Journal of Mathematical Imaging and Vision (2021) 63:528–540
https://doi.org/10.1007/s10851-021-01018-2

Analysing 'Simple' Image Registrations

Stephen Marsland[1] · Robert I. McLachlan[2] · Raziyeh Zarre[3]

Received: 31 July 2018 / Accepted: 14 January 2021 / Published online: 30 January 2021

Abstract
Processes such as growth and atrophy cause changes through time that can be visible in a series of medical images, following the hypothesis that form follows function. As was hypothesised by D'Arcy Thompson more than 100 years ago, models of the changes inherent in these actions can aid understanding of the processes at work. We consider how image registration using finite-dimensional planar Lie groups (in contrast to general diffeomorphisms) can be used in this process. The deformations identified can be described as points in the Lie algebra, thus enabling processes such as evolutionary change, growth, and deformation from disease, to be described in a linear space. The choice of appropriate Lie group becomes a modelling choice and can be selected using model selection; Occam's razor suggests that groups with the smallest number of parameters (which Thompson referred to as 'simple transformations') are to be preferred. We demonstrate our method on an example from Thompson of the cannon-bones of three hoofed mammals and a set of outline curves of the development of the human skull, with promising results.

Keywords Growth and form · Image registration · Lie groups

1 Introduction

Natural scientists classify organisms by their appearance and structure, known as form. The form of an organism is a characteristic of it that remains unchanged during the similarity transformations of translation, rotation, and scaling [8]. It is argued that 'form follows function', which suggests that by studying the form, it may be possible to infer the causes that underlie the variations of an organism over various timescales, such as growth, disease or evolution. These variations are measured and quantified, which may help us to discover processes of biological phenomena that are not identifiable from casual observation [3].

The first attempt to quantitatively study form was published by D'Arcy Thompson, most notably in his landmark book 'On Growth and Form' [24], the first edition of which was published over 100 years ago. In that book, Thompson deformed grids on which images were drawn by using 'simple' (i.e. low dimensional, often global) warps in order to make the appearance of different species match. In fact, many of the deformations that Thompson used appear to be affine, or what he calls 'isogonal' (i.e. Möbius), and thus are examples of planar Lie groups.

In biology and medicine, the change in the appearance of an organism or structure exhibits natural variation within a species as well as any underlying process of disease or evolution. Thus, it will be necessary to use large data sets and statistical methods to demonstrate conclusive results. Even in longitudinal studies of disease, where there are several images of the same person over time—which are potentially very useful for understanding the course of a disease—there will be substantial variety between individuals in the appearance and progression over time of the change process, which is complicated by the action of the normal growth or ageing process.

✉ Stephen Marsland
stephen.marsland@vuw.ac.nz

Robert I. McLachlan
r.mclachlan@massey.ac.nz

Raziyeh Zarre
raziyeh.zarre@waikato.ac.nz

[1] School of Mathematics and Statistics, Victoria University of Wellington, Wellington, New Zealand

[2] Institute of Fundamental Sciences, Massey University, Palmerston North, New Zealand

[3] Department of Mathematics and Statistics, University of Waikato, Hamilton, New Zealand

Thompson's work inspired the field of geometric morphometrics [3], where corresponding points identified in pairs of images (landmarks) are used to align the objects. The image is modified by the deformation induced by the transformation of these points, usually using thin-plate splines. However, these do not form a group, and statistical analysis in this field tends to be on the set of landmark points, not on the deformations.

Thompson's work is also often used as a motivation for diffeomorphic image registration (see, e.g. [15]); the diffeomorphisms are a group, albeit an infinite-dimensional one. One key difference is that Thompson explicitly stated that simple, preferably global, warps were to be preferred. Although diffeomorphic registration can potentially produce very good matches between images of objects, such as the brains of different individuals, it makes statistical analysis extremely difficult: while we have access to a Riemannian metric on the diffeomorphism group, the manifold is infinite dimensional and has curvature.

Within the field of diffeomorphic image registration of medical images there has been previous work considering the analysis of longitudinal development of disease over a period of time. This work has largely focussed on transporting the diffeomorphic trajectories of transformations into a common reference space, which enables the comparison of the vector fields directly. The most common approach is to use parallel transport, whereby vectors are transported along the geodesics, keeping the orientation of the vector unchanged. This has been applied to both the 'Large Deformation Diffeomorphic Metric Mapping' (LDDMM) [27] and 'Stationary Velocity Field' (SVF) [5,10] methods of constructing diffeomorphisms. In [9] a method for computing the parallel transport is given, which is based on Schild's ladder, a construction from general relativity. For example, one-parameter subgroups may be geodesic in the SVF framework [11]. The alternatives to parallel transport are to use coadjoint transport of the vector field directly [28] or to use the group properties to define a way to compose the transformations (known as transformation conjugation) [19].

All of these methods suffer from the problem that the entire diffeomorphic transformation has to be transported and compared. This requires a large number of computationally expensive registrations, where a lot of the deformation is quite possibly irrelevant to the process being studied. We argue that it is better to find the gross trajectories (in a simple group) first, and perform statistical analysis of the low-dimensional representations of longitudinal variation in the small group, before considering more expensive, infinite-dimensional image registrations. Our approach is to consider image registration in simple groups with a small number of parameters, in line with Thompson's stated aims in 'On Growth and Form'. This is an important consideration from the applications point of view: while any two shapes that are broadly (e.g. topologically) the same can be registered, close resemblance of a family of shapes under simple transformations (a low-dimensional model and the geometric significance of a group) is more significant. Further, the transformations themselves (as well as non-matching features) carry information.

There are two alternative approaches that can be used to simplify the analysis. One is to use strong regularisation in the registration in order to simplify the deformation that is identified; this is already a standard part of the optimisation process in LDDMM. However, while the deformation looks smoother, it is still necessary to analyse the entire deformation, which is parameterised by either the pixel grid, or a large number of landmarks on the image. Another possibility is to apply some form of statistical averaging across a set of deformations, again to simplify the warps. This has been done for landmarks, and leads to stochastic versions of the deformation equations [13,20,25]. Again, although this leads to smoother deformations, it does not reduce the underlying complexity of the allowable transformations. By explicitly reducing the number of free parameters in the deformation, we produce transformations that are far simpler, and therefore more amenable to statistical analysis, as well as providing more information about whether or not the model fit is good.

We propose a methodology that can assist in the statistical analysis of the underlying deformations by providing a low-dimensional space in which the images (and their underlying variation) are represented. Our approach is motivated by Thompson's focus on simple deformations. We return to the spirit of Thompson's work, and apply deformations chosen from planar Lie groups. While the quality of the match between the images is lower than using diffeomorphisms, the gross structures align well, and we can study the deformations between objects in a low-dimensional space (a subspace of the Lie algebra) using standard methods. In addition, by comparing the quality of the image match between different groups it is possible to select as a model the group that appears to best describe the underlying process. The use of model selection to identify the particular low-dimensional group being used provides further information about the underlying growth (or other deformation) process. In addition, the use of simple groups means that the models are simple both qualitatively and geometrically, and may possibly carry biologically relevant information, as argued by Petukhov [18].

We introduce our approach in Sect. 2, describing the method of image registration we use, introducing Lie groups briefly, and describing the statistical analysis methods that are then available, as well as the simple algorithm that enables us to interpolate between images. We then provide two examples in Sect. 3: the cannon-bones of the ox, sheep, and giraffe, and the development of the human skull. We have chosen these because they are the canonical examples from the literature, being based on an example of Thompson himself

[24], and of Petukhov [18], who considered the action of non-Euclidian groups as a method of growth for living bodies. They also reflect examples of evolutionary change, and growth.

2 Using Planar Lie Groups for Image Registration and Analysis

2.1 Image Registration

Given two images I and J referred to as the source and target, respectively, image registration consists of identifying a deformation that makes the difference between the transformed source and the target as small as possible, as measured in some norm. There are many methods of performing such registrations, with three primary differences being the choice of set from which the deformations can be chosen, the function by which image differences are measured, and the method of optimisation. References giving introductions to image registration methods include [12,17,21].

