

The Space-Time Map Applied to Drosophila Embryogenesis¹

Prem Janardhan, Martial Hebert
{Computer Science Dept., Robotics Institute}
Carnegie Mellon University
Pittsburgh, PA, USA
{prem@cs,hebert@ri}.cmu.edu

Katsushi Ikeuchi
Institute of Industrial Science
University of Tokyo
Tokyo, Japan
ki@cvt.iis.u-tokyo.ac.jp

Abstract

Many physical phenomena have complex structure in both space and time. To systematically understand these phenomena from images we need representations that unify the treatment of space and time. We create such a representation, the space-time map, for characterizing contour evolutions. Many types of information are computed and stored as facets of the map, registered to a space-time manifold generated by the evolution. We demonstrate our representation on the example of Drosophila embryogenesis in optical section. Changes in embryo shape are reflected as changes in the dye distribution along the deforming vitelline membrane contour. We extract a series of contours and create a two-dimensional space-time map. We track intensity on this map to obtain a velocity field. We extract space-time ridges and significant motions on this map, and use them along with prior knowledge to recognize the significant features and events of embryogenesis.

1. Introduction

In order to systematically analyze space-time phenomena such as human gesture and biological development, we need to build representations that unify or integrate the treatment of the space and time axes. Such representations will permit us to consider the entire phenomenon as a single space-time entity, thus making subsequent analysis simple and uniform. This paper describes the creation of such a space-time representation for contour-based changes, and its application to the analysis of Drosophila embryogenesis.

Unlike types of motion more typically studied in com-

puter vision, biological deformations are generally non-rigid, making them complex and difficult to analyze. The embryogenesis of Drosophila (i.e., the process by which the fertilized fruit fly egg develops into a larva) offers an interesting example of nonrigid biological deformation. The changes are drastic, yet stereotypical (different embryos all deform in the same characteristic way). This research aims to automatically understand these complex changes.

We build a system for automatically recognizing significant features and events of Drosophila embryogenesis. Such recognition is valuable because it can form the basis for quantitative embryogenetic measurements, provide a reference frame for other embryogenetic phenomena, and provide a means to distinguish normal and mutant embryos. The changes we plan to use have, in the past, been used for staging of embryogenesis [2]. We choose an imaging modality that makes the embryo surface clearly visible, namely, fluorescence imaging and intervittelline injection of vgal dye. Vgal dye is injected into the intervittelline space (space between the vitelline membrane that encases the embryo, and the embryo surface) where it becomes fluorescent. Thus, dye-containing gaps between the embryo surface and the vitelline membrane are highlighted. A sequence of slice images is obtained by optical section microscopy (defined in section 2). The vitelline membrane cross-sectional contour is extracted and tracked from the optical sections using a snake [5], and the image region contiguous to the contour processed, to extract an intensity profile representing the varying dye presence along the embryo boundary. We combine contours from successive images into a two-dimensional space-time map of the embryogenesis.

1. This research was partially supported by the National Science Foundation's High Performance Computing grant titled "Automated Interactive Microscope" (#BIR-9217091) and the National Science Foundation's Science and Technology Center grant titled "Center for Light Microscopy Imaging and Biotechnology" (#MCB-8920118). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of anyone else or the U.S. government

We use correlation to compute a velocity field on the map that describes the local motion of intensity on the map. A coarse-to-fine strategy, and consensus between neighboring points on the map, are employed to reduce the effects of noise. The output is a pair of dense velocity fields, in the forward and backward directions.

We detect rapid motions and intensity ridges on the space-time map, and recognize some that are biologically significant. The ridges are tracks of embryo surface indentations (many embryo surface features have been identified and named by biologists in the past). Recognition is performed using a combination of spatiotemporal context, provided by previous recognitions, and measured local properties. Recognition is thus hierarchical; being first effected on those ridges that are easily identifiable without the use of context, and later on those that require the use of the context provided by previous recognitions.

2. Biology

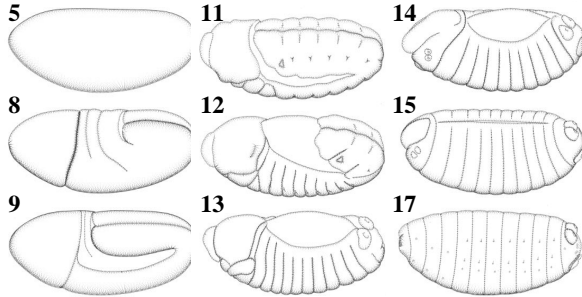


Figure 1: Sketches of the *Drosophila* embryo at numbered stages of embryogenesis, showing evolution of surface features. The staging method used is that of Campos-Ortega and Hartenstein (1985) [2].

Today, developmental biologists routinely collect movies of the embryogenesis of *Drosophila Melanogaster* (the fruit fly). Since living tissues are both transparent and able to tolerate the transmission of light, optical section microscopy is commonly used to study embryogenesis in cross-section. An optical section microscope is a microscope with a narrow depth of field, with which the desired optical section, or plane of the specimen, is imaged by adjusting the plane of focus. The embryo is stuck to a glass slide so that it cannot move, though it is free to deform. Fluorescent dye is commonly applied to highlight structures of interest prior to study. Subsequent to injection into an organism, a dye becomes active (capable of fluorescence) by reaction with natural biochemicals. When pumped with light, active dye fluoresces and illuminates internal structures. In particular, the injection of vgal dye into the intervittelline space permits the observation of morphological changes in the embryo surface in the course of the development of the embryo. During this development, furrows, invaginations

and ridges are observed to form, change in shape, and die. Many of these surface features have been named by biologists ([2][4], Figure 2).

