

Modes and Clustering for Time-Warped Gene Expression Profile Data

Xueli Liu¹ and Hans-Georg Müller^{2, *}

1. Departments of Human Genetics and Biomathematics, UCLA School of Medicine, Los Angeles, CA 90095
2. Department of Statistics, University of California, Davis, CA 95616

*To whom correspondence should be addressed.

ABSTRACT

Motivation: The study of the dynamics of regulatory processes has led to increased interest for the analysis of temporal gene expression level data. To address the dynamics of regulation, expression data are collected repeatedly over time. It is difficult to statistically represent the resulting high-dimensional data. When regulatory processes determine gene expression, time-warping is likely to be present, i.e., the sample of gene expression trajectories reflects variation not only in terms of the expression amplitudes, but also in terms of the temporal structure of gene expression.

Results: A nonparametric time-synchronized iterative mean updating technique is proposed to find an overall representation that corresponds to a mode of a sample of expression profiles, viewed as a random sample in function space. The proposed algorithm explores the application of previous work of Hall and Heckman (2002) to genome-wide expression data and provides an extension that includes random time-warping with the aim to synchronize time scales across genes. The proposed algorithm is universally applicable for the construction of modes for functional data with time-warping. We demonstrate the construction of mode functions for a sample of *Drosophila* gene expression data (Arbeitman *et al.*, 2002). The algorithm can be applied to define clusters among the observed trajectories of gene expression, without any kind of prior non-time-warped clustering, as illustrated in the numerical example.

Contact: mueller@wald.ucdavis.edu

INTRODUCTION

Rapid advances in the microarray technique drive the development of new methods to explore hidden biological processes inherent in the data. Current microarray gene expression experiments increasingly result in functional data, where each response can be viewed as a random trajectory of time-dependent expression levels, based on repeated measurements in time. These trajectories are often referred to as gene expression profiles. It is of interest to capture the variation of expression profiles both in amplitude and time progression across different genes.

Since gene expression is presumably regulated over time, for example in genes that regulate the cell cycle (Spellman *et al.*, 1998; Alter *et al.*, 2000) or other dynamic processes, it is to be expected that time-warping is present in gene expression profiles. Thus, stochastic models that incorporate individually varying time scales to describe the relationship among temporal gene expression data are of interest. Dynamic time warping, also called curve registration in engineering (Sakoe and Chiba, 1978), aims at aligning an individual curve or signal to a given template by warping the time axis of an individual iteratively, until the warped curve maximally coincides with the template in a suitable metric.

The application of dynamic time warping to gene expression level data has been pioneered by Aach and Church (2001). Two implementations of time-warping algorithms were presented, based on simple warping and on interpolating warping, aligning different series resulting from different experiments. While the data of one series can be used as the template for another series, the two series play symmetric roles. A template is often assumed to represent the overall characteristics of the population. In many situations, the target template is unknown and needs to be estimated, for example, through a current sample mean function as in the Procrustes method (Ramsay and Li, 1998). An alternative is a model-based approach to warping that does not rely on a template, such as the method proposed here.

The mode is the most common or frequent value in the sample space. It is a useful nonparametric summary of a data sample. In the univariate case, a mode is the location of the maximum of a density and can be identified easily by finding a relative maximum in a density estimate, for example, in a kernel density estimate (Müller, 1984; Silverman, 1986). In function space, a mode function also would correspond to a maximum of a density, except that it is the rather complex density belonging to a distribution of infinite-dimensional trajectories, which is a complicating factor.

Fukunaga and Hostetler (1975) proposed a simple nonparametric technique, the mean shift procedure, to estimate the gradient of a density function in the space \mathbf{R}^p , $p \geq 1$. The idea was generalized by Cheng (1995) and can be used to locate modes by moving upwards on the gradient. A finite dimensional version was discussed by Choi and Hall (1999), and some asymptotic results on the closely related notion of local moments were derived in Müller and Yan (2001). Comaniciu and Meer (2002) used mean shift for mode detection and clustering in a complex multimodal feature space. Hall and Heckman (2002), extending previous work of Gasser, Hall and Presnell (1998), provided a new angle to the problem of estimating and depicting the structure of a distribution of random functions by introducing the concept of mode functions and density ascent in a functional space.

In this paper, we combine mean updating ideas with time-warping to define mode functions for time-warped gene expression paths, with the aim of obtaining overall representations for samples of microarray gene expression trajectories. As a by-product, we obtain clustering of genes based on their expression profiles which is shown to provide a biologically meaningful classification for life cycle gene expression profiles obtained for *Drosophila* by Arbeitman et al. (2002).