In this paper, our registrations are performed using gradient descent on sum-of-squares difference (i.e. $RSS = \|I \circ \varphi_{IJ}^{-1} - J\|_2^2$, where φ_{IJ}^{-1} is the (inverse) transformation between the pair of images) between the pixel intensities, starting at the identity deformation, which is the standard approach. Images are blurred by convolving with a Gaussian filter of standard deviation four pixels in a first pass, and then registered without blurring in a second pass. As well as improving the registration results for full images, the blurring also enables line drawings to be registered, since the lines already overlap after the first registration. The quality of the registration from this process largely seems to be very good, although in some cases we do find it useful to shade the inside of objects

Where our approach differs from previous methods is in our choice of transformation sets. It is common to register images using either the similarity group (the set of translations, rotations, and scaling, being four dimensional for 2D images), the affine group (which includes shears, and is six dimensional for 2D images), or the full diffeomorphism group (the infinite-dimensional group of smooth functions with smooth inverses). Instead, we consider choices from the planar Lie groups, which are finite dimensional and relatively 'simple' in the sense of Thompson. There are 28 such groups, and they are classified in [6] up to the change of coordinates. The largest group in this set is that of invertible 2×2 matrices semi-direct product with the reals of dimension r, $\mathbb{R}^r \rtimes GL(2, \mathbb{C})$, i.e. the affine group for $r = 2$. The planar groups have a left action on images given by right composition with the inverse mapping, i.e. $I \cdot \varphi = I \circ \varphi^{-1}$.

2.2 Lie Groups

A Lie group $(G, \circ)$ is a group that is also a smooth manifold where the group operation $(g, h) \in G \times G \mapsto g \circ h \in G$ and inverse map $g \in G \mapsto g^{-1} \in G$ are both smooth and compatible with the manifold structure. The Lie algebra of a Lie group is the tangent space at the identity element of the group, denoted by T_eG. A Lie group is diffeomorphic to its Lie algebra in a neighbourhood of its identity, and this diffeomorphism is given by the exponential map. Hence the global group can be replaced by its local linearisation, which is easier to work with. In addition, a metric can be defined on a Lie group, which enables us to measure the distance between images. Planar Lie groups are Lie groups together with an action of the group on the plane.

We identify three groups of particular interest: PSL(2, $\mathbb{C}$) (the Möbius group), which acts on complex z by:

$$z \mapsto \frac{az + b}{cz + d},\ a, b, c, d, z \in \mathbb{C},\ ad - bc = 1, \tag{1}$$

and is 6 dimensional, since the constraint reduces the 4 complex variables to 3, the 8-dimensional projective special linear group PSL(3, $\mathbb{R}$), which acts on the plane by:

$$(x, y) \mapsto \left(\frac{a_1 x + b_1 y + c_1}{a_3 x + b_3 y + c_3}, \frac{a_2 x + b_2 y + c_2}{a_3 x + b_3 y + c_3} \right), \tag{2}$$

where $\det \begin{pmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ a_3 & b_3 & c_3 \end{pmatrix} = 1$, and the product of two copies of the linear-fractional group, PSL(2, $\mathbb{R}$) $\times$ PSL(2, $\mathbb{R}$), which acts on the plane by:

$$(x, y) \mapsto \left(\frac{a_1 x + b_1}{c_1 x + d_1}, \frac{a_2 y + b_2}{c_2 y + d_2} \right), \tag{3}$$

(where $a_1 d_1 - b_1 c_1 = a_2 d_2 - b_2 c_2 = 1$) and is 6 dimensional. All of these groups include the 4-dimensional set of Euclidean similarities as subgroups. Together with their subgroups, they comprise 20 of the 28 finite-dimensional planar Lie groups. We use them because they correspond to those used by the original authors in our examples (Sect. 3). However, the principles that we demonstrate are true for all of the groups, not just these examples.

The particular parameters of the transformation provide a coordinate representation of the warp necessary to construct each target image from the source in a group-specific space. We now consider how to use this information for the analysis of image registration data.

The purpose of choosing Lie groups for the registration is that they are manifolds as well as groups, and so there is a natural space in which to analyse the deformations that are identified for the registration, namely the Lie algebra.

The coordinates of the group parameters that transform the source into a particular target (up to the variations allowed by that group) can be used as a location in a d-dimensional space, where d is the dimension of the group. If a metric is then defined on that space, then distances between images can be defined in the group. In addition, series of images can be examined—either as indications of an underlying evolutionary or disease process acting through time, or to try to cluster related images—by plotting a line that fits through the data points, representing a trajectory that passes close to the data. It is then simple to interpolate between the images, extrapolate the further progression of the process (with the usual caveats about accuracy of extrapolation), or to plot longitudinal data from several different people and compare them.

As the transformation acts on the left, we use a right-invariant Riemannian metric (since this means that if a transformation is applied to a set of images, their pairwise distances (and relative transformations) do not change), which on matrix Lie groups can be written as:

$$\langle X, Y\rangle_{\mathcal{A}} = tr((X\mathcal{A}^{-1})^T Y\mathcal{A}^{-1}),$$

where $X, Y \in \mathfrak{g}$, the Lie algebra, $tr(\cdot)$ is the trace of the matrix, and the base point $\mathcal{A}$ is a matrix in the group.

2.3 Interpolation in the Algebra

We assume that we have a family of images of putatively related forms. Our approach consists of fitting a curve through the data in the Lie algebra and then mapping the curve into the group using the exponential map, yielding the following simple algorithm. Note that the group exponential and logarithm, exp and log, are taken at the identity throughout the paper.

Let $\{\varphi_i, i = 1, 2, 3, \dots\}$ be the given data in the group, i.e. the set of registrations of each image to the reference. If one of the images (j) is chosen to be the reference, then φ_j will be the identity transformation, mapping the image to itself, and $\varphi_i, i \neq j$ are the transformations from that image to each of the others. In the case of exact matching, it would not matter which image is chosen as the reference, since the transformations could be transformed between the images by computing the inverse transforms explicitly. However, in general, the registration is imperfect (inexact), so that the inverse mapping φ_i^{-1} is not the optimal transformation from the reference image back to image i.

There are a few ways to approach this problem. One is to ignore it and assume that the forward and inverse transformations are reasonably similar. This means that any of the images from the data set can be chosen and used. In the case of a growth or evolutionarily development data set, where there is an explicit time-ordering of the images, the hypothesis is that the development is the growth process presents as changing group parameters, likely monotonically. In this case, it is sensible to use the temporal information to choose the reference image, as either the first image, or the middle one (since the latter choice will be less different from the first and last elements of the time series than they are from each other, and so exhibit less extreme parameter values). The alternative, which is particularly appropriate for non-temporal data, is to use the mean of the transformed images as a moving target (so compute the pixel-wise mean of the images as the initial reference, register to that, and then update this reference with the mean of the transformed images, and iterate). This is a form of groupwise registration [14,22,26].

Regardless of how the transformations are found, the data analysis is then computed by the following process:

- Map φ_i into the Lie algebra using the log map: $\log(\varphi_i) = v_i$.
- Find a parameterised curve $E : \mathbb{R} \mapsto \mathfrak{g}$ passing near the v_i:

$$\min_{E\in\mathcal{E},t_i} \Sigma_i \|E(t_i) - v_i\|_2^2, \tag{4}$$

  where $\mathcal{E}$ is some set of curves, $\|\cdot\|$ is some metric on the Lie algebra, and t the parameters of E. (For a line through the origin, $E(t) = tA$ where $A \in \mathfrak{g}$.)
- Map the curve $E(t)$ into the group using the exponential map: $\exp(E(t))$.

2.4 Model Selection

The previous analysis presupposes that an appropriate group has been chosen. While in some cases there may be biological knowledge that informs this, in general the choice of group is a modelling choice. Applying Occam's razor, it is appropriate to use model selection to argue that the best group to use is the one with the lowest dimension that describes the transformations observed in the data well, i.e. it is a trade-off between model simplicity and the capacity of the model to describe the phenomena that underlie the data. We measure the first by the number of degrees of freedom (dimension) of the group and the second by the goodness-of-fit of the model to the data, which in the case of image registration corresponds to the residual sum-of-squares remaining after registration in the relevant group. Of course, care must be taken with this measure to ensure that the registration has not become stuck in a local minimum.

