



HAL
open science

Modelling a feather as a strongly anisotropic elastic shell

Jean Jouve, Victor Romero, Rahul Narain, Laurence Boissieux, Theodore Kim, Florence Bertails-Descoubes

► To cite this version:

Jean Jouve, Victor Romero, Rahul Narain, Laurence Boissieux, Theodore Kim, et al.. Modelling a feather as a strongly anisotropic elastic shell. ACM SIGGRAPH 2024, Jul 2024, Denver, United States. 10.1145/3641519.3657503 . hal-04633434

HAL Id: hal-04633434

<https://hal.science/hal-04633434>

Submitted on 3 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Modelling a feather as a strongly anisotropic elastic shell

Jean Jouve
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LJK
Grenoble, France
jean.jouve@inria.fr

Victor Romero
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LJK
Grenoble, France
victor.romero-gramegna@inria.fr

Rahul Narain
Indian Institute of Technology Delhi
New Delhi, India
narain@cse.iitd.ac.in

Laurence Boissieux
Univ. Grenoble Alpes, Inria
Grenoble, France
laurence.boissieux@inria.fr

Theodore Kim
Yale University
New Haven, United States of America
theodore.kim@yale.edu

Florence Bertails-Descoubes
Univ. Grenoble Alpes, Inria, CNRS,
Grenoble INP, LJK
Grenoble, France
florence.descoubes@inria.fr



(a) Initial Frame

(b) Final frame with our membrane model

(c) Final frame with an isotropic model

Figure 1: Renderings of two simulations of a feather where the base of the rachis is held in place (a) and then the tip of the vane is moved downwards in a direction parallel to the vane (b,c). Inset, the same scenario reproduced on a real feather. This scenario is run both with our new strongly anisotropic membrane model (b) and with an isotropic elastic material (c). Our model captures well the behaviour of the real feather (in particular the quasi-inextensibility of the membrane in the barb direction, causing the bending of the rachis), while the naive isotropic elasticity model is unable to.

ABSTRACT

Feathers exhibit a highly anisotropic behaviour, governed by their complex hierarchical microstructure composed of individual hairs (*barbs*) clamped onto a spine (*rachis*) and attached to each other through tiny hooks (*barbules*). Previous methods in computer graphics have approximated feathers as strips of cloth, thus failing to capture the particular macroscopic nonlinear behaviour of the feather surface (*vane*). To investigate the anisotropic properties of a feather vane, we design precise measurement protocols on real feather samples. Our experimental results suggest a linear strain-stress relationship of the feather membrane with orientation-dependent coefficients, as well as an extreme ratio of stiffnesses in the barb and barbule direction, of the order of 10^4 . From these findings we build a simple continuum model for the feather vane, where the

vane is represented as a three-parameter anisotropic elastic shell. However, implementing the model numerically reveals severe locking and ill-conditioning issues, due to the extreme stiffness ratio between the barb and the barbule directions. To resolve these issues, we align the mesh along the barb directions and replace the stiffest modes with an inextensibility constraint. We extensively validate our membrane model against real-world laboratory measurements, by using an intermediary microscale model that allows us to limit the number of required lab experiments. Finally, we enrich our membrane model with anisotropic bending, and show its practicality in graphics-like scenarios like a full feather and a larger-scale bird. Code and data for this paper are available at <https://gitlab.inria.fr/elan-public-code/feather-shell/>.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGGRAPH Conference Papers '24, July 27-August 1, 2024, Denver, CO, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0525-0/24/07.
<https://doi.org/10.1145/3641519.3657503>

ACM Reference Format:

Jean Jouve, Victor Romero, Rahul Narain, Laurence Boissieux, Theodore Kim, and Florence Bertails-Descoubes. 2024. Modelling a feather as a strongly anisotropic elastic shell. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers '24 (SIGGRAPH Conference Papers '24)*, July 27-August 1, 2024, Denver, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3641519.3657503>

1 INTRODUCTION

The unique mechanical behaviour and optical complexity of feathers gives birds their characteristic appearance. They have also been distinctive components of garments for centuries, from traditional Native American and First Nations decorations to modern-day accessories. Despite the fact that feathers are composed of hundreds to thousands of precisely aligned fibers, previous computer graphics works have uniformly approximated feathers as strips of isotropic cloth. While this approximation suffices for scenarios where the feathers interact in simple ways, it ignores a variety of visually important mechanical behaviors, such as anisotropic stiffness, and disparity in stiffness between the central spine (*rachis*) and its attached fibers (*barbs*).

Relying on experiments conducted at the micro-scale, we introduce a simple macroscopic feather model based on anisotropic elastic shell models that captures these phenomena. Our model has been carefully validated against real-world laboratory experiments, and is capable of capturing intricate and subtle motions that we have not seen elsewhere in the literature.

Our main contributions are as follows:

- A characterisation of the anisotropic in-plane elasticity of a feather, arising from a unique set of micro-scale experiments conducted on a large set of feather samples, which reveals, on the one hand, a simple linear strain-stress relationship of the feather membrane with orientation-dependent coefficients, and on the other hand, an extreme ratio of stiffnesses in the barb and barbule direction, of the order of 10^4 ;
- The derivation of a local linear orthotropic elastic membrane based on a minimal set of three parameters whose anisotropy falls far outside the range of usual models for cloth;
- A robust numerical implementation of this new strongly anisotropic membrane model, which addresses severe locking and ill-conditioning issues thanks to an adapted feather mesh and an inextensibility constraint in the barb direction;
- The careful validation, both quantitatively and qualitatively, of our model against micro-scale experiments and simulations conducted on feather samples.

