



HAL
open science

Numerical Modeling of Feathers

Jean Jouve

► **To cite this version:**

Jean Jouve. Numerical Modeling of Feathers. Modeling and Simulation. Université Grenoble Alpes, 2024. English. NNT: . tel-04907080

HAL Id: tel-04907080

<https://hal.science/tel-04907080v1>

Submitted on 22 Jan 2025

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.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

École doctorale : MSTII - Mathématiques, Sciences et technologies de l'information, Informatique

Spécialité : Mathématiques et Informatique

Unité de recherche : Laboratoire Jean Kuntzmann

Modélisation Numérique de Plumes

Numerical Modeling of Feathers

Présentée par :

Jean JOUVE

Direction de thèse :

Florence BERTAILS-DESCOUBES

DIRECTRICE DE RECHERCHE, CENTRE INRIA UNIVERSITE
GRENOBLE ALPES

Directrice de thèse

Victor ROMERO

STARTING FACULTY, INRIA

Co-encadrant de thèse

Rapporteurs :

MATHIEU DESBRUN

PROFESSEUR, ECOLE POLYTECHNIQUE

MIGUEL OTADUY

FULL PROFESSOR, UNIVERSIDAD REY JUAN CARLOS

Thèse soutenue publiquement le **21 novembre 2024**, devant le jury composé de :

MAUD MARCHAL,

PROFESSEURE DES UNIVERSITES, UNIVERSITE DE RENNES

Présidente

FLORENCE BERTAILS DESCOUBES,

DIRECTRICE DE RECHERCHE, CENTRE INRIA UNIVERSITE
GRENOBLE ALPES

Directrice de thèse

MATHIEU DESBRUN,

PROFESSEUR, ECOLE POLYTECHNIQUE

Rapporteur

MIGUEL OTADUY,

FULL PROFESSOR, UNIVERSIDAD REY JUAN CARLOS

Rapporteur

CATHERINE QUILLIET,

MAITRESSE DE CONFERENCES HDR, UNIVERSITE GRENOBLE ALPES

Examinatrice

BENOÎT ROMAN,

DIRECTEUR DE RECHERCHE, CNRS DELEGATION PARIS CENTRE

Examineur

ARIANE GAYOUT,

DOCTEURE EN SCIENCES, RIJKSUNIVERSITEIT GRONINGEN

Examinatrice

MARTIN PONCELET,

MAITRE DE CONFERENCES, ENS PARIS SACLAY

Examineur

Invités :

VICTOR ROMERO GRAMEGNA

STARTING FACULTY, CENTRE INRIA UNIVERSITE GRENOBLE ALPES



English 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 stiffness in the barb direction which is four order of magnitude greater than the stiffness in the barbule direction. 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, due to the extreme stiffness ratio between the barb and the barbule directions. To resolve this issue, we align the mesh along the barb directions. 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 such as the constrained motion of a full feather.

French abstract

Les plumes pénacées des oiseaux ont un comportement mécanique particulier qui émane de leur structure hiérarchique. Cette structure est constituée de tiges, appelées barbes, partant de d'un axe central, appelé rachis. Ces barbes s'accrochent entre elles à l'aide d'un système de crochets et goutières liées à un autre niveau de tiges attachés aux barbes, appelées barbules. Grâce à ce système d'accroche, les barbes forment une surface élastique, appelée étendard, qui a la particularité d'être fortement anisotrope. Le coefficient de raideur dans le sens des barbs est quatre ordres de grandeur plus grand que celui dans le sens des barbules. Cette aspect mécanique des plumes n'a jamais été pris en compte dans la modélisation des plumes dans l'industrie de l'informatique graphique. À la place, les plumes sont modélisées en tant que surface élastique isotrope. Dans le cadre de cette thèse, nous analysons le comportement élastique des plumes à travers des mesures précises sur de multiples tests de traction d'échantillons d'étendard de cygnes. De cette analyse, nous tirons une relation linéaire entre la contrainte et la déformation pour le comportement membranaire fortement anisotrope de la plume. Seulement, l'implémentation numérique de ce modèle à travers des éléments finis rencontre un problème de verrouillage. Nous résolvons ce problème en alignant les éléments avec la direction de plus grande raideur. Ce modèle est ensuite validé contre d'autres expériences de tractions, eux aussi sur des échantillons de cygnes. Dans le but de réduire le nombre d'expériences à faire, nous introduisons un modèle intermédiaire de l'étendard à une échelle plus basse et nécessitant moins d'expériences pour être validé. Notre modèle plus haut niveau peut ensuite être validé contre ce modèle, les expériences nécessaires étant produites à travers une simulation du modèle bas niveau. De plus, nous montrons que notre modèle peut être utilisé pour des applications en informatique graphique à travers l'implémentation préliminaire d'une plume complète que nous comparons au comportement d'une plume réelle ainsi qu'au modèle isotrope actuellement utilisé en informatique graphique.

Acknowledgement

Before going into what was scientifically done during my PhD thesis, I would like to thank everyone that helped me during the the time I have been in Grenoble.

I would like to thank my supervisors Florence Bertails-Descoubes and Victor Romero Gramagna who were making sure to be as available as possible to help me go forward during the thesis. They were quite understanding of my mental health problems, always taking the time to resolve disagreements without using their authority, but instead doing so through discussion. They prevented me from going on tangents which would have made me unable to present the work you will be reading here. I would also like to thank Theodore Kim and Rahul Narain with whom I worked during two and a half year of my PhD thesis and were a source of knowledge without which my work would have taken much longer to go through. I especially thank them for finding the time to meet remotely despite the great hours difference between India, the United States and France. Thibaut, the only other permanent researcher of our team, was also of a great help, especially in understanding the intricacies of the finite element methods and of elasticity theory.