When comparing the registrations between different groups it is important to remember that in many cases one of the groups will be a subgroup of the other, so that the models are nested. In this case, the residuals from registration in the group must necessarily be at least as small as the residu-

als from registration in the subgroup, since there are more degrees of freedom to allow a better match.

As previously discussed, it is unlikely that the registration will be perfect, since even if the process driving the deformation is well modelled by a particular group of transformations, there will also be local variations and imperfections, as well as other processes occurring. These variations, as well as the nature of images (which are based on a discrete, pixellated representation of the underlying continuous picture) and of the numerical methods mean that the registrations will not be perfectly symmetric (i.e. the registration from I to J will not be the precise inverse of the registration from J to I). It is hard to interpret the numbers in the distance function, since they differ from image to image and group to group. We therefore normalise the distance function and calculate the geometric mean for each pair of the images by computing:

$$D = 1 - \sqrt{\frac{\|I \circ \varphi_{IJ}^{-1} - J\|_2^2}{\|I - J\|_2^2} \frac{\|J \circ \varphi_{JI}^{-1} - I\|_2^2}{\|I - J\|_2^2}} \tag{5}$$

for source image I, target J and forward and backward transformations φ_{IJ} and φ_{JI}. If I and J match perfectly with φ then $D = 1$, and if $\|I \circ \varphi_{IJ}^{-1} - J\|_2^2 = \|I - J\|_2^2$ and $\|J \circ \varphi_{JI}^{-1} - I\|_2^2 = \|I - J\|_2^2$ then $D = 0$. For black and white images, the value of D tells us the amount of the overlap and match between the images.

The approach is then to perform the set of registrations using each potential group of interest, and then use model selection to select the group that best describes the transformations that explain the data, i.e. to consider the group of allowable transformations as a model of the data deformations.

Two of the most popular model selection criteria are the Akaike information criterion (AIC):

$$\text{AIC} = -2 \ln \mathcal{L} + 2p, \tag{6}$$

and the Bayesian information criterion (BIC):

$$\text{BIC} = -2 \ln \mathcal{L} + p \ln n. \tag{7}$$

Here, $\mathcal{L}$ is the likelihood of the errors in the estimated model (i.e. group), p is the number of parameters in the model, and n is the sample size. In the case that the model errors are independent and normally distributed, $\ln \mathcal{L} = -n|\ln(\text{RSS}/n)|$ (where RSS is the residual sum-of-squares error, i.e. the difference between the images after registration). A lower value of AIC (resp. BIC) is preferred. A rule of thumb is that a difference of AIC or BIC between two models of 2–6 is positive, 6–10 is strong, and more than 10 is very strong. Note that these criteria cannot be used to compare methods based on the full diffeomorphism group, since it has an infinite number of parameters.

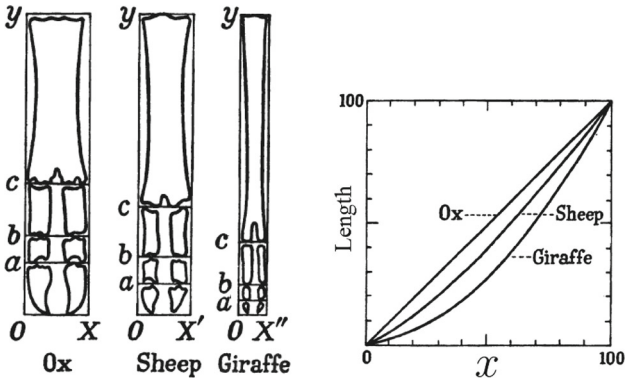

Fig. 1 *Left:* Cannon-bones of the ox, sheep and giraffe. The bones are scaled to have the same length in the y direction, and then corresponding points along that axis are identified. *Right:* Curves showing the cannon-bone length of the ox, sheep and giraffe as functions of the ox length (hence the identity relationship of the ox bone with itself). Both of these images are from [24]

A review of model selection in ecology [1] found that 84% of studies used AIC, 14% used BIC, and 2% used some other criterion. They argue that AIC is preferred when the 'true' model is extremely complex and essentially unknowable, and prediction errors are to be minimised, while BIC is preferred when the 'true' model is simple and can be in principle determined given enough data. We now use examples to demonstrate our method on an evolutionary process and a growth process. We have assumed in both cases that the true model is complex and not possible to find, and hence followed the suggestion in [1], and use AIC in this paper.

3 Examples

3.1 An Example from Thompson: Cannon-Bones

One of Thompson's examples in *On Growth and Form* [24] is a comparison of the cannon-bones (the bone in hoofed mammals that extends from the knee or hock to the fetlock) of the ox, sheep, and giraffe. He intended this to suggest a phylogenetic link between these animals: he identified that the fundamental difference between these bones in the three animals is their relative length and breadth, suggesting that they are closely related by evolution. In order to test this, he scaled them to a common length of 100, and then calculated the breadth of the sheep's bone as two-thirds of the breadth of the ox's and the breadth of the giraffe's as one-third of the breadth of the ox's. By identifying corresponding points o, a, b, c on each bone (shown in Fig. 1), he was then able to measure oa, ob, oc, and oy.

Taking these measurements in his three examples, Thompson considered three transformations: ox to ox (the identity transformation), ox to sheep, and ox to giraffe. He then plot-

ted the corresponding lengths of the four sections, with the x coordinate being the lengths for the ox, and the y coordinate being the lengths for the second animal. By drawing a line through each of these sets of points, Thompson was able to show how to transform the cannon-bone of one of the animals into another in a relatively simple way. This is an example of an image registration using landmarks, using the transformation set $Diff([0, 1]) \times Sim$ (i.e. an arbitrary rescaling of the four lengths independently after a one-dimensional horizontal scaling; here Sim is the group of similarities $x \mapsto ax + b$).

The question is whether or not this is the best choice of space of transformations. This question has been asked before, notably by Milnor in a talk he gave to the Institute for Advanced Studies [16]. He used an invariant (the one-dimensional cross-ratio, which is the restriction of the two-dimensional cross-ratio to real numbers) based on two sets of four marked points: o, a, b, c and a, b, c, y, as well as the two-dimensional cross-ratio. Milnor found that the one-dimensional cross-ratios were nearly equal, but the two-dimensional ones were not; he then registered the bones using Thompson's landmarks, using the groups $\mathrm{PSL}(2, \mathbb{C})$ and $\mathrm{PSL}(3, \mathbb{R})$. However, in neither case was the registration satisfactory. Instead, we consider the product group $\mathrm{PSL}(2, \mathbb{R}) \times \mathrm{PSL}(2, \mathbb{R})$, which we choose because the one-dimensional cross-ratio is preserved.

We also consider the simpler group $Sim \times \mathrm{PSL}(2, \mathbb{R})$, which scales the bones linearly in the x direction and uses the linear-fractional transformation in y. Note that this is a subgroup of the first group, and that in both groups, we can compose the transformations, and hence an approximation (because of the inexact matching) of the ox to giraffe transformation can be made by composing ox–sheep and sheep–giraffe; this is shown experimentally in Fig. 2.

The values of D (Eq. (5)) in the registrations of the bones in (first) $\mathrm{PSL}(2, \mathbb{R}) \times \mathrm{PSL}(2, \mathbb{R})$ and (second) its subgroup $Sim \times \mathrm{PSL}(2, \mathbb{R})$ are: for ox–sheep, 0.6433 and 0.6244, for ox–giraffe 0.9025 and 0.9005, and for sheep–giraffe 0.8873 and 0.8701, so there is very little difference between the two groups. These values indicate how well the images match; for example 90% of the ox is matched to the giraffe by registration. The registration between the ox and sheep leaves a relatively high residual, but it can be seen in Fig. 3 that the two images overlap well, although there are some differences at the boundaries of the images.

In Table 1, we compare the residuals when using each of the images as the source and the others as the targets for both groups. These show that the two groups provide very similar results for this data set; we will shortly apply model selection to choose between the groups. However, note that while the choice of source has some correlation with the residual, it is unclear precisely how to interpret this. We will demonstrate registration to the moving reference shortly, but in order to enable comparison with the results of Thompson shown in Fig. 1 we first use the ox as the source. Figure 3 shows the outline of the ox bone after transformation to the sheep and giraffe bones using both the group and subgroup. The group and subgroup differ only in their allowable x transformations, and Fig. 3 shows that the transformations found in the two groups are largely identical, being uniform scaling.

