# Learning Probabilistic Piecewise Rigid Atlases of Model Organisms via Generative Deep Networks

Amin Nejatbakhsh<sup>1</sup>, Neel Dey<sup>2</sup>, Vivek Venkatachalam<sup>3</sup>, Eviatar Yemini<sup>4</sup>, Liam Paninski<sup>1</sup>, and Erdem Varol<sup>6</sup>

<sup>1</sup> Departments of Neuroscience and Statistics, Columbia University <sup>2</sup> Computer Science and Artificial Intelligence Lab, MIT

Computer Science and Artificial Intelligence Lab, MI  $^{3}$ 

<sup>3</sup> Department of Physics, Northeastern University

<sup>4</sup> Department of Neurobiology, University of Massachusetts Chan Medical School <sup>5</sup> Department of Computer Science and Engineering, New York University

Abstract. Atlases are crucial to imaging statistics as they enable the standardization of inter-subject and inter-population analyses. While existing atlas estimation methods based on fluid/elastic/diffusion registration yield high-quality results for the human brain, these deformation models do not extend to a variety of other challenging areas of neuroscience such as the anatomy of C. elegans worms and fruit flies. To this end, this work presents a general probabilistic deep network-based framework for atlas estimation and registration which can flexibly incorporate various deformation models and levels of keypoint supervision that can be applied to a wide class of model organisms. Of particular relevance, it also develops a deformable piecewise rigid atlas model which is regularized to preserve inter-observation distances between neighbors. These modeling considerations are shown to improve atlas construction and key-point alignment across a diversity of datasets with small sample sizes including neuron positions in C. elegans hermaphrodites, fluorescence microscopy of male C. elegans, and images of fruit fly wings. Code is accessible at https://github.com/amin-nejat/Deformable-Atlas.

# 1 Introduction

Constructing biological atlases via image registration helps summarize normative patterns and variability within a target population. An atlas also provides a common coordinate system for image registration and segmentation, which can help decouple and quantify different sources of variability observed in the data [15,18,22,23]. However, while atlas estimation of structures such as the human brain are well served with existing registration techniques [4,9,10,12,19], this work argues that other domains of neuroscience require alternative models.

**Motivation.** While fluid, elastic, and diffusion-based deformation models are well-motivated for the human brain, other structures may benefit from piecewise rigid deformations. For example, fluorescence microscopy of the nematode C. elegans nervous system is of high interest for atlas construction [7,20,28,31,33],



Fig. 1: Schematic of generative model of atlas construction. A: Each observation  $(\mathbf{Z}_i)$  is modeled as a random draw from an atlas parametrized by  $\boldsymbol{\theta}$  and perturbed by transformation  $f_{\boldsymbol{\beta}_i}^{-1}$ . B: We infer atlas parameters  $(\boldsymbol{\theta})$  from observations  $(\mathbf{X}_i)$  by optimizing a neural network loss function that penalizes the distance of each transformed observation  $(\mathbf{Z}_i)$  to the latent atlas  $(\boldsymbol{\theta})$  and also learn the transformation model parameters  $(\boldsymbol{\beta}_i)$  that minimizes the loss.

where the shape and position of each neuron may individually deform rigidly while the number and function of neurons are conserved across individuals. Another model organism whose morphometry is of interest to neuroscience is the fruit fly *Drosophila melanogaster* [17]. In the fly, one suitable structure for atlas building is the wing whose inter-fly deformation can be well approximated by piece-wise rigid motion. However, probabilistic atlases that explore the structural variability amongst wings of different phenotypes and sexes has not yet been established. Further, atlas-building methods are *typically* hand-tailored to accommodate the specifications of a single organism, a single experimental condition, or the specifics of a developing/degenerating population [24], and are thus cannot be generally repurposed for experimentalists that require atlases for novel biological datasets that they curate [16].

**Contributed methods.** This paper provides a general probabilistic framework for building atlases for any model organism using point clouds and/or images. Herein, we demonstrate the utility for two disparate organisms: nematodes and the fruit flies. We model individual observations, e.g., neural positions in a particular worm or the wing shape in an individual fly, as drawing from a generative atlas perturbed by a deformation jointly estimated via registration networks. The proposed framework allows for incorporating arbitrary transformations and atlas distributions with an emphasis on piece-wise rigid transformations wellsuited to the considered model organisms. Further, we develop regularizers that encourage the conservation of inter-keypoint distances within an observation to prevent self-intersections. Lastly, it allows for flexible supervision-levels, ranging from full manual annotations of neural positions (supervised) to partial annotations (semi-supervised), and no annotations (unsupervised). **Experimental results.** We show three applications of our framework to model organisms under varying levels of supervision. First, we construct rigid and piecewise-rigid atlases of C. elegans neural positions from fully supervised datasets [33]. Then we build a semi-supervised atlas of male C.elegans using pixel intensities and partial annotations of neural positions [29]. Lastly, we show fully unsupervised atlas building of fruit fly wings from natural images [26]. To our knowledge, this work presents the first piecewise rigid atlases of C. elegans neural point clouds, C. elegans images, and D. melanogaster wing images.