All my fellow colleagues and friends, Émile, Alice, Mickaël, Octave and Sarah are also to be thanked for the time they spent listening to my technical issues and helping resolve them. But also for having made my time in Grenoble quite pleasant by accepting to play boardgames and going on hikes with me. I especially thank Alice with whom the one and half a year of being roommate was quite pleasant.

This brings me to thanking my current roommate and best friend Louka, thanks to which I spend most of my evenings having great discussions about politics and some more trivial things. My other close friends from Rennes Dylan, Gaël and Léon have also been a great help by welcoming in their home when I wanted a change of air and meet them again. Regarding this, I am grateful for the whole Logica team in which Dylan and Gaël work for letting me use their office and being very kind, allowing me to comfortably work remotely over there.

My parents Claude and Laurence as well as my sister Marion and her husband Thibault were also a great help when I needed a change of air. Thank you for welcoming me for multiple weeks at a time and also asking for me to come again, this helped a lot and gave a good boost to my self-confidence.

I also have a thought for the whole left activist community of Grenoble, for heightening my practical and theoretical political knowledge through multiple events and discussions. Among this group, I want to especially thank the group “Faut-il continuer la recherche ?” who gave me some opportunities to think

about the position of research in society.

Lastly, I would like to thank my psychologist Paul Sabino who followed me during my three years and helped me a great deal with my difficult mental health problem.

Contents

Contents	9
1 Introduction	11
1.1 Introduction	11
1.2 Summary of our approach	14
1.3 Notations	15
2 Modelling Method	19
2.1 Scale of the phenomena vs. scale of modelling	19
2.2 Material model identification	22
2.3 Membrane modelling	26
3 Lab experiments	31
3.1 Scenarios	31
3.2 Setup	35
3.3 Results and observations	38
4 Macroscopic model for the membrane	47
4.1 Analytical Model	47
4.2 Numerical model	49
4.3 Numerical locking	52
5 Calibration and validation	57
5.1 Avoiding manual work	58
5.2 Mesoscale Model	59
5.3 Membrane model	64
5.4 Discussion	67
6 Model of a full feather	71
6.1 Rachis	71
6.2 Non-uniformity of barb orientation	72
6.3 Flexion modelling	73
6.4 Parameter calibration	75
6.5 Energy penalty for bending	78
6.6 Qualitative Validation	79
7 Conclusion	83
A Elasticity Theory	87
A.1 Poisson's ratio	87

A.2	Elasticity tensor change of basis	88
A.3	Size effect	89
B	ARCSim	93
B.1	ARCSim implementation	93
C	Experimental scenarios	97
C.1	Intermediate model and transverse scenario	97
C.2	Equipment	99
D	French overall summary	101
D.1	Introduction	101
D.2	Méthode expérimentale	102
D.3	Modèle	103
D.4	Verrouillage	106
D.5	Validation	106
D.6	Conclusion	108
E	French chapter summary	111
E.1	Introduction	111
E.2	Méthode de modélisation	111
E.3	Expériences	112
E.4	Modélisation	112
E.5	Calibration and validation	112
E.6	Plume complète	112
	Bibliography	115

Chapter 1

Introduction

1.1 Introduction

Feathers are central to the life of birds. They play multiple roles: among others, they help their thermoregulation, flight, and camouflage. There are two types of feathers which play different roles. *Peannaceous* feathers, which are crucial for rain protection and flight, are found in the outer plumage of the bird and also on the bird tail and wing. *Down* feathers can be found underneath the peannaceous feathers. They play an important role in the thermoregulation of the bird. Due to their different roles, those feathers possess different structures. From the naked eye, the difference is striking: a peannaceous feather forms a curved surface (see Figure 1.1a) while a down feather looks like a bundle of hair (see Figure 1.1b).

The distinct look between peannaceous and down feathers comes from a disparity in their underlying structure. Both down and peannaceous feathers consist of a hierarchical structure starting with a thick shaft called the *rachis*. On either side of this rachis emanate hundreds of thinner rods called *barbs*, each of them having hundreds of *barbules* branching out. The dissimilarity between the two types of feathers comes from this last level of the hierarchical



(a) Photo of a peannaceous feathers from Ducasse (2004).



(b) Photo of a down feather from Hagens (2021).




Figure 1.2 – 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 the barbules on the other side do not, allowing barbs to reversibly attach and detach (right). The attached barbs form a surface called the *vane*.

structure. Some barbules of peannaceous feathers possess hooks, allowing the barbs to attach together and forming this way two surfaces on either side of the rachis called *vanes*. In contrast, in down feathers, no hooks can be found, preventing the formation of vanes that gives the characteristic look of peannaceous feathers.

