

Comparative Analysis of Spatial Patterns of Gene Expression in *Drosophila melanogaster* Imaginal Discs

Cyrus L. Harmon¹, Parvez Ahammad², Ann Hammonds¹,
Richard Weiszmann³, Susan E. Celniker³, S. Shankar Sastry², and
Gerald M. Rubin¹

¹ Department of Molecular and Cell Biology,

University of California, Berkeley, Berkeley, CA 94720

² Department of Electrical Engineering and Computer Science,

University of California, Berkeley, Berkeley, CA 94720

³ Department of Genome Sciences,

Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Abstract. Determining the precise spatial extent of expression of genes across different tissues, along with knowledge of the biochemical function of the genes is critical for understanding the roles of various genes in the development of metazoan organisms. To address this problem, we have developed high-throughput methods for generating images of gene expression in *Drosophila melanogaster* imaginal discs and for the automated analysis of these images. Our method automatically learns tissue shapes from a small number of manually segmented training examples and automatically aligns, extracts and scores new images, which are analyzed to generate gene expression maps for each gene. We have developed a reverse lookup procedure that enables us to identify genes that have spatial expression patterns most similar to a given gene of interest. Our methods enable us to cluster both the genes and the pixels that of the maps, thereby identifying sets of genes that have similar patterns, and regions of the tissues of interest that have similar gene expression profiles across a large number of genes.

Primary keyphrases: Genomic imaging, Gene expression analysis, Clustering. **Secondary keyphrases:** Microarray data analysis, Imaginal discs

1 Introduction

We have developed a high-throughput pipeline for generating images of imaginal discs, stained with labeled probes that hybridize to individual genes, and an automated system for analyzing a database of such images and for building an atlas of gene expression patterns. Using these tools, we construct consensus shape models of individual imaginal disc types which are used to automatically extract imaginal disc shapes from new images and to align these features to consensus shape models. In parallel, we extract the stain pattern for each disc which is aligned to and overlaid on the consensus shape model, thereby yielding

a representation of the spatial extent of expression of a given gene in the context of the consensus shape. These patterns are averaged to produce consensus gene expression representations for each gene. The consensus representations are assembled in a database and clustered to identify both genes and regions of the discs with similar patterns of gene expression. New images can be used to search the database to identify known patterns similar to a query pattern, which can be automatically extracted from a new image. The data-flow of our approach is shown in Figure 1.

Previous work has identified the vast majority of the genes in *Drosophila* [1, 2], shown how a large numbers of genes are dynamically expressed across the entire organism throughout the life cycle [3], and shown how genes are expressed across space and time in embryos [4]. DNA-microarrays have enabled parallel, high-throughput analysis of gene expression from many different tissues under different conditions or at different time points [5, 6]. However, these techniques do not provide spatial information about where these genes are expressed within these tissues.

Image processing techniques have been applied to the problem of quantitative gene expression analysis by analyzing the hybridization of fluorescent probes to specific spots on DNA microarrays [6, 7], and, recently, fully-automated approaches to this problem have been developed [8].

Large-scale studies of patterns of gene expression in *Drosophila* have been performed using DNA microarrays both on whole organisms [9, 3] and individual tissues such as imaginal discs [10, 11]. Klebes et al. compared differential gene expression in different imaginal discs and between imaginal discs and non-disc tissue. Butler et al. manually dissected imaginal discs and were able to identify transcripts that were enriched in specific compartments of the wing discs [10]. However, these studies yield little information about the precise spatial patterns of gene expression.

2 Pipeline Overview

Our pipeline, shown in Figure 1, begins with the acquisition and purification of the biological material, followed by steps to select appropriate genes for inclusion in the project. The next phase of the pipeline entails generation of labeled probes for specific genes, hybridization of these probes to appropriate pools of imaginal discs, and imaging of the resulting stained discs. The data analysis portion of the pipeline consists of two main sections, a learning process in which a small number of manually segmented examples are used to learn the shapes of the imaginal discs and an automated analysis phase in which the learned shapes are used to automatically isolate instances of imaginal discs in the complete set of images. These aligned images are then scored for stain, and used to produce consensus maps for each gene. The consensus maps are then used for a variety of clustering and analysis procedures.

