



Alignment of Multiple Configurations Using Hierarchical Models

Yann RUFFIEUX and Peter J. GREEN

We describe a method for aligning multiple unlabeled configurations simultaneously. Specifically, we extend the two-configuration matching approach of Green and Mardia (2006) to the multiple configuration setting. Our approach is based on the introduction of a set of hidden locations underlying the observed configuration points. A Poisson process prior is assigned to these locations, resulting in a simplified formulation of the model. We make use of a structure containing the relevant information on the matches, of which there are different types to take into account. Bayesian inference can be made simultaneously on the matching and the relative transformations between the configurations. We focus on the particular case of rigid-body transformations and Gaussian observation errors. We apply our method to a problem in chemoinformatics: the alignment of steroid molecules. Supplementary materials are available online.

Key Words: Chemoinformatics; Markov chain Monte Carlo; Matching; Rigid-body transformation; Shape analysis; Steroids.

1. INTRODUCTION

In many scientific disciplines one is confronted with the problem of comparing objects. Typically, the scientist locates a number of characteristic points, called landmarks, which correspond on the objects of a given population. For example, a landmark might be a recognizable location on a given organism, such as the corner of an eye, the tip of a finger, or the meeting of two sutures on a skull. Numerous techniques have been studied over the years for the geometrical comparison of objects when the landmarks are labeled, that is, when the point correspondences between the objects under study have been established. Now if the landmark configurations are unlabeled, so that the correspondences between the points of each configuration are unknown, then our problem also becomes one of *matching*: identifying and labeling corresponding landmarks.

Yann Ruffieux is Assistant Researcher, FSB IMA STAT Station 8, EPFL—Swiss Federal Institute of Technology, 1015 Ecublens, Switzerland (E-mail: yann.ruffieux@epfl.ch). Peter J. Green is Professor, School of Mathematics, University of Bristol, Bristol BS8 ITW, U.K. (E-mail: P.J.Green@bristol.ac.uk).

© 2009 American Statistical Association, Institute of Mathematical Statistics,
and Interface Foundation of North America

Journal of Computational and Graphical Statistics, Volume 18, Number 3, Pages 756–773
DOI: 10.1198/jcgs.2009.07048

A number of methods have been developed for the alignment of unlabeled point configurations, in various contexts. In image analysis, for instance, Cross and Hancock (1998) and, more recently, Lin, Zhu, and Wang (2007) used graph theory techniques for the matching of point sets representing two-dimensional images, whereas Chui and Rangarajan (2000) considered the use of non-rigid-body transformations. The alignment problem has also attracted a great deal of interest from the chemoinformatics community (Lemmen and Lengauer 2000). In this field it is a common assumption that structurally similar molecules have similar activities. This assumption has led to the development of quantitative structure-activity relationship (QSAR) analysis which seeks to quantify the link between the chemical structure and the observed properties of a molecule. In drug design, for example, a subject of prime interest is the local interaction between a small molecule (the ligand) and a given protein receptor. If the geometrical structure of the receptor is known, then established methods such as docking can be applied to specify the protein-ligand interaction. However, in most cases this structure is unknown, meaning the drug designer must rely on a study of the similarity (or diversity) in available ligands. The alignment of the molecules is a first important step toward such a study.

We focus specifically on generalizing the approach of Green and Mardia (2006), who described a Bayesian methodology for aligning two point configurations. We wish to extend this methodology to deal with an arbitrary number of configurations. Independent pairwise comparison of the configurations could be an option, but would be somewhat cumbersome and incoherent. Unless all of the configurations are treated simultaneously in a single model, there is loss of information, for example about the “noise-free” locations of the matched points. In Section 3.4 we see some evidence of the empirical impact of this. The elegance of the pairwise model makes this extension natural and relatively straightforward. The problem of matching multiple configurations has also been addressed by Dryden, Hirst, and Melville (2007); see the Discussion section of our article for further details.

This article is organized as follows. In Section 2 we treat the simultaneous alignment of multiple point configurations. We describe a hierarchical Bayesian model for this task, and propose a Markov chain Monte Carlo algorithm for making inference on the model in the case of rigid-body transformations between the configurations. In Section 3 we consider an application of our approach to the matching of three steroid molecules. Finally in Section 4 we make an assessment of our methods and suggest directions for future work.

2. HIERARCHICAL MODELING OF MULTICONFIGURATION ALIGNMENT

In this section we consider a hierarchical model for matching multiple configurations simultaneously. We closely follow the approach of Green and Mardia's (2006) two-configuration method, though we must now allow for the possibility of many types of matches.

2.1 THE ALIGNMENT PROBLEM

Suppose we are given C configurations $x^{(1)}, x^{(2)}, \dots, x^{(C)}$, whose points are recorded in d -dimensional real space: for $c = 1, 2, \dots, C$, write $x^{(c)} = \{x_j^{(c)}, j = 1, 2, \dots, n_c\}$, where $x_j^{(c)} \in \mathbb{R}^d$ and n_c is the number of points in configuration $x^{(c)}$. The labeling is assumed to be arbitrary, thus providing no initial information on the correspondences between points. We wish to align the C configurations simultaneously by establishing these correspondences and filtering out the relative transformations between the configurations.

We introduce a set of hidden locations $\mu = \{\mu_i\} \subset \mathbb{R}^d$. These can be interpreted as the “true” locations of the configuration points, so that the latter are noisy observations of the former. Specifically, define the labeling arrays $\xi^{(1)}, \xi^{(2)}, \dots, \xi^{(C)}$, which link the index of an observation to that of its corresponding hidden point. In particular, $\xi_j^{(c)}$ is the index of the μ -point underlying the observation $x_j^{(c)}$. Assume that a hidden location is observed at most once in each configuration, and that it may remain unobserved. Thus the elements within each $\xi^{(c)}$ are distinct, and a μ -point may generate anywhere between zero and C configuration points.

Now suppose each configuration goes through some transformation before being observed. For $c = 1, 2, \dots, C$, let $\mathcal{A}^{(c)}$ be the transformation bringing the points of the $x^{(c)}$ configuration *back* to the reference frame defined by the μ -points. Our C -configuration alignment model can be written as

$$\mathcal{A}^{(c)}x_j^{(c)} = \mu_{\xi_j^{(c)}} + \varepsilon_j^{(c)} \quad \text{for } j = 1, 2, \dots, n_c, c = 1, 2, \dots, C. \quad (2.1)$$

The random error vector $\varepsilon_j^{(c)}$ is assumed to have density $f^{(c)}$ and to be independent of the μ -points and of all the other errors.

Our primary objective with this model is to *match* observations of the same μ -point. Formally we wish to find maximal sets of configurations $\{x^{(i_1)}, x^{(i_2)}, \dots, x^{(i_k)}\}$ and indices $\{j_1, j_2, \dots, j_k\}$ such that $\xi_{j_1}^{(i_1)} = \xi_{j_2}^{(i_2)} = \dots = \xi_{j_k}^{(i_k)}$. There is an abuse of language here in that a match may involve points taken from more than two configurations.