We also add a plausible anisotropic bending energy for simulating full-size feathers subject to various motions and constraints. Despite its apparent simplicity, our method, implemented within the ARCSim simulator [Pfaff et al. 2014], proves to capture real-world scenarios more faithfully than standard isotropic models or mildly anisotropic models (see Figure 1 and our accompanying video).

2 RELATED WORK

In addition to their unique mechanical properties, feathers possess a variety of rich geometric appearances and motions, which makes them especially attractive for computer graphics applications such as bird animations and decorative objects.

2.1 Related work in graphics

The layout and grooming of multiple feathers [Liu et al. 2015, 2017; Baron and Patterson 2019] and feather precursors [Chuong et al. 2000] such as scales [Landreneau and Schaefer 2010] has been a subject of interest for many years. Film production houses have built

many custom systems, including ones from ILM [Bowline and Kačić-Alesić 2011], DreamWorks Animation [Weber and Gornowicz 2009; Augello et al. 2019], Animal Logic [Heckenberg et al. 2011], and MPC [Haapaoja and Genzwürker 2019]. Across all of these systems, the feathers are modeled as isotropic strips of cloth; none undertake an in-depth examination of anisotropic feather mechanics. The individual fibers on a feather were modeled by Streit and Heidrich [2002], but to our knowledge, no attempt was made to simulate the model. Simulations of feathers in the wind [Wei et al. 2003] have aimed to capture tumbling dynamics. Bird flight animation controllers [Wu and Popović 2003] also operate at the cloth level, and some even assume featherless dragon wings [Won et al. 2017].

We will build on previous works that use experimental measurements to guide both model design and parameter tuning. Both Wang et al. [2011] and Miguel et al. [2012] measured forces and displacements at multiple points of a cloth sample, and then ran a gradient descent optimisation to obtain the parameters for an orthotropic, in-plane stretching model. Both Clyde et al. [2017] and Sperl et al. [2022] used simple stretching experiments to measure the force and displacement at just one end of the sample, and also used gradient descent to fit the data to cloth parameters. Volino et al. [2009] also ran simple stretching experiments with a few measurements, but instead of running an optimisation, derived model parameters directly from the measurements. Unfortunately, most of these experiments do not transfer to feathers.

On the rendering side, a variety of works have attempted to quantify and capture the visually complex, iridescent, anisotropic reflectances that arise from the highly aligned fibers in feathers [Chen et al. 2002; Harvey et al. 2013; Huang et al. 2022; Baron et al. 2022]. Our work on the simulation of feathers complements these investigations into feather appearance.

2.2 Studies of feather mechanics

In mechanical engineering, some works have measured the physical parameters of feathers. However, they either focus on the rachis [Purslow and Vincent 1978; Macleod 1980] or compute the Young's modulus of the underlying feather material, β -keratin [Bonser and Purslow 1995; Wool 2011]. Both of these measurements are insufficient for our purposes.

To the best of our knowledge, models from physics, mechanics and biology represent each fiber (barb) individually, which makes simulating a full feather computationally prohibitive [Kovalev et al. 2014; Chen et al. 2016; Zhang et al. 2018]. Taking inspiration from the experimental setup of Kovalev et al. [2014] and Zhang et al. [2018], we instead design our own experiments for studying the stretching behaviour of feathers. We will use the insights from these experiments to guide the design of our anisotropic shell model.

3 FEATHER STRUCTURE AND MECHANICS

A feather consists of a thick shaft called the *rachis* from which thinner rods called *barbs* emanate on either side, forming two sheet-like structures called vanes (see Figure 2). Both the rachis and the barbs taper in thickness from base to tip, leading to significant variation in rigidity across the vane [Purslow and Vincent 1978]. Each barb has many *barbules* which can hook onto barbules from




Figure 2: Hierarchical structure of a feather. The central shaft is the *rachis*, and the hundreds of rods that branch off it are *barbs* (left). Each barb has hundreds of *barbules*. The barbules along one side possess hooks, while those on the other do not, allowing barbs to reversibly attach and detach (right). The attached barbs form a surface called the *vane*.

adjacent barbs in the vane, in a manner similar to Velcro. Applied forces can form or break individual connections between barbules.

This multiscale structure gives rise to unique mechanical behaviours. For example, suppose each barb’s barbules are hooked onto those of its neighbours. The resulting cohesion between barbs causes the vane to behave as a continuous sheet, but with highly anisotropic and spatially varying elastic properties: it is far stiffer in the orientation parallel to the barbs and closer to the rachis. If the vane is subject to sufficient tension perpendicular to the barbs, the connections between barbules can detach, resulting in an apparent fracture of the vane parallel to the barb orientation.

We refer to this continuum, sheet-like view of feather mechanics as the “macroscopic” view, in contrast to the “microscopic” interactions of individual barbs and barbules. In the following, we aim to characterise the behaviour of the feather at the macroscopic scale and derive a simple anisotropic shell model that captures the most salient properties of a feather.

4 MEASUREMENT OF IN-PLANE ELASTICITY

To study the behaviour of the feather vane at the macroscopic scale, we propose systematic tensile tests on real feather samples at various orientations. To the best of our knowledge, this is the first time that such tests have been performed on real feathers.

4.1 Experimental setup

Performing reproducible tensile tests on real feathers is particularly challenging given the variability of mechanical parameters between species, and the highly heterogeneous nature of biological materials. To avoid species variability, we perform experiments on commercially available 25 cm swan feathers, as well as on a few guinea fowls¹ (see Figure 3, top). From each feather we cut rectangular patches with three different barb orientations, and to have a statistical distribution of the mechanical response, we take sample patches from different regions of a feather.

¹We preferred to work with swan feathers because these show the most regular barb distribution, hence they serve as a model biological system to test our assumptions. Our validation protocol only uses swan samples.

