

Neuron matching in *C. elegans* with robust approximate linear regression without correspondence

Amin Nejatbakhsh, Erdem Varol

Department of Neuroscience, Department of Statistics, Zuckerman Institute
Grossman Center for the Statistics of Mind, Columbia University, New York, NY USA

mn2822@columbia.edu, ev2430@columbia.edu

Abstract

We propose methods for estimating correspondence between two point sets under the presence of outliers in both the source and target sets. The proposed algorithms expand upon the theory of the regression without correspondence problem to estimate transformation coefficients using unordered multisets of covariates and responses. Previous theoretical analysis of the problem has been done in a setting where the responses are a complete permutation of the regressed covariates. This paper expands the problem setting by analyzing the cases where only a subset of the responses is a permutation of the regressed covariates in addition to some covariates possibly being adversarial outliers. We term this problem robust regression without correspondence and provide several algorithms based on random sample consensus for exact and approximate recovery in a noiseless and noisy one-dimensional setting as well as an approximation algorithm for multiple dimensions. The theoretical guarantees of the algorithms are verified in simulated data. We demonstrate an important computational neuroscience application of the proposed framework by demonstrating its effectiveness in a *Caenorhabditis elegans* neuron matching problem where the presence of outliers in both the source and target nematodes is a natural tendency. Open source code implementing this method is available at <https://github.com/amin-nejat/RRWOC>.

1. Introduction

Point set registration is one of the central problems in computer vision that involves the optimization of a transformation that aligns two sets of point clouds [34, 29]. Point set registration have been applied in numerous fields including but not limited to robotics [39], medical imaging [3], object recognition [11], panorama stitching [4] and computational neuroscience [7]. The types of allowable transformations and energy functions utilized in the cost function have

differentiated varying methods [6, 25, 40, 15, 18, 2, 24, 29, 17, 28, 12, 8, 23, 22, 36]. In general, point set registration methods employ an iterative strategy of solving the transformation and updating the matching which works well in practice but there are no guarantees for reaching the global optima [9]. Only a few methods have provided approximate globally optimal solutions [37, 40]. These methods rely on severe constraints of the transformation domains, such as the 3D rotation group $SO(3)$, in order to employ branch and bound techniques on discretizations.

Theoretical analysis of the recovery guarantees of point set registration has not been performed for a general number of dimensions until recently when it was termed as *unlabelled sensing* by [32] as a problem with duality connections with the well-known problem of compressed sensing [10]. In this problem, similar to linear regression, the response signal is modeled as a linear combination of a set of covariates. However, the correspondence of the responses to the covariates is modeled as having been shuffled by an unknown permutation matrix. For this reason, the problem has also been termed as *linear regression with shuffled labels* [1], *linear regression with an unknown permutation* [27], *homomorphic sensing* [31] or *linear regression without correspondence* (RWOC) [16], the latter of which will be used to refer to the problem herein. Although RWOC is, in general, an NP-hard problem [27], there have been several advances in recent years to propose signal to noise ratio (SNR) bounds for recovery of the permutation matrix and the regression coefficients [27, 33]. Conversely, the same works have also analyzed the SNR and sampling regime by which no recovery is possible.

Nevertheless, the computer vision community has attempted to solve the point set registration problem through consideration of outliers and missing correspondences, which are typically encountered in real-world applications. A common technique used in point set registration to robustify the optimization against outliers is to employ random sampling consensus (RANSAC) subroutines [14, 30, 35]. The main advantages of RANSAC are that the randomiza-

tion procedure employed can severely reduce the computational cost of an otherwise combinatorial search.

Motivated by applications in computational neuroscience such as matching the neuronal populations of *Caenorhabditis elegans* (*C. elegans*) across different nematodes, we aim to unify the ideas presented in RWOC literature and robust point set registration methods to provide provably approximate solutions to the RWOC problem in the presence of outliers and missing measurements commonly encountered in fluorescence microscopy data. Robustly and automatically matching and identifying neurons in *C. elegans* could expedite the post-experimental data analysis and hypothesis testing cycle [7, 20, 26, 38].

1.1. Main contributions

The main contributions presented in this paper are the introduction of randomized algorithms for the recovery of the regression coefficients in the RWOC problem that takes into account noise, missing data, and outliers. Hsu et al. [16] provide algorithms for the noisy case without generative assumptions; their algorithm takes into account square permutation matrices, which assumes that the entire signal is captured in the responses and does not take into account any missing correspondences or outliers. Unnikrishnan et al. [32, 33] provide combinatorial existence arguments. Tsakiris et al. [31] provide an algorithm that takes into account missing correspondences or outliers but not both. Our method is designed for the practical purpose of matching point clouds that may have noisy measurements, missing correspondences, and outliers. Missing data can be thought of as outliers in the source point set, but they can have different interpretations. For example, if the goal is to register an image onto an already existing atlas, then the parts of the atlas that are not present in the image are called missing data. The assumption is that the atlas contains a complete set of objects while the image could be missing some parts for reasons such as incomplete field of view, mutant defects, individual differences, etc. This is undoubtedly the case in the application domain of neuron tracking and matching in biological applications where structures of interest might be missing from the field of view or other unrelated confounding biological structures might exist and potentially be captured by the detection algorithms. Specifically, we demonstrate the efficacy of the proposed method in the identification and tracking of in-vivo (*C. elegans*) neurons where it is possible that some neurons are missing and adversarial objects that might be confused as neurons are present.

In summary, our contributions are four-fold:

1. We introduce the notion of "robust" regression without correspondence (rRWOC) that models missing correspondences between responses and covariates as well as completely missed associations in the form of outliers and missing data. In contrast with standard

point set registration methods, we further consider the case of adversarial outliers.

