecture to treat the sides evenly, this is the most challenging situation for WeaveNet. Nonetheless,

⁵1.362, 2.534, 3.746, 4.694 in Seq and 2.406, 6.478, 11.956, 18.706 in Bal for $N = 3, 5, 7, 9$.

⁶We early-stopped the training for DBM-18f/b at 80k due to a sudden overfit after the epoch.

279 the proposed split batch normalization (with the side-identifiable code) achieved similar performance
 280 to GS and PowerBalance. We show the performance drop with the fully symmetric version in [B.2](#)
 281 of the appendix, which is also interesting from the ethical viewpoint. It is noteworthy that the
 282 model trained with $N = 20$ performs well even with $N = 30$, which indicates that the method has
 283 generalizability for size difference.

284 **Generalization ability for different distributions** A learning-based method should have a certain
 285 generalizability for input distribution shifts. To test the ability, we evaluated the performance of
 286 models trained with UU, DD, and GG on test sets of different distributions.

Table 3: The generalizability of WeaveNet (trained/tested with $N = 30$).

train		test			Avg.
WN-60f		UU	DD	GG	
UU	$pSEq$	18.30	25.81	29.09	<u>21.10</u>
	Stably Matched (%)	98.10	94.90	93.60	95.53
DD	$pSEq$	171.27	10.52	77.36	86.38
	Stably Matched (%)	2.80	99.00	0.10	33.97
GG	$pSEq$	21.38	12.85	27.39	20.54
	Stably Matched (%)	97.30	98.10	98.00	97.80

Table 4: Average SEq (\downarrow) and Bal (\downarrow) at $N = 100$.

$100 \times 100, UU$	SEq	Bal
GS	1259.39	1709.53
PolyMin	153.35	952.85
DACC	194.65	988.02
Power Balance	49.41	909.73
WN-80f/b+Hungarian		
$pSEq/pBal$	257.99	1145.36
SEq/Bal	68.36	919.75
Stably Matched (%)	89.4	80.8

287 Table [3](#) shows the results. Remarkably, there is a contrast between the model trained with DD and the
 288 others. The model with DD could hardly satisfy the one-to-one stable matching constraint when tested
 289 on UU/GG, and resulted in poor $pSEq$ scores. In contrast, the model with GG achieved satisfying
 290 $pSEq$ scores on UU/DD. Since GG generates preference lists based on a common preference score
 291 (i/N for i -th agent) with noise, agents in GG tend to have similar preference lists (i.e., hard to assign
 292 optimally). A model trained with such hard samples works well even for the test samples drawn from
 293 other distribution. UU has also performed well owing to its non-biased sampling strategy. On the
 294 other hand, DD worst performed due to its highly biased generation strategy. From these results, we
 295 confirmed that WeaveNet has certain robustness in the distribution shift as long as training samples
 296 are competitive enough.

297 5.3 Demonstration with $N = 100$

298 We further demonstrate the capability of WeaveNet under a larger size of problem instances, $N = 100$.
 299 In this case, we found that WN-80f and WN-80b failed to yield one-to-one matchings for 13.4%
 300 and 19.8%, respectively (see the Table [9](#) in [B.2](#) for details). To compensate for this problem, we
 301 applied the Hungarian algorithm ([Kuhn, 1955](#)) to surely binarize \hat{m} into a one-to-one matching.
 302 Table [4](#) shows WeaveNet’s relatively good SEq and Bal scores. Even with the help of the Hungarian
 303 algorithm, they were strongly penalized in $pSEq$ and $pBal$ due to the poor stable matching rate. In
 304 other words, we can potentially fill the large gap by better constraining the output.

305 Since this work is just a pilot study toward a practical differential assignment solver, there is still a lot
 306 of space for improvement. The proposed test protocol with stable matching will facilitate it since we
 307 can freely adjust the difficulty of the problem to develop and enhance the methods continuously.