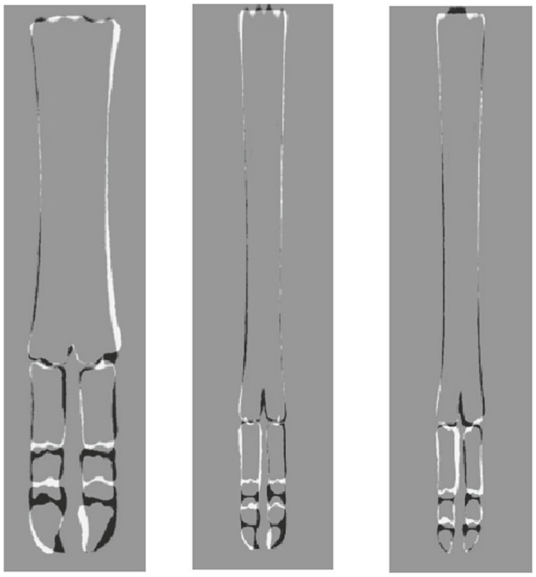

Fig. 2 Discrepancy between **a** transformed ox (to sheep) and sheep, **b** transformed ox (to giraffe) and giraffe, **c** transformed sheep (to giraffe) and giraffe using the group $\mathrm{PSL}(2, \mathbb{R}) \times \mathrm{PSL}(2, \mathbb{R})$

Figure 4 shows the progression of the registration in the subgroup $Sim \times \mathrm{PSL}(2, \mathbb{R})$ when the reference image is chosen to be the pixel-wise mean of the set of images rather than one of the images. The figure shows the mean of the three, which progresses from being blurry to clear as the registration runs. The residuals after the five iterations (although there is actually very little change in values after just three iterations) are 1116 for the ox, 751 for the sheep, and 1119 for the giraffe, while the values of D are 0.79, 0.90, and 0.90, respectively; the latter are very similar to those of registering to the fixed true image of the ox bone.

This first finite-dimensional registration supports Thompson's idea of simple transformation between related forms, but we still need to decide which group to use for registration. In the present example, the models are nested, so the residuals from registration in the group must necessarily be at least as small as the residuals from registration in the subgroup. The question becomes how much reduction in residual is needed to justify a more complex model. In this example the difference between the two AIC values ≈ 2, and so the simpler model is preferred (in both the registration to the ox, and the registration to the mean). The results for the simpler $Sim \times \mathrm{PSL}(2, \mathbb{R})$ model, registered to the ox, are shown

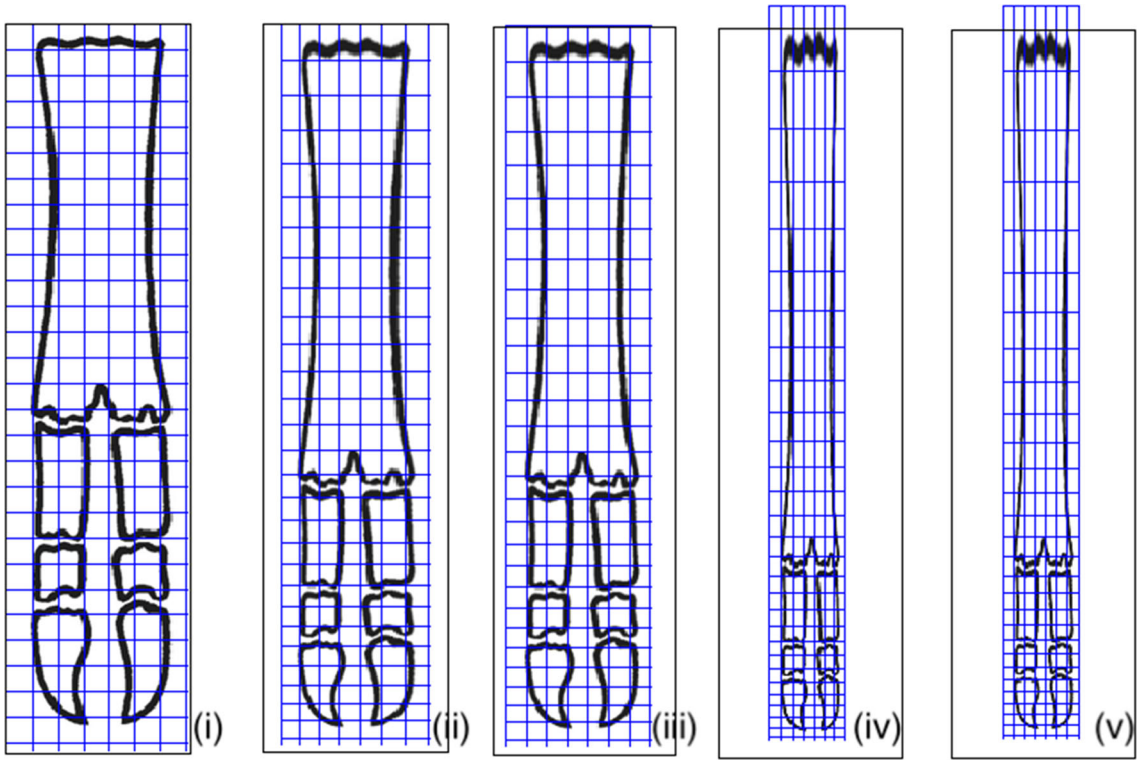

Fig. 3 The ox cannon-bone with a super-imposed rectangular grid (**a**), and then the transformation of the ox bone and the rectangular grid to the sheep using (**b**) the group $\mathrm{PSL}(2,\mathbb{R}) \times \mathrm{PSL}(2,\mathbb{R})$ and (**c**) subgroup $Sim \times \mathrm{PSL}(2,\mathbb{R})$, followed by the transformation to the giraffe bone using the same group (**d**) and subgroup (**e**)

in Fig. 5; they are very close to Thompson's registration in Fig. 1.

Of course, it would be possible to obtain a better registration using a larger group, but the most likely groups are infinite dimensional, namely $Diff(\mathbb{R}) \times Diff(\mathbb{R})$ and $Diff(\mathbb{R}^2)$. Since 5 parameters have already explained 90% of the difference between the bones, it seems unlikely from the point of view of model selection that such drastically more complex models would be preferred. However, in a more detailed study (ideally with a much larger data set) it would be possible to perform statistical model selection if an appropriate distribution for the data could be identified. An alternative method of describing image deformations, such as using B-splines [17] could also be considered; it would allow larger sets of deformations to be applied, albeit at the cost of losing a natural metric on the space of transformations and the qualitative properties inherent to Lie transformation groups.

Turning now to interpolation between the bones in the Lie group, we label the three transformations we have computed as φ_i, ($i = 1$: ox–ox, $i = 2$: ox–sheep, $i = 3$: ox–giraffe), which in full generality can be written as:

$$\varphi_i(x, y) = \left(s_i x + t_i, \frac{a_i y + b_i}{c_i y + d_i}\right).$$

The translation part of $\varphi_i(x, \cdot)$ can be ignored, because it only relates to the position of the image in the plane, hence $t_i = 0$. Furthermore, the boundary conditions ($\varphi_i(x, 0) = (x, 0)$ and $\varphi_i(x, 1) = (x, 1)$) and the fact that the determinant

Table 1 Residuals after registration (as a percentage of the original sum-of-squares pixel intensities differences) when registering the cannon-bones in the product groups $\mathrm{PSL}(2,\mathbb{R}) \times \mathrm{PSL}(2,\mathbb{R})$ and $Sim \times \mathrm{PSL}(2,\mathbb{R})$

Source | Target	Ox		Sheep		Giraffe	
x-transform	$\mathrm{PSL}(2,\mathbb{R})$	Sim	$\mathrm{PSL}(2,\mathbb{R})$	Sim	$\mathrm{PSL}(2,\mathbb{R})$	Sim
Ox	0	0	30%	31%	5%	5%
Sheep	42%	45%	0	0	7%	9%
Giraffe	18%	19%	18%	19%	0	0