2. We introduce a polynomial-time algorithm to find the exact solution for the one-dimensional noiseless rRWOC and the approximate solution in the noisy regime.
3. We introduce a randomized approximately correct algorithm that is more efficient than pure-brute force approaches in multiple dimensional rRWOC.
4. We demonstrate the computational neuroscience application of our approach to point-set registration problems in the context of automatically matching and identification of the cellular layout of the nervous system of the nematode *C. elegans*.

1.2. Paper organization

In section 2, we introduce our statistical regression model (rRWOC) that accounts for permuted correspondences, outliers, and noise. We then demonstrate the added computational complexity of recovery of rRWOC in contrast with simple linear regression and RWOC in a one-dimensional case in section 3.1. In section 3.2, we provide a randomized algorithm for the rRWOC problem in multiple dimensions with convergence bounds. Lastly, in section 4.1, we verify the theoretical recovery guarantees in simulated experiments and in section 4.2 show the neuroscience application of the proposed algorithms in the *C. elegans* neuron matching problem.

2. Regression model

First, we introduce notation. Let $\mathbf{X} = [\mathbf{x}_1 | \mathbf{x}_2 | \dots | \mathbf{x}_m]^T \in \mathbb{R}^{m \times d}$ and $\mathbf{Y} = [\mathbf{y}_1 | \mathbf{y}_2 | \dots | \mathbf{y}_n]^T \in \mathbb{R}^{n \times d}$ denote two d -dimensional point sets consisting of m and n points, respectively. Let us call \mathbf{X} the reference or source set. Let \mathbf{Y} denote the target set which may contain outliers and missing correspondences. Note that the points in \mathbf{X} that are missing correspondences in \mathbf{Y} can be seen as outliers in the source set, hence justifying our claim that we model outliers in both the source and target sets.

Let the set of indices $\mathcal{I} = \{i_1, \dots, i_{|\mathcal{I}|}\} \subseteq [n]$ denote the indices of \mathbf{y}_j which are inliers. Conversely, let $\mathcal{O} = \{o_1, \dots, o_{|\mathcal{O}|}\} \subseteq [n]$ denote set of indices of \mathbf{y}_j which are outliers. By construction, these sets are a disjoint partition of the entire index set of target points: $\mathcal{I} \cup \mathcal{O} = [n]$ and $\mathcal{I} \cap \mathcal{O} = \emptyset$. Let $\mathbf{\Pi} \in \mathcal{P}^{n \times m}$ denote a possibly unbalanced permutation matrix where there are at most $\min\{n, m\}$ ones placed such that no row or column has more than a single one. All other entries are zeroes. Let $\pi(i)$ denote the location of the one in the i th row of the permutation




Figure 1. Demonstration of various problem settings of regression without correspondence. **A:** Full set of hidden correspondences between source and target multisets. **B:** Missing correspondences in the target set. **C:** Unstructured outliers in the target set. **D:** Adversarial outliers in the target set – this setting imposes a theoretical ceiling of 50% outliers in the target set. However, in practice, more than 50% ratio of unstructured outliers can be handled.

matrix Π . Next, let $\beta \in \mathbb{R}^{d \times d}$ denote the regression coefficients and $\epsilon \sim \mathcal{N}(0, \nu I)$ denote zero-mean Gaussian noise. Lastly, let $U[\mathcal{C}]$ denote the uniform distribution within some closed convex set \mathcal{C} . Given these definitions, we can define the **robust regression without correspondence (rRWOC)** model as

$$\begin{aligned} y_{i_j} &= \mathbf{x}_{\pi(i_j)} \beta + \epsilon && \text{for } i_j \in \mathcal{I} \\ y_{o_l} &\sim U[\mathcal{C}] && \text{for } o_l \in \mathcal{O} \end{aligned} \quad (1)$$

Note that the bias terms in the regression can be modeled by padding \mathbf{x} with a constant column of ones.

In contrast with linear regression, where the sole objective is to recover the coefficients β , the two-fold objective of RWOC is to recover the correct permutation matrix Π , and the regression coefficients β . To add to the complexity of the problem, the three-fold objective of rRWOC is to recover the inlier set \mathcal{I} , the permutation Π , and the coefficients β .

3. Algorithms

To aid in the recovery of the solution in rRWOC, we introduce the following assumption.

Assumption 1 (Maximal inlier set). *For point sets \mathbf{X} , \mathbf{Y} , there exists a triple $\{\mathcal{I}^*, \beta^*, \Pi^*\}$ that is maximal in the sense that $n \geq |\mathcal{I}^*| \geq |\mathcal{I}'|$ such that any other triple $\{\mathcal{I}', \beta', \Pi'\}$ is not considered to be the underlying regression model.*

Assumption 1 allows the identifiability of whether a given hypothetical index set can be considered to be the true underlying inlier set or not. In practical terms, suppose we generate simulated data with n points in \mathbf{Y} of which $k > n/2$ are outliers generated uniformly and the

remainder generated with respect to a coefficient $\beta^{\mathcal{I}}$ such that $\mathbf{Y}_{[\mathcal{I}]} = \mathbf{X}_{\pi(\mathcal{I})} \beta^{\mathcal{I}} + \epsilon^{\mathcal{I}}$. There may be cases such that uniformly generated "outliers", $\mathbf{Y}_{[\mathcal{O}]}$, are structured such that there exists a coefficient $\beta^{\mathcal{O}}$ and permutation $\Pi^{\mathcal{O}}$ such that $\mathbf{Y}_{[\mathcal{O}]} = \mathbf{X}_{\pi(\mathcal{O})} \beta^{\mathcal{O}} + \epsilon^{\mathcal{O}}$ where $\text{Var}(\epsilon^{\mathcal{I}}) \geq \text{Var}(\epsilon^{\mathcal{O}})$. In this case, $\beta^{\mathcal{O}}$ is identifiable but not verifiable as "correct." In practical terms, assumption 1 puts a ceiling on the maximum proportion of outliers that any regression without correspondence algorithm can handle. In a simplest example, if the target point set consists of two duplicate copies of rotated and transformed source point set, it is impossible to identify the correct matching. However, if one of the duplicates has less points, then we can invoke the principle of the maximal inlier set to identify the correct target set. See figure 1 for a visualization.