Later we will assume that each $\mathcal{A}^{(c)}$ is an affine transformation, made up of a linear transformation matrix $A^{(c)} \in \mathbb{R}^{d \times d}$ and a translation vector $\tau^{(c)} \in \mathbb{R}^d$ so that $\mathcal{A}^{(c)}x_j^{(c)} = A^{(c)}x_j^{(c)} + \tau^{(c)}$ for $j = 1, \dots, n_c$ and $c = 1, \dots, C$. We will require constraints on the $\mathcal{A}^{(c)}$ to ensure identifiability of the model; we discuss this in Section 2.5.

2.2 HIERARCHICAL MODELING: PRELIMINARIES

We will regroup the matches in a parameter \mathcal{M} . How these matches are represented is irrelevant for the moment; one might wish to use binary matrices, as did Dryden, Hirst, and Melville (2007) and Green and Mardia (2006), or write each match as an index array containing the labels of the matched points and those of the configurations involved. For \mathcal{M} to be consistent with the model described in Section 2.1, we must in particular ensure that a given point is involved in exactly one (maximal) match, with the convention that an unmatched point is itself a trivial match of size 1. In this sense \mathcal{M} can be seen as a partition on the set of all observed configuration points, under the constraint that no two points from a given configuration be in the same subset. We stress the fact that the elements of \mathcal{M}

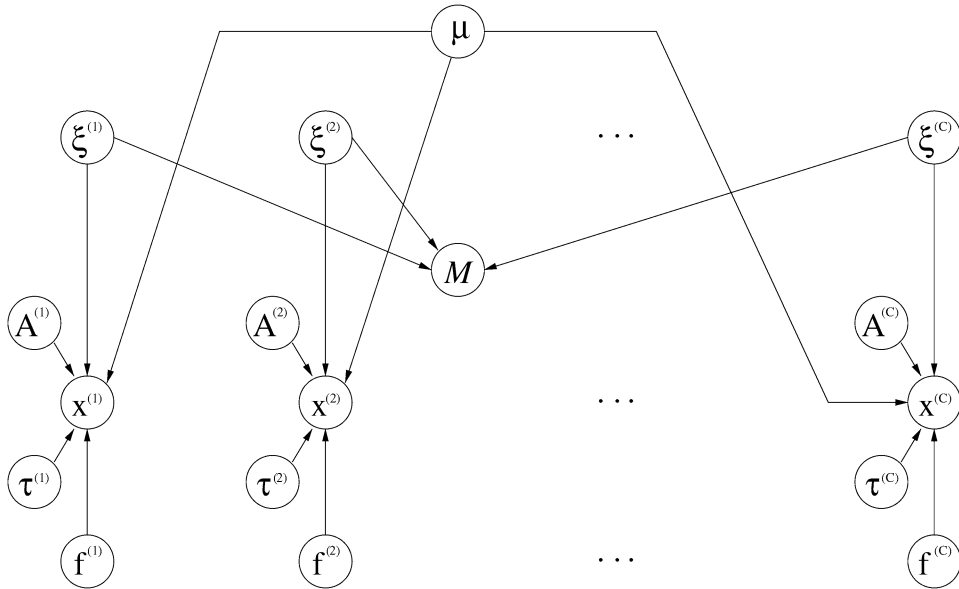


Figure 1. Directed acyclic graph of the hierarchical model.

refer to the indices of the matched points and give no information about their position. The directed acyclic graph (DAG) of our hierarchical model, including this new parameter \mathcal{M} , is displayed in Figure 1.

We categorize the matches contained in \mathcal{M} according to their “type.” Consider a generic set $I \subset \{1, 2, \dots, C\}$ of configuration indices, with $I \neq \emptyset$. This set corresponds to a type of match: for example, if $C = 3$, then $I = \{2, 3\}$ refers to a match involving a point from the $x^{(2)}$ configuration and a point from the $x^{(3)}$ configuration but none from the $x^{(1)}$ configuration. We call I -match a match involving *exactly* the configurations whose index is included in I . If $I = \{i_1, i_2, \dots, i_K\}$, an I -match can be represented by an index array (j_1, j_2, \dots, j_K) such that $\xi_{j_1}^{(i_1)} = \xi_{j_2}^{(i_2)} = \dots = \xi_{j_K}^{(i_K)}$ and such that $c \notin I$ implies $\xi_l^{(c)} \neq \xi_{j_k}^{(i_k)}$ for all $l = 1, 2, \dots, n_c$ and all $k = 1, 2, \dots, K$. We write $|I|$ as the number of configuration indices in I . If $|I| = 1$, our I -match is in fact an unmatched point; as stated above this will also be treated as a type of match.

2.3 POISSON PROCESS ASSUMPTION AND PRIOR DISTRIBUTION FOR THE MATCHES

We make the prior assumption that the μ -points follow a multivariate Poisson process with constant rate λ over a region $V \subset \mathbb{R}^d$ of volume v . Recall that each μ -point generates a number of observations or remains unobserved. For I as defined in the previous section, let q_I be the probability that a given hidden location generates an I -match. For instance, if $C = 3$, then $q_{\{1,3\}}$ is the probability that a particular μ -point is observed in the $x^{(1)}$ and $x^{(3)}$ configurations but not in the $x^{(2)}$ configuration. Thus the probability of a hidden location remaining unobserved is $1 - \sum_I q_I$. Assume also that the matches are generated independently from μ -point to μ -point, based on the same probabilities q_I . A very useful

consequence of these assumptions is that our global Poisson process can be partitioned into 2^C *thinned* Poisson processes: for fixed I , the set of μ -points which have generated an I -match is itself a Poisson process with rate λq_I ; furthermore, this process will be independent of the other processes of the partition.

We define the parameterization

$$q_I = \rho_I \cdot \prod_{c \in I} q_{\{c\}}, \tag{2.2}$$

where $\rho_I = 1$ if $|I| = 1$. This type of parameterization has been treated in other contexts, such as that of regression with binary response (Ekholm, Smith, and McDonald 1995) and genetic map functions (Speed 2005). The parameter ρ_I is sometimes called the dependence ratio or coincidence coefficient. Here it can be seen as a relative measure of how likely an I -match is to occur a priori.

Now we wish to assign a prior distribution to the match structure \mathcal{M} , based on the Poisson process assumptions described earlier. Let L_I be the number of I -matches contained in \mathcal{M} . Given $\{n_c, c = 1, 2, \dots, C\}$, we must have

$$L_{\{c\}} = n_c - \sum_{\{I: |I| \geq 2, I \ni c\}} L_I \quad \text{for } c = 1, 2, \dots, C. \tag{2.3}$$

The prior distribution for \mathcal{M} can be written as

$$p(\mathcal{M}) = p(\mathcal{M} \mid \{L_I\}) \cdot p(\{L_I\}). \tag{2.4}$$

From the Poisson process assumption on the μ -points, the counts L_I are independent Poisson variables with means $\lambda v q_I$. Using (2.2) and (2.3), we find that the prior distribution for the match counts has the form

$$\begin{aligned} p(\{L_I\}) &\propto \prod_I (\lambda v q_I)^{L_I} / \prod_I L_I! \\ &\propto \prod_I \left\{ \frac{\rho_I}{(\lambda v)^{|I|-1}} \right\}^{L_I} / \prod_I L_I!, \end{aligned} \tag{2.5}$$

so the q_I parameters conveniently cancel.

