A Visual Analytics Approach to Study Anatomic Covariation

Max Hermann*
Bonn University, Germany

Anja C. Schunke†
MPI for Evolutionary Biology

Thomas Schultz*
Bonn University, Germany &
MPI for Intelligent Systems

Reinhard Klein*
Bonn University, Germany

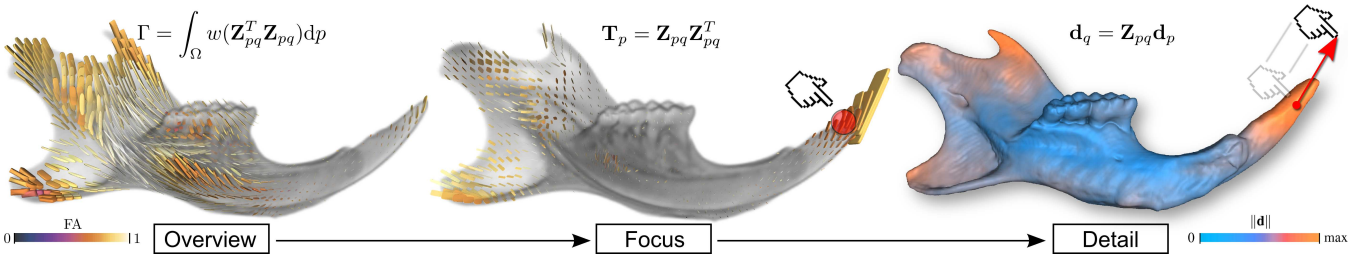

Figure 1: The proposed visual analysis of covariation is performed on three levels of abstraction, demonstrated on a mouse mandible: A static overview (left) provides guidance to candidate points exhibiting non-trivial covariation patterns with the remaining shape. For a particular point p a focus visualization (center) reveals the underlying covariation pattern to p. By interactively dragging around p, details of the correlated shape variation can be investigated in a dynamic animation (right), uncovering also specific directional dependencies of covariation.

Abstract

Gaining insight into anatomic covariation helps the understanding of organismic shape variability in general and is of particular interest for delimiting morphological modules. Generation of hypotheses on structural covariation is undoubtedly a highly creative process, and as such, requires an exploratory approach. In this work we propose a new local anatomic covariance tensor which enables interactive visualizations to explore covariation at different levels of detail, stimulating rapid formation and (qualitative) evaluation of hypotheses. The effectiveness of the presented approach is demonstrated on a μCT dataset of mouse mandibles for which results from the literature are successfully reproduced, while providing a more detailed representation of covariation compared to state-of-the-art methods.

Index Terms: J.3 [Computer Applications]: Life and Medical Sciences—Biology and genetics; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

1 Introduction

Visualization is a core tool in understanding shape variability in organisms from medical imaging data. In this paper we are particularly concerned with shape covariations and novel visual analytics methods to reveal the (complex) interdependency between specific structural parts and the shape as a whole. In biology and anthropology the concept of morphological integration and modularity, particularly of skull and mandible, has gained increasing interest over the last decade (see [21], and references therein). Differences in covariation within and between certain parts of a morphological structure are described, aiming at a better understanding of morphology in general as well as of underlying developmental, functional, or genetic constraints.

At the core of this paper stands a new *local* shape covariance tensor derived from the model-based deformation framework of Blanz et al. [4]. This new tensor summarizes covariation between a specific point and the remainder of the shape and thus generalizes the global shape variance tensor used by Kindlman et al. [17]. A key observation is that for linear shape models (which are probably the most common type) the covariation between a point p and all other points q on the shape is a linear relationship as well, and from that we can define our tensor. Although linear analysis yields efficient visualization algorithms for a single tensor field, there are as many tensor fields as voxels in the dataset. Even for image data as small as 64^3 voxels, manual examination of all possible tensor fields becomes prohibitive. On the other hand, a single tensor field is already a simplified summary of the underlying linear relationship. Therefore additional (visual) navigation strategies are required.

Following the visualization mantra "overview, zoom & filter, then details-on-demand" [32] we provide three different visualizations with decreasing level of abstraction to make the huge amount of tensor fields accessible and enable effective dissemination and detailed investigation of complex shape covariation patterns. The complete pipeline is illustrated in Fig.1. An initial overview visualization highlights strength and directional dependency of covariation associated with certain areas on the structure. Thereby it provides guidance to potentially interesting substructures for further exploration. Each tensor field of a candidate region can in turn be looked at in greater detail in a second tensor glyph visualization. This gives insight into the overall strength and proportions of covariation between a candidate region and global shape, but omits directional dependencies. To examine the latter at the level of individual shape variations, a third visualization provides the possibility to perform model-based deformation interactively in a *click-and-drag* style, inspired by recent works in facial animation.

In combination, the presented three visualizations provide an intuitive and informative interface to visual analysis of shape covariance. We demonstrate our approach on a standard model organism and structure, the rodent mandible, which received considerable attention in research on biological modularity. This allows us to perform a comparison to state-of-the-art methods and a validation of our findings with results from literature. In summary, we consider the following points to constitute the main contributions of this work:

- Introduction of a new *local* shape covariance tensor, especially suited for visual analysis.

*e-mail: {hermann,schultz,klein}@cs.uni-bonn.de
†e-mail: schunke@evolbio.mpg.de

- A hierarchical set of three visualizations ranging from overview to detailed exploration, based on analysis and integration of the local covariance tensor.
- Comparison with state-of-the-art methods performed on a standard model organism, investigated for the first time (to the best of the authors knowledge) at this degree of detail.

2 Related work

Remarkably, the advent of modern shape analysis started off with purely visual methods [34], whose mathematical formalization as statistical shape models many years later was celebrated as a "revolution" in Morphometrics [1]. Nowadays efficient visualization algorithms operating on these models have enabled interactive exploration systems [20, 9, 14]. However, their emphasis is on navigation in shape space and less on exploratory analysis of shape covariation. This section will first review the two shape models targeted at, before relating to other interactive approaches.

2.1 Shape models

This paper works with linear shape models, split here into two major classes, distinguished by their representation of shape.

Point distribution models (PDM) [10, 11] can be considered state-of-the-art in applications in morphology of organisms, particularly Zoology and Anthropology [5], as well as in many medical imaging applications (see [13] and references therein). A PDM represents a shape by a set of points (landmarks) specified on corresponding (homologous) positions on all shapes under investigation. Most of the related interactive systems operate on PDM.

Statistical deformation models (SDM) [2], a major tool in computational anatomy [12], try to include all information gained from medical imaging data. To this end shape variation is described via dense deformation fields, mapping the mean shape to each particular individual of the dataset. Although computationally more expensive, efficient representations based for instance on a B-spline parameterization [30] or carefully performed dimensionality reduction [14] are available. Since a SDM includes more shape details than a PDM, it is an ideal choice for exploratory visualization of shape variability and we will present and demonstrate our methods in this setting. Even though, the underlying approach could as well be applied to a PDM.