Equipped with the rRWOC model and assumption 1, we now demonstrate the progressive increase in the complexity of recovery of ordinary linear regression, RWOC, and rRWOC in one-dimension.

3.1. Optimal regression in $d = 1$

Linear regression in one-dimension with known correspondences, no offset term and no outliers can be obtained in $O(n)$ time using the univariate normal equation: $\beta_{OLS} = \frac{\sum_i^n y_i x_{\pi(i)}}{\sum_i^n x_{\pi(i)}^2}$. On the other hand, RWOC in the one-dimensional case with *no noise* can be solved in $O(n \log(n))$ steps via the method of moments and a simple sorting operation. Namely, first, the regressor β_{RWOC} can be estimated using the ratio of the first moments of the covariates to the responses:

$$\beta_{RWOC} = \frac{\sum_{i=1}^n y_i}{\sum_{i=1}^n x_i} \quad (2)$$

Algorithm 1 One dimensional robust regression without correspondence - Exhaustive approach

Input: Reference set: $\{x_1, \dots, x_m\}$, target set: $\{y_1, \dots, y_n\}$, outlier margin: ν

Require: $k < \frac{n}{2}$ (number of outliers)

```

1: for  $i = 1, \dots, n$  do
2:   for  $j = 1, \dots, m$  do
3:     Compute  $\beta^{i,j} = y_i/x_j$ 
4:     Compute linear assignment [21]:
        $\Pi^{i,j} \leftarrow \arg \min_{\Pi \in \mathcal{P}^{n \times m}} \|\mathbf{x}\beta^{i,j} - \Pi^T \mathbf{y}\|_2^2$ 
5:     Compute hypothetical inliers:
        $\mathcal{I}^{i,j} = \{l : |x_{\pi^{i,j}(l)}\beta^{i,j} - y_l| \leq \nu\}$ 
6:   end for
7: end for
8: return  $(i^*, j^*) = \arg \max_{(i,j)} |\mathcal{I}^{i,j}|$ ,  $\mathcal{I}^* = \mathcal{I}^{i^*,j^*}$ ,  $\Pi^* =$ 
 $\Pi^{i^*,j^*}$ ,  $\beta^* \leftarrow \frac{\sum_{l \in \mathcal{I}^*} y_l x_{\pi^*(l)}}{\sum_{l \in \mathcal{I}^*} x_{\pi^*(l)}^2}$ 

```

and then the permutation can be recovered using the rearrangement inequality [5],

$$\min_{\Pi} \sum_{i=1}^n (y_i - \hat{y}_{\pi(i)})^2 = \sum_{i=1}^n (y_{(i)} - \hat{y}_{(i)})^2 = \quad (3)$$

$$\|\Pi_y \mathbf{y} - \Pi_{\hat{y}} \hat{\mathbf{y}}\|_2^2 \longrightarrow \Pi_{RWOC} = \Pi_y^T \Pi_{\hat{y}}$$

where $y_{(i)}$ denotes sorted y_i and $\hat{y}_{(i)}$ denotes sorted $x_i \beta_{RWOC}$ and Π_y and $\Pi_{\hat{y}}$ denote the permutation matrices that capture the sorting operations.

In the case with outlier elements in \mathbf{y} , the problem is non-trivial, even in one dimension, since sorting does not allow the identification of outliers¹. To solve the one dimensional rRWOC, we introduce algorithm 1 which recovers the triplet $\{\mathcal{I}^*, \beta^*, \Pi^*\}$ in an exhaustive fashion.

Proposition 1 (Correctness of Algorithm 1). *Suppose there exist $n - k$ inliers in \mathbf{y} and that $k < n/2$. Then algorithm 1 yields the correct regression coefficient $\beta^* = \beta$ with probability 1 for noiseless data and with high probability for noisy data with an appropriately selected margin parameter ν .*

Proof. (The full proof is included in supplementary material) The overview of the proof is as follows. In the noiseless case, if $j = \pi(i)$ then $\beta^{i,j} = \frac{y_i}{x_j} = \beta^*$. The projection $\mathbf{x}\beta^{i,j}$ maps all reference points to their exact corresponding reference points. Thus the Hungarian algorithm will yield these as the assignments since they incur minimal cost. Therefore, we will have $|\mathcal{I}^{i,j}| \geq n - k$. The cardinality of inliers is lower bounded and not equal to $n - k$ since outlier points may by chance be transformed to points in \mathbf{y} as well. Contrarily, suppose the transformation $\beta^{i,l}$ for $l \neq \pi(i)$ yields

¹See supplementary material section 5 for a toy example experiment.

