Bisected graph matching improves automated pairing of bilaterally homologous neurons from connectomes

Benjamin D. Pedigo[1*], Michael Winding[2], Carey E. Priebe[1], Joshua T. Vogelstein[1]
[1]Johns Hopkins University, [2]University of Cambridge, *corresponding author: bpedigo@jhu.edu

Abstract. Graph matching algorithms attempt to find the best correspondence between the nodes of two networks. These techniques have previously been used to match individual neurons in nanoscale connectomes – in particular, to find pairings of neurons across hemispheres. However, since graph matching techniques deal specifically with two isolated networks, they have only utilized the ipsilateral (same hemisphere) subgraphs when performing the matching. Here, we present a modification to a state-of-the-art graph matching algorithm which allows it to solve what we call the bisected graph matching problem. This modification allows us to also use connections between the brain hemispheres when predicting neuron pairs. Via simulations and real connectome datasets we show that when edge correlation is present between the contralateral (between hemisphere) subgraphs, this approach improves matching accuracy. We expect that our proposed method will improve future endeavors to accurately match neurons across hemispheres in connectomes, and be useful in other applications where the bisected graph matching problem arises.

1 Introduction

Graph matching is a widely used optimization technique whereby one can find a matching between the nodes in one network and those in another. Solving the graph matching problem yields a matching (i.e. the correspondence between the nodes of the two networks) which minimizes edge disagreements between the two networks. The graph matching problem has found uses in fields as disparate as computer vision [1], biometrics [1], social networks [2], and natural language processing [3], to name just a few.

Most important for this work is the use of graph matching techniques to find homologous neuron pairs across brain hemispheres [4, 5]. Connectomes–maps of neural connectivity–can naturally be represented by networks, wherein a node represents a neuron and an edge represents synapses from one neuron to another [6, 7]. Previous works used graph matching techniques to find neuron pairings between hemispheres based on the observed connectivity [4, 5]. The graph matching problem by its very formulation is concerned with two separate networks; as such, previous applications of graph matching to find homologous neuron pairings across hemispheres have considered one network to be the set of nodes and edges within the left hemisphere, and the other network to be defined likewise for the right hemisphere. In other words, they have only considered the **ipsilateral** connections which connect within a brain hemisphere, and ignored the **contralateral** connections which cross brain hemispheres. Contralateral connections are quite common in connectomes studied thus far: in subset of the larval *Drosophila melanogaster* (vinegar fly) brain connectome published thus far, an edge picked at random from the network has about a 35% chance of being a contralateral connection [8–28]. It is natural to wonder, then, whether these connections can be used successfully to improve automated neuron pairing.

Here, we show that rather than ignoring the contralateral connections for the purposes of predicting neuron pairs, they can be explicitly included in the optimization by generalizing graph matching to a single network which has been split into two parts. We demonstrate via simulation that when sufficient edge correlations exist between the contralateral subgraphs, our proposed method provides an improvement in matching accuracy. We then show that this methodology indeed improves matching accuracy in our motivating example of a bilateral brain by comparing our algorithm to traditional graph matching on five connectome datasets.

2 From graph matching to bisected graph matching

First, consider the graph matching problem from the perspective of attempting to predict neuron pairs between brain hemispheres (thus adopting the terminology of left/right, etc.), though the techniques described here could be applied more generally. For now, consider the case where both hemispherse have exactly n nodes, though methods for graph matching with an unequal number of nodes have been described [29]. Let A_{LL} be the $n \times n$ adjacency matrix for the subgraph of connections from left hemisphere to left hemisphere neurons, and let A_{RR} be defined likewise for the right hemisphere. The graph matching problem can then be written as

$$\min_{P \in \mathcal{P}} \|A_{LL} - PA_{RR}P^T\|_F^2, \tag{2.1}$$

where the set of permutation matrices on n nodes is denoted by $\mathcal{P}$. This objective function measures the number of edge disagreements for an unweighted network, or the norm of the weight disagreements for a weighted network. By minimizing this quantity over the set of permutations, one can find a matching between the networks under which the edge structure we observe appears similar.

We are interested in some similar measure which also includes the fact that we want the contralateral connections, under some matching, to appear similar. To formalize this, let A_{LR} be the adjacency matrix for the subgraph of connections from left hemisphere to right hemisphere neurons, and let A_{RL} be defined likewise for the connections from the right to the left. We add a term to the graph matching objective function which measures the disagreement between the contralateral subgraphs under some permutation of the nodes of the right hemisphere:

$$\min_{P \in \mathcal{P}} \|A_{LL} - PA_{RR}P^T\|_F^2 + \|A_{LR}P^T - PA_{RL}\|_F^2. \tag{2.2}$$

We call the problem in Equation 2.2 the **bisected graph matching problem** (illustrated in Figure 1). With this formulation, the graph matching problem can be seen as a special case of the bisected graph matching problem, since the objective function in Equation 2.2 reduces to that of Equation 2.1 in the special case where A_{LR} and A_{RL} are both the zero matrix. Note that this problem is also distinct from the multiplex graph matching problem described in Pantazis et al. [30], as the contralateral subgraphs require only a permutation of their rows *or* their columns (not both) to maintain the correct structure of the adjacency matrix.