The proposed procedure can provide valuable insights into the appropriateness of the transformation models. For example, the contained analyses address whether worm posture can be modeled using rigid or piecewise rigid motion and the constructed fly wing atlases enable morphometric comparisons across genotypes.

## 1.1 Related Work

C. elegans statistical atlases. A number of atlases of neural positions in the C. elegans hermaphrodite have been introduced, utilizing a variety of shape and pose models [6,21,25,30,31]. However, atlas construction of the male C. elegans nervous system introduces further challenges due to having more neurons and higher density [14], additional ganglia enclosing these neurons [27], and greater variability in their neuronal and gangliar positions [29]. Therefore, existing atlas models of hermaphrodite neuron positions do not necessarily generalize to males.

**Piecewise rigid registration.** Piecewise rigid deformations are of key interest when modeling articulated structures (such as bones [32], joints, or *C. elegans* neurons) where local movements may be linear and global motion may be deformable. Theoretical treatments for such registration models have appeared in [1,2,8], yet to our knowledge, this paper is the first to construct atlases with such models, especially in a probabilistic deep network framework.

Atlas building. In practice, several atlas building methods alternate between registering observations to a template and updating the template with a pointwise intensity and/or shape average of the aligned observations [3,4,19]. More recent deep network-based methods [9,12,13] instead explicitly synthesize templates via regularized registration objectives without averaging. This latter method typically yield sharper and more interpretable estimates, and this is the approach that we follow and modify towards probabilistic atlas construction with deformation models well-suited for model organisms.

Most relevant to our work, atlas construction via statistical inference with deep networks has been done for diffusion-regularized registration for large-scale 3D neuroimages in [9]. The proposed framework is distinct in that: (1) We specifically use piecewise rigid deformations for atlas construction on model organisms. As our flow fields are explicitly constrained, we do not require diffusion regularization; (2) Our framework constructs atlases across varying levels of keypoint supervision (un/semi/fully-supervised) which is crucial for model organisms; (3)

3

4 A. Nejatbakhsh et al.

We develop an inter-keypoint distance conserving regularization to avoid selfintersections; (4) The datasets considered here are of much lower sample sizes and do not admit training the large registration and synthesis networks of [9].

## 2 Methods

Setup. We denote the atlas as a latent variable  $Z \in \mathbb{R}^D$  following the distribution  $P_{\theta}(Z)$  (Fig.1A). Both X (the observations) and Z random variables can be high-dimensional or low dimensional depending on the application. For example, an atlas constructed using point clouds is lower dimensional than an atlas that is constructed using image intensities [31]. Given the atlas, i.e. a distribution over the random variable Z, the observations  $X_i$  are samples from the prior  $Z_i$  that are warped by transformation  $f_{\beta_i} \in \mathcal{F}$  where  $\mathcal{F}$  is a function class of feasible deformations between the atlas and observations and  $\beta_i$  are the parameters of the deformation for the *i*th observation. In this work,  $\mathcal{F}$  is the space of rigid or piecewise rigid transformations.

**Inference and optimization.** With  $\mathcal{F}$  and the functional form of  $P_{\theta}$  prespecified, our goal is to solve the inverse problem for parameters  $\theta, \beta_{1:n}$ . In general terms, we write a probabilistic cost function informed by our statistical model and optimize it w.r.t.  $\theta, \beta_{1:n}$  as:

$$\begin{split} \mathcal{L}(\boldsymbol{\theta},\boldsymbol{\beta}_{1:n}) &= \log P_{\boldsymbol{\theta},\boldsymbol{\beta}_{1:n}}(\boldsymbol{X}_{1:n},\boldsymbol{Z}_{1:n}) = \sum_{i=1}^{n} \log P_{\boldsymbol{\beta}_{i}}(\boldsymbol{X}_{i}|\boldsymbol{Z}_{i}) + \log P_{\boldsymbol{\theta}}(\boldsymbol{Z}_{i}) \\ & \text{where } \boldsymbol{Z}_{i} \sim P_{\boldsymbol{\theta}}(\boldsymbol{Z}) \quad \text{and} \quad \boldsymbol{X}_{i}|\boldsymbol{Z}_{i} \sim P(\boldsymbol{X}|\boldsymbol{Z}_{i}) = f_{\boldsymbol{\beta}_{i}}(\boldsymbol{Z}_{i}) + \epsilon_{i} \end{split}$$