Algorithm 2 One dimensional robust regression without correspondence - Randomized approach

Input: Reference set: $\{x_1, \dots, x_m\}$, target set: $\{y_1, \dots, y_n\}$, δ (probability of success), outlier margin: ν

Require: $k < \frac{n}{2}$ (number of outliers)

```

1: for  $t = 1, \dots, q$  do
2:   Sample  $i \sim [n]$  and sample  $j \sim [m]$ 
3:   Compute  $\beta^t = y_i/x_j$ 
4:   Compute linear assignment [21]:
        $\Pi^t \leftarrow \arg \min_{\Pi \in \mathcal{P}^{n \times m}} \|\mathbf{x}\beta^t - \Pi^T \mathbf{y}\|_2^2$ 
5:   Compute hypothetical inliers:
        $\mathcal{I}^t = \{l : |x_{\pi^t(l)}\beta^t - y_l| \leq \nu\}$ 
6: end for
7: return  $t^* = \arg \max_t |\mathcal{I}^t|$ ,  $\mathcal{I}^* = \mathcal{I}^{t^*}$ ,
 $\Pi^* = \Pi^{t^*}$ ,  $\beta^* \leftarrow \frac{\sum_{l \in \mathcal{I}^*} y_l x_{\pi^*(l)}}{\sum_{l \in \mathcal{I}^*} x_{\pi^*(l)}^2}$ 

```

a larger hypothesized inlier set $\mathcal{I}^{i,l}$, such that $|\mathcal{I}^{i,l}| > |\mathcal{I}^{i,j}|$ then this means that there are more points in $\mathbf{x}\beta^{i,l}$ that are closer to \mathbf{y} than $\mathbf{x}\beta^{i,j}$, contradicting the assumption that $n - k$ is the maximal inlier set. \square

The time complexity of algorithm 1 can be analyzed as follows. The main computational cost is due to linear assignment which incurs a cost of $O(\max\{m, n\}^3)$ if [19] variant is used. Linear assignment is repeated mn times. If m and n are of the same order, then algorithm 1 has complexity $O(n^5)$.

However, if the ratio of inliers to outliers is relatively high, then it is possible to use randomization procedures like RANSAC [14, 30] to speed up the algorithm to yield the correct regression coefficient with high probability. This is demonstrated in algorithm 2.

Proposition 2 (Correctness of Algorithm 2). *Suppose there are $n - k$ inliers in \mathbf{x} and that $k < n/2$. In $q \geq \frac{\log(1-\delta)}{\log(1-\frac{n-k}{mn})}$ iterations, algorithm 2 yields the correct regression coefficient $\beta^* = \beta$ with probability $\delta \in (0, 1)$ for an appropriately selected margin parameter ν .*

Proof. The success of algorithm 1 relies on the fact that the exhaustive search eventually hits a tuple (i, j) such that $j = \pi(i)$ which yields the correct regression coefficient. Therefore, when randomly sampling $(i, j) \sim [n] \times [m]$, the probability of choosing a corresponding pair is $\frac{n-k}{n} \frac{1}{m}$. The probability of iterating q times such that no correct correspondence is selected is $(1 - (n - k)/(nm))^q = (1 - \delta)$ where δ is the desired success rate. Taking logs yields, $q = \frac{\log(1-\delta)}{\log(1-(n-k)/(nm))}$ \square

Algorithm 3 Robust regression without correspondence - Randomized approach

Input: $\mathbf{X} = [\mathbf{x}_1 | \dots | \mathbf{x}_m]^T \in \mathbb{R}^{m \times d}$ (reference points), $\mathbf{Y} = [\mathbf{y}_1 | \dots | \mathbf{y}_n]^T \in \mathbb{R}^{n \times d}$ (target points), δ (probability of success), ν (outlier margin)

Require: $k < \frac{n}{2}$ (number of outliers)

- 1: **for** $t = 1, \dots, q$ **do**
 - 2: Sample $\mathbf{i} = (i_1, \dots, i_d) \sim [n]^d$ w/o replacement
 - 3: Sample $\mathbf{j} = (j_1, \dots, j_d) \sim [m]^d$ w/o replacement
 - 4: Compute $\beta^t = \arg \min_{\beta} \|\mathbf{X}_{[\mathbf{j}]} \beta - \mathbf{Y}_{[\mathbf{i}]} \|_F^2$
 - 5: Compute linear assignment via [21]:
 $\Pi^t \leftarrow \arg \min_{\Pi \in \mathcal{P}^{m \times n}} \|\mathbf{X} \beta^t - \Pi \mathbf{Y}\|_F^2$
 - 6: Compute hypothetical inliers:
 $\mathcal{I}^t = \{l : \|\mathbf{x}_{\pi^t(l)} \beta^t - \mathbf{y}_l\|_2 \leq \nu\}$
 - 7: **end for**
 - 8: **return** $t^* = \arg \max_t |\mathcal{I}^t|$, $\mathcal{I}^* = \mathcal{I}^{t^*}$,
 $\Pi^* = \Pi_{\mathcal{I}^*}^{t^*}, \beta^* \leftarrow \arg \min_{\beta} \|\mathbf{X}_{\pi^*(\mathcal{I}^*)} \beta - \mathbf{Y}_{\mathcal{I}^*}\|_F^2$
-

The time complexity of randomized algorithm 2 is $O\left(\frac{\log(1-\delta)}{\log(1-(n-k)/n^2)} n^3\right)$.

3.2. Randomized approximation algorithm ($d \geq 2$)

The exhaustive approach for the $d \geq 2$ dimensional case requires $\binom{n}{d} \binom{m}{d}$ d -subset comparisons of \mathbf{X}, \mathbf{Y} in order to guarantee hitting correct (in the noiseless case) or approximately correct (in the noisy case) regression coefficients, with complexity $O(m^d n^d)$. However, especially in higher dimensions, the randomized procedure enables a substantial reduction of iterations to yield a high probability correct triplet of inlier set, permutation, and regression coefficients. The randomized algorithm for rRWOC in $d \geq 2$ is demonstrated in algorithm 3. Random ordered d -tuples of reference and target point sets are sampled and are used to align the remainder of the point set. The number of hypothetical inliers for each hypothetical correspondence is assessed by checking whether the transformed reference points are arbitrarily close to a target point. With high probability, if correct a d -tuple correspondence is captured, the number of transformed reference points matching a target point will be high (Figure 1 top), otherwise it will result in a partial coverage (Figure 1 bottom).