Given this notion of what it means to find a good matching between the hemispheres, a state-of-the-art graph matching technique can be adapted to solve this new bisected graph matching problem. Introduced in Vogelstein et al. [31], the Fast Approximate Quadradic (FAQ) algorithm first relaxes the (discrete) graph matching problem to the relaxed graph matching problem, allowing the tools of continuous optimization to be used. FAQ then uses the Frank-Wolfe method [32] to attempt to minimize Equation 2.1.

The Frank-Wolfe method finds a search direction by minimizing a first-order Taylor series of the objective function, requiring that we compute the objective function's gradient with respect to its argument, P. The gradient of Equation 2.2 with respect to P (see Section 9 for details) is

$$\nabla f(P) = -\left(A_{LL}PA_{RR}^T + A_{LL}^T PA_{RR} + A_{LR}P^T A_{RL}^T + A_{RL}^T P^T A_{LR}\right) \tag{2.3}$$

By substituting this new gradient calculation into the FAQ algorithm and keeping the rest of the algorithm the same, FAQ can be adapted to solve the bisected graph matching problem. We provide a full description of the modified FAQ algorithm in Section 6. For the remainder of the paper, graph matching (GM) refers to the use of FAQ, while bisected graph matching (BGM) refers to the use of FAQ as modified above (Section 6).

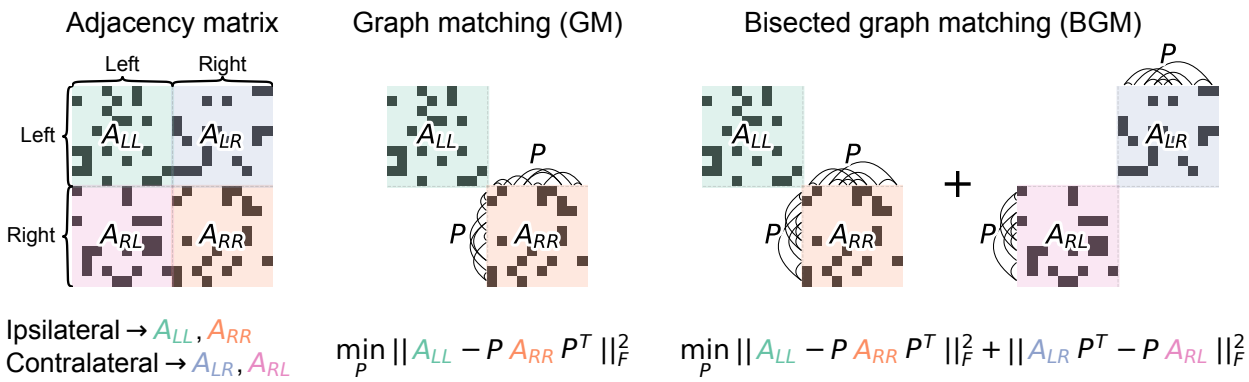

Figure 1: Schematic describing graph matching (GM) and bisected graph matching (BGM). Both aim to find a matching (which can be represented by a permutation, P) of the nodes of one hemisphere with respect to the other. GM does so by minimizing the norm of the edge differences between ipsilateral (A_{LL} and A_{RR}) subgraphs under some matching. BGM aims to jointly minimize the norm of edge differences between ipsilateral *and* contralateral (A_{LR} and A_{RL}) subgraphs under the same matching applied to both. Note that for the contralateral subgraphs, a permutation of the right hemisphere nodes amounts to permuting the columns (for A_{LR}) *or* the rows (for A_{RL}), but not both.

3 Simulation

Here, we demonstrate that this approach improves matching accuracy in simulated datasets when there is sufficient correlation in the contralateral subgraphs. To understand how this correlation affects the usefulness of bisected graph matching, we created simulated data using the correlated Erdős-Rényi model ($CorrER$). The correlated Erdős-Rényi model is a special case of the correlated stochastic block model introduced in Lyzinski et al. [33]. Briefly, a pair of networks is distributed $CorrER(n, p, \rho)$ if both networks marginally are distributed as Erdős-Rényi models [34, 35] with n nodes and global connection probability p, but the edges of the two networks have Pearson correlation ρ. Here, we use the version of this model for a directed network to more closely resemble nanoscale connectome data, which has directed edges. We used the correlated Erdős-Rényi model to construct a simulation of a "bilateral" network as follows:

1. The ipsilateral subgraphs were sampled from a correlated Erdős-Rényi model ($CorrER$):

$$A_{LL}, A_{RR} \sim CorrER(10, 0.3, 0.8)$$

2. Independently from the ipsilateral networks, the contralateral subgraphs were sampled from a correlated Erdős-Rényi model:

$$A_{LR}, A_{RL} \sim CorrER(10, 0.2, \rho_{contra})$$

3. The full network was defined as

$$A = \begin{bmatrix} A_{LL} & A_{LR} \\ A_{RL} & A_{RR} \end{bmatrix}$$

4. To simulate an unknown correspondence between the nodes of the left and right hemispheres, we applied a random permutation (P_{rand}) to the nodes of the "right hemisphere" in each sampled network:

$$A_{input} = \begin{bmatrix} I_n & 0 \\ 0 & P_{rand} \end{bmatrix} A \begin{bmatrix} I_n & 0 \\ 0 & P_{rand} \end{bmatrix}^T = \begin{bmatrix} A_{LL} & A_{LR}P_{rand}^T \\ P_{rand}A_{RL} & P_{rand}A_{RR}P_{rand}^T \end{bmatrix}$$

A_{input} is the network which was input to the matching algorithms, obscuring the true matching from the true alignment in order to evaluate their performance.