Now make the prior assumption that, conditional on the match counts $\{L_I\}$, the distribution for \mathcal{M} is uniform. In other words, consider as equally likely each match arrangement which is consistent with the counts. The number of such arrangements is

$$\prod_{c=1}^C n_c! / \prod_I L_I!,$$

as can be seen using a recursion argument. Using (2.4) and (2.5), it follows that the prior distribution for \mathcal{M} has the form

$$p(\mathcal{M}) \propto \prod_I \left\{ \frac{\rho_I}{(\lambda v)^{|I|-1}} \right\}^{L_I}. \tag{2.6}$$

The case $C = 2$ matches with Green and Mardia's (2006) expression for the prior distribution of their "matching matrix."

2.4 JOINT MODEL

We now seek to compute the joint likelihood of \mathcal{M} and $\mathcal{A} = \{\mathcal{A}^{(1)}, \mathcal{A}^{(2)}, \dots, \mathcal{A}^{(C)}\}$ given the set of configurations $X = \{x^{(1)}, x^{(2)}, \dots, x^{(C)}\}$.

Fix $I = \{i_1, i_2, \dots, i_K\}$ and let $\{x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_K}^{(i_K)}\}$ be the points of a given I -match in \mathcal{M} , where of course $K = |I|$. From (2.1), we find that

$$p(x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_K}^{(i_K)} \mid \mathcal{A}, \mu, \xi^{(1)}, \xi^{(2)}, \dots, \xi^{(C)}) = \prod_{k=1}^K |A^{(i_k)}| f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu_{\xi_{j_1}^{(i_1)}}),$$

where $|A|$ denotes the absolute value of the determinant of the matrix A . Now consider the set of μ -points which have generated an I -match—we mentioned that this set follows a Poisson process. Given \mathcal{M} , and therefore given L_I , the points of this set are uniformly distributed over the region V . As a result the contribution of the matched points defined earlier to our likelihood is

$$p(x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_K}^{(i_K)} \mid \mathcal{A}, \mathcal{M}) = v^{-1} \int_V \prod_{k=1}^K |A^{(i_k)}| f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu) d\mu.$$

The above integration will be carried out over \mathbb{R}^d . We are thus ignoring the edge effects from the boundary of V ; this is valid if V is taken large enough relative to the support of the error densities $f^{(c)}$.

Suppose $I = \{i_1, i_2, \dots, i_{|I|}\}$ and let S_I be the set of I -matches contained in \mathcal{M} . The elements of S_I are written as index arrays of the form $(j_1, j_2, \dots, j_{|I|})$, with the convention that $\{x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_{|I|}}^{(i_{|I|})}\}$ is the corresponding set of matched points. The contribution of the I -matches to the likelihood is

$$v^{-L_I} \prod_{(j_1, \dots, j_{|I|}) \in S_I} \int_{\mathbb{R}^d} \prod_{k=1}^{|I|} |A^{(i_k)}| f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu) d\mu.$$

Multiplying over all match types, the full likelihood of \mathcal{A} and \mathcal{M} can be seen to be

$$p(X \mid \mathcal{A}, \mathcal{M}) = \left(v^{-\sum_I L_I} \prod_{c=1}^C |A^{(c)}|^{n_c} \right) \times \prod_I \prod_{(j_1, \dots, j_{|I|}) \in S_I} \int_{\mathbb{R}^d} \prod_{k=1}^{|I|} f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu) d\mu. \tag{2.7}$$

We introduce prior distributions $p(A^{(c)})$ and $p(\tau^{(c)})$ for the transformation parameters, for $c = 1, 2, \dots, C$. These priors are left undefined for the time being. The parameters λ, v , and ρ_I are treated as fixed. From (2.6) and (2.7), the joint posterior distribution has the form

$$p(\mathcal{A}, \mathcal{M} \mid X) \propto \prod_{c=1}^C \{ p(A^{(c)}) p(\tau^{(c)}) |A^{(c)}|^{n_c} \} \times \prod_I \prod_{(j_1, \dots, j_{|I|}) \in S_I} \frac{\rho_I}{\lambda^{|I|-1}} \int_{\mathbb{R}^d} \prod_{k=1}^{|I|} f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu) d\mu. \tag{2.8}$$

Here and elsewhere, the “ \propto ” symbol indicates proportionality with respect to the variables to the left of the conditioning sign. Thus the μ -points and labeling arrays have been effectively integrated out; the relevant information contained in these parameters is captured by the structure \mathcal{M} . Note also that the volume v plays no role in our posterior distribution.

Now assume the error densities $f^{(c)}$ are centered Gaussian densities with covariance matrices all equal to $\sigma^2 I_d$. In this case the integrals in (2.8) can be written in closed form: for a given set of points $\{x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_{|I|}}^{(i_{|I|})}\}$, define

$$\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_{|I|}}^{(i_{|I|})}) = \sum_{k=1}^{|I|} \|\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - c\|^2,$$

where

$$c = \frac{1}{|I|} \sum_{k=1}^{|I|} \mathcal{A}^{(i_k)} x_{j_k}^{(i_k)}$$

and $\|\cdot\|$ is the Euclidean norm. Thus $\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})})$ is a measure of the deviation in the transformed points $\{\mathcal{A}^{(i_1)} x_{j_1}^{(i_1)}, \dots, \mathcal{A}^{(i_{|I|})} x_{j_{|I|}}^{(i_{|I|})}\}$. With this notation and the Gaussian assumption for the errors, one finds that

$$\int_{\mathbb{R}^d} \prod_{k=1}^{|I|} f^{(i_k)}(\mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} - \mu) d\mu = |I|^{-d/2} (2\pi\sigma^2)^{-d(|I|-1)/2} \times \exp\left\{-\frac{1}{2\sigma^2} \gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, x_{j_2}^{(i_2)}, \dots, x_{j_{|I|}}^{(i_{|I|})})\right\}. \quad (2.9)$$

This identity is also valid if $|I| = 1$, because $\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}) = 0$. Now a prior distribution $p(\sigma^2)$ can be introduced and the variance parameter σ^2 incorporated in the model (2.8).

2.5 INFERENCE WITH MARKOV CHAIN MONTE CARLO

We wish to make inference on the parameters of the model (2.8), given the data configurations X . The parameters of interest are the error variance σ^2 , the translations $\tau^{(c)}$, the transformation matrices $A^{(c)}$, and the set of matches \mathcal{M} . The ratios $\rho_I/\lambda^{|I|-1}$ will be considered as fixed hyperparameters, estimated through some other method.

