

Lory 1.0: Model-based Estimation of Local Shape Deformations

Software compilation: 03/11/2014

Manual revision: 03/11/2014

© Eladio Marquez 2012-2014 / © The Mathworks, Inc. 1984-2013
emarquez@bio.fsu.edu
Dept. Biological Science
Florida State University

Citation:

Márquez, E.; Cabeen, R.; Woods, R.; Houle, D. 2012. The measurement of local variation in shape. *Evolutionary Biology*. *In press*. DOI: 10.1007/s11692-012-9159-6

LORY implements the methods described in Marquez et al. (2012) for the estimation of shape variables that measure *local differences in shape* between individual samples and their Procrustes mean. “Local” refers herein to infinitesimal differences in shape mapped continuously over entire configurations; these differentials give rise to shape variables that differ from other variables such as Procrustes residuals or Partial Warp scores both in their intent and interpretation. LORY’s main output consists of shape variables that can be individually interpreted as local changes in infinitesimal *area* relative to a reference configuration (i.e., the sample Procrustes mean). Throughout this guide, we refer to these changes as *local deformations*.

Overview

When comparing the shape of two landmark configurations, the entirety of information about their differences is contained in the full set of landmark coordinates. Whole-shape information is often useful to estimate summary statistics derived from shape space metrics, such as mean shapes, pairwise distances, sample variances, or overall sample dimensions. For many practical applications of shape analysis, properties of whole configurations are less relevant than those of partitions within them, including their patterns of statistical association. These applications are rather focused on questions pertaining to localized (i.e., “focal”) features. To illustrate this distinction between global and local features, consider a study that asks whether whole shapes are more different among a group species than expected according to some hypothesis. Methods based on shape-space metrics would not only allow estimating appropriate statistics to address this question, but would also provide the necessary information to build *visualizations* to help investigators determine the spatial distribution of differences and similarities among sampled species.

A natural follow up to such visualizations would be testing for interspecific differences at deliberately chosen regions within configurations. For instance, one could wish to investigate whether species differ with respect to a set of traits but not another one. The geometric morphometrics literature contains many examples where answers to questions like these are sought by partitioning a configuration in multiple landmark blocks, and re-applying with each the same methods used for whole shape configurations. This seemingly natural solution, however, ignores the fact that shape is a property of configurations; that each of a set of landmark partitions belongs to a different shape space, and therefore that subsets of landmarks in isolation are not generally comparable to the “same” landmarks when they are considered as part of a

larger configuration. As a general rule, shape deformations are interpreted as local only in relationship to other features of the same configuration, which leads to a seemingly contradictory requirement of a global sampling in order to estimate local features. In practical terms, this means that landmarks outside of a partition affect how the shape of the partition in question is estimated and interpreted, i.e., the shape of a partition is context-dependent.

The context-dependence of estimates of local shape variation has an effect not only on how deformation visualizations are interpreted, but also on the actual values of landmark residuals. For example, a shape-invariant region located in the proximity of a variable region would generally undergo a rigid “displacement” of its landmarks (see Figs. 4 and 5 of Marquez et al. 2012). Should we choose to consider these landmarks as variables without re-superimposing the partition, we would incorrectly conclude that the region is variable and internally correlated, although it is only so in relation to a separate (i.e., unsampled) region. On the other hand, re-superimposing the partition would lead to a better assessment of its shape, but would lead to a loss of contextual information, as superimposition spreads landmark variation over all landmarks and thus destroys any pattern of local variation influenced by surrounding regions.

The approach discussed in Marquez et al. (2012) and implemented in LORY does not require partitioning of configurations to estimate local shape information, and thus avoids most of these problems. As discussed below, this approach uses landmarks not directly as data, but as *parameters* that inform how continuous local deformations are spatially distributed in configurations. Under this approach, shapes become *functions* that are only discretized as variables via *evaluation* of such functions at fixed, pre-determined locations.

Basic rationale

What does it mean for shapes to be represented as functions? Consider the plot in Fig. 1, depicting the shape of a configuration (with respect to a reference configuration) using a typical deformation grid.

Upon being asked to interpret these deformations, we would intuitively list a series of features, such as

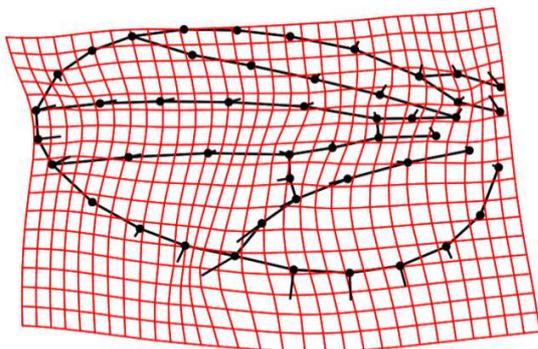


Figure 1. Deformation grid depicting implied shape changes between two configurations.

local expansions and contractions, both relative to each other and to the reference configuration of choice. By doing this, we attempt to synthesize the complexity embodied by a direction of variation spanning multiple dimensions. In general, interpretations of this sort are not easily stated in terms of individual landmarks; instead, describing shape deformations associated to a specific region often requires considering the behavior of landmarks beyond this region before deciding whether a particular change is a contraction, an expansion, a combination or neither of them. The information we need to devise a proper interpretation is thus contained in the entirety of the configuration.

During the process of interpretation, we implicitly make assumptions about the shape changes undergone by regions in between landmarks. This is true irrespective of the number or density of landmarks we have sampled. A localized expansion, for instance, would seem a reasonable interpretation for a region whose

bounding landmarks seem to point outwards relative to each other. However, this pattern could also be consistent with a potentially large number of reasonable alternative hypotheses, such a scenario whereby only the outermost ridge of this focal region is undergoing an expansion.

From these considerations we can derive two important conclusions; namely, (1) individual landmarks do not contain *quantitative* information about local shape, and (2) a single, fixed configuration of landmarks can be *consistent* with multiple shape deformation scenarios. Therefore, in order to estimate local shape differences, we should be able to integrate information derived from whole landmark configurations by way of (preferably) *known* or knowable functions, as opposed to focusing on isolated subsets of landmarks. Landmark coordinates, as well as their means, variances, and covariances, have no meaning as shape descriptors, and in fact can be misleading, as evinced when arithmetic operations are applied to entire landmark coordinates with barely any effect on their distributional moments, but often causing dramatic effects on the shape to which they belong. For example, a simple sign inversion applied to a coordinate will only change the sign of its mean, leaving its variance and other moments unchanged. When considered in relation to other landmarks, however, a sign inversion can imply a shift between a contraction and an expansion, or even between a variable and an invariant shape (see Fig. 4 of Marquez et al. 2012).

The preceding discussion shows that *measuring* shape, and particularly local shape, may be more appropriately approached as a decision process in which multiple competing models are compared for optimal fit to a given sample of landmark configurations. A discussion of exactly *how* such comparisons are carried out is outside of the scope of this paper and software, where we focus on the more basic problem of how to estimate local shape differences given a model. A general discussion on the opportunity that a model-based approach presents for biological studies of shape variation is presented in Marquez et al. 2012.

It should be noted that a model-based approach is not the *only* way to correctly measure shape. As discussed above, information about differences between two shapes is fully contained in the whole configuration of landmarks, and similarly information about shape variation is fully described by the variance-covariance matrix computed from a sample of (whole) configurations, as well as by any full-rank basis upon which these configurations are projected (e.g., Principal Warps). Neither of these approaches, however, can be used to produce local shape estimates.

Setting aside the model selection issue for now, two questions that remain are (1) how are landmarks used to parameterize these models, and (2) how are modeled functions used to generate variables amenable of statistical analysis. A single family of methods, spatial interpolation, provides an answer to both of these questions.

Interpolation functions

Interpolation is the mathematical process of assigning a value to an unobserved feature based on information provided by observed ones. In the context of shape analysis, interpolation refers to the prediction or estimation of deformations at unsampled locations based on patterns of deformation observed at sampled landmarks. Interpolation is used implicitly during ordinary interpretation of shape deformations, or explicitly when a mathematical function is used to model deformations of regions devoid of landmarks. In either case, interpolation assumes that information is transferred smoothly from sampled landmarks to inter-landmark spaces, i.e., shape deformation information is *ergodic*. A corollary of these

premises is that information *quality* is largely a function of the number of sampled features, i.e., the density of landmark sampling. This notably contrasts with typical geometric morphometric applications, where sampling focuses chiefly on landmark *quality* (e.g., strict homology), as opposed to quantity. Whereas homology concerns are still important to ensure landmark-wise comparability across a sample, an interpolation-focused approach benefits from extensively using “fuzzy” homologous features, such as curves, surfaces, and volumes, whether they are sampled as functions or as collections of pseudo- and semi-landmarks.

