

Affine Bases for Affine Spaces

Supplementary Material

1. Proofs

In the following, $\mathbf{Q}_0, \mathbf{Q}_1 \in \mathbb{R}^{d \times d}$ denote symmetric matrices representing distinct k -flats. Recall that the kernel of such a matrix \mathbf{Q} spans the flat, and the span of \mathbf{Q} spans the normals. We sort the eigenvalues in *ascending* order, so \mathbf{Q} has eigenvalues

$$0 = \lambda_1 = \dots = \lambda_k < \lambda_{k+1} = \dots = \lambda_d = 1. \quad (1)$$

Affine combinations of such matrices (denoted as \mathbf{Q}^* in the main paper) have the following form for the two given matrices:

$$\mathbf{Q}(\mu) = (1 - \mu)\mathbf{Q}_0 + \mu\mathbf{Q}_1. \quad (2)$$

By $\lambda_i(\mu)$, we denote the i th eigenvalue of $\mathbf{Q}(\mu)$ (also in ascending order). We start by proving Lemma 1 from the main paper. Recall:

Lemma 1. *Given two $\mathbf{Q}_0, \mathbf{Q}_1$ that are part of a QDF representation of two flats that are not orthogonal, the matrices $\mathbf{Q}(\mu)$ do not have the eigenvalue μ (except for $\mu \in \{0, 1\}$).*

Proof. We first recall that $\mathbf{Q}_0, \mathbf{Q}_1$ only have the eigenvalues 0 and 1. Assume \mathbf{v} is an eigenvector of $\mathbf{Q}(\mu)$ corresponding to μ . Write $\mathbf{v} = \mathbf{v}' + \mathbf{v}''$, where \mathbf{v}' is in the span of \mathbf{Q}_1 , i.e., $\mathbf{Q}_1\mathbf{v}' = \mathbf{v}'$, and \mathbf{v}'' is in the null space of \mathbf{Q}_1 . Because of the eigenvalue property, we have

$$\mathbf{Q}(\mu)(\mathbf{v}' + \mathbf{v}'') = \mu(\mathbf{v}' + \mathbf{v}'') \quad (3)$$

and by plugging into the definition of $\mathbf{Q}(\mu)$, we get

$$\mathbf{Q}(\mu)(\mathbf{v}' + \mathbf{v}'') = (1 - \mu)\mathbf{Q}_0\mathbf{v} + \mu\mathbf{v}'. \quad (4)$$

Combining this we find

$$\mu\mathbf{v}'' = (1 - \mu)\mathbf{Q}_0\mathbf{v}. \quad (5)$$

Any product with \mathbf{Q}_0 generates a vector in the span of \mathbf{Q}_0 , but \mathbf{v}'' is in the null space of \mathbf{Q}_1 . So unless μ or $1 - \mu$ are zero, either \mathbf{v}'' is in the span of \mathbf{Q}_0 , or it maps \mathbf{v}' to zero – in either case, the flats represented by $\mathbf{Q}_0, \mathbf{Q}_1$ are orthogonal to each other, in contradiction to our assumption. \square

Then, Lemma 2 uses this to conclude that the eigenspace corresponding to the k smallest eigenvalues (i.e, the linear part of the flat defined by the matrix) is uniquely defined.

Lemma 2. *The eigenvalues λ_k and λ_{k+1} of $\mathbf{Q}(\mu)$ are different (assuming non-orthogonal flats and ascending eigenvalues, as above).*

Proof. Eigenvalues are continuous functions of the coefficients of the matrix and the coefficients of $\mathbf{Q}(\mu)$ are continuous in μ . Moreover, by the previous Lemma, $\lambda_i(\mu) \notin \{\mu, 1 - \mu\}$. Since

$$\lambda_i(0) = \lambda_i(1) = \begin{cases} 0 & i \leq k \\ 1 & i > k \end{cases} \quad (6)$$

we get

$$\lambda_i(\mu) \begin{cases} < \min\{\mu, 1 - \mu\} & i \leq k \\ > \max\{\mu, 1 - \mu\} & i > k \end{cases} \quad (7)$$

by continuity. In particular, $\lambda_k(\mu) < \lambda_{k+1}(\mu)$. \square

This is enough to show that the eigenspaces of $\mathbf{Q}(\mu)$ are an injective function of μ .

