J Math Imaging Vis (2009) 34: 222–233
DOI 10.1007/s10851-009-0143-4

On the Topology and Geometry of Spaces of Affine Shapes

David Groisser · Hemant D. Tagare

Published online: 31 March 2009

Abstract We define the *space of affine shapes of k points in* $\mathbf{R}^n$ to be the topological quotient of $(\mathbf{R}^n)^k$ modulo the natural action of the affine group of $\mathbf{R}^n$. These spaces arise naturally in many image-processing applications, and despite having poor separation properties, have some topological and geometric properties reminiscent of the more familiar Procrustes shape spaces Σ_n^k in which one identifies configurations related by an orientation-preserving Euclidean similarity transformation. We examine the topology of the connected, non-Hausdorff spaces Sh_n^k in detail. Each Sh_n^k is a disjoint union of naturally ordered strata, each of which is homeomorphic in the relative topology to a Grassmannian, and we show how the strata are attached to each other. The top stratum carries a natural Riemannian metric, which we compute explicitly for $k > n$, expressing the metric purely in terms of "pre-shape" data, i.e. configurations of k points in $\mathbf{R}^n$.

Keywords Shape space · Affine shape

1 Introduction

Procrustes shape spaces Σ_n^k of k labeled points ("landmarks") in $\mathbf{R}^n$ [5, 11, 12, 22] have proven useful in many statistical and non-statistical applications. Several authors have suggested an extension of Procrustes shape spaces to affine-shape spaces [15, 18, 21, 23], which we denote Sh_n^k. These are spaces of configurations of k labeled point-landmarks in $\mathbf{R}^n$ modulo the action of $\mathrm{Aff}(n)$, the affine group of $\mathbf{R}^n$. Affine-shape spaces arise very naturally in many image-processing applications. Assuming a pinhole camera model, images are formed under a perspective transformation. If an object is distant from the camera, the perspective transformation can be approximated closely by an affine transformation [8]. Hence, many problems involving object recognition or shape statistics of object landmarks in images are naturally posed in an affine-shape space. For example, creation of image mosaics and panoramas from individual images [15, 24], the recognition of object boundaries and silhouettes [1–3, 14, 25, 26], and the estimation of structure from motion [4] can benefit from the mathematical framework of affine shapes [10]. Statistical analysis on manifolds [9, 19] that attempts to compensate for affine transforms can also benefit from this analysis. In medical image analysis, affine transformations are often used to model the fitting of an anatomical template (or an atlas) to an image [6, 13, 20]. By creating a distance in an affine-shape space of landmarks that are common to the template and the image, it should be possible to evaluate the quality of the assumption that the affine transform is suitable for the task.

Motivated by these applications we turn to analyzing affine-shape spaces. Some properties of these spaces are already known or are mentioned by other authors [15, 18, 21, 23]. In particular, it is known that each affine-shape space is a disjoint union of a naturally ordered strata, each of which is homeomorphic in the relative topology to a Grassmannian. However, these spaces are non-Hausdorff, and their topology does not appear to have been analyzed

D. Groisser
Department of Mathematics, University of Florida, Gainesville, FL 32611-8105, USA
e-mail: groisser@math.ufl.edu

H.D. Tagare (✉)
Department of Diagnostic Radiology, Department of Electrical Engineering, Yale University, New Haven, CT 06520, USA
e-mail: hemant.tagare@yale.edu

in detail,[1] which we do in this paper. The top stratum also carries a natural Riemannian metric which provides a distance between affine-shapes. This intrinsic Riemannian metric coincides (up to a constant factor) with the extrinsic Riemannian metric used in [18], but in this paper we provide explicit formulas for the metric directly in terms of pre-shape data rather than in terms of an embedding of the Grassmannian into a Euclidean space.

In spite of the differences between the affine-shape spaces Sh_n^k and Procrustes shape spaces Σ_n^k (the former are not even Hausdorff, while the latter are not just Hausdorff, but cell complexes; see [12]), the Sh_n^k have some topological and geometric properties that are reminiscent of the Σ_n^k. At the topological level, both carry a natural stratification by the dimension of the affine span of a configuration that represents a given shape (see [12] for the Procrustes case). There are also natural embeddings $\Sigma_n^k \hookrightarrow \Sigma_n^{k+1}$ and $Sh_n^k \hookrightarrow Sh_n^{k+1}$ given by appending to a configuration its centroid. These stratifications and embeddings are preserved by the natural projections $\Sigma_n^k \to Sh_n^k$. At the geometric level, the way in which we produce a metric on the top stratum of Sh_n^k is by Riemannian submersion, analogously to what one does on the nonsingular part of Σ_n^k. Later (Remark 4.3) we will see that there is another geometric similarity. Some additional comparisons between the Sh_n^k and Σ_n^k are given in the latter part of Sect. 4.

2 Notation and Preliminaries

For $n \geq 1$ let $GL(n) = GL(n, \mathbf{R})$ be the group of invertible $n \times n$ real matrices, and $GL_+(n)$ the subgroup of matrices with positive determinant. To avoid extra notation, we will often identify elements of $GL(n)$ with the associated linear transformations of $\mathbf{R}^n$. We denote by $\mathrm{Aff}(n)$ the group of invertible affine transformations of $\mathbf{R}^n$, and by $\mathrm{Aff}_+(n)$ the subgroup group of $\mathrm{Aff}(n)$ whose elements preserve orientation. (Thus $\mathrm{Aff}(n)$ is generated by $GL(n)$ and translations; for $\mathrm{Aff}_+(n)$ the group $GL_+(n)$ replaces $GL(n)$.)

For $k \geq 1$ and $n \geq 0$, the *space of affine pre-shapes* of k labeled points in $\mathbf{R}^n$ is simply $(\mathbf{R}^n)^k := \mathbf{R}^n \times \mathbf{R}^n \times \cdots \times \mathbf{R}^n$, the space of lists of k points in $\mathbf{R}^n$. Any transformation $T : \mathbf{R}^n \to \mathbf{R}^n$ acts on such a list by simultaneous motion of the points in the list: $T(v_1, \ldots, v_k) := (T(v_1), \ldots, T(v_k))$. In particular, $\mathrm{Aff}(n)$ and $\mathrm{Aff}_+(n)$ act on $(\mathbf{R}^n)^k$. We define the *space of affine shapes* (respectively, *oriented affine shapes*) *of k points in* $\mathbf{R}^n$ to be the quotient space $Sh_n^k := (\mathbf{R}^n)^k/\mathrm{Aff}(n)$ (resp., $\widetilde{Sh}_n^k := (\mathbf{R}^n)^k/\mathrm{Aff}_+(n)$), endowed with the quotient topology. Note that in contrast to what is usually done for Procrustean shape spaces, we do not exclude lists $(v_1, \ldots, v_k)$ in which all the v_i are identical, which in each quotient Sh_n^k and $\widetilde{Sh}_n^k$ become a single point that we denote $\bar{\mathbf{0}}$. In the Procrustean case, it is necessary to exclude this singular point in order to get a nice shape-space, but in the case of affine shapes, this singular point fits quite systematically into the general picture we will present.

The first step in determining the topology of Sh_n^k and $\widetilde{Sh}_n^k$ is to quotient $(\mathbf{R}^n)^k$ by the translation subgroups of $\mathrm{Aff}(n)$ and $\mathrm{Aff}_+(n)$. Let $H^{n,k} := \{(v_1, \ldots, v_k) \in (\mathbf{R}^n)^k \mid \sum_i v_i = 0\}$, an $n(k-1)$-dimensional subspace of $(\mathbf{R}^n)^k$ that is preserved by the action of $GL(n)$ and but is not preserved by any nontrivial translation. For every $\mathbf{v} = (v_1, \ldots, v_k) \in (\mathbf{R}^n)^k$ there exists a unique translation $T_{\mathbf{v}}$ for which $\mathrm{Cen}(\mathbf{v}) := T_{\mathbf{v}}(\mathbf{v}) \in H^{n,k}$, and the map $\mathbf{v} \to T_{\mathbf{v}}$ is continuous. (The *centering map* Cen is simply the map $(v_1, \ldots, v_k) \mapsto (v_1 - v_0, \ldots, v_k - v_0)$, where v_0 is the "center of mass" $(\sum_i v_i)/k$.)

Thus there is a natural 1-1 correspondence between Sh_n^k and $H^{n,k}/GL(n)$, and similarly between $\widetilde{Sh}_n^k$ and $H^{n,k}/GL_+(n)$. Using the continuity of $\mathbf{v} \mapsto T_{\mathbf{v}}$, it is not hard to show that these correspondences are homeomorphisms. Thus

$$Sh_n^k \cong H^{n,k}/GL(n), \qquad \widetilde{Sh}_n^k \cong H^{n,k}/GL_+(n) \tag{2.1}$$

(throughout this paper, in the context of topological spaces "$\cong$" means "homeomorphism" or "is homeomorphic to").

The actions of $GL(n)$ and $GL_+(n)$ on $\mathbf{R}^n$ are not *proper* (there exist orbits that are not closed subsets of $\mathbf{R}^n$). Quotient spaces of non-proper group actions are notoriously ill-behaved. We recall the reason: suppose a group G acts on a topological space X, that $p \in X$ is a point whose orbit $\mathcal{O}_p$ is not closed in X, and that $q \in X$ lies in the closure of $\mathcal{O}_p$ but not in $\mathcal{O}_p$ itself. Then, letting $\bar{p}, \bar{q}$ denote the images of p, q in the quotient X/G, every open neighborhood of $\bar{q}$ contains $\bar{p}$. In particular X/G is not Hausdorff. This general fact applies to spaces of affine shapes (or oriented affine shapes) in $\mathbf{R}^n$ for all $n \geq 1$. As we will see later, every shape $\bar{p}$ represented by a pre-shape p contained in a proper affine

[1] As we shall see, Sh_n^k can be naturally identified with a quotient space that arises from a "bad" linear action of the group $GL(n, \mathbf{R})$. Understanding such quotients, at least when the group is $SL(n, \mathbf{C})$ acting on complex projective space, is a central problem in Geometric Invariant Theory, a subfield of algebraic geometry. These quotients arise in constructions much more general than the one in this paper; *cf.* [17]. Stratification is an omnipresent phenomenon, and the Geometric Invariant Theory literature shows that the strata carry beautiful structures—e.g. they are projective varieties, as are the strata of Sh_n^k. However, it is not easy to tell from the geometric-invariant theorists' elegant presentation of their framework, which uses tools and language from both algebraic and differential geometry, just which facts about Sh_n^k follow directly from what is in this literature. In particular it is not easy to extract information about how the strata are glued together, and the topology community does not seem to be aware of the work done by the algebraic geometers in this area. In this paper we provide a direct, self-contained derivation of the topology on Sh_n^k, without appealing to any abstract machinery with which workers in the shape-space field may be unfamiliar.

subspace of $\mathbf{R}^n$ is a singular point of shape space, in the sense that there are points $\bar{q}$ in shape space for which every open neighborhood of $\bar{p}$ contains $\bar{q}$.