Landmark-based interpolation *per se* does not produce local shape data. Instead, interpolation functions allow us to describe configuration shapes *globally*, so that we can choose the sampling scale *a posteriori*. Working with functions offers a great deal of versatility compared to working with landmarks, both in terms of the scale at which hypotheses can be tested, and in the type of measurements that can be used for these tests: functions, as well as their derivatives and integrals can be directly used as data or evaluated at arbitrary locations, providing thus a convenient way to combine them into traditional (e.g., landmark-based) morphometrics applications. This is precisely what LORY is intended for: to fit interpolation functions to samples of landmarks and to provide point estimates of shape deformation by evaluating these functions at sensibly chosen locations.

There are several ways to implement spatial interpolation, the most relevant of which are discussed in Marquez et al (2012). Most of these approaches, notably splines, kriging, and finite element methods, become mathematically indistinguishable under certain conditions. In its current version, LORY implements 2-D *interpolating splines* based on two widely used interpolation functions, namely Thin-plate Splines (TPS; Bookstein 1992) and Elastic Body Splines (EBS; Davis 1995, 1997). Figs. 2 and 3 contrast outcomes from the two models using two visualization styles. TPS models landmark displacements as *bending* (i.e., irreversible) deformations of rigid objects, and therefore global deformations are “easier” (less penalized) than local ones. EBS, on the other hand, models landmark displacements as *elastic* (i.e., reversible, spring-like) deformations of soft objects, so that localized

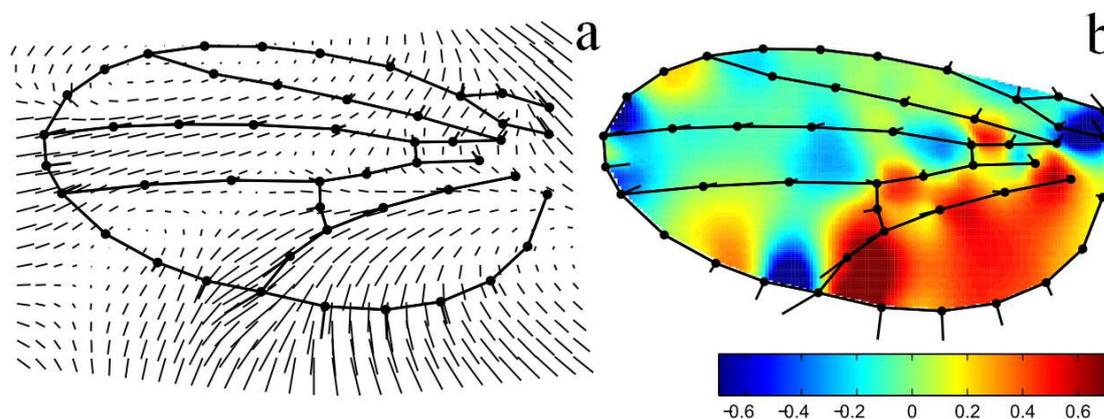


Figure 2. Alternative visualizations of shape deformation: a. velocity ("quiver") plot, with vectors drawn on a regular grid denoting implied deformation; b. "parrot" plot where hues represent relative proportional local changes in area (red hues: expansions, blue hues: contractions).

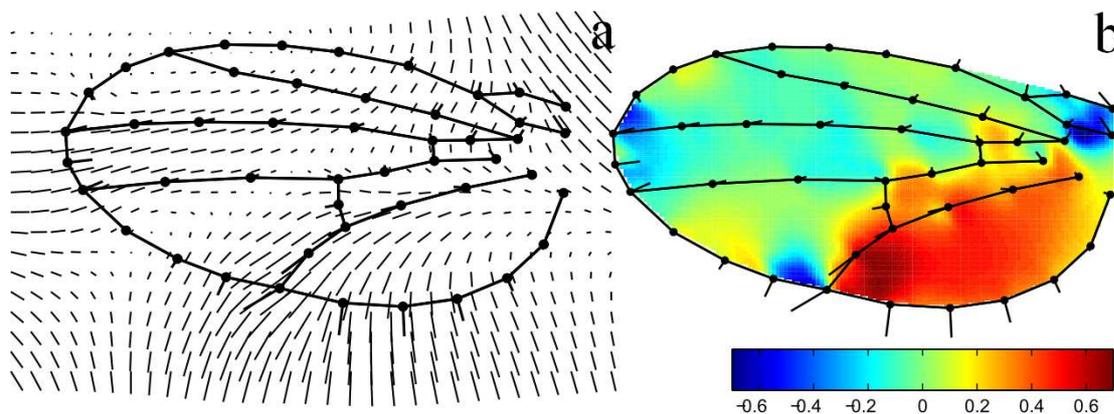


Figure 3. Elastic-body splines deformation plotted using two alternative visualization styles: a. velocity ("quiver") plot, with vectors drawn on a regular grid denoting implied deformation; b. "parrot" plot where hues represent relative proportional local changes in area (red hues: expansions, blue hues: contractions).

deformations are not heavily penalized. Accordingly, TPS will tend to spread relatively localized landmark displacements further into unsampled regions than EBS does.

Figures 2 and 3 show two visualization styles implemented in LORY. Fig. 2a and 3a show velocity ("quiver") plots where a 2-D spline is interpolated for each node of a regular grid and shown in the same style as a typical true landmark deformation. In these plots, each node is associated to an x - y -element vector, and just like actual landmarks, their usefulness as local shape estimators is rather limited, in that each vector estimates a displacement at a single point, without integrating information provided by its surroundings. One way in which such integration is accomplished is via local evaluation of the Jacobian matrix of the interpolation function (Woods 2003, Marquez et al 2012). These Jacobians contain first-order partial derivatives of interpolating splines, and thus describe the orientation of a deformation function at any point within its domain. Whereas a vector spline represents a point displacement devoid of actual local shape information, a Jacobian contains information needed to determine whether such displacement is part of a local contraction or expansion, as well as to estimate the magnitude of the deformation.

Jacobian matrices are, however, relatively impractical as measurement variables because of their sheer dimensionality (for 2-D and 3-D data, each Jacobian is a 2×2 and 3×3 matrix, respectively), and thus it is often preferable to use their determinants (or log-determinants) instead. Jacobian determinants contain only part of the information in Jacobian matrices; specifically, whereas matrices provide a full directional description of deformations at a point, determinants only tell whether *area* (or volume, in the 3-D case) is increasing or decreasing in the infinitesimal region surrounding this point, relative to a reference configuration. We can therefore ask if the disadvantage of losing detailed local information outweighs the practicality offered by a single summary variable. In Marquez et al (2012), we argue for using determinants, given their standalone interpretability as local expansions/contractions, as long as a reasonably high density of them is sampled simultaneously. The resulting "field" of scalar measurements can be visualized as heat maps, such as those shown in Fig. 2b. In these *parrot plots*, intensity of red hues represents magnitude of local expansions, whereas blue hues represent local contractions (greens are approximately invariant regions, with respect to the reference configuration). Clearly, expansions/contractions are not synonymous to "shape," but this is actually seen as an advantage as these

two behaviors, when consider over the entire configuration, *define* shape, and at the same time remain *individually interpretable*.

Computational details for Jacobians are provided in Woods (2003) and Marquez et al (2012). Like shape, Jacobian matrices exist in a Riemannian space and therefore LORY projects them onto a tangent space centered at the Procrustes mean configuration. Outputs from LORY include Jacobian matrices and base-2 logarithms of their determinants (simply termed “Jacobians”). These values can be interpreted as proportional (with respect to the mean) changes in local area (see Fig. 2b); e.g., $\log_2\text{-Jacobian} = 0$ represents no change, $\log_2\text{-Jacobian} = 0.5$ represents a local halving in surface area, and $\log_2\text{-Jacobian} = 1$ represents a local doubling in surface area. Although LORY can currently deal only with 2-D data, all of these methods have been developed and are readily applicable to 3-D data as well.

Evaluation of functions

In order to analyze shape functions using traditional multivariate methods, these functions have to be evaluated at specific locations to yield Jacobians that can be treated as ordinary shape variables. In the context of biological morphometrics, the question is *where* to locate such evaluation points. Normally, homology across configurations provides the necessary criteria to ensure comparability and meaningfulness of variables, and applications discussed herein are no exception to this. Shape spaces and their metrics are built upon the premise that transformations among distinct shapes are topologically correct, and methods presented herein depend upon the same assumptions that underlie shape spaces.

Assuming that landmark homology is not problematic, the question is how to define evaluation (i.e., inter-landmark) points that can be assumed to be approximately homologous. Because any such homology criterion must be derived from actual landmarks, Delaunay triangulation seems an optimal candidate to find inter-landmark regions that can be uniquely and unambiguously identified on a sample of configurations. Similarly, a variety of point features can be defined with respect to the triangles found using this algorithm (e.g., incenters, circumcenters, Voronoi vertices, centroids, etc.), and these points can in turn become function evaluation locations or the vertices for further passes of the Delaunay algorithm, to create an increasingly finer evaluation mesh (see Figure 3 of Marquez et al 2012). The specific way in which LORY implements these ideas are discussed below.