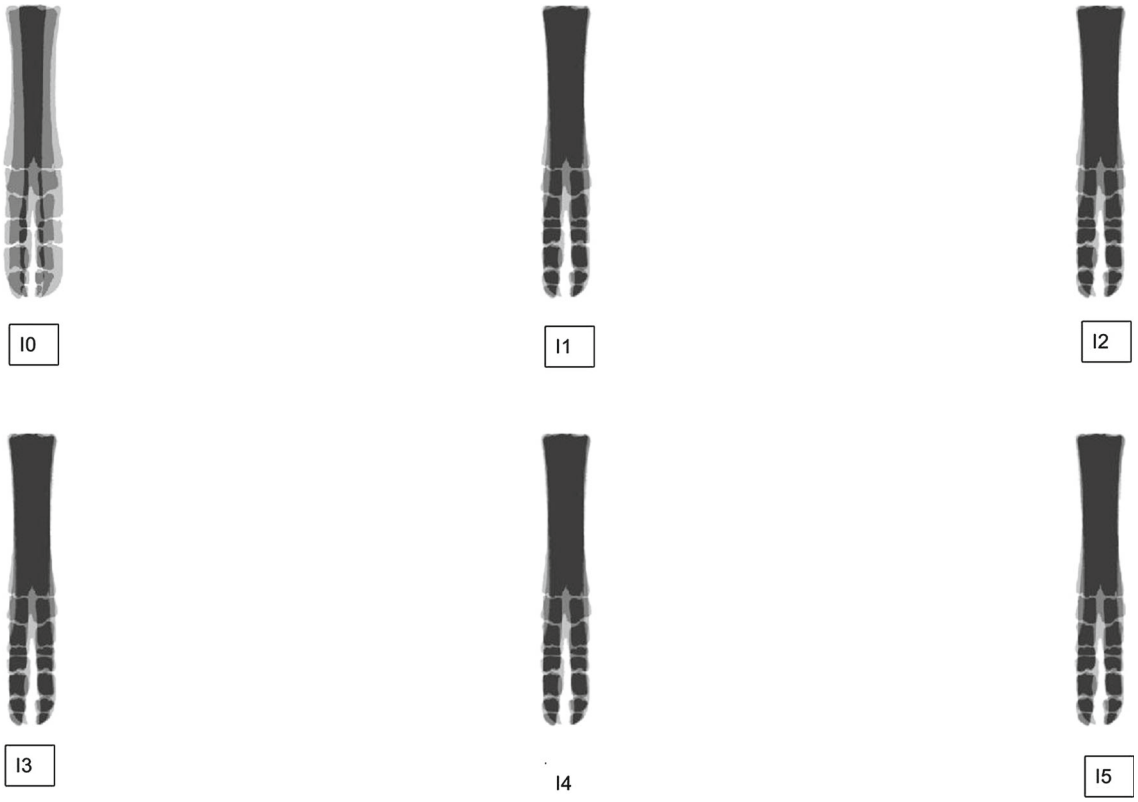

Fig. 4 Iterative improvement of the pixel-wise mean of the three images, which is used as the reference image, as the registration iterates. $I0$ is before any registration. After three iterations the algorithm has largely converged

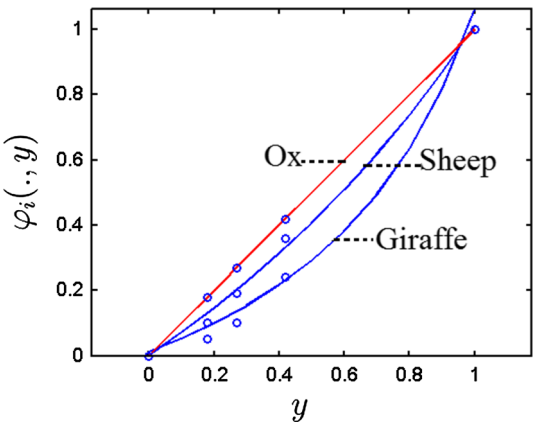

Fig. 5 Vertical dimension of the ox, sheep and giraffe cannon-bones as functions of the ox vertical dimension, as determined from an image registration in $Sim \times \mathrm{PSL}(2, \mathbb{R})$. These results may be compared to Thompson's 'by hand' registration in Fig. 1

of the transformation =1 mean that:

$$\frac{a_i(0) + b_i}{c_i(0) + d_i} = 0 \Rightarrow b_i = 0,$$
$$\det \begin{pmatrix} a_i & 0 \\ c_i & d_i \end{pmatrix} = 1 \Rightarrow d_i = \frac{1}{a_i},$$
$$\frac{a_i(1) + 0}{c_i(1) + \frac{1}{a_i}} = 1 \Rightarrow c_i = a_i - \frac{1}{a_i}.$$

Hence:

$$\varphi_i(x, y) = \left(s_i x, \frac{a_i y}{(a_i - \frac{1}{a_i})y + \frac{1}{a_i}} \right),$$

which corresponds to transformations in the product group of the linear-fractional transformations $\mathrm{SL}(2, \mathbb{R}) \times \mathrm{SL}(2, \mathbb{R})$ by:

$$\varphi_i(x, y) \leftrightarrow \left(\begin{pmatrix} \sqrt{s_i} & 0 \\ 0 & \frac{1}{\sqrt{s_i}} \end{pmatrix}, \begin{pmatrix} a_i & 0 \\ a_i - \frac{1}{a_i} & \frac{1}{a_i} \end{pmatrix} \right).$$

We identify φ_i with the corresponding matrix. The general form of a Lie algebra element corresponding to φ_i is then:

$$v_i = \log(\varphi_i) = \left(\begin{pmatrix} \beta_i & 0 \\ 0 & -\beta_i \end{pmatrix}, \begin{pmatrix} \alpha_i & 0 \\ 2\alpha_i & -\alpha_i \end{pmatrix} \right).$$

Figure 6 shows the data $\{(\beta_i, \alpha_i), i = 1, 2, 3\}$ in $\mathbb{R}^2$ and the fitted line in the Lie group L (which is not the

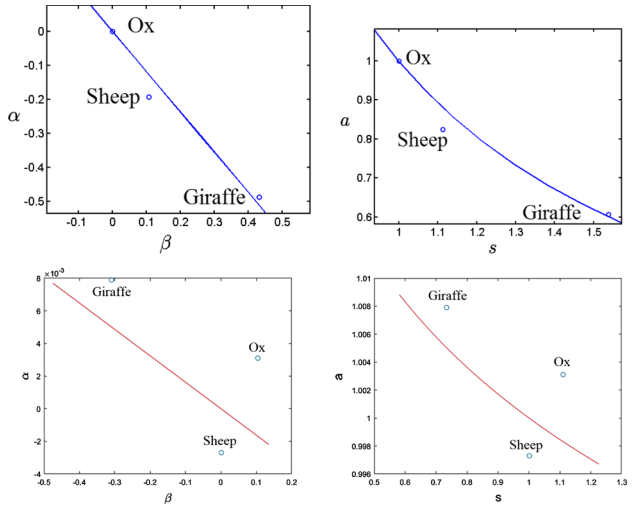

Fig. 6 The three data points and (*left*) a line L fitted through them *in the Lie algebra* of $Sim \times \mathrm{PSL}(2, \mathbb{R})$, (*right*) the curve $\exp(L)$ in group parameters in the group for *top:* registration to the ox, and *bottom:* registration to the mean

least-squares line fitted to the data points) as well as the corresponding curve in the group $Sim \times \mathrm{PSL}(2, \mathbb{R})$, for registration to both the ox and the mean. We performed a leave-one-out test on the data by computing the interpolation between two of the bones and comparing the interpolated version of the third with the ground truth. The mean-square error of both is smaller: for the first 0.018, while for the second it is 0.08; in both case, it suggests that the method is reliable. Figure 7 shows some equally spaced interpolations between the three sample bones. Since there are only three data points it is not surprising that the data fit very well. However, the interpolations also look like bones.

We now move on to consider a different data set, which is a product of growth rather than evolutionary change.