To describe the nonsingular points of Sh_n^k and $\widetilde{Sh}_n^k$ we introduce the following notation: $\mathbf{R}_*^{n,k} := \{\mathbf{v} = (v_1, \dots, v_k) \in (\mathbf{R}^n)^k \mid \nexists$ affine subspace of $\mathbf{R}^n$ of dimension $< n$ containing all the $v_i\}$, $H_*^{n,k} := H^{n,k} \cap \mathbf{R}_*^{n,k}$, $Sh_{n,*}^k := \mathbf{R}_*^{n,k}/\mathrm{Aff}(n)$, $\widetilde{Sh}_{n,*}^k := \mathbf{R}_*^{n,k}/\mathrm{Aff}_+(n)$. The reader may easily verify the following facts:

1. The following are equivalent: (i) $(v_1, \dots, v_k) \in \mathbf{R}_*^{n,k}$; (ii) for each fixed $j \in \{1, \dots, k\}$ the vectors $\{v_i - v_j\}_{i=1}^k$ span $\mathbf{R}^n$; and (iii) the vectors $\{v_i^{\mathrm{cen}}\}_{i=1}^k$ span $\mathbf{R}^n$, where $(v_1^{\mathrm{cen}}, \dots, v_k^{\mathrm{cen}}) = \mathrm{Cen}(\mathbf{v})$.
2. $H_*^{n,k} = \{(v_1, \dots, v_k) \in H^{n,k} \mid \nexists$ proper subspace of $\mathbf{R}^n$ containing all the v_i (equivalently, the $\{v_i\}$ span $\mathbf{R}^n$)$\}$
3. Since the affine span of the vectors $\{v_i\}$ comprising an element $\mathbf{v} \in \mathbf{R}^{n,k}$ is of dimension at most $k-1$, if $k \le n$ then the sets $\mathbf{R}_*^{n,k}$ and $H_*^{n,k}$ are empty.
4. The centering map induces homeomorphisms $Sh_{n,*}^k \cong H_*^{n,k}/GL(n)$, $\widetilde{Sh}_{n,*}^k \cong H_*^{n,k}/GL_+(n)$.

We will see in Sect. 3 that $Sh_{n,*}^k$ and $\widetilde{Sh}_{n,*}^k$ are smooth manifolds in a natural way; we think of them as the sets of "smooth points" of Sh_n^k and $\widetilde{Sh}_{n,*}^k$. We define the *singular set* $\mathrm{Sing}_n^k \subset Sh_n^k$ (respectively, $\widetilde{\mathrm{Sing}}_n^k \subset \widetilde{Sh}_n^k$) to be the complement of $Sh_{n,*}^k$.

Before determining the topology of Sh_n^k for general k, n it is instructive to work out a simple example:

Example 2.1 (Shapes in $\mathbf{R}^1$) Since $GL(1) = \mathbf{R}_* := \mathbf{R} - \{0\}$, if $k \ge 2$ we have $Sh_{1,*}^k = H_*^{k,1}/\mathbf{R}_* \cong \mathbf{RP}^{k-2}$, since the dimension of $H^{k,1}$ is $k-1$; if $k = 1$ then $Sh_{1,*}^k$ is empty. Sing_1^k consists only of the 1-point shape $\bar{\mathbf{0}}$. Thus, as a set

$$Sh_1^k = \{\bar{\mathbf{0}}\} \coprod \begin{cases} \emptyset & \text{if } k = 1, \\ Sh_{1,*}^k \cong \mathbf{RP}^{k-2} & \text{if } k \ge 2. \end{cases} \tag{2.2}$$

The induced topology on the subset $Sh_{1,*}^k$ (for $k \ge 2$) is the usual topology on $\mathbf{RP}^{k-2}$. Furthermore for every $\bar{\mathbf{v}} \in Sh_{1,*}^k$ there exists an open neighborhood $U(\bar{\mathbf{v}})$ that excludes $\bar{\mathbf{0}}$, since for any centered pre-shape $\mathbf{v} \in H_*^{k,1}$ there exists an open neighborhood $U(\mathbf{v})$ that excludes the pre-shape $\mathbf{0}$. However, since $\mathbf{0}$ is in the closure of every orbit of $GL(1)$ on $H^{k,1}$, *every* open neighborhood of $\{\bar{\mathbf{0}}\}$ includes *every* point of $Sh_{1,*}^k$. Said another way, the only open neighborhood of $\bar{\mathbf{0}}$ in Sh_1^k is the entire space Sh_1^k.

Similarly, we find

$$\widetilde{Sh}_1^k = \{\bar{\mathbf{0}}\} \coprod \begin{cases} \emptyset & \text{if } k = 1, \\ \widetilde{Sh}_{1,*}^k \cong S^{k-2} & \text{if } k \ge 2 \end{cases} \tag{2.3}$$

(where S^{k-2} is the $(k-2)$-sphere) with the same sort of topology that Sh_1^k has. Observe that for this value of n the natural projection $\widetilde{Sh}_n^k \to Sh_n^k$ is two-to-one on the smooth set and one-to-one on the singular set. We will see later that this is the case in general.

Remark 2.2 Of course $\widetilde{Sh}_{1,*}^k$ is the same as the Procrustes shape space Σ_1^k, since the group of orientation-preserving, origin-preserving, similarity transformations of $\mathbf{R}^1$ is exactly $GL_+(\mathbf{R}^1)$.

3 The Smooth Sets $Sh_{n,*}^k$, $\widetilde{Sh}_{n,*}^k$ for General n

The space $(\mathbf{R}^n)^k$ can be naturally identified with $\mathbf{R}^{n,k}$, the space of real $n \times k$ matrices. Rather than introduce notation for this natural isomorphism, we will simply regard elements $\mathbf{v} \in (\mathbf{R}^n)^k$ as living simultaneously in $\mathbf{R}^{n,k}$, and regard $\mathbf{R}_*^{n,k}$ and $H_*^{n,k}$ as subspaces of $\mathbf{R}^{n,k}$. However, when we wish to emphasize characteristics of an element $\mathbf{v} \in (\mathbf{R}^n)^k$ as a matrix rather than as a list of vectors, we will write $\mathbf{v}_{\mathrm{mat}}$ for the associated matrix. Temporarily using "$\cdot$" to denote the action of $GL(n)$ on $(\mathbf{R}^n)^k$, for $A \in GL(n)$ we have $(A \cdot \mathbf{v})_{\mathrm{mat}} = A\mathbf{v}_{\mathrm{mat}}$, i.e. the action carries over into simple matrix-multiplication. Therefore we will henceforth simply write "$A\mathbf{v}$" instead of "$A \cdot \mathbf{v}$".

For all $\mathbf{v} \in \mathbf{R}^{n,k}$, let $W_{\mathbf{v}} \subset \mathbf{R}^k$ denote the row space of $\mathbf{v}_{\mathrm{mat}}$. A more abstract characterization of $W_{\mathbf{v}}$ that will be conceptually useful later is as follows. For any finite-dimensional vector space Z let Z^* denote the dual space $\mathrm{Hom}(Z, \mathbf{R})$. Then there are natural isomorphisms $Z^k \cong Z \otimes \mathbf{R}^k \cong \mathrm{Hom}(Z^*, \mathbf{R}^k)$. Let $L : Z^k \to \mathrm{Hom}(Z^*, \mathbf{R}^k)$ denote the composite isomorphism, and for $\mathbf{v} \in Z^k$ write $L_{\mathbf{v}} = L(\mathbf{v})$; thus $L_{\mathbf{v}}$ is a map $Z^* \to \mathbf{R}^k$. Taking $Z = \mathbf{R}^n$, the row space $W_{\mathbf{v}}$ is simply $\mathrm{image}(L_{\mathbf{v}})$, and if $\mathbf{v} = (v_1, \dots, v_k)$ then $L_{\mathbf{v}}$ is given explicitly by $L_{\mathbf{v}}(z) = (\langle z, v_1 \rangle, \dots, \langle z, v_k \rangle)$, where $\langle \cdot, \cdot \rangle$ denotes the dual pairing between Z^* and Z.

Let $\mathbf{R}_0^k \subset \mathbf{R}^k$ be the subspace consisting of all vectors whose components sum to 0; thus $\dim(\mathbf{R}_0^k) = k-1$. Assume for now that $k > n$, so that $H_*^{n,k}$ is nonempty. If $\mathbf{v} \in H^{n,k}$, then every row of $\mathbf{v}_{\mathrm{mat}}$ lies in $\mathbf{R}_0^k$. If $\mathbf{v} \in H_*^{n,k}$, then the columns of $\mathbf{v}_{\mathrm{mat}}$ span an n-dimensional space, hence so do the n rows (equivalently, $\mathrm{image}(L_{\mathbf{v}}) \subset \mathbf{R}_0^k$ is n-dimensional); thus they are a linearly independent set in $\mathbf{R}_0^k$. Since left-multiplication by an invertible matrix does not change the row space, we have $W_{A\mathbf{v}} = W_{\mathbf{v}}$ for all $A \in$

$GL(n)$. For any linear transformation $\mathbf{R}^n \to \mathbf{R}^n$ (respectively, any $n \times n$ matrix B), letting $B^* : (\mathbf{R}^n)^* \to (\mathbf{R}^n)^*$ denote the map dual to B (resp., the dual of the map "left-multiplication by B"), we have $L_{B\mathbf{v}} = L_{\mathbf{v}} \circ B^*$. In our current situation, since A is an isomorphism so is A^*, and therefore $\text{image}(L_{A\mathbf{v}}) = \text{image}(L_{\mathbf{v}})$.

Recall that for any finite-dimensional vector space X, the (generalized) *Stiefel manifold* $V_n(X)$ of n-frames in X ($n \le K := \dim(X)$) is the set of linearly independent n-tuples[2] of vectors in X, topologized as a subset of X^n, in which it is open. For each n-frame $q \in V_n(X)$, let $S_q \subset X$ be its span—an n-dimensional subspace having q as a basis. The group $GL(n)$ acts on $V_n(X)$ in an obvious way, and for $A \in GL(n)$ we have $S_{Aq} = S_q$ (A just changes the basis of S_q). The assignment $q \mapsto S_q$ is a map from $V_n(X)$ to $G_n(X)$, the Grassmannian of n-dimensional subspaces of X, and exhibits $G_n(X)$ as the quotient $V_n(X)/GL(n)$.

Thus the map that assigns to $\mathbf{v} \in H_*^{n,k}$ its n-tuple of rows is a diffeomorphism $H_*^{n,k} \to V_n(\mathbf{R}_0^k)$, and induces a smooth map $\tau_n^k : H_*^{n,k} \to G_n(\mathbf{R}_0^k)$, $\tau_n^k(\mathbf{v}) = W_{\mathbf{v}}$, that satisfies $\tau_n^k(A\mathbf{v}) = \tau_n^k(\mathbf{v})$. Hence τ_n^k induces a well-defined map $\bar{\tau}_n^k : Sh_{n,*}^k = H_*^{n,k}/GL(n) \to G_n(\mathbf{R}_0^k)$ that one can easily check is a homeomorphism. This homeomorphism, together with the natural smooth structure of the Grassmanian $G_n(\mathbf{R}_0^k)$, endows $Sh_{n,*}^k$ with the structure of a smooth manifold; thus $\bar{\tau}_n^k$ becomes a diffeomorphism $Sh_{n,*}^k \to G_n(\mathbf{R}_0^k)$. Any isomorphism $\mathbf{R}_0^k \to \mathbf{R}^{k-1}$ then gives rise to a diffeomorphism $Sh_{n,*}^k \to G_n(\mathbf{R}^{k-1})$, a manifold of dimension $n(k-1-n)$. Note that $G_1(\mathbf{R}^{k-1}) = \mathbf{RP}^{k-2}$, the space seen earlier in Example 2.1.