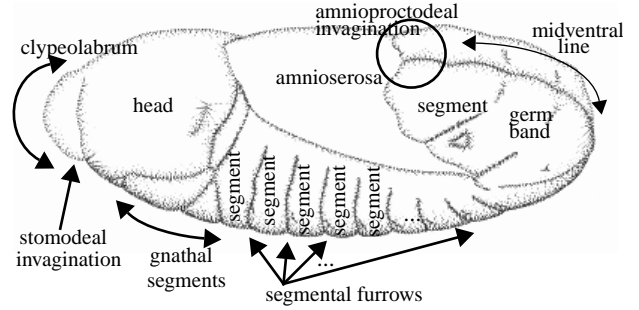


Figure 2: Stage 12 *Drosophila* embryo with labelled parts

For our experiments, we use 20X microscope movies taken at a non-superficial sagittal optical section. A sagittal section is a section parallel to the plane of bilateral symmetry of the organism. A non-superficial section is a relatively deep section; one in which embryo tissue predominates, and dye is present only in small pockets; the presence of dye is then more informative of the shape of the embryo. In these images, the vitelline membrane appears as an oval slightly deforming contour, with the embryo cross-sectional outline inside the membrane being also roughly oval, but possessing numerous small indentations that change drastically with time. Wherever the embryo surface (and hence, cross-sectional outline) is indented, bright dye collects in the resulting space between the embryo and the vitelline membrane, and is observed. The intervittelline space appears as a sort of bright belt of varying (or vanishing) brightness and thickness just inside the vitelline membrane.

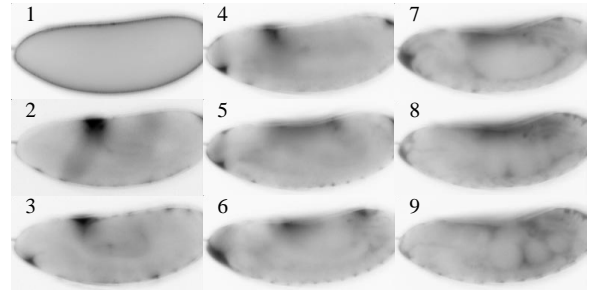


Figure 3: Sequence of frames from an intervittelline injection vgal movie. The frames are equally spaced in time and span about 12 hours in real time

3. Contour Extraction

Many important changes in the shape of the embryo shape are reflected as changes in the shape of its cross-sectional contour. Due to the small and cusped nature of the indentations present in this contour, it is difficult to stably extract its exact shape using a snake. We therefore choose

to extract the vitelline membrane cross-sectional contour instead, since this has a relatively smooth shape. We then make primary use of the image brightness in the vicinity of this contour, rather than the contour shape, for our analysis. Image brightness is caused by the presence of dye. Bright dye fills the intervittelline space more or less uniformly. Since the vitelline membrane only flexes slightly, the presence of dye in the neighborhood of the vitelline membrane is decided by embryo surface shape. Thus, our analysis effectively uses surface shape.

Since the vitelline membrane only flexes slightly, and the embryo is stuck fast to a glass slide, tracking its cross-sectional contour is easily accomplished automatically by a standard image snake [5], operating upon the gradient magnitude image (the gradients are calculated using smoothed derivative of Gaussian operators). Initialization of the snake is performed by hand on the first frame.

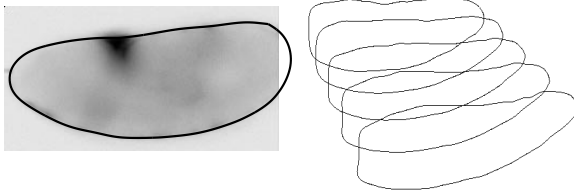


Figure 4: Vitelline membrane contour extracted from a vgl image, and a series of these contours.

4. Intensity Profile

In a given optical section, the intervittelline space is manifest as a belt whose width varies due to the presence of embryo surface features. As the embryo surface evolves, the width, and consequently, the (total) brightness, of this belt changes. Changes in embryo surface features can be characterized using the corresponding changes in the image intensity within the belt. We estimate this intensity using a belt detector and store it as an intensity profile along the vitelline membrane contour. In this intensity profile, indentations correspond to intensity peaks. Due to noise, out-of-focus effects and the varying width of the intervittelline belt, we need to focus on stable computation of the intensity of the intervittelline belt. Our method processes all the intensity within a relatively wide region surrounding the contour (chosen wide enough to encompass the entire intervittelline belt) to produce its estimate of the contrast of the belt with respect to its surroundings.

To deal with the varying width of the intervittelline belt, we use belt detectors with different widths and combine the results. We also vary the position of the belt center along the contour normal by a few pixels in either direction (for convenience, we refer to this variable, the position of the belt center, as “phase”) relative to the position of the con-

tour. We process the samples as follows: if the sampled phases are P_1, P_2, \dots, P_m and the sampled widths are W_1, W_2, \dots, W_m , and the belt detector output (at a given contour point) for a phase P_i and width W_j is b_{ij} , we first calculate $c = \min_{i,j}(b_{ij})$, and subtract it off from each b_{ij} . Then, we calculate the composite response B as the *RMS* (root mean square) of the resulting values; i.e., $B = \text{RMS}_{ij}(b_{ij} - c)$. The use of the *RMS* provides an estimate that is stable with respect to the width and phase sampling employed, since it makes use of many responses. It also has the desirable property of weighting large responses more than small ones (when responses are larger than unity, as in our case). Subtracting the minimum value gives negative responses low weight. Since these do not indicate dye presence, they are not significant. See Figures 11,12 for detection output.

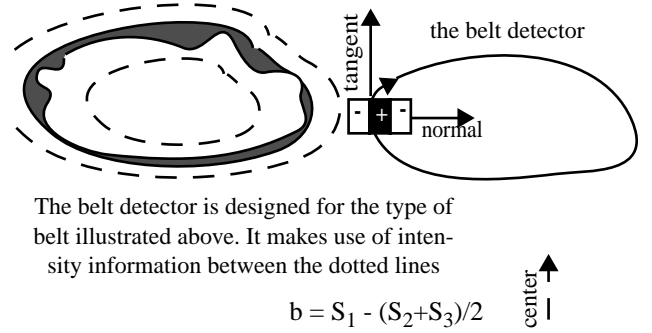


Figure 5: This belt detector calculates, for each contour point, a measure of the local contrast along the contour normal. S_1 , S_2 and S_3 are the intensity sums within the areas marked. b is the detector output for a particular width and phase (see text). The belt center is slightly inside the contour so as to be more likely to be centered over the intervittelline belt.