Intended applications

Local estimates of shape deformation can be used in any context where landmarks are appropriate. They are, nonetheless, particularly suited for two general families of questions, namely those interested in statistical properties of *parts* within configurations, and those dealing with specific models to explain shape and shape variation. The latter category is at presently just starting to take flight with disciplines such as computational morphodynamics, whose focus has not yet turned to the study of variation, whereas the former involves longstanding topics related to variational properties of individual parts (e.g., character evolution) as well as the statistical association among parts (e.g., heterochrony, morphological integration, modularity). LORY provides estimates of both Jacobian matrices and determinants, while offering a handful of alternative methods to determine evaluation sites.

Like most software for geometric morphometrics analysis, LORY has a strong visualization component that can be used in connection to local shape analyses or as a standalone feature. Deformations can be visualized as vectors, regular and tessellated grids, via parrot plots, quiver plots, and as animations, making LORY a useful tool for general interpretation of shape differences and for exploratory analyses of

local shape differences. The following sections explain how to use the software and offer pointers on how to interpret some of its outputs.

Preliminaries

Setup

LORY executable, Lory.exe, is coded and compiled in Matlab 7.9 (R2009b), using Matlab Compiler 7.11. In order to be able to run LORY, the corresponding Matlab Component Runtime (MCR) libraries must be installed. Currently, there are versions of LORY only for Windows 32- and 64-bits, which should run in the corresponding Windows versions up to 7. These versions should be matched to the appropriate (32- and 64-bit) MCR libraries. These libraries are installed via the program MCRInstaller.exe, which can be downloaded from elsewhere in the web. We have made this file available at <http://www-personal.umich.edu/~emarquez/morph>. Note that LORY may not work with versions of libraries other than 7.11.

After libraries are properly installed, LORY can be saved and run from anywhere in your computer. If you have the correct libraries version and are experiencing problems running LORY (e.g., you get an error message regarding a missing DLL file), then you should modify your Windows path so that LORY can find the MCR libraries when run. You should be able to find instructions for your specific Windows version by googling questions such as *how to add a directory to the Windows PATH environmental variable*. You should then add the folder `\MATLAB Component Runtime\711` to Windows path.

Data formats

At the moment LORY accepts two data formats: basic James Rohlf's TPS format, and space- or tab-delimited "XYCS" matrix format, which contain one row per specimen, and coordinates are arranged in {x,y}-pairs (i.e., $x_1, y_1, x_2, y_2, \dots$), with an (optional) additional column containing centroid size. Regarding TPS files, LORY can only read SCALE= tags, and might ignore (or fail upon finding) other specialized tags, such as CURVES=. Even though LORY cannot read TPS-formatted curves, files containing them can be transformed into XYCS-formatted matrices using the program SemiThinner, also available at <http://www-personal.umich.edu/~emarquez/morph>.

Loaded landmark data do not have to be already superimposed, as LORY does a Procrustes alignment as the first step upon successful import. LORY also offers an option to remove extra landmarks added during digitization to provide a scale reference (i.e., "ruler" landmarks). LORY will scale the data either to this ruler or to the SCALE= tag in TPS files, giving priority to the former whenever both are found. Note, however, that scale is not relevant in any of the analyses provided in this release of LORY.

Using LORY

Overview of the Interface

LORY's interface features are shown in Fig. 3 using the same numbers used in the descriptions and explanations that follow. Components 1-3 control for data input, 4-7 define how evaluation points for interpolations are chosen, 9 offers two interpolation function options, 10-12 control data outputs, 13-17 control graphical outputs, and 18 indicates a message area.

In addition to the main interface, a **console** screen opens along with LORY. This screen has to remain open while using LORY, and it will contain, for the most part, error messages. It is a good idea to send any error message reported in this screen when contacting us for support.

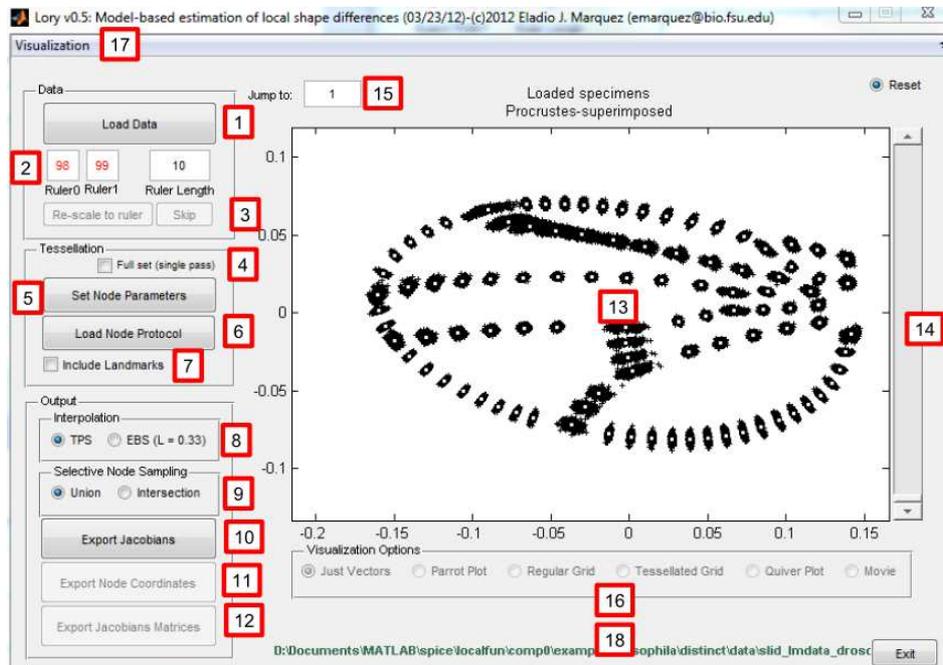
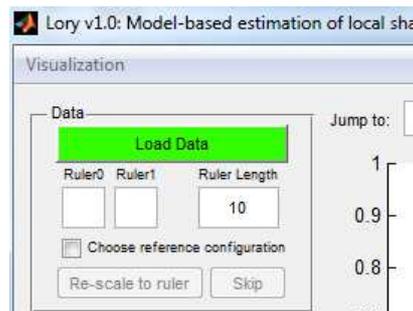


Figure 3. LORY main interface. See text for legend explanation.

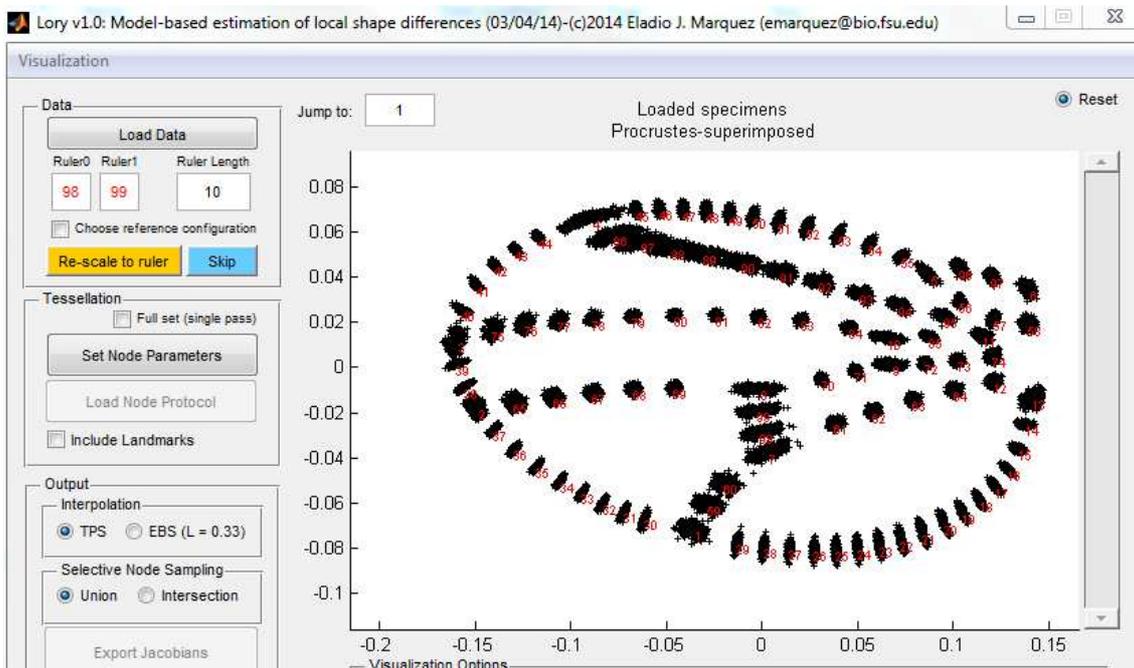
Loading data

The following assumes that you have formatted your TPS or XYCS files according to LORY requirements. When first opening the program, you will notice that the *Load data* button (1) has been highlighted to indicate that this is the only available function:



Upon successfully reading your data file, data contents are graphed in the plot area (13), and the user is prompted to choose between two options for data post-processing:

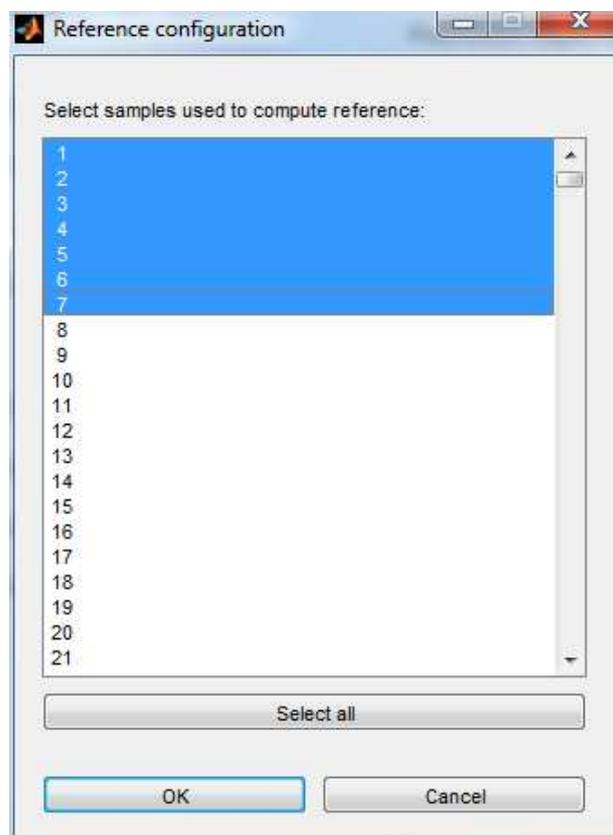
If loaded data contain two extra points for scaling (ruler landmarks), the coordinates corresponding to these landmarks and the length of the ruler should be entered in marked spaces (2) and the *Re-scale to ruler* button (3) should be pressed. If no scale landmarks are included, but every point in the data set is to be considered as a landmark, then *Skip* (3) should be pressed instead.



If the option *Choose reference configuration* (3) is checked, when either of these buttons is pressed, a new list window appears (shown on the right) prompting the user to select one or more among the loaded configurations to use in the calculation of the reference configuration. If more than one configuration is selected, Lory uses the mean of these specimens as the reference with respect to which all deformations and Jacobians are computed. Whether a single or multiple configurations are used to calculate the reference, Lory proceeds first by carrying out a Generalized Procrustes alignment (i.e., superimposing to the Procrustes mean), and *then* it re-superimposes the data to the new reference.

Visualizations

LORY's visualization style panel (16) offers six types of (largely equivalent) visualizations, listed below. Each can be used to plot individual configurations as landmarks only or as landmarks connected by a custom

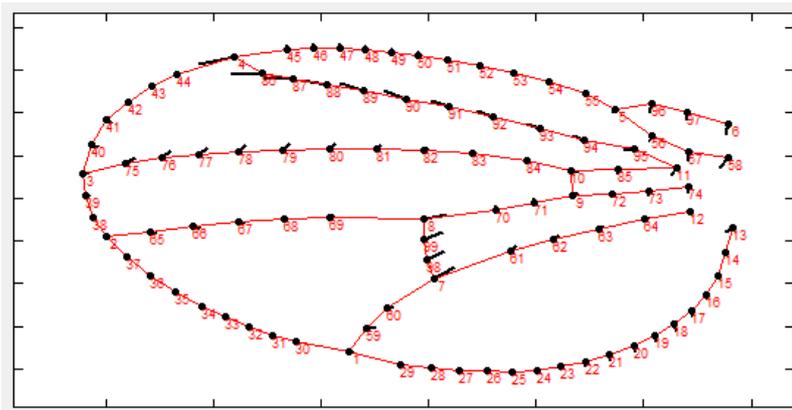


wireframe. Examples shown here include one such wireframe, see below for instructions on how to use this feature.

Plot styles


Just vectors

In this plot, differences between each individual and the Procrustes mean are plotted as vectors on the latter. Landmark numbering follows ordering in the original data file, and is included as reference for verification and protocol creation.




Parrot plots

These plots are produced via interpolation of landmark displacements throughout inter-landmark regions. For the chosen interpolant function (currently, TPS or EBS, selectable from panel (8)), the Jacobian is evaluated at each node of a triangulated grid built upon parameters set by user, and then interpolated throughout a variable-resolution pixel grid using cubic interpolation, after which each “pixel” in the grid is assigned a color according to a customizable color map, set to range between also customizable boundaries. Customizations are accessible from the visualization menu (17), and are detailed below. Each color represents a value for the log-2 determinant of the Jacobian, i.e., 0 represents no change in local area between individual and reference, -0.5 represents a local halving of the area, +1 a local doubling, and so on.




Regular grids

By default, regular grid deformations are plotted on top of parrot plots. These grids are the typical visualization in geometric morphometric applications, where the chosen interpolation function is evaluated at each node of the grid. The density of the grid is partly customizable. Other visual attributes, such as color, thickness, or coverage are fully customizable from the visualization menu **(17)**.




Tessellated grids

In contrast to regular grids, tessellated grids are defined by the nodes of triangulated landmarks. LORY offers a number of customizable options for generation of these grids, accessible from the Tessellation panel **(4-7)** and discussed below.




Quiver plots

These plots are equivalent to regular grid plots, except that instead of a grid, actual splines at nodes are drawn. Customization options for this plot are the same as for regular grids.



Movies


This option produces a loop animation of the reference-specimen deformation. This is the only plot style that requires a wireframe, and if none has been loaded to memory, LORY will request one before showing the movie. Upon selecting a wireframe protocol file (see below for details), the following options window is displayed:



Upon clicking OK, the movie will be played back and **automatically saved** as a GIF file to the specified location using these parameters.

Superimposed data

An additional plot style provided by LORY consists of the entire Procrustes-superimposed data set, which is always accessible by dragging the specimen selection slider (**14**) to its bottom.



Switching between individual configurations

The specimen selection slider (14) can be used to browse across individual specimen deformations. Alternatively, the *Jump to* box (15) can be used to directly access any arbitrary specimen. Individual configurations are numbered in the same order they are in the original data set; specimen numbers are printed atop the plot area for each specimen.

Using wireframes

With the exception of movies, wireframes are optional for all plots. They can be loaded and toggled on/off from the Visualization menu (17) option “Use wireframe protocol.” A correctly-formatted protocol contains two columns and as many rows as “wires” connecting two points. Each row contains the landmark number of the two endpoints of a wire, where landmarks are numbered as in the original data set, as shown by plots with the *Just vectors* style (16). Therefore, not all landmarks need to be mentioned in a protocol, but all landmarks in a protocol must be present in the data. An example data set and wireframe protocol files have been included with this release.

Deformation parameters

Accessible from the Visualization menu (17), selecting *Modify visualization parameters* opens a series of prompts where users can specify the following graphical properties.

- *Deformation magnification*: signed value applied as a product to individual deformations. Currently allows values between -100 and 100; use -1 to simply revert the direction of the deformation. Default: 1 (i.e., original scale)
- *Extent of white space*: proportion of white space within graphs beyond the most extreme point in all directions (including grids). Default: 1.2 (i.e., 20% of white space).
- *Extent of grid*: proportion of grid that extends beyond the most extreme point along the shortest direction of a configuration. Default: 1.2 (i.e., 20%).

- *Number of columns in grid*: sets the number of columns in regular grids, while the number of rows is automatically calculated to ensure grid cells are square. Default: 30.
- *Triangulation iterations for interpolation*: this parameter controls the resolution of the triangulated grid used to build parrot plots. The first iteration of the Delaunay triangulation uses all loaded landmarks and semi-landmarks, whereas additional iterations use the centroids of triangles found in previous ones to find increasingly resolved triangles. Centroids of triangles found in the final iteration are used as evaluation points for the chosen interpolation function (currently, TPS or EBS), selected in (8). To visualize the grid set by this option, match the number of tessellation passes in (5) to this number, set both thresholds to a very low number (e.g., 1e-16), and select *Tessellated Grid* from the plot style panel (16). Note that setting this parameter to a high value may substantially increase the computation time for parrot plots. Default: 1.
- *Vector magnification*: multiplier for vector length. This factor does not apply to grids or parrot plots, whereas for vectors it stacks on top of deformation magnification. Default: 1 (i.e., original scale).
- *Pcolor plot resolution*: resolution of the pixel grid used to evaluate Jacobians during parrot plot creation. This is the number of pixel along a row of the grid (number of rows is automatically calculated). Leave this field blank to omit showing a parrot plot altogether. Default: 150.
- *Colorbar maximum/minimum*: these are the upper and lower limits in actual deformation units for the color map scale in parrot plots. If either of these limits is left blank, LORY finds optimal (but asymmetrical) limits. Default: [-0.3,0.3] (i.e., highest color intensity is reached at 30% increase and decrease in local area with respect to reference configuration).
- *Poisson constant*: this parameter (labeled as L in LORY's interface) is used by EBS to fine-tune the elasticity of the model to deformations. Values closer to 0 correspond to low-resistance, high-elasticity models where local deformations are favored, whereas values closer to 0.5 correspond to rigid materials, where global deformations are favored, resulting in deformations that are similar to those obtained from TPS. Default: 0.33.