The unwieldy aspect of the joint distribution (2.8) makes it difficult to use conventional analytic or numerical estimation methods in this context. An attractive possibility here is to use Markov chain Monte Carlo (MCMC) simulation. We simulate a Markov chain by updating the parameters in sweeps, in such a way that the underlying transition kernel of the chain verifies detailed balance, with (2.8) as the stationary, or limiting, distribution. The sampled chain can be used as a basis for inference, provided it has reached equilibrium. For an accessible introduction to MCMC methods, see for example Green (2001), whereas Robert and Casella (2004) gave a more detailed account.

To simplify our method, we will make the assumption that the transformation matrices $A^{(c)}$ are rotation matrices. Thus we are concentrating on rigid-body transformations, and the point configurations can be seen as elements of a size-and-shape space (Dryden and Mardia 1998).

We will use a Gibbs sampling scheme for updating the transformation and error variance parameters. The matches will be updated with a Metropolis–Hastings jump. The C++ implementation of the algorithm (with R interface), including instructions and functions for postprocessing, can be found as part of the supplementary material on the JCGS website.

Updating the Continuous Parameters

For the parameters σ^2 , $\tau^{(c)}$, and $A^{(c)}$, for $c = 1, 2, \dots, C$, conditionally conjugate priors can be found which result in full conditional distributions of the same form. This will make updating these parameters relatively straightforward. The conjugacy assumptions are not particularly restrictive here: in practice we do not expect to make use of strong prior information on the continuous parameters.

We assign an inverse gamma prior distribution to σ^2 ; in particular we set $\sigma^{-2} \sim \Gamma(a, b)$, where a and b are respectively the shape and rate parameters of the gamma distribution. From (2.8) and (2.9), the full conditional distribution of σ^{-2} is

$$(\sigma^{-2} \mid \mathcal{A}, \mathcal{M}, X) \sim \Gamma(\tilde{a}, \tilde{b}),$$

where

$$\tilde{a} = a + \frac{d}{2} \sum_I L_I (|I| - 1)$$

and

$$\tilde{b} = b + \frac{1}{2} \sum_I \sum_{(j_1, \dots, j_{|I|}) \in S_I} \gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})}).$$

Thus the error variance can be updated using a Gibbs sampler step, that is, by simulating from the full conditional inverse gamma distribution.

Set $\mathcal{A}^{(-c)} = \{\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(c-1)}, \mathcal{A}^{(c+1)}, \dots, \mathcal{A}^{(C)}\}$. We choose to assign Gaussian priors to the translation parameters. For $c = 1, 2, \dots, C$, suppose a priori that $\tau^{(c)} \sim \mathcal{N}_d(\mu^{(c)}, \eta_c^2 I_d)$. Using (2.8),

$$(\tau^{(c)} \mid \sigma^2, \mathcal{A}^{(-c)}, A^{(c)}, \mathcal{M}, X) \sim \mathcal{N}_d\left(\frac{1/\eta_c^2 \mu^{(c)} + 1/\sigma^2 m_c}{1/\eta_c^2 + 1/\sigma^2 w_c}, \frac{1}{1/\eta_c^2 + 1/\sigma^2 w_c} I_d\right),$$

with

$$m_c = \sum_{I: I \ni c} \sum_{(j_1, \dots, j_{|I|}) \in S_I} \frac{1}{|I|} \left\{ \left(\sum_{k: i_k \neq c} \mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} \right) - (|I| - 1) A^{(c)} x_{j_{k(c)}}^{(c)} \right\}$$

and

$$w_c = \sum_{I: I \ni c} \left(\frac{|I| - 1}{|I|} \right) L_I,$$

and where the subindex $k(c)$ is such that $x^{(i_{k(c)})} = x^{(c)}$. The translation parameters are thus also updated using a Gibbs move.

We can also find conjugate priors for the rotation matrices $A^{(c)}$, though this is less obvious. For $c = 1, 2, \dots, C$, set $p(A^{(c)}) \propto \exp\{\text{tr}(F_c^T A^{(c)})\}$ for some $d \times d$ matrix F_c . The fact that we are concentrating on rotation matrices means that $|A^{(c)}| = 1$ and $(A^{(c)})^{-1} = (A^{(c)})^T$. A somewhat involved calculation yields

$$p(A^{(c)} \mid \sigma^2, \mathcal{A}^{(-c)}, \tau^{(c)}, \mathcal{M}, X) \propto \exp[\text{tr}\{(F_c + S_c)^T A^{(c)}\}],$$

where $\text{tr}(\cdot)$ is the trace operator and S_c the $d \times d$ matrix

$$S_c = \frac{1}{\sigma^2} \sum_{I: I \ni c} \sum_{(j_1, \dots, j_{|I|}) \in S_I} \frac{1}{|I|} \left\{ \left(\sum_{k: i_k \neq c} \mathcal{A}^{(i_k)} x_{j_k}^{(i_k)} \right) - (|I| - 1) \tau^{(c)} \right\} \left(x_{j_k(c)}^{(c)} \right)^T.$$

The conditionally conjugate distribution $p(A) \propto \exp\{\text{tr}(F^T A)\}$ is called the matrix Fisher distribution, and is well known in directional statistics (Mardia and Jupp 2003, p. 289). Rather than updating the rotation matrices themselves, we will work on the corresponding rotation angles. For example, if $d = 2$, we define the angle $\theta^{(c)}$, whereas if $d = 3$, we have the three generalized Euler angles $\theta_{12}^{(c)}, \theta_{23}^{(c)}$, and $\theta_{13}^{(c)}$. Green and Mardia (2006, pp. 241–242) described how the angles can be updated when assuming a conjugate matrix Fisher prior for the rotation matrices, in the case where d is 2 or 3.

For simplicity, we consider here only the case where the $A^{(c)}$ are uniformly distributed and mutually independent a priori. This is achieved by assigning zero matrices to the F_c earlier. It is then true that the relative rotations $(A^{(c_1)})^T \cdot A^{(c_2)}$ are uniform and mutually independent for $c_2 \neq c_1$ and fixed c_1 . So without loss of generality, we can impose the identifying constraint that $\mathcal{A}^{(1)}$ be fixed as the identity transformation. This is the same as saying that the first data configuration lies in the same frame as the hidden point locations. The choice of which configuration to “fix” will not have a material effect on MCMC inference outside of sampling error: with uniform rotation and diffuse translation priors, the posterior distribution is invariant to this choice.

Updating the Matches

The matches will be updated with a Metropolis–Hastings jump. We write

$$\mathcal{M} = \{(t_1^1, t_2^1, \dots, t_C^1), (t_1^2, t_2^2, \dots, t_C^2), \dots, (t_1^K, t_2^K, \dots, t_C^K)\},$$

with $K = \sum_I L_I$. Each C -tuple $(t_1^k, t_2^k, \dots, t_C^k)$ represents a match, t_c^k being the index of the point from the $x^{(c)}$ configuration involved in the match. If a given configuration is not involved in the match, a “–” flag is inserted at the appropriate position. For instance, if $C = 3$, the 3-tuple $(2, 4, 1)$ refers to a match between $x_2^{(1)}, x_4^{(2)}$, and $x_1^{(3)}$, whereas $(-, 2, 1)$ is a match between $x_2^{(2)}$ and $x_1^{(3)}$, with no $x^{(1)}$ -point involved. We also include unmatched points in this list: $(1, -, -)$ indicates that $x_1^{(1)}$ is unmatched, for example.