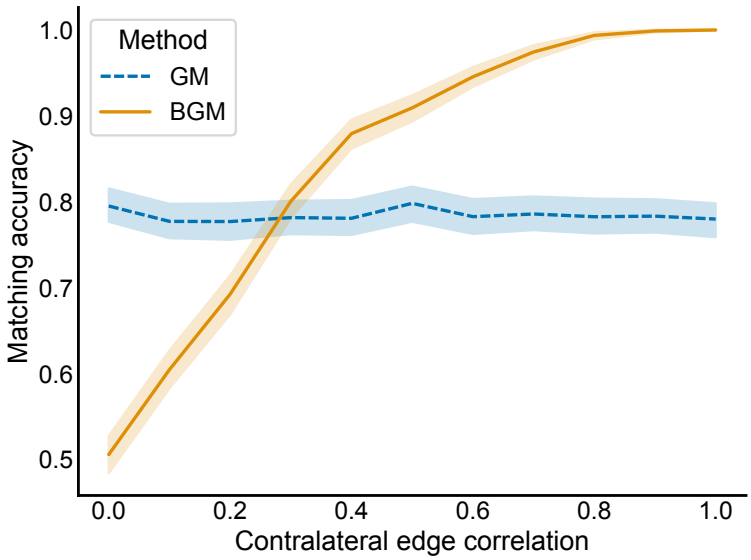

Figure 2: Performance of graph matching (GM) using the FAQ algorithm [31] and bisected graph matching (BGM) (this work) on a simulated dataset constructed such that the ipsilateral and contralateral connections both come from correlated Erdős-Rényi models (Section 3). Each network had 10 nodes per side, the ipsilateral connection density was 0.3, the ipsilateral edge correlation was 0.8, and the contralateral connection density was 0.2. We varied the contralateral edge correlation from 0 to 1, and for each value, we simulated 1000 networks and ran both algorithms on the same data with the same initialization. Lines show the mean matching accuracy, shaded regions show 95% confidence intervals. As the correlation in the contralateral connections increases, including them in the optimization becomes more helpful.

We varied the value of ρ_{contra} from zero (A_{LR} and A_{RL} have no correlation, and thus are not helpful for matching) to one (A_{LR} and A_{RL} are isomorphic, providing extremely helpful information for matching). For each value of ρ_{contra}, we simulated 1,000 networks. For each network, we ran the graph matching (GM) and bisected graph matching (BGM) algorithms to attempt to uncover the correct permutation which would realign the left and right hemispheres. For both algorithms, we used default parameters, and one initialization for each algorithm. For each run of each algorithm, we examined the matching accuracy, which is the percentage of nodes correctly matched.

Figure 2 shows the match accuracy for both algorithms as a function of ρ_{contra}. For low values of ρ_{contra}, using bisected graph matching actually degrades performance. When $\rho_{contra} = 0$, the match accuracy drops by ~29%. This is unsurprising, as in this case the contralateral subgraphs are effectively noise with respect to the correct matching between the left and right. However, as ρ_{contra} increases, bisected graph matching eventually outperforms graph matching. For this simulation, when ρ_{contra} is greater than 0.4, the accuracy for bisected graph matching is higher (by more than ~22% when $\rho_{contra} \geq 0.9$).

Whether or not bisected graph matching improves accuracy is determined by many factors, including the correlation in contralateral edge structure studied here in this simple simulation. We next sought to see whether this bisected graph matching would be helpful in our motivating example of matching neurons between two sides of a nervous system.

4 Connectomes

We examined the performance of both graph matching algorithms on a set of real connectome datasets. To ensure we could evaluate the performance of both algorithms, we restricted our analysis to connectomes for which pairings of individual neurons between sides of the nervous system were already known. We studied the (chemical) connectomes of both a hermaphrodite and a male *Caenorhabditis elegans* worm [37], the pharynges of two *Pristionchus pacificus* worms [36], and a subset of a larval *Drosophila melanogaster* [8–28]. For all three datasets, neuron pairings across sides of the nervous system are not complete - indeed, some neurons appear only on one side of the organism or exactly in the center [37]. Thus, we restricted our analysis to the subset of neurons which were present as a bilaterally homologous pair and for which this pairing was known. We then ensured that the remaining set of nodes was fully (weakly) connected for each dataset, removing nodes not part of the largest connected component. Table 1 shows the number of nodes and edges remaining for each of the connectome datasets considered here. We treat each network as weighted (by synapse count) and directed (since the direction of chemical synapses is known).

For each connectome, we predicted each neuron's pair on the other side of the nervous system by applying either the graph matching or bisected graph matching algorithms. We ran 50 initializations from the barycenter (an uninformative initialization, see Section 6). For each initialization, we ran both algorithms with default parameters, and measured the matching accuracy with respect to the known pairing of neurons.

We observed that for all the connectomes studied here, bisected graph matching improves matching performance (Figure 3), sometimes dramatically so. For all five connectomes, the match ratio for the bisected graph matching algorithm was significantly higher ($p < 0.0005$ for each dataset, two-sided Wilcoxon rank sum tests). Matching accuracy increased by ~44% and ~20% for the two *P. pacificus* samples, ~29% for the hermaphrodite *C. elegans*, ~15% for the male *C. elegans*, and ~22% for the maggot subset, respectively. These results demonstrate the practical utility of our proposed algorithm for improving bilaterally homologous pair prediction in connectomes.