Normalization: Four different processes are known to make the contour brightness distribution non-stationary. First, due to reaction with an enzyme in the intervittelline space, the dye brightness gradually increases over time. Second, absorption of dye by embryo cells causing a reduction in observed brightness. Third, “neutral density” filters are automatically introduced into the light path as necessary, to keep image intensity within the dynamic range of the CCD. Fourth, due to intermittent malfunctioning of the camera shutter, occasional contours exhibit very low brightness. Contours falling under the fourth case are identified by their characteristically low brightness and removed. If we examine median contour brightness (Figure 6), we observe a gradual drift up or down due to the first two effects, with sharp jumps wherever the neutral density filter was used.

It is desirable to make the contour brightness distribu-

tion stationary, so that we can apply brightness thresholds in a consistent manner over time. In order to achieve this, we divide the brightness of every pixel by the median brightness of its contour, thus flattening the plot of Figure 6. The median is experimentally observed to produce satisfactory results whereas normalization by using, e.g., the max brightness of the contour does not. For the result, see Figure 10.

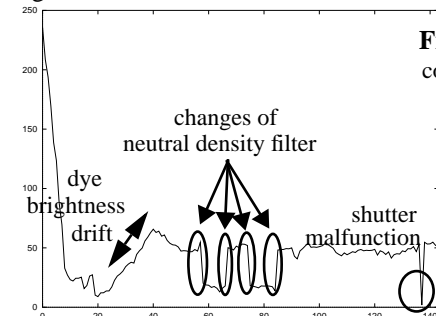


Figure 6: Median contour brightness vs time

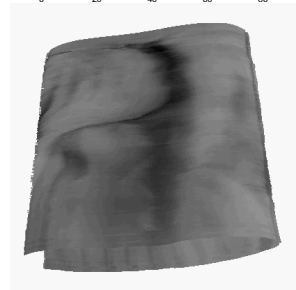


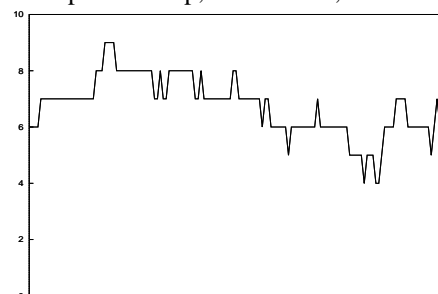
Figure 7: A view of the space-time manifold that represents the embryogenesis, with the extracted intensity profile mounted onto it.

5. Space-Time Map

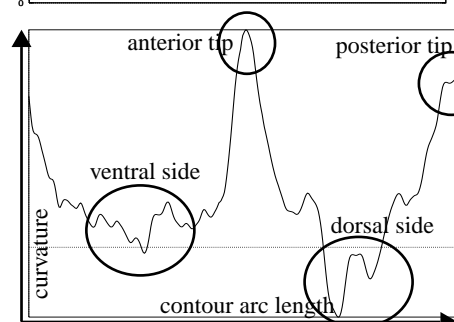
Space-Time Manifold: We now have a series of contours, with attached intensity profiles. To extract the spatiotemporal features of the contour evolution, we need to create a spatiotemporal representation (creating such a representation also permits simpler conceptualization); this entails ensuring that the contours have uniform sampling properties, and establishing correspondence between successive contours. We reparametrize the contours by arc length to uniformize their sampling. Each contour contains 400 elements in our experiments. We perform contour alignment to identify a reference (alignment) point per contour. The intensity profiles are observed to move relative to the alignment point; later, we describe tracking of the intensity. These contours constitute a sampling of a spatiotemporal manifold that represents the embryogenetic contour evolution (Figure 7). From the manifold, we can extract spatiotemporal features and motions of the evolution. All subsequent processing, in concept, occurs on this manifold. The manifold is a compact representation of a complex evolution occurring in $R^3 \times T$. For ease of viewing, all results are displayed on a flattened SXT version of this manifold that we call the space-time map (since it is used as a map of the embryogenesis).

Alignment: Contour shape changes with time, but remains roughly oval. A coarse alignment is performed based on a calculation of the principal axis of the oval. The point at which this axis intersects the contour is used as the coarse alignment point. A more precise local alignment based on contour curvature (measured as vertex angle) is performed (we minimize the difference between successive contour curvature distributions).

Curvature Features: The oval contour shape exhibits certain invariant curvature features: an anterior tip and a posterior tip that are curvature peaks (the anterior tip is more convex), and a dorsal and ventral side that are curvature troughs (the ventral side is more convex). These features are large and relatively insensitive to noise. Reversals of contour element order are performed as necessary to ensure the following order of presentation along the contour: posterior tip, ventral side, anterior tip, dorsal side.



The difference between the coarse and fine alignment (in contour elements) plotted against time.



The anterior and posterior tip and dorsal and ventral side of the embryo are identified using curvature.

Figure 8: Alignment and calculation of curvature features

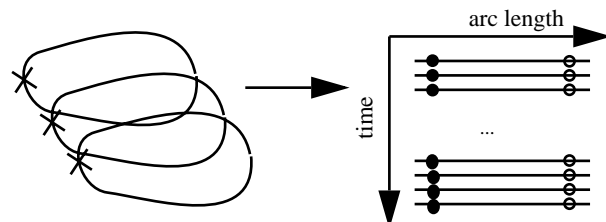


Figure 9: Converting the space-time manifold to a space-time map

Making the Map: Next, we create the space-time map by “cutting” each contour at its alignment point and making it a row of the map. The map is an image-like object in which each row represents data collected from a single

time point. Onto this map, we compute and register many facets, each facet containing a different type of information about the contour evolution.