This velcro-like structure of peannaceous feather gives them some interesting mechanical properties. First the surface formed by those feathers on the wing allows them to produce lift for the purpose of the flight. Also, when the vane is under great stress, the barbs dissipate the stored energy by unhooking the connections made by the barbules instead of breaking (see F. Zhang et al. (2018b) for a video describing this process). A bird can, at a later time, reattach the unhooked barbules by passing their beaks through the feather or by shaking their wings during the upkeep of their plumage (Zhao et al. 2020a). This makes peannaceous feathers quite the robust appendage. Another interesting mechanical aspect of the peannaceous feather is found in their *membrane behaviour*, meaning their mechanical behaviour that doesn't involve bending. The membrane behaviour of the feather vane is mostly driven by its hierarchical structure: axial tension along the barb is analogous to the very stiff extension of a beam while tension in the transverse direction is more akin to beam flexion which is known to be substantially less stiff than the former. Hence, the stiffness of the membrane elastic behaviour is highly dependent on the orientation, i.e. it is *anisotropic*. Moreover, this anisotropy is strong as there are four orders of difference between the highest and lowest stiffness. For comparison, other highly anisotropic materials such as carbon reinforced epoxy have at most two order of magnitude of difference between the maximum and minimum stiffness (Mirdehghan 2021; Toolbox 2003). This might or might not play a role in the bird day to day life — a more thorough biomechanical literature review would be necessary to have a clear idea of the state of the art on the question. Nonetheless, the study of this behaviour would be valuable for understanding how the feather deforms.

In addition to their peculiar mechanical behaviour, feathers are central to the typical appearance of birds (see Figure 1.3). Moreover, they are also used in apparel and decoration throughout human history (see Figure 1.4).



Figure 1.3 – Feathers are so characteristic of bird’s appearance that most people are weirded out by the absence of feather on a part of the bird body. Image from (Mazur 2008).



Figure 1.4 – Feathers are central to the apparel and traditions of some cultures. Here are presented two apparels from the Wayana people located in the south-eastern part of Guiana (Image from (left) (Museum 1907) (right) (Museum 1972)).

Hence, their appearance in multiple movies through the use of special effect and animation. To avoid the manual animation of hundreds of feathers on a bird body, animators sometimes use simulator to give a first draft of the movement which they will then iterate over to reach the wanted artistical vision. To the best of my knowledge, feature film studios have modelled feathers as simple strips of isotropic cloth (Bowline and Kačić-Alesić 2011; Weber and Gornowicz 2009; Augello et al. 2019; Heckenberg et al. 2011; Haapaoja and Genzwürker 2019). Such models are not able to reproduce scenarios where the strong anisotropy of the vane and the repairing behaviour are key (see Figure 1.5).

Though most of the mechanical behaviours of feathers are clearly understood from a qualitative point of view, little endeavour has been put toward their modelling. Efforts have been focused on measuring the Young Modulus of β -keratin — the material of the feather — (Bonser and Purslow 1995; Wool 2011) and of the rachis (P.P. Purslow and Vincent 1978; Macleod 1980). Higher level models that describe the elastic, plastic and, fracture interactions between barbs due to barbules interlocking have also been devised (Kovalev et al. 2014; F. Zhang et al. 2018a). Among the available work on feathers, those last models are what come the closest to providing useful tools for the modelling of the




Figure 1.5 – Renders of the first (left) and final frame (right) of the deformation of a feather. The bottom images show a simulation of a full feather using an isotropic model as is usually done in the film industry. The top images display the same deformation demonstrated on a real feather. The visible unrealistic deformation of the vane displays the limitation of an isotropic model for the modelling of a pennaceous feather. See Chapter 6 for details on this simulation.

dynamic appearance of feathers. Unfortunately, their suggested models would require more than a million degrees of freedom for the modelling of one full feather, making them unpractical.

Hence, to further the state of the art of feather modelling both in biomechanics and computer graphics, we aimed at bringing closer the possibility of a predictive model of the dynamics of a full feather which takes into account both fracture, reparability and strong anisotropy. Moreover, we also consider the computational efficiency of such a model for the need of animation where the animator has to do multiple iterations with the simulator before reaching the final motion. While this target was not reached during the course of the thesis, we took some steps toward it through the modelling of the membrane behaviour of the vane, focusing on its elastic behaviour and leaving any plastic or fracture aspect for future works. Furthermore, we did some preliminary work on the qualitative modelling of a full feather elastic behaviour, taking into account membrane and flexural behaviour for both the rachis and the vane.

Succinctly, the main contributions of my Ph.D. thesis are:

- An experimental characterisation of the membrane mechanical behaviour.
- A simple elastic membrane model for the feather vane.
- The validation of the vane membrane model against experiments through a two-step validation.
- A qualitative model for a full feather (rachis and vane) along with some qualitative validation.

1.2 Summary of our approach

In Chapter 3, we describe our experimental setup for analysing the behaviour of the feather vane. The experimental scenarios chosen for this analysis are

simple traction experiments performed on samples of feather vane in different orientations. The analysis of the resulting stress-strain curve shows that there is at least four order of magnitude difference between the stiffness in two orientations, which demonstrates a strong anisotropy within the feather vane. Moreover, we also observe a linear relationship between stress and strain. This last observation is used as a base for the derivation of our vane membrane model.

In Chapter 4, we derive a simple model for the membrane behaviour of the feather vane based on the observations made in the previous chapter. The resulting model is actually very close to the one used in computer graphics for the modelling of clothes, the difference being the substantially stronger anisotropy present in the feather vane compared to cloth. This model is implemented numerically using a linear triangular finite element method. However, this implementation is subject to locking emanating from the strong anisotropy. This numerical locking is resolved by aligning the triangular elements with the barbs.