2.2 Related interactive approaches

In the last years a rich set of interactive visual analysis tools for landmark data and PDM's has been released [20, 9, 27, 24, 37] and recently, a first system for the exploration of SDM's [14]. All systems feature a principal component analysis (PCA) model for exploration, while some additionally provide support for canonical variates analysis (CVA) and regression tools like partial-least squares (PLS). CVA can be used to analyze covariation with respect to specific groups of individuals in the dataset, providing an orthogonal parameter space similar to PCA for exploration. A PLS regression can be calculated between two pre-selected landmark groups (shape parts) to visually analyze covariation between them in a 2D correlation plot [20, 37]. Once a hypotheses on co-varying groups of individuals or landmarks is formed, CVA and PLS methods provide a valuable tool for further exploration. To generate hypotheses on shape covariation in the first place, only the PCA model remains as an exploration vehicle.

Inspiration for the presented interaction method comes from works in facial animation. Direct manipulation approaches, where a 3D face model is predicted from manipulation of only a small subset of the vertices, are found for blend shapes [25] and part-based models [33]. Although these works are a great source of inspiration and we will internally use the same least-squares minimization [4], the ultimate goals of artistic animation and visual analysis of shape variation are quite contrary.

3 Preliminaries

The general approach to formalize shape difference is based on the notion of transformations, describing mappings of one shape into the other. Decomposing these transformations enables to factor out differences (in shape representation) due to different "motions". After factoring out similarity transformations in our case, eventually a common coordinate frame arises in which the remaining elastic deformations constitute the shape differences of interest. This is the starting point for a statistical deformation model (SDM) which describes a normal distribution of elastic deformations, typically identified with displacement fields, around an average template serving as the mean shape.

3.1 Statistical deformation model

3D images of the same modality serve as input data, showing the shape of an organism in different individuals. Let $\{\mathscr{I}_1,\dots,\mathscr{I}_n\}$ denote the set of 3D input images treated here as scalar functions $\mathscr{I}_i\colon \Omega\to\mathbb{R}$ defined over a common spatial domain $\Omega\subset\mathbb{R}^3$. The deformations between the template image $\hat{\mathscr{I}}$ and each individual are described by coordinate mappings $\phi_i\colon \Omega\to\Omega$ such that $\hat{\mathscr{I}}(x)\approx\mathscr{I}_i(\phi_i(x))$ for $x\in\Omega$. Note that thereby all ϕ_i share the template's coordinate system as their base domain. Each mapping ϕ_i is given as $\phi_i := \mathrm{Id}+\mathbf{d}_i$ with a displacement vector field $\mathbf{d}_i\colon \Omega\to\mathbb{R}^3$ and Id denoting the identity transform. The vector fields $\mathbf{d}_i$ encode the displacement from the template image and can thus be represented as vectors $\mathbf{d}_i\in\mathbb{R}^{3m}$ where m is the number of discrete image samples, i.e. voxels in $\hat{\mathscr{I}}$.

Linear model and synthesis The displacement fields constitute the data matrix $\mathbf{X}=[\mathbf{d}_1,\dots,\mathbf{d}_n]\in\mathbb{R}^{3m\times n}$ for further analysis. Based on this data the SDM establishes a linear model

$$\mathbf{X}=\mathbf{BC} \tag{1}$$

with coefficients $\mathbf{C}\in\mathbb{R}^{n\times n}$ and a basis $\mathbf{B}\in\mathbb{R}^{3m\times n}$. In the navigation process, novel deformations ϕ are synthesized from the SDM via

$$\phi=\mathrm{Id}+\mathbf{d}=\mathrm{Id}+\mathbf{Bc}\,, \tag{2}$$

where $\mathbf{c}=(c_1,\dots,c_n)^T$ should be chosen with $c_i\in[-3\sigma_i,+3\sigma_i]$ conforming to the standard deviation σ_i of the underlying normal distribution model.

Estimating an average template We use generalized Procrustes analysis [11] to estimate a mean shape. The algorithm works iteratively and is initialized with one of the individuals from the dataset, which is refined stepwise until convergence. In each step, all images are registered against the current estimate $\mathscr{I}'$, yielding intermediate deformation fields $\mathbf{d}'_i$. The average deformation $\overline{\mathbf{d}'}=\frac{1}{n}\sum_{i=1}^n\mathbf{d}'_i$ describes the remaining bias towards the initialization, since by definition $\overline{\mathbf{d}'}=0$ for the true mean. To gain a refined estimate which is closer to the true mean, $\mathscr{I}'$ is deformed according to $\overline{\mathbf{d}'}$ to $\mathscr{I}'$. The procedure is repeated until $\overline{\mathbf{d}'}$ is close to zero. After convergence, the final estimate and deformation fields facilitate the template $\hat{\mathscr{I}}$ and data matrix $\mathbf{X}$ for the SDM.

3.2 PCA model

The linear model (1) of an SDM is chosen to have decorrelated and orthogonal basis vectors. Those are the result of a principal component analysis (PCA), which basically diagonalizes the sample covariance matrix. Since we established our data matrix $\mathbf{X}$ with respect to the average template, the first moment of the data is zero and the covariance estimate is $\frac{1}{n-1}\mathbf{XX}^T$. Since in our setting $n\ll 3m$ the rank of the covariance matrix is at most $n'=n-1$. A diagonalization $\mathbf{XX}^T=\mathbf{US}^2\mathbf{U}^T$ is found via singular value decomposition of $\mathbf{X}=\mathbf{USV}^T$, where $\mathbf{U}$ is the set of orthonormal eigenvectors, $\mathbf{S}^2$ a diagonal matrix containing the eigenvalues and $\mathbf{V}$ an

orthogonal matrix. As basis we will use the set of scaled eigenvectors $\mathbf{B} = n'^{-1/2}\mathbf{U}\mathbf{S}$ yielding coefficients $\mathbf{C} = n'^{1/2}\mathbf{S}^{+}\mathbf{U}^T\mathbf{X} = \mathbf{V}^T$, where $\mathbf{S}^{+}$ is a pseudo-inverse, only inverting the non-zero diagonal entries of $\mathbf{S}$.

PCA estimates a normal distribution [15] on the displacement fields modeled as a random variable. By construction, the corresponding coefficients $\mathbf{c}$ follow a centered normal distribution of unit variance and the probability density takes the form

$$p(\mathbf{c}) = (2\pi)^{-n'/2} e^{-\frac{1}{2}\|\mathbf{c}\|^2}.$$

In the following we will make use of the negative log-likelihood of the above

$$-\log p(\mathbf{c}) = \frac{1}{2}\|\mathbf{c}\|^2 + \text{const.} \tag{3}$$

as a measure for penalizing unlikeliness of a displacement field.

3.3 PLS analysis

A state-of-the-art approach to assess shape covariation is to perform a partial least-squares analysis (PLS), pioneered by Tucker [35] and Wold [40] and introduced to morphometrics by Rohlf and Corti [29]. It requires a precise hypothesis about a two-block separation of the shape. By reordering the rows of our original data matrix we can represent the two blocks as $\mathbf{X}^T = \left[\mathbf{X_1}^T \mathbf{X_2}^T\right]$. PLS is technically similar to PCA, but instead of the covariance matrix $\mathbf{C} = \mathbf{X}\mathbf{X}^T$, the cross-covariance matrix $\mathbf{C}_{12} = \mathbf{X_1}\mathbf{X_2}^T$ is diagonalized. Since $\mathbf{C}_{12}$ is no longer symmetric, this yields not an eigenvalue but a singular value decomposition $\mathbf{C}_{12} = \mathbf{L}\mathbf{S}_{12}\mathbf{R}^T$ with different left and right singular vectors. Where the first PCA eigenvector $\mathbf{u}$ maximizes the covariance $|\mathrm{Cov}(\mathbf{X}^T\mathbf{u}, \mathbf{X}^T\mathbf{u})|$, the first pair of singular vectors maximizes the cross-covariance between the two blocks $|\mathrm{Cov}(\mathbf{X_1}^T\mathbf{l}, \mathbf{X_2}^T\mathbf{r})|$. The same holds for further pairs of singular vectors in the respective orthogonal subspaces. Extensions to three and more blocks exist [6], although lacking a closed form solution. By stacking singular vector pairs, full deformation fields for the complete shape can be constructed for visualization purposes. The different size of the blocks has to be taken into account though [23].