For the high-throughput pipeline, we needed a sufficient source of imaginal discs such that we would have a number of each type of imaginal disc for each

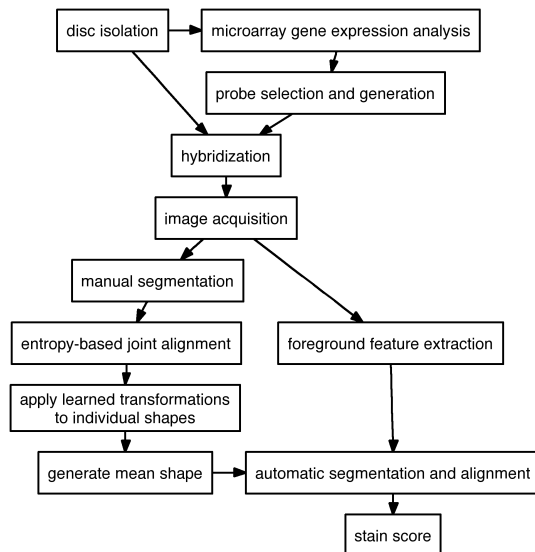


Fig. 1: Overview of the high-throughput process for determining spatial patterns of expression of genes in *Drosophila* imaginal discs.

probe. Moreover, it was desirable to be able to work in 96-well plates, following the protocol of the Berkeley *Drosophila* Genome Project (BDGP) *Drosophila* embryo gene expression pattern project [4]. We adapted the imaginal disc mass-isolation protocol from Eugene et al. [12] for use in a 96-well format. We collected eggs from Canton S stocks and grew larvae until the wandering third instar larval stage, at which point the larvae were harvested and the discs isolated. The mass-isolation procedure typically yielded 500,000 discs, including wing, leg, eye-antennal, haltere and genital discs, but not the smaller disc types, as these were not recovered in our purification process. Approximately 100,000 discs were used per 96-well plate, yielding on the order of 1000 discs per probe.

2.1 Gene Selection and Microarray Gene Expression Analysis

While it would be desirable to examine the spatial pattern of every gene, this is prohibitively expensive and time-consuming. For the current project, we had the resources to produce probes for only a few hundred genes. Many genes have no detectable spatial expression pattern, or, on the other hand, are ubiquitously expressed and we would like to avoid making probes for these genes. Therefore, we have developed a protocol to examine the over 13,000 genes in the *Drosophila* genome using gene expression microarrays to analyze the levels of expression across multiple tissues. The central idea is that by examining the quantitative level of gene expression across multiple tissues, including *Drosophila* embryos at multiple time points, distinct classes of imaginal discs, and adults, we could iden-

tify genes that were differentially expressed in imaginal discs and therefore more likely to have a non-trivial spatial pattern of gene expression than genes that were expressed at a constant level across all tissues. In particular, we identified genes that were expressed at a higher level in certain imaginal disc types, relative to the other disc types, or that were more highly expressed in all imaginal discs relative to the embryo and adult samples.

2.2 Disc Annotation Database

In order to facilitate the process of automatic segmentation, the operator records meta-data about the images, as they are being captured, in a database [4]. The annotations include the orientation of the peripodial epithelium and the handedness of the disc, from which we deduce whether or not the disc needs to be reflected about the vertical axis with respect to the canonical orientation.

All imaginal discs except the genital discs occur in pairs, one on each side of the larval body, yielding a left disc and a right disc of each type. When the discs are placed on a microscope slide for imaging, they can be found in one of two orientations, with the peripodial epithelium, which exists on only one side of the disc, either on the top or the bottom. The combination of the handedness of the disc and the orientation of the peripodial epithelium gives us 4 possibilities for the combined state and orientation of the disc. We make a simplifying assumption and assume that the left and right discs are mirror images of each other and that the stain pattern of a gene in a right disc will be equivalent to the mirror image of the stain pattern of that gene from the left disc of the same type. This assumption gives us two handedness/orientation combinations that are considered to be in the canonical orientation, and two combinations that are mirror images of these. The shapes corresponding to the two mirror image handedness/orientation combinations are automatically reflected about the vertical axis by the congealing pipeline after manual segmentation, to bring the images into the canonical orientation.