Following this, in Chapter 5, we proceed to the validation of this model. Unfortunately, the scenarios presented in Chapter 3 used for the study of the feather vane are not sufficient for a proper validation of our model. Instead of designing and performing new experiments, we choose to introduce a lower scale intermediate model which can be validated on the available experiments thanks to some geometry considerations. Then, using this validated intermediate model, we proceed to validate our previously introduced vane model thanks to experiments generated numerically using the intermediate model.

Finally, since we are considering applications of our work in computer graphics, we demonstrate the potential applicability of our membrane model for the modelling of a full feather in Chapter 6. Here, for simplicity, we model the feather as one single heterogeneous shell: a central band of material represents the rachis and the material on either side of the rachis model the vane. The vane is modelled by our model for the membrane behaviour and the flexural behaviour is modelled through an ad hoc anisotropic bending model. To evaluate the usefulness of this model, it is compared against the video of a real feather and against a model similar to the isotropic model usually used in the industry for the modelling of feathers.

1.3 Notations

General rules

- Vectors and other higher order tensor are written in italic bold: $\mathbf{f}, \boldsymbol{\sigma}, \mathbf{E}$;
- Scalar are written in italic: Φ_B, d, ρ ;
- The sets of real number is written as \mathbb{R} ;
- Other sets are written with uppercase italic Greek letters: Γ_t, Ω .

Symbols

Table 1.1 – Notations common to differential geometry.

Notation	Description
$A \times \Gamma$	Cartesian product of the sets A and Γ .
$A \subset \Gamma$	The set A is a subset of Γ .
$\mathbf{A}^{\frac{1}{2}}$	Matrix square root of a positive semi-definite matrix \mathbf{A} .
\mathbf{A}^T	Transpose of matrix \mathbf{A} .
\mathbf{AB}	Matrix product of matrices \mathbf{A} and \mathbf{B} .
\mathbf{A}^2	Square of the matrix \mathbf{A} .
$\mathbf{E} : \boldsymbol{\varepsilon}$	Double contraction between tensors \mathbf{E} and $\boldsymbol{\varepsilon}$.
$\nu_v \sim 1$	The scalar ν_v is of the same order of magnitude as 1.
$\nu_w \ll 1$	The scalar ν_w is much smaller than 1.
$\nabla \cdot \mathbf{F}\boldsymbol{\sigma}$	Divergence of $\mathbf{F}\boldsymbol{\sigma}$.
$\frac{\partial \boldsymbol{\varphi}}{\partial \mathbf{r}}$	Jacobian of the function $\boldsymbol{\varphi}$ with respect to its argument that was previously designated as \mathbf{r} .
$\mathbb{R}, \mathbb{R}^2, \mathbb{R}^3$	Real numbers; $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$; $\mathbb{R}^3 = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$.
$\mathbf{0}$	Vector zero.
\mathbf{R}	Rotation matrix.

Table 1.2 – Notations used throughout the document.

Notation	Description
Ω	Reference state of an elastic material.
$\boldsymbol{\varphi}$	Deformation of an elastic material.
\mathbf{u}	Displacement.
$\ddot{\mathbf{u}}$	Acceleration.
\mathbf{r}	Material point of an elastic material, always lies in Ω .
\mathbf{F}	Deformation gradient, i.e. $\frac{\partial \boldsymbol{\varphi}}{\partial \mathbf{r}}$.
ψ	Surface energy density.
ψ_m	Membrane surface energy density.
ψ_f	Flexural surface energy density.
$\boldsymbol{\eta}$	Second order tensors that measures bending.
y-axis, x-axis	Axis aligned and orthogonal with the direction of elongation within the considered scenarios.
v-axis, w-axis	Axis aligned and orthogonal with the barbs.
$\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{v}}, \hat{\mathbf{w}}$	Unit vector in the direction of increasing value along the x-axis, y-axis, v-axis, and w-axis.
$\boldsymbol{\varepsilon}$	Green-Lagrange strain tensor.
ε_{xy}	Components of the second order tensor $\boldsymbol{\varepsilon}$ along x-axis and y-axis, i.e. $\varepsilon_{xy} = \hat{\mathbf{x}}^T \boldsymbol{\varepsilon} \hat{\mathbf{y}}$.
ε_{\max}	Maximal strain considered in the experiments.
$\tilde{\varepsilon}_{yy}$	Strain measured in the experiments under a uniform deformation assumption.
λ	Stretching ratio, ratio between current length and length at rest.
λ_{\max}	Maximal stretch ratio considered in the experiments.
$\tilde{\lambda}$	Measured stretch ratio.

Table 1.2 – Notations used throughout the document (continued).

Notation	Description
$\boldsymbol{\sigma}$	Second Piola-Kirchhoff stress tensor.
$\tilde{\sigma}_{yy}$	Stress measured in the experiments under a uniform deformation assumption.
$\tilde{\sigma}_{yy}^s, \tilde{\sigma}_{yy}^r$	Measured stress predicted by the intermediate model with non-rigid and rigid barb bands.
\mathbf{f}	A force.
\mathbf{p}	Linear force density.
dl	A line element.
f	Force measured in the experiments.
d	Displacement measured in the experiments.
W, H, b	Width, height and thickness of a sample.
a	Aspect ratio of the samples, i.e. $\frac{H}{W}$.
$\Gamma_t, \Gamma_b, \Gamma_s$	Set of material points forming the top, bottom and sides edges of a sample.
ν_v	Poisson's ratio arising from an elongation along the v-axis
\mathbf{E}	Elasticity tensor.
E_{ww}, E_{vv}, E_{vw}	Transverse, longitudinal and shearing modulus.
α	Stiffness ratio: $\frac{E_{vv}}{E_{ww}}$
ρ	Surface mass density.
E_b, E_B	In the intermediate model, Young's moduli of the bands representing the barbule and the barbs.
T	In the intermediate model, distance between the centres of two band representing the barbs.
w_B	In the intermediate model, width of bands representing the barbs.
Φ_B	In the intermediate model, surface fraction occupied by the bands representing the barbs.
$\Delta_{ }, \Delta_{\perp}$	Slopes of stress-strain curves in the longitudinal and transverse scenario.
θ	Angle between the barbs and the x-axis.
w_b, w_e	Width of the rachis at its base and end.
e	Target edge size in the meshing for the full feather model.
A	Area of a triangle.
D	Bending modulus.
$D_{ }, D_{\perp}$	Bending modulus for flexion transverse to and along barbs.
η	Curvature.
η_0	Curvature at rest.
β	Angle between an edge and the local barb orientation.