Suppose that \mathcal{M} is the current list of matches in the MCMC algorithm. The jump proposal proceeds as follows:

- with probability q we choose to *split* a C -tuple; in this case we draw an element uniformly at random in the list \mathcal{M} .
 - If the C -tuple drawn corresponds to an unmatched point, we do nothing;
 - otherwise we split it into two C -tuples; for instance, $(2, 3, 1)$ can be split into $(2, -, -)$ and $(-, 3, 1)$. In general there will be many potential splits: here we could have chosen to split $(2, 3, 1)$ into $(-, 3, -)$ and $(2, -, 1)$. Suppose the match to be split is an I -match. Then there are $B_I = 2^{|I|-1} - 1$ ways to split this match. We select one of these splits uniformly at random.

- With probability $1 - q$ we choose to *merge* two C -tuples; in this case we select two distinct elements uniformly at random from \mathcal{M} .
 - If the two C -tuples drawn contain a common configuration, for example, $(j_1, k, -)$ and $(j_2, -, -)$, then we do nothing;
 - otherwise we merge the C -tuples; for example, $(j, k, -)$ and $(-, -, l)$ become (j, k, l) , whereas $(-, k, -)$ and $(-, -, l)$ become $(-, k, l)$.

The split and merge operations defined above form a complementary reversible pair. Clearly, all possible match arrangements can be explored using these two operations only.

The acceptance probability of a jump is readily worked out from (2.8). Suppose the proposal is to split an I -match into an I' -match and an I'' -match, such that $I = I' \cup I''$ and $I' \cap I'' = \emptyset$. Reverting to the algebraic representation, suppose we are splitting $(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})})$ into $(x_{j_1'}^{(i_1')}, \dots, x_{j_{|I'|}}^{(i_{|I'|})})$ and $(x_{j_1''}^{(i_1'')}, \dots, x_{j_{|I''|}}^{(i_{|I''|})})$. The acceptance probability for this proposal is $\min\{1, p_S\}$, where

$$p_S = \left(\frac{\rho_{I'} \rho_{I''} \lambda}{\rho_I} \right) \times \left(\frac{2\pi\sigma^2 |I|}{|I'| |I''|} \right)^{d/2} \times \frac{2(1-q)B_I}{q(K+1)} \\ \times \frac{\exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1'}^{(i_1')}, \dots, x_{j_{|I'|}}^{(i_{|I'|})})\} \exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1''}^{(i_1'')}, \dots, x_{j_{|I''|}}^{(i_{|I''|})})\}}{\exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})})\}}.$$

This acceptance probability is also valid when at least one of the new matches after the split is an unmatched point—recall that $\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}) = 0$ and that $\rho_I = 1$ if $|I| = 1$. Now suppose we attempt to merge $(x_{j_1'}^{(i_1')}, \dots, x_{j_{|I'|}}^{(i_{|I'|})})$ and $(x_{j_1''}^{(i_1'')}, \dots, x_{j_{|I''|}}^{(i_{|I''|})})$ into $(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})})$. The acceptance probability for this jump is $\min\{1, p_M\}$, where

$$p_M = \left(\frac{\rho_I}{\rho_{I'} \rho_{I''} \lambda} \right) \times \left(\frac{|I'| |I''|}{2\pi\sigma^2 |I|} \right)^{d/2} \times \frac{qK}{2(1-q)B_I} \\ \times \frac{\exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1}^{(i_1)}, \dots, x_{j_{|I|}}^{(i_{|I|})})\}}{\exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1'}^{(i_1')}, \dots, x_{j_{|I'|}}^{(i_{|I'|})})\} \exp\{-1/(2\sigma^2)\gamma_{\mathcal{A}}(x_{j_1''}^{(i_1'')}, \dots, x_{j_{|I''|}}^{(i_{|I''|})})\}}.$$

To speed up the exploration of the parameter space, we will typically make several match jump proposals within each sweep of the MCMC algorithm.

3. APPLICATION ALIGNING STEROID MOLECULES

3.1 THE DATA

For this example we select $C = 3$ steroid molecules from the CoMFA database, which can be accessed at <http://www2.ccc.uni-erlangen.de/services/steroids>. This database has become a benchmark for testing computer-assisted drug design methods, thanks in large part to Cramer, Paterson, and Bunce's (1988) use of it in their Comparative Molecular Field

Analysis (CoMFA). Ever since the publication of that paper, these molecules have been used as a training set for various 3D quantitative structure-activity relationship (QSAR) methods (see Coats 1998). The three molecules, plus the two used later in this section, may also be obtained from the supplementary material on the *JCGS* website.

The three molecules that we have chosen here are aldosterone, cortisone, and prednisolone, which we label $x^{(1)}$, $x^{(2)}$, and $x^{(3)}$, respectively. Each of these molecules contains $n_1 = n_2 = n_3 = 54$ arbitrarily labeled atoms in $d = 3$ dimensional space. We wish to align these molecules using the methodology described in this article.

Here we have seven types of matches to deal with, including the unmatched types: {1}, {2}, {3}, {1, 2}, {2, 3}, {1, 3}, and {1, 2, 3}. For simplicity we will drop the brackets and commas from the I sets when appropriate, so that for example $L_{\{1,2\}}$ and $\rho_{\{2,3\}}$ are written as L_{12} and ρ_{23} , and a 13-match is a match involving the first and third configurations but not the second.

3.2 RESULTS

The MCMC algorithm was launched without setting any initial matches. We set $\rho_{12}/\lambda = \rho_{23}/\lambda = \rho_{13}/\lambda = 31.25$, and $\rho_{123}/\lambda^2 = 3,660$; see Section 3.3 for an explanation of how we chose these values. By convention ρ_1 , ρ_2 , and ρ_3 are fixed at 1. The error variance parameters were set to $a = 1$ and $b = 0.1$. The transformation priors were rendered largely noninformative by setting $\mu^{(2)} = \mu^{(3)} = 0$ and $\eta_2 = \eta_3 = 10$ and by assigning the zero matrix to F_2 and F_3 . The sampler was run for 50,000 sweeps, the first 10,000 being discarded as burn-in; 50 match proposals were made per sweep and we set $q = 0.5$ as the probability of choosing a merge in the Metropolis–Hastings step. This run took around 25 seconds on a Pentium 4 processor. More generally, for a fixed number of iterations, we found the computing time of the algorithm to increase with C in a roughly linear fashion.