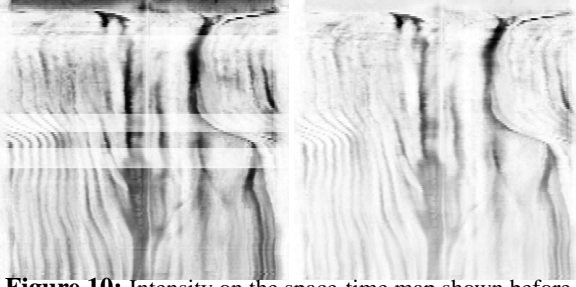


Figure 10: Intensity on the space-time map shown before and after normalization.

6. Velocity Field

Due to sparse temporal sampling, tracking is needed to provide correspondence between successive intensity profiles. Tracking makes it possible to extract spatiotemporal features, and motion information. We choose intensity-based tracking rather than feature-based tracking because we desire dense connectivity information to perform better motion detection. We delay feature detection until after tracking.

Denote the profile at time t by S_t , and the entire map by S . We attempt to match each window w of S_t with a window w' of S_{t+1} using correlation. The output of our tracker is a pair of velocity vector fields V_f and V_b over points of S . At any point (s, t) of S , they provide velocity vectors v_f and v_b respectively such that $(s + v_f, t + 1)$ and $(s + v_b, t - 1)$ are the images of (s, t) under the computed forward and backward motion respectively at (s, t) .

We use the SSD for comparing intensity windows. For each pixel (s, t) in S , we calculate SSD of a window centered on (s, t) with windows centered on pixels in the range $([s-w \dots s+w], t+1)$. An appearance match is signalled by a minimum in the resulting profile. Generally, there are multiple minima in the profile, and sometimes, the lowest one is not the right one, due to large signal deformation or presence of multiple features of similar appearance. Using a larger correlation window reduces both kinds of errors by reducing appearance ambiguity. However, a large comparison window may contain multiple features with conflicting velocities, resulting in an imprecise velocity estimate. Thus, changing the window size trades off accuracy for precision.

So we use a coarse-to-fine strategy: multiple comparison windows, each generating a different SSD profile (we use windows w_3, w_4, w_5, w_6 that generate SSD profiles

P_1, P_2, P_3, P_4 and have spatial widths $2^n - 1$, for $n = 3, 4, 5, 6$ respectively, the temporal width remaining constant at 3). The large windows provide accurate, if imprecise velocity estimates, while the small windows provide precision. We first smooth each SSD profile by averaging it with its spatiotemporal neighbors. We then locate the deepest minimum in P_6 and track this minimum down through P_5, P_4 and P_3 . When tracking a minimum m down from level k to level $k-1$, we first establish a search window for the minimum m to identify the range within which to search at level $k-1$. The left limit of this basin is established heuristically to be whichever of the following two options lies closer to m : 1) the point to the left of m at which the SSD first rises to twice the value at m itself, or: 2) the position of the closest maximum neighboring m on the left side. The right limit is established in a symmetrical manner. Spatiotemporal consensus between neighboring velocity estimates, via a median filter, is finally used to correct some remaining errors. See Figure 13 for example results.

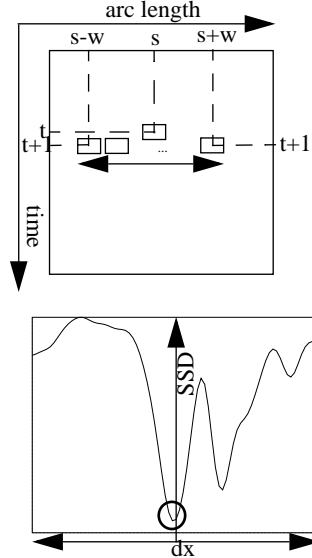


Figure 11: obtaining the velocity field: a window centered on (s, t) is correlated with windows centered on $([s-w \dots s+w], t+1)$

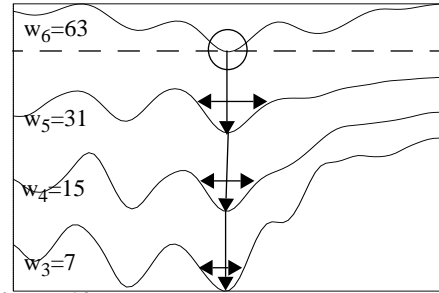


Figure 12: The lowest minimum in the coarsest scale correlation profile is tracked down to the finest scale to obtain the velocity. The tracking search window (shown with a double-headed arrow) is monotonically reduced as we move from coarse to fine

7. Ridge Detection

A *ridge* has been defined as the top of a mountain range or as the boundary between watersheds in some landscape [6]. In our case, the landscape is the intensity profile mounted on the space-time map. Each row of the map contains prominent intensity bumps which correspond to furrows or invaginations on the embryo surface. The tracks of these features over time are observed as ridges. We want to detect these ridges because they correspond to important biological entities. Other researchers [7] have worked on the

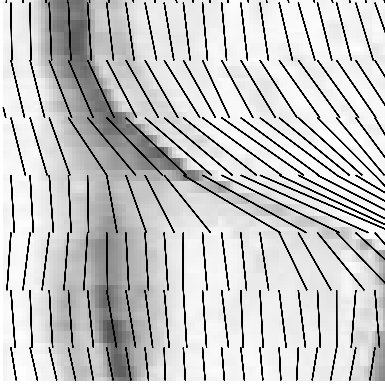


Figure 13: The forward velocity field

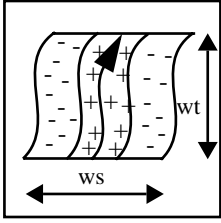


Figure 14: The feature detector window conforms to the local velocity field to obtain a better estimate.