308 6 Conclusion

309 This paper proposed a novel differential assignment solver, *WeaveNet*, and an evaluation protocol on
 310 two *strongly NP-hard* variants of stable matching. In the experiments, we demonstrated the advantage
 311 of *set encoder* and the two-stream architecture of WeaveNet against the other learning-based methods.
 312 These techniques also achieved a better performance than the state-of-the-art algorithmic method
 313 when $N = 20$ and a comparative performance when $N = 30$. Furthermore, the asymmetric variants,
 314 *split batch normalization* with the side-identifiable code, enabled the method to work even with
 315 the strongly biased dataset of UD. We also confirmed that the proposed method does not work at
 316 $N = 100$, which will be an immediate task for this new field of differential assignment solver. We
 317 hope that this work becomes a starting point to open a new vista for real-world assignment problems.

318 **References**

- 319 Ravindra K Ahuja, Arvind Kumar, Krishna C Jha, and James B Orlin. Exact and heuristic algorithms
320 for the weapon-target assignment problem. *Operations research*, 55(6):1136–1146, 2007.
- 321 Lukas Brozovsky and Vaclav Petricek. Recommender system for online dating service. In *Proceedings*
322 *of Conference Znalosti 2007*, Ostrava, 2007.
- 323 Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation
324 using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern*
325 *Recognition (CVPR)*, July 2017.
- 326 Piotr Dworzczak. Deferred acceptance with compensation chains. In *Proceedings of the 2016 ACM*
327 *Conference on Economics and Computation*, pages 65–66, 2016.
- 328 Patrick Emami, Panos M. Pardalos, Lily Elefteriadou, and Sanjay Ranka. Machine learning methods
329 for data association in multi-object tracking. *ACM Computing Surveys*, 53(4), 2020.
- 330 Tomás Feder. A new fixed point approach for stable networks and stable marriages. *Journal of*
331 *Computer and System Sciences*, 45(2):233–284, 1992.
- 332 Tomás Feder. *Stable networks and product graphs*. American Mathematical Society, 1995.
- 333 David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American*
334 *Mathematical Monthly*, 69(1):9–15, 1962.
- 335 Daniel Gibbons, Cheng-Chew Lim, and Peng Shi. Deep learning for bipartite assignment problems.
336 In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2318–
337 2325, 2019.
- 338 Sushmita Gupta, Sanjukta Roy, Saket Saurabh, and Meirav Zehavi. Balanced stable marriage: How
339 close is close enough? In *Workshop on Algorithms and Data Structures*, pages 423–437, 2019.
- 340 Dan Gusfield and Robert W Irving. *The stable marriage problem: structure and algorithms*. MIT
341 press, 1989.
- 342 Dan Gusfield. Three fast algorithms for four problems in stable marriage. *SIAM Journal on*
343 *Computing*, 16(1):111–128, 1987.
- 344 Robert W Irving, Paul Leather, and Dan Gusfield. An efficient algorithm for the “optimal” stable
345 marriage. *Journal of the ACM*, 34(3):532–543, 1987.
- 346 Kazuo Iwama, Shuichi Miyazaki, and Hiroki Yanagisawa. Approximation algorithms for the sex-
347 equal stable marriage problem. *ACM Transaction on Algorithms*, 7(1), 2010.
- 348 Akiko Kato. Complexity of the sex-equal stable marriage problem. *Japan Journal of Industrial and*
349 *Applied Mathematics*, 10(1):1, 1993.
- 350 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of*
351 *International Conference on Learning Representations*, 2015.
- 352 Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics*
353 *quarterly*, 2(1-2):83–97, 1955.
- 354 Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for
355 semi-supervised learning. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages
356 3538–3545, 2018.
- 357 Shira Li. *Deep Learning for Two-Sided Matching Markets*. Bachelor’s thesis, Harvard University,
358 2019.
- 359 Eric McDermid and Robert W Irving. Sex-equal stable matchings: Complexity and exact algorithms.
360 *Algorithmica*, 68(3):545–570, 2014.
- 361 Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node
362 classification. In *Proceedings of International Conference on Learning Representations*, 2020.