From inspection of the parameter traces and posterior likelihood produced by the MCMC sampler, we conclude that the chain has reached equilibrium. See the supplementary material on the *JCGS* website for a graphical representation of some of these traces. The error variance and translations were estimated using the sample posterior means based on a subsample of 2000 after burn-in; we found $\bar{\sigma}^2 = 0.0076$, $\bar{\tau}^{(2)} = (-1.224, -0.639, -0.786)^T$, and $\bar{\tau}^{(3)} = (-0.796, -0.444, -0.640)^T$. The rotation matrices were estimated by taking their respective sampled polar parts (see Green and Mardia 2006, p. 248), giving us the estimates

$$\widehat{A}^{(2)} = \begin{pmatrix} 0.967 & 0.136 & -0.216 \\ -0.166 & 0.977 & -0.131 \\ 0.193 & 0.163 & 0.968 \end{pmatrix}, \quad \widehat{A}^{(3)} = \begin{pmatrix} 0.888 & 0.186 & -0.420 \\ -0.141 & 0.980 & 0.137 \\ 0.438 & -0.063 & 0.897 \end{pmatrix}.$$

The posterior sample means of the match counts were $(\bar{L}_{12}, \bar{L}_{23}, \bar{L}_{13}, \bar{L}_{123}) = (4.46, 5.59, 1.14, 42.70)$.

We choose to estimate \mathcal{M} by ranking the matches by order of their sample posterior probability, and selecting the k most frequent, say. We must of course ensure that the matches in the corresponding estimator $\widehat{\mathcal{M}}$ are compatible, that is, that no point is involved in more than one match. Setting a lower threshold of 0.5 for the posterior probabilities of the selected matches will ensure a coherent $\widehat{\mathcal{M}}$. However, we may “miss” some

matches in the process. Indeed, consider a MCMC output in which the matches $(l, j, -)$ and $(-, j, k)$ both appear with frequency 0.4 and (l, j, k) appears with frequency 0.2. If we set a lower probability threshold of 0.5, none of these three matches will be selected, which appears counterintuitive. To avoid this problem one might wish to generalize Green and Mardia's (2006) loss function approach to the multiple configuration setting. There are various ways to do this; however, all seem to lead to an awkward constrained optimization problem. Such problems can be set up as linear programs, but scale badly with problem size.

Here we find that 47 matches of size 2 or greater have probability higher than 0.9 and 54 have probability higher than 0.5. Of these 54 matches, 44 are 123-matches, 4 are 12-matches, 5 are 23-matches, and one is a 31-match. The three molecules are aligned graphically in Figure 2. Notice that in the top right corner of the figure, there appears to be a nonrandom observation error in the matched points. This might be a result of assigning too large a value to ρ_{123}/λ^2 , as will be seen in Section 3.3. It could also be the consequence of systematic model error; the assumption of rigid-body transformations might be invalid, for instance.

To study the vulnerability of the sampler to local modes, the following experiment was conducted: 100 independent MCMC runs were launched, all with the hyperparameters fixed at the values given above. After 50,000 sweeps of a run, the posterior likelihood (2.8) was computed and compared to a threshold value established from earlier runs, such as the one described above. If the likelihood was lower than this threshold, then the sampler

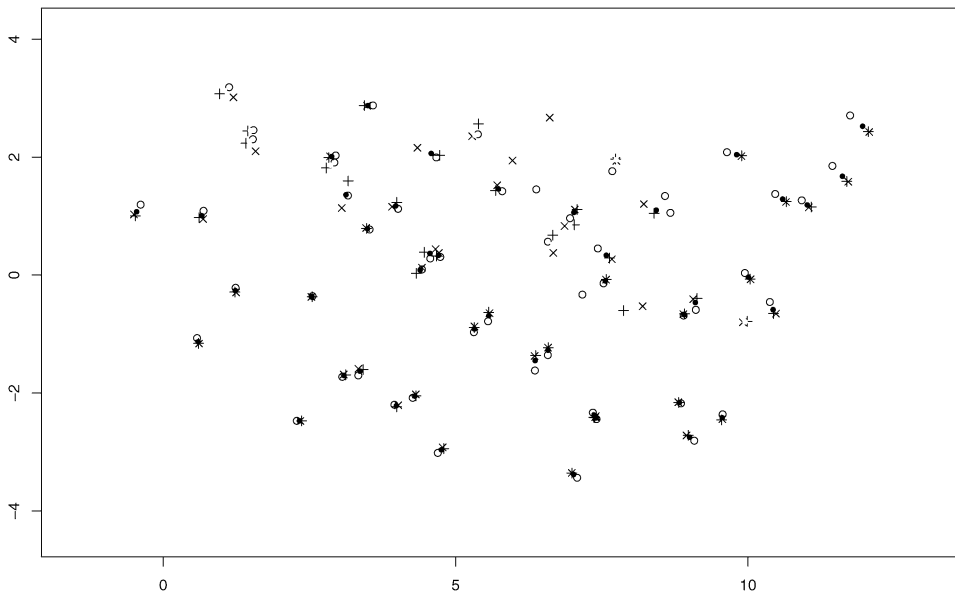


Figure 2. Aligned molecules from Section 3.2: the full transformations are estimated from a MCMC subsample of size 2000, and are filtered out from the data. The observations are then projected onto the principal components plane. The 'o' symbols represent the $x^{(1)}$ configuration (aldosterone), the '+' symbols the $x^{(2)}$ configuration (cortisone), and the 'x' symbols the $x^{(3)}$ configuration (prednisolone). The solid dots correspond to the centers of the 123-matches (black) and jk -matches (gray).

was adjudged to have become trapped in a local mode. Of the 100 runs, 91 passed this test and thus were deemed to have found their way to the main mode of the distribution. Of course these favorable results might be more a consequence of the nature of the data than of the robustness of our method. In case of difficulty we might, for instance, locate the carbon rings of the molecules using a graph-theory-based algorithm (see Dryden, Hirst, and Melville 2007, p. 246). Then our sampler could be initialized based on the matched carbon rings and the appropriate shape registration. Another option would be to adapt Lin, Zhu, and Wang's (2007) "strong seeds" approach to initialize the matching algorithm. This involves identifying seeds in the configurations and growing them into composite candidates for matching, via a branch-and-bound algorithm. The result is a much-reduced parameter space for the MCMC sampler to explore. It would be an interesting challenge to transpose this approach to our problem of aligning multiple molecules, but we choose to leave it for future work.

3.3 PRIOR SETTINGS

We now briefly study the effect of the hyperparameter values on the MCMC inference. In the case where no prior information on the transformation parameters is available, it is convenient to set a uniform prior on the rotation matrix and on the directions of the translations, as specified in Section 3.2. Furthermore, we typically select the variances η_2^2 and η_3^2 to be large enough so that the resulting translations encompass the configurations. We study in a little more detail the effect of the hyperparameters a , b , and $\rho_I/\lambda^{|I|-1}$.