Proposition 3. For $q \geq \frac{\log(1-\delta)}{\log\left(1 - \frac{\binom{m-k}{d}}{\binom{m}{d} \binom{n}{d}}\right)}$, algorithm 3 re-

covers β^* and Π^* and the set of inliers for the noiseless case with probability $(1 - \delta)$ using arbitrarily small ν . For sufficiently small noise variance and appropriately chosen ν , algorithm 3 recovers approximate β^* with high probability.

Proof. Analogous to the analysis of algorithm 2, the probability of drawing d inliers out of n points with k outliers in \mathbf{Y} is $\frac{\binom{n-k}{d}}{\binom{n}{d}}$. The probability of matching the drawn inliers with the d corresponding sampled reference points in \mathbf{X} is $\frac{1}{\binom{m}{d}}$. Probability that any draw is not going to match is $1 - \frac{\binom{n-k}{d}}{\binom{m}{d} \binom{n}{d}}$. The probability that q draws will be incorrect is $\left(1 - \frac{\binom{n-k}{d}}{\binom{m}{d} \binom{n}{d}}\right)^q$. If we set this to be the probability of failure $(1 - \delta)$, we then have the estimate for the number of draws we need to make as $q(\delta, n, m, k) \geq \log(1 - \delta) / \log\left(1 - \frac{\binom{n-k}{d}}{\binom{m}{d} \binom{n}{d}}\right)$ \square

The complexity of algorithm 3 can be analyzed as follows. In each inner loop, the regression coefficient solution requires $O(d^3)$ time, the Hungarian algorithm requires $O(nmd)$ to compute the input distance matrix and then $O(\max\{n, m\}^3)$ to optimize the permutation matrix. The rest of the operations are $O(d)$. Therefore, the overall time complexity is

$$O\left(\frac{\log(1-\delta)}{\log\left(1 - \frac{\binom{m-k}{d}}{\binom{m}{d} \binom{n}{d}}\right)} (d^3 + nmd + \max\{n, m\}^3)\right). \quad (4)$$

3.2.1 Margin parameter (ν) selection

Both of the proofs of the noiseless and the noisy cases of proposition 1 rely on knowledge of the true regression coefficient and the noise variance in order to estimate the margin coefficient ν and output the optimal regression coefficient with high probability. However, in practice, as in many RANSAC-like robust regression settings, these parameters cannot be known apriori, and ν is typically determined via empirical heuristics and or cross-validation [14].

In the noiseless case, an appropriate heuristic is choosing ν arbitrarily small since the correct regression should yield zero residual. However, for the noisy case, if available, supervised data should be used with known correspondences to estimate the actual dispersion of point correspondences.

4. Numerical results

To verify the theoretical guarantees of the proposed algorithms, simulated data in 3 dimensions was generated in both noisy and noiseless regimes. Furthermore, iterative solutions of β and Π were obtained to demonstrate the suboptimality of local minima found using block coordinate descent for this non-convex problem.

The neuroscience application of rRWOC was demonstrated in the context of point set matching of neurons of *C. elegans* worms recorded using fluorescence microscopy imaging. The matching accuracy with respect to ground




Figure 2. **Left:** Recovery rate (colorbar) vs. missing data ratio (y-axis) vs. SNR (x-axis), **Middle:** Recovery rate (colorbar) vs. outlier ratio (y-axis) vs. SNR (x-axis), **Right:** Recovery rate (y-axis) vs. outlier ratio (x-axis), blue: rRWOC, red: ICP, yellow: randomized rRWOC ($\delta = 0.9$), purple: randomized rRWOC ($\delta = 0.6$)




Figure 3. 2D projection of 3D fluorescence microscopy image of *C.elegans* head in Yemini et al. dataset. Superimposed annotation points denote neuron locations. Outliers are detections that do not correspond to neurons and missing data are undetected neurons.

truth was assessed for rRWOC as well as a robust variant of the iterative closest point (ICP) algorithm [6] known as trimmed ICP [9]. We also compared to the state of the art algorithm for regression without correspondence, termed homomorphic sensing (HS) [31].

Computational setup and code: All experiments were performed on an Intel i5-7500 CPU at 3.40GHz with 32GB RAM. MATLAB code for 3D versions of algorithm 3 are included in supplementary material along with sample *C. elegans* neuron point clouds.

4.1. Simulated data

Three dimensional source point set \mathbf{X} was generated by sampling $\mathbf{x}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3)$ for $j = 1, \dots, J$ where $J \in [20, \dots, 40]$. A random transformation β was obtained by computing the QR factorization of a 3×3 random gaussian matrix \mathbf{M} such that $\mathbf{QR} = \mathbf{M}$, taking the orthonormal rotation component \mathbf{Q} . This was randomly scaled by a factor between $s = [0.5, 1.5]$ so that $\beta = s\mathbf{Q}$. For $k \in [1, \dots, 19]$, $20 - k$ inlier target points were generated by transforming a random $20 - k$ subset of \mathbf{X} by β and adding gaussian noise with varying σ^2 : $\mathbf{Y}_{\mathcal{I}} = \mathbf{X}_{\pi(\mathcal{I})}\beta + \epsilon$. Furthermore, k points in \mathbf{Y} were randomly uniformly sampled from the convex hull of the $20 - k$ inlier points: $\mathbf{Y}_{\mathcal{O}} \sim \mathcal{U}[\mathcal{C}(\mathbf{Y}_{\mathcal{I}})]$. This procedure yielded two unordered multisets, $\mathbf{X} \in \mathbb{R}^{J \times 3}$ and $\mathbf{Y} \in \mathbb{R}^{20 \times 3}$. Using these unordered multisets as input to rRWOC, the regression coefficients $\hat{\beta}$ were estimated. If $\|\hat{\beta} - \beta\|_F \leq 1e - 3$, the event

was considered a correct recovery, otherwise a failure. The margin parameter ν was set to be $\nu = \sigma$. Also, using the randomized algorithm 3, the success probability parameter was set to $\delta = 0.9$.