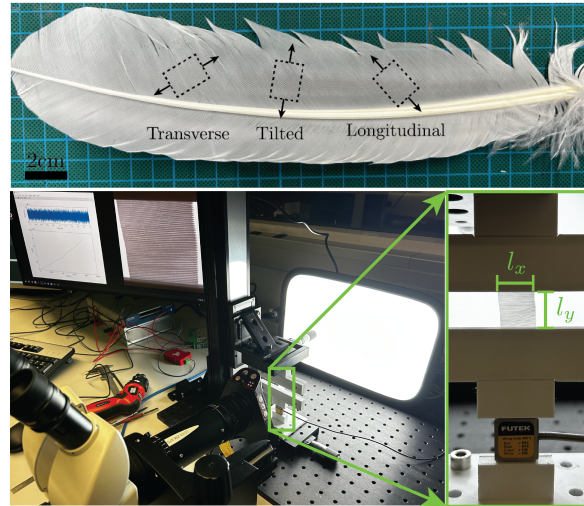


Figure 3: Top: Real feather and example patches. Bottom left: Experimental setup for traction tests. Bottom right: Zoom on the traction experiment.

When subjected to load-unload cycles we observe hysteresis, which is inherent to the frictional nature of the connection between barbules. However, we limit our work here to a non-dissipative model, and only present pure traction experiments (without cycling) from which we can obtain elastic parameters.

In anticipation of validating our macroscopic model against a microscopic model where barbs are represented explicitly (see subsection 8.1), we strive to measure through our tensile tests both the barb and barbule stiffnesses, as well as their geometrical distribution. Towards this aim, we devise an experimental setup that allows us to measure force-displacement relationships with high accuracy (Figure 3, bottom). See supplementary document, section 2 for more details.

4.2 Methodology

We test, in pure traction, patches cut with three different orientations (see images in Figure 13, and top row of Figure 9). To avoid damaging the feather sample in the cutting process, we use a paper guillotine. Samples are then clamped to two edges: during traction, one remains fixed and attached to the force sensor while the other, attached to the motorized stage, is pulled up. Each run is performed in consecutive steps up to a target maximal strain. In each step we displace the motor to a given position, stop, and make a measurement for 2 second at 5 KHz. Hence, one data point is the average of 10,000 measures, allowing us to make a highly sensitive force estimation.

Each sample is characterised by its width l_x , height l_y , thickness h , and barb orientation θ with respect to the sample edges. We assume the thickness and barb orientation to be uniform across the sample. The width is measured on the sample once it is cut, while the height is computed from the position of the motor at the beginning of a traction test. To estimate one thickness value per sample, we compute the average over several measurements taken using a micrometer on different regions of the sample patch. To avoid any difference in mechanical response due to damage by the




Figure 4: Surface coverage of barbs and barbules. Left: the segmented barbs (red) and barbules (green). Right: the original image of the patch region. Everywhere, the vane is locally symmetric with respect to the axis tangent to the barbs (-----).

micrometer, we perform thickness measurements after the traction test is done.

4.3 Observations

The results of our force-displacement measurements are shown in the inset in Figure 13. As expected, we see significant variability from sample to sample. To define global mechanical parameters for the elastic response of feathers, we use the Green-Lagrange strain tensor,

$$\epsilon = \frac{1}{2} \left(\left(\frac{d}{l_y} + 1 \right)^2 - 1 \right) \quad (1)$$

and the second Piola-Kirchhoff stress tensor,

$$\sigma = \frac{l_y}{l_x(d + l_y)} \frac{F}{h} \quad (2)$$

where F is the measured force and d is the measured displacement. For simplicity, we will refer to the Green-Lagrange strain tensor as the *strain tensor* and the second Piola-Kirchhoff stress tensor the *stress tensor*.

In Figure 13, we show that the data collapses remarkably well (for a heterogeneous biological material), whether for the swan species or the guinea fowl species, and we remark that the average thickness h is the key scaling factor. A first important result is the *linear* stress-strain relationship that we can observe, in all three orientations, prior to fracture. A second important result is the extreme ratio that we observe, both for swan and guinea fowl feathers, between the barb and barbule stiffnesses, of the order of 10^4 . It is also worth noticing that the traction response becomes stiffer as we reach the 90° orientation, and fracture nucleates very early in the tilted experiment driven by the shearing between barbs. The fracture is not brittle, which indicates significant energy dissipation prior to fracture, possibly due to inter-barbule friction.

Finally, we estimate the density of barbs and barbules by analysing 50 zoomed images from different regions on a feather vane (Figure 4 (right)). Intensity thresholding is performed to segment the barbs from the barbules (green and red, respectively, in Figure 4 (left)). The threshold value is chosen manually per image to avoid including the dark middle parts corresponding to the interlaced barbules.

From these measurements, we computed the average surface density of barbs and barbules to be $\Phi_B = 0.16 \pm 0.03$ and $\Phi_b = 0.84 \pm 0.03$ respectively.

5 A SIMPLE ANISOTROPIC VANE MODEL

Inspired by our experimental results on uniform vane samples, we propose a simple continuum model for a feather vane that relies on an anisotropic, membrane model that requires only three macroscopic parameters.

5.1 Uniform barb orientation

We first restrict ourselves to the case where all barbs share the same orientation, similar to our experiments in section 4. Experimental results show a linear relation between the stress and the strain tensors for tensile deformations at 0° , 90° , and 45° . We extrapolate this behaviour to all deformations and model the elasticity using a linear relationship between stress and strain. In Voigt's notations, this reads as

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix} = \underbrace{\begin{bmatrix} E_{11} & E_{12} & E_{13} \\ E_{12} & E_{22} & E_{23} \\ E_{13} & E_{23} & E_{33} \end{bmatrix}}_E \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ 2\epsilon_{xy} \end{bmatrix}. \quad (3)$$

For this model to respect energy conservation, the matrix E , commonly called the *elasticity tensor*, must be symmetric. The elasticity tensor generally has six parameters, but we can reduce them to three using the following geometric and physical considerations specific to feathers.