SYSTEM AND METHODS

Concepts and Notation

“Landmark registration” and the “Procrustes algorithm” are two common versions of warping methods (Kneip and Gasser, 1992; Ramsay and Li, 1998). Wang and Gasser (1999) proposed an approach using penalty functions that measure the misalignment among sample curves. A time-synchronizing method was proposed in Liu and Müller (2002). The observed curves are assumed to be elements of a warped time space \mathcal{W} and to be generated by a latent bivariate stochastic process in a time-synchronized space \mathcal{S} , where one component corresponds to a random time-warping function, and the other component corresponds to a random amplitude function. The time transformation is constrained to be monotone increasing. However, monotone functions do not form a vector space, but rather a convex space. Thus it is natural to consider convex functional operations.

When presented with the observed random trajectories $\tilde{Y} \in \mathcal{W}$, such as a sample of gene expression profiles, transforming to a time-synchronized space poses an identifiability problem, as for the same $\tilde{Y} \in \mathcal{W}$, there may exist several latent processes that generate the same observed process (Liu and Müller, 2002). By specifying time-synchronizing maps, this identifiability problem can be circumvented. Specifically, for each observed trajectory $\tilde{Y} \in \mathcal{W}$, we may use a simple area-under-the-curve method to obtain time-synchronizing functions

$$\varphi_{\tilde{Y}}(x) = \int_0^x |\tilde{Y}(s)| ds / \int_0^T |\tilde{Y}(s)| ds, \quad (1)$$

which map each observed trajectory \tilde{Y} to a time synchronizing function $\varphi_{\tilde{Y}} : [0, T] \mapsto [0, 1]$. This leads to the latent bivariate process $(X, Y) \in \mathcal{S}$, via $X(t) = \varphi_{\tilde{Y}}^{-1}(t)$, $Y(t) = \tilde{Y}(\varphi_{\tilde{Y}}^{-1}(t))$, $t \in [0, 1]$.

On the other hand, an observed random trajectory $\{(x, \tilde{Y}(x)), x \in [0, T]\} \in \mathcal{W}$, is generated from a latent bivariate process $\{(X(t), Y(t)), t \in [0, 1]\} \in \mathcal{S}$ through the warping mapping $\psi : \mathcal{S} \mapsto \mathcal{W}$,

$$\psi : \{(X(t), Y(t)), t \in [0, 1]\} \mapsto \{(x, \tilde{Y}(x)), x \in [0, T]\}, \quad (2)$$

defined by $\tilde{Y}(x) = Y(X^{-1}(x))$, where $X^{-1}(\cdot)$ denotes the inverse of $X(\cdot)$.

As discussed above, a convex operation is needed to incorporate the time-synchronizing warping into the mean updating model. Given two observed processes $\tilde{Y}_1, \tilde{Y}_2 \in \mathcal{W}$, and a fixed constant $0 \leq \pi \leq 1$, define a convex sum of \tilde{Y}_1, \tilde{Y}_2 by:

$$\pi \tilde{Y}_1 \oplus (1 - \pi) \tilde{Y}_2 = \psi\{\pi \psi^{-1}(\tilde{Y}_1) + (1 - \pi) \psi^{-1}(\tilde{Y}_2)\}, \quad (3)$$

where ψ is given in (2).

The functional convex sample mean of n observed random processes $\tilde{Y}_1, \dots, \tilde{Y}_n$ is accordingly defined by successively extending these operations to more than two components, leading to

$$\tilde{Y}_{\oplus} = \bigoplus_{j=1}^n \frac{1}{n} \tilde{Y}_j, \quad (4)$$

referred to as the *functional convex average*. The functional convex average is obtained by first time-synchronizing the curves, performing a conventional averaging operation on the time-synchronized versions, and then transforming back to the warped time space. A corresponding distance measure for the distance between two curves is then provided by (incorporating the suggestion of a reviewer, for a discussion of other distances see Liu and Müller, 2002)

$$D(\tilde{Y}_1, \tilde{Y}_2) = \left[\int (Y_1(t) - Y_2(t))^2 dt \right]^{1/2}. \quad (5)$$

Drosophila gene expression profiles

The proposed methodology is applied to a subset of the *Drosophila* life cycle gene expression data of Arbeitman *et al.* (2002). About one-third of the *Drosophila* genes (Adams *et al.* 2000) throughout the life cycle – from fertilization to aging adults – have been biologically characterized, and we use this subsample in order to compare our results with the previous biological characterization. The experiment used cDNA microarrays with the aim to analyze the expression levels of 4028 genes during 66 sequential times throughout lifetime, in order to characterize genes in terms of their differential function during the fly’s life cycle.

The gene expression level recorded at each time point is a relative abundance, by comparing the signal of the experimental sample at each time to the signal of a common reference sample. The data spans embryonic, larval, pupal and adult stages. Among the 66 sample time points, there were 1-hour overlapping periods during the early embryonic phases. We treat the time points as equally spaced which led to well-defined shapes of the gene expression profiles. We only use the first 54 time points since there is large variation at older ages.