3.2 An Example from Petukhov: Human Skull Growth

A related investigation of biological change was considered by Petukhov, who in [18], in the spirit of the Erlangen project, looked at the role of non-Euclidean symmetry groups in biological growth. One of the examples used by Petukhov is the growth of a human skull. We used the data shown in Fig. 8, which is a craniofacial growth series of an immature human from [4], to validate Petukhov's work. The method of roentgencephalometry was used to align the set of X-ray tracings based on a physiological landmark (the sella located in the middle of pituitary fossa) and orientates the images on the line that stretches from sella to nasion; for more information, see [8]. Our first question is whether or not Petukhov's modelling choice (the Möbius group) is a good one. We study this by considering other modelling choices as well, and testing them on a small data set. While the data set is small, it suffices for demonstration purposes.

Following the discussion in Sect. 2.3 concerning time-series data, the third (middle) skull is taken as the source and registered with the others using $\mathrm{PSL}(2, \mathbb{C})$ and $\mathrm{PSL}(3, \mathbb{R})$. Figure 9 shows the discrepancy between transformed source and the targets. It can be seen that $\mathrm{PSL}(2, \mathbb{C})$ has produced a better matching of the jaw of the human skull than $\mathrm{PSL}(3, \mathbb{R})$, while $\mathrm{PSL}(3, \mathbb{R})$ has done a better job of matching the head shape. Looking at the quality of the registrations, it seems that either group provides a reasonable model of human growth. Using AIC, $\mathrm{PSL}(2, \mathbb{C})$ is preferred as it has fewer parameters, but given the shortage of data, we progress by using both groups, and seeking curves in their Lie algebras that fit the growth of the skull.

Let φ_i, ψ_i, $i = 1, 2, 3, 4, 5$ be the transformations that were obtained by registration in $\mathrm{PSL}(2, \mathbb{C})$ and $\mathrm{PSL}(3, \mathbb{R})$, respectively (note that φ_3 and ψ_3 will be the identity element since they are registering the image to itself). The groups $\mathrm{PSL}(2, \mathbb{C})$ and $SL(2, \mathbb{C})$ are homomorphic, as are $\mathrm{PSL}(3, \mathbb{R})$ and $SL(3, \mathbb{R})$. We use this relationship by fitting a curve in the matrix group and then mapping the curves homomorphically into the projective group.

We consider standard bases for the Lie algebras $\mathfrak{sl}(2, \mathbb{C})$ and $\mathfrak{sl}(3, \mathbb{R})$ and the standard Euclidian metric in those bases. In the standard basis of $\mathfrak{sl}(2, \mathbb{C})$, we omit the vectors relating to rotation and translation, because the 'shape' of an object (in the sense of shape space [7]) is invariant under these actions. In $\mathfrak{sl}(3, \mathbb{R})$, we ignore only translation. This projection is equivalent to using a metric on the Lie algebra that is zero in the similarity components.

Now, we need to choose a model to fit the data in the Lie algebra. A polynomial of degree four can be fitted perfectly through the data, but we prefer a simpler model, which will not overfit (this point will be considered further later). One-parameter subgroups are the simplest model that can describe the points in the algebra, and hence the relationship between the human skulls at different ages. They correspond to straight lines passing through the origin in the Lie algebra; in order to add one extra degree of freedom, we also considered straight lines that do not pass through the origin.

Let Γ_1 and Σ_1 be the lines that pass through the origin, and Γ_2 and Σ_2 be the lines that do not, in $\mathfrak{sl}(2, \mathbb{C})$ and $\mathfrak{sl}(3, \mathbb{R})$, respectively, with the Euclidian metric on the subspace of the Lie algebra. The residuals for the fitted lines are 0.0123 for the non-0 intercept line and 0.0254 for the 0 intercept using $\mathfrak{sl}(2, \mathbb{C})$ and 0.0376 and 0.0552, respectively, using $\mathfrak{sl}(3, \mathbb{R})$.

The group exponential maps of these lines ($\exp(\Gamma_1)$, $\exp(\Gamma_2)$, $\exp(\Sigma_1)$ and $\exp(\Sigma_2)$) are the curves fitting through the points representing the skulls.

For the Möbius registration, we now provide the relevant computations; they are very similar for the case of $\mathrm{PSL}(3, \mathbb{R})$.

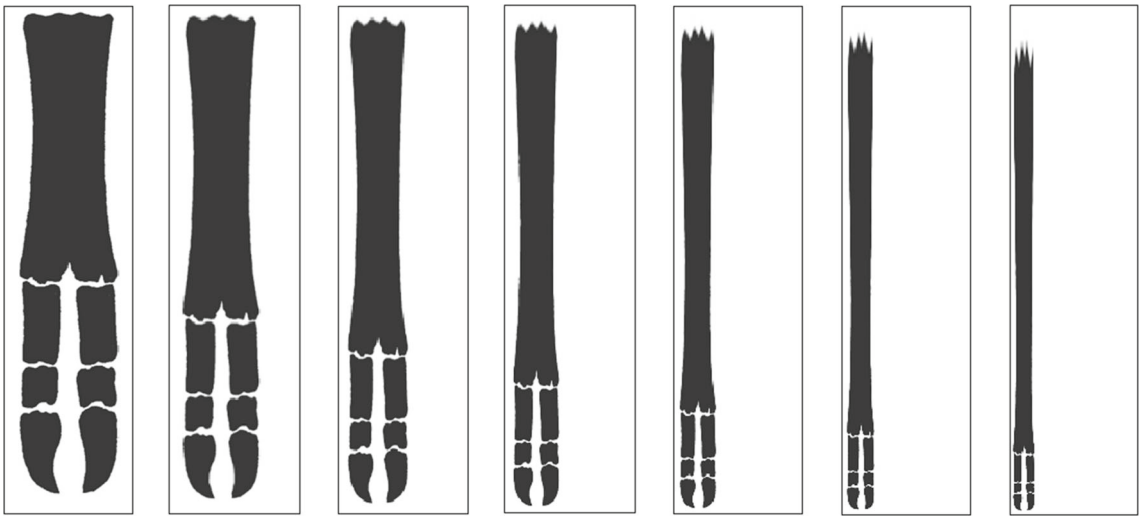

Fig. 7 Transformation of ox foot by $\exp(L(t))$; $E(t)$ is the line fitted through $\log(\varphi_i)$

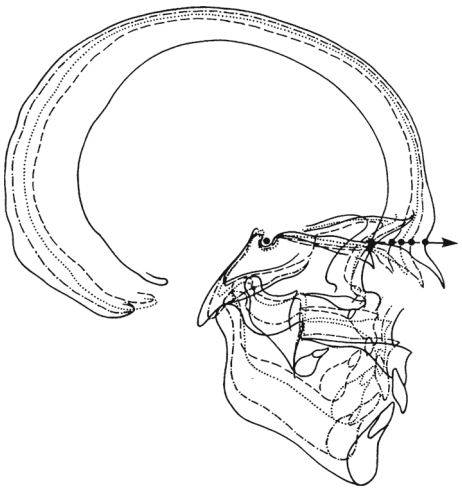

Fig. 8 Traditional craniofacial growth series based on longitudinal cephalometric radiographs, taken from [8]

Let $\varphi_k(z) = \frac{(S_k+iR_k)z+T_k^1+iT_k^2}{(A_k+iB_k)z+D_k^1+iD_k^2}$ be a transformation in PSL$(2, \mathbb{C})$ that is obtained by the registration in the Möbius group, with corresponding matrix in $SL(2, \mathbb{C})$:

$$\begin{pmatrix} S_k + iR_k & T_k^1 + iT_k^2 \\ A_k + iB_k & D_k^1 + iD_k^2 \end{pmatrix}.$$

Then for $\zeta_k = \log(\varphi_k)$:

$$\zeta_k = \begin{pmatrix} s_k + ir_k & t_k^1 + it_k^2 \\ a_k + ib_k & -(s_k + ir_k) \end{pmatrix}. \tag{8}$$

We consider a standard basis for the Lie algebra, $\mathfrak{sl}(2, \mathbb{C})$ (matrices in $\mathfrak{sl}(2, \mathbb{C})$ have trace equal to zero):

$$v_1 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, v_2 = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, v_3 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$
$$v_4 = \begin{pmatrix} 0 & i \\ 0 & 0 \end{pmatrix}, v_5 = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, v_6 = \begin{pmatrix} 0 & 0 \\ i & 0 \end{pmatrix}. \tag{9}$$

We calculate the norm of sub-transformations: scale, rotation, translation and nonlinear transformation $(\frac{z}{cz+1})$ in φ_k by calculating the norm of their corresponding vectors in the Lie algebra. For example, if $\zeta_k = a_1v_1 + a_2v_2 + a_3v_3 + a_4v_4 + a_5v_5 + a_6v_6$, then the norm of scale is $tr((a_1v_1)(a_1v_1)^*)^{\frac{1}{2}}$. The norms of the rotation and translation parts are very close to zero, and the 'shape' of an object is invariant under rotation and translation [7]. Therefore, we ignore these parts of ζ_k, and so $\zeta_k = \log(\varphi_k) = \begin{pmatrix} s_k & 0 \\ a_k + ib_k & -s_k \end{pmatrix}$, and $\varphi_k = \begin{pmatrix} S_k & 0 \\ A_k + iB_k & \frac{1}{S_k} \end{pmatrix}$.