So far, PLS analysis provides *the* major visual tool to study shape covariation (e.g. [22] and references in Sec. 6). However, it does not allow inspection of covariation at a finer level than the selected blocks, which have to be selected a-priori, and individual covariation patterns remain superimposed in the singular vectors.

3.4 Global anatomic covariance field

An interesting alternative to PCA visualization to convey the overall variability contained in a dataset is that of anatomic covariance fields presented by Kindlmann et al. [17]. This work uses a glyph-based visualization of the covariance tensor at each point in the dataset. For the set of displacement vectors $\{\mathbf{d}_{p,1}, \ldots, \mathbf{d}_{p,n}\}$ at point $p \in \Omega$ with $d_{p,i} \in \mathbb{R}^3$ the 3×3 covariance tensor is defined as

$$\mathbf{T}_{\text{global}}(p) = \frac{1}{n-1}\sum_{i=1}^{n} \mathbf{d}_{p,i}\mathbf{d}_{p,i}^T \tag{4}$$

We adopt the same visualization for our local covariance tensor (10). To distinguish this global approach from our analysis of local interactions, we will refer to (4) as *global covariation* in the following.

4 Linear analysis of covariation

Our approach builds upon the least-squares framework introduced by Blanz et al. [4] for model-based deformation, which we briefly summarize in the following sub-section, before introducing a local covariance and overview tensor in Sec. 4.2 and 4.3.

4.1 Model-based deformation

For linear shape models Blanz et al. [4] propose a model-based deformation framework which allows the user to produce plausible shapes via displacement of single vertices. The "plausibility" of a shape is defined via its probability according to the Gaussian normal distribution underlying all PCA shape models, see Sec. 3.2.

In a least-squares optimization, the coefficients $\mathbf{c}$ in PCA space of a plausibly deformed shape are estimated where the displaced vertices $\mathbf{d}_p$ (the user edit) are considered as soft-constraints and the optimization is additionally regularized by the shape's probability, penalizing improbable outcomes. Using a squared data term and the negative log-likelihood (3) the optimization can be expressed via the following energy functional

$$E(\mathbf{c}) = \|\mathbf{d}_p - \mathbf{B}_p\mathbf{c}\|_2^2 + \gamma\|\mathbf{c}\|_2^2 \tag{5}$$

where $\mathbf{d}_p$ is the stacked vector of displacements and $\mathbf{B}_p$ is the shape space basis matrix reduced to the rows corresponding to the displaced vertex coordinates in $\mathbf{d}_p$. At the constrained positions p the minimizer of $E(\mathbf{c})$ will try to match the edited shape $\mathbf{d}_p$ while the remaining vertices are deformed to yield a highly probable shape, depending on the choice of the regularization parameter $\gamma \in \mathbb{R}$. In this work we take on the view that the result of the above optimization gives us the shape variation correlated to the edit $\mathbf{d}_p$, at least for a fixed γ. The regularization parameter γ is chosen via cross-validation in a leave-one-out fashion benchmarked on the average reconstruction error. In our applications the resulting error function is found to be smooth and in particular stable around the optimum allowing for a robust choice of γ.

4.2 A linear operator for point-wise covariation

The same model-based deformation framework (5) can directly be applied to statistical deformation models, where instead of vertices now displacement vectors are constrained. Restricting to the special case of a single constrained position p, we now derive a linear operator for covariation between the point p and any other point $q \in \Omega$. The displacement predicted from (5) at q to a variation at p is what we interpret in this paper as the correlated shape change, or *interaction*, between p and q.

Once optimal coefficients $\mathbf{c}_{\text{opt}}$ minimizing Eq. (5) are found, the corresponding displacement $\mathbf{d}_q$ at q co-varying with the change $\mathbf{d}_p$ at p is given by

$$\mathbf{d}_q = \mathbf{B}_q\mathbf{c}_{\text{opt}} \quad \text{where} \quad \mathbf{c}_{\text{opt}} := \arg_{\mathbf{c}} \min E(\mathbf{c}). \tag{6}$$

To solve for $\mathbf{c}_{\text{opt}}$ consider the fact that a minimum of $E(\mathbf{c})$ must have vanishing gradient

$$\frac{\partial}{\partial \mathbf{c}} E = \mathbf{B}_p^T\mathbf{B}_p\mathbf{c} - \mathbf{B}_p^T\mathbf{d}_p + \gamma\mathbf{c} \stackrel{!}{=} 0$$

yielding a linear system

$$(\mathbf{B}_p^T\mathbf{B}_p + \gamma\mathbf{I}) = \mathbf{B}_p^T\mathbf{d}_p$$

of the form $\mathbf{A}\mathbf{c} = \text{const.}$ which can be solved by direct inversion since the $n \times n$ system matrix $\mathbf{A}$ is symmetric and

$$\mathbf{c}_{\text{opt}} = \mathbf{A}^{-1}\mathbf{B}_p^T\mathbf{d}_p. \tag{7}$$

Expanding (7) into (6) we arrive at

$$\mathbf{d}_q = \mathbf{B}_q(\mathbf{B}_p^T\mathbf{B}_p + \gamma\mathbf{I})^{-1}\mathbf{B}_p^T\mathbf{d}_p = \mathbf{Z}_{pq}\mathbf{d}_p \tag{8}$$

where the relationship between displacements at p and q is given by the 3×3 matrix $\mathbf{Z}_{pq}$. Note that $\mathbf{Z}_{pp} \neq \mathbf{I}$ in general because the regularization term penalizes improbable edits. It can also happen that the edit is not contained in the span of $\mathbf{B}$, meaning it has zero probability. One should further keep in mind that the relationship is not symmetric and $\mathbf{Z}_{pq} \neq \mathbf{Z}_{qp}^{-1}$ in general.