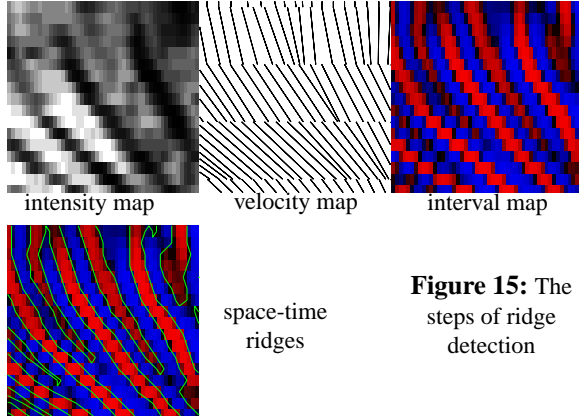


Figure 15: The steps of ridge detection

problem of detecting ridges and crest lines in multidimensional data; our work presently differs, but we will consider incorporating some of these ideas in the future.

Our ridge detector computes local intensity contrast on the space-time map, making use of the velocity field to improve the result. The ridge detector window conforms to the velocity field in the neighborhood of the given pixel,

calculating at each pixel, the difference between the average intensity of the central region, and the average intensity in the two regions that surround it. For a given pixel, the velocity-warped window consists of a spatial neighborhood of the pixel and points generated by the forward and backward motion of that neighborhood over some time interval.

Each row of the ridge detector output is first thresholded to remove low values, and then segmented into positive and negative intervals. The velocity field is then used to track the positive intervals forward and backward. The interval tracks created by this linking operation are the space-time ridges we use. The elements of ridge detection are graphically detailed in Figure 15.

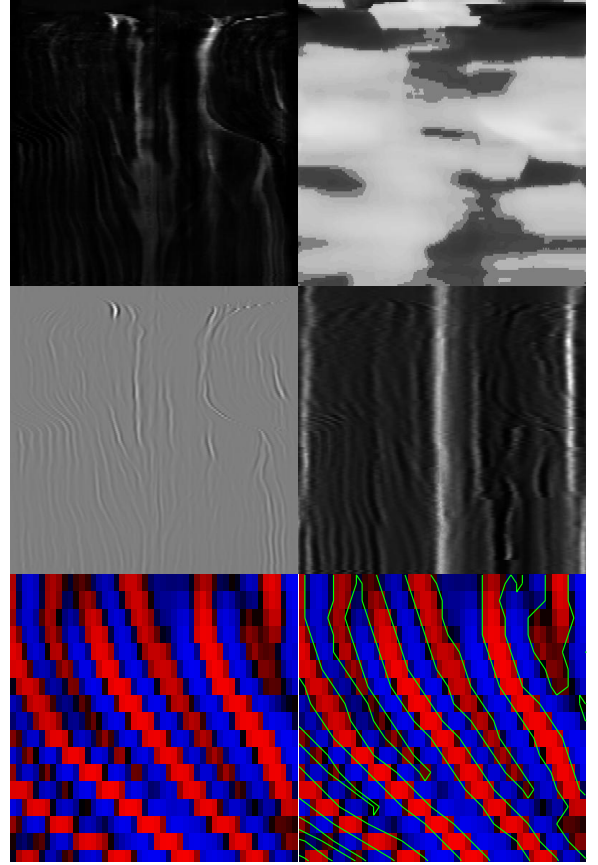


Figure 16: Some different facets of the space-time map. In raster order: intensity, velocity, ridgeness, segmented intervals, segmented ridges (last two shown clipped for clarity)

Final Representation--Multi-Facet Space-Time Map: The space-time map has many different facets (intensity, curvature, 2D spatial position, forward and backward velocity, ridgeness, segmented intervals, segmented ridges) containing different types of information, all registered to the space-time manifold generated by the deforming contour. When performing space-time recognition tasks, we

use information from all of these facets as required.

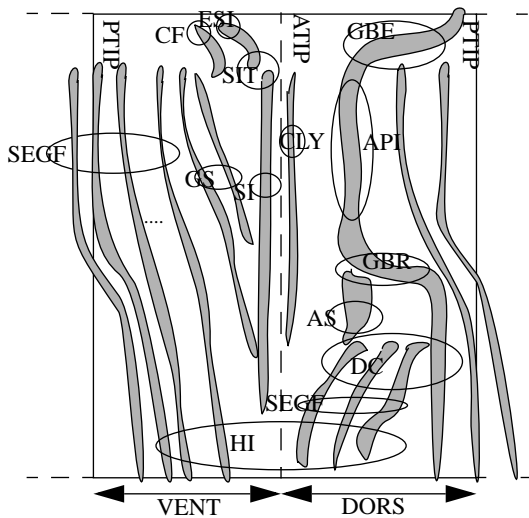


Figure 17: Schematic space-time map of Drosophila Embryogenesis showing the main features and events

API	Amnioproctodeal Invagination	GS	Gnathal Segments
AS	Amnioserosa	HI	Head Involution
CF	Cephalic Furrow	SEGF	Segmental Furrows
CLY	Clypeolabrum	ATIP	Anterior Tip
DC	Dorsal Closure	PTIP	Posterior Tip
ESI	Early SI	DORS	Dorsal Side
GBE	Germ Band Extension	VENT	Ventral Side
GBR	Germ Band Retraction		
SI	Stomodaeal Invagination		
SIT	SI Tightening		

8. The Task

To *detect* an object is to discern its presence from the available evidence. To *recognize* an object is to identify that object from past experience or knowledge [1]. The difference between detection and recognition is that recognition makes use of prior knowledge, while detection proceeds merely on the basis of the available evidence. For example: we recognize, rather than detect, chairs since, in addition to looking at an image, we need to apply substantial prior knowledge of the appearance of chairs. We detect, rather than recognize, edges, since no prior knowledge is needed to identify an edge. Presumably, one first detects low-level entities (e.g., edges, ridges, temporal events), and then uses them along with prior knowledge to recognize higher-level entities (objects, locations, processes). Our task of identifying the important features/events (space-time ridges and motions) of embryogenesis is one of recognition rather than detection, because it involves choosing a specific individual from multiple alternatives generated by

a detection process.