This procedure was repeated 100 times for varying $k = 1, \dots, 19$, varying σ^2 and varying $J = 20, \dots, 40$ to assess the empirical recovery rate as a function of outlier amount, SNR, and missing correspondences in the target, respectively. The recovery rates vs. outlier ratio, and SNR can be seen in figure 2-middle. The recovery rates vs. missing data ratio and SNR can be seen in figure 2-left. Lastly, the comparison of the recovery rate of exhaustive and randomized rRWOC versus iterative closest point can be seen in figure 2-right.

These empirical results demonstrate that for a sufficiently high SNR and outlier ratio of less than 50%, the proposed algorithm yields almost perfect recovery rates. Furthermore, the comparisons with iterative closest point algorithm (ICP) show that rRWOC is much more robust to outliers than ICP since the inclusion of any outliers results in failure of ICP to recover the true transformation.

4.1.1 Unstructured vs. adversarial outliers

The evaluation of the compared algorithms in the presence of unstructured outliers as well as adversarial ones was done on the *fish* point cloud dataset. Experimental details are described in the supplementary material sections 3 and 4.




Figure 4. **Left: A:** Unaligned point sets of reference *C. elegans* neuron positions (red) and target neuron positions (green) **B:** Alignment with coherent point drift algorithm [25], **C:** Alignment with iterative closest point algorithm [9], **D:** Alignment with proposed algorithm 3. **Right:** Margin parameter (ν) estimation in the *C. elegans* dataset.

	Method	<i>TP</i>	<i>FP</i>	<i>FN</i>	<i>ACC</i>	<i>F1</i>	<i>PREC</i>	<i>REC</i>	<i>MD</i>
<i>C. elegans</i> Head	rRWOC	135±28	57±23	60±28	0.53±0.15	0.69±0.13	0.70±0.12	0.69±0.14	2.63±0.27
	ICP [9]	41±58	151±58	153±59	0.15±0.23	0.21±0.30	0.21±0.30	0.21±0.30	4.18±1.59
	CPD [25]	5±2	188±4	190±2	0.01±0.01	0.03±0.01	0.03±0.01	0.03±0.01	11.13±0.34
	HS [31]	110±24	70±23	80±23	0.45±0.15	0.60±0.13	0.50±0.09	0.53±0.11	3.2±0.34
<i>C. elegans</i> Tail	rRWOC	33±6	10±6	11±6	0.61±0.17	0.75±0.14	0.76±0.14	0.74±0.13	2.14±0.31
	ICP [9]	2±1	42±1	43±1	0.02±0.01	0.04±0.02	0.04±0.02	0.04±0.02	7.83±1.66
	CPD [25]	3±1	41±1	42±1	0.03±0.02	0.04±0.02	0.04±0.02	0.04±0.02	7.43±1.32
	HS [31]	36±4	9±4	10±5	0.65±0.13	0.78±0.12	0.72±0.13	0.82±0.12	1.9±0.32
Fish Unstructured	rRWOC	28 ± 13	18 ± 13	27 ± 13	0.42 ± 0.21	0.61 ± 0.29	0.67 ± 0.31	0.56 ± 0.26	0.12 ± 0.01
	ICP [9]	2 ± 1	45 ± 1	54 ± 1	0.02 ± 0.01	0.04 ± 0.02	0.05 ± 0.02	0.04 ± 0.02	0.23 ± 0.05
	CPD [25]	1 ± 2	46 ± 2	55 ± 2	0.01 ± 0.03	0.02 ± 0.06	0.02 ± 0.06	0.02 ± 0.05	0.12 ± 0.00
	HS [31]	14 ± 0	33 ± 0	42 ± 0	0.17 ± 0.01	0.30 ± 0.02	0.33 ± 0.02	0.28 ± 0.01	0.22 ± 0.00
Fish Adversarial	rRWOC	28 ± 13	18 ± 13	26 ± 13	0.43 ± 0.28	0.62 ± 0.28	0.67 ± 0.30	0.57 ± 0.26	0.14 ± 0.07
	ICP [9]	0 ± 18	47 ± 18	55 ± 18	0 ± 0.34	0 ± 0.40	0 ± 0.43	0 ± 0.37	0.30 ± 0.05
	CPD [25]	0 ± 0	47 ± 0	55 ± 0	0 ± 0.00	0 ± 0.01	0 ± 0.01	0 ± 0.01	0.10 ± 0.02
	HS [31]	15 ± 8	32 ± 8	40 ± 8.5129	0.19 ± 0.11	0.33 ± 0.18	0.36 ± 0.19	0.31 ± 0.17	0.19 ± 0.11

Table 1. Transformation recovery and permutation recovery by rRWOC, ICP, CPD and HS algorithms in the *C. elegans* and *fish* dataset. TP = true positive, FP = false positive, TN = true negative, FN = false negative, ACC = accuracy, F1 = F1 score, PREC = precision, REC = recall, MD = mean distance

In this scenario, we observed that rRWOC and HS both outperformed ICP, and CPD and performed similarly in the unstructured outlier scenario. However, in the adversarial outlier scenario rRWOC outperformed HS due to explicit modeling of the outliers and missing correspondences where HS only models one of these two corruptions. The quantitative results can be found in table 1 as well as in figures 1 and 2 in the supplementary material.