Additional plot appearance options

The following width and color options are accessible from Visualization menu (17), following *Modify plot appearance*.

- *Grid color and width*: applies to both regular and tessellated grids.
- *Vectors color and width*: applies to vectors drawn on landmarks, tessellation nodes, and quiver plots.
- *Wireframe color and width*: applies to static and animated plots.
- *Landmark markers color and size*: applies only to landmark markers.
- *Include arrows*: applies to vectors drawn on landmarks, tessellation nodes, and quiver plots.

Acceptable line width values are positive numbers, whereas acceptable colors are listed along with each item. Note that literal entries consist of single literal shortcuts, as indicated in the Windows prompt.

Finally, color maps/schemes are also changed from the Visualization menu, item *Choose color scheme for parrot plots* (Default: "Jet") Notice that some schemes (e.g., "Jet", "HSV") are multidirectional, where the intensity of multiple colors is charted, whereas others (e.g., "Gray", "Copper") are one-dimensional,

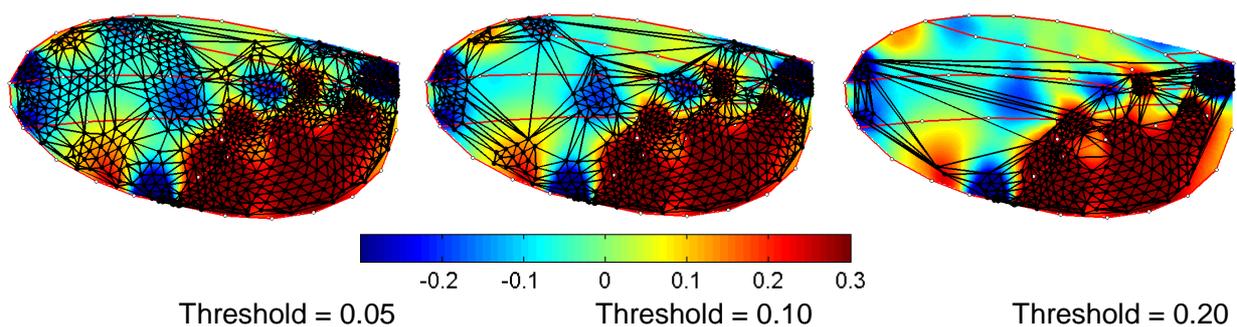
Note that although all triangles are based on the original landmarks, the landmarks themselves are not used as nodes. This is intentional, as nodes are *not raw data*, but an attempt to discretize evaluation points for a continuous function in a way that respects topological relationships between actual data points. In general, it is not advisable to evaluate interpolation functions at the landmarks, because a derivative (including Jacobians) are not defined or are ill defined at landmarks for many functions. LORY, however, allows including landmarks among triangulation nodes by checking the *Include Landmarks (7)* option. Any interpolation function will reproduce the same value observed at landmarks when evaluated at that same landmark.

Node selection

For any given data set, node protocol, and number of tessellation passes, LORY will compute a unique set of nodes that can be completely or partially included in outputs. To get sampling with reasonable low redundancy, one could use a small number of passes, or alternatively optimally select informative nodes from a highly resolved grid. The latter is the basic principle behind LORY's implementation of node selection.

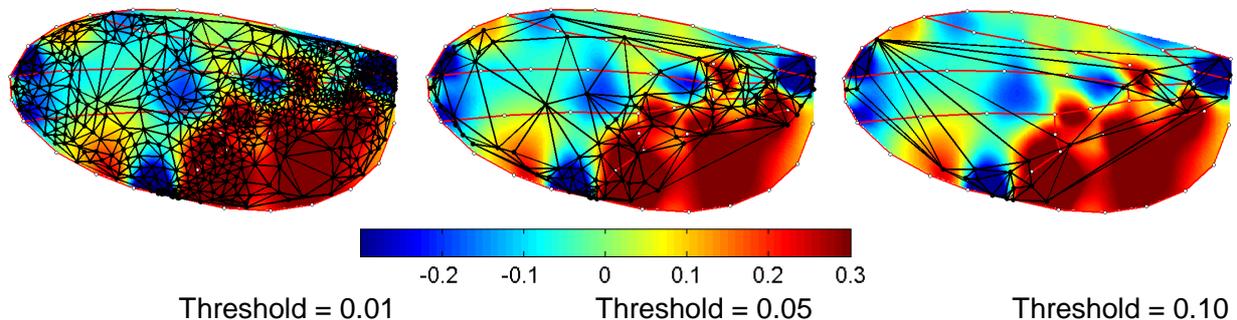
LORY applies two filters using thresholds modifiable through *Set Node Parameters (5)*, which seek to minimize *within-configuration* redundancy and to maximize *between-configuration* variation, respectively. The two filters are applied independently, so that it is possible to apply them individually. To ignore a filter, its threshold should be set to a very low (but positive) value (e.g., $1e-16$). Ignoring *both* filters leads to unselected nodes, such as those depicted above, and this is the setting to which LORY defaults when the *Full set (single pass) option (4)* is checked.

Both filters are evaluated in terms of local area change (i.e., \log_2 of Jacobian determinants). LORY applies the sample (i.e., *between-configuration*) filter first, which simply ignores all local area change values below the set threshold. For example, the following figures illustrate the same 3-pass tessellation shown above, after applying a sample filter with thresholds 0.05, 0.1, and 0.2 (removing, respectively, regions with lower than 5%, 10%, and 20% absolute changes in local area), while keeping the within-configuration threshold at $1e-16$:



The second filter (i.e., *within-configuration*) compares *adjacent* nodes and keeps those whose proportional area change (i.e., \log_2 of the Jacobian determinant) is larger than specified by the corresponding threshold. The algorithm first sorts nodes in descending order of local deformation, and then systematically processes them individually eliminating adjacent nodes that are too similar. The net effect is a node distribution that tends to favor regions of relatively strong local change, which presumably possess higher information content. The following figures result from applying within-

configuration filters of 0.01, 0.05, and 0.1 to the 3-pass tessellation shown above, while keeping the sample threshold at 1e-16.



In addition to controlling thresholds directly, node filters can be globally modified by altering the magnitude of the deformation magnification via the *Visualization (17)* menu, item *Modify visualization parameters*.

Interpolation functions

In its current version, LORY implements two models (8), thin-plate splines (TPS) and elastic body splines (EBS), as described by Dryden and Mardia (1998; see also Bookstein 1992), and Davis *et al.* (1995, 1997), respectively. Differences between estimates derived from the two models are particularly acute for configurations with ostensible localized deformations, which TPS tends to distribute globally whereas EBS tends to contain locally. As discussed above, EBS' Poisson parameter provides some degree of freedom to fine-tune the localness of deformations predicted by the model. Marquez *et al.* (2012) provide further discussion about possible future directions regarding the implementation of alternative interpolation functions.

Selective node sampling and output

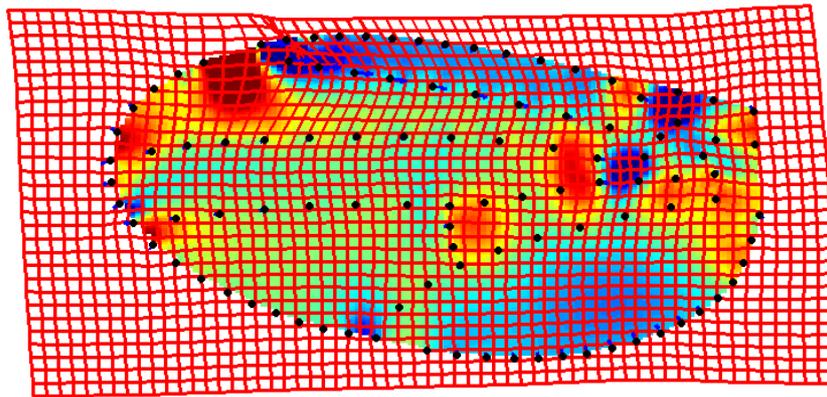
Node filtering, as described above, is applied individually to each configuration in a sample. To build a sample of Jacobian matrices and determinants, LORY ensures that the *same* nodes, as defined during triangulation by their topological relation to reference nodes, are sampled for all specimens in the sample. LORY currently produces three kinds of output, namely an $n \times k$ matrix of Jacobian determinants (10), where n is sample size and k the number of Jacobians sampled per configuration, a $2n \times 2k$ block-matrix of 2×2 Jacobian matrices (12), and an $n \times 2k$ matrix of coordinates for nodes corresponding to each sampled Jacobian (11).