4.3 A new local shape covariance tensor

For a given displacement at point p in the shape, Eq. (8) predicts the most probable corresponding displacement at q. To visualize this relationship encoded in $\mathbf{Z}_{pq}$ we resort to statistical covariance analysis. There, covariance structure of a random vector $\mathbf{x}$ is defined as $\Sigma(\mathbf{x}) = \mathbb{E}\left\{(\mathbf{x}-\mathbb{E}(\mathbf{x}))(\mathbf{x}-\mathbb{E}(\mathbf{x}))^T\right\}$ with expectation $\mathbb{E}$. Assuming that input displacements $\mathbf{d}_p$ are random vectors drawn from a distribution with covariance $\Sigma(\mathbf{d}_p)$, the covariance structure at q follows directly from the linearity of expectation:

$$\Sigma(\mathbf{d}_q) = \Sigma(\mathbf{Z}_{pq}\mathbf{d}_p) = \mathbf{Z}_{pq}\Sigma(\mathbf{d}_p)\mathbf{Z}_{pq}^T \tag{9}$$

For our directed interaction we want an unbiased estimate of $\Sigma(\mathbf{d}_q)$ and assume thus a prior of isotropic covariance at p. By setting therefore $\Sigma(\mathbf{d}_p) = \mathbf{I}$ we arrive at the following simple definition of a local covariance tensor at q for a fixed p:

$$\mathbf{T}_p(q) = \mathbf{Z}_{pq}\mathbf{Z}_{pq}^T \tag{10}$$

From the properties of the linear operator, i.e. $\mathbf{Z}_{pp} \neq \mathbf{I}$ and $\mathbf{Z}_{pq} \neq \mathbf{Z}_{qp}^{-1}$, it follows that $\mathbf{T}_p(q)$ describes a one-sided, directed interaction from p towards q and in general a posterior covariance $\mathbf{T}_p(p) \neq \mathbf{I}$ at p.

Efficient local tensor sampling To speed up the computation of local tensor fields we split $\mathbf{Z}_{pq}$ into two factors, depending each solely on p and q, respectively:

$$\mathbf{Z}_{pq} = \mathbf{B}_q\mathbf{Z}_p \quad \text{where} \quad \mathbf{Z}_p = (\mathbf{B}_p^T\mathbf{B}_p + \gamma\mathbf{I})^{-1}\mathbf{B}_p^T$$

Thereby the matrix inversion in $\mathbf{Z}_p$ has only to be done once on positioning the probe, while the associated tensor field (10) can be sampled at the cost of a matrix multiplication per sample.

4.4 Measuring directional interaction strength

Since a covariance tensor $\mathbf{T}$ is symmetric it can be diagonalized as $\mathbf{T} = \mathbf{R}\Lambda^2\mathbf{R}^T$ into a rotation $\mathbf{R}$ from the unit to an eigenvector basis and a diagonal matrix $\Lambda^2 = \mathrm{diag}(\lambda_1,\lambda_2,\lambda_3)$ of corresponding eigenvalues, sorted in descending order $\lambda_1 \geq \lambda_2 \geq \lambda_3$. The ratios between the eigenvalues delineate the anisotropy of the local shape variation [38].

For an overview visualization we want to assess which directions of displacement $\mathbf{d}_p$ at a point p will result in large deformations on the remaining shape, describing potentially interesting covariation. The strength of interaction inflicted by $\mathbf{d}_p$ can be measured by averaging the squared magnitude of the displacement responses $\mathbf{d}_q = \mathbf{Z}_{pq}\mathbf{d}_p$ via

$$\eta_p(\mathbf{d}_p) = \frac{1}{|\Omega|}\sum_{q\in\Omega}\|\mathbf{Z}_{pq}\mathbf{d}_p\|^2 = \frac{1}{|\Omega|}\sum_{q\in\Omega}\mathbf{d}_p^T\mathbf{Z}_{pq}^T\mathbf{Z}_{pq}\mathbf{d}_p.$$

The average of the occurring quadratic forms $\mathbf{Z}_{pq}^T\mathbf{Z}_{pq}$,

$$\Gamma'(p) = \frac{1}{|\Omega|}\sum_{q\in\Omega}\mathbf{Z}_{pq}^T\mathbf{Z}_{pq}$$

is already a tensor representation of interaction strength. We want to stress that this is not an average of the local covariance tensor but $\mathbf{Z}_{pq}^T\mathbf{Z}_{pq}$. The spectrum of the average Γ' informs on which editing directions $\mathbf{d}_p$ impact the shape stronger (large λ_i) and weaker (small λ_i). In the derivation, η_p treats all edit directions as equally important, irregardless of their probability due to the shape model. Since improbable directions are of lesser interest for analysis of covariation we weight Γ' with the global covariance tensor. The weighted variant takes the form

$$\Gamma(p) = \mathbf{T}_{\text{global}}^T(p)\,\Gamma'(p)\,\mathbf{T}_{\text{global}}(p) \tag{11}$$

and is the tensor used for overview visualization shown in Fig. 2. To exaggerate tensors with smaller trace we take the square-root of eigenvalues for scaling the glyphs.

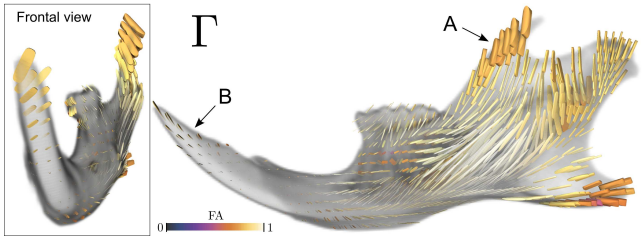

Figure 2: Overview visualization of *Mus* dataset. Glyph size indicates strength of the interaction pattern $\mathbf{T}_p$ associated with a particular point (compare Fig. 8 to point B). Principal axes inform on edit directions $\mathbf{d}_p$ with strong responses (compare Fig. 6 to point A).

5 Visualization Methods

5.1 Direct deformation visualization

Visualizing a synthesized deformation (2) requires that the corresponding mapping ϕ is applied to the template $\hat{\mathscr{I}}$, producing a deformed shape. Computing the actual deformed volume is computationally expensive, thus we resort to direct space warping techniques [7], applied on-line during rendering. The basic idea, due to Barr [3], is that casting straight rays into the scene will show the undeformed object, while deforming the rays with a mapping ϕ^{-1} will result in an image of the object, as if it was deformed with ϕ. The same idea is realized in our GPU raycaster.

In raycasting, a ray is traversed through positions x in the template's coordinate system. To yield a deformed image $\mathscr{I}$ from the template $\hat{\mathscr{I}}$, the ray is displaced according to $\mathscr{I}(x) = \hat{\mathscr{I}}(\phi^{-1}(x))$. The inverse mapping is approximated here by negating the original displacement, $\tilde{\phi}^{-1} = \mathrm{Id} - \mathbf{d}$. For a maximum displacement magnitude of γ the approximation error amounts to $\tilde{\phi}^{-1}(\phi(x)) = \mathrm{Id} + o(\gamma^2)$ [1] which we found sufficiently accurate for our homogeneous dataset exhibiting mostly small-scale and very smooth displacements. The basis functions $\mathbf{B}$ used in the computation of ϕ are kept in GPU memory and the linear combination (2) is evaluated in a shader, exploiting trilinear interpolation capability of modern graphics hardware.

Our raycaster provides direct volume rendering (DVR) and indirect isosurface rendering. While DVR makes interior structures accessible (see video), it requires additional techniques to be employed for color-coding due to color-mixing along a cast ray. For simplicity we encode additional information like displacement magnitude only in the isosurface rendering. Picking a point on the shape for model-based editing is also realized in the raycaster by selecting the first hit point on an arbitrary, user-selected isosurface. Our implementation achieves on average 11 fps during editing on a Intel Core2 Q6600 CPU at 2.4GHz equipped with a Nvidia GTX 460 graphics card.