For the purpose of illustration, we study two previously characterized clusters of genes: Transient early zygotic genes and co-expressed genes enriched for muscle-specific genes. Excluding genes with incomplete information, one has 19 genes from the first cluster and 13 genes from the second cluster. Figure 1 displays gene expression trajectories for these two clusters of genes: The original expression profiles for all 32 genes are assembled in the upper, and the kernel smoothed versions in the lower panel. The curves in the upper panel are obtained by connecting the measurements with straight lines. In the lower panel, we use bandwidth 1.5 time units and a Gaussian kernel with a local linear kernel smoother.

It can be seen that each gene has its unique profile and progresses at its own rate, resulting in varying peak sizes and peak locations. The complex nature of the variation seen for individual gene expression trajectories suggests both amplitude and time variability, and provides the motivation to combine a time-synchronizing algorithm with the mean updating algorithm to obtain functional modes. These mode functions serve as summaries for gene expression trajectories.

In Figure 2, we display expression profiles for gene CG15634 in the upper and gene CG8606 in the lower panel, along with the trajectories that fall into spherical neighborhoods. These neighborhoods are defined by the distance D (5) in function space, centered around the above profiles and with radii 1.8201 and 1.50, respectively. Note that gene CG15634 has far fewer neighbors than gene CG8606, even though the neighborhood for the latter has a larger radius.

ALGORITHM UNDER TIME-SYNCHRONIZATION

The iterative mean updating algorithm in Hall and Heckman (2002) employs a weighted mean updating algorithm: A sequence of updates is constructed that point towards the direction of steepest ascent in the density surface in function space. This sequence ultimately will converge to the mode function. This functional iterative mean updating algorithm however does not take the time variability into account. To include a time-synchronizing step, we simply replace the ordinary sums that are used in this algorithm by the functional convex sums that are defined in (3), (4).

For observed gene expression profiles $\tilde{Y}_1, \dots, \tilde{Y}_n$, where n is the number of genes, we first time-synchronize these genes by the convex synchronization method described above, replacing the usual L^2 distance with the distance D (5). In the numerical examples, we choose

$$\tilde{W}_h(\tilde{Y}_1, \tilde{Y}_2) = e^{-D^2(\tilde{Y}_1, \tilde{Y}_2)/(2h^2)}, \quad (6)$$

where h determines the size of the local neighborhood. To implement the mean update step that originates at a given profile $\tilde{Y}_0 \in \mathcal{W}$, we first obtain the weights

$$W_{ih}(\tilde{Y}_0, \tilde{Y}_i) = \tilde{W}_h(\tilde{Y}_0, \tilde{Y}_i) / \sum_{j=1}^n \tilde{W}_h(\tilde{Y}_0, \tilde{Y}_j), \quad (7)$$

where $i = 1, \dots, n$. Since all $\tilde{W}_h(\tilde{Y}_0, \tilde{Y}_i)$, $i = 1, \dots, n$ are non-negative, it can be verified that (7) is in $[0, 1]$ and $\sum_{i=1}^n W_{ih}(\tilde{Y}_0, \tilde{Y}_i) = 1$ for a given sample of curve $\tilde{Y}_1, \dots, \tilde{Y}_n \in \mathcal{W}$ and any chosen $\tilde{Y}_0 \in \mathcal{W}$.

Now for any random element $\tilde{Y}_0 \in \mathcal{W}$, denote

$$\tilde{T}_h(\tilde{Y}_0) = \bigoplus_{i=1}^n W_{ih}(\tilde{Y}_0, \tilde{Y}_i) \tilde{Y}_i, \quad (8)$$

and

$$\tilde{\nu}_h(\tilde{Y}_0) = \tilde{T}_h(\tilde{Y}_0) - \tilde{Y}_0. \quad (9)$$

Here $\tilde{T}_h(\tilde{Y}_0)$ is the mean update of \tilde{Y}_0 , and $\tilde{\nu}_h(\tilde{Y}_0)$ is the step from the starting point \tilde{Y}_0 to the mean update $\tilde{T}_h(\tilde{Y}_0)$. As starting points for the mean update iterations, we successively substitute the i -th expression path, \tilde{Y}_i , $i = 1, \dots, n$, for \tilde{Y}_0 at the initial step. At the $(j+1)$ th step, we substitute the current mean updates $\tilde{Y}_{i,j}$ for \tilde{Y}_0 , $i = 1, \dots, n$. The iteration is then constructed as

$$\tilde{Y}_{i,0} = \tilde{Y}_i, \quad i = 1, \dots, n,$$

and

$$\tilde{Y}_{i,j+1} = \tilde{Y}_{i,j} + \epsilon \tilde{\nu}_h(\tilde{Y}_{i,j}), \quad (10)$$

for $j \geq 0$. Here, ϵ is a positive small scaling factor. As $\epsilon \rightarrow 0$, this sequence becomes a density ascent line, in the same way as in the original algorithm of Hall and Heckman (2002).