Theorem 1. *The flats represented by $\mathbf{Q}(\mu)$ and $\mathbf{Q}(\mu')$ for $\mu \neq \mu'$ are different if \mathbf{Q}_0 and \mathbf{Q}_1 represent different flats and are not orthogonal.*

Proof. Recall that the flat being defined is spanned by the sum of eigenspaces corresponding to the smallest k eigenvalues. Let \mathbf{E} span this space for $\mathbf{Q}(\mu)$, then $\mathbf{E}\mathbf{s}$ is a vector in this space. Assume it gets mapped to $\mathbf{E}\mathbf{s}'$ by $\mathbf{Q}(\mu)$. If $\mathbf{Q}(\mu')$ has the same eigenspace, the vector $\mathbf{E}\mathbf{s}$ also gets mapped into the span of \mathbf{E} , albeit to a different representation \mathbf{s}'' . Now consider the two equations

$$\begin{aligned} (1 - \mu)\mathbf{Q}_0\mathbf{E}\mathbf{s} + \mu\mathbf{Q}_1\mathbf{E}\mathbf{s} &= \mathbf{E}\mathbf{s}' \\ (1 - \mu')\mathbf{Q}_0\mathbf{E}\mathbf{s} + \mu'\mathbf{Q}_1\mathbf{E}\mathbf{s} &= \mathbf{E}\mathbf{s}'' \end{aligned} \quad (8)$$

and multiply the first with μ' , the second with μ and subtract. We find that $\mathbf{E}\mathbf{s}$ is also contained in the eigenspace of \mathbf{Q}_1 . By a similar argument, it is also contained in \mathbf{Q}_0 . This implies that \mathbf{Q}_0 and \mathbf{Q}_1 have the same eigenspace, hence represent the same flat. \square

2. Experiment

To give better understanding about the experiment, we provide concrete numbers of the calculations involved and give more examples of different images rendered.

2.1. Base Simplex

The viewing axis is the y -axis that has $\mathbf{a} = (0, 1, 0)^\top$ as spanning vector and $\mathbf{n}_1 = (1, 0, 0)^\top$ and $\mathbf{n}_2 = (0, 0, 1)^\top$ as normal vectors. This is one of many choices of spanning and normal vectors, but the (\mathbf{Q}, \mathbf{r}) -representation is the same for any possible (orthonormal) choice. Its (\mathbf{Q}, \mathbf{r}) -representation is

$$\mathbf{Q} = \mathbf{n}_1\mathbf{n}_1^\top + \mathbf{n}_2\mathbf{n}_2^\top = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (9)$$

and $\mathbf{r} = \mathbf{0}$, as the y -axis passes through the origin. Flattened as a 12-vector, we get

$$(\mathbf{Q}, \mathbf{r}) = (1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0)^\top. \quad (10)$$

As the gradient, we get

$$\nabla(\mathbf{Q}, \mathbf{r}) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sqrt{2} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sqrt{2} \end{bmatrix}, \quad (11)$$

where each column is one flattened partial derivative. Accounting for the symmetry of the \mathbf{Q} -matrix, each column of the gradient is set to have the same length.

Scaling factor The gradient is multiplied with the matrix \mathbf{S}_4 that represents an origin-centered regular 4-simplex, which yields a 12×5 matrix, each column is one direction in the tangent space of the (\mathbf{Q}, \mathbf{r}) -representation of the y -axis. Adding any non-zero multiple of any of the columns to the y -axis' representation will yield to coordinates outside the curved manifold of affine lines. Thus, they have to be projected back onto the manifold of QDFs of lines in 3-space, the result is the five lines that form the initial simplex.

Throughout our experiments, we noticed that the result is very sensitive to the choice of scaling factor. We do not yet have found a systematic way to find suitable scaling factors other than a grid or bisection search.