We take an alternating approach for optimizing  $\mathcal{L}$  where we iteratively optimize  $\mathcal{L}$  w.r.t.  $\boldsymbol{\theta}$  and  $\boldsymbol{\beta}_{1:n}$ . Given our estimate of the values  $\boldsymbol{\beta}_{1:n}$  denoted by  $\hat{\boldsymbol{\beta}}_{1:n}$  we find the best fit  $\hat{\boldsymbol{\theta}}$  to the data in the following way:

$$\hat{\boldsymbol{\theta}} = \max_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta} | \hat{\boldsymbol{\beta}}_{1:n}) = \max_{\boldsymbol{\theta}} \sum_{i=1}^{n} \log P_{\boldsymbol{\theta}}(f_{\hat{\boldsymbol{\beta}}_{i}}^{-1}(X_{i}))$$
(1)

Notice that here we are trying to find the sufficient statistics of  $P_{\theta}$  from known observations  $\mathbb{Z}_{1:n}$ . For the case of multivariate normal distributions where  $\theta = \{\mu, \Sigma\}$ , the maximum likelihood estimate (MLE) is the empirical mean and covariance. However, our formulation allows for incorporating arbitrarily complex distributions where we solve the MLE problem using stochastic variational inference in the parameter space. This is facilitated by probabilistic programming where generic algorithms for MLE and MAP estimation are provided.

On the other hand if we have a reasonable estimate of  $\theta$  then in order to update our estimates of  $\beta_{1:n}$  we need to solve the following for each *i*:

$$\hat{\boldsymbol{\beta}}_{\boldsymbol{i}} = \max_{\boldsymbol{\beta}_i \in \mathcal{F}} \mathcal{L}(\boldsymbol{\beta}_i | \hat{\boldsymbol{\theta}}) = \max_{\boldsymbol{\beta}_i \in \mathcal{F}} \log P_{\hat{\boldsymbol{\theta}}}(f_{\boldsymbol{\beta}_i}^{-1}(X_i))$$
(2)

Analytical solutions might exist for specific  $\mathcal{F}$ , but in general specific algorithms are required for particular choices of  $\mathcal{F}$  as demonstrated in the experiments.

Amortized learning of transformations. Instead of optimizing this loss function with respect to the parameters  $\{\beta_{1:n}\}$  directly, we solve an alternative amortized optimization problem. We parameterize  $\beta_i = nn_{\phi}(X_i)$  where  $nn_{\phi}$  is a neural network with weights  $\phi$  and optimize the loss w.r.t. the parameters of the neural net. We parameterize rigid and piecewise rigid transformations using 3 angles and 3 translation parameters per piece. The optimization is performed using the Adam optimizer. The updates of  $\beta_{1:n}$  are performed by backpropagating the gradients of  $\phi$  while the maximum likelihood estimation of  $\theta$  is performed with the stochastic variational inference module of Pyro [5].

A full instantiation of the model requires an observation model  $P(\mathbf{X}|\mathbf{Z})$ , transformation parameters  $\boldsymbol{\beta}_{1:n}$ , transformation function  $f_{\boldsymbol{\beta}}$ , transformation regularization  $\mathcal{R}(\boldsymbol{\beta})$ , and a prior model  $P_{\boldsymbol{\theta}}(\mathbf{Z})$ . In the experiments section, we provide various instantiations of the model consistent with the assumption of their corresponding datasets, as summarized in Table 1. For each instantiation, the loss is determined by the choice of its components and is used to update the transformation  $(\boldsymbol{\beta}_{1:n})$  and prior parameters  $(\boldsymbol{\theta})$ .

## 3 Experiments

Our experiments are presented below and are split into fully-supervised, semisupervised, and unsupervised settings acting on point cloud and/or image representations. To benchmark registration error when ground truth landmarks are available, we use the commonly used target to registration error (TRE) which measures the  $l_2$  distance between moved keypoints and target keypoints.

#### 3.1 Supervised atlas of hermaphrodite C. elegans neuron positions

We used a public dataset of five point clouds of hermaphrodite *C. elegans* tail neurons with 42 neurons per worm [33]. Each neuron in worm *i* has a 3D location denoted by  $\mathbf{p}_{i,n} \in \mathbb{R}^3$  and an RGB color denoted by  $\mathbf{c}_{i,n} \in \mathbb{R}^3$ . Therefore, the observation for each worm consists of positions and colors of all its neurons  $\mathbf{X}_i = \{(\mathbf{p}_{i,n}, \mathbf{c}_{i,n}), n = 1, \dots, 42\}$ . Every neuron in the worm body corresponds to a ganglion, therefore for every neuron in the dataset we also have a label determining which ganglion that neuron corresponds to. The ganglia provide a natural grouping of the neurons that move together in space, making our proposed PR model suitable for the registration of the point clouds. Our regularization ensures that the distances between neighboring ganglia are approximately preserved after alignment (Eq. 14) as depicted in Fig. 2B.