- 363 Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets
364 for 3D classification and segmentation. In *Proceedings of IEEE Conference on Computer Vision
365 and Pattern Recognition*, July 2017.
- 366 Nikolaos Tziavelis, Ioannis Giannakopoulos, Katerina Doka, Nectarios Koziris, and Panagiotis Karras.
367 Equitable stable matchings in quadratic time. In *Proceedings of Advances in Neural Information
368 Processing Systems*, pages 457–467, 2019.
- 369 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
370 networks? In *Proceedings of International Conference on Learning Representations*, 2019.
- 371 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and
372 Alexander J Smola. Deep sets. In *Proceedings of Advances in Neural Information Processing
373 Systems*, volume 30, pages 3391–3401, 2017.

374 **Checklist**

375 1. For all authors...

- 376 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contri-
377 butions and scope? [Yes] The experimental results in Section 5 correspond to the main claims,
378 which are summarized as the contribution list in Section 1.
- 379 (b) Did you describe the limitations of your work? [Yes] Our method works best among any
380 baselines when $N \leq 20$, comparative to the state-of-the-art algorithmic baseline when $N = 30$,
381 but poorly when $N = 100$. See Section 5.3.
- 382 (c) Did you discuss any potential negative societal impacts of your work? [Yes] We briefly
383 discussed the fairness/unfairness achieved by our method in the paragraph “Asymmetric variant
384 with split batch normalization” in Section 4.1. Namely, the asymmetric variant can better
385 optimize the objective than the symmetric one (which ensures mathematically equal treatment
386 for sides A and B) but harms the equal treatment of the two sides.
- 387 (d) Have you read the ethics review guidelines and ensured that your paper conforms to them?
388 [Yes]

389 2. If you are including theoretical results...

- 390 (a) Did you state the full set of assumptions of all theoretical results? [N/A] We have no main
391 theoretical results. The only theoretical discussion is for the computational cost of WeaveNet in
392 A.2, whose assumption is the network shape explained in the paper.
- 393 (b) Did you include complete proofs of all theoretical results? [N/A] We have no main theoretical
394 results. The only theoretical discussion is for the computational cost of WeaveNet in A.2, where
395 we provided enough detailed explanation as a proof.

396 3. If you ran experiments...

- 397 (a) Did you include the code, data, and instructions needed to reproduce the main experimental
398 results (either in the supplemental material or as a URL)? [Yes] We have included it in the
399 supplemental material.
- 400 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were
401 chosen)? [Yes] We described the training details in Section 5 “Training protocol” and Section
402 A.5.
- 403 (c) Did you report error bars (e.g., with respect to the random seed after running experiments
404 multiple times)? [Yes] We presented figures with the error bars in Section A.4 of the appendix,
405 which demonstrated the stable behavior of the proposed method against the random seed. For
406 the other part, we have multiple settings, and we observed a stable trend in the results.
- 407 (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs,
408 internal cluster, or cloud provider)? [Yes] We described it in the appendix, Section C.

409 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- 410 (a) If your work uses existing assets, did you cite the creators? [Yes] See Section C in the appendix.
- 411 (b) Did you mention the license of the assets? [Yes] See Section C.
- 412 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We
413 provides a code to random-generate the problem instances, which contains no personal or any
414 other sensitive information, but only the distribution parameters of the LibimSeti dataset.
- 415 (d) Did you discuss whether and how consent was obtained from people whose data you’re
416 using/curating? [Yes] See Section C.
- 417 (e) Did you discuss whether the data you are using/curating contains personally identifiable
418 information or offensive content? [Yes] See Section C.

419 5. If you used crowdsourcing or conducted research with human subjects...

- 420 (a) Did you include the full text of instructions given to participants and screenshots, if applicable?
421 [N/A] We used neither crowd-sourcing nor human subjects.
- 422 (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB)
423 approvals, if applicable? [N/A]
- 424 (c) Did you include the estimated hourly wage paid to participants and the total amount spent on
425 participant compensation? [N/A]