5.2 Tensor field visualization

For the visualization of covariance tensor fields we use the same glyph-based technique as Kindlmann et al. [17]. In a glyph based approach, an effective visualization is achieved by scaling and rotating a geometrical primitive, typically a sphere, according to $\mathbf{R}$ and Λ. Instead of a sphere, we follow Kindlmann et al. and use a superquadric glyph [16] because of its qualities for visual disambiguation of linear, planar and spherical shaped tensors, relevant to our visual analysis.

[1] *Proof:* By definition $f := \tilde{\phi}^{-1}(\phi(x)) = \phi(x) - d(\phi(x)) = x + d(x) - d(x+d(x))$. Taylor expansion of the last term yields $f = x - \nabla d(x)d(x) + o(|d(x)|^2$ where $\nabla d(x)$ denotes the Jacobi matrix. Rewriting $d(x) = \gamma g(x)$ with $|\nabla g(x)g(x)| \leq 1$ gives an upper bound $|\nabla d(x)d(x)| \in o(\gamma^2)$. Utilizing $|d(x)| \leq \gamma$ we finally arrive at $f = x + o(\gamma^2)$. □

The geometric encoding of tensor properties is accompanied by color-coding the glyph according to local anisotropy. In deformation models the anisotropy informs on the amount of directional dependency of shape variation and is particularly helpful for our local analysis in understanding the exact interaction between two points. A suitable measure in this context is fractional anisotropy, defined as $\mathrm{FA} := \left(\frac{3}{2}\frac{(\lambda_1-\mu)^2+(\lambda_2-\mu)^2+(\lambda_3-\mu)^2}{(\lambda_1+\lambda_2+\lambda_3)^2}\right)^{1/2}$, which interpolates between the spherical case ($\lambda_1=\lambda_2=\lambda_3$ and $FA=0$) and the linear one ($\lambda_1 \gg \lambda_2 \approx \lambda_3 = 0$ and $FA=1$), irregardless of tensor norm $\|\mathbf{T}\|_F$. The final visualization is produced by sampling the tensor field on a regular grid and placing a corresponding superquadric glyph with color mapped FA at each sample position.

For the local covariance visualization the tensor field $\mathbf{T}_p$ serves as input, for which a probe is positioned in 3D at a point p of interest in advance. The overview visualization is generated from the precomputed tensor field Γ. For all visualizations and computations the image domain Ω is thresholded to the Hounsfield range of bone structures. Additional methods like glyph packing [18] and halos [31] could be applied to enhance visual disambiguation of the often observed highly anisotropic patterns.

6 Visual Analysis of the Mouse Mandible

The mouse mandible is a standard model for morphometric analyses, including the study of morphological integration and modularity, partly because of its relatively simple structure [22, 41, 8, 28]. Irregardless of this simplicity, finding modules and modeling their interaction is considered a difficult task. In the majority of studies so far, hypotheses of the existence and position of modules have been formulated a priori and tested subsequently. PLS analysis is the primary method to investigate correlation and provide visualizations of shape covariation.

Structure of the mouse mandible Each separate half of the mandible consists of a single bone with three processes in the back, a row of molar teeth and the single incisor, whose posterior end lies far back in the mandible, below or behind the coronoid process, see Fig. 3. A common functional segmentation is the separation into two parts, the frontal region, bearing incisor and molar teeth, and the rear processes with muscle attachments. However it is speculated that there could be an additional set of modules at a smaller scale [22] and in recent work several finer scale subdivisions [41, 8, 28] are addressed.

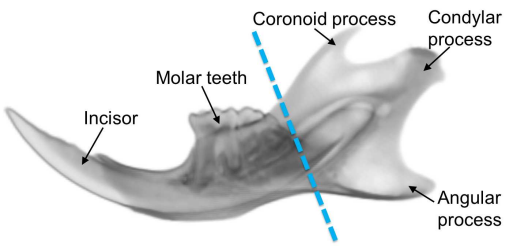

Figure 3: Anatomical parts of the mouse mandible referred to in this work. A common subdivision [22] into 2 functional subunits is shown.

6.1 Dataset and preprocessing

The analysis is conducted on left mandibles of 30 house mice (*Mus musculus*) whose skull μCT scans are reduced to an image resolution of $200 \times 200 \times 400$ voxels and semi-automatically segmented into upper skull and mandibles [39]. Damaged left mandibles were replaced with their mirrored right counterpart. Differences due to translation, rotation and scale are factored out with an image based similarity alignment. The alignment is optimized based on L2 intensity error on histogram equalized images using the elastix toolbox [19]. Elastic registration is performed with the symmetric log-domain diffeomorphic demons algorithm [36]. Bootstrapping the template is started with an arbitrary reference individual. Because of the homogenity of the dataset at hand, 3 iterations of GPA were sufficient to reach convergence. In the PCA model 93% of shape variability is captured by the first 5 components, see Fig. 4. For synthesis and model-based editing the PCA model was reduced to these first five modes, while all tensor computations were performed on the full model.

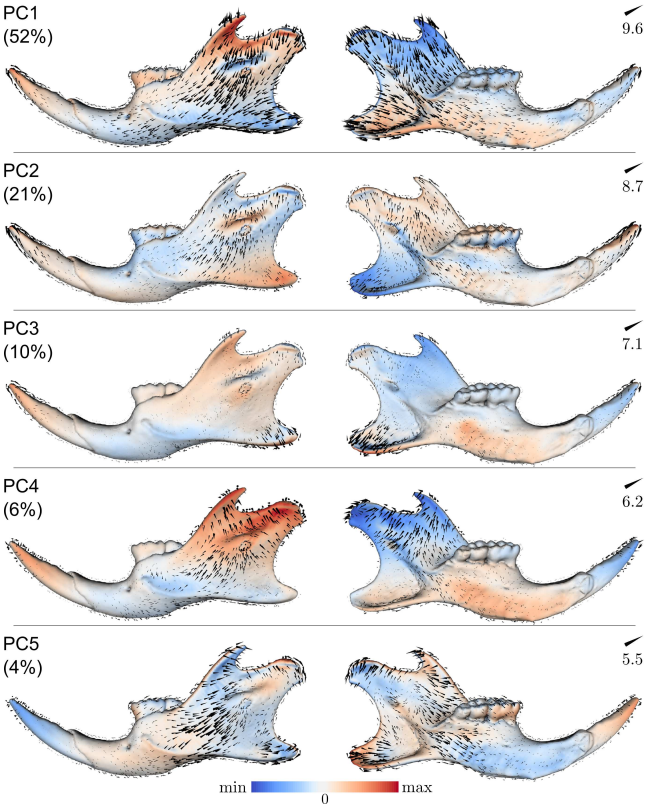

Figure 4: PCA eigenmodes of *Mus* dataset capturing 93% of the total variance. Shown are the vectorfields on a representative iso-surface decomposed into a surface orthogonal and tangential part, visualized color-coded and as vector glyphs respectively with glyphs of max. magnitude scaled to same length. (Visualization based on [42].)

6.2 Global analysis and PLS

For later comparison, we will first describe results achieved with previous methods of related work in morphometrics, applied to our dataset. Particularly PCA [41, 26], PLS [8] and a combination of both [28, 22] was used. Additionally we will apply the global covariance tensor [17].