Model	Source	$\alpha_{\mathbf{Q}}$ [$\times 0.005$]	$\alpha_{\mathbf{r}}$ [$\times 0.005$]
BUNNY	Stanford	2	2
SPOT	Crane et al. [1]	2.4	0.55
DAVID	Hattori et al. [2]	3.65	0.4
HORSE	CyberWare	3.65	0.35
SUIT MAN	Hattori et al. [2]	3.8	0.38
UTAH TEAPOT	Martin Newell	2.4	0.5

Table 1. Scaling factorts for our rendering examples, given as multiples of 0.005.

2.2. More Examples

Fig. 1 shows more examples of meshes that are rendered via our triangulation method. Similar to the BUNNY example, we build a regular simplex around the y -axis. This time, we use two separate scaling factors $\alpha_{\mathbf{Q}}$ and $\alpha_{\mathbf{r}}$ for scaling the \mathbf{Q} - and \mathbf{r} -derivatives (first two and last two columns of the gradient), respectively. The used parameter values are shown in Tab. 1.

References

- [1] Keenan Crane, Ulrich Pinkall, and Peter Schröder. Robust fairing via conformal curvature flow. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013. 3
- [2] Shota Hattori, Tatsuya Yatagawa, Yutaka Ohtake, and Hiromasa Suzuki. Learning self-prior for mesh denoising using dual graph convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 3