It is natural to place the curves that converge to the same limit into one cluster. This provides an effective classification method for gene expression data. Caution is needed in cases where the limit may be a local maximum or saddle point (Hall and Heckman, 2002). In the numerical examples, we choose the scalar ϵ and the bandwidth h subjectively. We chose $\epsilon = .1$, which provided a reasonable balance between accuracy and speed of convergence (Hall and Heckman, 2002). Different bandwidths may result in different numbers of mode functions. Our numerical and application results however are robust over a range of bandwidths. We choose bandwidths which lead to well-interpretable results. We will refer to the mode function obtained from Hall and Heckman (2002) as the “unsynchronized functional mode”, and the mode function from the proposed method as the “synchronized functional mode”.

IMPLEMENTATION

We implemented the two algorithms – the unsynchronized mean updating algorithm and the synchronized mean updating algorithm – to obtain the unsynchronized functional mode and synchronized functional mode separately, for a sample of *Drosophila* life cycle gene expression profiles. The algorithms can be visualized by tracing the “ascent paths” that originate at a specific expression profile and end at the corresponding time-synchronized mode function.

Figure 3 provides a comparison of the two types of mode function. In the upper panel are the mode functions as obtained by applying the proposed synchronized mean updating algorithm and in the lower panel are the mode functions as obtained by applying the unsynchronized algorithm of Hall and Heckman (2002). In both panels, the solid curve is the mode function for the transient zygotic gene cluster and the

dashed curve is the mode function for the co-expressed muscle-specific gene cluster. We find that while the transient zygotic mode function is almost the same for the two methods, there are differences for the muscle-specific mode functions. The unsynchronized version has two peaks around the times 27 and 35 units respectively, while the time-synchronized version has only one peak around time unit 30.

Interestingly, if we change the bandwidth in the smoothing step from 1.5 time units to 3 time units, which is also a reasonable bandwidth to use, the unsynchronized mode function becomes more similar to the synchronized version. It then also has a single peak at around time unit 30. Therefore the synchronized mode is shape-invariant in response to such bandwidth changes while the unsynchronized mode is not. It appears then that the synchronized mode is more robust to small changes in individual expression profiles than the unsynchronized version. The time-synchronized version has a tendency to coalesce closely neighboring peaks in expression profiles into one overall peak, due to the time acceleration that the warping method affords in the neighborhood of a peak.

Applying the time-synchronized algorithm for clustering, we classify genes according to the mode function to which their expression profile converges. The observed misclassification error when comparing with the biological classification in Arbeitman et al. (2002) is quite low for both methods (one misclassified gene for the time-synchronized and none for the unsynchronized algorithm). We demonstrate the ascent paths originating at individual gene expression profiles, connecting them with the corresponding mode function, for one gene from each of the two clusters.

In Figure 4, the upper panel demonstrates the path originating at gene CG8606 from the transient zygotic cluster and the lower panel the path originating at gene CG9432 from the muscle-specific cluster. In these panels, the dotted curve is the smoothed original gene expression profile, and the dashed curve is the respective mode function. The solid curves in between are trajectories that lie on certain points on the respective ascent paths. A total of 2000 iterations of the mean update algorithm were carried out; displayed are the resulting trajectories corresponding to iteration steps 20, 40, 60, 80, 100, and 2000. These serve as intermediate points to visualize the ascent paths towards the mode function.

DISCUSSION

Comparison of genes that exhibit similar overall expression may not be valid because different conditions over which they are observed may not be appropriately synchronized. Systematic and effective approaches to the alignment of time-course gene expression data are of interest in the post-genomic era. The increasing prevalence of gene expression profile data provides motivation and opportunity to develop and investigate time-synchronizing algorithms, extending the pioneering work on warping for

gene expression data by Aach and Church (2001).

In this study, we developed a time-synchronized mean-shift algorithm for *Drosophila* life cycle gene expression profiles. We implemented two mean updating algorithms for gene expression profiles: The unsynchronized mean updating algorithm of Hall and Heckman (2002), and the time-synchronized version that we propose here. This version incorporates individual time transformations and thus takes the time variability that is inherent in gene expression data into account. These algorithms were then applied to compute the resulting mode functions for the gene expression profiles and to find two clusters with similar expression profiles. These clusters consist of genes whose profiles are in the domain of attraction of the same mode function.