2.3 Shape Representation and Learning

The results of our manual segmentation process are a set of binary image masks that represent shapes from individual images. From these shapes we learn a shape representation in the form of a canonical binary image in which pixels are either 1 for foreground or 0 for background, or a canonical grayscale image in which pixels values are between 0 and 1 and where the value represents the probability that a given pixel is on in an element of that particular shape. This simple shape model for each imaginal disc type serves as the reference map upon which comparisons regarding patterns of gene expression will be made. While there are many possible approaches to this task, including parameterizing the curve that defines the borders of the shapes, we use a simple non-parametric method called Congealing [13, 14], that iteratively learns a set of affine transformations that minimize the overall pixelwise entropy of a stack of images to learn a canonical shape model for each disc type [15].




Fig. 2: (a) Raw RGB image file of a wing imaginal disc. (b) Discrete approximation of the second derivative of a wing imaginal disc image, produced by convolution with a Laplacian filter kernel. (c) The previous image after a round of morphological closing by dilation and erosion. (d) Resulting image after applying the 3-4 distance transform.




Fig. 3: (a) Canonical wing disc shape model. (b) 3-4 distance transform of the wing disc shape model.

3 Parallel Alignment and Automatic Feature Extraction

Our approach to identifying imaginal discs in images is based on extraction of simple foreground features and subsequent alignment of the foreground features to a shape model by optimizing the parameters of an affine transformation between the foreground features and the shape model. This approach was initially proposed by Barrow et al. [16] for use with manually extracted point features which were aligned to a traced boundary. In order to avoid becoming trapped in local minima during the optimization process, a process known as “Chamfering” is used to generate a map of the distance of foreground features to the nearest edge. Borgefors called this map of the distance to the foreground features a “Distance Transform” and proposed a hierarchical method for performing this style of alignment in a coarse-to-fine manner [17].

3.1 Foreground Feature Extraction

Foreground biological material has substantially more variability than the background. While the measured intensity values of the background change across the image, these values generally change smoothly. The biological material present in the image, on the other hand, contains substantial pixel intensity variability due to changes in the scattering and absorption of light by the material. Therefore, we use the absolute value of a discrete approximation of the second

derivative of the image, produced by discrete convolution with a laplacian filter kernel [18], to detect the foreground biological material. An example image is shown in Figure 2(b). After applying the Laplacian operation, the image is then morphologically closed by first dilating the image, then eroding the image using a square structuring element. The resulting image is shown in Figure 2(c).

The Distance Transform is a transformation of an image in which edge (or foreground) pixels are assigned a value of zero and non-edge (or non-foreground) pixels are assigned a value proportional to the distance to the nearest edge pixel [17, 19]. The true Euclidean distance is somewhat computationally expensive. Fortunately, Borgefors provides an efficient approach, known as the 3-4 Distance Transform, for approximating the true Euclidean distance using a two-pass, forward and backward, algorithm that analyzes the backward and forward 8-neighbors, respectively, of each pixel and computes an approximation of the distance of each pixel to the nearest edge pixel. A visual inspection of this approach and the L_1 -based approximation provided by Soille [19] shows that the 3-4 Distance Transform yields substantially better results.

Just as we apply the distance transform to the target image, we apply the distance transform to the model to produce a smoothed-out representation of the model, suitable for template matching. The results of the distance transformed model can be seen in Figure 3.

3.2 Distance Metrics

There are multiple instances in our pipeline where it is necessary to compare one template with another template, and assign a score based on a distance measure that measures how dissimilar they are to each other. Depending on the scenario, the template can be a real-valued vector or a real-valued or binary two-dimensional image patch. Common distance choices include the L_1 and L_2 distance metrics. However the L_1 and L_2 distances are not invariant to scalings and shifts in image intensities. Therefore, we used the Normalized Cross-Correlation distance[20] to compare images. To compare two image patches, X and Y , where $X, Y \in \mathbf{R}^{M \times N}$, we use the following formula:

$$\text{NCC}(X, Y) = \frac{\sum_{i=1}^M \sum_{j=1}^N (X_{i,j} - \bar{X})(Y_{i,j} - \bar{Y})}{\sqrt{\sum_{i=1}^M \sum_{j=1}^N (X_{i,j} - \bar{X})^2 \sum_{i=1}^M \sum_{j=1}^N (Y_{i,j} - \bar{Y})^2}} \quad (1)$$