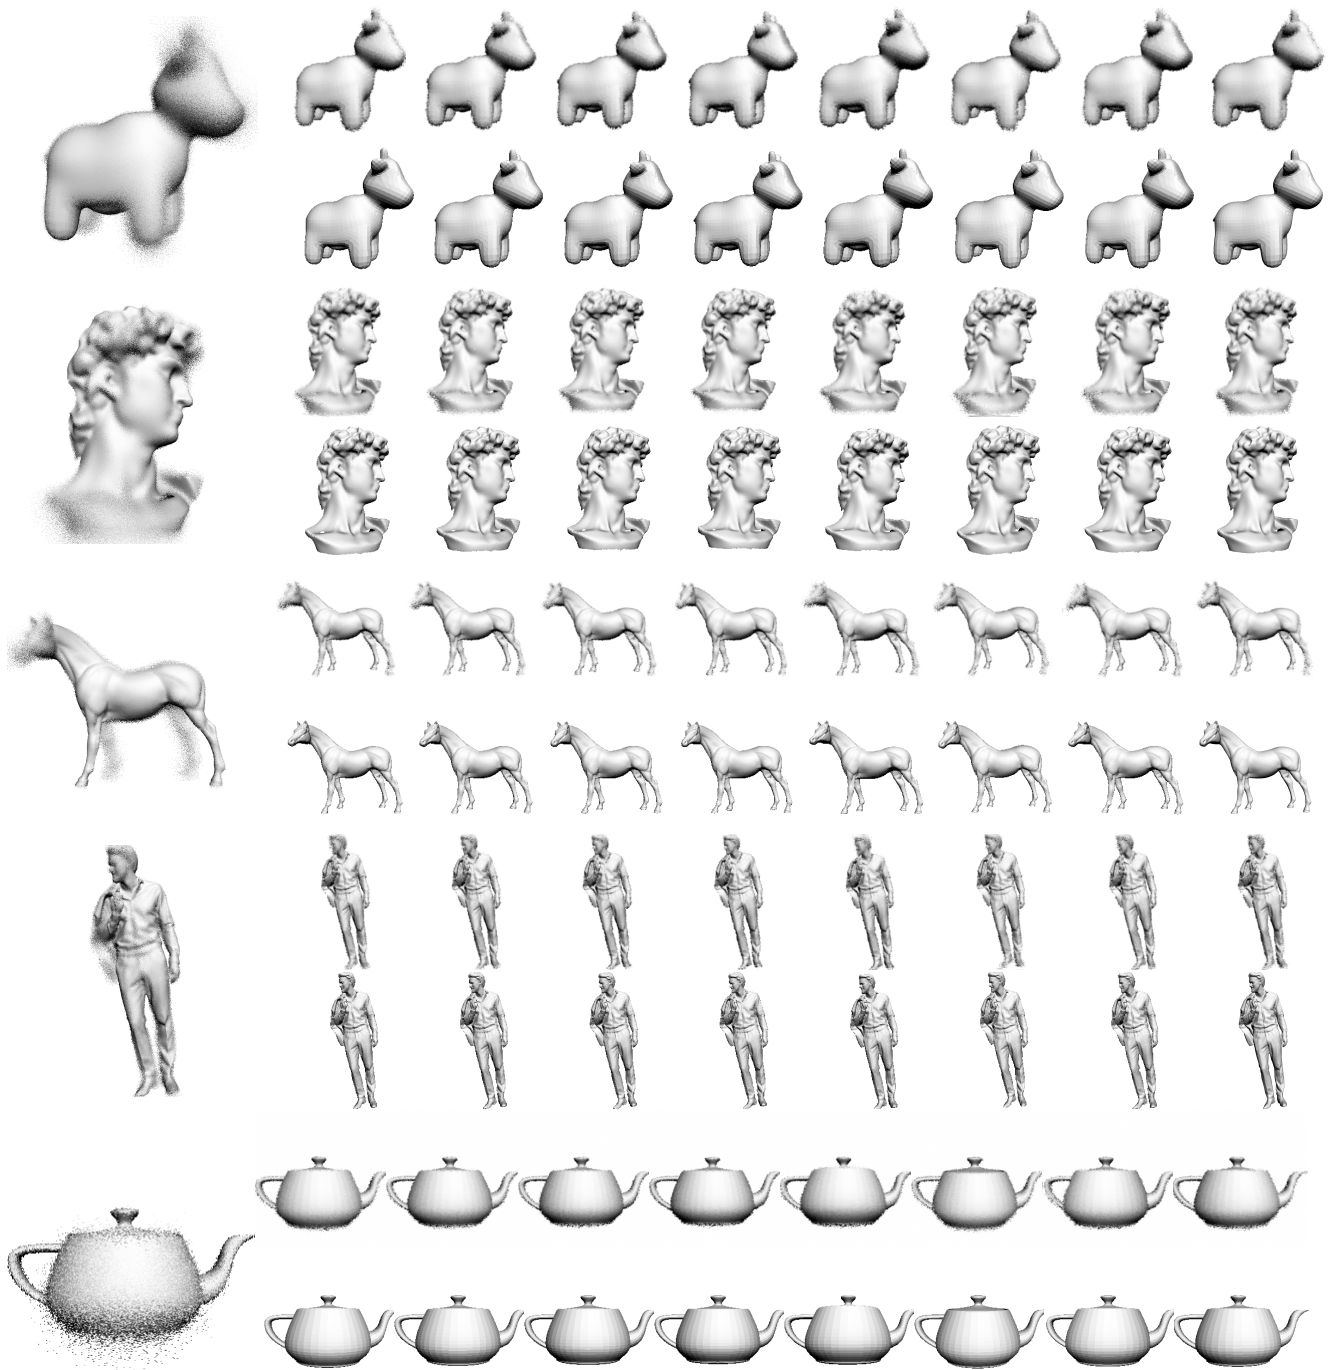


Figure 1. Similar to the Bunny in the main paper, each model is normalized to $[-1, 1]^3$, rotated and sampled on a fixed 1000^2 raster screen in the xy -plane with eye positions chosen uniformly random in $[-1, 1]^2 \times [-4.5, -1.5]$. The resulting gray values (displayed on the fixed grid on the left) are associated to the vertex of a Delaunay triangulation in the space of affine lines. The top rows show reconstructions for varying focal lengths (first four images) and rotating the image screen and eye by $\pm 10^\circ$ along the major axes (last four images). The respective bottom rows show the ground truth. Reconstructions are at the resolution 250^2 and use four-times oversampling.