We can do a similar analysis for $\widetilde{Sh}_{n,*}^k$. In this case we consider the manifold $\tilde{V}_n(\mathbf{R}_0^k)$ of *oriented* frames, introducing an equivalence relation on $V_n(\mathbf{R}_0^k)$ in which two frames are equivalent if and only some element of $GL_+(n)$ carries one to the other. This divides $V_n(\mathbf{R}_0^k)$ into two disjoint sets labeled by orientation. Letting $\widetilde{G}_n(\mathbf{R}_0^k)$ denote the Grassmannian of oriented n-planes in $\mathbf{R}_0^k$ (a double-cover of $G_n(\mathbf{R}_0^k)$) we have $\widetilde{G}_n(\mathbf{R}_0^k) = \tilde{V}_n(\mathbf{R}_0^k)/GL_+(n)$. Thus τ_n^k induces a homeomorphism $\tilde{\tau}_n^k : \widetilde{Sh}_{n,*}^k = H_*^{n,k}/GL(n) \to \widetilde{G}_n(\mathbf{R}_0^k)$, which we use to endow $\widetilde{Sh}_{n,*}^k$ with the structure of a manifold.

Thus we have now proven the following theorem:

Theorem 3.1 *If $k > n$ then $Sh_{n,*}^k$ is naturally diffeomorphic (via $\bar{\tau}_n^k$) to $G_n(\mathbf{R}_0^k)$, and $\widetilde{Sh}_{n,*}^k$ is naturally diffeomorphic (via $\tilde{\tau}_n^k$) to $\widetilde{G}_n(\mathbf{R}_0^k)$.*

Remark 3.2 The approach we have taken so far has one inelegant feature: it seems to rely on a choice of linear coordinate system in $\mathbf{R}^n$, despite the fact that the shape space itself is independent of such choices, as evidenced by our final identifications of $Sh_{n,*}^k$ with $G_n(\mathbf{R}_0^k)$. At the cost of abstraction, one can avoid this inelegance as follows. Replace $\mathbf{R}^n$ with Z_{aff}, an n-dimensional real affine space with underlying vector space Z. An affine pre-shape of k points is then any element of $(Z_{\text{aff}})^k$. Let $H^k(Z) = \{(v_1, \dots, v_k) \in Z^k \mid \sum_i v_i = 0\}$, and $H_*^k(Z) = \{(v_1, \dots, v_k) \in H^k(Z) \mid \text{span}\{v_i\} = Z\}$. There is a natural "centering map" $(Z_{\text{aff}})^k \to H^k(Z)$ that is completely independent of any choice of origin in Z_{aff}. The shape space $Sh^k(Z)$, defined as $Z_{\text{aff}}^k/\text{Aff}(Z_{\text{aff}})$, is then homeomorphic to $H^k(Z)/GL(Z)$, where $\text{Aff}(Z_{\text{aff}})$ is the group of invertible affine transformations of Z_{aff} and $GL(Z)$ the group of invertible linear transformations of Z. For $\mathbf{v} \in H^k(Z)$ define $L_{\mathbf{v}} : Z \to \mathbf{R}^k$ as before. Then $\text{image}(L_{\mathbf{v}}) \subset \mathbf{R}_0^k$, and if $\mathbf{v} \in H_*^k(Z)$ then $L_{\mathbf{v}}$ is injective, so that $\dim(W_{\mathbf{v}}) = n$. Thus the map $\tau : \mathbf{v} \mapsto \text{image}(L_{\mathbf{v}})$ sends $H_*^k(Z)$ to $G_n(\mathbf{R}_0^k)$. For any $A \in GL(Z)$, let A^* denote the dual map in $GL(Z^*)$; by our earlier argument $\text{image}(L_{A\mathbf{v}}) = \text{image}(L_{\mathbf{v}} \circ A^*) = \text{image}(L_{\mathbf{v}})$. Thus $\hat{\tau}$ descends to a map $\bar{\tau} : H_*^k(Z)/GL(Z) \to G_n(\mathbf{R}_0^k)$. One can again check that this map is homeomorphism, and then proceed as before to endow $Sh_*^k(Z) := H_*^k(Z)/GL(Z)$ with the smooth structure of $G_n(\mathbf{R}_0^k)$.

4 Stratification of Sh_n^k

We have already "reduced" the quotient $(\mathbf{R}^n)^k/\text{Aff}(n)$ to $H^{n,k}/GL(n)$, by first modding out by translations, and will be using further reductions to analyze the singular sets in Sh_n^k and $\widetilde{Sh}_n^k$. Because we are dealing with a non-proper group action, it may be worthwhile to remind the reader that only a very minimal set of hypotheses is needed to ensure that such reductions are not merely bijections, but homeomorphisms:

Lemma 4.1 (Reduction Principle) *Suppose a group G acts on a topological space X by continuous transformations and that there is a subset $Y \subset X$ with the following properties*:

1. *For all $x_0 \in X$, there exists an open neighborhood U_{x_0} of x_0 and a (not necessarily unique or continuous) function $g^{(x_0)} : U_{x_0} \to G$ such that the map $x \mapsto g^{(x_0)}(x) \cdot x$ is a continuous map $U_{x_0} \to Y$.*
2. *For all $y \in Y$, if $g \in G$ satisfies $g \cdot y \in Y$, then there exists $g_y \in G_Y$ with $g_y \cdot y = g \cdot y$, where $G_Y = \{h \in G \mid h \cdot Y \subset Y\}$.*

Then the assignment $x \mapsto g^{(x)}(x) \cdot x$ induces a well-defined map $\phi : X/G \to Y/G_Y$, independent of choices of the maps $g^{(x)}$, given by $\phi(G \cdot x) = G_Y \cdot (g(x) \cdot x)$, and ϕ is a homeomorphism with respect to the quotient topologies.

[2] We use the convention of [16, §12], in which the frames are not required to be orthonormal.

Proof Left to reader, with the reminder that quotient maps for group-actions are always open maps. □

The complement Sing_n^k of $Sh_{n,*}^k$ in Sh_n^k is the quotient by $GL(n)$ of the set of elements $\mathbf{v} = (v_1, \dots, v_k) \in H_*^{n,k}$ for which $\dim(\mathrm{span}\{v_i\}) \le n-1$, and hence for which the $\{v_i\}$ are contained in some $(n-1)$-dimensional subspace $P \subset \mathbf{R}^n$. Every such P can be carried by some $A_P \in GL(n)$ to a fixed copy of P_0 of $\mathbf{R}^{n-1} \subset \mathbf{R}^n$, which we will take to be the subspace given by the condition "last coordinate $= 0$" and, in a small enough neighborhood of P in $G_{n-1}(\mathbf{R}^n)$, A_P may be chosen to depend continuously on P. The subgroup of $GL(n)$ preserving this fixed $\mathbf{R}^{n-1}$, hence preserving $Y := H^{n,k} \cap (P_0)^k \subset H^{n,k}$, is

$$G_Y = \left\{ \begin{pmatrix} A' & b \\ 0 & c \end{pmatrix} \mid A' \in GL(n-1), b \in \mathbf{R}^{n-1}, c \in \mathbf{R}_* \right\}. \tag{4.1}$$

Given $\mathbf{v} = (v_1, \dots, v_k) \in Y$ and $A \in GL(n)$ for which $A\mathbf{v} \in Y$, there exists $A_\mathbf{v} \in G_Y$ for which $A_\mathbf{v}\mathbf{v} = A\mathbf{v}$ (choose a basis $\{z_i\}_1^m$ of $\mathrm{span}\{v_1, \dots, v_k\}$, extend to a basis $\{z_i\}_1^n$ of $\mathbf{R}^n$ with $z_i \in P_0$ for $i < n$, and define $A_\mathbf{v} z_i = A z_i$ for $i \le m$, $A_\mathbf{v} z_i = z_i$ for $i > m$). In the notation of (4.1), the action on P_0 of the element of G_Y is simply left-multiplication by A' (after identifying P_0 with $\mathbf{R}^{n-1}$ by dropping the last coordinate). Thus using Lemma 4.1 we obtain a natural homeomorphism

$$\sigma_{n,n-1}^k : \mathrm{Sing}_n^k = H^{k,n-1}/GL(n-1) \to Sh_{n-1}^k. \tag{4.2}$$

By definition, we have $Sh_n^k = Sh_{n,*}^k \coprod \mathrm{Sing}_n^k$, and we express certain aspects of the topology of this decomposition by writing

$$\begin{aligned} Sh_n^k &\cong' Sh_{n,*}^k \coprod \mathrm{Sing}_n^k && (4.3)\\ &\cong' Sh_{n,*}^k \coprod Sh_{n-1}^k && (4.4)\\ &\cong' G_n(\mathbf{R}_0^k) \coprod Sh_{n-1}^k && (4.5) \end{aligned}$$

where "$\cong'$" means only that the left-hand side is a disjoint union of sets, each of which, in the relative topology, is homeomorphic to one of the sets on the right-hand side. As already seen in Example 2.1 with $j = 1$, Sh_j^k is *not* homeomorphic to $\mathrm{Sing}_j^k \coprod Sh_{j-1}^k$ topologized as the union of independent topological spaces.

For $0 \le j \le n$ let us define

$$H_j^{n,k} = \{\mathbf{v} \in H^{n,k} \mid \mathrm{rank}(\mathbf{v}_{\mathrm{mat}}) = j\},$$

$$Sh_{n,j}^k = H_j^{n,k}/GL(n) \subset Sh_n^k$$

(so $H_n^{n,k}$ and $Sh_{n,n}^k$ are now alternative names for $H_*^{n,k}$ and $Sh_{n,*}^k$). The discussion above leads us to the following theorem, part (a) of which (minus the indicated maps) is stated without proof as Theorem 2.1 in [18].

Theorem 4.2 *Let $k \ge 1$, $n \ge 0$. (a) If $k > n$ then*

$$\begin{array}{lcccccccccr}
Sh_n^k \cong' & Sh_{n,n}^k & \coprod & Sh_{n,n-1}^k & \coprod \cdots \coprod & Sh_{n,1}^k & \coprod & Sh_{n,0}^k & & & (4.6)\\
& \updownarrow \mathrm{id.} & & \updownarrow \sigma_{n,n-1}^k & & \updownarrow \sigma_{n,1}^k & & \updownarrow \sigma_{n,0}^k & & & \\
\cong' & Sh_{n,*}^k & \coprod & Sh_{n-1,*}^k & \coprod \cdots \coprod & Sh_{1,*}^k & \coprod & Sh_0^k & & & (4.7)\\
& \updownarrow \tau_n^k & & \updownarrow \tau_{n-1}^k & & \updownarrow \tau_1^k & & \updownarrow \tau_0^k & & & \\
\cong' & G_n(\mathbf{R}_0^k) & \coprod & G_{n-1}(\mathbf{R}_0^k) & \coprod \cdots \coprod & G_1(\mathbf{R}_0^k) & \coprod & G_0(\mathbf{R}_0^k) & & & (4.8)\\
& \updownarrow & & \updownarrow & & \updownarrow & & \updownarrow & & & \\
\cong' & G_n(\mathbf{R}^{k-1}) & \coprod & G_{n-1}(\mathbf{R}^{k-1}) & \coprod \cdots \coprod & G_1(\mathbf{R}^{k-1}) & \coprod & G_0(\mathbf{R}^{k-1}) & & & (4.9)
\end{array}$$

where $\sigma_{n,j}^k : Sh_{n,j}^k \to Sh_{j,*}^k$ is obtained by iterating the procedure giving the decomposition (4.3)–(4.4), and where all of the maps indicated with "$\updownarrow$" are homeomorphisms from the upper space to the lower space. The natural projection $\widetilde{Sh}_n^k \to Sh_n^k$ restricts to a double-cover $\widetilde{Sh}_{n,*}^k \to Sh_{n,*}^k$, and to a bijection $\widetilde{\mathrm{Sing}}_n^k \to \mathrm{Sing}_n^k$. Hence $\widetilde{Sh}_n^k$ has a decomposition similar to (4.7)–(4.9), the only difference being that the first spaces appearing on the right-hand sides are replaced by $\widetilde{Sh}_{n,n}^k \cong \widetilde{Sh}_{n,*}^k \cong \widetilde{G}_n(\mathbf{R}_0^k) \cong \widetilde{G}_n(\mathbf{R}^{k-1})$.