Chapter 2

Modelling Method

2.1 Scale of the phenomena vs. scale of modelling

To make the modelling of the feather and the validation of the model practical, we have to choose the scale size of the phenomena we are interested in. In our case we chose to focus on the scale just above the barbs. This means that we are not interested in the movement of the individual barbs nor of anything smaller than the barbs but of the overall movement of the vane. There are two reasons behind this choice of scale. First, the overall movement of the vane is the predominant visual aspect of the feather dynamics to the naked eye, i.e. it is the *macroscopic* scale. Hence, it is of most importance in the context of computer graphics applications¹. Second, considering the vane scale makes our aim of a quantitative validation much easier. In fact, validating that our model properly predicts the movement of the feather at this scale rather than at a lower scale would require less precise and pricey equipments.

Choosing the scale size of the phenomena we are interested in does not necessarily fix the scale size of the model: the only thing that it defines is the scale of the validation. The significance of this distinction can be seen for example in the recent work on knitted fabrics in computer graphics. Here, the phenomenon of interest is the movement of the knitted surface. The movement of the scale below, i.e. the scale of individual fibres, is often quite subtle especially in some garments where they are almost invisible to the naked eye (see Figure 2.1). Considering this, knitted fabrics were initially modelled using the same methods used for woven cloth. However, this approximation is insufficient to properly predict the behaviour of knitted cloth (see Figure 2.2).

To get closer to knitted garments, a popular and accurate solution is to model the cloth at a lower scale than the phenomena of interest, the strand level, trading computation time for accuracy² (Kaldor et al. 2008; Cirio et al. 2017). Hence, in the context of the modelling of feathers, though the phenom-

¹Movement of smaller scales can also have their importance depending on the proximity between the camera and the feather, but we believe the scale of the vane to still remain the one most commonly encountered.

²More recent work have managed produced accurate homogeneous model for knitted fabrics through the homogenisation of the strand level simulations (Sperl, Narain, et al. 2020). However, this does not invalidate the necessity to consider different modelling scales, since a strand level model was still needed for the homogenisation process.



Figure 2.1 – While we are mostly used to think of fabric made of big yarn when knitted garment are mentioned (left and right). For some knitted cloth, the knits, and yarn are practically invisible to the naked eye (middle). (Image from Kozachenko (2020))



Figure 2.2 – Extension test on fabric knitted with the stockinette pattern. Those tests are performed with different tool: with a real fabric (left); with the cloth model that was used before the introduction of yarn-level models; and with a yarn-level model. Here, the cloth model is not able to reproduce the curling that appears on the left of the real fabric, showing its limitation for the simulation of knitted garment. Images are taken from Kaldor et al. (2008).

ena of interest is at the scale of the feather vane, it is relevant to consider all lower scales and compare their pros and cons.

Due to the hierarchical structure of the feather, there are three relevant modelling scales in the feather: the barbule scale, the barb scale, and the vane scale. For each of these possible modelling scales, we have to consider how easy it would be build an accurate model, and how easy it would be to provide initial data for a numerical simulation.

Modelling a feather at the barbule scale is attractive. Indeed, just by using




Figure 2.3 – Renderings of barbules taken from F. Zhang et al. (2018a). Barbules with hooks are leaf like and ends with thin rod from which the hooks grow out of (left). Barbules that form a gutter onto which hooks grab also have a leaf like structure, however, instead of ending with hooks, they end with small spines (right). This complex shape of the barbules would make it difficult to generate the appropriate geometry for a simulation of the feather vane at the barbules level.

existing models for elasticity and friction, you would be able to capture every phenomenon of interest: elasticity, plasticity, fracture, and reparability. However, prescribing the geometry of a feather would not be trivial as barbules have a rather complex shape (see Figure 2.3)³. Moreover, there is near a million barbules per feather, and each would require multiple degrees of freedom, making the computation cost of the modelling at this scale deterring.

The barb scale is more reasonable in terms of geometry prescription — barbs being only thin rods — and in terms of computation cost — only around 800 barbs per feathers. Although a model at this scale could use existing efficient one dimensional thin rods model for the rachis and the barbs (Bertails et al. 2006; Bergou et al. 2008), a new model would still need to be devised for the interaction between barbs. While Kovalev et al. (2014) have built a model for this interaction for the 2D case, the barbules are represented individually: this becomes prohibitive when simulating a full feather.