5 Discussion

5.1 Summary

We proposed a simple generalization of the graph matching problem, which incorporates any connections *between* the two sets of nodes being matched. We then showed how this problem could be solved by using a new objective function in the framework of a state-of-the-art graph matching algorithm, FAQ [31]. In simulations, we saw that as the strength of the correlation between the contralateral subgraphs increases, these connections become more useful to include in the matching process. By running both graph matching and bisected graph matching on five connectome datasets, we provided compelling evidence that for practical purposes in neuroscience, including these contralateral connections in the optimization is beneficial.

5.2 Limitations

As we showed in simulation in Section 3, bisected graph matching is only likely to improve matching accuracy in connectomes when there is sufficient correlation between the contralateral subgraphs. For a new organism (or possibly even just a new sample), this will not be known in practice–one would need to know the node matching to compute such an edge correlation, and if that information was known, there would be little reason to match the networks in the first place. Domain knowledge as to the nature of the contralateral connections in an organism's brain may be important when choosing whether to include them in the matching as described in this work.

We demonstrated how to generalize only the original FAQ algorithm to the bisected graph matching problem; however, many extensions to FAQ exist. Fishkind et al. [29] introduced "seeds" to the algorithm, which allow a practitioner to incorporate a known partial matching of nodes into the optimization.

Dataset	Citation(s)	# nodes	# edges	% contra. synapses
P. pacificus pharynx 1	[36]	18	35	30
P. pacificus pharynx 2	[36]	22	33	32
C. elegans hermaphrodite	[37]	286	2,838	41
C. elegans male	[37]	360	2,482	37
D. melanogaster larva subset	[8–28]	1,240	32,564	31

Table 1: Summary of the connectome datasets studied here. The number of nodes, number of edges edges, and percentage of contralateral synapses are with respect to the datasets after processing to select fully connected networks composed of neurons who have bilateral pairs, as described in Section 4. The number of nodes per hemisphere and (and the number of known pairs) is always half the total number of nodes since only paired neurons are considered.

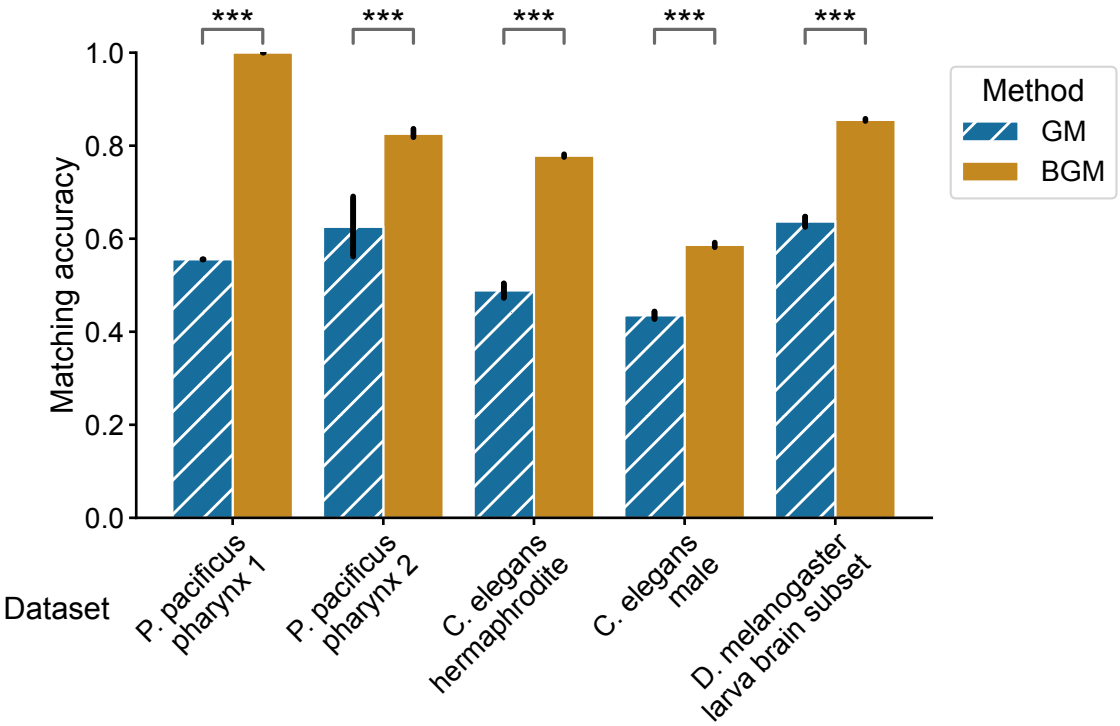

Figure 3: Performance of graph matching (GM) and bisected graph matching (BGM) on five bilateral connectome datasets. For each comparison, we performed 50 initializations, and from each initialization we ran both the graph matching and the bisected graph matching algorithms. The mean match accuracy is shown for both methods on each dataset. Error bars show 95% confidence intervals for the mean. For every dataset, bisected graph matching provided a performance improvement over graph matching which ignores contralateral connections. $***$ indicates a significant difference where $p < 0.0005$ (two-sided Wilcoxon signed rank test).