The first cluster of transient early zygotic genes has been biologically identified as playing an important role in early embryonic development and patterning by self-organizing map (SOM) analysis (Tamayo et al., 1999). These genes are highly expressed mainly during the crucial period of development of cellularization. The mode function that characterizes this cluster and is displayed as solid function in Fig. 3 (upper panel) clearly supports the interpretation of early activity, with a second minor increase of activity visible in the adult stage. The second cluster of co-expressed genes enriched for muscle-specific genes was biologically identified by a hierarchical clustering analysis (Eisen et al., 1998). This set of genes is known to have an expression pattern that coincides with larval muscle development, and the dashed curve in the upper panel Fig. 3 indicates that the mode function in the time interval considered essentially consists of one major peak of expression at around 30 time units, providing an estimate for the timing of maximal expression. The local neighborhoods depicted in Fig. 2 (especially upper panel) demonstrate differences in the timing of some genes within the same cluster, indicating differences in the activation of genes along the developmental time scale.

The strength of the proposed methods is that besides smoothness of the expression profiles and the existence of well-defined modes in function space no other assumptions are needed. Both the mean-shift algorithm as well as the time transformations that are employed to handle the time variability are non-parametric and do not require strong model assumptions. Combining both features, time variability is naturally taken into account. Time-synchronized mode functions provide a robust description of shape. The resulting clustering method has potential for functional genomics due to its nonparametric nature and since it takes into account entire gene expression profiles and their time variability. We introduced here a distance measure between gene expression profiles that incorporates time warping. The development of additional distance measures is of interest since an appropriate choice of a distance has immediate ramifications for the detection of outliers and subsequent clustering and classification.

ACKNOWLEDGEMENTS

This work was supported in part by NSF Grants DMS 99-71602 and DMS 02-04869. We also thank two anonymous referees for their suggestions which led to numerous improvements in the paper.

REFERENCES

- Aach, J. and Church, G.M. (2001) Aligning gene expression time series with time warping algorithms. *Bioinformatics*, **17** 495–508.
- Adams, M.D. *et al.* (2000) The genome sequence of *Drosophila melanogaster*. *Science*, **287** 2185–2195.
- Arbeitman, M.N., Furlong, E.E.M., Imam, F., Johnson, E., Null, B.H., Baker, B.S., Krasnow, M.A., Scott, M.P., Davis, R.W., and White, K.P. (2002) Gene expression during the life cycle of *Drosophila melanogaster*. *Science*, **297** 2270–2274.
- Choi, E. and Hall, P. (1999) Data sharpening as a prelude to density estimation. *Biometrika*, **86** 941–947.
- Cheng, Y. (1995) Mean shift, mode seeking, and clustering. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **17** 790–799.
- Comaniciu, D. and Meer, P. (2002) Mean shift: a robust approach toward feature space analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **24** 1–18.
- Eisen, M.B., Spellman, P.T., Brown, P.O., Botstein, D. (1998) Cluster analysis and display of genome-wide expression patterns. *Proc. Natl. Acad. Sci. U.S.A.*, **95**, 14863–14868.
- Fukunaga, K. and Hosteler, L.D. (1975) The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans. Information Theory*, **21**, 32–40.
- Gasser, T., Hall, P. and Presnell, B. (1998) Nonparametric estimation of the mode of a distribution of random curves. *J. Roy. Statist. Soc. Ser. B* **60** 681–691.

- Hall, P. and Heckman, N.E. (2002) Estimating and depicting the structure of a distribution of random functions. *Biometrika*, **89** 145–158.
- Kneip, A. and Gasser, T. (1992) Statistical tools to analyze data representing a sample of curves. *Ann. Statist.*, **16** 82–112.
- Liu, X. and Müller, H.G. (2002) Functional convex averaging and synchronization for time-warped random curves. Submitted.
- Müller, H.G. (1984) Smooth optimum kernel estimators of densities, regression curves and modes. *Ann. Statist.* **12** 766–774.
- Müller, H.G. and Yan, X. (2001) On local moments. *J. Multivariate Analysis* **76**, 90–109.
- Ramsay, J.O. and Li, X. (1998) Curve registration. *J. Roy. Statist. Soc. Ser. B* **60** 351–363.
- Sakoe, H. and Chiba, S. (1978) Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoustics, Speech, and Signal Processing*, **26** 43–49.
- Silverman, B.W. (1986) Density Estimation. Chapman and Hall, London.
- Spellman, P.T., Sherlock, G., Zhang, M.Q., Iyer, V.R., Anders, K., Eisen, M.B., Brown, P.O., Botstein, D. and Futcher, B. (1998) Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces Cerevisiae* by microarray hybridization, *Molecular Biology of the Cell*, **9** 3273–3297.
- Tamayo, P., Slonim, D., Mesirov, J., Kitareewan, Q.Z.S., Dmitrovsky, E., Lander, E.S., Golub, T.R. (1999) Interpreting patterns of gene expression with self-organizing maps: Methods and application to hematopoietic differentiation. *Proc. Natl. Acad. Sci. U.S.A.*, **96**, 2907–2912.
- Wang, K. and Gasser, T. (1999) Synchronizing sample curves nonparametrically. *Ann. Statist.* **27** 439–460.