Finally, for the vane scale almost all the necessary modelling tools already exist. Indeed, elasticity, plasticity, and fracture have been thoroughly studied in the context of homogeneous surfaces. To have a complete modelling of the phenomena at the vane scale, only the healing model would be missing. Moreover, choosing the vane scale should allow us to use fewer degrees of freedom than would be necessary compared to a model at the barb scale⁴. Modelling at the vane scale is not without downsides compared to the other scales. More model parameters are required to take into account the effect

³Obtaining the geometry of a known feather could easily be done through tomography or similar techniques. But when one does not want the geometry to match one specific real feather, as would often be the case in computer graphics, the task would become substantially more complicated.

⁴A strip of surface elements should be enough, to model multiple barbs. Hence, the number of necessary degrees of freedom would probably be divided by the number of barbs per strips.

Scale	Barbules	Barbs	Vane
Geometry Complexity	+++	+	+
Computational Cost	+++	++	+
Experimental setup	+++	++	+
Model design difficulty	+	+++	++

Table 2.1 – Purely subjective and a priori evaluation, those might differ widely in practice. This table aims at providing a summary of the different pros and cons of the modelling of the feather vane at different scales discussed in this section.

of the underlying complex geometry that is abstracted away. And it is less obvious how to find the right model that is able to take into account those effects. In contrast, in the barbules scale only a simple elasticity and friction model would be required.

In addition to comparing the impact of the modelling scale on computation cost and modelling difficulty, we also consider the difficulty of the analysis of the scale through experiments. While the validation is done at the vane scale as it is the scale of the phenomenon of interest, we still want our model to emanate from experimental observations. And, the smaller the scale we would need to analyse, the more precise the equipments need to be and the more tedious the setting up of experiments becomes.

In light of these differences between the modelling scales — summarized in Table 2.1 — and due to the performance requirement of computer graphics applications as well as our lack of very precise measurement devices, we model the feather behaviour at the vane scale, i.e. the macroscopic scale. We will show through a quantitative validation in Chapter 5 that our choice of scale is sufficient for our purpose.

2.2 Material model identification

Finding the right elastic model for a new material is far from being a new problem. The general problem of material model identification has been studied for as long as people have wanted to predict and measure the world around us, and for elasticity this most likely started with the work of Hook (1678), and continues to be studied, even now (Ghiba et al. 2017). Hence, due to the long-standing nature of this problem, there are multiple possible methods available to us to find the right vane model for our feather. We describe in this section some of these methods and specify the one we use among them. A more involved discussion on material model identification can be found in the introduction of Flaschel (2023).

Here, we make a clear distinction between *material model identification* and *parameter identification*. The former consists in finding the formal description the model with any dimensional quantity left to be set. The latter is the process of fixing those dimensioned parameters that were left unset. For example, the material model identification for a one dimensional spring stretched by a factor

λ might result in the following quadratic internal energy model,

$$k(\lambda - 1)^2, \quad (2.1)$$

where k is a parameter that should be set later on through the process of parameter identification. As stated earlier, this section only deals with material model identification, parameter identification being discussed later in Chapter 5.

Conventional method

The oldest method for material model identification is more of an art than a method, but is still widely used nowadays. It consists in guessing the material model from experimental observations (see Figure 2.6a). This process can sometimes be tedious as the initial guess might be demonstrated to be wrong through parameter identification if one is unable to find parameter values that allow the model to predict all experimental observations. In this case, another guess must be made, and parameter identification is needed once again, increasing the amount of work necessary with each guess. Most often, this method of finding a material model is not all guessing, thankfully, some geometric arguments such as symmetry within the material can sometime help find the proper model.

Homogenisation and dimensionality reduction

Sometimes, a material model, that is more complex than necessary, might already be available for the system considered. This unnecessary complexity can usually be identified by the model predicting some deformation that are negligible or irrelevant to the application considered. For example, modelling a thin rod through a full three-dimensional model might be superfluous as the cross-section does not deform much in most applications (see Figure 2.4). Another example could be modelling every small cavity in an elastic porous material, this might be unuseful when one is only interested in the macroscopic deformation of the material (see Figure 2.5). In those cases, a simpler and more efficient material model can be derived from the original model through one of two similar process: *dimensionality reduction*, which reduces the necessary number of dimension for modelling a system, this is what one would do to find a simpler model for the thin rod (see Audoly and Pomeau (2010) for a detail of the process); and *homogenisation* which, as its name suggests, homogenise away the small heterogeneity present in the original model to produce a new model which behaves the same statistically (Miled et al. 2011)

These two methods are quite alike in how they are applied: First, one has to identify small parameters in the original material model. Those small parameters usually emanate from the difference between the scale of what is modelled and the scale of the phenomenon of interest. In the thin rod example, this parameter would usually be the ratio between the size of the cross-section and the length of the rod. For the porous material, that would most likely be the average size of a cavity over the size of the system. Then, the new material model can be worked out by analysing how the original model behaves as this small parameter goes to zero. Both of those methods' protocol is schematised in Figure 2.6b.




Figure 2.4 – Modelling the deformation of the full volume of a thin rod through three-dimensional elasticity theory (left) is usually unnecessary, since the cross-section stays relatively undeformed. Instead, a simpler one-dimensional model requiring less degrees of freedom can be obtained through dimensionality reduction (right) (Bertails et al. 2006; Bergou et al. 2008).




Figure 2.5 – Modelling every cavity in a porous medium (left) might be superfluous when one is only interested in the overall deformation of the considered system. Thankfully, an equivalent homogeneous model (right) can be obtained by averaging the behaviour emanating from the presence of the cavities resulting in a model which required fewer degrees of freedom to predict the deformation of the system.