where \bar{X} denotes the mean value of the template X . The value returned by NCC always ranges between -1 and +1, irrespective of the size of the template. When $\text{NCC}=+1$, the two templates match perfectly. If the templates X and Y take only binary values 0, 1, we define the set of pixels on in both templates as $A = \sum_{i=1}^M \sum_{j=1}^N \mathbf{1}(\mathbf{X}_{i,j} = \mathbf{1}, \mathbf{Y}_{i,j} = \mathbf{1})$, the sets of pixels on in either template as $B = \sum_{i=1}^M \sum_{j=1}^N \mathbf{1}(\mathbf{X}_{i,j} = \mathbf{1}, \mathbf{Y}_{i,j} = \mathbf{0})$, and $C = \sum_{i=1}^M \sum_{j=1}^N \mathbf{1}(\mathbf{X}_{i,j} = \mathbf{0}, \mathbf{Y}_{i,j} = \mathbf{1})$, and the set of pixels off in both templates as $D = \sum_{i=1}^M \sum_{j=1}^N \mathbf{1}(\mathbf{X}_{i,j} = \mathbf{0}, \mathbf{Y}_{i,j} = \mathbf{0})$ where $\mathbf{1}(p)$ is an indicator function that takes value 1 if the condition p is sat-

ified. The Jaccard metric can then be defined in terms of A, B, C, D .

$$\text{Jaccard}(X, Y) = \frac{A}{A + B + C} \quad (2)$$

The NCC measure is equivalent to the Jaccard metric when $D \gg (A + B + C)$. While the Jaccard metric explicitly avoids computing distance between the features that are absent in both X and Y templates – and is thus desirable, it doesn’t extend well to real-valued templates.

3.3 Coordinate Descent from Multiple Starting Configurations

Given the distance-transformed model, an image corresponding to the identified foreground features, and an appropriate distance metric, we seek an affine transformation that, when applied to the target image, minimizes the distance between the (affine) transformed target features and the model. We use a reasonable set of starting configurations, such as the identity affine transform, and 90 degree rotations of the identity transform, and perform coordinate descent on the parameters of the affine transformation. This procedure quickly yields an acceptable match between the target and the imaginal disc model over 85% of the time. The procedure can fail to find a good match when the image is cluttered or contains multiple discs. Poorly performing instances are manually removed from the atlas prior to further processing.

While the correct alignment between two images generally yields a local minimum, this is not the only minimum to be found in the space described by the parameters of the affine transformation. In congealing, we avoid settling on an incorrect local minimum which may be found by reducing the scale of both images towards zero or infinity, essentially matching two all black or all white images, by rescaling the affine transformations in each iteration of the algorithm. In the case of the distance transforms of the foreground features and the model, we occasionally find that a better score is found via an incorrect alignment of two images that have undergone substantial transformations to achieve this alignment. To address this problem, we constrain the affine transformations such that the log-space parameters for x- and y-scaling are never allowed to go below -0.35 or above 0.35, constraining the minimum size of the figure to be 70% of the original size and the maximum to be 142% of the original size. Similarly, the log scale x- and y-shear parameters are never allowed to go below -0.2 or above 0.2. Finally, as an additional constraint on the transformations, we never allow the absolute value of the difference between the x-scale and y-scale log-space parameters to go over 0.2. This drives the alignment procedure to local minima that are more likely to represent the true biologically relevant alignment, even if the procedure would otherwise find a lower-scoring alignment by using a radical deformation of the shapes.

3.4 Stain Scoring

Imaginal discs are stained with digoxigenin labeled mRNA or DNA probes. The labeled discs were then treated with anti-digoxigenin coupled to alkaline phos-




Fig. 4: (a) An image of *Drosophila melanogaster* wing imaginal discs with areas of heavy stain. (b) The result from the stain-scoring algorithm using Equation 3 producing clipped areas in the center where no stain is reported. (c) The improved stain-scoring result obtained by using Equation 4.

phatase, an enzyme that catalyzes a reaction that generates a blue dye from a colorless substrate. The presence of labeled probe is indicated by a blue color, resulting from the absorption of non-blue photons by the dye. We have developed a quantitative stain-scoring heuristic that uses the difference between the pixel intensity values of the blue channel and the average of the red and green channels. We compute the stain intensity for a given pixel as shown below.

$$s_{ij} = b_{ij} - \frac{r_{ij} + g_{ij}}{2} \quad (3)$$

As shown in Figure 4, this approach yields satisfactory results except in areas of very high staining, at which point the pixel values become very dark gray with no discriminating values between the blue channel and red and green channel average. Therefore, we set a minimum value for the blue channel and treat values in the red and green channels beneath this threshold as evidence of staining and use the difference between the value in the average of the red and green channels and the minimum as the value for the stain intensity in these areas.