Global methods like PCA give us an impression of the overall variability contained in the dataset. Fig. 4 shows the relevant PCA eigenmodes. We observe that processes and incisor are each influenced by several modes, so the PCA does not reveal a clear separation into modules. A more concise overview of shape variation is provided by the anatomic covariance field in Fig. 5. It can be read from the glyph pattern that the three processes have different principal directions of variation. At the posterior part of the incisor, inside the mandible, a strongly anisotropic region is visible following nicely the incisor curvature. Note that the global covariance

field shows pointwise variation of the dataset but does *not* convey information on covariance between different points on the shape.

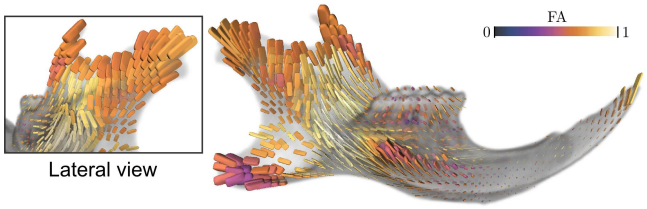

Figure 5: Global anatomic covariance field [17] for *Mus* dataset. A lateral view is shown in the inset for comparison with Fig. 10.

Results of a PLS analysis are visualized in Fig. 11. The visible similarity in PLS1 of the rear processes to (minus) PC1 is also encountered in landmark analyses [22] and can probably be attributed to the fact that the processes constitute a great proportion of global variation. The PLS2 modes specifically indicate an interaction between coronoid and condylar process and between incisor and angular process. Modes beyond PLS1 are harder to interpret because they only represent covariation orthogonal to previous mode pairs.

6.3 Results

In this section we validate our approach by reproducing hypotheses on module delimitations and interactions from literature [22, 41, 8, 28]. Additionally new observations of our analysis are described, pointing to a finer-scaled hypothesis on module segmentation. Our analysis follows the pipeline illustrated in Fig. 1, going back-and-forth between the different views.

Starting from the overview shown in Fig. 2 one can identify several candidate regions. Among them, in agreement with literature [22], the rear processes and tip of the incisor (visible in frontal view) exhibit strong impact on covariation, as indicated by glyph sizes. Additionally the posterior end of the incisor inside the mandible also shows up prominently. The latter could not be observed in previous approaches based on landmark data and focusing on the outline of mandible shape; therefore it is excluded in the following as it can not serve for validation.

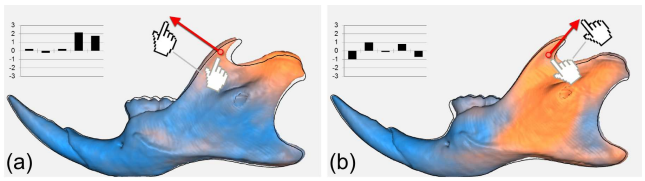

Figure 6: Model-based editing on *Mus* dataset at coronoid process. Color coded from cool to warm (0 max) is the displacement magnitude from the template, whose silhouette is overlaid. Inset bar plots describe the first 5 PC coefficients in units of standard deviations. Observe that varying the coronoid position is strongly associated with the condylar process (a), while changing its length impacts the rear of the mandible at a larger scale (b).

A main advantage of our approach is that one can assess covariation at the level of individual displacements. This enables to distinguish directional dependencies in covariation as for instance found at the coronoid process. The overview suggests at least two different covariation patterns: Principal direction of the corresponding glyph at (A) in Fig.2 predicts a stronger response on variation of length (vertical) versus position (horizontal) of the process. The specific shape variations are quite different as an investigation in detail view shows, illustrated in Fig. 6. Editing the coronoid process also reveals a coupled interaction to the adjacent condylar process. The existence of an interaction between these processes was known already [41] and is confirmed by PLS analysis, but the specific pattern of covariation visualized in focus view in Figs. 7 and 10 provides a new level of detail not seen previously.

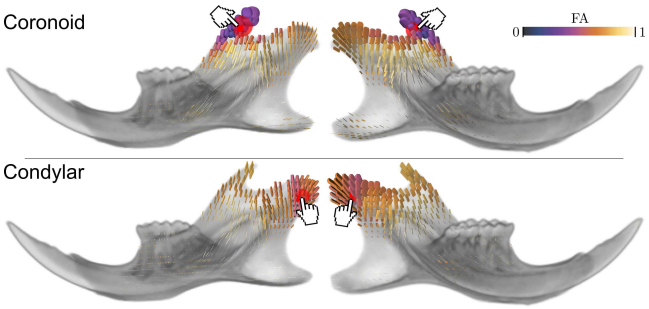

Figure 7: Local covariance fields for probes at coronoid and condylar processes, probing points highlighted in red. A strong interaction in-between the two processes becomes apparent and shows the diminishing impact on angular process and incisor. This finding is in agreement with PLS analysis and results of Zelditch et al. [41].

Looking further at the condylar process in the overview suggests a separation between a tip part and a proximal part, closer to the mandible center. While the pattern near the tip is quite homogeneous, the proximal one is more diverse. In fact, editing the condylar process shows different reactions depending on the exact position of the probe on the process. This suggests a more fine-scaled analysis, proposed also in recent work [8, 28], with the here presented methods.

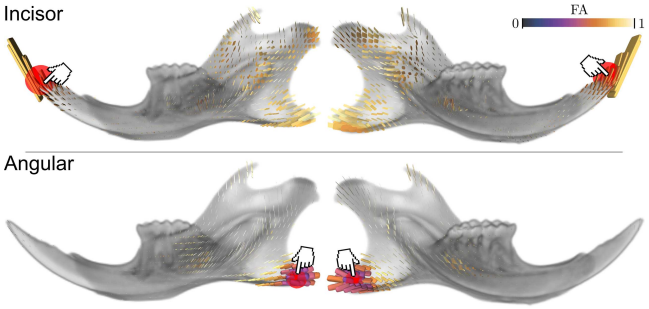

Figure 8: Local covariance fields for probes at incisor and angular process, probing points highlighted in red. Note that the incisor tensor field is scaled by a factor of 2.5 compared to all other fields, attributed to the smaller overall interaction strength (see also Fig. 2). While the incisor probe produces a global covariation pattern, the covariation response of the angular process is locally concentrated.

Examining incisor and angular process in focus view provides good examples of two extremes of global and local patterns, visible in Fig. 8. The angular probe is locally concentrated but still exhibits highly directed ($\lambda_1 \gg \lambda_2 \approx \lambda_3 \approx 0$) interaction with the upper two processes. The strength of interactions for incisor and angular probes again reproduce PLS findings. Since the complete incisor structure spans throughout the mandible, an effect on the rear processes is expected on variation in the incisor tip. Editing the tip of the incisor as shown in Fig. 9 reveals further that the associated variation can be decomposed into two parts, depending roughly on the principal axes of the local tensor near the incisor tip. Pulling in directions of the principal axis (a),(b) inflicts a deformation at the upper coronoid and the lower angular process, while pulling orthogonal (c) to the axis has associated variation at the central condylar process and the rear part of the mandible. This is another example

of directional dependencies, which can not be explored with previous methods.