**Statistical model:** We experiment with two different deformation classes for the spatial component of the point clouds, namely rigid (R) and regularized piecewise rigid (PR) warps (Eq. 9). The function class for the color component is a simple **softmax** operator normalizing the transformed RGB colors to sum to one (Eq. 11). For the prior distribution over the positions and colors, we considered two statistical models. We chose the prior distribution over the positions

6 A. Nejatbakhsh et al.



Fig. 2: Supervised positional and color atlas of tail neurons of hermaphrodite *C. elegans.* A: Model description for point clouds of neuron positions and colors in 3 dimensions. An MLP processes the input data and outputs piecewise rigid (PR) parameters  $\beta_i$  per piece. B: Motivation of PR for worm point clouds. PR aligns ganglia (denoted by  $g_1, g_2, g_3$ ) between the atlas (gray worm) and observed point clouds (blue worms) while maintaining the distances between the neighboring ganglia. C: Train errors and cross-validated TRE ( $\pm$  standard error) for 4 different models (PR: piecewise rigid, R: rigid, D: Dirichlet, N: normal). Notice that all models exhibit overfitting due to the small sample size. D, E: Learned atlas using PR-Dirichlet (D) and PR-Normal (E) models. Small dots indicate individuals' neural positions, larger dots indicate mean positions in the atlas and ellipses indicate one standard deviation of mass. The PRD model projects the mean colors into 3 distinct ones while the PRN model preserves a more detailed description of mean colors.



Fig. 3: Semi-supervised image atlas of tail neurons of hermaphrodite and male *C. elegans.* A: Model description for image data. Images are pushed through a CNN which regresses PR transformation parameters and flow field. The flow field is regularized using the Jacobian of the transformation on a grid of points and the position of the landmark points. Transformations ( $\beta_{1:n}$ ) and atlas ( $\theta$ ) parameters are optimized under the log-likelihood cost (red box). B: Out-ofsample alignment error decreases with increasing number of neural annotations for the male dataset. The hermaphrodite dataset exhibits less deformations and hence the TRE does not improve by adding more annotations. C: Superposition of unaligned NeuroPAL [33] strain hermaphrodite worms used as input. D, E: Atlases constructed using a widely-adopted neuroimaging atlas estimation method (SyGN [4]), highlighting how human brain-specific models do not extend to the considered model organisms. F-H: Results from our semi-supervised piecewise rigid atlas estimation framework with a varying number of annotations per worm, showing improved alignment and realistic deformations.

to be multivariate normal as suggested by previous work [6,31]. However, for the color distribution we experimented with Normal and Dirichlet distributions<sup>6</sup> (Eq. 16, 17). For registration amortization, we use a fully connected architecture for  $\phi$  (Fig. 2A).

**Results.** In Fig. 2, we illustrate the atlas parameters  $\theta$  and aligned point clouds  $Z_{1:n}$  as well the uncertainties and the training and testing errors. We learned the atlas and the registration network using 4 worms and tested whether the registration model is capable of aligning the test image to the learned atlas. In Fig. 2 the training TRE refers to the error between the atlas points and aligned

<sup>&</sup>lt;sup>6</sup> Dirichlet is an appropriate choice for color distribution as its samples sum to one and prior information about expected color can be encoded in its parameters  $\alpha$ .

points corresponding to the training worms while the test TRE measures this error for the test worms when aligned using the trained registration model. As we are using a very small sample size (n = 4) we expect the test error to be higher than the training error. Our results suggest that the rigid model with Dirichlet color distribution (RD) achieves the best test error (5-fold cross-validated) in the fully-supervised point cloud setting but all 4 models achieve comparable performances in terms of training error.

## 3.2 Semi-supervised atlas of C. elegans images w/ partial keypoints

We now showcase the flexibility of our framework in a semi-supervised setting by applying it to images (instead of point clouds) and using partial landmark annotations to guide the transformation. We use a dataset of 5 images of hermaphrodite *C. elegans* and 12 images of male *C. elegans* along with landmarks corresponding to the locations of neurons in those images.

Statistical model. Formally,  $\mathbf{X}_i = {\mathbf{I}_i, \mathbf{p}_{i,n}}$  where  $\mathbf{I}_i \in \mathbb{R}^{W \times H \times D \times C}$ ,  $n = 1, \ldots, 42, \mathbf{p}_{i,n} \in \mathbb{R}^3$ , and where W, H, D, C are the width, height, depth, and the number of channels. Here, PR transformation operates both on images and landmarks (Eq. 10). The semi-supervision is achieved by regularizing the transformation parameters to align the landmarks (Eq. 12). We have further regularization using the Jacobian of the transformation to ensure the feasibility of the transformation (Eq. 15). The prior distribution is the image space is pixel-wise standard normal (Eq. 18). Instead of a fully connected architecture for  $\boldsymbol{\phi}$  here we used a convolutional neural network to extract image features that are useful for registration. The pictorial description of the model is illustrated in Fig. 3A.