At any point on the vane of a feather, the vane is locally symmetric around the axis collinear to the barbs (See Figure 4). This symmetry, usually called *orthotropy*, removes some components from the elasticity tensor [Milton 2002, p. 37], namely $E_{13} = 0$ and $E_{23} = 0$.

Moreover, the stretched sample in the transverse experiment does not show any deflection along the x -axis. Since no forces are applied on the side of the sample, this implies that feather vanes have zero Poisson's ratio, that is, $E_{12} = 0$. This last property is expected from the stiffness of barbs. When the barbules are stretched, they cannot compress the barbs they are attached to, which prevents deformation in the direction transverse to the deformation.

With these parameters removed, our analytical vane model is left with only three parameters, giving,

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} E_{xx} & 0 & 0 \\ 0 & E_{yy} & 0 \\ 0 & 0 & E_{xy} \end{bmatrix} \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ 2\epsilon_{xy} \end{bmatrix}, \quad (4)$$

where we renamed the parameters of the elasticity tensor to be closer to the notations of the stress and strain that they relate. We will refer to E_{xx} as the *longitudinal modulus*, E_{yy} as the *transverse modulus* and E_{xy} as the *shear modulus*.

5.2 Non-uniform barb orientation

The barb orientation is not uniform over a vane, and therefore we cannot simply choose a coordinate system in which the elasticity tensor is of the above form at all points on the vane. Instead, given the barb orientation θ at a point, we perform a change of basis of

the previously defined tensor \mathbf{E} so that the stiffest orientation, E_{xx} here, is aligned with θ . Then, we can derive stretching forces from the areal energy density

$$\frac{1}{2} h \boldsymbol{\epsilon} : \mathbf{E}(\theta) : \boldsymbol{\epsilon}, \quad (5)$$

where h is the thickness of the vane and $\mathbf{E}(\theta)$ is the rotated elasticity tensor.

This elasticity model is close to the usual linear orthotropic elastic materials seen in cloth simulation up to the following differences: our model is heterogeneous across the shell and, due to the extreme stiffness ratio between barb and barbules and the structure of the vane, the range of parameters value needed to properly model feathers is widely different to the range of parameters used to model cloth. Anticipating results of subsection 8.1, we checked that the behaviour obtained with a standard orthotropic cloth material taken from Wang et al. [2011] is indeed very different from ours, see also our accompanying video.

6 NUMERICS: LOCKING, ILL-CONDITIONING

We implemented our anisotropic model numerically in the freely available thin elastic plate simulator ARCSim [Pfaff et al. 2014], to which we added a full nonlinear solver based on Ipopt [Wächter and Biegler 2006] (see supplementary document, section 3). In this implementation, triangular and linear Lagrange elements are used.

Unfortunately, two serious issues rapidly became apparent in our simulations. First, we identified the presence of *numerical locking*, i.e. the higher the stiffness ratio between the two material directions, the slower the convergence of the finite element model² (see supplementary document, section 4). When using the extreme value of 10^4 , which corresponds to our experimentally measured (see section 4) barb/barbule stiffness ratio, the cost of convergence becomes prohibitive (see Figure 5 —●—). Second, this extreme ratio also makes the problem *ill-conditioned*, which considerably impacts the convergence of the nonlinear solver.

Mesh alignment. To mitigate the locking phenomenon, we take inspiration from the work of [Yu et al. 2006] and [ten Thije and Akkerman 2008], which consider a material with two very stiff directions relative to its shear stiffness. One of their solutions is to align the mesh elements along the directions of high stiffness. We apply this in our context, and ensure that each triangular element has at least one edge aligned along the barb direction. As expected, this alignment considerably improves convergence (see Figure 5 —■—).

Constraint in the barb direction. Aligning the mesh with the stiffest direction mitigates the locking phenomenon, but does not prevent the problem of ill-conditioning at large stiffness ratios. As a result, an undesirably large number of iterations are needed for the non-linear solver to converge (see Figure 6 — and —). Our solution is to directly introduce the limit case of inextensibility along the stiffest direction. We replace the energy component in the high stiffness direction with a constraint that prevents extension along

²This locking phenomenon is not specific to the ARCSim implementation, as we have checked that it also arises when using a thin membrane model implemented within the FENICS library [Hale et al. 2018].




Figure 5: Stress σ for a deformation of 5 % in the tilted scenario, as a function of the number of elements. The stiffness ratio is fixed to our experimental value $\frac{E_{xx}}{E_{yy}} = 10^4$. The convergence for simulations using an arbitrary mesh (—●—) is considerably worse than aligned meshes both without (—■—) and with (—●—) the inextensibility constraint. The reference value (---) is the one obtained with an unconstrained simulation based on an aligned mesh consisting of 128×10^3 elements.




Figure 6: Left: Stress-strain curves for the tilted scenario simulated with an aligned mesh, both with and without the inextensibility constraint. The unconstrained calculations were made with 50 (—) and 100 (—) non-linear iterations. The constrained version (---) used at most 50 non-linear iterations. Inset: due to poor conditioning, unconstrained results tend to oscillate. Right: with the inextensibility constraint, oscillation-free solutions are computed at much lower cost.

the same direction (see supplementary document, section 5 for more details). This modification does not degrade the convergence of the elements (see Figure 5) and allows us to remove oscillatory artifacts that arise from the poor conditioning (see Figure 6).