Pantazis et al. [30] introduced a version of FAQ for multilayer networks, which may be helpful when considering different kinds of synaptic connections such as gap junction vs. chemical [37] or axo-axonic vs. axo-dendritic. Further, other approaches to matching neurons in neuroscience do not use connectivity at all: Costa et al. [38] introduced an algorithm for matching neurons on the basis of morphology which has been widely used on connectomic reconstructions. In practice, it is likely that the best neuron pairings will be achieved by a joint optimization which considers morphology, multilayer edges, seeds, and the contralateral connections as proposed in this work. We did not explore this possibility here, but it remains an intriguing future pursuit.

5.3 Outlook

As more connectomes are mapped both from new organisms or for more individuals of the same species, the tools provided here will accelerate the process of finding correct pairings of neurons between the brain hemispheres. These neuron pairings appear to be a fundamental property of the invertebrate nervous systems studied in connectomics thus far. Finding these neuron pairs is important for understanding the stereotypy in an organism (i.e. how similar is the connectivity of the left and the right) [37, 39–41]. Additionally, these neuron pairs can be useful for statistical approaches which leverage a one-to-one correspondence of nodes across networks [42, 43]. We also note that the bisected graph matching problem (and the analogous version of the quadradic assignment problem, which is equivalent to the graph matching problem up to a sign change [31]) may arise in other settings where one wishes to match nodes in a single graph which can be split into two parts and some level of symmetry exists between them.

6 Methods

The FAQ algorithm relaxes the (discontinuous) search for a permutation matrix to the convex hull of this set, the set of doubly stochastic matrices, $\mathcal{D}$ [31]. FAQ then proceeds by using the Frank-Wolfe algorithm to iteratively update its search for a suitable doubly stochastic matrix which maps one adjacency matrix to another. This allows the tools of continuous optimization to be used. After a suitable doubly stochastic solution has been found, FAQ then projects back onto the set of permutation matrices by solving the linear assignment problem. Below, we detail the BGM-via-Frank-Wolfe algorithm (referred to simply as BGM in the text), which simply adapts this procedure by replacing the objective function and its gradient to solve the bisected graph matching problem as described in Section 2. We refer the interested reader to Vogelstein et al. [31] for further details on the original algorithm, and to Fishkind et al. [29] for many interesting extensions. We also note that implementations of the FAQ algorithm are available in SciPy [44] and graspologic [45].

Algorithm 1 Bisected Graph Matching (BGM) via Frank-Wolfe. To recover the FAQ algorithm [31] (GM), simply set A_{LR}, A_{RL} to the zero matrix.

Input: Adjacency matrices for each of the four subgraphs: $A_{LL}, A_{RR}, A_{LR}, A_{RL} \in \mathbb{R}^{n\times n}$.

Initialize: $P_{(0)} \in \mathcal{D}$, barycenter ($P_{(0)} = \frac{1}{n}\mathbf{1}_n \times \mathbf{1}_n^\intercal$) unless otherwise specified

for $i = 1, 2, 3, \ldots$ (while stopping criterion not met) **do**

1. Compute $\nabla f(P_{(i)}) = -\left(A_{LL}P_{(i)}A_{RR}^T + A_{LL}^T P_{(i)} A_{RR} + A_{LR}P_{(i)}^T A_{RL} + A_{RL}^T P_{(i)}^T A_{LR}\right)$
2. Compute $Q_{(i)} \in \operatorname{argmin} \operatorname{tr}\left(Q^\intercal \nabla f(P_{(i)})\right)$ over $Q \in \mathcal{D}$ via linear assignment problem solver, e.g. Hungarian algorithm [46]
3. Compute step size $\alpha^{(i)} \in f(\alpha P_{(i)} + (1-\alpha)Q_{(i)})$, for $\alpha \in [0, 1]$
4. Set $P_{(i+1)} = \alpha P_{(i)} + (1-\alpha)Q_{(i)}$

end for

return $\hat{Q} \in \operatorname{argmin} \operatorname{tr}\left(Q^\intercal \nabla f(P_{(final)})\right)$ over $Q \in \mathcal{P}$ via linear assignment problem solver.

7 Code and data

Analyses relied on graspologic [45], NumPy [47], SciPy [44], Pandas [48], and NetworkX [49]. Plotting was performed using matplotlib [50] and Seaborn [51]. All code used for this paper can be found at github.com/neurodata/bgm and viewed as a JupyterBook [52] at docs.neurodata.io/bgm. The source code is archived at doi.org/10.5281/zenodo.6561550 (this version at zenodo.org/record/6564779).

8 Acknowledgements

We thank Thomas Athey for helpful comments.

B.D.P. was supported by the NSF Graduate Research Fellowship (Grant no. DGE1746891). J.T.V. was supported by the NSF CAREER Award (Grant no. 1942963). J.T.V. was supported by the NSF NeuroNex Award (Grant no. 2014862). J.T.V and C.E.P. were supported by the NIH BRAIN Initiative (Grant no. 1RF1MH123233-01).

References

[1] Donatello Conte, Pasquale Foggia, Carlo Sansone, and Mario Vento. Thirty years of graph matching in pattern recognition. International journal of pattern recognition and artificial intelligence, 18 (03):265–298, 2004.