**Results.** We use a subset of keypoints for registration and atlas construction and hold out other key points for benchmarking generalization. The test keypoints are chosen to cover different parts of the image to provide a full picture of the registration quality (Fig. 3). We vary the level of semi-supervision by using  $\{6, 11, 17\}$  landmarks per image and observed that test TRE drops with more landmarks as expected (Fig. 3B). Further, the alignment parameters lead to more biologically feasible transformations when we include more landmarks (Fig. 3F-H). We then compare our method against the *de facto* standard deformable atlas construction technique for human brains (SyGN) [4] with and without affine prealignment. As SyGN does not make use of landmarks or piecewise rigid deformations, it fails to align the neurons (especially in lower density posterior regions) and does not yield biologically plausible atlases (Fig. 3D,E). We observe that for the hermaphrodite dataset the PR and R models achieve comparable performances. However, for the male dataset, where images contain denser subsets of neurons in smaller regions, the PR model outperforms R when more landmarks are included as expected.

## 3.3 Unsupervised atlas of transgenic D. melanogaster wings

We used a public dataset [26] of 128 2D fruit fly images from 4 genotypes (egfr, samw, star, tkv) and 2 sexes to infer a latent atlas image that represents an



Fig. 4: Unsupervised atlas of fruit fly wing: Atlas: We infer a latent canonical atlas wing in the pixel space without the use of any markers or annotations using 128 example images of fruit fly wings in varying poses. Pre-rigid Unregistered, Unregistered: pre-aligned using rigid transformation. The unregistered image is shown in panels before and after rigid alignment. egfr\_F\_L tkv\_M\_L: Averaging wing images of different genotypes and sexes enables a visual comparison of morphological differences between these groups. M vs. F: Pointwise t-statistics (q<0.05) between males and females yields a heatmap that shows that females have more mass in the the medial part the wing than males.

average wing corrected for postural differences by a piecewise rigid motion model. This dataset can be organized into pairs  $X_i = (I_i, Y_i)$  where  $I_i \in \mathbb{R}^{W \times H}$  and  $Y_i \in \{\text{efgr, samw, star, tkv}\} \times \{\text{male, female}\}.$ 

Statistical model. The observation model is the same as section 3.2 with landmarks removed, i.e., we only rely on image intensities for learning the atlas (Eq. 13). Similar to [9] our framework allows for incorporating genotypedependent parameters and learning conditional atlases (Eq. 19). To do this, we chose a prior of the form  $P_{\theta}(\mathbf{Z}|\mathbf{Y}_i) = \mathcal{N}(\mathbf{Z}; \boldsymbol{\mu}(\mathbf{Y}_i))$ .

**Results.** The resulting image atlas is illustrated in Fig. 4. Using the atlas coordinate framework, we performed a pixelwise t-test on the aligned wings of females and males to observe statistically significant differences in the wing tip density in the medial part of the wing. Furthermore, our results show morphological differences between genotypes which matches domain knowledge.

# 4 Discussion

*Limitations.* Some limitations exist in the presented work and will be addressed in the future: (1) The sample sizes considered in the *C. elegans* experiments are small and cannot yield broadly generalizable templates. As more data is publicly released, we will retrain our models to lower atlas bias. (2) Unlike our point cloud experiments, our image experiments do not yield pixelwise uncertainties. We will modify our atlas prior in future work to remedy this.

**Conclusions.** This work developed a general probabilistic framework to compute piecewise rigid point cloud and image atlases in novel imaging datasets of