7 SIMULATING A FULL FEATHER

Similar to previous works on shell simulation [Grinspun et al. 2003; Chen et al. 2018], we represent a feather as a triangle mesh with internal membrane and bending forces. We model the rachis as a distinguished thin region of the feather mesh, with a highly




Figure 7: In the mesh of the feather, the rachis (orange) is distinguished from the vane (light blue) and has a different material from the vane. The barb orientation is given as splines by the user (dark blue), those are used to generate a mesh aligned with those orientation.

stiff isotropic material (see Figure 7). We chose this approach for simplicity to remain within the sheet simulation framework. One could alternatively model it as a stiff elastic rod [Bertails et al. 2006; Bergou et al. 2008; Bertails 2009] coupled with the vane, although this would require solving the coupled dynamics.

In the rest of this section we discuss the simulation model used for the rest of the feather sheet.

Non-uniform anisotropy. As mentioned earlier, the barb orientation is non-uniform across the vane. This non-uniformity is supplied by the user through splines in texture coordinate. Since we need the mesh elements within the vane to be aligned with the barb direction, they are generated such that most of them have at least one edge aligned with the splines (See Figure 7 and supplementary document, section 6, for more details).

Anisotropic membrane. The barb orientation field is chosen to be piecewise constant per faces. Since our membrane energy density defined in Equation 5 is orthotropic per faces, it is already handled by ARCSim.

Anisotropic bending. Bending forces are derived from an energy associated to each edge,

$$\frac{A_1 + A_2}{6} EI(\beta) (\kappa - \kappa_0)^2, \quad (6)$$

where A_1 and A_2 are the area of the triangles incident to the edge, κ and κ_0 are respectively the curvature and rest curvature at the edge, and $EI(\beta)$ is the bending modulus. The curvature is computed through the method described in Grinspun et al. [2003]. The bending modulus depends on the angle β made between the edge and the local barb orientation in reference space, which we call the *tilt angle*.

The directional variation of the bending modulus is parametrized by the longitudinal bending modulus EI_{\perp} (bending the barbs themselves) and the transverse bending modulus EI_{\parallel} (bending the vane keeping the barbs straight). The relation between the bending modulus and the tilt angle is chosen such that $EI(0) = EI_{\parallel}$ and $EI(\frac{\pi}{2}) = EI_{\perp}$:

$$EI(\beta) = \frac{EI_{\parallel} - EI_{\perp}}{2} \cos(2\beta) + \frac{EI_{\parallel} + EI_{\perp}}{2}. \quad (7)$$

For feathers, EI_{\parallel} is chosen much smaller than EI_{\perp} so that deformations that avoid bending of barbs are favored.

Since the barb orientation is defined constant on faces, we have two tilt angles β_1 and β_2 , one for each face. The tilt angle is obtained




Figure 8: The microscale model is composed of bands of isotropic elastic material that are alternately stiff (blue) and supple (red). Both materials have zero Poisson's ratio. The Young modulus of the stiff bands is denoted by E_B , and that of supple bands by E_b . This model is also parametrized by the distance between barb centers d_B and the width of barbs w_B , which allows us to define the barb surface density as $\Phi_B = \frac{w_B}{d_B}$.

through an area based average of both angles,

$$\beta = \frac{A_1 \beta_1 + A_2 \beta_2}{A_1 + A_2}. \quad (8)$$

8 VALIDATION AND RESULTS

We validate our macroscopic membrane model quantitatively against real and virtual traction tests of feather samples conducted at the microscale, before presenting some visual comparisons on a full feather. Finally, we show that our approach extends to the realistic simulation of a larger scale bird with flapping wings. Please watch our accompanying video for animated results. Our implementation is released freely at <https://gitlab.inria.fr/elan-public-code/feather-shell/> as an updated version of ARCSim.

The method used to choose the parameters is detailed in supplementary document, section 7.

8.1 Quantitative validation on feather samples

Our local orthotropic membrane model is characterised by three independent parameters, E_{xx} , E_{yy} , and E_{xy} . While the first two parameters can be measured on longitudinal and transverse traction tests respectively³, the last one, E_{xy} , requires a tilted traction test (involving shearing) to be estimated. This would require us to design yet another tilted traction experiment so as to perform an independent validation of the model.

Given the difficulty performing multiple tilted traction tests on feather samples at different angles, we choose not to validate our membrane model directly on real experiments. Instead, we introduce an intermediary, so-called *microscale* numerical model, validated on a unique experimental tilted traction test, which is then meant to serve as a reference for validating our macroscale model at any angle (in the limit of admissible angles, see below).

Microscale model. The microscale model reproduces the microstructure of the vane at the barb scale, through a non-homogeneous shell consisting of stiff bands representing the barbs, connected by supple bands representing the barbules. The materials of both bands are isotropic elastic with zero Poisson's ratio (see Figure 8).

The parameters of this model are measured from the data gathered in our experiments. The ratio between the width of barbs w_B

³Actually, with the inextensibility constraint, E_{xx} does not need to be estimated anymore and is simply set to the same value as E_{yy} .

and the distance between barbs d_B is set to the measured surface density of barbs Φ_B , i.e. $\Phi_B = 0.16 \pm 0.03$.

This leaves two unknown parameters, the Young modulus of barbs E_B , and that of barbules E_b . Among our experiments (longitudinal, transverse, and tilted traction tests), the tilted experiment yields the most complex deformation. Hence, we leverage it as a demanding, independent validation target while the first two experiments are used to measure E_b and E_B .

In the longitudinal scenario, it can be shown that the slope⁴ of the microscale model $\Delta_{||}$ is related to the material parameters through

$$\Delta_{||} = \Phi_B E_B + (1 - \Phi_B) E_b. \quad (9)$$

The slope of the longitudinal experiment determines $\Delta_{||}$, leaving one degree of freedom between E_b and E_B .