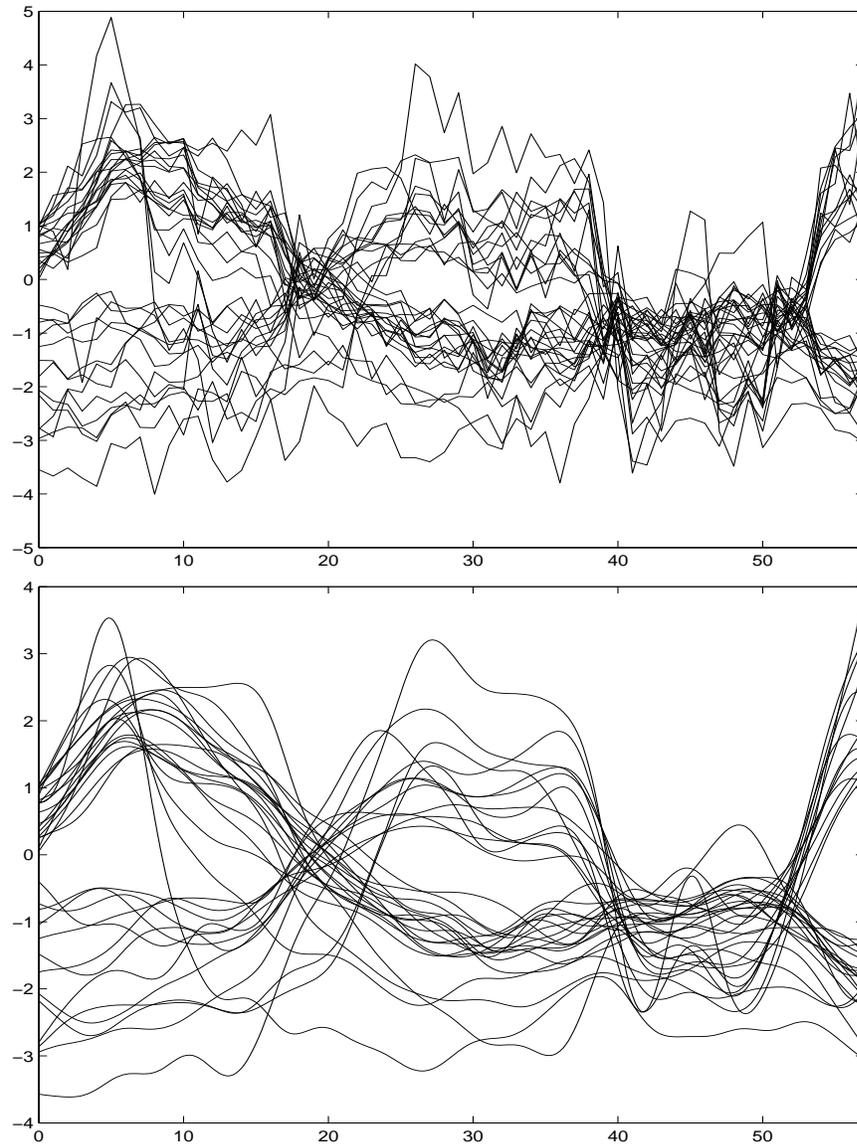


Figure 1: *Gene expression profiles for the 32 classified genes from the study of Arbeitman et al. (2002) for Drosophila. Upper panel: Linearly interpolated observed data. Lower panel: Smoothed data, using bandwidth 1.5 time units and a Gaussian kernel. The x-axis is time with 57 units and the y-axis corresponds to normalized gene expression levels.*