Unless both thresholds are ignored and the full set of nodes sampled for all configurations are used as evaluation nodes, it is likely that distinct subsets of nodes are selected for different configurations. LORY offers two approaches to deal with these disparities, namely set *union* or *intersection* (9). When *Union* is selected, LORY samples and saves every node present in 5% or more of the sample for all individuals, whereas when *Intersection* is selected, only the subset of nodes present in all individual is included in the output. In general, set union produces more reliable results than set intersection, since it is not uncommon for the latter result in an empty set, especially when the density of the interpolation grid is high. When threshold values are too high for LORY to find sufficient evaluation points for a particular configuration during computation of individual Jacobians, the following message is displayed in the console:

```
Not enough evaluation points defined based on current set of parameters.
```

This will lead to a further, similar error when using set intersections as the rule for sample Jacobian computations, whereas it will rarely lead to such an error when using set unions.

A high density of sampled landmarks and semi-landmarks can also lead to interpolation functions to become undefined for individual configurations bearing highly deformed local regions, with respect to the reference configuration. As an example, consider the following TPS deformation grid:

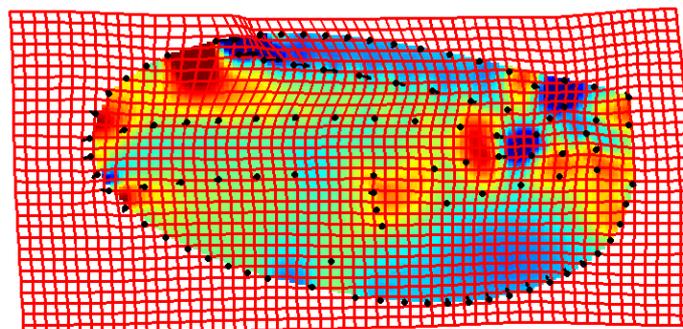


Note at the upper margin of the shape that the grid collapses or “flips,” which leads to the function to become locally non-derivable in Euclidean space, such that Jacobians and local area changes are undefined at any evaluation point in that region. The occurrence of this effect is function-dependent, and elasticity-based models like EBS are in general more tolerant, although not entirely immune, to this sort of pattern.

When LORY finds such disruptions during computation of Jacobians, the following message is displayed in the console:

```
Warning: some of the Jacobian determinants are negative in configuration #N.  
Only real part has been retained, which may produce inaccurate results.
```

This warning should not be ignored, because incorrectly computed Jacobians can and will introduce misleading errors that may become prominent in statistical analyses. There are several ways to deal with this problem, such as manually deleting the offending individual configuration or evaluation point or reducing the magnitude of the deformation. An automatic and generally less intrusive solution, however, would be reducing the density of landmarks or semi-landmarks surrounding highly variable landmarks. Applying this latter solution to the comparison above yields the following pattern:



Note that local deformations are virtually unaffected by this change.

Worked example

Problem and data

In the following example, we briefly illustrate some of the uses of Jacobians by exploring patterns of intraspecific variation and interspecific differences between wing shapes of two closely related *Drosophila* species, namely *D. melanogaster* and *D. willistoni*. The data for this example consists of 289 individual wing configurations, 149 *melanogaster* and 140 *willistoni*, all of which correspond to female specimens. A total of 14 landmarks and 85 semi-landmarks are sampled from each configuration; all landmarks and 32 semi-landmarks are used as reference nodes for triangulation. Node parameters are set to *Full set (single pass)*, which corresponds to one iteration of the Delaunay algorithm without any selection of evaluation nodes. Figure 4 shows sample mean shapes for both species. A portion of this data set has been packaged with LORY (with file name `Drosophila_wing_example.dat`), along with a node protocol (`refnodes.txt`), a wireframe protocol (`wireframe.txt`), and group labels for each configuration in the data file (`specsexgrp.txt`). Using these data and protocols, LORY finds 74 nodes for evaluation, which are the basis for the following analyses.

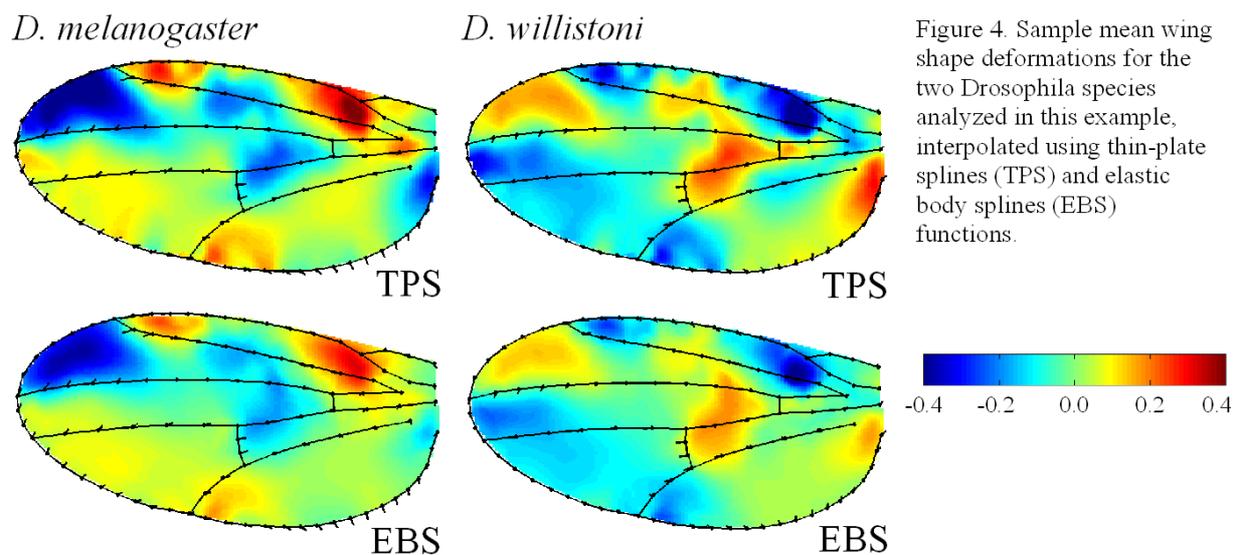


Figure 4. Sample mean wing shape deformations for the two *Drosophila* species analyzed in this example, interpolated using thin-plate splines (TPS) and elastic body splines (EBS) functions.

Interspecific differences and intraspecific variation

Figure 5 maps mean differences and corresponding Bonferroni-corrected ANOVA significance labels between *melanogaster* and *willistoni* based on TPS and EBS interpolation. For visualization purposes, the interspecific differences have been interpolated using cubic splines based on observed differences at the 74 nodes. As expected, differences are slightly more pronounced in TPS data, but significant nodes obtained from EBS data are not a mere subset of significant TPS nodes, reflecting the fact that both functions differ in more complex ways than the mere degree of spread of the strain imposed by local deformations.