In the transverse scenario, we are unable to obtain a similarly simple expression relating the material parameters to the stress-strain data. Instead, we determine the parameter values by fitting a numerical simulation to the observed slope in the transverse experiment. We set up the microscale model in ARCSim and perform the transverse experiment numerically using different values of E_b (fixing E_B as well as using Equation 9), record the slope of the computed stress-strain relationship, and perform bisection to find the parameter values for which the slope matches that of the experimental data. The measured parameters are listed in Table 1.

In Figure 12 (top), we show the comparison between real and virtual traction tests using our microscale model. As expected, our model shows excellent agreement with the longitudinal and transverse tests. More remarkably, the agreement remains excellent on the tilted scenario (right), which is an independent test *not* used for measurement. Given these good results, we now consider our microscale model to be sufficiently predictive so that it can be used reasonably as a reference for validating our macroscale model.

Note that we sometimes observe that barbules rearrange themselves as the material is sheared in some tilted scenario performed in the lab (see supplementary document, section 12, for more details). Those plastic (irreversible) events cause a drop in stress and are not taken into account by our model which is purely elastic. In our validation, we specifically chose an experiment where those plastic events do not appear, as dealing with them is out of the scope of this paper and is left for future work.

Macroscale model. The macroscale model parameters are measured from the microscale model. The longitudinal and transverse moduli E_{xx} and E_{yy} are equal to the slopes in the longitudinal and transverse scenarios, respectively, as simulated with the microscale model. This leaves the shear modulus E_{xy} to be measured⁵.

Using the same method used for microscale model, we measure this parameter through bisection, performing macroscale model

⁴In our measurement protocols, we do not use directly stress-strain data points, instead we use the slope of those sets obtained through linear regression. Hence, when the slope of an experiment or a simulation is mentioned, this denotes the slope obtained through a linear regression of the generated stress-strain data points.

⁵An ideal and elegant solution would have consisted in retrieving this modulus analytically from the stress-strain curve in the tilted experiment. However we were unable to do so because finite-size effects on the border of the sample prevented us from finding a closed-form solution to the associated elasticity problem. The previous work of Volino et al. [2009] did this for cloth by using samples with a large size aspect ratio, which reduces the impact of finite-size effects. Such a protocol is not transferable to feathers, as the size of our samples are constrained by the vane size.

Φ_B	E_b [N m ⁻²]	E_B [N m ⁻²]	E_{yy} [N m ⁻²]	E_{xy} [N m ⁻²]
0.13	4.0×10^4	2.8×10^9	4.7×10^4	2.9×10^4
0.16	3.7×10^4	2.3×10^9	4.5×10^4	3.1×10^4
0.19	3.6×10^4	1.9×10^9	4.6×10^4	3.4×10^4

Table 1: Microscale and macroscale model measurements. The macroscale measurements were obtained by fitting to the microscale model parameters.

simulations with different values of E_{xy} to match the slope with that of the microscale model in the tilted scenario.

In Figure 11, we eventually show that the macroscale model, whose parameters have been fitted from the microscale model on the longitudinal, transverse and tilted traction tests, matches the same microscale model on traction scenarios with barbs oriented at various angles. This shows the validity of our model. The measured parameters for the macroscale model are listed in Table 1.

In addition to this quantitative validation, we also visually compare, on the tilted traction test, the geometry of our macroscale model, microscale model, and real experiment (see Figure 9). Remarkably, we observe not only the same contours of the samples during extension, but also the very similar appearing of a shear band (in the experiment and in the macroscale model), which increases in width as the aspect ratio of the feather sample is increased. These observations nicely complete our quantitative agreements on forces with qualitative agreements on shapes.

8.2 Qualitative results

Qualitative validation on a full feather. In Figure 1 we visually compare footage of a real feather with results from a simulation using our macroscale model. Our model is run with the previously measured membrane parameters obtained for a surface density of barbs $\Phi_B = 0.16$ (see Table 1) and with bending moduli $EI_{||} = 10^{-7}$ N m and $EI_{\perp} = 2.5 \times 10^{-4}$ N m.

In this scenario, a feather is held at the base of its rachis while its tip is moved downwards, in a direction parallel to the vane. We compare our simulator to a standard isotropic elasticity model. The isotropic model uses a Young modulus E equal to the transverse modulus E_{yy} of our model and a zero Poisson ratio. In the isotropic simulation, the bending is simulated the same as with our model. In addition, in our accompanying video we also perform simulations using a mildly anisotropic model with no inextensibility constraint (see supplementary document, section 9). Results show that unlike the isotropic and mildly anisotropic models, our solution matches the behaviour of the real feather. In particular, it allows to recover the bending of the rachis, while other models are unable to. Moreover, the isotropic simulation took 3 min to be completed, compared to 5 min for the anisotropic one, on the same computer. This shows that our model speed has the same order of magnitude as a usual ARCSim simulation.

Larger scale scenario: the flapping wing example. Finally, we apply our model to a large scale scenario featuring a bird with two flapping wings made of 10 feathers each, see Figure 10. To simulate one animation cycle of the bird, 432 time-steps of $2.8 \mu\text{s}$ each were necessary to handle collision. Each cycle took 8.5 hours to be computed with our model.

9 CONCLUSION

In this work, we have presented a first, validated macroscale model for simulating the in-plane deformation of a bird feather realistically, efficiently and robustly. Thanks to our laboratory experiments, we found that the feather’s vane obeys a simple linear orthotropic constitutive law, albeit with an extreme ratio of stiffness between the two principal directions of the material. This led us to model the vane as a strongly anisotropic elastic membrane. Based on the ARCSim software, we have incorporated strong anisotropy in the membrane while addressing numerical locking issues and ill-conditioning through an adapted mesh and an inextensibility constraint. Our novel anisotropic membrane model was quantitatively validated against real experiments. In addition, we have shown that feathers simulated with our membrane model, together with anisotropic bending stiffness, visually match the behaviour of real feathers. In the future, we plan to enrich our feather model by carefully validating our anisotropic bending model, as well as incorporating fracture and non-uniform thickness of the membrane. We believe our work can be inspirational for all kinds of materials involving extreme anisotropy, such as fibrous media that can be found in biological tissues and reinforcement structures.