	Setting (EqNo.)	Description
Data	Full (3) Semi (4) None (5)	$ \begin{aligned} &\{(\boldsymbol{p}_{i,n}, \boldsymbol{c}_{i,n}), n = 1, \dots, 42\} \\ &\{\boldsymbol{I}_i, \boldsymbol{p}_{i,n}, n = 1, \dots, 42\}  \boldsymbol{I}_i \in \mathbb{R}^{W \times H \times D \times C}  \boldsymbol{p}_{i,n} \in \mathbb{R}^3 \\ &(\boldsymbol{I}_i, \boldsymbol{Y}_i), \boldsymbol{I}_i \in \mathbb{R}^{W \times H}, \boldsymbol{Y}_i \in \{\text{efgr, samw, star, tkv}\} \times \{\text{male, female}\} \end{aligned} $
Obs.	Full (6) Semi (7) None (8)	$\begin{split} &\delta(\boldsymbol{p}_i; f_{\boldsymbol{\beta}_i^p}(\boldsymbol{p}_i))\delta(\boldsymbol{c}_i; f_{\boldsymbol{\beta}_i^c}(\boldsymbol{c}_i)) \exp(-\mathcal{R}(\boldsymbol{\beta}_i)) \\ &\delta(\boldsymbol{I}_i, \boldsymbol{p}_i; f_{\boldsymbol{\beta}_i}(\boldsymbol{I}_i, \boldsymbol{p}_i)) \exp(-\mathcal{R}(\boldsymbol{\beta}_i)) \\ &\delta(\boldsymbol{I}_i; f_{\boldsymbol{\beta}_i}(\boldsymbol{I}_i)) \exp(-\mathcal{R}(\boldsymbol{\beta}_i)) \end{split}$
Par.	Full (9) Semi/None (10)	$ \begin{array}{l} \boldsymbol{\beta}_{i}^{c} = \{\}  \boldsymbol{\beta}_{i}^{p} = \{\boldsymbol{T}_{i}^{1:K}, \boldsymbol{\Phi}_{i}^{1:K}\},  \boldsymbol{T}_{i}^{k} \in [-1, 1]^{3}, \boldsymbol{\Phi}_{i}^{k} \in [-\pi, \pi]^{3} \\ \boldsymbol{\beta}_{i} = \{\boldsymbol{T}_{i}^{1:K}, \boldsymbol{\Phi}_{i}^{1:K}\},  \boldsymbol{T}_{i}^{k} \in [-1, 1]^{3}, \boldsymbol{\Phi}_{i}^{k} \in [-\pi, \pi]^{3} \end{array} $
Func.	Full (11) Semi (12) None (13)	$ \begin{split} & f_{\boldsymbol{\beta}}(\boldsymbol{p},\boldsymbol{c}) = \left[ f_{\boldsymbol{\beta}_{p}}(\boldsymbol{p}), f_{\boldsymbol{\beta}_{c}}(\boldsymbol{c}) \right] = \left[ \mathtt{R}(\boldsymbol{\Phi}^{g(p)})\boldsymbol{p} + \boldsymbol{T}^{g(p)}, \mathtt{softmax}(\boldsymbol{c}) \right] \\ & f_{\boldsymbol{\beta}}(\boldsymbol{I},\boldsymbol{p}) = \left[ \mathtt{FF}_{\boldsymbol{\beta}} \circ \boldsymbol{I}, \mathtt{FF}_{\boldsymbol{\beta}}(\boldsymbol{p}) \right]  \mathtt{FF}(\boldsymbol{p}) = \sum_{k=1}^{K} w_{k}(\boldsymbol{p}) \big( \mathtt{R}(\boldsymbol{\Phi}_{k})\boldsymbol{p} + \boldsymbol{T}_{k} \big) \\ & f_{\boldsymbol{\beta}}(\boldsymbol{I}) = \mathtt{FF}_{\boldsymbol{\beta}} \circ \boldsymbol{I},  \mathtt{FF}(\boldsymbol{p}) = \sum_{k=1}^{K} w_{k}(\boldsymbol{p}) \big( \mathtt{R}(\boldsymbol{\Phi}_{k})\boldsymbol{p} + \boldsymbol{T}_{k} \big) \end{split} $
$\operatorname{Reg.}$	Full (14) Semi/None (15)	$ \begin{aligned} \mathcal{R}(\boldsymbol{\beta}_i) &= \sigma_{\boldsymbol{\beta}} \sum_{i,j \in \mathcal{G}} \left\  D(\boldsymbol{\mu}_i, \boldsymbol{\mu}_j) - D(\boldsymbol{\mu}_i^a, \boldsymbol{\mu}_j^a) \right\ ^2 \\ \mathcal{R}(\boldsymbol{\beta}) &= \sigma_{\boldsymbol{\beta}} \sum_{k=1}^{K} \det \mathcal{J}_{\boldsymbol{\beta}}(\boldsymbol{q}_k) \end{aligned} $
Prior	Full-1 (16) Full-2 (17) Semi (18) None (19)	$\begin{aligned} P_{\theta}(\boldsymbol{Z}) &= P(\boldsymbol{p})P(\boldsymbol{c}) = \mathcal{N}(\boldsymbol{p}; \boldsymbol{\mu}_{\boldsymbol{p}}, \sigma_{\boldsymbol{p}}\mathbf{I})\mathcal{N}(\boldsymbol{c}; \boldsymbol{\mu}_{\boldsymbol{c}}, \sigma_{\boldsymbol{c}}\mathbf{I}) & \text{MVN-MVN} \\ P_{\theta}(\boldsymbol{Z}) &= P(\boldsymbol{p})P(\boldsymbol{c}) = \mathcal{N}(\boldsymbol{p}; \boldsymbol{\mu}_{\boldsymbol{p}}, \sigma_{\boldsymbol{p}}\mathbf{I})\text{Dir}(\boldsymbol{c}; \boldsymbol{\mu}_{\boldsymbol{c}}) & \text{MVN-Dir} \\ P_{\theta}(\boldsymbol{I}_{i}, \boldsymbol{p}_{i}) &= \mathcal{N}(\boldsymbol{I}_{i}; \boldsymbol{\mu}^{I}, \sigma_{I}\mathbf{I})\mathcal{N}(\boldsymbol{p}_{i}; \boldsymbol{\mu}^{\boldsymbol{p}}, \sigma_{\boldsymbol{p}}\mathbf{I}) \\ P_{\theta}(\boldsymbol{Z} \boldsymbol{Y}_{i}) &= \mathcal{N}(\boldsymbol{Z}; \boldsymbol{\mu}(\boldsymbol{Y}_{i})) \end{aligned}$