[2] Ali Saad-Eldin, Benjamin D Pedigo, Carey E Priebe, and Joshua T Vogelstein. Graph matching via optimal transport. arXiv preprint arXiv:2111.05366, 2021.

[3] Kelly Marchisio, Youngser Park, Ali Saad-Eldin, Anton Alyakin, Kevin Duh, Carey Priebe, and Philipp Koehn. An analysis of euclidean vs. graph-based framing for bilingual lexicon induction from word embedding spaces. arXiv preprint arXiv:2109.12640, 2021.

[4] Jaewon Chung, Eric Bridgeford, Jesús Arroyo, Benjamin D Pedigo, Ali Saad-Eldin, Vivek Gopalakrishnan, Liang Xiang, Carey E Priebe, and Joshua T Vogelstein. Statistical connectomics. Annual Review of Statistics and Its Application, 8:463–492, 2021.

[5] Daniel L Sussman, Youngser Park, Carey E Priebe, and Vince Lyzinski. Matched filters for noisy induced subgraph detection. IEEE transactions on pattern analysis and machine intelligence, 42 (11):2887–2900, 2019.

[6] Joshua T Vogelstein, Eric W Bridgeford, Benjamin D Pedigo, Jaewon Chung, Keith Levin, Brett Mensh, and Carey E Priebe. Connectal coding: discovering the structures linking cognitive phenotypes to individual histories. Current opinion in neurobiology, 55:199–212, 2019.

[7] Danielle S Bassett and Olaf Sporns. Network neuroscience. Nature neuroscience, 20(3):353–364, 2017.

[8] Ellie S Heckscher, Aref Arzan Zarin, Serge Faumont, Matthew Q Clark, Laurina Manning, Akira Fushiki, Casey M Schneider-Mizell, Richard D Fetter, James W Truman, Maarten F Zwart, et al. Even-skipped+ interneurons are core components of a sensorimotor circuit that maintains left-right symmetric muscle contraction amplitude. Neuron, 88(2):314–329, 2015.

[9] Tomoko Ohyama, Casey M Schneider-Mizell, Richard D Fetter, Javier Valdes Aleman, Romain Franconville, Marta Rivera-Alba, Brett D Mensh, Kristin M Branson, Julie H Simpson, James W Truman, et al. A multilevel multimodal circuit enhances action selection in drosophila. Nature, 520 (7549):633–639, 2015.

[10] Matthew E Berck, Avinash Khandelwal, Lindsey Claus, Luis Hernandez-Nunez, Guangwei Si, Christopher J Tabone, Feng Li, James W Truman, Rick D Fetter, Matthieu Louis, et al. The wiring diagram of a glomerular olfactory system. Elife, 5:e14859, 2016.

[11] Akira Fushiki, Maarten F Zwart, Hiroshi Kohsaka, Richard D Fetter, Albert Cardona, and Akinao Nose. A circuit mechanism for the propagation of waves of muscle contraction in drosophila. Elife, 5:e13253, 2016.

[12] Tihana Jovanic, Casey Martin Schneider-Mizell, Mei Shao, Jean-Baptiste Masson, Gennady Denisov, Richard Doty Fetter, Brett Daren Mensh, James William Truman, Albert Cardona, and Marta Zlatic. Competitive disinhibition mediates behavioral choice and sequences in drosophila. Cell, 167(3):858–870, 2016.

[13] Philipp Schlegel, Michael J Texada, Anton Miroschnikow, Andreas Schoofs, Sebastian Hückesfeld, Marc Peters, Casey M Schneider-Mizell, Haluk Lacin, Feng Li, Richard D Fetter, et al. Synaptic transmission parallels neuromodulation in a central food-intake circuit. Elife, 5:e16799, 2016.

[14] Maarten F Zwart, Stefan R Pulver, James W Truman, Akira Fushiki, Richard D Fetter, Albert Cardona, and Matthias Landgraf. Selective inhibition mediates the sequential recruitment of motor pools. Neuron, 91(3):615–628, 2016.

[15] Katharina Eichler, Feng Li, Ashok Litwin-Kumar, Youngser Park, Ingrid Andrade, Casey M Schneider-Mizell, Timo Saumweber, Annina Huser, Claire Eschbach, Bertram Gerber, et al. The complete connectome of a learning and memory centre in an insect brain. Nature, 548(7666): 175–182, 2017.

[16] Stephan Gerhard, Ingrid Andrade, Richard D Fetter, Albert Cardona, and Casey M Schneider-

Mizell. Conserved neural circuit structure across drosophila larval development revealed by comparative connectomics. Elife, 6:e29089, 2017.

[17] Ivan Larderet, Pauline MJ Fritsch, Nanae Gendre, G Larisa Neagu-Maier, Richard D Fetter, Casey M Schneider-Mizell, James W Truman, Marta Zlatic, Albert Cardona, and Simon G Sprecher. Organization of the drosophila larval visual circuit. Elife, 6:e28387, 2017.

[18] Suguru Takagi, Benjamin Thomas Cocanougher, Sawako Niki, Dohjin Miyamoto, Hiroshi Kohsaka, Hokto Kazama, Richard Doty Fetter, James William Truman, Marta Zlatic, Albert Cardona, et al. Divergent connectivity of homologous command-like neurons mediates segment-specific touch responses in drosophila. Neuron, 96(6):1373–1387, 2017.