ACKNOWLEDGEMENTS

We would like to thank the Inria invited “external researcher program” that allowed to fund Rahul Narain’s visit at Inria Grenoble in fall 2023. We are also grateful to all the anonymous reviewers for their constructive comments.

REFERENCES

Nicholas Augello, David Tonnesen, and Arunachalam Somasundaram. 2019. Hummingbird: Dreamworks Feather System. In *ACM SIGGRAPH Talks*. Article 60, 2 pages.

Jessica Baron, Daljit Singh Dhillon, N Adam Smith, and Eric Patterson. 2022. Microstructure-based appearance rendering for feathers. *Computers & Graphics* 102 (2022), 452–459.

Jessica Baron and Eric Patterson. 2019. Procedurally generating biologically driven feathers. In *Computer Graphics International Conference*. Springer, 342–348.

Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete Elastic Rods. *ACM Trans. Graph.* 27, 3, Article 63 (Aug. 2008), 12 pages.

Florence Bertails. 2009. Linear Time Super-Helices. *Computer Graphics Forum (Proceedings of Eurographics)* (2009).

Florence Bertails, Basile Audoly, Marie-Paule Cani, Bernard Querleux, Frédéric Leroy, and Jean-Luc Lévêque. 2006. Super-Helices for Predicting the Dynamics of Natural Hair. In *ACM Transactions on Graphics*.

R Bonser and P Purslow. 1995. The Young’s modulus of feather keratin. *Journal of Experimental Biology* 198, 4 (1995), 1029–1033.

Stephen D. Bowline and Zoran Kačić-Alesić. 2011. Dynamic, Penetration-free Feathers in Rango. In *ACM SIGGRAPH Talks*. Article 35, 1 pages.

Hsiao-Yu Chen, Arnav Sastry, Wim M. van Rees, and Etienne Vouga. 2018. Physical Simulation of Environmentally Induced Thin Shell Deformation. *ACM Trans. Graph.* 37, 4, Article 146 (jul 2018), 13 pages. <https://doi.org/10.1145/3197517.3201395>

Qiang Chen, Stanislav Gorb, Alexander Kovalev, Zhiyong Li, and Nicola Pugno. 2016. An analytical hierarchical model explaining the robustness and flaw-tolerance of the interlocking barb-barbule structure of bird feathers. *EPL (Europhysics Letters)* 116, 2 (2016), 24001.

Yanyun Chen, Yingqing Xu, Baining Guo, and Heung-Yeung Shum. 2002. Modeling and Rendering of Realistic Feathers. *ACM Trans. Graph.* 21, 3 (July 2002), 630–636.

Cheng-Ming Chuong, Rajas Chodankar, Randall B Wideltz, and Ting-Xin Jiang. 2000. Evo-Devo of feathers and scales: building complex epithelial appendages. *Current Opinion in Genetics & Development* 10 (2000), 449–456.

David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Modeling and data-driven parameter estimation for woven fabrics. *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2017).

Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete Shells. (2003), 62–67.

Rasmus Haapaaja and Christoph Genzwürker. 2019. Mesh-driven Generation and Animation of Groomed Feathers. In *ACM SIGGRAPH Talks*. Article 61, 2 pages.

J. S. Hale, M. Brunetti, S. Bordas, and C. Maurini. 2018. Simple and extensible plate and shell finite element models through automatic code generation tools. *Computers & Structures* 209 (2018), 163–181.

Todd Alan Harvey, Kimberly S. Bostwick, and Steve Marschner. 2013. Directional reflectance and milli-scale feather morphology of the African Emerald Cuckoo, *Chrysococcyx cupreus*. *Journal of The Royal Society Interface* 10, 86 (2013), 20130391.

Daniel Heckenberg, Damien Gray, Bryan Smith, Jonathan Wills, and Chris Bone. 2011. Quill: Birds of a Feather Tool. In *ACM SIGGRAPH Talks*. Article 34, 1 pages.

Weizhen Huang, Sebastian Merzbach, Clara Callenberg, Doekele Stavenga, and Matthias Hullin. 2022. Rendering Iridescent Rock Dove Neck Feathers. In *Proceedings of SIGGRAPH*. 1–8.

Alexander Kovalev, Alexander E. Filippov, and Stanislav N. Gorb. 2014. Unzipping bird feathers. *Journal of The Royal Society Interface* 11, 92 (2014).

Eric Landreneau and Scott Schaefer. 2010. Scales and Scale-like Structures. *Computer Graphics Forum (Proceedings of Eurographics)* 29, 5 (2010), 1653–1660.

Le Liu, Xiaosheng Li, Yanyun Chen, Xuehui Liu, Jian J. Zhang, and Enhua Wu. 2015. An Efficient Feathering System with Collision Control. *Computer Graphics Forum (Proceedings of Pacific Graphics)* 34, 7 (2015), 279–288.

Le Liu, Xuehui Liu, Bin Sheng, Yanyun Chen, and Enhua Wu. 2017. Incremental collision-free feathering for animated surfaces. *The Visual Computer* 33, 6 (Jun 2017), 883–890.

G.D. Macleod. 1980. Mechanical Properties of Contour Feathers. *Journal of Experimental Biology* 87, 1 (1980), 65–72.

Eder Miguel, Derek Bradley, B. Thomaszewski, B. Bickel, Wojciech Matusik, Miguel A. Otaduy, and Steve Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Computer Graphics Forum* 31 (2012).