(b) *If $k \le n$ then $\widetilde{Sh}_n^k = Sh_n^k \cong Sh_{k-1}^k$, "$=$" meaning that the natural projection $\widetilde{Sh}_n^k \to Sh_n^k$ is a bijection.*

Proof (a) Equation (4.6) is a tautology. In the procedure used above to analyze Sing_n^k, let R be the map that carried the complement of $H_*^{n,k}$ in $H^{n,k}$ to Y. Then $R(\mathbf{v})$ is obtained from $\mathbf{v}$ by multiplying by an invertible matrix, which does not change the rank ($\mathrm{rank}((R(\mathbf{v}))_{\mathrm{mat}}) = \mathrm{rank}(\mathbf{v}_{\mathrm{mat}})$). Dropping the last coordinate of points in P_0 maps $(R(\mathbf{v}))_{\mathrm{mat}}$ to an $(n-1) \times k$ matrix by dropping a row of zeroes, which again does not change the rank. Thus the subset of Sh_n^k which the homeomorphism $\sigma_{n,n-1}^k$ identifies with $Sh_{n-1,*}^k$ is exactly $Sh_{n,n-1}^k$. Iterating, we get the statement that for all $j \le n-1$, $\sigma_{n,j}^k$ is a homeomorphism $Sh_{n,j}^k \to Sh_{j,*}^k$. The assertion that the maps τ_j^k are homeomorphisms follows follow from iterating (4.3)–(4.5); the passage to (4.9) is then obvious since $\mathbf{R}_0^k \cong \mathbf{R}^{k-1}$.

We have already seen in Theorem 3.1 that $\widetilde{Sh}_{n,*}^k$ is naturally homeomorphic to $\widetilde{G}_n(\mathbf{R}_0^k)$, and tracing through our identifications it is clear that the natural projection $\widetilde{Sh}_{n,*}^k \to Sh_{n,*}^k$ induces the canonical covering map $\widetilde{G}_n(\mathbf{R}_0^k) \to G_n(\mathbf{R}_0^k)$.

The analysis of the singular set $\widetilde{\mathrm{Sing}}_n^k$ proceeds almost identically to the analysis of Sing_n^k preceding the theorem. Every $(n-1)$-dimensional subspace of $\mathbf{R}^n$ can be carried by $GL_+(n)$ to our fixed copy P_0 of $\mathbf{R}^{n-1} \subset \mathbf{R}^n$, but the subgroup G_{Y+} of $GL_+(n)$ preserving P_0 is the subgroup of G_Y, the group in (4.1), for which $\mathrm{sign}(\det(A')) = \mathrm{sign}(c)$, so both signs of $\det(A')$ can occur. Thus the residual action of G_{Y+} on $\mathbf{R}^{n-1}$ is that of the *full* group $GL(n-1)$, so the orbit of any $\mathbf{v} \in H^{n,k} \cap Y^k$ under G_{Y+} is the same as its orbit under G_Y. Hence the natural projection $\widetilde{\mathrm{Sing}}_n^k \to \mathrm{Sing}_n^k$ is a canonical bijection.

(b) For $k \le n$, every $\mathbf{v} \in H^{n,k}$ lies in some $(k-1)$-dimensional subspace of $H^{n,k}$. Thus the same analysis as for the singular sets in the case $k > n$ now leads to a homeomorphism $Sh_n^k \to Sh_{k-1}^k$ and to a canonical bijection $\widetilde{Sh}_n^k \to Sh_n^k$. □

Of course, Theorem 4.2 does not tell us the full topology of Sh_n^k, since it says nothing about how the strata are glued together. That discussion we postpone till the next section.

Remark 4.3 The homeomorphisms $\sigma_{n,j}^k$ and τ_j^k in Theorem 4.2 are canonical, independent of any choices, whereas the homeomorphisms taking us from (4.8) to (4.9) depend on a choice of isomorphism $\mathbf{R}_0^k \to \mathbf{R}^{k-1}$. This is completely analogous to the situation for the Procrustes shape-spaces Σ_2^k, which are *canonically* identifiable with $\mathbf{P}(\mathbf{C}_0^k)$ (the complex projectivization of $\mathbf{C}_0^k$) and non-canonically identifiable with $\mathbf{CP}^{k-2}$. In the Procrustean case, this is geometrically significant: the Procustes-Riemannian metric (see [12]) comes from the metric on $\mathbf{C}_0^k$ *as a subspace of* $\mathbf{C}^k$. Any isometry $\mathbf{C}_0^k \to \mathbf{C}^{k-1}$ will carry Σ_2^k isometrically to $\mathbf{CP}^{k-2}$ with a Fubini-Study metric, but there is no *canonical* isometry; from a geometric standpoint one must avoid the temptation to use the non-isometric isomorphisms $\mathbf{C}_0^k \to \mathbf{C}^{k-1}$ obtained by dropping a coordinate. In Sect. 6, when we discuss metrics on the $Sh_{n,*}^k$, the same principle will be important; we must use the *natural* identification $Sh_{n,*}^k \cong G_n(\mathbf{R}_0^k)$ and not an unnatural identification with $G_n(\mathbf{R}^{k-1})$.

Remark 4.4 One can get the false impression from (4.9) that Sh_n^k is not connected and that the strata, each of which is compact in the relative topology, are its connected components. However, Sh_n^k is connected, being the image of the connected space $H^{n,k}$ under a continuous map.

In view of the identification of the strata in (4.7) in terms of matrix ranks, it is natural to order the strata by declaring $Sh_{n,n}^k > Sh_{n,n-1}^k > \cdots > Sh_{n,1}^k > Sh_{n,0}^k$ (similarly for the $\widetilde{Sh}_{n,j}^k$); thus we refer to $Sh_{n,*}^k$ as the top stratum of Sh_n^k and to $Sh_{n,0}^k$ as the bottom stratum, and the meaning of "higher" and "lower" when relating strata to each other is clear. In the same vein, for $\bar{\mathbf{v}} \in Sh_n^k$, define the *level* of $\bar{\mathbf{v}}$, written level($\bar{\mathbf{v}}$), to be that j for which $\bar{\mathbf{v}} \in Sh_{n,j}^k$; thus $\bar{\mathbf{v}}$ lies in a higher stratum than $\bar{\mathbf{w}} \iff \mathrm{level}(\bar{\mathbf{v}}) > \mathrm{level}(\bar{\mathbf{w}})$. We make some observations:

1. For a fixed k, the largest $n < k$ is $k-1$. Thus the stratifications (4.7)–(4.9) are contained in the stratification of Sh_{k-1}^k, for which we have the following picture:

$$
\begin{array}{lccccccc}
Sh_{k-1}^k & \cong' G_{k-1}(\mathbf{R}^{k-1}) \coprod \cdots \coprod & G_j(\mathbf{R}^{k-1}) & \coprod \cdots \coprod G_1(\mathbf{R}^{k-1}) & \coprod G_0(\mathbf{R}^{k-1}) \\
\text{dimensions:} & 0 & j(k-1-j) & 1(k-2) & 0
\end{array}
\tag{4.10}
$$

The dimensions of the strata are symmetric relative to the middle stratum or two middle strata, depending on the parity of k, and are maximal in the middle. Thus *the top stratum is not always the stratum of largest dimension*: if $n > k/2$, there is always a singular stratum of dimension greater than that of the smooth set; if $n = k/2$ then the smooth stratum and the top singular stratum have equal dimensions. Note that the next-to-lowest stratum is just $\mathbf{RP}^{k-2}$, as seen in Example 2.1, written in different notation.

In the case of Procrustes shape-spaces, the rank-stratification of $H^{n,k}$ as $\coprod_j H_j^{n,k}$ also produces a corresponding stratification of Σ_n^k as $\coprod_j \Sigma_{n,j}^k$. In contrast to the "dimensional anomaly" that we see for strata in the affine case for $k \le 2n$, the top Procrustes stratum $\Sigma_{n,\max(n,k-1)}^k$ always has the largest dimension, and the dimensions of the strata strictly decrease as j decreases from $\max(n, k-1)$ to 1. For $1 \le j \le \max(n, k-1)$, the union of the strata up through level j form a cell-subcomplex of the cell-complex Σ_n^k.

2. The bottom stratum of Sh_n^k always consists of a single point, namely $\bar{\mathbf{0}} \in Sh_n^k$. While this shape is excluded in the definition of Procrustes shape-spaces, we see that it fits quite naturally into the discussion of spaces of affine shapes; indeed, (4.7)–(4.9) and (4.10) would seem incomplete without it.

One can, of course, define an "augmented Procrustes shape space" $\widehat{\Sigma}_n^k$ as the quotient-space obtained by letting rotations and dilations of $\mathbf{R}^n$ act on all of $H^{n,k}$ rather than just on $H^{n,k} - \{\mathbf{0}\}$, thereby including the single-point shape $\bar{\mathbf{0}}$ into Procrustean analysis. However, while $\widehat{\Sigma}_n^k = \Sigma_n^k \coprod \{\bar{\mathbf{0}}\}$ as a point-set, this decomposition is not a homeomorphism: the only open neighborhood of $\bar{\mathbf{0}}$ is the whole space $\widehat{\Sigma}_n^k$. Thus, including the single-point shape

Table 1 Decomposition of Sh_2^k and $\widetilde{Sh}_2^k$ for $1 \le k \le 5$. The strata are listed in decreasing order; the first is the smooth set. For $k > 2$, since the singular strata are the same for Sh_2^k and $\widetilde{Sh}_2^k$, we list only the top stratum of $\widetilde{Sh}_2^k$. The identifications of $\widetilde{G}_2(\mathbf{R}^4)$ with $S^2 \times S^2$ and of $G_2(\mathbf{R}^4)$ with $(S^2 \times S^2)/\mathbf{Z}_2$, where the nontrivial element of $\mathbf{Z}_2$ acts by the antipodal map simultaneously on each S^2, are well-known facts that can be proven by examining the action of the unit quaternions on $\mathbf{R}^4$

k	Sh_2^k	$\widetilde{Sh}_{2,*}^k$
1	point	same as Sh_1^1
2	$Sh_2^2 \cong Sh_1^2 \cong' G_1(\mathbf{R}^1) \coprod G_0(\mathbf{R}^0)$	same as Sh_2^2
	$\cong'$ point $\coprod$ point	
	dimensions : 0, 0	
3	$G_2(\mathbf{R}^2) \coprod G_1(\mathbf{R}^2) \coprod G_0(\mathbf{R}^2)$	$\widetilde{G}_2(\mathbf{R}^2) \cong$ two points
	$\cong'$ point $\coprod \mathbf{R}P^1 =$ circle $\coprod$ point	
	dimensions : 0, 1, 0	
4	$G_2(\mathbf{R}^3) \coprod G_1(\mathbf{R}^3) \coprod G_0(\mathbf{R}^3)$	$\widetilde{G}_3(\mathbf{R}^2) \cong S^2$
	$\cong' \mathbf{RP}^2 \coprod \mathbf{RP}^2 \coprod$ point	
	dimensions : 2, 2, 0	
5	$G_2(\mathbf{R}^4) \coprod G_1(\mathbf{R}^4) \coprod G_0(\mathbf{R}^4)$	$\widetilde{G}_2(\mathbf{R}^4) \cong S^2 \times S^2$
	$\cong' (S^2 \times S^2)/\mathbf{Z}_2 \coprod \mathbf{R}P^3 \coprod$ point	
	dimensions : 4, 3, 0	

into the Procrustes picture drastically alters the topology: the (new) bottom stratum is attached to all higher strata very differently from the way the higher strata are attached to each other, and no union of strata that includes the bottom stratum is a cell complex. Σ_n^k is a metric space while $\widehat{\Sigma}_n^k$ is not. In contrast, if we delete the point $\overline{\mathbf{0}}$ from Sh_n^k, the topology of what remains is no more nicely behaved than the topology of the full space Sh_n^k (unless $k = 2$, in which case what remains is a single point).