$$s_{ij} = \max(b_{\min}, b_{ij}) - \frac{r_{ij} + g_{ij}}{2} \quad (4)$$

Ideally, we would be able to use the measured stain intensity as a direct measure of the abundance of a particular transcript at a given location. Unfortunately, this is not the case as there is substantial variability for the effectiveness of individual probes. While the current method does not yield directly comparable measures of abundance that are valid for comparisons between different genes, the measures of abundance for a single gene, over the spatial extent of the disc should be more directly comparable. In order to maximize the measured differences in expression of a single gene across the entire disc, one can normalize the recorded values, such that the highest measured stain value is assigned to a predetermined value, and other values scaled, linearly, such that the highest measured value is assigned to the highest value on the new scale, and the zero values on the original scale are assigned to zero values on the new scale. By performing this kind of normalization, we expand the dynamic range of the measured pixel values for given genes.

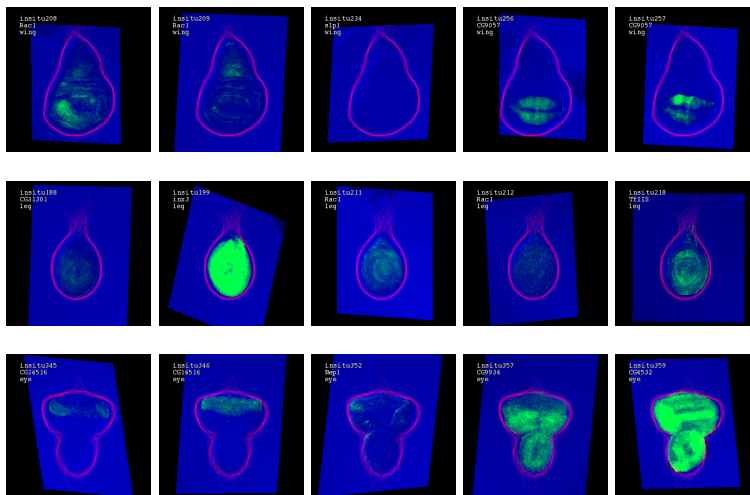


Fig. 5: Images of *Drosophila melanogaster* wing (first row), leg (second row), and eye/antenna (third row) imaginal discs, automatically segmented and aligned to the model. The original grayscale image is shown in the blue channel, the outline of the disc model is shown in red and the scored stain is shown in green.

In Figure 5, we show images of imaginal discs that have been automatically aligned, extracted from background. We have aligned over 800 images from over 130 different genes across the four main imaginal disc types.

4 Gene Expression Maps and Map Clustering

We would like to construct maps of gene expression that incorporate multiple samples, when possible. Two simple approaches are to use the mean and the median of the available maps. For each gene and disc-type pair, if there is only one image, we treat this image as the map. If there is more than one image, we construct maps where the value of each pixel is the median of the values from each aligned, stain-scored image at the given pixel, which are shown in Figure 6. In addition to the median, we construct maps of the pixelwise standard deviation of the stain scored image, yielding a measure of the variability at for a given pixel across the samples (data not shown).

4.1 Reverse Lookup to Find Similar Expression Patterns

Having built a data set of images and maps, one can interrogate new images to automatically extract a disc of a given type, score the image for stain and compare the stain pattern to either the individual stain patterns in the database or, perhaps more importantly, the median maps for each gene, thereby allowing one to discover to which genes the new pattern is most similar. In Figure 7 we




Fig. 6: Gene expression maps of imaginal discs in *Drosophila melanogaster*. The maps are made by taking the median expression value at each pixel from the stack of images for each gene. (a) Maps of MESR3, dpp, drl, CG4914, and CG9057 in the wing. (b) Maps of CG9747, nub, dpp, fd96Cb, and drm in the leg. (c) Maps of Rapgap1, EG:EG0007.7, dpp, btd, and Traf1 in the eye/antenna disc.