G.W. Milton. 2002. *The theory of composites*. Cambridge University Press.

Tobias Pfaff, Rahul Narain, Juan Miguel de Joya, and James F. O’Brien. 2014. Adaptive tearing and cracking of thin sheets. *ACM Transactions on Graphics (TOG)* 33 (2014), 1 – 9.

P.P. Purslow and J.F. V. Vincent. 1978. Mechanical Properties of Primary Feathers From the Pigeon. *Journal of Experimental Biology* 72, 1 (02 1978), 251–260. <https://doi.org/10.1242/jeb.72.1.251> arXiv:<https://journals.biologists.com/jeb/article-pdf/72/1/251/1343731/251.pdf>

Georg Sperl, Rosa M. Sánchez-Banderas, Manwen Li, Chris Wojtan, and Miguel A. Otaduy. 2022. Estimation of Yarn-Level Simulation Models for Production Fabrics. *ACM Transactions on Graphics (TOG)* 41, 4 (2022).

L. Streit and W. Heidrich. 2002. A Biologically-Parameterized Feather Model. *Computer Graphics Forum (Proceedings of Eurographics)* 21, 3 (2002), 565–573.

R.H.W. ten Thije and R. Akkerman. 2008. Solutions to intra-ply shear locking in finite element analyses of fibre reinforced materials. *Composites Part A: Applied Science and Manufacturing* 39, 7 (2008), 1167–1176. <https://doi.org/10.1016/j.compositesa.2008.03.014>

Pascal Volino, Nadia Magnenat-Thalmann, and François Faure. 2009. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.* 28 (2009), 105:1–105:16.

Andreas Wächter and Lorenz T. Biegler. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106 (2006), 25–57. <https://api.semanticscholar.org/CorpusID:14183894>

Huamin Wang, Ravi Ramamoorthi, and James F. O’Brien. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. *ACM Transactions on Graphics* 30, 4 (July 2011), 71:1–11. <http://graphics.berkeley.edu/papers/Wang-DDE-2011-08/> Proceedings of ACM SIGGRAPH 2011, Vancouver, BC Canada.

Andrew J. Weber and Galen Gornowicz. 2009. Collision-free Construction of Animated Feathers Using Implicit Constraint Surfaces. *ACM Trans. Graph.* 28, 2, Article 12 (May 2009), 8 pages.

Xiaoming Wei, Ye Zhao, Zhe Fan, Wei Li, Suzanne Yoakum-Stover, and Arie Kaufman. 2003. Blowing in the Wind. In *Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 75–85.

Jungdam Won, Jongho Park, Kwanyu Kim, and Jehee Lee. 2017. How to Train Your Dragon: Example-guided Control of Flapping Flight. *ACM Trans. Graph.* 36, 6, Article 198 (Nov. 2017), 13 pages.

Mingjiang Zhan; Richard P. Wool. 2011. Mechanical properties of chicken feather fibers. *Polymer Composites* 32 (2011). Issue 6.

Jia-chi Wu and Zoran Popović. 2003. Realistic Modeling of Bird Flight Animations. *ACM Trans. Graph.* 22, 3 (July 2003), 888–895.

Xiaobo Yu, Bruce Cartwright, Damian McGuckin, Lin Ye, and Yiu-Wing Mai. 2006. Intra-ply shear locking in finite element analyses of woven fabric forming processes. *Composites Part A: Applied Science and Manufacturing* 37, 5 (2006), 790–803. <https://doi.org/10.1016/j.compositesa.2005.04.024>

F. Zhang, L. Jiang, and S. Wang. 2018. Repairable cascaded slide-lock system endows bird feathers with tear-resistance and superdurability. *Proceedings of the National Academy of Sciences* 115 (2018), 10046 – 10051.




Figure 9: Tilted scenario with increasing sample aspect ratio (left to right) executed in the lab (top), in the simulation with microscale model (middle), and with the macroscale model (bottom). At fixed aspect ratio, each traction test depicts a shear band which increases in width as the sample is deformed. The shear band is remarkably similar between the real experiments and our macroscale model. This similarity demonstrates the capacity of our model to capture well the geometry of the deformed feather.



Figure 10: Bird flight simulation using our feather model, coupled with contact.




Figure 11: Comparison between slopes of the stress-strain curves of the macroscale model and the microscale model in traction scenarios with barbules oriented at different angles. The slopes are nondimensionalized through scaling by the Young modulus of the barbules in the microscale model E_b . Comparison is done using the measured barb density $\Phi_B = 0.16$, as well as the two extreme values of the error range ± 0.03 .




Figure 12: Comparison of the microscale model behaviour with the behaviour observed in experiments. Comparison is done between the stress-strain curves measured in our three traction experiments. Parameters of the microscale model were chosen to minimise the distance to every data point of the experiments in both the transverse scenario (left) and the longitudinal scenario (middle). The microscale model parameters resulting from our fitting procedure are listed in Table 1. Simulation of the tilted scenario with those parameters is in agreement with the experiments despite the large measured standard deviation on the barb density Φ_B (right).




Figure 13: Experimental stress-strain measurements for feather patches with two different orientations: transverse (left), longitudinal (right). Patches were taken from feathers of swans (---) and guinea fowls (-.-.-). Force-displacement curves are inset of the stress-strain curves, showing how the computation of the stress and strain collapses the curve of each bird species together, see Equation 2 and Equation 1 respectively. The experiments themselves are shown in the bottom row. Note the linear elastic constitutive law of the membrane, as well as the extreme stiffness ratio, of the order of 10^4 , between the barb direction (measured with the longitudinal test) and the barbule direction (measured with the transverse test). The measured dimensions of samples can be found in Table 1 in the supplemental document.