4.2. Neuron matching of *C. elegans*

For this application, we have used the publicly available *C. elegans* fluorescence imaging dataset of Nguyen et al. [26] found at <http://dx.doi.org/10.21227/H2901H> as well as the neuronal position dataset provided in [38]. The worm *C. elegans* is a widely known model organism for studying the nervous system due to the known

structural connectome of the 302 neurons it contains. The data provided 3D z-stack images of the head of 14 worms that each consists of approximately 185 to 200 neurons captured under confocal microscopy using fluorescent tagged protein GFP. In figure 3, the depth-colored 2D projection of an image frame can be seen superimposed with annotation points delineating the locations of neurons. Figure 3 also highlights the need for a method of matching and aligning worm point clouds that is robust to outliers or missing associations. Here, we define outliers as points where there is no neuron present and define missing data as neurons with no detection present.

Of the 14 datasets of the head neurons of *C. elegans* worms, random pairs were drawn to be the source and target point sets. From the remaining worms, the positional covariance of each neuron was estimated using the supervised




Figure 5. *C. elegans* neuron identification problem. **A**: Source point cloud of neuron locations and positional covariances. **B**: The target image with outliers (indicated by large non-neuronal green objects) and the resulting neuron identifications.

alignment method of [13]. Since the positional variance of each neuron was uniquely identified using training data, we used variable margin parameters for rRWOC such that $\nu_l = \max_{i=1,2,3} \lambda_i(\Sigma_l)$ where Σ_l is the covariance matrix of the l th neuron and $\lambda_i(\cdot)$ denotes the i th eigenvalue. Randomized RWOC (algorithm 3) was deployed with $\delta = 0.9$. The results were compared with iterative closest point (ICP) [6] as well as coherent point drift (CPD) [25] algorithms.

The demonstration of the *C. elegans* application of rRWOC is seen in figure 5. Here the source point set is a statistical atlas neuron positions [38] and the target point set is neuron detections which may be corrupted by non-neuronal outliers. The outcome is that the detected neurons are identified correctly using the proposed algorithm.

The recovery rates in terms of recovering the transformation β^* as well as the permutation Π^* , are summarized in table 1. In general, rRWOC was able to recover both the transformation and permutation better than ICP and CPD, which tend to be initialization-dependent as well as HS which is a global method. In all of the experiments, ICP and CPD were initialized with random rotation. rRWOC and HS are invariant to initialization since they are not descent-based methods. HS performs slightly better in the tail of *C. elegans* than rRWOC since the tail dataset tends to have fewer outliers which HS is more sensitive to. Contrarily, rRWOC does better than HS in the head since there are more outliers and missing correspondences.

4.3. Discussion

In this paper, we expanded on the linear regression without correspondence model [33, 1, 16, 27] to account for missing data and outliers. Furthermore, we provided sev-

eral exact and approximate algorithms for the recovery of regression coefficients under noiseless and noisy regimes. The proposed algorithms are combinatorial at worst with variable dimensions. However, randomization procedures make the average-case complexity in constant dimension tractable given enough tolerance for failure. We provided several theoretical guarantees for exact recovery and running time complexity. Furthermore, we empirically demonstrated the recovery rates of the proposed algorithms in simulated and biological data. A future algorithmic direction is to employ branch and bound techniques found in [31] to reduce the computational complexity of the brute force nature of the algorithms.

One of the crucial parameters in our algorithm is the margin parameter of ν that heavily influences the computation of the inlier sets. In geometric terms, the ν parameter should correspond to a physical radius such that if a target point is within that radius relative to a reference point, it is more likely to be a match than not be one. In the case of *C. elegans* where we have access to a statistical atlas that describes the positional covariance of each neuron (see figure 5A), we can quantify this radius as the half-width full maximum (HWF) of the Gaussian distribution of neuron positions given by the relation $f = \sigma\sqrt{2\log 2}$ where σ is the average neural positional standard deviation. The cross-validation experiments in figure 4 demonstrate that the optimal ν does indeed correspond to this value.

Acknowledgements: We acknowledge the following funding sources: NSF NeuroNex Award DBI-1707398, The Gatsby Charitable Foundation, NIBIB R01 EB22913, DMS 1912194, Simons Foundation Collaboration on the Global Brain.