To add to our list of concrete examples, Table 1 lists the spaces Sh_2^k for small values of k. Note that in general $\widetilde{Sh}_n^k - \{\overline{\mathbf{0}}\}$ is a quotient of the Procrustes shape-space Σ_n^k. In particular, $\widetilde{Sh}_2^3 - \{\overline{\mathbf{0}}\}$ is a quotient of $\Sigma_2^3 \cong S^2$, which at first seems impossible looking at the table entry for $\widetilde{Sh}_2^3$, until one remembers that the quotient is not Hausdorff. In the model of Σ_2^3 as $\mathbf{CP}^1 \cong S^2$, the equator $\mathbf{RP}^1$ corresponds to collinear configurations, the open northern hemisphere corresponds to triangles in one orientation class, and the open southern hemisphere to triangles in the other orientation class. Under the projection $\Sigma_2^3 \to Sh_2^3$, each open hemisphere gets mapped to one of the points in $\widetilde{G}_2(\mathbf{R}^2)$, while the equator gets mapped bijectively to the stratum $Sh_{2,1}^3 \cong \mathbf{RP}^1$.

5 Gluing the Strata Together

For simplicity, in this section and the next we confine our analysis to the spaces Sh_n^k; for $\widetilde{Sh}_n^k$ the reader can easily make the necessary minor adjustments, which of course are only relevant to the top stratum.

To discuss how the strata of Sh_n^k are attached to each other, we make the following definition.

Definition 5.1 Let X be a topological space. For all $p \in X$, define $\text{blur}(p)$ to be the intersection of all open neighborhoods of p in X. Call p *blurry* if $\text{blur}(p) \neq \{p\}$. (b) For $Y \subset X$, define $\text{blur}(Y) = \bigcup_{p \in Y} \text{blur}(p)$.

In other words, $\text{blur}(p)$ consists of those points in X that cannot be separated from p by an open set containing p. The authors have searched in vain for pre-existing terminology for this concept. However, recall that a topological space X is said to be T_0 if, given two distinct points in X, for at least one of the points there exists an open set containing that point while excluding the other; X is said to be T_1 if, given distinct points in X, for *each* point there exists an open set containing that point but excluding the other. In a T_1 space, $\text{blur}(p) = \{p\}$, so "blur" is worth defining only for spaces that have very poor separation properties—which, as we will see, is the case for the spaces Sh_n^k.

Blurs and closures are intimately related to each other. As (5.1) below shows, for points the notions of blur and closure are in some sense dual to each other. This duality does not hold for sets; the implication in part (a) of the lemma below is only one-directional, metric spaces providing simple counterexamples to the other direction.

Lemma 5.2 *Let X be a topological space.*

(a) *For all $Y \subset X$, $\text{closure}(Y) \supset \{q \in X \mid \text{blur}(q) \cap Y \neq \emptyset\}$. Hence if $Y, Z \subset X$, then $\text{blur}(Z) \cap Y \neq \emptyset \Rightarrow Z \cap \text{closure}(Y) \neq \emptyset$.*
(b) *For all $p, q \in X$, $\text{closure}(\{p\}) = \{q \in X \mid p \in \text{blur}(q)\}$. In other words*

$$p \in \text{blur}(q) \iff q \in \text{closure}(\{p\}). \tag{5.1}$$

Proof

(a) Let $Y \subset X$. Using "prime" to denote complementation, $q \in \text{closure}(Y)' \Rightarrow q \in U$ for some open set U not intersecting Y. Hence $\text{blur}(q) \subset U \subset Y'$, so $\text{blur}(q) \cap Y = \emptyset$. Therefore $\text{closure}(Y)' \subset \{q \in X \mid \text{blur}(q) \cap Y = \emptyset\}$, so the first the assertion of part (a) lemma follows. The second assertion follows easily from the first.

(b) If $q \in X$ and $p \notin \text{blur}(q)$ then there exists an open neighborhood of q not containing p, and hence $q \in \text{closure}(\{p\})'$. Thus $\text{closure}(\{p\})' \supset \{q \in X \mid p \notin \text{blur}(q)\}$, so $\text{closure}(\{p\}) \subset \{q \in X \mid p \in \text{blur}(q)\}$; but from (a), $\text{closure}(\{p\}) \supset \{q \in X \mid p \in \text{blur}(q)\}$. □

In the stratified space Sh^k_n, blurs and closures stratify as well; for $0 \le j \le n$ we define $\text{blur}_j(p) := \text{blur}(p) \cap Sh^k_{n,j}$ and $\text{closure}_j(p) := \text{closure}(p) \cap Sh^k_{n,j}$.

In a stratified Hausdorff space, the way that the strata are attached is determined by intersections of strata with the closures of other strata. In a space that is not even T_1, blurs as well as closures are needed to describe how the strata are attached (under any intuitively reasonable definition of "attached"). We shall see below that for $p \in Sh^k_n$, $\text{blur}(p)$ connects p to all higher strata but no lower ones, while for $\text{closure}(\{p\})$ the behavior is exactly the opposite.

We now describe explicitly the blurs of points and strata in Sh^k_n. For $k > 1$ and $n \ge 0$ let $\iota : Sh^k_n \to \coprod_{0\le j\le \max(n,k-1)} G_j(\mathbf{R}^k_0)$ be the map given by $\iota(\bar{\mathbf{v}}) = \text{image}(L_{\mathbf{v}}) = $ row-space of $\mathbf{v}$, where $\mathbf{v} \in H^{n,k}$ is a representative of $\bar{\mathbf{v}} \in Sh^k_n$. For $0 \le j \le \max(n, k-1)$ let $\iota^k_{n,j} := \iota|_{Sh^k_{n,j}} : Sh^k_{n,j} \to G_j(\mathbf{R}^k_0)$. Each map $\iota^k_{n,j}$ is always a homeomorphism: if $k > n$ then, in the notation of in Theorem 4.2, $\iota^k_{n,j} = \tau^k_j \circ \sigma^k_{n,j}$, while if $k \le n$ then $\iota^k_{n,j}$ is the map $\tau^k_j \circ \sigma^k_{k-1,j}$ precomposed with the homeomorphism $Sh^k_{n,j} \cong Sh^k_{k-1,j}$ given by Theorem 4.2(b). (In fact each $\iota^k_{n,j}$ is a diffeomorphism, by our definition of the smooth structure on each stratum.)

Lemma 5.3 *Let $k \ge 1, n \ge 0$, and let $\bar{\mathbf{v}}, \bar{\mathbf{w}} \in Sh^k_n$. Then the subspace $\iota(\bar{\mathbf{w}}) \subset \mathbf{R}^k_0$ contains the subspace $\iota(\bar{\mathbf{v}}) \subset \mathbf{R}^k_0$ if and only if there exists a linear transformation $A : \mathbf{R}^n \to \mathbf{R}^n$, not necessarily invertible, such that $\mathbf{v} = A\mathbf{w}$ for some (and hence all) representatives $\mathbf{v}, \mathbf{w}$ of $\bar{\mathbf{v}}, \bar{\mathbf{w}}$.*

Proof If $\mathbf{v} = A\mathbf{w}$ for some representatives $\mathbf{v}, \mathbf{w}$ of $\bar{\mathbf{v}}, \bar{\mathbf{w}}$ and some $A \in \text{Hom}(\mathbf{R}^n, \mathbf{R}^n)$, then, as discussed in Sect. 3, we have $L_{\mathbf{v}} = L_{\mathbf{w}} \circ A^*$, and hence $\text{image}(L_{\mathbf{v}}) \subset \text{image}(L_{\mathbf{w}})$, i.e. $\iota(\bar{\mathbf{v}}) \subset \iota(\bar{\mathbf{w}})$.

Conversely, suppose that $\iota(\bar{\mathbf{v}}) \subset \iota(\bar{\mathbf{w}})$; i.e. $\text{image}(L_{\mathbf{v}}) \subset \text{image}(L_{\mathbf{w}})$ for representatives $\mathbf{v}, \mathbf{w}$ of $\bar{\mathbf{v}}, \bar{\mathbf{w}}$. Then if $\{y_i\}$ is a basis of $(\mathbf{R}^n)^*$, there exist $\{y'_i\}^n_1$ such that $L_{\mathbf{v}}(y_i) = L_{\mathbf{w}}(y'_i)$, $1 \le i \le n$. Define a linear map $A^* : (\mathbf{R}^n)^* \to (\mathbf{R}^n)^*$ by declaring $A^*(y_i) = y'_i$ and extending linearly. Then $L_{\mathbf{v}} = L_{\mathbf{w}} \circ A^*$. But if $A : \mathbf{R}^n \to \mathbf{R}^n$ is the map dual to A^*, we also have $L_{A\mathbf{w}} = L_{\mathbf{w}} \circ A^*$. Hence $L_{\mathbf{v}} = L_{A\mathbf{w}}$, implying that $\mathbf{v} = A\mathbf{w}$. □

As a consequence of Lemma 5.3, we have the following:

Proposition 5.4 *Let $k \ge 1$, $0 \le j \le n$, $0 \le i \le n$, and let $\bar{\mathbf{v}} \in Sh^k_{n,j}$. Let $\bar{\mathbf{w}} \in Sh^k_n$ and let $\mathbf{v}, \mathbf{w} \in H^{k,n}$ be representatives of $\bar{\mathbf{v}}, \bar{\mathbf{w}}$. Then $\bar{\mathbf{w}} \in \text{blur}(\bar{\mathbf{v}})$ if and only if there exists a linear transformation $A : \mathbf{R}^n \to \mathbf{R}^n$ (not necessarily invertible) carrying $\mathbf{w}$ to $\mathbf{v}$, and we have*

$$\iota(\text{blur}_i(\bar{\mathbf{v}})) = \{P \in G_i(\mathbf{R}^k_0) \mid \text{the } i\text{-plane } P \text{ contains the } j\text{-plane } \iota(\bar{\mathbf{v}})\}. \tag{5.2}$$

Thus if $j = \text{level}(\bar{\mathbf{v}})$, then $\text{blur}_j(\bar{\mathbf{v}}) = \{\bar{\mathbf{v}}\}$, and $\text{blur}(\bar{\mathbf{v}})$ intersects all *higher strata—those with level greater than j—and* no *lower strata. In particular, no point in the top stratum is blurry, every point not in the top stratum is blurry, and*

$$\text{blur}(\bar{\mathbf{v}}) \subset \{\bar{\mathbf{v}}\} \coprod \left(\coprod_{i > \text{level}(\bar{\mathbf{v}})} Sh^k_{n,i} \right). \tag{5.3}$$

Proof Let $\bar{\mathbf{w}} \in Sh^k_{n,i}$ and let $\mathcal{O}_{\bar{\mathbf{v}}}, \mathcal{O}_{\bar{\mathbf{w}}} \subset H^{k,n}$ be the $GL(n)$-orbits representing $\bar{\mathbf{v}}, \bar{\mathbf{w}}$. Then $\bar{\mathbf{w}} \in \text{blur}(\bar{\mathbf{v}})$ if and only if the closure of $\mathcal{O}_{\bar{\mathbf{w}}}$ in $H^{k,n}$ intersects $\mathcal{O}_{\bar{\mathbf{v}}}$, i.e. if and only if there exists $\mathbf{w} \in \mathcal{O}_{\bar{\mathbf{w}}}$ and a sequence $\{A_m\}$ in $GL(n)$ with $\{\mathbf{w}_m := A_m\mathbf{w}\}$ converging to a point $\mathbf{v} \in \mathcal{O}_{\bar{\mathbf{v}}}$. (Note that we do not assume that the sequence $\{A_m\}$ itself converges.)