**Table 1:** Instantiations of our model in Fully Supervised (Full), Semi-Supervised (Semi), and Unsupervised (None) settings. Each statistical model consists of the dataset (Data), observation model (Obs.), warp parameters (Par.), transformation function (Func.), warp regularization (Reg.), and prior/atlas distribution (Prior). Notation: I is an image, p is a 3D point, c is an RGB color, Y is a discrete label,  $\delta$  is the delta distribution,  $\beta$  denotes the transformation parameters,  $\Phi$  is a vector of rotation angles and R creates a rotation matrix with  $\Phi$ , T is the set of translations, FF is a flow field parameterized by  $\beta$ ,  $w_{1:K}$  are scalar weights summing to one,  $\mu$  is the centroid of positions within a ganglion,  $\mathcal{G}$  is the ganglion neighborhood graph, and  $\mathcal{J}$  is the Jacobian. We use g(p) to denote the ganglion that the neuron positioned at point p corresponds to. The positional parameters described above determine a piecewise rigid transformation in the spatial domain. The prior  $\mathcal{R}(\beta)$  regularizes warps and ensures smoothness across neighboring ganglia.

model organisms such as C. elegans and fruit flies. As new imaging modalities emerge to capture different views of the nervous systems of model animals, we expect that the flexibility of our framework will be valuable for standardizing downstream analyses. We also expect piecewise rigid atlas construction to apply to applications such as motion correction [11,34] and kinematics modeling [32].

Acknowledgements Paninski: NSF NeuroNex DBI-1707398, Gatsby Charitable Foundation, DMS 1912194, Simons Foundation Collab. on the Global Brain. Yemini: Klingenstein-Simons Fellowship in Neuroscience, Hypothesis Fund. Dey: NIH NIBIB NAC P41EB015902, NIBIB 5R01EB032708. Varol: 1K99MH128772-01A1. Venkatachalam: Burroughs Wellcome Fund and NIH R01 NS126334.

11

# References

- Arsigny, V., Commowick, O., Ayache, N., Pennec, X.: A fast and log-euclidean polyaffine framework for locally linear registration. Journal of Mathematical Imaging and Vision 33(2), 222–238 (2009)
- Arsigny, V., Pennec, X., Ayache, N.: Polyrigid and polyaffine transformations: A new class of diffeomorphisms for locally rigid or affine registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 829–837. Springer (2003)
- 3. Avants, B., Gee, J.C.: Geodesic estimation for large deformation anatomical shape averaging and interpolation. Neuroimage **23**, S139–S150 (2004)
- Avants, B.B., Yushkevich, P., Pluta, J., Minkoff, D., Korczykowski, M., Detre, J., Gee, J.C.: The optimal template effect in hippocampus studies of diseased populations. Neuroimage 49(3), 2457–2466 (2010)
- Bingham, E., Chen, J.P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P.A., Horsfall, P., Goodman, N.D.: Pyro: Deep universal probabilistic programming. J. Mach. Learn. Res. 20, 28:1–28:6 (2019)
- Bubnis, G., Ban, S., DiFranco, M.D., Kato, S.: A probabilistic atlas for cell identification (2019)
- Choe, K.P., Strange, K.: Molecular and genetic characterization of osmosensing and signal transduction in the nematode caenorhabditis elegans. The FEBS Journal 274(22), 5782–5789 (2007)
- Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P.Y., Ayache, N., Malandain, G.: An efficient locally affine framework for the smooth registration of anatomical structures. Medical Image Analysis 12(4), 427–441 (2008)
- Dalca, A., Rakic, M., Guttag, J., Sabuncu, M.: Learning conditional deformable templates with convolutional networks. Advances in neural information processing systems 32 (2019)
- 10. Davis, B.C., Fletcher, P.T., Bullitt, E., Joshi, S.: Population shape regression from random design data. International journal of computer vision **90**(2), 255–266 (2010)
- Dey, N., Messinger, J., Smith, R.T., Curcio, C.A., Gerig, G.: Robust non-negative tensor factorization, diffeomorphic motion correction, and functional statistics to understand fixation in fluorescence microscopy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 658–666. Springer (2019)
- Dey, N., Ren, M., Dalca, A.V., Gerig, G.: Generative adversarial registration for improved conditional deformable templates. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3929–3941 (2021)
- Ding, Z., Niethammer, M.: Aladdin: Joint atlas building and diffeomorphic registration learning with pairwise alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20784–20793 (2022)
- 14. Emmons, S.W., Sternberg, P.W.: Male development and mating behavior (2011)
- Greitz, T., Bohm, C., Holte, S., Eriksson, L.: A computerized brain atlas: construction, anatomical content, and some applications. Journal of computer assisted tomography 15(1), 26–38 (1991)
- Heckscher, E.S., Long, F., Layden, M.J., Chuang, C.H., Manning, L., Richart, J., Pearson, J.C., Crews, S.T., Peng, H., Myers, E., Doe, C.Q.: Atlas-builder software and the eNeuro atlas: resources for developmental biology and neuroscience. Development 141(12), 2524–2532 (06 2014)