References

- [1] Abubakar Abid, Ada Poon, and James Zou. Linear regression with shuffled labels. *arXiv preprint arXiv:1705.01342*, 2017.
- [2] Dror Aiger, Niloy J Mitra, and Daniel Cohen-Or. 4-points congruent sets for robust pairwise surface registration. In *ACM transactions on graphics (TOG)*, volume 27, page 85. Acm, 2008.
- [3] Michel A Audette, Frank P Ferrie, and Terry M Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical image analysis*, 4(3):201–217, 2000.
- [4] Jean-Charles Bazin, Yongduek Seo, Richard Hartley, and Marc Pollefeys. Globally optimal inlier set maximization with unknown rotation and focal length. In *European Conference on Computer Vision*, pages 803–817. Springer, 2014.
- [5] Edwin F Beckenbach and Richard Bellman. *Inequalities*, volume 30. Springer Science & Business Media, 2012.
- [6] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992.
- [7] Greg Bubnis, Steven Ban, Matthew D DiFranco, and Saul Kato. A probabilistic atlas for cell identification. *arXiv preprint arXiv:1903.09227*, 2019.
- [8] Alvaro Parra Bustos, Tat-Jun Chin, Frank Neumann, Tobias Friedrich, and Maximilian Katzmann. A practical maximum clique algorithm for matching with pairwise constraints. *arXiv preprint arXiv:1902.01534*, 2019.
- [9] Dmitry Chetverikov, Dmitry Svirko, Dmitry Stepanov, and Pavel Krsek. The trimmed iterative closest point algorithm. In *Object recognition supported by user interaction for service robots*, volume 3, pages 545–548. IEEE, 2002.
- [10] David L Donoho et al. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [11] Bertram Drost, Markus Ulrich, Nassir Navab, and Slobodan Ilic. Model globally, match locally: Efficient and robust 3d object recognition. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 998–1005. Ieee, 2010.
- [12] Olof Enqvist, Klas Josephson, and Fredrik Kahl. Optimal correspondences from pairwise constraints. In *2009 IEEE 12th international conference on computer vision*, pages 1295–1302. IEEE, 2009.
- [13] Georgios Dimitrios Evangelidis and Radu Horaud. Joint alignment of multiple point sets with batch and incremental expectation-maximization. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1397–1410, 2018.
- [14] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [15] Anders Hast, Johan Nysjö, and Andrea Marchetti. Optimal ransac-towards a repeatable algorithm for finding the optimal set. 2013.
- [16] Daniel J Hsu, Kevin Shi, and Xiaorui Sun. Linear regression without correspondence. In *Advances in Neural Information Processing Systems*, pages 1531–1540, 2017.
- [17] Piotr Indyk, Rajeev Motwani, and Suresh Venkatasubramanian. Geometric matching under noise: Combinatorial bounds and algorithms. In *SODA*, pages 457–465, 1999.
- [18] Sandy Irani and Prabhakar Raghavan. Combinatorial and experimental results for randomized point matching algorithms. *Computational Geometry*, 12(1-2):17–31, 1999.
- [19] Roy Jonker and Ton Volgenant. Improving the hungarian assignment algorithm. *Operations Research Letters*, 5(4):171–175, 1986.
- [20] Dagmar Kainmueller, Florian Jug, Carsten Rother, and Gene Myers. Active graph matching for automatic joint segmentation and annotation of c. elegans. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 81–88. Springer, 2014.
- [21] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [22] Haggai Maron and Yaron Lipman. (probably) concave graph matching. In *Advances in Neural Information Processing Systems*, pages 408–418, 2018.
- [23] Nicolas Mellado, Dror Aiger, and Niloy J Mitra. Super 4pcs fast global pointcloud registration via smart indexing. In *Computer Graphics Forum*, volume 33, pages 205–215. Wiley Online Library, 2014.
- [24] David M Mount, Nathan S Netanyahu, and Jacqueline Le Moigne. Efficient algorithms for robust feature matching. *Pattern recognition*, 32(1):17–38, 1999.
- [25] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.
- [26] Jeffrey P Nguyen, Ashley N Linder, George S Plummer, Joshua W Shaevitz, and Andrew M Leifer. Automatically tracking neurons in a moving and deforming brain. *PLoS computational biology*, 13(5):e1005517, 2017.
- [27] Ashwin Pananjady, Martin J Wainwright, and Thomas A Courtade. Linear regression with an unknown permutation: Statistical and computational limits. In *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 417–424. IEEE, 2016.
- [28] Jonathan Pokrass, Alexander M Bronstein, Michael M Bronstein, Pablo Sprechmann, and Guillermo Sapiro. Sparse modeling of intrinsic correspondences. In *Computer Graphics Forum*, volume 32, pages 459–468. Wiley Online Library, 2013.
- [29] Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. *IEEE transactions on visualization and computer graphics*, 19(7):1199–1217, 2013.
- [30] Philip HS Torr and Andrew Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding*, 78(1):138–156, 2000.

- [31] Manolis Tsakiris and Liangzu Peng. Homomorphic sensing. In *International Conference on Machine Learning*, pages 6335–6344, 2019.
- [32] Jayakrishnan Unnikrishnan, Saeid Haghighatshoar, and Martin Vetterli. Unlabeled sensing: Solving a linear system with unordered measurements. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 786–793. IEEE, 2015.
- [33] Jayakrishnan Unnikrishnan, Saeid Haghighatshoar, and Martin Vetterli. Unlabeled sensing with random linear measurements. *IEEE Transactions on Information Theory*, 64(5):3237–3253, 2018.
- [34] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011.
- [35] Heng Yang and Luca Carlone. A polynomial-time solution for robust registration with extreme outlier rates. *arXiv preprint arXiv:1903.08588*, 2019.
- [36] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *arXiv preprint arXiv:2001.07715*, 2020.
- [37] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE transactions on pattern analysis and machine intelligence*, 38(11):2241–2254, 2016.
- [38] Eviatar Yemini, Albert Lin, Amin Nejatbakhsh, Erdem Varol, Ruoxi Sun, Gonzalo E Mena, Aravinthan DT Samuel, Liam Paninski, Vivek Venkatachalam, and Oliver Hobert. Neuropal: A neuronal polychromatic atlas of landmarks for whole-brain imaging in *c. elegans*. *bioRxiv*, page 676312, 2019.
- [39] Ji Zhang and Sanjiv Singh. Visual-lidar odometry and mapping: Low-drift, robust, and fast. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2174–2181. IEEE, 2015.
- [40] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *European Conference on Computer Vision*, pages 766–782. Springer, 2016.