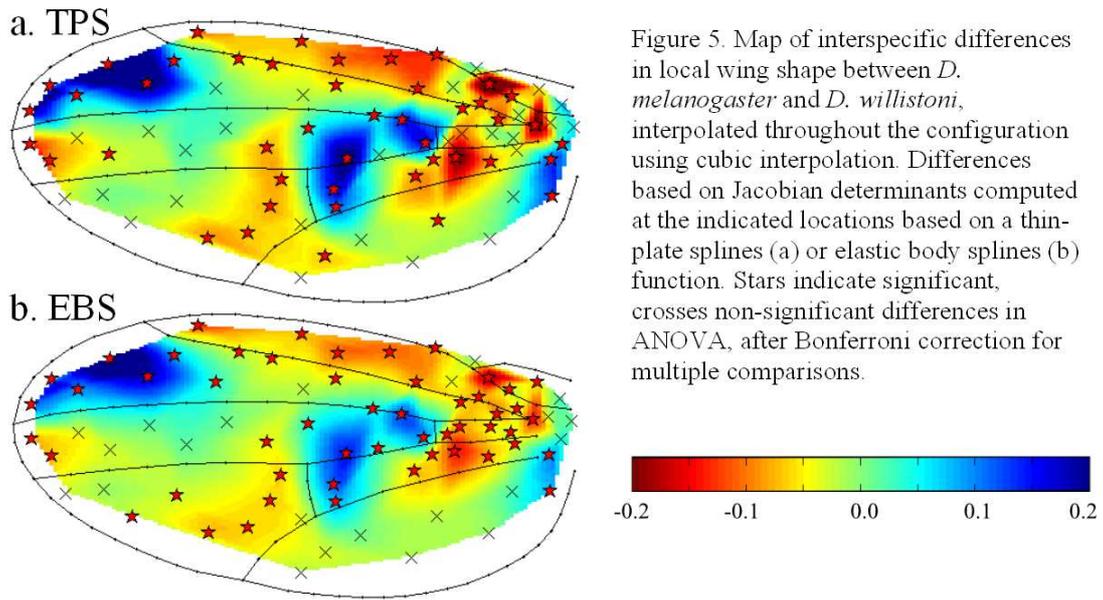
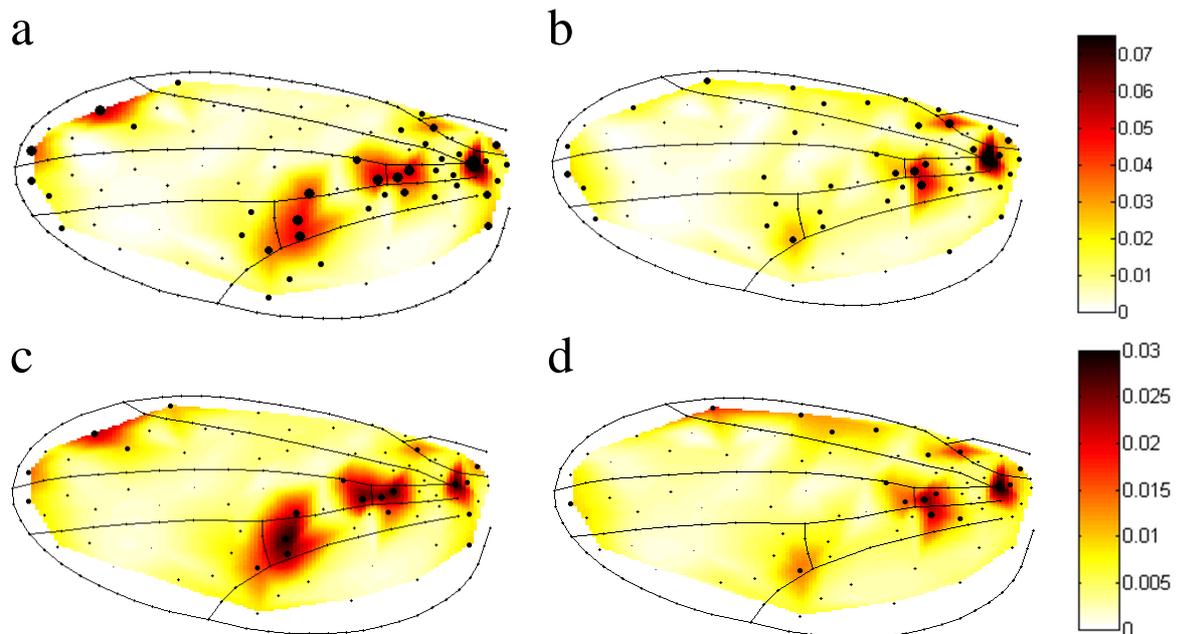


Figure 6 maps intraspecific variances at each node for both species and interpolation functions. Despite the difference in magnitude of variances produced by the two models, variances are spatially auto-correlated between models ($r = 0.94$ for *willistoni*, $r = 0.95$ for *melanogaster*). Plots show both differences and similarities in the spatial distribution of local shape variation between the two species. Interspecific spatial correlations are $r = 0.77$ with TPS data and $r = 0.74$ with EBS data, but a plot of these variances (Fig. 7) reveals the latter to be more evenly distributed and therefore potentially more reliable than those based on TPS, which seem more influenced by extreme values. This is consistent with the higher sensitivity of TPS to highly localized deformations.



splines (a, b) and elastic body splines (c, d) interpolation functions. Variances have been interpolated throughout the wing using a cubic interpolant. Marker size for evaluation nodes is proportional to the variance estimated at each point.

We can investigate the association between interspecific differences and intraspecific variances (Fig. 8). Correlations between squared mean interspecific differences and intraspecific variances are $r = 0.61$ (TPS), 0.52 (EBS) for *melanogaster*, and $r = 0.42$ (TPS), 0.37 (EBS) for *willistoni*. The weaker association in *willistoni* is accounted for by the relatively low variation in the region surrounding the posterior crossvein and the region between L2 and L3 veins, both of which vary strongly in *melanogaster*. Based on these observations, we could explore questions such as whether regions that simultaneously show high intraspecific variance and have co-diverged across species are also highly integrated within species (note that a proper test for such a hypothesis should include a large number of species and corresponding phylogenetic information). A simple way to do this is to select subsets of nodes with high variance and disparity and ask whether they are more strongly correlated mutually than they are to other nodes. Figure 9 draws up such scenarios for both species and models used here, as connectivity graphs with edges linking nodes in the top 75% percentiles for intraspecific variance and interspecific squared mean difference.

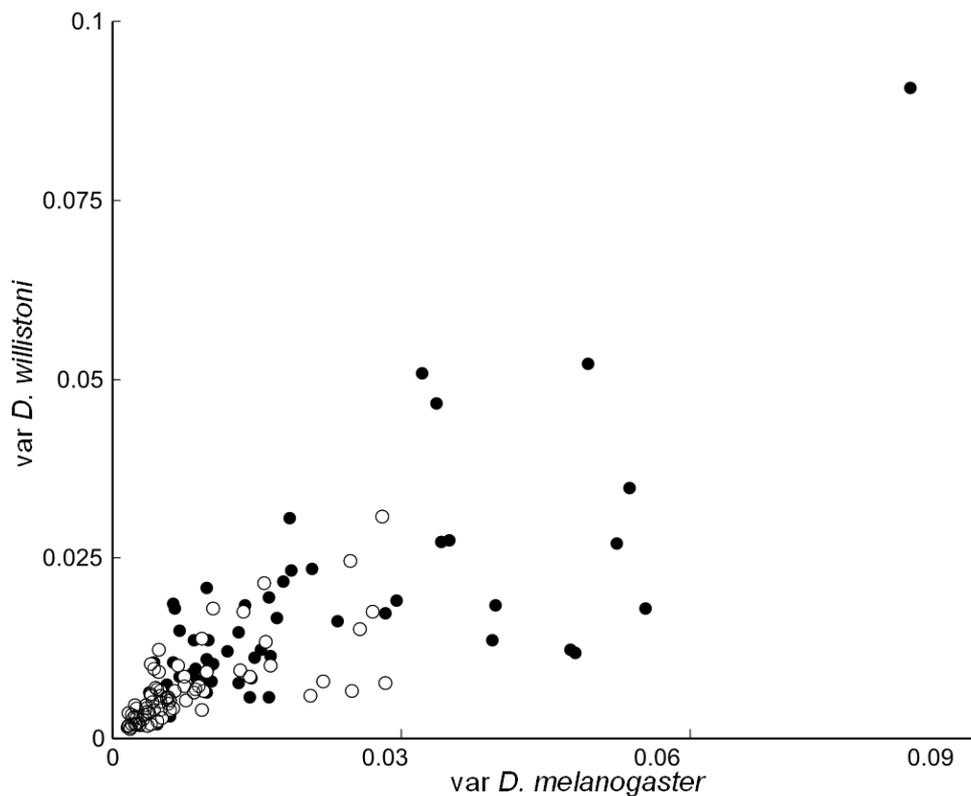


Figure 7. Comparison of variances in local wing shape deformation between two *Drosophila* species. Variances computed from Jacobian determinants derived from thin-plate splines (closed circles) and elastic body splines (open circles) functions.

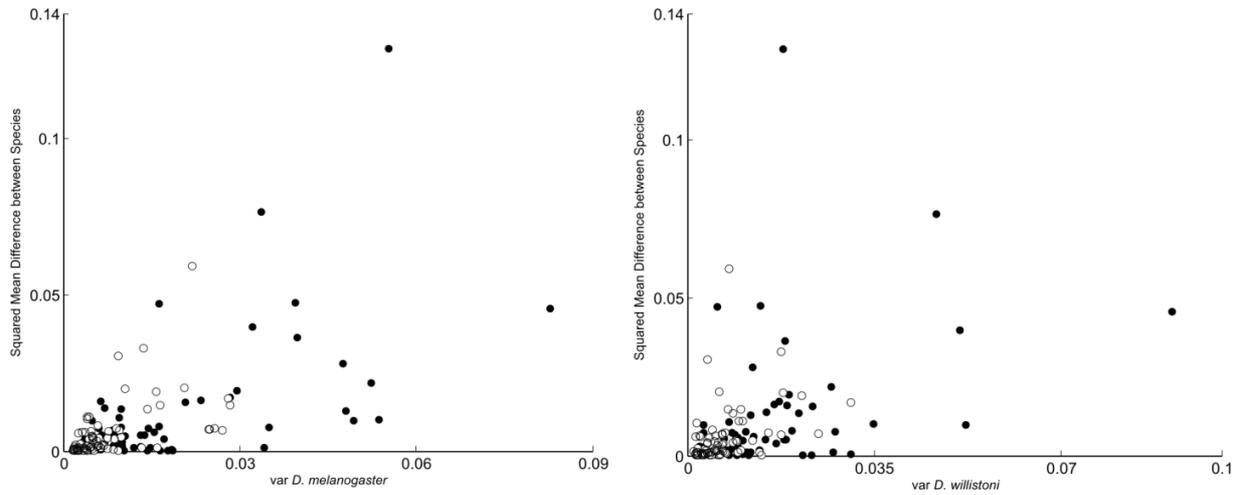


Figure 8. Association between variances and squared mean differences of local wing shape deformations in *D. melanogaster* (left) and *D. willistoni* (right). Deformations computed as Jacobian determinants derived from thin-plate splines (closed circles) and elastic body splines (open circles).