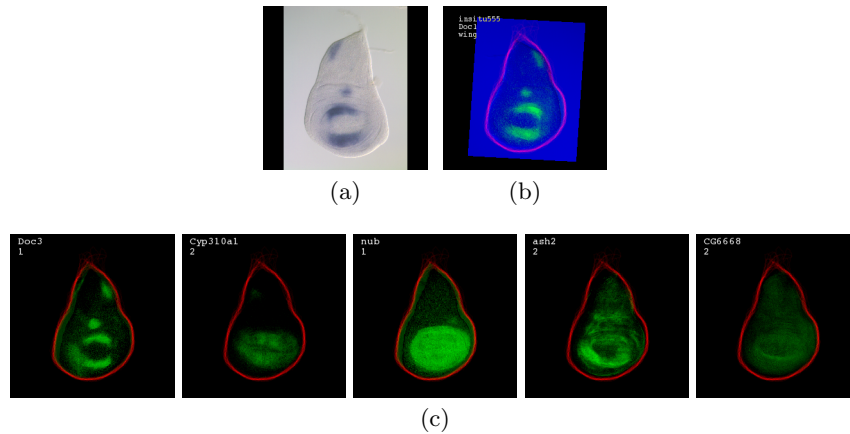


Fig. 7: Reverse lookup procedure to search the database of average gene expression patterns most similar to the Doc1 gene. (a) Image of a wing imaginal disc stained for Doc1 gene expression. (b) Automatically aligned and extracted stain pattern for Doc1. (c) Average gene expression maps of Doc3, Doc2, nub, Cyp310a1 and Pepck.




Fig. 8: Comparative gene expression maps. The expression patterns of multiple genes in Wing imaginal discs are shown overlaid on a common reference map. (a) MESR3 (red), CG9057 (green) and Sp558 (blue). (b) BG:DS00180-3 (red), drm (green) and Timp (blue).

see a new image which is then automatically segmented, extracted and scored for stain. The resulting stain pattern is then searched against the median maps using a simple L_2 distance metric and the 5 resulting maps most similar to this pattern are shown.

4.2 Comparative Maps of Multiple Genes

By registering the images to a common global template for each shape, we are able to make comparisons between expression patterns determined in different experiments and we are able to create overlay maps that contain a representation of the spatial patterns of more than one gene. While one can use biological techniques to image multiple genes simultaneously [21, 22, 23], one must know which genes to look for. Using our methods, we can computationally construct virtual comparative maps of arbitrary sets of genes. Example overlay maps are shown in Figure 8.

In addition to the reverse lookup and image overlay capabilities, we can look for common features between different genes by clustering them together using any number of standard clustering techniques. We have chosen to use simple clustering techniques to perform two classes of clustering, a clustering of the genes and a clustering of vectors of gene expression for all of the genes that can be computed for each pixel.

4.3 Gene Clustering

The k-means clustering algorithm clusters n points in a d -dimensional space around k cluster centers [24]. We denote the data points x_n and establish a variable z_n which contains the value of the current cluster in which point x_n resides. The assignment to clusters stored in z is computed by the following equation:

$$z_n = \underset{j}{\operatorname{argmin}} \|x_n - \mu_j\| \quad (5)$$

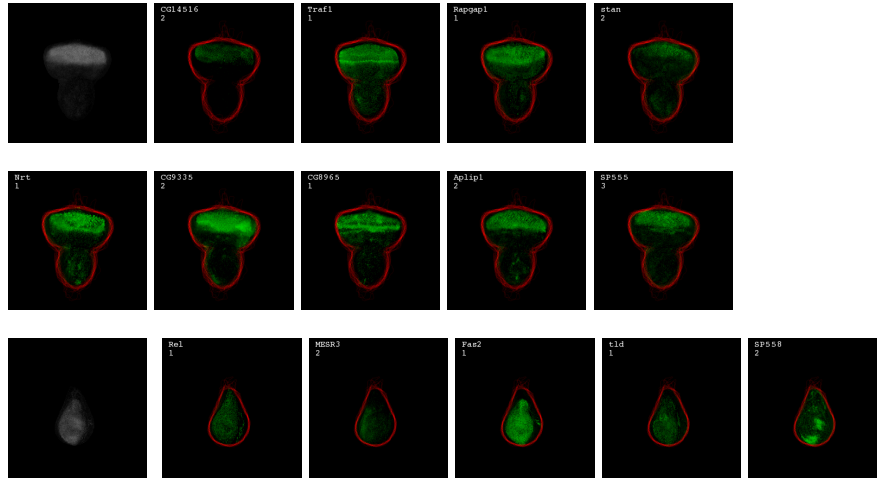


Fig. 9: Representative clusters from gene clustering. The first column contains a cluster center from an eye/antenna disc cluster (top row) and a haltere disc cluster (bottom row). Other columns contain cluster members for the given cluster centers: CG14516, Traf1, Rapgap1, stan, Nrt, CG9335, CG8965, Aplip1, and SP555 for the eye/antenna cluster; Rel, MESR3, Fas2, tld, and SP558 for the haltere cluster.