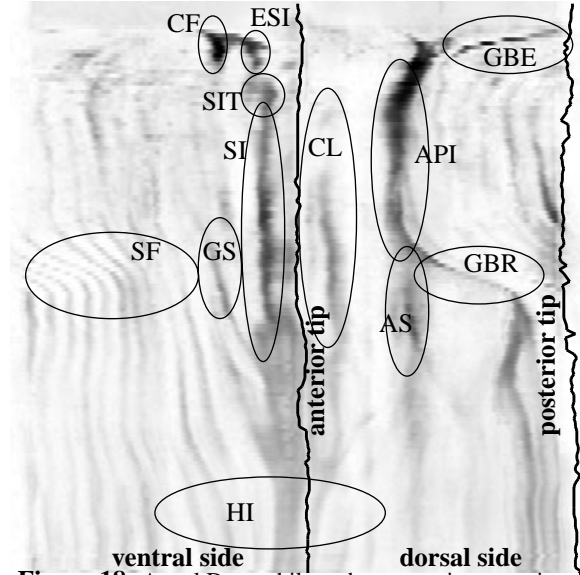


Figure 18: A real Drosophila embryogenesis space-time map showing observable events. See Figure 17 for abbreviations

A complex phenomenon or object may consist of several separately recognizable or detectable parts. One act of recognition or detection may assist in others. For example, human face recognition may be used to guide the recognition of objects known to be likely to co-appear with faces, such as arms. It may be difficult to separately recognize human arms, since they are not very remarkable in appearance, but the recognition of the (more distinctive) face provides information regarding the likely position and size of the arms, which we can exploit. In recognizing a complex phenomenon, the information provided by the recognition each part should help in the recognition of the other parts and the whole. Such a recognition framework will make it simpler to deal well with noise and occlusion.

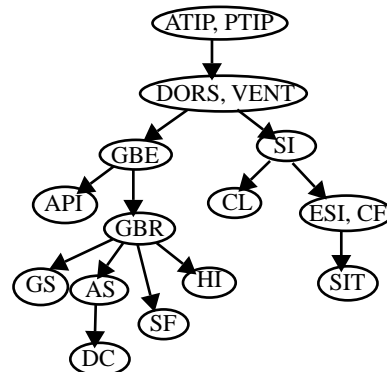


Figure 19: The partial order currently used for recognition of the features and events of embryogenesis. The abbreviations used are those of Figure 17.

Figures 17 and 18 indicate several major features and events of Drosophila embryogenesis that we wish to detect from the space-time map. The identity of these features and

events arise as a combination of their spatiotemporal context (position relative to other features and events) and their intrinsic properties (brightness, velocity, length). Some features, such as the stomodeal invagination, are recognizable with relatively little context (the position of the anterior and posterior tip is sufficient). Some, such as the amnioserosa and gnathal segment, require more context such as the position of germ band retraction. In the absence of context, many features with similar properties will either be impossible to distinguish, or require more complicated comparisons due to greater ambiguity, with larger chance of error. We form a partial ordering of features and events based on spatiotemporal context. If the recognition of A provides context necessary for the recognition of B, then A appears earlier in the ordering than B. Each act of recognition not only locates the position of a feature or event, but also offers constraint about the position of features and events that fall later in the ordering, that facilitates or enables their recognition. At present, this ordering is manually generated, but in future, we plan a systematic analysis that may lead to its automation. Below, we describe how features and events are recognized.

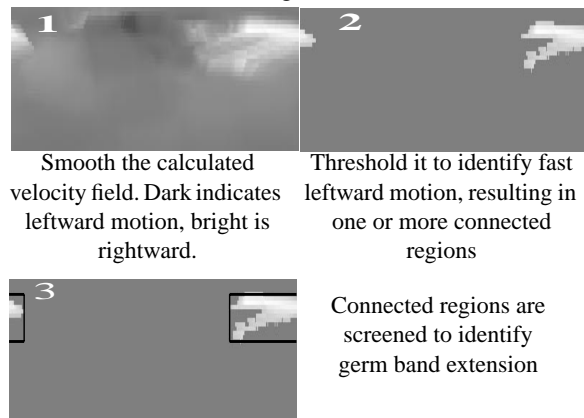


Figure 20: Recognizing germ band extension (GBE)

Important Motions: The germ band is a strip of tissue on the ventral side of the embryo that undergoes dramatic elongation and contraction during the course of embryogenesis, causing the tip of the germ band to extend from its initial position at the posterior tip, towards the anterior tip along the dorsal side of the embryo, and then retract to its original position during the contraction. Germ band extension and retraction are the most important motions of embryogenesis.

Germ Band Extension: On a space-time map of embryogenesis, germ band extension appears as a rapid leftward motion of the germ band on the dorsal side of the embryo. It is recognized as follows: the velocity field is averaged (instead of a rectangular averaging window, a velocity-warped window similar to the ridge detection win-

dow described in Section 7 is used) to remove small velocity variations. We then threshold the smoothed velocity field, compute connected components over surviving regions, and retain components over which the velocity integral (calculated as a sum over the smoothed velocity field) is less than a threshold. If multiple components are found, the one generating the least sum, found on the dorsal side is chosen as a germ band extension representative. Thus, germ band extension is essentially a large region of the dorsal side of the map which has a large leftward velocity. The thresholds used here and elsewhere in this work were determined by measurements on real data. Generally, different thresholds were used for different features/events, but the same threshold was successfully used for a given feature/event over all the data sets used. See Figure 20 for a graphical explanation of germ band extension.