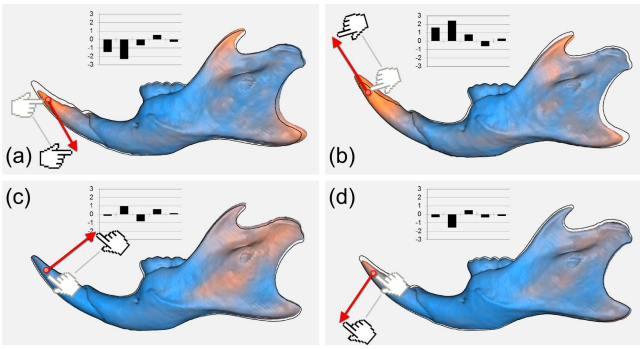

Figure 9: Model-based editing on *Mus* dataset at incisor, same legend as in Fig. 6. Pulling the incisor in directions (a) and (b) leads to the expected elongation / foreshortening, seemingly correlated with the shape of the tips of the rear processes. Dragging the incisor orthogonally (c)-(d) keeps the front nearly rigid showing minor interaction in the posterior part.

The anisotropy pattern in the focus view in Fig. 10 also nicely illustrates some properties of our local tensor, setting it apart from the global covariance field. Anisotropy near the probe is expected to be similar to that of the probe, indicating strong local covariation (typical for stiff bone structures). That anisotropy is higher in the process opposite of the probe indicates a structured covariation. Note that these patterns are distinct from the global anatomic covariation field shown in the inset in Fig. 5. From this we conclude that our system allows one to focus on particular components of the global pattern attributed to specific inter-part interactions.

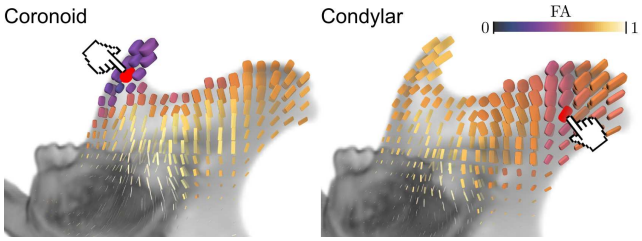

Figure 10: Close-ups of local covariance fields for coronoid and condylar processes show their consistent interaction. While the local region close to each probe (highlighted in red) exhibits more isotropic covariation, the reaction pattern at the other process is more directed. This reveals which part of the global covariance field can be attributed to this particular interaction between the two processes.

7 Summary and Conclusion

Starting from an existing method for shape animation and reconstruction, we developed a complete visual analysis pipeline for effective exploration of shape covariation in statistical deformation models. To this end a new local covariance tensor was derived and integrated, yielding two novel visualizations accompanying the model-based deformation of Blanz et al. [4]. Together, all three methods complement each other in a visual analysis of interactions between different parts of the shape at varying degrees of detail.

Exploratory analysis on a scientific dataset was carried out on the backdrop of morphological integration, illustrating the utility of the visual analysis in the search for module boundaries. Comparison to state-of-the-art techniques (PLS, PCA and global anatomic covariance tensor) showed that the presented approach enables an exploration at a finer level of detail including for the first time directional dependencies. Reproducing several results from recent works in morphometry can be taken as sanity check here. Further applications in morphometric studies are expected a) on datasets with higher diversity consisting of different species and b) on more complex structures like the rodent skull. It would also be interesting to see how effective the approach is in uncovering other, non-anatomic sources of shape variation like (structured) registration or reconstruction errors. Realtime sampling and interactive filtering of local covariance fields is subject to future work.

Acknowledgements

This work was supported in part by Deutsche Forschungsgesellschaft within the priority program SPP1335. The authors thank Rainer Hutterer, Zoologisches Forschungsmuseum Alexander Koenig, Bonn, for loan of specimens and Vitalis Wiens for preprocessing the μCT datasets.

References

[1] D. C. Adams, F. J. Rohlf, and D. E. Slice. Geometric morphometrics: Ten years of progress following the revolution. *Italian Journal of Zoology*, 71(1):5–16, 2004.

[2] J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston. Identifying global anatomical differences: Deformation-based morphometry. *Human Brain Mapping*, 6(5-6):348–357, 1998.

[3] A. H. Barr. Global and local deformations of solid primitives. In *ACM Siggraph Computer Graphics*, volume 18, pages 21–30. ACM, 1984.

[4] V. Blanz, A. Mehl, T. Vetter, and H.-P. Seidel. A statistical method for robust 3d surface reconstruction from sparse data. *Int. Symp. 3D Data Processing Visualization and Transmission*, pages 293–300, 2004.

[5] F. L. Bookstein. *Morphometric tools for landmark data: Geometry and Biology*. Cambridge University Press, 1991.

[6] F. L. Bookstein, P. Gunz, P. Mitterœcker, H. Prossinger, K. Schæfer, and H. Seidler. Cranial integration in *Homo*: singular warps analysis of the midsagittal plane in ontogeny and evolution. *Journal of Human Evolution*, 44(2):167–187, 2003.

[7] T. Brunet, K. E. Nowak, and M. Gleicher. Integrating dynamic deformations into interactive volume visualization. *Eurographics/IEEE VGTC Symp. on Visualization (EUROVIS)*, pages 219–226, 2006.

[8] G. Burgio, M. Baylac, E. Heyer, and X. Montagutelli. Exploration of the genetic organization of morphological modularity on the mouse mandible using a set of interspecific recombinant congenic strains between c57bl/6 and mice of the *Mus spretus* species. *G3: Genes—Genomes— Genetics*, 2(10):1257–1268, 2012.

[9] S. Busking, C. P. Botha, and F. H. Post. Dynamic Multi-View Exploration of Shape Spaces. *Computer Graphics Forum*, 29(3):973–982, 2010.

[10] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models: Their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan. 1995.

[11] I. Dryden and K. Mardia. *Statistical Shape Analysis*. Wiley, 1997.

[12] U. Grenander and M. I. Miller. Computational anatomy: an emerging discipline. *Q. Appl. Math.*, LVI(4):617–694, Dec. 1998.

[13] T. Heimann and H.-P. Meinzer. Statistical shape models for 3d medical image segmentation: A review. *Medical Image Analysis*, 13(4):543 – 563, 2009.

[14] M. Hermann, A. C. Schunke, and R. Klein. Semantically steered visual analysis of highly detailed morphometric shape spaces. In *1st IEEE Symposium on biological data visualization (BioVis)*, pages 151–158, 2011.

[15] I. Jolliffe. *Principal component analysis*. Wiley Online Library, 2005.

[16] G. Kindlmann. Superquadric tensor glyphs. In *Proceedings EG/IEEE Symposium on Visualization (VisSym'04)*, pages 147–154, 2004.

[17] G. Kindlmann, D. Weinstein, A. Lee, A. Toga, and P. Thompson. Visualization of anatomic covariance tensor fields. *Engineering in Medicine and Biology Society (IEMBS '04).*, 1:1842–1845, 2004.

[18] G. Kindlmann and C.-F. Westin. Diffusion tensor visualization with glyph packing. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1329–1336, 2006.