Since this process is fully analytical, it requires certain mathematical skills and is usually far from trivial. However, some work has been done on its automation through symbolic computation (Audoly and Lestringant 2023).

Data-driven methods

The following discussion on data-driven methods is highly inspired by the introduction of Flaschel (2023).

A last set of method, which we could call *data-driven*, aims at automating the two previous problem, i.e. finding a model from observations or finding a more efficient model from a complex model. To do so, these methods will represent the material model through a functional with a tremendous amount of parameters, such as a neural network (Ghaboussi et al. 1991), splines (Sussman and Bathe 2009), or others. Thanks to this, those models can pretty much represent any function. Meaning that the material model can be made to match any set of experimental observations through an automated parameter identification. Since all the work is done during parameter identification, those methods do not quite fit in what we call material model identification. However, due to their generality, their comparison with material identification method is quite relevant. See Figure 2.6c for a schematised description of the associated protocol.

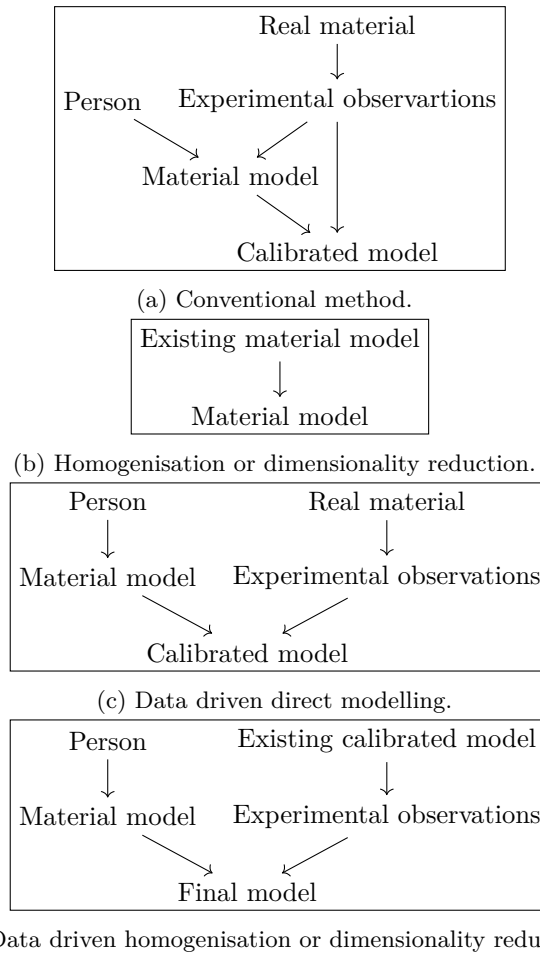


Figure 2.6 – Schematics of the different material and parameter identification methods described in this chapter.

While these methods do reduce the effort necessary for material model identification, this effort is transferred to the production of the substantial amount of experimental data necessary for setting all the parameters. This effort can be discharged to a computer if the problem at hand is homogenisation (Z. Zhang et al. 2023) or dimensionality reduction (Sperl, Narain, et al. 2020) as the data production can be automated thanks to the already available model (see Figure 2.6d). Also, in the case of solid mechanics, a considerable amount of data can be generated from a single experiment by measuring the full displacement of the considered specimen through Digital Image Correlation.

Unfortunately, models resulting from data-driven methods are difficult to interpret and analyse through analytical means compared to conventional models. This prevents the investigation of material behaviours in general cases, the proof of existence of solutions, the study of the impact each parameter has on the overall behaviour of the material, and so on. Flaschel et al. (2021) overcome this issue by choosing a model which is the sum of usual conventional model

for elasticity, in addition they enforce a sparse number of non-zero parameters which results in a somewhat simple model which can be properly analysed using usual analytical tools.

Discussion

In our work, material model identification was done through the conventional method. We could have considered the analytical homogenisation of a model of the underlying structure of the feather vane. However, it was not clear to us how this problem should be tackled. We did use a homogenisation procedure on a barb scale model for the simpler process of identifying the parameters of the vane scale model we chose. In this case, the material is fixed, and the parameter of the vane scale model are fitted on the lower scale model without much use of analytical tool.

Data-driven methods were not considered, since we wanted our model to be as simple and interpretable as possible. We could have used the automated model discovery from Flaschel et al. (2021) to ensure interpretability while having a better guarantee that the resulting model would match our experimental data. However, we were not aware of their work at the time, hence, we leave the application of their method to the feather vane for future work.

2.3 Membrane modelling

As discussed in the previous section, we model the feather as an elastic system. More specifically, by *elastic*, we mean that the behaviour of the material doesn't depend on its history, i.e. its current state is completely described by its current geometrical deformation and is independent on how the material was deformed to get there. With this in mind, an elasticity model will simply relate the current deformation of the solid to the resulting internal forces induced by this deformation.

To predict the forces within the considered system, a general three-dimensional elasticity model needs a description of the deformation of any point within its volume. This is often inefficient when working with systems which are thin in one dimension such as the vane of feathers. Indeed, this thinness allows the whole deformation to be described only by the one of the mean surface, which means that fewer degrees of freedom are necessary to describe the state of the system (see Figure 2.7). The theory that allow this simplification is called shell theory (Ciarlet 2021b) and its simplification to surfaces which have a planar rest state is called plate theory (Ciarlet 2021a). For our modelling of the feather vane, we naturally employ these theories.