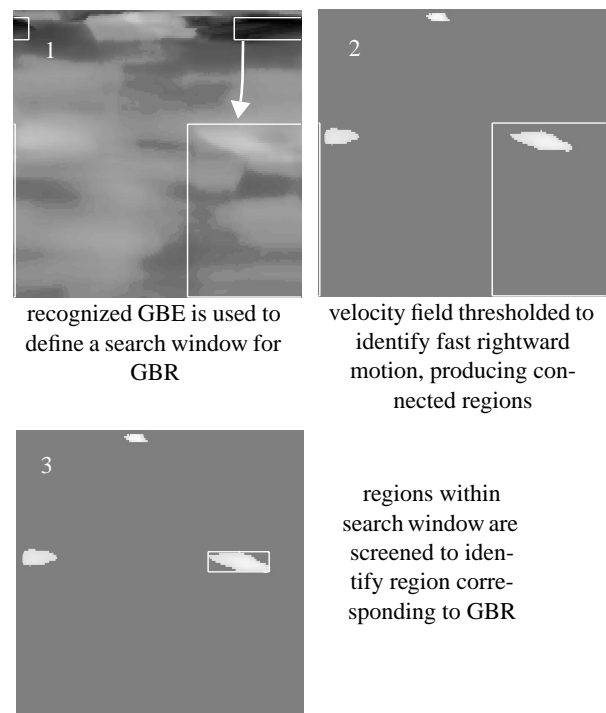


Figure 21: Recognizing germ band retraction (GBR) by making use of detected GBE.

Germ Band Retraction: Germ band retraction is a rapid rightward motion occurring within a search window on the dorsal side at least 50 time units after extension. If extension has already been recognized, we use its position to create a search window for retraction. We then perform in-stream averaging of the velocity field, apply a lower threshold to the result, calculate the connected components and check for components with velocity integral of at least 4000.0 within the search window. See Figure 21 for a graphical explanation of the recognized germ band retrac-

tion.

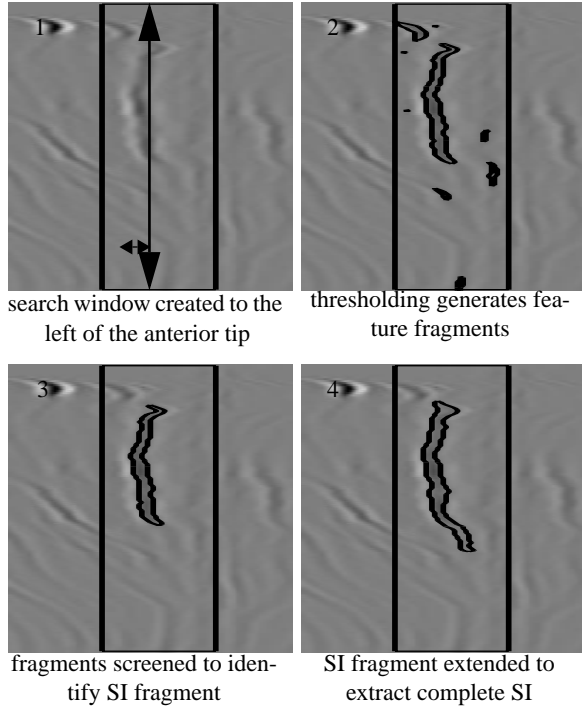


Figure 22: Recognition of the Stomodeal Invagination (SI) based on ridge detection output

Space-Time Ridges: Important space-time ridges of embryogenesis currently recognized include the stomodeal invagination, the clypeolabrum, the cephalic furrow, the early stomodeal invagination, the amnioserosa, and the gnathal segment. Each of these recognitions make use of thresholds $th1$ and $th2$ calculated using the intervals facet as follows: The highest ridgeness value within each positive interval is collected, these values sorted and the 70th percentile is used as $th1$. $th2$ is half of $th1$. The process for recognizing each of these ridges follows the same pattern: A search window is defined based on prior knowledge and prior recognition. The ridge facet of the map is thresholded using $th1$ and intervals recalculated and tracked as described in Section 7. The result is a set of segmented ridge fragments corresponding to bright ridges. We use fragment properties to screen the fragments, identifying the one corresponding to the ridge we are searching for. If multiple fragments pass the screening, a disambiguation condition is applied to recover exactly one fragment. On obtaining the sought fragment, we extend it to extract the entire ridge. The extension process picks the brightest interval in the fragment and tracks it first forward, then backward in time on the intervals facet of the map calculated using the threshold $th2$. Ridge length is calculated relative to the sampling rate, with the sampling rate of a reference data set set to unity. See Figures 22 and 23 for

detection results.

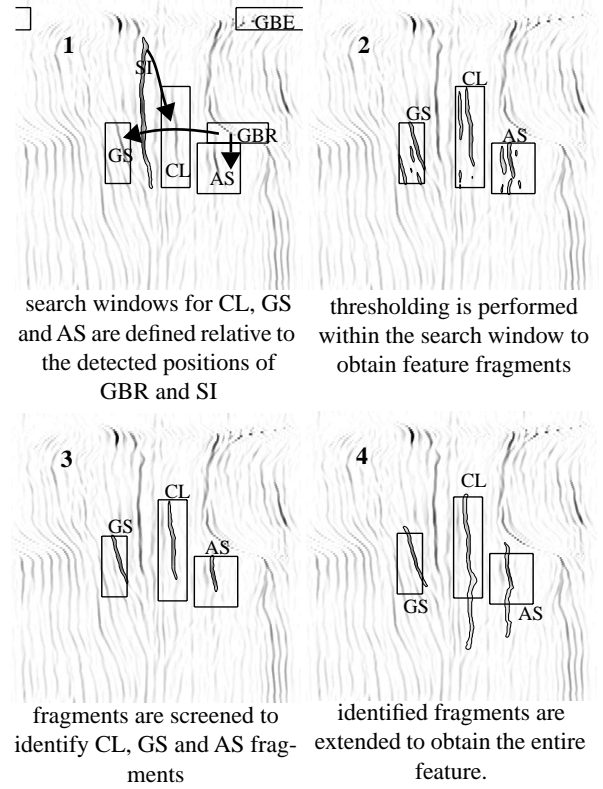


Figure 23: Detection of the gnathal segment (GS) and amnioserosa (AS) use already detected germ band retraction (GBR) to create their search windows. Detection of the clypeolabrum (CL) similarly makes use of the already detected stomodeal invagination (SI).