First suppose that there is such a sequence $\{A_m\}$. Then the image of $L_{\mathbf{w}_m}$ is the m-plane $\iota(\bar{\mathbf{w}})$, independent of i. Hence if $y \in (\mathbf{R}^n)^*$, $L_{\mathbf{w}_m}(y) \in \iota(\bar{\mathbf{w}})$ for all i, so $\lim_{i\to\infty}(L_{\mathbf{w}_m}(y)) \in \iota(\bar{\mathbf{w}})$. But the linear maps $L_{\mathbf{w}_m}$ converge to the linear map $L_{\mathbf{v}}$, so $L_{\mathbf{v}}(y) \in \iota(\bar{\mathbf{w}})$. (In terms of matrices, the foregoing says simply that the row space of the $n \times k$ matrix $A_i\mathbf{w}$ is the same for all i, so the rows of the limit matrix lie in this subspace as well.) Thus $\iota(\mathbf{v}) = \text{image}(L_{\mathbf{v}}) \subset \iota(\bar{\mathbf{w}})$; i.e. the m-plane $\iota(\bar{\mathbf{w}})$ contains the j-plane $\iota(\mathbf{v})$.

Conversely, suppose $\iota(\bar{\mathbf{w}}) \supset \iota(\bar{\mathbf{v}})$. By Lemma 5.3, $\mathbf{v} = A\mathbf{w}$ for some $A \in \text{Hom}(\mathbf{R}^n, \mathbf{R}^n)$. For $\epsilon \ne 0$ sufficiently small, $A + \epsilon I$ is invertible (the characteristic polynomial of A has only finitely many roots). Thus there exists a sequence ϵ_m tending to 0, with $A + \epsilon_m I$ invertible $\forall i$. Then $\lim_{i\to\infty}(A + \epsilon_m I)\mathbf{w}$ is a sequence in $\mathcal{O}_{\bar{\mathbf{w}}}$ whose limit is $\mathbf{v} \in \mathcal{O}_{\mathbf{v}}$. Hence $\bar{\mathbf{w}} \in \text{blur}(\bar{\mathbf{v}})$. Since ι is a homeomorphism $Sh^k_{n,i} \to G_i(\mathbf{R}^k_0)$, every $P \in G_i(\mathbf{R}^k_0)$ is $\iota(\bar{\mathbf{w}})$ for some $\bar{\mathbf{w}} \in Sh^k_{n,i}$, and (5.2) follows. □

Corollary 5.5 *Let $\bar{\mathbf{v}} \in Sh^k_{n,j}$. For $0 \le j \le i \le \max(k-1, n)$, $\text{blur}_i(\bar{\mathbf{v}})$ is homeomorphic to $G_{i-j}(\mathbf{R}^{k-1-j})$; for other values of i, $\text{blur}_i(\bar{\mathbf{v}})$ is empty.*

Proof Let V be a finite-dimensional vector space, let $H \subset V$ be a subspace of dimension j, and let X be the set of i-dimensional subspaces of V containing H. Then there is a natural bijection $X \leftrightarrow G_{i-j}(V/H)$. If we give X the relative topology as a subset of the Grassmannian $G_j(V)$, it is not hard to show that this bijection is a homeomorphism. If $\dim V = k-1$, an isomorphism $V/H \cong \mathbf{R}^{k-1-j}$ then yields a homeomorphism from X to $G_{i-j}(\mathbf{R}^{k-1-j})$. Applying this with $V = \mathbf{R}_0^k$, the result follows from Proposition 5.4. □

Remarks 5.6

(1) It is clear from (5.2) that $\text{blur}(\text{blur}(\bar{\mathbf{v}})) = \text{blur}(\bar{\mathbf{v}})$ for all $\bar{\mathbf{v}} \in Sh_n^k$, and hence that $\text{blur}(\text{blur}(U)) = \text{blur}(U)$ for all $U \subset Sh_n^k$. This property, reminiscent of closures, is a general feature of quotients by group actions; (5.2) just illustrates this principle.

(2) The blur of the bottom stratum $Sh_{n,0}^k = \{\bar{\mathbf{0}}\}$ is the entire space Sh_n^k (this is clear even without the analysis in this section, but again (5.2) provides a nice illustration). More generally, it follows from (5.2) that for each stratum $Sh_{n,j}^k$,

$$\text{blur}(Sh_{n,j}^k) = \coprod_{i \geq j} Sh_{n,i}^k. \quad (5.4)$$

Consequently, the only open set in Sh_n^k that contains Sing_n^k is the entire space Sh_n^k.

(3) From (5.2), we can also immediately compute $\text{blur}(\bar{\mathbf{v}}) \cap \text{blur}(\bar{\mathbf{w}})$ for any $\bar{\mathbf{v}}, \bar{\mathbf{w}}$, and (from Corollary 5.5) see that its intersection with any $Sh_{n,j}^k$, if nonempty, is yet another Grassmannian. Grassmannians truly abound in Sh_n^k.

(4) For the special case $n = 2$, if $k > 2$ we always have exactly three strata (cf. Table 1). In this case it is projective spaces that abound: for all $\bar{\mathbf{v}}$ in the stratum $Sh_{2,1}^k \cong \mathbf{RP}^{k-2}$, we have $\text{blur}(\bar{\mathbf{v}}) = \{\bar{\mathbf{v}}\} \cup (\text{blur}_2(\bar{\mathbf{v}}))$, and $\text{blur}_2(\bar{\mathbf{v}}) = \text{blur}(\bar{\mathbf{v}}) \cap Sh_{2,2}^k \cong \mathbf{RP}^{k-3}$.

Another corollary of Proposition 5.4 is the following.

Corollary 5.7 *For all $k \geq 2$ and $n \geq 1$, Sh_n^k is T_0 but not T_1.*

Proof As noted earlier, in a T_1 space X $\text{blur}(p) = \{p\}$ $\forall p \in X$. However, in Sh_n^k—which has at least two points since $k \geq 2$ and $n \geq 1$—$\text{blur}(\{\bar{\mathbf{0}}\}) = Sh_n^k$. Thus Sh_n^k is not T_1.

Now let $\bar{\mathbf{v}}, \bar{\mathbf{w}}$ be distinct points in Sh_n^k; without loss of generality assume $\text{level}(\bar{\mathbf{v}}) \geq \text{level}(\bar{\mathbf{w}})$. Let $\mathbf{v}, \mathbf{w} \in H^{n,k}$ be representatives of $\bar{\mathbf{v}}, \bar{\mathbf{w}}$, and first assume $\text{level}(\mathbf{w}) < \text{level}(\mathbf{v})$. Since for all j, the set $\{\mathbf{u} \in H^{n,k} \mid \text{rank}(\mathbf{u}) \geq j\}$ is open, there is an open set $U \subset H^{n,k}$ that contains $\mathbf{v}$ but not $\mathbf{w}$. If $\text{level}(\mathbf{w}) = \text{level}(\mathbf{v})$, then again there exists an open set $U \subset H^{n,k}$ containing $\mathbf{v}$ but not $\mathbf{w}$. In either case the image $\bar{U}$ of U in Sh_n^k is an open set in the quotient topology (since quotient maps by group-actions are open maps) containing $\bar{\mathbf{v}}$ but not $\bar{\mathbf{w}}$. Hence Sh_n^k is T_0. □

An amusing fact is that the two-point space Sh_1^2 is a naturally occurring instance of the simplest topological space that is T_0 but not T_1.

Having analyzed the blurs of interest in Sh_n^k, we move on to closures.

Proposition 5.8 *Let $k \geq 1$. Then the following are true.*

(a) *For $0 \leq i \leq n$, $0 \leq j \leq n$, and all $\bar{\mathbf{v}} \in Sh_{n,j}^k$,*

$$\iota(\text{closure}_i(\bar{\mathbf{v}})) = \{P \in G_i(\mathbf{R}_0^k) \mid \text{the } j\text{-plane } \iota(\bar{\mathbf{v}}) \text{ contains the } i\text{-plane } P\}. \quad (5.5)$$

Hence

$$\text{closure}(\{\bar{\mathbf{v}}\}) \subset \{\bar{\mathbf{v}}\} \coprod \Big(\coprod_{i < \text{level}(\bar{\mathbf{v}})} Sh_{n,i}^k \Big), \quad (5.6)$$

"opposite" to the stratification of $\text{blur}(\bar{\mathbf{v}})$ *(cf.* (5.2) *and* (5.3)).

(b) *For $0 \leq j \leq n$, $\coprod_{i \geq j} Sh_{n,i}^k$ is open; equivalently, for $0 \leq j \leq n$, $\coprod_{i \leq j} Sh_{n,i}^k$ is closed.*

(c) *For $0 \leq j \leq n$,* $\text{closure}(Sh_{n,j}^k) = \coprod_{i \leq j} Sh_{n,i}^k$ *(again "opposite" to the corresponding relation for blurs,* (5.4)). *In particular, for $j > 0$, although $Sh_{n,j}^k \cong G_j(\mathbf{R}^{k-1})$ is compact in the relative topology it is not closed as a subset of Sh_n^k.*

(d) *$\{\bar{\mathbf{0}}\}$ is closed, and is the* only *closed point in Sh_n^k. No point in Sh_n^k is open.*

(e) *For all $\bar{\mathbf{v}} \in Sh_n^k$,* $\text{boundary}\{\bar{\mathbf{v}}\} = \text{closure}(\{\bar{\mathbf{v}}\})$.

Proof

(a) This follows from (5.1), (5.2), and the fact that $\iota_{n,i}^k : Sh_{n,i}^k \to G_i(\mathbf{R}_0^k)$ is bijective.

(b) $\coprod_{i \geq j} Sh_{n,i}^k$ is the image, under the quotient map $H^{n,k} \to Sh_n^k$, of the open set $\{\mathbf{v} \in H^{n,k} \mid \text{rank}(\mathbf{v}) \geq j\}$, hence is open in the quotient topology.

(c) Using Lemma 5.2(a) and Proposition 5.4,

$$\text{closure}(Sh_{n,j}^k) \supset \{\bar{\mathbf{w}} \in Sh_n^k \mid \text{blur}(\bar{\mathbf{w}}) \cap Sh_{n,j}^k \neq \emptyset\} \supset \coprod_{i \geq j} Sh_{n,i}^k.$$

But by part (b) $\coprod_{i \geq j} Sh_{n,i}^k$ is closed, so we have the opposite inclusion as well.

(d) The second statement in part (b), with $j = 0$, implies that $Sh_{n,0}^k = \{\bar{\mathbf{0}}\}$ is closed. For every other $\bar{\mathbf{v}} \in Sh_n^k$, $\text{level}(\bar{\mathbf{v}}) > 0$, and hence $\iota(\bar{\mathbf{v}})$ contains some lower-dimensional subspace. It then follows from (5.5) that closure($\{\bar{\mathbf{v}}\}$) is strictly larger than $\{\bar{\mathbf{v}}\}$. It is easily seen that no orbit of $GL(n)$ in $H^{n,k}$ is an open set, and hence no point of Sh_n^k is open in the quotient topology.