- 12 A. Nejatbakhsh et al.
- Houle, D., Govindaraju, D.R., Omholt, S.: Phenomics: the next challenge. Nature reviews genetics 11(12), 855–866 (2010)
- Jones, A.R., Overly, C.C., Sunkin, S.M.: The allen brain atlas: 5 years and beyond. Nature Reviews Neuroscience 10(11), 821–828 (2009)
- Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23, S151–S160 (2004)
- Kaiser, M., Hilgetag, C.C.: Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. PLoS computational biology 2(7), e95 (2006)
- Long, F., Peng, H., Liu, X., Kim, S.K., Myers, E.: A 3d digital atlas of c. elegans and its application to single-cell analyses. Nature Methods 6(9), 667–672 (2009)
- Roland, P., Graufelds, C., Wăhlin, J., Ingelman, L., Andersson, M., Ledberg, A., Pedersen, J., Åkerman, S., Dabringhaus, A., Zilles, K.: Human brain atlas: for highresolution functional and anatomical mapping. Human Brain Mapping (1994)
- Scheffer, L.K., Meinertzhagen, I.A.: The fly brain atlas. Annual review of cell and developmental biology 35, 637–653 (2019)
- Schuh, A., Makropoulos, A., Robinson, E.C., Cordero-Grande, L., Hughes, E., Hutter, J., Price, A.N., Murgasova, M., Teixeira, R.P.A., Tusor, N., et al.: Unbiased construction of a temporally consistent morphological atlas of neonatal brain development. bioRxiv p. 251512 (2018)
- 25. Skuhersky, M., Wu, T., Yemini, E., Boyden, E., Tegmark, M.: Toward a more accurate 3d atlas of c. elegans neurons. bioRxiv (2021)
- Sonnenschein, A., VanderZee, D., Pitchers, W.R., Chari, S., Dworkin, I.: An image database of drosophila melanogaster wings for phenomic and biometric analysis. GigaScience 4(1), s13742–015 (2015)
- Sulston, J.E., Horvitz, H.R.: Post-embryonic cell lineages of the nematode, caenorhabditis elegans. Developmental biology 56(1), 110–156 (1977)
- Szigeti, B., Gleeson, P., Vella, M., Khayrulin, S., Palyanov, A., Hokanson, J., Currie, M., Cantarelli, M., Idili, G., Larson, S.: Openworm: an open-science approach to modeling caenorhabditis elegans. Frontiers in computational neuroscience 8, 137 (2014)
- Tekieli, T., Yemini, E., Nejatbakhsh, A., Wang, C., Varol, E., Fernandez, R.W., Masoudi, N., Paninski, L., Hobert, O.: Visualizing the organization and differentiation of the male-specific nervous system of c. elegans. Development (2021)
- 30. Toyoshima, Y., Wu, S., Kanamori, M., Sato, H., Jang, M.S., Oe, S., Murakami, Y., Teramoto, T., Park, C., Iwasaki, Y., Ishihara, T., Yoshida, R., Iino, Y.: An annotation dataset facilitates automatic annotation of whole-brain activity imaging of c. elegans. bioRxiv (2019). https://doi.org/10.1101/698241
- Varol, E., Nejatbakhsh, A., Sun, R., Mena, G., Yemini, E., Hobert, O., Paninski, L.: Statistical atlas of c. elegans neurons. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 119–129 (2020)
- 32. Wustenberg, R.: Carpal bone rigid-body kinematics by log-euclidean polyrigid estimation (2022)
- Yemini, E., Lin, A., Nejatbakhsh, A., Varol, E., Sun, R., Mena, G.E., Samuel, A.D., Paninski, L., Venkatachalam, V., Hobert, O.: Neuropal: a multicolor atlas for whole-brain neuronal identification in c. elegans. Cell 184(1), 272–288 (2021)
- 34. Yu, J., Nejatbakhsh, A., Torkashvand, M., Gangadharan, S., Seyedolmohadesin, M., Kim, J., Paninski, L., Venkatachalam, V.: Versatile multiple object tracking in sparse 2d/3d videos via diffeomorphic image registration. bioRxiv (2022)