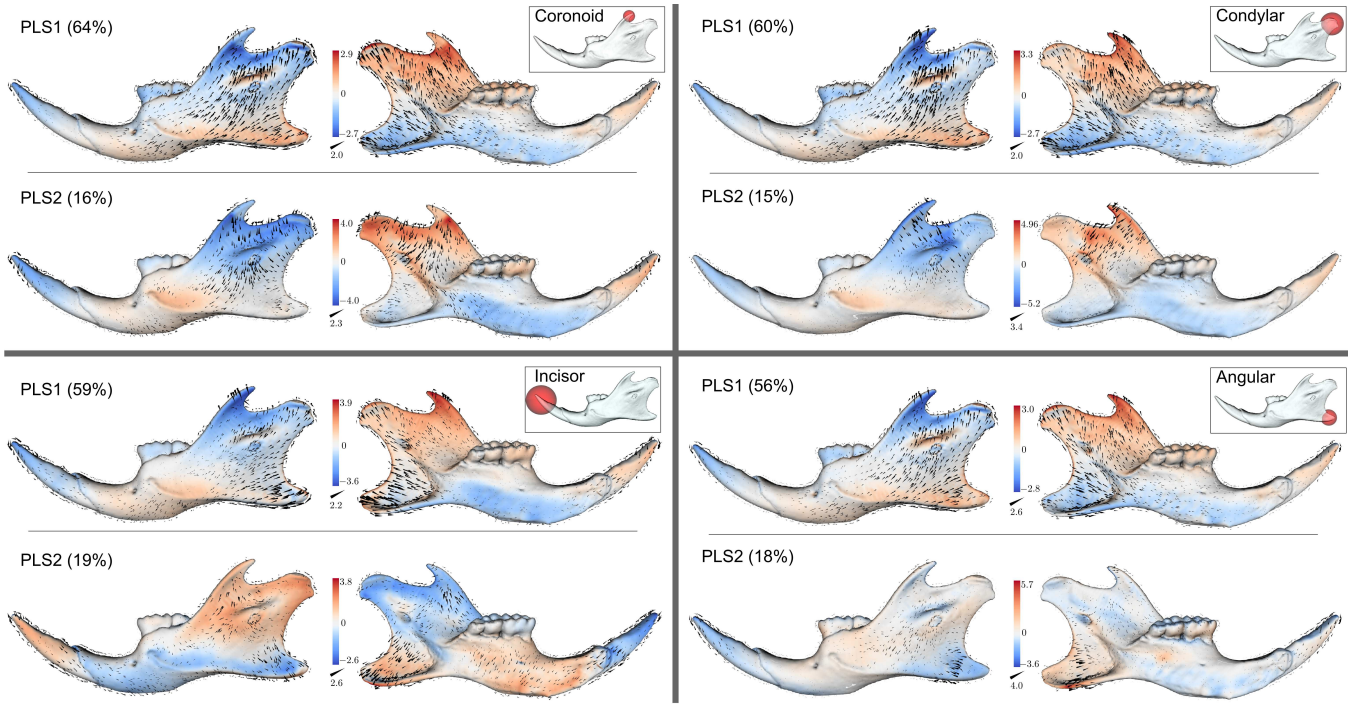

Figure 11: PLS analysis on *Mus* dataset with four different segmentations into white and red blocks as indicated in the insets, visualization design as in Fig. 4. Note the similarity of PLS1 for the rear processes which resembles the first mode PC1 (up to sign) of the global analysis. The variation patterns of PLS2 shows strong covariation between coronoid and condylar processes (top row) while the incisor affects more the angular process (bottom row).

[19] S. Klein, M. Staring, K. Murphy, M. A. Viergever, J. P. Pluim, et al. Elastix: a toolbox for intensity-based medical image registration. *IEEE transactions on medical imaging*, 29(1):196–205, 2010.

[20] C. P. Klingenberg. MorphoJ: an integrated software package for geometric morphometrics. *Molecular Ecology Resources*, 11, 2010.

[21] C. P. Klingenberg. Cranial integration and modularity: insights into evolution and development from morphometric data. *Hystrix, the Italian Journal of Mammalogy*, 24(1):43–58, 2013.

[22] C. P. Klingenberg, K. Mebus, and J.-C. Auffray. Developmental integration in a complex morphological structure: how distinct are the modules in the mouse mandible? *Evolution & Development*, 5(5):522–531, 2003.

[23] C. P. Klingenberg and S. D. Zaklan. Morphological integration between developmental compartments in the drosophila wing. *Evolution*, 54(4):1273–1285, 2000.

[24] J. Krieger and N. MacLeod. A repository of morphometric tools and online morphometric workbenches. *Webpage:* http://www.morphotools.net/index.html (Software packages and Mathematica notebooks, page retrieved on 2013/03/16).

[25] J. Lewis and K.-i. Anjyo. Direct manipulation blendshapes. *IEEE Computer Graphics and Applications*, 30(4):42 –50, 2010.

[26] L. R. Monteiro, V. Bonato, and S. F. Dos Reis. Evolutionary integration and morphological diversification in complex morphological structures: mandible shape divergence in spiny rats (Rodentia, Echimyidae). *Evolution & development*, 7(5):429–439, 2005.

[27] P. O'Higgins and N. Jones. Morphologika 2.5 software package. *Webpage:* http://sites.google.com/site/hymsfme/resources (page retrieved on 2013/03/16).

[28] S. Renaud, P. Alibert, and J.-C. Auffray. Modularity as a source of new morphological variation in the mandible of hybrid mice. *BMC evolutionary biology*, 12(1):141, 2012.

[29] F. J. Rohlf and M. Corti. Use of two-block partial least-squares to study covariation in shape. *Systematic Biology*, 49(4):740–753, 2000.

[30] D. Rueckert, A. F. Frangi, and J. A. Schnabel. Automatic construction of 3d statistical deformation models using non-rigid registration. Proc. MICCAI '01, pages 77–84, 2001.

[31] T. Schultz and G. Kindlmann. Superquadric glyphs for symmetric second-order tensors. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE Visualization)*, 16(6):1595–1604, 2010.

[32] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on*, pages 336–343, 1996.

[33] J. R. Tena, F. De la Torre, and I. Matthews. Interactive region-based linear 3d face models. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 30:76:1–76:10, 2011.

[34] D. W. Thompson. *On Growth and Form*. Dover, 1917.

[35] L. R. Tucker. An inter-battery method of factor analysis. *Psychometrika*, 23(2):111–136, 1958.

[36] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Symmetric log-domain diffeomorphic registration: A demons-based approach. Proc. MICCAI '08, pages 754–761, 2008.

[37] G. W. Weber and F. L. Bookstein. EVAN toolbox. *Webpage:* http://evan-society.org/node/42 (Open Source package, page retrieved on 2013/03/16).

[38] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM '97*, page 1742, 1997.

[39] V. Wiens. Volumetric segmentation of complex bone structures from medical imaging data using reeb graphs. In *Central European Seminar on Computer Graphics for Students*, pages 113–120, 2013.

[40] H. Wold. Path models with latent variables: The NIPALS approach. *Quantitative Sociology: International perspectives on mathematical and statistical model building*, 1975.

[41] M. L. Zelditch, A. R. Wood, R. M. Bonett, and D. L. Swiderski. Modularity of the rodent mandible: integrating bones, muscles, and teeth. *Evolution & development*, 10(6):756–768, 2008.

[42] C. P. E. Zollikofer and M. S. Ponce De Léon. Visualizing patterns of craniofacial shape variation in *Homo sapiens*. In *Biological Sciences 269(1493)*, pages 801–807, 2002.