(e) Since $\{\bar{\mathbf{v}}\}$ is not open, the closure of $\{\bar{\mathbf{v}}\}'$ is the entire space Sh_n^k, so boundary$\{\bar{\mathbf{v}}\}$ = closure($\{\bar{\mathbf{v}}\}$). □

6 Geometry of Sh_n^k

Since the space Sh_n^k is not even T_1 it is very far from being metrizable. However, thinking of the top stratum $Sh_{n,*}^k$ as representing "most" of Sh_n^k—which, for dimensional reasons, one should only do if $k > 2n$ (see (4.10) and the surrounding discussion)—it is certainly reasonable to try to put a "nice" distance function on $Sh_{n,*}^k$, or, as we will do, a Riemannian metric. Of course, since each stratum is a perfectly nice manifold one could do this with the lower strata as well, but there seems little point since the strata cannot be joined in a metrically compatible way.

We assume for the rest of this section that $k > n$.

There is at least one natural Riemannian metric we can put on $Sh_{n,*}^k$, namely the one induced by the natural identification $Sh_{n,*}^k \cong G_n(\mathbf{R}_0^k)$ and the standard Riemannian metric on this Grassmannian. Before diving in, however, we pause to review the very different roles played by $\mathbf{R}^n$ and $\mathbf{R}^k$ in Sh_n^k. As we have already observed, $\mathbf{R}^n$ can be replaced by any n-dimensional affine space without changing the canonical identification (4.9) of Sh_n^k with a disjoint union of Grassmannians $G_j(\mathbf{R}_0^k)$. Choices of origin of the affine n-space and basis of the underlying vector space do not affect this identification. No metric or inner-product structure on $\mathbf{R}^n$ entered the identification, so certainly no such structure should be reflected in a natural metric on $Sh_{n,*}^k$.

However, $\mathbf{R}^k$, in contrast to $\mathbf{R}^n$, does not enter the shape-space picture as any old k-dimensional vector space; it comes to us *specifically* as $\mathbf{R}^k$. Thus any object naturally associated with $\mathbf{R}^k$ induces an object *naturally* associated with Sh_n^k. In particular this applies to the standard inner product on $\mathbf{R}^k$ and to the Riemannian metric that it induces, in a standard way we now review, on $G_n(\mathbf{R}_0^k)$.

Let E be a finite-dimensional inner-product space and let $n \le \dim(E)$. The standard Riemannian metric g on $G_n(E)$ is constructed as follows. For $X \in G_n(E)$ there is a natural identification $T_X G_n(E) \cong \mathrm{Hom}(X, X^\perp)$, which arises because every $Y \in G_n(X)$ sufficiently close to X is the "orthogonal graph over X" of a unique linear map $S : X \to X^\perp$. (Here by "orthogonal graph over X" we mean $\{\alpha + S\alpha \mid \alpha \in X\}$.) The inner product on E also induces an inner product on $\mathrm{Hom}(X, X^\perp)$; given two elements S_1, S_2 of this space, their inner product is $g_X(S_1, S_2) := \mathrm{tr}(S_1^\dagger S_2)$, where $S_1^\dagger : X^\perp \to X$ is the adjoint of S_1 with respect to the inner products on X and $X^\perp$ inherited from E. An alternative characterization of the inner product g_X is that the squared norm of $S : X \to X^\perp$ is the sum of the squares of the matrix coefficients of S, where the matrix is taken with respect to orthonormal bases in both X and $X^\perp$. When $E = \mathbf{R}^m$ with the standard inner product, the Riemannian metric g coincides, up to scale, with the metric on $G_n(\mathbf{R}^m) \cong O(m)/(O(n) \times O(m-n))$ induced by Riemannian submersion from the standard bi-invariant Riemannian metric on the orthogonal group $O(m)$ (the invariant metric coinciding at the identity with the negative of the Killing form). This is reminiscent of the situation for the smooth subsets of Procrustes shape spaces, where the Procrustes-Riemannian metric is also determined by Riemannian submersion; cf. [12]. However, in the Procrustean context, the conformal structure of $\mathbf{R}^n$, rather than just its affine structure, enters the metric (through the identification of the shape-spaces as quotients of spheres), in contrast to the affine-shape context.

Although a formula for g on $G_n(\mathbf{R}^{k-1})$ can be written down with relative ease, it is important now to remember that $Sh_{n,*}^k$ is *not* canonically $G_n(\mathbf{R}^{k-1})$, because there is no canonical isomorphism $\mathbf{R}_0^k \to \mathbf{R}^{k-1}$. In order to pull back to $G_n(\mathbf{R}_0^k)$ the metric on $G_n(\mathbf{R}^{k-1})$ written in terms of coordinates on $\mathbf{R}^{k-1}$, we would need an isometry $\mathbf{R}_0^k \to \mathbf{R}^{k-1}$, and again there is no canonical choice. Therefore our goal will be to produce a formula for the metric (which we will still call g) on $G_n(\mathbf{R}_0^k)$ purely in terms of pre-shape data. We will do this via the description of g above, taking $E = \mathbf{R}_0^k$ with the inner product inherited from the standard inner product on $\mathbf{R}^k$, after performing some preliminary computations for general E.

Notation 6.1

(i) Let $(\cdot,\cdot)$ denote the inner product on E.

(ii) For $X \in G_n(E)$ let $\mathrm{proj}_{X^\perp} : E \to X^\perp \subset E$ denote orthogonal projection onto $X^\perp$. Also let $P_{X^\perp}$ be the $k \times k$ matrix, with respect to the standard basis of $\mathbf{R}^k$, of the orthogonal projection $\mathbf{R}^k \to X^\perp \subset E$, and for $\mathbf{w} \in H^{n,k}$ define $\mathrm{proj}^n_{X^\perp}(\mathbf{w}) = \mathbf{w}P_{X^\perp}$. Since $P_{X^\perp}$ is symmetric, $\mathrm{proj}^n_{X^\perp}(\mathbf{w})$ is the matrix in $H^{n,k}$ whose ith row is the image of the ith row of $\mathbf{w}$ under $\mathrm{proj}_{X^\perp}$.

(iii) For any $m \ge 1$, given two lists of vectors $\alpha := \{\alpha_i\}_i^m$ and $\beta := \{\beta_i\}_1^m$ in E, let $M^{\alpha,\beta}$ be the $m \times m$ matrix whose entries are given by

$$(M^{\alpha,\beta})_{ij} = (\alpha_i, \beta_j). \tag{6.1}$$

Given $X \in G_n(E)$, any linear maps $S_1, S_2 : X \to E$, and any basis α of X, we can compute the matrix B of $S_1^\dagger S_2$ with respect to α as follows:

$$\sum_k B_{ki}\alpha_k = S_1^\dagger S_2 \alpha_i$$

$$\Rightarrow \quad \left(\alpha_j, \sum_k B_{ki}\alpha_k\right) = (\alpha_j, S_1^\dagger S_2 \alpha_i)$$

$$\Rightarrow \quad (M^{\alpha,\alpha}B)_{ji} = (S_1\alpha_j, S_2\alpha_i) = (M^{S_1\alpha,S_2\alpha})_{ji}$$

$$\Rightarrow \quad B = (M^{\alpha,\alpha})^{-1}M^{S_1\alpha,S_2\alpha} \tag{6.2}$$

where $S_i\alpha := (S_i\alpha_1, \ldots, S_i\alpha_n)$. Thus the inner product g_X on $T_X(G_n(E)) = \mathrm{Hom}(X, X^\perp)$ is given by

$$g_X(S_1, S_2) = \mathrm{tr}((M^{\alpha,\alpha})^{-1}M^{S_1\alpha,S_2\alpha}) \tag{6.3}$$

for any basis α of X.

We will express this in terms of shape-space data using the standard inner product on the exterior powers of the inner-product space E. This inner product, defined by declaring $\{e_{i_1} \wedge e_{i_2} \wedge \cdots \wedge e_{i_m} \mid i_1 < i_2 < \cdots < i_m\}$ to be an orthonormal basis of $\bigwedge^m(E)$, where $\{e_i\}$ is an orthonormal basis of E, satisfies

$$(\alpha_1 \wedge \alpha_2 \wedge \cdots \wedge \alpha_m, \beta_1 \wedge \beta_2 \wedge \cdots \wedge \beta_m) = \det(M^{\alpha,\beta}) \tag{6.4}$$

where $\alpha = (\alpha_1, \ldots, \alpha_m)$, $\beta = (\beta_1, \ldots, \beta_m)$. Thus for a basis α of $X \in G_n(\mathbf{R}^k_0)$, we have

$$\det(M^{\alpha,\alpha}) = \|\alpha_1 \wedge \alpha_2 \wedge \cdots \wedge \alpha_n\|^2. \tag{6.5}$$

Also note that if $\{\alpha_i\}_1^m, \{\beta_i\}_1^m$ are arbitrary vectors in a subspace $X \subset E$, and η_1, η_2 are arbitrary vectors in $X^\perp$, then

$$(\alpha_1 \wedge \cdots \wedge \alpha_m \wedge \eta_1, \beta_1 \wedge \cdots \wedge \beta_m \wedge \eta_2) = (\alpha_1 \wedge \cdots \wedge \alpha_m, \beta_1 \wedge \cdots \wedge \beta_m)(\eta_1, \eta_2). \tag{6.6}$$

Now take $E = \mathbf{R}^k_0$ and let $\pi : H^{n,k}_* \to G_n(\mathbf{R}^k_0)$ be the canonical projection (the map that takes $\mathbf{v} \in H^{n,k}_*$ to the row space of $\mathbf{v}_{\mathrm{mat}}$). Let $\mathbf{v} \in H^{n,k}_*$, let $X = \pi(\mathbf{v}) \in G_n(\mathbf{R}^k_0)$, and let $\alpha_1, \ldots, \alpha_n$ be the rows of $\mathbf{v}_{\mathrm{mat}}$ (a basis of X). Also let $\mathbf{w} \in H^{n,k}$, let $\{\beta_i\}_1^n$ be the rows of $\mathbf{w}_{\mathrm{mat}}$, let $\eta_i = \mathrm{proj}_{X^\perp}(\beta_i)$, and let $\mathbf{u} = \mathrm{proj}^n_{X^\perp}(\mathbf{w})$. Consider the curve $t \mapsto \pi(\mathbf{v}+t\mathbf{w}) \in G_n(\mathbf{R}^k_0)$. For small t, $\{\alpha_i + t\beta_i\}$ is a basis of the n-plane $\pi(\mathbf{v} + t\mathbf{w})$; hence so is $\{\alpha_i + t\eta_i\}$. Therefore $\pi(\mathbf{v} + t\mathbf{w})$ is the orthogonal graph over X of the unique linear map $S_{\mathbf{v},\mathbf{w},t} : X \to X^\perp$ determined by $S_{\mathbf{v},\mathbf{w},t}(\alpha_i) = t\eta_i$. Note that $S_{\mathbf{v},\mathbf{w},t} = S_{\mathbf{v},\mathrm{proj}^n_{X^\perp}(\mathbf{w}),t} = S_{\mathbf{v},\mathbf{w}P_{X^\perp},t}$. Writing $S_{\mathbf{v},\mathbf{w}} = S_{\mathbf{v},\mathbf{w},1}$ and letting $\pi_{*\mathbf{v}}\mathbf{w} \in T_X(G_n(\mathbf{R}^k_0))$ denote the image under π of the tangent vector $\frac{d}{dt}(\mathbf{v}+t\mathbf{w})|_{t=0} \in T_\mathbf{v}H^{n,k}_*$, we therefore have

$$\pi_{*\mathbf{v}}\mathbf{w} = S_{\mathbf{v},\mathbf{w}P_{X^\perp}} = \pi_{*\mathbf{v}}(\mathbf{w}P_{X^\perp}). \tag{6.7}$$