The left of Fig. 10 shows the data (s_k, a_k, b_k) and the fitted lines Γ_1 and Γ_2, while the right shows the data (S_k, A_k, B_k) and the curves $\exp(\Gamma_1)$, $\exp(\Gamma_2)$, respectively. These are the best fitted models in two different representational spaces of this data.

It is not always easy to see visually how well the linear model fits the data. We therefore calculated the *coefficient of determination* (R^2) and the *adjusted coefficient of determination* $(\bar{R}^2)$. The following table gives the values of R^2 and $\bar{R}^2$ of the fitted lines in $\mathfrak{sl}(2, \mathbb{C})$ and $\mathfrak{sl}(3, \mathbb{R})$. As can be seen in the table, the value of $\bar{R}^2$ for the three models is 0.8, which suggests that the fitted lines represent the data well. There is no significant difference between the fits in the two groups.

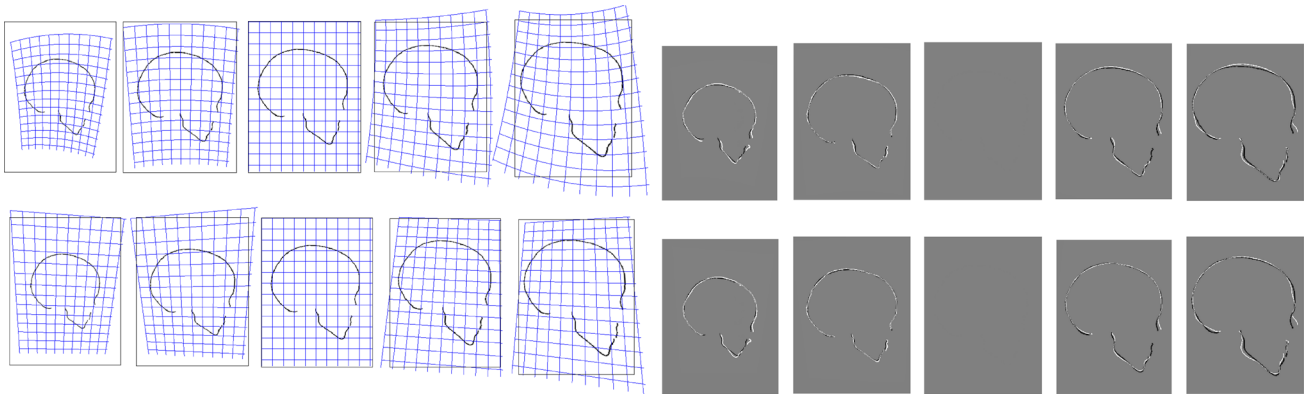

Fig. 9 *Left:* Grids showing the deformation applied to the middle skull in the time sequence in order to register it to each skull in groups PSL(2, $\mathbb{C}$) (top) and PSL(3, $\mathbb{R}$) (below). *Right:* The discrepancy between the middle skull in the sequence after registration to each skull in the two groups (hence, the perfect registration in the middle column is the registration of the skull to itself). It can be seen that the transformations do seem to follow a trajectory, with the degree of deformation increasing through the sequence

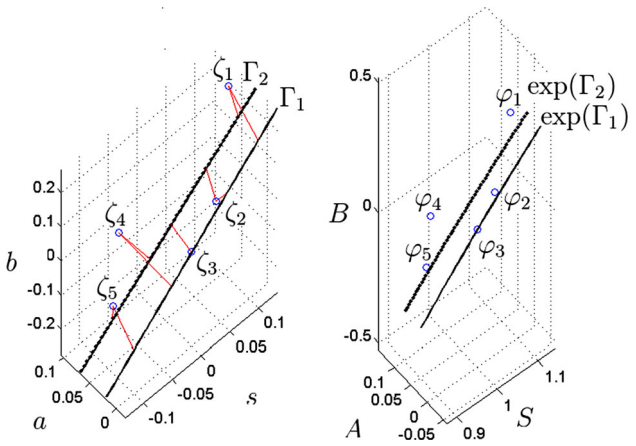

Fig. 10 Fits of the skull registrations in PSL(2, $\mathbb{C}$). On the left, the points correspond to the positions of the 5 images in the three-dimensional parameter space remaining once rotation and translation have been taken out. Γ_1 is a line passing through the origin fitted through ζ_k, while Γ_2 is not constrained to pass through the origin. They map by the group exponential into the group, producing the parameters plotted on the right

$\mathfrak{sl}(2,\mathbb{C})$	R^2	$\bar{R}^2$	$\mathfrak{sl}(3,\mathbb{R})$	R^2	$\bar{R}^2$
Γ_1	0.79	0.73	Σ_1	0.85	0.80
Γ_2	0.90	0.80	Σ_2	0.90	0.80

So far, we have shown that the linear models represent the data well. Therefore, we generate some possible human skulls by interpolation using both $\exp(\Gamma_2(t))$ and $\exp(\Sigma_2(t))$ and applying the transformation to the source (middle skull) in order to generate a new skull; these are shown in the middle and right of Fig. 11, respectively. Both provide reasonably good matches to the data.

In order to test whether or not the straight line models did indeed provide good matches, we also increased the degree of

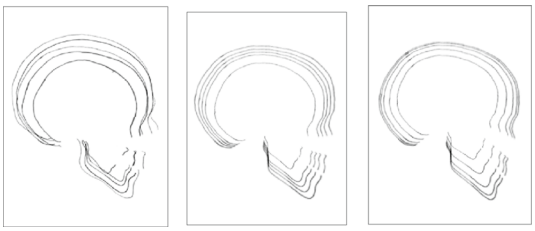

Fig. 11 **a** The original skulls. **b** Generating skulls using the curve $\exp(\Gamma_2)$ in PSL(2, $\mathbb{C}$). **c** Generating skulls using the curve $\exp(\Sigma_2)$ in PSL(3, $\mathbb{R}$)

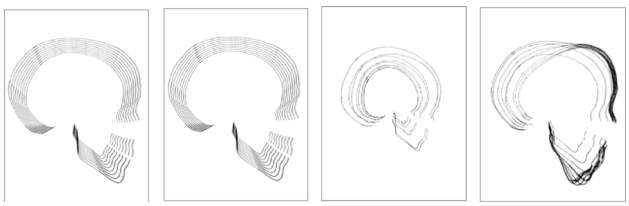

Fig. 12 Skulls generating by the curves in PSL(2, $\mathbb{C}$). Exponential of the fitted **a** line with zero intercept, **b** line with nonzero intercept, **c** quadratic curve, **d** cubic curve in $\mathfrak{psl}(2,\mathbb{C})$

the model, using a quadratic and cubic function to fit the data with the Möbius transformation. The residuals of the four models are: cubic: 0, quadratic: 0.0037, line (nonzero intercept): 0.0123, line (zero intercept): 0.0254. As the degree of the model increases, so their residual decreases, but the predictions made by the models gets worse, a clear sign of overfitting; this is shown in Fig. 12, where the interpolated skulls along the growth curve are generated by the four models. It can be seen that the cubic and quadratic models provide very unlikely looking skulls. Therefore, it seems that cubic and quadratic are not good models to describe human skull growth. We used a leave-one-out method to check the reliability of the models in the prediction of new data. We compute

the Mean Squared Error (MSE) of the leave-one-out errors for the models, which suggest that the straight line method is the most reliable: cubic: 0.0079, quadratic: 0.0214, line (nonzero intercept): 0.0085, line (zero intercept): 0.0119.