[19] Anita Burgos, Ken Honjo, Tomoko Ohyama, Cheng Sam Qian, Grace Ji-eun Shin, Daryl M Gohl, Marion Silies, W Daniel Tracey, Marta Zlatic, Albert Cardona, et al. Nociceptive interneurons control modular motor pathways to promote escape behavior in drosophila. Elife, 7:e26016, 2018.

[20] Arnaldo Carreira-Rosario, Aref Arzan Zarin, Matthew Q Clark, Laurina Manning, Richard D Fetter, Albert Cardona, and Chris Q Doe. Mdn brain descending neurons coordinately activate backward and inhibit forward locomotion. Elife, 7:e38554, 2018.

[21] Anton Miroschnikow, Philipp Schlegel, Andreas Schoofs, Sebastian Hueckesfeld, Feng Li, Casey M Schneider-Mizell, Richard D Fetter, James W Truman, Albert Cardona, and Michael J Pankratz. Convergence of monosynaptic and polysynaptic sensory paths onto common motor outputs in a drosophila feeding connectome. Elife, 7:e40247, 2018.

[22] Ibrahim Tastekin, Avinash Khandelwal, David Tadres, Nico D Fessner, James W Truman, Marta Zlatic, Albert Cardona, and Matthieu Louis. Sensorimotor pathway controlling stopping behavior during chemotaxis in the drosophila melanogaster larva. Elife, 7:e38740, 2018.

[23] Tihana Jovanic, Michael Winding, Albert Cardona, James W Truman, Marc Gershow, and Marta Zlatic. Neural substrates of drosophila larval anemotaxis. Current Biology, 29(4):554–566, 2019.

[24] Brandon Mark, Sen-Lin Lai, Aref Arzan Zarin, Laurina Manning, Heather Q Pollington, Ashok Litwin-Kumar, Albert Cardona, James W Truman, and Chris Q Doe. A developmental framework linking neurogenesis and circuit formation in the drosophila cns. Elife, 10:e67510, 2021.

[25] Aref Arzan Zarin, Brandon Mark, Albert Cardona, Ashok Litwin-Kumar, and Chris Q Doe. A multilayer circuit architecture for the generation of distinct locomotor behaviors in drosophila. Elife, 8: e51781, 2019.

[26] Claire Eschbach, Akira Fushiki, Michael Winding, Casey M Schneider-Mizell, Mei Shao, Rebecca Arruda, Katharina Eichler, Javier Valdes-Aleman, Tomoko Ohyama, Andreas S Thum, et al. Recurrent architecture for adaptive regulation of learning in the insect brain. Nature Neuroscience, 23(4):544–555, 2020.

[27] Sebastian Hückesfeld, Philipp Schlegel, Anton Miroschnikow, Andreas Schoofs, Ingo Zinke, André N Haubrich, Casey M Schneider-Mizell, James W Truman, Richard D Fetter, Albert Cardona, et al. Unveiling the sensory and interneuronal pathways of the neuroendocrine connectome in drosophila. Elife, 10:e65745, 2021.

[28] Claire Eschbach, Akira Fushiki, Michael Winding, Bruno Afonso, Ingrid V Andrade, Benjamin T Cocanougher, Katharina Eichler, Ruben Gepner, Guangwei Si, Javier Valdes-Aleman, et al. Circuits for integrating learned and innate valences in the insect brain. Elife, 10:e62567, 2021.

[29] Donniell E Fishkind, Sancar Adali, Heather G Patsolic, Lingyao Meng, Digvijay Singh, Vince Lyzinski, and Carey E Priebe. Seeded graph matching. Pattern recognition, 87:203–215, 2019.

[30] Konstantinos Pantazis, Daniel L Sussman, Youngser Park, Carey E Priebe, and Vince Lyzinski. Multiplex graph matching matched filters. In 2019 IEEE International Conference on Big Data (Big Data), pages 4921–4927. IEEE, 2019.

[31] Joshua T Vogelstein, John M Conroy, Vince Lyzinski, Louis J Podrazik, Steven G Kratzer, Eric T Harley, Donniell E Fishkind, R Jacob Vogelstein, and Carey E Priebe. Fast approximate quadratic programming for graph matching. PLOS one, 10(4):e0121002, 2015.