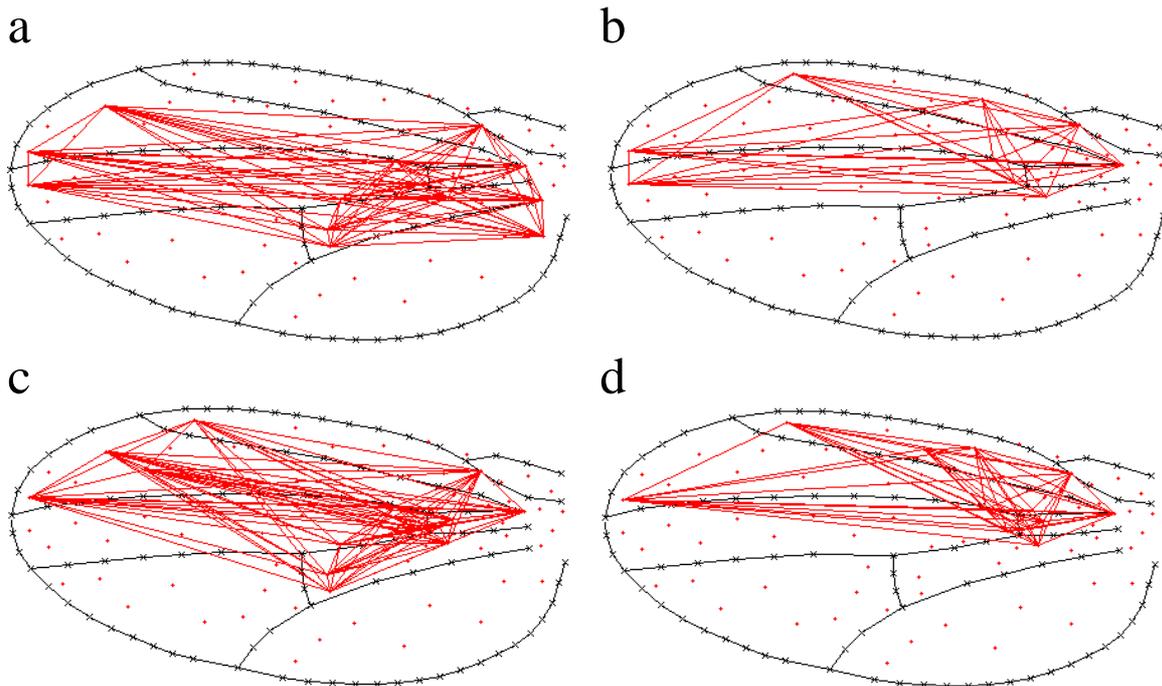


Figure 9. Connectivity graphs depicting a scenario where evaluation nodes with high intraspecific variance and squared mean interspecific differences are mutually integrated more strongly than they are to other nodes. Linked nodes are within the 75% percentile of variances and squared mean differences of local wing shape for *D. melanogaster* (a, c), and *D. willistoni* (b, c), based on Jacobian determinants computed at the indicated nodes using thin-plate splines (a, b) and elastic body splines (c, d) interpolation functions.

One way to test for statistical correlations across the wing that takes advantage of the availability of scalar measurements of local deformation is the analysis of conditional independence among traits in the context

of graph theory. These methods test for conditional independence between pairs of variables after partialling out the effect of every other variable, using a maximum likelihood approach (see Whittaker 1990, Magwene 2001 for details). A deviance statistic is computed for each link representing an association between two traits in a model where all possible links are included, whose significance is tested against a χ^2 distribution with one degree of freedom.

For the *Drosophila* data sets illustrated here, the full model includes 2701 links among 74 nodes, and deviance analyses find 311, 284, 274, and 283 significant links for the *melanogaster* TPS, *melanogaster* EBS, *willistoni* TPS, and *willistoni* EBS data sets, respectively. The minimum partial correlation associated to these links is 0.341 for *melanogaster* and 0.351 for *willistoni*. Restricting visualizations to partial correlations of 0.5 and above to facilitate interpretation, we obtain the graphs depicted in Figure 10. These graphs show some differences among the depicted data sets, but also a predominant common pattern whereby few links are supported between nodes located on distinct compartments across the anterior-posterior boundary (i.e., anterior and posterior to L3 vein), whereas most of the links that cross this boundary are observed at the distal and proximal extremes of the wing. Whereas investigating the biological significance of these results is outside of the scope of this guide, we can at least conclude that there is little support in these data for the hypothesis that co-evolving and variable traits are more integrated than expected by chance alone.

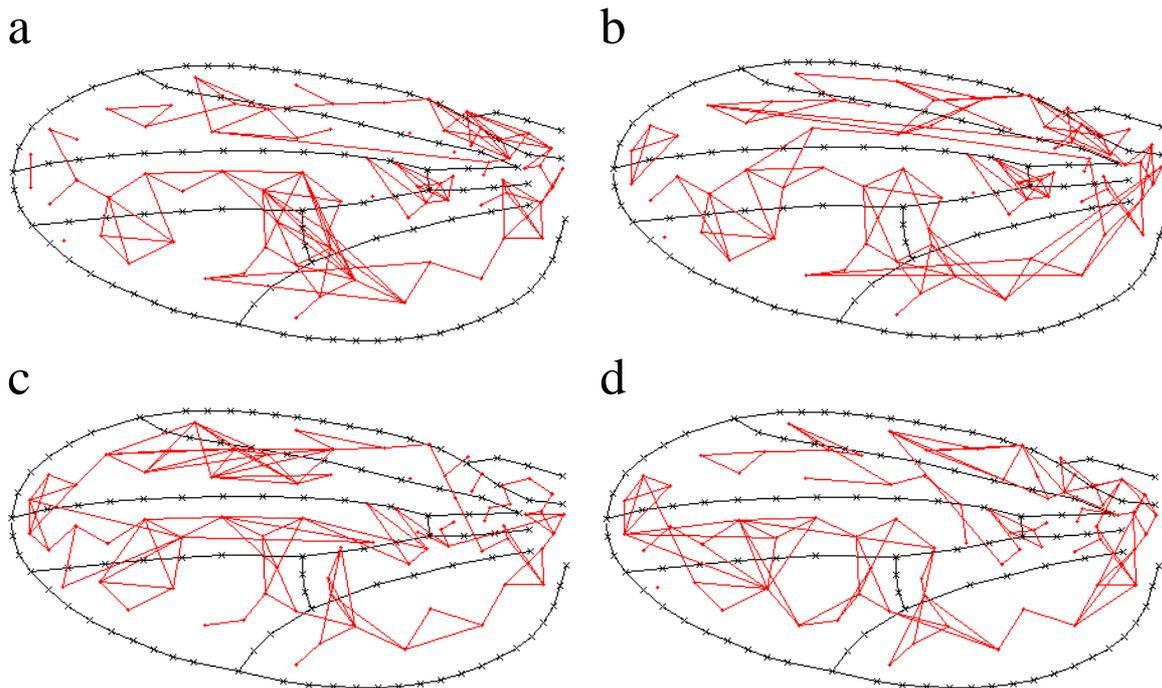


Figure 10. Connectivity graphs depicting links between pairs of nodes with a partial correlation of 0.5 and above. Partial correlations computed for each pair after removing the effect of every other node. According to a maximum likelihood deviance test, all links shown remain highly significant after a Bonferroni correction. Results shown for species *D. melanogaster* (a, c), and *D. willistoni* (b, c), based on Jacobian determinants computed at the indicated nodes using thin-plate splines (a, b) and elastic body splines (c, d) interpolation functions.

Cited works

- Bookstein FL. 1992. Morphometric Tools for Landmark Data: Geometry and Biology. Cambridge: Cambridge Univ Press.
- Davis MH, Khotanzad A, Flamig DP, Harms SE. 1995. Elastic Body Splines: A physics based approach to coordinate transformation in medical image matching. Proc 8th IEEE Symp Comp Med Syst 81-88.
- Davis MH, Khotanzad A, Flamig DP, Harms SE. 1997. A physics-based coordinate transformation for 3-D image matching. IEEE Trans Med Imag 16:317-328.
- Dryden IL, Mardia KV. 1998. Statistical Shape Analysis. Chichester: J. Wiley.
- Magwene PM. 2001. New tools for studying integration and modularity. Evolution 55:1734-1745
- Márquez E, Cabeen R, Woods R, Houle D. 2012. The measurement of local variation in shape. Evolutionary Biology. *In press*.
- Whittaker J. 1990. Graphical models in applied mathematical multivariate statistics. Wiley: New York.
- Woods RP. 2003. Characterizing volume and surface deformations in an atlas framework: Theory, applications, and implementation. Neuroimage 18:769-788.