Fig. 10: Wing, leg, haltere and eye/antenna pixel clusters with 32 clusters (color-coded) in each image. This figure shows which pixels in the imaginal disc are similar to each other based on the gene expression profiles across all the genes measured in that particular disc.

where $\|\cdot\|$ is an appropriate distance metric. The new cluster means μ are computed by:

$$\mu_i = \frac{\sum_n \delta(z_n, n) x_n}{\sum_n \delta(z_n, n)} \quad (6)$$

where δ is the Kronecker delta function. Thus, for a given a set of gene maps, x , with x_n indicating the n th gene map, and μ_j representing the j th cluster, we run the k-means clustering as in Equations 5 and 6. In our experiments, we found that using normalized cross correlation (NCC) produced better results than using L_2 distance metric. Some representative clusters are shown in Figure 9.

4.4 Pixel Clustering

Given a set of gene maps, we want to cluster the pixels based on a vector of gene expression across all of the genes at each pixel. Given the set of all pixels P and $p \in P$, we want to cluster the x^p vectors. We refer to the cluster centers as π_j . To cluster the pixels, instead of the genes, we use the k-means clustering equations as follows:

$$z_n = \operatorname{argmin}_{j \in \{1, \dots, P\}} \|x^p - \pi_j\| \quad (7)$$

where $\|\cdot\|$ is an appropriate distance metric for comparing the vectors, in our experiments we used the L_2 distance. The new cluster means π are computed by:

$$\pi_i = \frac{\sum_p \delta(z_p, p) x^p}{\sum_p \delta(z_p, p)} \quad (8)$$

The results of pixel clustering of wing, leg, haltere and eye/antenna imaginal discs with 32 clusters are shown in Figure 10.

5 Discussion

In this paper, we have presented a method of generating a large number of spatial patterns of gene expression in *Drosophila melanogaster* imaginal discs, and for using shapes learned from the data, rather than using a single exemplar, as a global model of the shape of interest, to which a set of patterns can be aligned. We have presented methods for automatically detecting an imaginal disc in an image and for aligning this image to the learned shape model, for automatically scoring these images for stain, and for registering these stain patterns to the global model. Our representation of stain patterns as a quantitative measure of gene expression, aligned to a global model, enables us to efficiently cluster both the patterns of the genes themselves, and the regions of the tissues, as represented by the pixels in the global model. Finally, we have developed a reverse-lookup procedure, that enables us to take a new image, stained for a gene of interest, and to search our database of patterns to find genes with similar spatial patterns of gene expression.

Using our methods, we have determined the patterns of over 130 genes in some or all of the four largest and most well-characterized imaginal disc types, the wing, leg, haltere and eye/antenna discs. Our parallel automated alignment and feature extraction method works adequately over 85% of the time, but may fail when presented with images containing multiple imaginal discs or substantial amounts of biological clutter.

Our method of pixel clustering represents a novel way of viewing spatial gene expression data. By clustering the pixels, we are able to identify regions of the discs that are, on the basis of their gene expression profiles of on the order of 100 genes, more similar to the other pixels in that region than to the other regions. The pixel clusters identified are reminiscent of the zones of development identified in the classical imaginal disc fate maps in the literature. We look

forward to exploring the relationships between these clusters and genetically-identified regions of the discs for which the eventual fates of the cells have been determined.

Future work may include extending the procedure to be able to identify multiple discs in a single image, to be more robust to biological noise such as the presence of trachea or other larval material in the image, and to use hierarchical clustering methods for the gene and pixel clustering.

Acknowledgments

We thank Amy Beaton, Ben Berman and Michael Eisen for helpful discussions and Pavel Tomancak for development of the BDGP in situ image tracking database. We thank David Forsyth and Erik Learned-Miller for helpful discussions regarding Congealing.