In the above runs we set $a = 1$, thus assigning an exponential prior distribution to $1/\sigma^2$. The second hyperparameter b determines the rate of this distribution: larger b should result in larger variance in the observation errors, and thus more variability in the matching. Conversely, the smaller the value of b , the closer together a set of (transformed) points will have to be, to be considered as a candidate for a match. Increasing $b = 0.1$ by a factor of 10 will double the posterior mean of σ^2 . Inference on the matches is only slightly affected: a few two-way matches are replaced by 123-matches. Also, this increase generates a local mode problem, as some of the runs become entangled in a minor mode for an indefinite time. This is to be expected, because by increasing b we are allowing the sampler to explore additional alignments. Reducing $b = 0.1$ to $b = 0.01$ has little effect on either the matching or the posterior mean of the variance. However, reducing it further seems to create a second major mode in the posterior distribution, causing the algorithm to switch continually between two alignments. This new alignment is very similar to the first, except that approximately 10 of the 123-matches are replaced by 23-matches. The likely explanation for this is that in reducing b , we have become less tolerant toward matching, and thus have split several "borderline" 123-matches.

Now we study the influence of the ratios $r_I = \rho_I/\lambda^{|I|-1}$ on the matching. Recall that each hyperparameter r_I appears in the prior distribution (2.6) for the matches; we expect that increasing r_I will result in more I -matches being accepted in the algorithm. These ratios may be estimated by taking advantage of prior "guesses" one might have on the number of matches of each type. When such information is available, as is often the case

Table 1. Sample mean of the match counts and estimate of σ^2 in the case $C = 3$, for four different sets of prior guesses.

Prior specifications				Inference				
\tilde{L}_{12}	\tilde{L}_{23}	\tilde{L}_{13}	\tilde{L}_{123}	\tilde{L}_{12}	\tilde{L}_{23}	\tilde{L}_{13}	\tilde{L}_{123}	$\hat{\sigma}^2$
8	8	8	30	4.46	5.59	1.14	42.70	7.55×10^{-3}
25	5	5	20	7.32	4.81	0.74	40.90	7.24×10^{-3}
5	25	5	20	5.61	14.99	1.06	32.27	4.72×10^{-3}
5	5	25	20	4.21	4.74	2.14	42.70	7.71×10^{-3}

in practice, the argument of Green and Mardia (2001, p. 250) can be extended to the multivariate distribution (2.5). Suppose we have established the guesses $\{\tilde{L}_I\}$ for the match counts; if we set

$$r_I = \tilde{L}_I \cdot v^{|I|-1} / \prod_{c \in I} \tilde{L}_{\{c\}},$$

then the resulting prior distribution for the counts will have a unique mode in $\{\tilde{L}_I\}$. The value for the volume v must be determined from the data, but this is not usually difficult to do. For example, the ratios chosen in Section 3.2 are based on the guesses $\tilde{L}_{12} = \tilde{L}_{23} = \tilde{L}_{13} = 8$ and $\tilde{L}_{123} = 30$, with v fixed at 250. In Table 1 we consider four scenarios for the guesses, including that of Section 3.2. Displayed for each case is the sample posterior mean of the match counts and the σ^2 estimate. The latter value can be seen as a rough measure of the deviation in the matched points after transformation; it should not be used to choose between prior sets of values, however. It is interesting to note that the third set of guesses, which favors strict 23-matches, brings about a discernible change in the inference. The “borderline” 123-matches mentioned earlier have been replaced by 23-matches, giving a more precise alignment.

3.4 MULTIPLE VERSUS PAIRWISE MATCHING

We briefly consider the gain of using our multiple matching approach rather than aligning the configurations independently by pairs. For this purpose we add two further steroid molecules 11-deoxycorticosterone and 17 α -hydroxyprogesterone ($x^{(4)}$ and $x^{(5)}$, respectively) to the three described earlier; both contain 54 atoms.

First we treat the pairwise alignment of molecule $x^{(1)}$ to molecules $x^{(2)}$, $x^{(3)}$, $x^{(4)}$, and $x^{(5)}$, respectively. In Table 2 we display some of the inference obtained when aligning $x^{(1)}$ and $x^{(2)}$, for different prior scenarios. The pairwise alignments of $x^{(3)}$, $x^{(4)}$, $x^{(5)}$ to $x^{(1)}$ give mostly similar results in terms of number of inferred matches; we choose not to display them here. In particular the mean numbers of unmatched points for the case $\tilde{L}_{12} = 30$ are all between 6 and 10. Table 3 contains the results when aligning the five molecules simultaneously using our multiple-configuration method. Once \tilde{L}_{12345} drops below 35, a fair portion of the first molecule becomes disengaged from the other four (this was already apparent when aligning three molecules; see Section 3.3). From Table 2 we see that in the pairwise case one would have to set \tilde{L}_{12} to be as low as 10 to obtain an alignment similar

Table 2. Sample match count means and estimate of σ^2 for five different sets of prior values, in the context of matching $x^{(2)}$ to $x^{(1)}$.

Prior specifications		Inference		
\tilde{L}_{12}	\tilde{L}_1	\bar{L}_{12}	\bar{L}_1	$\hat{\sigma}^2$
30	24	47.48	6.52	9.01×10^{-3}
25	29	45.72	8.28	8.36×10^{-3}
20	34	42.23	11.77	6.99×10^{-3}
15	39	36.55	17.45	4.77×10^{-3}
10	44	35.07	18.93	4.33×10^{-3}

to the multiple one. This suggests that, in this context at least, the pairwise approach has a proclivity for overmatching. For reference the five-way alignment with $\tilde{L}_{12345} = 30$ is displayed in Figure 3.

The above comparison confirms that the inclusion of two or three additional configurations may have a positive impact on the alignment inference. One might understand this as a ‘‘borrowing of strength’’ of sorts; further configurations provide further information on the number and location of implied μ -points, information which can in turn be exploited in the alignment of the initial configurations. Clearly, there is no way to take advantage of this information if the molecules are aligned by pairs.

4. DISCUSSION

In this article we have seen that the two-configuration matching approach of Green and Mardia (2006) generalizes readily to the multiconfiguration context. We believe that our fully Bayesian approach offers general and flexible inference, and that it can be adapted to deal with alignment problems in various contexts. The methodology was applied to the matching of three steroid molecules, with promising results: with this ‘‘easy’’ dataset, the sampler seemed to have little difficulty avoiding the anticipated local mode problem.

The problem of aligning multiple molecules has also been treated by Dryden, Hirst, and Melville (2007); their approach is similar to ours, in that a hierarchical model is constructed

Table 3. Sample mean of some of the match counts and estimate of σ^2 in the case $C = 5$, for five prior sets of values.