9. Conclusion

This report has described the creation of a representation, the space-time map, for characterizing contour evolutions, and its application to *Drosophila* embryogenesis. We gather a series of contours representing the deforming shape, resample them by arc length, align them, compute intensity profiles, and combine the contours into a space-time manifold. We calculate and register onto this manifold, additional information about the evolution, such as velocity, curvature, space-time ridges, etc., with each type of information being stored as a separate facet or layer. We cut the manifold along contour alignment points and flatten it into an arclength vs time flat, multi-facet, space-time map. Higher-level tasks such as recognition make use of all the different available facets. By storing the information in a co-registered manner we are able to easily index all the different facets to obtain local information whenever processing any point of the manifold. This is important because many types of information are usually necessary for high-level tasks such as recognition. Using this repre-

sensation, we are able to automatically recognize the significant space-time ridges and motions of *Drosophila* embryogenesis. For a data flow diagram, see Figure 25.

10. The Results

We had 11 data sets available to us. Detection results on only one of these is displayed, for reasons of space, in Figure 24. These data sets were collected by Dr. Jonathan Minden of the biology department at Carnegie Mellon University, and his Ph.D. student, Ruria Namba. The algorithms described in this document were developed in consultation with Jon, Ruria and other Ph.D. students of Jon's lab: Todd Pazdera, Mustafa Unlu and Sidney Cambridge. We are grateful for the ready advice and assistance received from all of the above people during the course of this research. The data sets had a variety of different problems: Some suffered from tilting of the embryo and resulting occlusion of some features. Some had low sampling rate, making tracking more difficult. Some suffered from temporary occlusion caused by bubbles. In all of the observed cases, the processing was successful despite these problems.

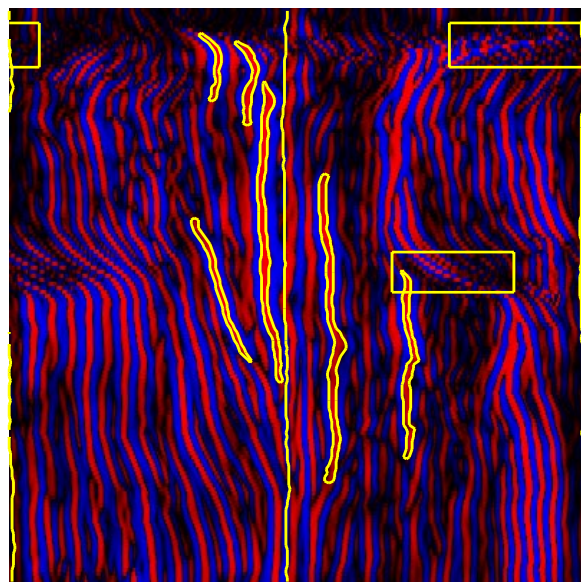


Figure 24: The space-time map of *Drosophila* embryogenesis with detected features and events superimposed.

11. References

- [1] *The American Heritage Dictionary of the English Language*. Boston: Houghton Mifflin, 1992.
- [2] J.A. Campos-Ortega and V. Hartenstein, *The Embryonic Development of Drosophila Melanogaster*. Springer-Verlag, 1995.

[3] T. Darrell and A. Pentland, Space-Time Gestures. *Proceedings, IEEE Conference on Computer Vision and Pattern Recognition, 1993*, pp. 335-340.

[4] V. Hartenstein, *Atlas of Drosophila Development*. Cold Spring Harbor Laboratory Press, 1993.

[5] M. Kass, A. Witkin and D. Terzopoulos, Snakes: Active Contour Models, *International Journal of Computer Vision*, 1(4), 1988, pp. 321-331.

[6] J.J. Koenderink and A.J. Van Doorn, Local Features of Smooth Shapes: Ridges and Courses, *Proceedings, SPIE Conference on Geometric Methods in Computer Vision II*, July 1993, pp. 2-13.

[7] O. Monga, R. Legagne and R. Deriche, Extraction of the zero crossings of the curvature derivatives in volumic 3D medical images: a multi-scale approach. *Proceedings, IEEE Conference on Computer Vision and Pattern Recognition*, June 1994, pp 852-5.

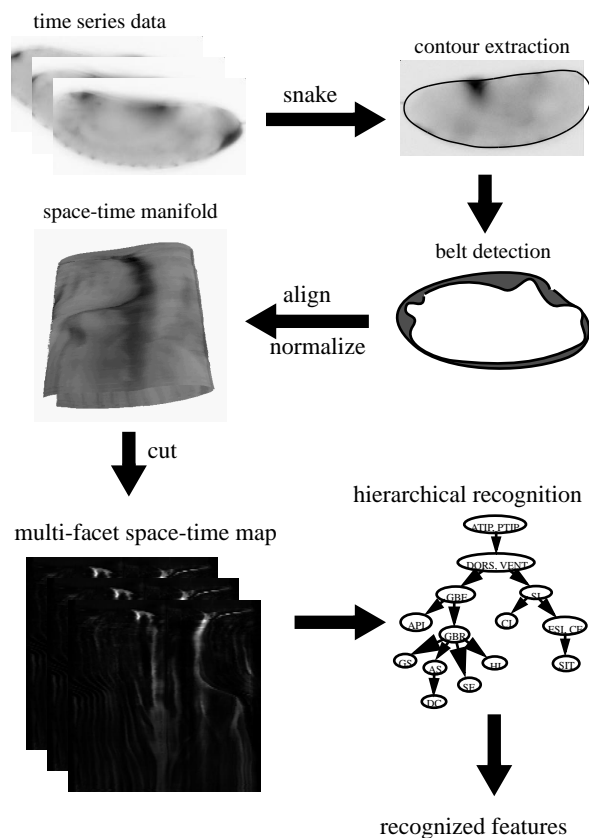


Figure 25: Flow of data