References

- [1] M. Adams, *et al.*, *Science* **287**, 2185 (2000).
- [2] S. Misra, *et al.*, *Genome Biol* **3**, 1 (2002).
- [3] M. N. Arbeitman, *et al.*, *Science* **297**, 2270 (2002).
- [4] P. Tomancak, *et al.*, *Genome Biology* **3**, 1 (2002).
- [5] R. Lipshutz, S. Fodor, T. Gingeras, D. Lockhart, *Nat Genet* **21**, 20 (1999).
- [6] M. B. Eisen, P. O. Brown, *Methods Enzymol* **303**, 179 (1999).
- [7] G. Kamberova, S. Shah, *DNA Array Image Analysis – Nuts and Bolts* (DNA Press LLC, 2002).
- [8] A. N. Jain, *et al.*, *Genome Res* **12**, 325 (2002).
- [9] K. P. White, S. A. Rifkin, P. Hurban, D. S. Hogness, *Science* **286**, 2179 (1999).
- [10] A. Klebes, B. Biehs, F. Cifuentes, T. Kornberg, *Genome Biol* **3**, RESEARCH0038 (2002).
- [11] M. J. Butler, *et al.*, *Development* **130**, 659 (2003).
- [12] O. Eugene, A. Yund, J. W. Fristrom, *Tissue Culture Association Manual* **5**, 1055 (1979).
- [13] E. Miller, N. Matsakis, P. Viola, *IEEE Conference on Computer Vision and Pattern Recognition* (2000), pp. 464–471.
- [14] E. Miller, Learning from One Example in Machine Vision by Sharing Probability Densities, Ph.D. thesis, Massachusetts Institute of Technology (2002).
- [15] P. Ahammad, C. L. Harmon, A. Hammonds, S. S. Sastry, G. M. Rubin, *CVPR (2)* (IEEE Computer Society, Washington, DC, USA, 2005), pp. 755–760.
- [16] H. G. Barrow, J. M. Tenenbaum, R. C. Boles, H. C. Wolf, *IJCAI* (1977), pp. 659–663.
- [17] G. Borgefors, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**, 849 (1988). Published.
- [18] R. C. Gonzalez, R. E. Woods, *Digital Image Processing* (Prentice Hall, 2002), second edn.
- [19] P. Soille, *Morphological Image Analysis* (Springer-Verlag, Berlin, Heidelberg, New York, 1999).
- [20] J. Lewis, Fast normalized cross-correlation (1995).

- [21] J. W. O'Neill, E. Bier, *Biotechniques* **17**, 870, 874 (1994).
- [22] G. Hauptmann, *Methods* **23**, 359 (2001).
- [23] D. Kosman, *et al.*, *Science* **305**, 846 (2004).
- [24] R. O. Duda, P. E. Hart, D. G. Stork, *Pattern Classification* (John Wiley and Sons, 2001), second edn.

A Appendix

A.1 Foreground Feature Extraction

Algorithm A.1: EXTRACT-FOREGROUND-FEATURES(*img*)

```

laplacian ← LAPLACIAN(img)
absLaplacian ← ABS(laplacian)
dilated ← DILATE(absLaplacian, SquareStructuringElement)
eroded ← ERODE(dilated, SquareStructuringElement)
foregroundFeatures ← DISTANCE-TRANSFORM(eroded)
return (foregroundFeatures)

```

A.2 Aligning Target Image to Model

Algorithm A.2: ALIGN-TARGET-TO-MODEL(*target*, *model*)

```

feat ← EXTRACTFOREGROUNDFEATURES(target)
for each x ∈ initialConfigurationList
  for iteration ← 0 to maxIterations
    origScore ← NCC(feat, model)
    for x ∈ x
      x- ← x − stepSize
      stepDownImage ← AFFINETRANSFORMIMAGE(feat, x-, x)
      stepDownScore ← NCC(stepDownImage, model)
      if stepDownScore > origScore
        then x ← x-
      x+ ← x + stepSize
      stepUpImage ← AFFINETRANSFORMIMAGE(feat, x+, x)
      stepUpScore ← NCC(stepUpImage, model)
      if stepUpScore > MAX(origScore, stepDownScore)
        then x ← x+
    x ← UPDATE(x, x)
  alignedTarget ← AFFINETRANSFORMIMAGE(feat, x)
return (alignedTarget)

```