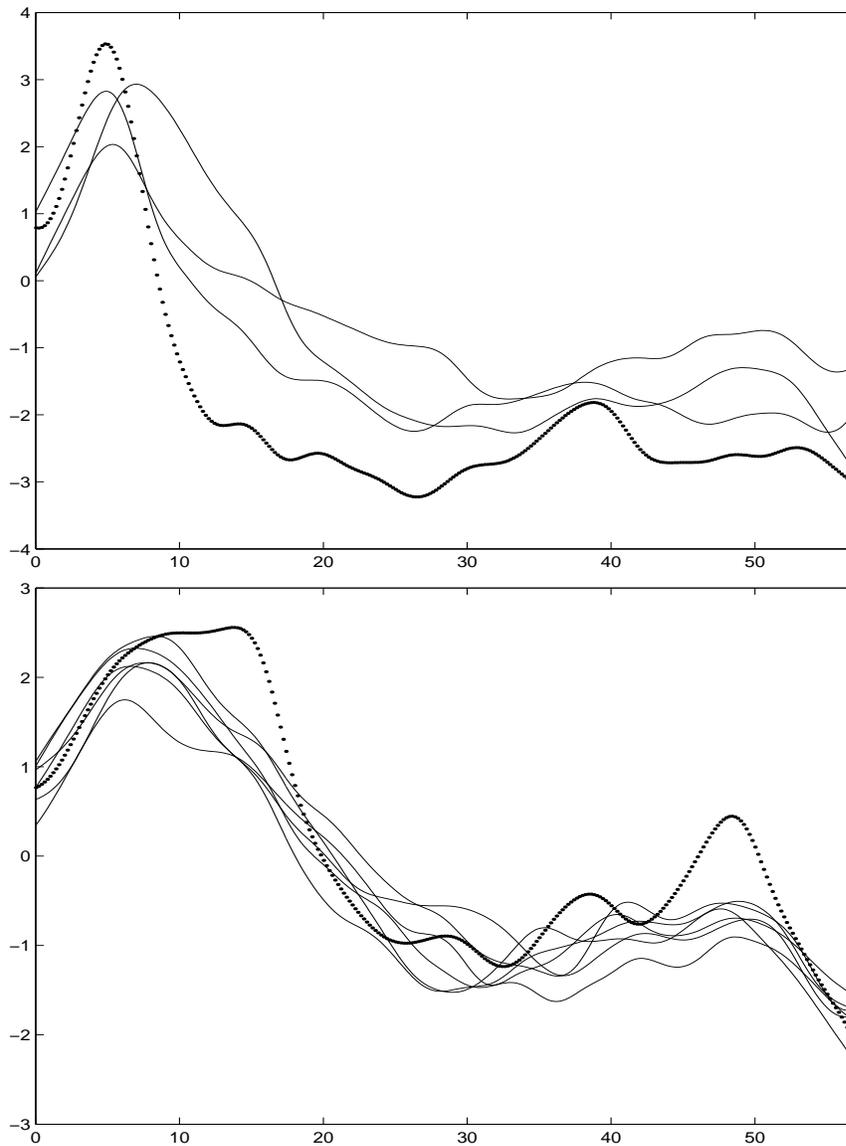


Figure 2: *Two selected gene expression profiles and their spherical neighborhoods in terms of the distance D (5). Upper panel: The spherical neighborhood of radius 1.8201 for the expression profile of gene CG15634 (dotted) contains the expression profiles for genes CG2916, CG14025, and CG9506 (solid). Lower panel: The spherical neighborhood of radius 1.5 for the expression profile of gene CG8606 (dotted) contains the expression profiles for genes CG1078, CG16901, CG7269, CG10417, CG7626, and CG4602 (solid).*