It follows that if $\mathbf{w}'$ is a second matrix in $H^{n,k}$ then

$$g_X(\pi_{*\mathbf{v}}\mathbf{w}, \pi_{*\mathbf{v}}\mathbf{w}') = g_X(\pi_{*\mathbf{v}}(\mathbf{w}P_{X^\perp}), \pi_{*\mathbf{v}}(\mathbf{w}'P_{X^\perp})) = g_X(S_{\mathbf{v},\mathbf{w}P_{X^\perp}}, S_{\mathbf{v},\mathbf{w}'P_{X^\perp}}). \tag{6.8}$$

Proposition 6.2 *Let* $\mathbf{v}, \mathbf{w}, \mathbf{w}', \{\alpha_i\}, \{\beta_i\}, \{\eta_i\}$ *be as above. Also let* β'_i *be the* i*th row of* $\mathbf{w}'$, *and* $\eta'_i = \mathrm{proj}_{\pi(\mathbf{v})^\perp}(\beta'_i)$, $1 \le i \le n$. *Then*

$$g_{\pi(\mathbf{v})}(\pi_{*\mathbf{v}}\mathbf{w}, \pi_{*\mathbf{v}}\mathbf{w}') = \frac{\sum_{i,j}(\alpha_1 \wedge \cdots \wedge \eta_i \wedge \cdots \wedge \alpha_n, \alpha_1 \wedge \cdots \wedge \eta'_j \wedge \cdots \wedge \alpha_n)}{\|\alpha_1 \wedge \cdots \wedge \alpha_n\|^2} \tag{6.9}$$

$$= \frac{1}{\|\alpha_1 \wedge \cdots \wedge \alpha_n\|^4} \times \left\{\sum_{i,j}(-1)^{i+j}(\alpha_1 \wedge \cdots \widehat{\alpha_i} \cdots \wedge \alpha_n, \alpha_1 \wedge \cdots \widehat{\alpha_j} \cdots \wedge \alpha_n) \times (\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \beta_i, \alpha_1 \wedge \cdots \wedge \alpha_n \wedge \beta'_j)\right\} \tag{6.10}$$

where the hats in (6.10) *denote omission, and where* $\alpha_1 \wedge \cdots \wedge \eta_i \wedge \cdots \wedge \alpha_n$ *in* (6.9) *means* $\alpha_1 \wedge \cdots \wedge \alpha_{i-1} \wedge \eta_i \wedge \alpha_{i+1} \wedge \cdots \wedge \alpha_n$.

Proof (a) Using (6.4), (6.5), and the usual "cofactor" formula for the inverse of a matrix, we have

$$(M^{\alpha,\alpha})^{-1}_{ij} = (-1)^{i+j}\frac{(\alpha_1 \wedge \cdots \widehat{\alpha_i} \cdots \wedge \alpha_n, \alpha_1 \wedge \cdots \widehat{\alpha_j} \cdots \wedge \alpha_n)}{\|\alpha_1 \wedge \alpha_2 \wedge \cdots \wedge \alpha_n\|^2}. \tag{6.11}$$

Using (6.3) and (6.8), we then obtain

$$g_{\pi(\mathbf{v})}(\pi_{*\mathbf{v}}\mathbf{w}, \pi_{*\mathbf{v}}\mathbf{w}') = \frac{1}{\|\alpha_1 \wedge \cdots \wedge \alpha_n\|^2} \times \left\{\sum_{i,j}(-1)^{i+j}(\alpha_1 \wedge \cdots \widehat{\alpha_i} \cdots \wedge \alpha_n, \alpha_1 \wedge \cdots \widehat{\alpha_j} \cdots \wedge \alpha_n) \times (\eta_i, \eta'_j)\right\}. \tag{6.12}$$

Since $\eta_i, \eta'_i \in \mathrm{span}(\{\alpha_i\})^\perp$, (6.6) reduces (6.12) to (6.9). Next, observe that the right-hand side of (6.10) does not change if we replace β_i, β'_j by η_i, η'_j. But after making this replacement, (6.6) reduces the right-hand side of (6.10) to the right-hand side of (6.12). □

We have included both formulas (6.9) and (6.10) since the second formula is expressed directly in terms of preshape space data, with no need for an orthogonal projection, while the first has the advantage of being simpler.

Using the Riemannian metric g, one can obtain a distance-function on the stratum $Sh^k_{n,*}$ the usual way, but we have not found a closed-form expression as simple as the formulas in Proposition 6.2. We can, however, write down a less explicit formula, making use of a known formula for the geodesic-distance function on Grassmannians, as presented in [7], p. 337. For two points $X, Y \in G_n(\mathbf{R}^k)$, let $\mathbf{v}, \mathbf{w} \in \mathbf{R}^{n,k}$ be $n \times k$ matrices representing X, Y respectively, with the columns of each of $\mathbf{v}, \mathbf{w}$ assumed *orthonormal*. Then the diagonal matrix in a singular-values decomposition of the $n \times n$ matrix $\mathbf{v}\mathbf{w}^t$ is of the form $\mathrm{diag}(\cos\theta_1, \ldots, \cos\theta_n)$, where $0 \le \theta_i \le \pi/2$, and the geodesic distance between X and Y in $G_n(\mathbf{R}^k)$ is

$$d(X,Y) = d(\pi^{n,k}(\mathbf{v}), \pi^{n,k}(\mathbf{w})) = \left(\sum_i \theta_i^2\right)^{1/2} \tag{6.13}$$

where $\pi^{n,k} : \mathbf{R}_*^{n,k} \to G_n(\mathbf{R}^k)$ is the natural projection.

Next, observe that for any finite-dimensional inner-product space E and any subspace E_0 of dimension $\ge n$, it follows from our definition of Riemannian metrics on $G_n(E_0)$ and $G_n(E)$ that the embedding $G_n(E_0) \hookrightarrow G_n(E)$ is *isometric*. (This can also be seen directly at the level of distance-functions: writing $E = E_0 \oplus E_0^\perp$ and choosing an adapted orthonormal basis of E, if all the columns of $\mathbf{v}, \mathbf{w}$ lie in E_0 and $\hat{\mathbf{v}}, \hat{\mathbf{w}}$ are the corresponding $n \times \dim(E_0)$ matrices, then $\hat{\mathbf{v}}\hat{\mathbf{w}}^t = \mathbf{v}\mathbf{w}^t$.) In particular this holds with $E_0 = \mathbf{R}_0^k$ and $E = \mathbf{R}^k$. Thus, given arbitrary pre-shapes $\mathbf{v}', \mathbf{w}' \in H_*^{n,k}$, letting $\mathbf{v}, \mathbf{w}$ be representatives of the same shapes but with orthonormal columns (obtained e.g. by Gram-Schmidt), the distance in $Sh_{n,*}^k$ between $\pi(\mathbf{v}'), \pi(\mathbf{w}')$ is given by (6.13).

In closing we note, once again, that there *does not exist* a distance function on any open set in the full space Sh_n^k compatible with the quotient topology, since this space is not even T_1.

References

1. Berthilsson, R., Astrom, K., Heyden, A.: Reconstruction of curves in R3, using factorization and bundle adjustment. In: The Proceedings of the Seventh IEEE International Conference on Computer Vision, pp. 674–679 (1999)
2. Bala, E., Cetin, A.E.: Computationally efficient wavelet affine invariant functions for shape recognition. IEEE Trans. Pattern Anal. Mach. Intell. **26**(8), 1095–1099 (2004)
3. Betke, M., Makris, N.C.: Recognition, resolution, and complexity of objects subject to affine transformations. Int. J. Comput. Vis. **44**(1), 5–40 (2001)
4. Berthilsson, R., Heyden, A., Sparr, G.: Recursive structure and motion from image sequences using shape and depth spaces. In: Proceedings IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 444–449 (1997)
5. Carne, T.K.: The geometry of shape spaces. Proc. Lond. Math. Soc. **16**, 407–432 (1989)
6. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. Comput. Vis. Image Underst. **89**(2–3), 114–141 (2003)
7. Edelman, A., Arias, T.A., Smith, S.T.: The geometry of algorithms with orthogonality constraints. SIAM J. Matrix Anal. Appl. **20**(2), 303–353 (1998)
8. Faugeras, O.: Three Dimensional Computer Vision. MIT Press, Cambridge (1993)
9. Fletcher, P.T., Venkatasubramanian, S., Joshi, S.: Robust statistics on Riemannian manifolds via the geometric median. In: Proceedings CVPR (2008)
10. Hagedoorn, M., Veltkamp, R.C.: Reliable and efficient pattern matching using an affine invariant metric. Int. J. Comput. Vis. **31**(2/3), 203–225 (1999)
11. Kendall, D.G.: Shape manifolds, Procrustean metrics, and complex projective spaces. Bull. Lond. Math. Soc. **16**, 81–121 (1984)
12. Kendall, D.G., Barden, D., Carne, T.K., Le, H.: Shape and Shape Theory. Wiley, Chichester (1999)
13. Klein, G.J., Reutter, R.W., Huesman, R.H.: Four-dimensional affine registration models for respiratory-gated PET. IEEE Trans. Nucl. Sci. **48**, 756–760 (2001)
14. Lamdan, Y., Schwartz, J.T., Wolfson, H.J.: Affine invariance model based object recognition. IEEE Trans. Robot. Autom. **6**, 578–589 (1990)
15. Mardia, K.V., Patrangenaru, V., Derado, G., Patrangenaru, V.P.: Reconstruction of planar scenes from multiple views using affine and projective shape manifolds. In: 2003 IEEE Workshop on Statistical Signal Processing, pp. 298–301, October 2003
16. Milnor, J.W., Stasheff, J.D.: Characteristic Classes. Princeton University Press, Princeton (1974)
17. Ness, L.: A stratification of the null cone via the moment map. Am. J. Math. **106**, 1281–1329 (1984)
18. Patrangenaru, V., Mardia, K.V.: Affine shape analysis and image analysis. In: Aykroyd, R.G., Mardia, K.V., Langdon, M.J. (eds.) Stochastic Geometry, Biological Structure and Images, pp. 57–62 (2003)
19. Pennec, X.: Intrinsic statistics on Riemannian manifolds: basic tools for geometric measurements. J. Math. Imaging Vis. **25**(1), 127–154 (2006)
20. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Imaging **18**, 712–721 (1999)
21. Sepiashvili, D., Moura, J.M.F., Ha, V.H.S.: Affine permutation symmetry: invariance and shape space. In: 2003 IEEE Workshop on Statistical Signal Processing, pp. 307–310 (2003)
22. Small, C.G.: The Statistical Theory of Shape. Springer, New York (1996)
23. Sparr, G.: Simultaneous reconstruction of scene structure and camera locations from uncalibrated image sequences. In: Proceedings of the 13th International Conference on Pattern Recognition, vol. 1, pp. 328–333, August 1996
24. Triggs, B.: Autocalibration from planar scenes. In: Burkhardt, H., Neumann, B. (eds.) Proceedings of the Fifth European Conference on Computer Vision. Lecture Notes in Computer Science. Springer, Berlin (1998)
25. Tieng, Q.M., Boles, W.W.: Wavelet-based affine invariant representation: a tool for recognizing planar objects in 3D space. IEEE Trans. Pattern Anal. Mach. Intell. **19**(8), 846–857 (1997)
26. Yang, Z., Cohen, F.S.: Image registration and object recognition using affine invariants and convex hulls. IEEE Trans. Image Process. **8**(7), 934–946 (1999)