[32] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. Naval research logistics quarterly, 3(1-2):95–110, 1956.
[33] Vince Lyzinski, Daniel L Sussman, Donniell E Fishkind, Henry Pao, Li Chen, Joshua T Vogelstein, Youngser Park, and Carey E Priebe. Spectral clustering for divide-and-conquer graph matching. Parallel Computing, 47:70–87, 2015.
[34] Paul Erdos, Alfréd Rényi, et al. On the evolution of random graphs. Publ. Math. Inst. Hung. Acad. Sci, 5(1):17–60, 1960.
[35] Edgar N Gilbert. Random graphs. The Annals of Mathematical Statistics, 30(4):1141–1144, 1959.
[36] Daniel J Bumbarger, Metta Riebesell, Christian Rödelsperger, and Ralf J Sommer. System-wide rewiring underlies behavioral differences in predatory and bacterial-feeding nematodes. Cell, 152 (1-2):109–119, 2013.
[37] Steven J Cook, Travis A Jarrell, Christopher A Brittin, Yi Wang, Adam E Bloniarz, Maksim A Yakovlev, Ken CQ Nguyen, Leo T-H Tang, Emily A Bayer, Janet S Duerr, et al. Whole-animal connectomes of both caenorhabditis elegans sexes. Nature, 571(7763):63–71, 2019.
[38] Marta Costa, James D Manton, Aaron D Ostrovsky, Steffen Prohaska, and Gregory SXE Jefferis. Nblast: rapid, sensitive comparison of neuronal structure and construction of neuron family databases. Neuron, 91(2):293–311, 2016.
[39] Nadine Randel, Réza Shahidi, Csaba Verasztó, Luis A Bezares-Calderón, Steffen Schmidt, and Gáspár Jékely. Inter-individual stereotypy of the platynereis larval visual connectome. Elife, 4: e08069, 2015.
[40] Daniel Witvliet, Ben Mulcahy, James K Mitchell, Yaron Meirovitch, Daniel R Berger, Yuelong Wu, Yufang Liu, Wan Xian Koh, Rajeev Parvathala, Douglas Holmyard, et al. Connectomes across development reveal principles of brain maturation. Nature, 596(7871):257–261, 2021.
[41] Philipp Schlegel, Alexander Shakeel Bates, Tomke Stürner, Sridhar R Jagannathan, Nikolas Drummond, Joseph Hsu, Laia Serratosa Capdevila, Alexandre Javier, Elizabeth C Marin, Asa Barth-Maron, et al. Information flow, cell types and stereotypy in a full olfactory connectome. Elife, 10: e66018, 2021.
[42] Avanti Athreya, Donniell E Fishkind, Minh Tang, Carey E Priebe, Youngser Park, Joshua T Vogelstein, Keith Levin, Vince Lyzinski, and Yichen Qin. Statistical inference on random dot product graphs: a survey. The Journal of Machine Learning Research, 18(1):8393–8484, 2017.
[43] Minh Tang, Avanti Athreya, Daniel L Sussman, Vince Lyzinski, Youngser Park, and Carey E Priebe. A semiparametric two-sample hypothesis testing problem for random graphs. Journal of Computational and Graphical Statistics, 26(2):344–354, 2017.
[44] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. Nature Methods, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
[45] Jaewon Chung, Benjamin D Pedigo, Eric W Bridgeford, Bijan K Varjavand, Hayden S Helm, and Joshua T Vogelstein. Graspy: Graph statistics in python. J. Mach. Learn. Res., 20:158–1, 2019.
[46] Harold W Kuhn. The hungarian method for the assignment problem. Naval research logistics quarterly, 2(1-2):83–97, 1955.
[47] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren

Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. Nature, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL https://doi.org/10.1038/s41586-020-2649-2.

[48] Wes McKinney et al. Data structures for statistical computing in python. In Proceedings of the 9th Python in Science Conference, volume 445, pages 51–56. Austin, TX, 2010.

[49] Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.

[50] John D Hunter. Matplotlib: A 2d graphics environment. Computing in science & engineering, 9 (03):90–95, 2007.

[51] Michael L. Waskom. seaborn: statistical data visualization. Journal of Open Source Software, 6 (60):3021, 2021. doi: 10.21105/joss.03021. URL https://doi.org/10.21105/joss.03021.

[52] Executable Books Community. Jupyter book, February 2020. URL https://doi.org/10.5281/zenodo.4539666.

9 Appendix

Here, we show how the modified gradient (with respect to the permutation P) is computed for bisected graph matching. We start with the objective function for the bisected graph matching problem:

$$f(P) = \|A_{LL} - PA_{RR}P^T\|_F^2 + \|A_{LR}P^T - PA_{RL}\|_F^2$$

This can be written as the sum of a term for the ipsilateral subgraphs ($f_I(P)$) and one for the contralateral subgraphs ($f_C(P)$):

$$f(P) = f_I(P) + f_C(P)$$

Rewriting the first term:

$$f_I(P) = \|A_{LL} - PA_{RR}P^T\|_F^2 = tr((A_{LL} - PA_{RR}P^T)^T(A_{LL} - PA_{RR}P^T))$$

$$f_I(P) = \|A_{LL}\|_F^2 + \|A_{RR}\|_F^2 - 2tr(A_{LL}PA_{RR}^TP^T)$$

And likewise for the second:

$$f_C(P) = \|A_{LR}P^T - PA_{RL}\|_F^2 = tr((A_{LR}P^T - PA_{RL})^T(A_{LR}P^T - PA_{RL}))$$

$$f_C(P) = \|A_{LR}\|_F^2 + \|A_{RL}\|_F^2 - 2tr(A_{LR}^TPA_{RL}P)$$

We can drop terms unrelated to P and constants (since they will not affect the optimization) and returning to the full objective function we have:

$$f(P) = -tr(A_{LL}PA_{RR}^TP^T) - tr(A_{LR}^TPA_{RL}P)$$

Then, taking the derivative with respect to P for the first term:

$$\nabla f_I(P) = -A_{LL}PA_{RR}^T - A_{LL}^TPA_{RR}$$

Similarly for the second term:

$$\nabla f_C(P) = -A_{LR}P^TA_{RL}^T - A_{RL}^TP^TA_{LR}$$

So for the full objective function, the gradient is:

$$\nabla f(P) = -\left(A_{LL}PA_{RR}^T + A_{LL}^TPA_{RR} + A_{LR}P^TA_{RL}^T + A_{RL}^TP^TA_{LR}\right)$$