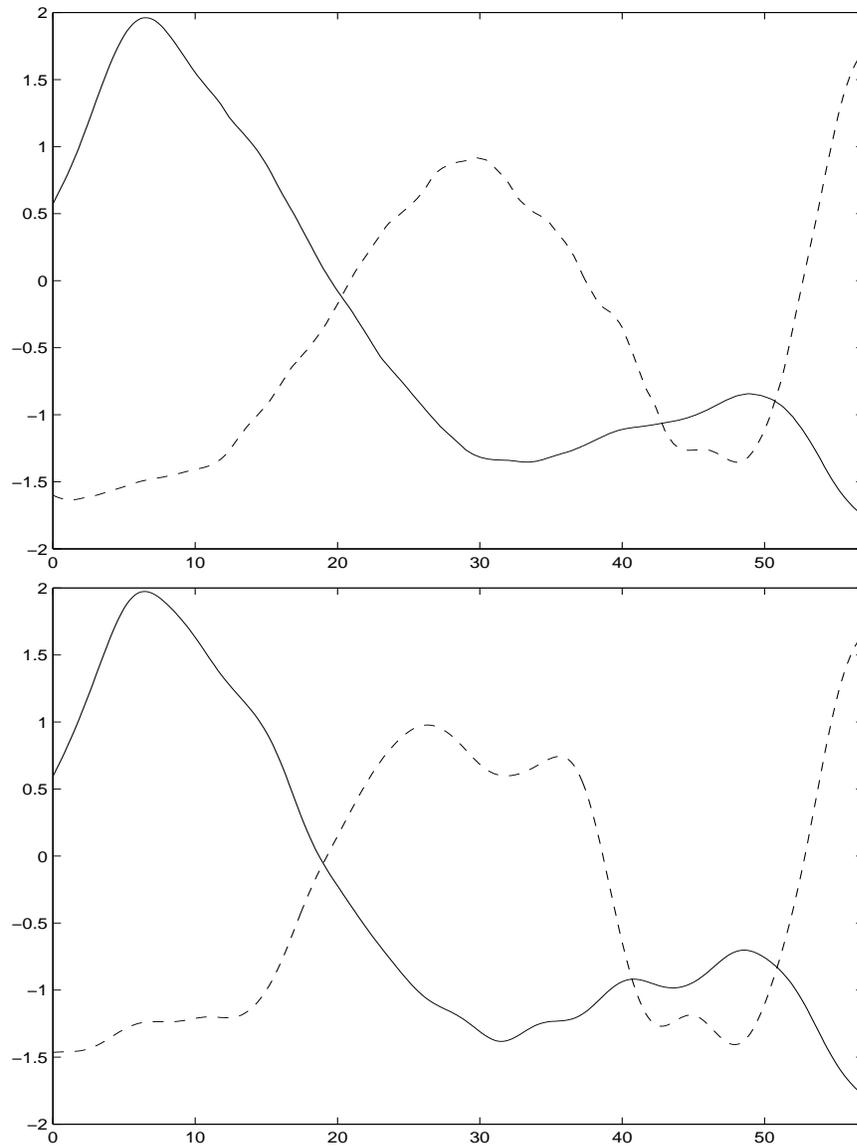


Figure 3: *Mode functions for Drosophila life cycle gene expression profiles. Upper panel: Synchronized mode functions. Lower panel: Unsynchronized mode functions, obtained by the Hall and Heckman (2002) algorithm. For both panels, the solid curve is the mode function for the transient zygotic gene cluster, and the dashed curve the mode function for the co-expressed muscle-specific gene cluster (Arbeitman et al, 2002).*

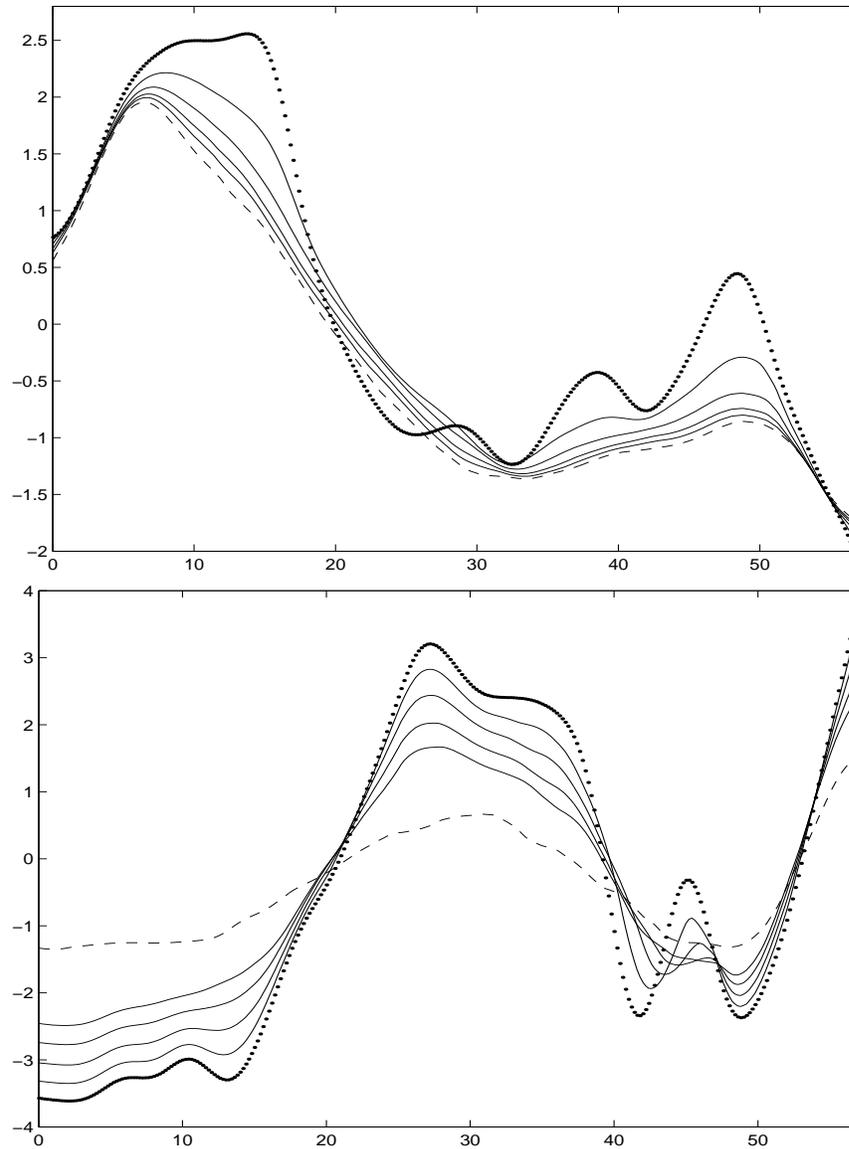


Figure 4: *Illustration of ascent paths to the time-synchronized mode function. Upper panel: Path originating at gene CG8606; lower panel: Path originating at gene CG9432. The dotted curves correspond to the original smoothed gene expression profiles, the dashed curves to the time-synchronized mode function, and the solid curves in-between correspond to selected trajectories that are positioned on the paths.*