Prior specifications		Inference		
\tilde{L}_{12345}	\tilde{L}_1	\bar{L}_{12345}	\bar{L}_1	$\hat{\sigma}^2$
39.9	0.1	41.44	5.80	4.32×10^{-3}
35	1	35.36	10.20	2.98×10^{-3}
30	2	22.07	20.96	1.46×10^{-3}
25	5	21.55	21.96	1.36×10^{-3}
20	10	20.86	22.66	1.29×10^{-3}

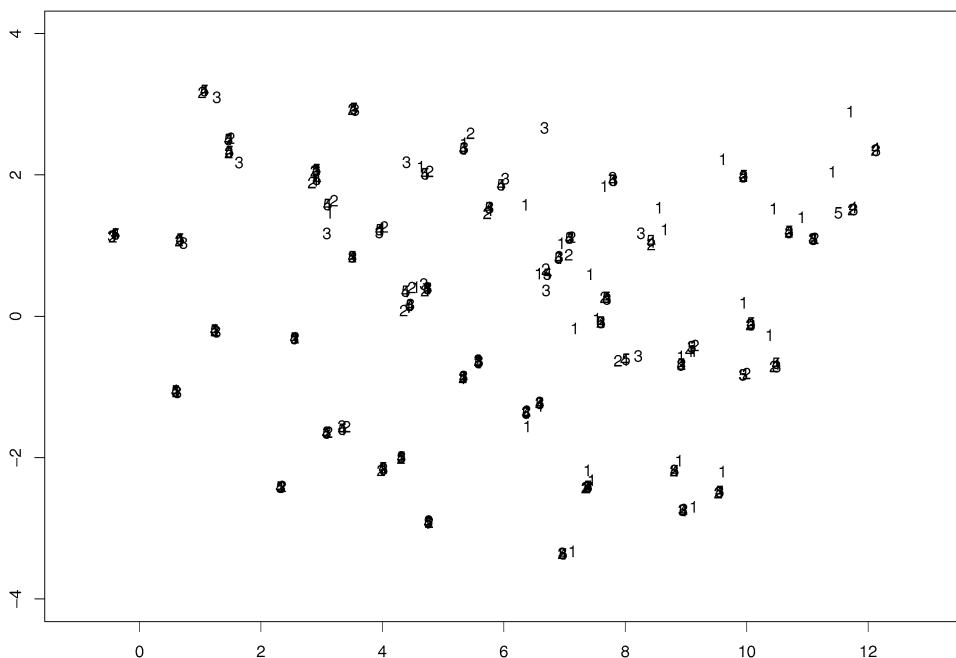


Figure 3. Multiple alignment of the five molecules from Section 3.4: the full transformations are estimated from a MCMC subsample of size 2000, and are filtered out from the data. The points are then projected onto the principal components plane, and are labeled according to the number of the configuration they belong to.

and a hidden reference molecule defined. However, the hidden points are not integrated out, and the transformations are maximized out using Procrustean registration techniques. Furthermore, only C “types” of matches are considered in their model (compared to our $2^C - C - 1$): the alignment is made pairwise between each observed point configuration and the hidden molecule. As far as we know, ours is the only available method which models multiple matchings of different types, in a general and mathematically rigorous manner. In terms of computation speed, the methodology of Dryden, Hirst, and Melville (2007) would probably be more efficient than the one proposed in this article when C is large. So the choice of method might depend upon the number of configurations to be aligned and the extent to which one wished to retain full statistical efficiency and control the prior match specifications.

An important aspect of alignment which is not addressed in this article is that of *marking*. In many contexts, additional information on the observations is available. For example, Dryden, Hirst, and Melville (2007) included “marks” on each atom of the molecules to be aligned; these marks may contain information influencing the matching, such as partial charge and van der Waals radius. In a similar vein, Green and Mardia (2006) included the possibility of coloring the observations, to model the possibility that points of the same color are more likely to be matched a priori. Thus knowledge of amino acid types can be used advantageously for the matching of active sites in proteins. Incorporating such information on the points may make the inference more clear-cut, by reducing multimodality in the posterior distribution.

It would be interesting to consider applications that assume nonrigid or even nonlinear transformations between the configurations. Our model allows for such transformations, but the implementation would have to be suitably adapted. The same can be said regarding the use of non-Gaussian observation errors and of different prior distributions for the parameters.

SUPPLEMENTAL MATERIALS

Data Sets: For use with R, a list containing the five steroids treated in the paper, in this order: aldosterone, cortisone, prednisolone, 11-deoxycorticosterone, 17 α -hydroxyprogesterone. (five_steroids.RData)

Computer Code: Contains the shared C++ code (main.so), the R code (MultAlign.R), and instructions for running the program (MultAlign.txt). (MAlign_code.tar.gz)

Supplementary Documents: Three figures complementing the article. (Supp_Figures.pdf)

ACKNOWLEDGMENTS

The authors thank Kanti Mardia for introducing them to this problem and for stimulating discussions, as well as Anthony Davison and John "Mac" McDonald for their helpful comments and suggestions.

[Received April 2007. Revised August 2008.]

REFERENCES

- Chui, H., and Rangarajan, A. (2000), "A New Algorithm for Non-Rigid Point Matching," *IEEE Conference on Computer Vision and Patter Recognition*, 2, 44–51.
- Coats, E. A. (1998), "The CoMFA Steroids as a Benchmark Dataset for Development of 3D QSAR Methods," *Perspectives in Drug Discovery and Design*, 12/13/14, 199–213.
- Cramer, R. D., III, Patterson, D. E., and Bunce, J. D. (1988), "Comparative Molecular Field Analysis (CoMFA): I. Effect of Shape on Binding of Steroids to Carrier Proteins," *Journal of the American Chemical Society*, 110, 5959–5967.
- Cross, A. D. J., and Hancock, E. R. (1998), "Graph Matching With a Dual-Step EM Algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20 (11), 1236–1253.
- Dryden, I. L., and Mardia, K. V. (1998), *Statistical Shape Analysis*, Chichester: Wiley.
- Dryden, I. L., Hirst, J. D., and Melville, J. L. (2007), "Statistical Analysis of Unlabelled Point Sets: Comparing Molecules in Chemoinformatics," *Biometrics*, 63, 237–251.
- Ekholm, A., Smith, P. W. F., and McDonald, J. W. (1995), "Marginal Regression Analysis of Multivariate Binary Response," *Biometrika*, 82 (4), 847–854.
- Green, P. J. (2001), "A Primer on Markov Chain Monte Carlo," in *Complex Stochastic Systems*, eds. O. E. Barndorff-Nielsen, D. R. Cox, and C. Klüppelberg, London: Chapman & Hall, pp. 1–62.
- Green, P. J., and Mardia, K. V. (2006), "Bayesian Alignment Using Hierarchical Models, With Applications in Protein Bioinformatics," *Biometrika*, 93 (2), 234–254.
- Lemmen, C., and Lengauer, T. (2000), "Computational Methods for the Structural Alignment of Molecules," *Journal of Computer-Aided Molecular Design*, 14, 215–232.

- Lin, L., Zhu, S.-C., and Wang, Y. (2007), "Layered Graph Match With Graph Editing," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8.
- Mardia, K. V., and Jupp, P. E. (2003), *Directional Statistics*, Chichester: Wiley.
- Robert, C. P., and Casella, G. (2004), *Monte Carlo Statistical Methods* (2nd ed.), New York: Springer.
- Speed, T. P. (2005), "Genetic Map Functions," in *Encyclopedia of Biostatistics* (2nd ed.), New York: Wiley.