Using the residuals of the registration, we computed the AIC for PSL(2, $\mathbb{C}$) as 16.2, and for PSL(3, $\mathbb{R}$) as 17.8. This suggests that there is very little difference in the information provided by the two models.

4 Conclusion

The simple transformations advocated by Thompson in 'On Growth and Form' have metamorphosed into diffeomorphic image registration. While this has many useful features, there are times when the added complexity does not provide benefits. First, the registration itself can provide explicit information about the growth process itself, and second, the use of simple groups gives access to additional statistical analysis methods. In addition, it is substantially quicker and easier to perform registrations in low-dimension groups.

In this paper, we have developed a method based on the planar Lie groups that provides access to interpolation methods in the corresponding Lie algebra to the group where registration is performed, and demonstrated that this group can be selected using statistical model selection methods. We have demonstrated this on two simple data sets which are classical examples of evolutionary development (Thompson's cannon-bones of three hoofed mammals) and biological growth (the human skulls considered by Petukhov). The results using these small and simple data sets are very promising and show that our approach has benefits as an initial step for understanding growth processes. Our examples have considered three groups, but the methods that we have demonstrated will work for any of the 28 planar Lie groups, and also to products of them.

In this paper, we have provided case studies based on single sets of examples, and there is a need to perform larger-scale experiments based on more data. This will enable the model selection approach to be used to identify the correct modelling group far more accurately. Two benefits of doing this will be to further investigate the work of Thompson and Petukhov: in the more than 100 years since Thompson first proposed his theory of the method of transformations (in [23]) it has not been fully tested, as was observed by [2]. Further, Petukhov postulates that the Möbius group forms the basis of many examples of the development process; our method will also allow this to be verified.

Another possible modelling choice is in the metric that is used. We can consider different metrics and how they can be chosen appropriately. This is a question of data fitting, and is somewhere where the low-dimensional representations implicit in this approach have significant benefits, compared to trying to fit parameters for the metric in diffeomorphic image registration, which is extremely data intensive. Finally, we are also in the process of combining registrations in different groups to better understand the growth process, and how a variety of different processes can act simultaneously.

Acknowledgements This research was supported by the Royal Society of New Zealand Marsden Fund (Grant No. MAU0908).

References

1. Aho, K., Derryberry, D., Peterson, T.: Model selection for ecologists: the worldviews of AIC and BIC. Ecology **95**(3), 631–636 (2014)
2. Arthur, W.: D'Arcy Thompson and the theory of transformations. Nat. Rev. Genet. **7**(5), 401–406 (2006). https://doi.org/10.1038/nrg1835
3. Bookstein, F.L.: Morphometric Tools for Landmark Data: Geometry and Biology. Cambridge University Press, Cambridge (1997)
4. Broadbent, B.H., Golden, W.H.: Bolton Standards of Dentofacial Developmental Growth. Mosby, St. Louis (1975)
5. Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N.: Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI, pp. 297–304 (2009)
6. González-López, A., Kamran, N., Olver, P.J.: Lie algebras of vector fields in the real plane. Proc. Lond. Math. Soc. **64**(3), 339–368 (1992)
7. Kendall, D.G.: Shape manifolds, Procrustean metrics, and complex projective spaces. Bull. Lond. Math. Soc. **16**(2), 81–121 (1984)
8. Lele, S.R., Richtsmeier, J.T.: An Invariant Approach to Statistical Analysis of Shapes. CRC Press, Boca Raton (2001)
9. Lorenzi, M., Ayache, N., Pennec, X.: Schild's ladder for the parallel transport of deformations in time series of images. In: Information Processing in Medical Imaging (IPMI), pp. 463–474 (2011)
10. Lorenzi, M., Frisoni, G., Ayache, N., Pennec, X.: Mapping the effects of aβ 1- 42 levels on the longitudinal changes in healthy aging: hierarchical modeling based on stationary velocity fields. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI), pp. 663–670 (2011)
11. Lorenzi, M., Pennec, X.: Geodesics, parallel transport & one-parameter subgroups for diffeomorphic image registration. Int. J. Comput. Vis. **105**(2), 111–127 (2013)
12. Mani, V.R.S., Rivazhagan, S.: Survey of medical image registration. J. Biomed. Eng. Technol. **1**(2), 8–25 (2013)
13. Marsland, S., Shardlow, T.: Langevin equations for landmark image registration with uncertainty. SIAM J. Imag. Sci. **10**(2), 782–807 (2017)
14. Marsland, S., Twining, C., Taylor, C.: A minimum description length objective function for groupwise non-rigid image registration. Image Vis. Comput. **26**(3), 333–346 (2008)
15. Miller, M.I., Trouvé, A., Younes, L.: Hamiltonian systems and optimal control in computational anatomy: 100 years since D'Arcy Thompson. Annu. Rev. Biomed. Eng. **17**(1), 447–509 (2015). https://doi.org/10.1146/annurev-bioeng-071114-040601
16. Milnor, J.: Geometry of growth and form: commentary on D'Arcy Thompson (2010). A talk at the Institute for Advanced Study in Princeton. Available at https://www.ias.edu/ideas/2010/milnor-geometry-of-growth
17. Modersitzki, J.: Fair: Flexible Algorithms for Image Registration. SIAM, New Delhi (2009)

18. Petukhov, S.V.: Non-Euclidean geometries and algorithms of living bodies. Comput. Math. Appl. **17**(4), 505–534 (1989)
19. Rao, A., Chandrashekara, R., Sanchez-Hortiz, G., Mohiaddin, R., Aljabar, P., Hajnal, J., Puri, B., Rueckert, D.: Spatial transformation of motion and deformation fields using nonrigid registration. IEEE Trans. Med. Imag. **23**, 1065–1076 (2004)
20. Sommer, S., Arnaudon, A., Kuhnel, L., Joshi, S.: Bridge simulation and metric estimation on landmark manifolds. In: Graphs in Biomedical Image Analysis, Computational Anatomy and Imaging Genetics, Lecture Notes in Computer Science, pp. 79–91. Springer (2017)
21. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: a survey. IEEE Trans. Med. Imag. **32**(7), 1153–1190 (2013)
22. Tang, Z., Wu, Y., Yong, F.: Groupwise registration of MR brain images with tumors. Phys. Med. Biol. **62**(17), (2017)
23. Thompson, D.W.: Morphology and mathematics. Trans. R. Soc. Edinb. L **857–895**, (1915)
24. Thompson, D.W.: On Growth and Form, 2nd edn. Cambridge University Press, Cambridge (1942)
25. Trouve, A., Vialard, F.X.: Shape splines and stochastic shape evolutions: a second order point of view. Q. Appl. Math. **70**(2), 219–251 (2012)
26. Wu, G., Wang, Q., Jia, H., Shen, D.: Feature-based groupwise registration by hierarchical anatomical correspondence detection. Hum. Brain Mapp. **33**, 253–271 (2012)
27. Younes, L.: Jacobi fields in groups of diffeomorphisms and applications. Q. Appl. Math. **65**(1), 113–134 (2007)
28. Younes, L., Qiu, A., Winslow, R., Miller, M.: Transport of relational structures in groups of diffeomorphisms. J. Math. Imag. Vis. **32**, 41–56 (2008)

Stephen Marsland is a Professor of Mathematics and Data Science at Victoria University of Wellington, New Zealand. He studied at Oxford University and the University of Manchester, graduating with a Ph.D. in 2002. He moved to Massey University, New Zealand, in 2004, and then to Victoria University of Wellington in 2018. He is interested in geometry and its applications, particularly machine learning, signal processing (birdsong analysis), and image analysis.

Robert I. McLachlan is a Distinguished Professor in Applied Mathematics at Massey University, New Zealand. He studied at Caltech with Herb Keller and worked at the University of Colorado, Boulder, and ETH Zurich, before joining Massey University in 1994. He is interested in applications of geometry to all areas of applied mathematics, including geometric numerical integration, PDEs, dynamical systems, and image processing.

Raziyeh Zarre is a Senior Tutor in Mathematics at Waikato University, New Zealand. She studied at Massey University in Palmerston North, New Zealand. Her research interests are in applications of mathematics to image analysis.