Models lying within shell theory usually have two components, a membrane component which describe the in-plane behaviour of the shell, and a flexural component which describe how the shell bends. In our work, we decided to focus on the modelling and validation of a model for only the membrane behaviour. Hence, in this section, we will show that the modelling of the membrane behaviour can be tackled without considering the modelling of the flexural behaviour. Instead of discussing this using shell theory, we use the simpler plate theory. The different aspects discussed should be extendable to the more general case of shell theory.




Figure 2.7 – Three-dimensional elasticity theory would consider the deformation of the whole volume when modelling a thin structure (left) while plate and shell theory only consider the deformation of the mean surface (right).

Separation of modelling between flexural and membrane components

To discuss the reasoning behind our focus on membrane behaviour, we first make a high level description of plate theory.

Describing the deformation of a plate means giving the position of the different material points that make it up. Hence, we consider the set of material points $\Omega \subset \mathbb{R}^2$ which also gives an arbitrary rest state of the plate. This is also called the *reference state*. With this, the deformation is given by a function $\varphi : \Omega \rightarrow \mathbb{R}^3$ that maps any given material point $\mathbf{r} \in \Omega$ to its current position in space $\varphi(\mathbf{r}) \in \mathbb{R}^3$, i.e., it describes how the system is deformed from its rest state. We will often use the displacement $\mathbf{u}(\mathbf{r}) = \varphi(\mathbf{r}) - \mathbf{r}$ in some expressions instead of the deformation φ as it can result in simpler equations.

Since the deformation of a localised part of an elastic material does not affect remote regions, the response of the material at a point $\mathbf{r} \in \Omega$ only depends on the derivative of the deformation at this point $\frac{\partial \varphi}{\partial \mathbf{r}}, \frac{\partial^2 \varphi}{\partial \mathbf{r}^2}, \dots$ at the point of interest. The first derivative is often called *deformation gradient* and is denoted by \mathbf{F} .

Based on this description of the deformation, most elastic model prefer to use a potential energy density $\psi(\mathbf{F}, \frac{\partial \mathbf{F}}{\partial \mathbf{r}}, \dots)$ from which the forces can be derived rather than to describe these forces directly. This ensures energy conservation and gives access to a myriad of useful analytical tools for the study of elasticity. Such models are called *hyperelastic*.

In most elasticity models for plate and shell theory, the energy density can be separated in the sum of two energy densities, one for the membrane behaviour ψ_m and another for the flexural behaviour ψ_f ⁵. Formally, this means that the energy takes the following form

$$\psi\left(\mathbf{F}, \frac{\partial \mathbf{F}}{\partial \mathbf{r}}, \dots\right) = \psi_m(\mathbf{F}) + \psi_f\left(\boldsymbol{\eta}\left(\mathbf{F}, \frac{\partial \mathbf{F}}{\partial \mathbf{r}}, \dots\right)\right),$$

⁵This separation is present for all model implemented in commonly used software. For example, all model implemented in FENICS-SHELLS (Jack S Hale et al. 2018), Kirchhoff-Love plate (Ciarlet 2021a); von-Kármán plates (Audoly and Pomeau 2010); Reissner-Mindlin plate (Ciarlet 2021a); and Naghdi shell (Antman 2005), can be separated into a flexural and membrane model. Some unconventional systems cannot be modelled with such a separation, this is the case of shell that have a non-uniform distribution of material across the thickness (Caillerie and Sanchez-Palencia 1995)

where the membrane energy ψ_m depends only on the deformation gradient⁶ \mathbf{F} , and $\boldsymbol{\eta}$ is a second order tensor, which measures bending⁷. Considering this specific form of the energy, it seems natural to find one model for the membrane behaviour and another for the flexural behaviour, rather than trying to characterize the whole plate model at once.

The order in which those models are more easily defined can be analysed by noticing that any bending measure $\boldsymbol{\eta}$ is zero if and only if the deformation of the plate is in plane. And, naturally, the flexural energy ψ_f is zero if and only if the bending measure $\boldsymbol{\eta}$ is zero. Hence, as long as the deformation is in plane, the energy density of the plate coincides with its membrane energy, i.e.

$$\psi\left(\mathbf{F}, \frac{\partial \mathbf{F}}{\partial \mathbf{r}}, \dots\right) = \psi_m(\mathbf{F}).$$

Moreover, since any deformation gradient \mathbf{F} can be obtained through an in-plane deformation⁸, the membrane energy ψ_m can be completely characterized through in-plane experiments. In light of this and of the difficulty presented by the validation of the membrane model, we chose to focus on the membrane behaviour and leave the meticulous characterisation of the flexural behaviour for future work.

Property of the membrane model

We now discuss some properties that any membrane energy ψ_m should satisfy. A more detailed introduction of those properties can be found in Audoly and Pomeau (2010) or Ciarlet (1988) for a particularly rigorous mathematical derivation of them.

Green-Lagrange strain

The membrane energy ψ_m is usually not made to depend directly on the deformation gradient \mathbf{F} . This is because an elastic energy needs to be invariant by any rigid motion of the considered system, meaning that for any rotation \mathbf{R} , we have

$$\psi_m(\mathbf{R}\mathbf{F}) = \psi_m(\mathbf{F}). \quad (2.2)$$

Instead, the membrane energy ψ_m only depends on the symmetric part $\mathbf{S} = (\mathbf{F}^T \mathbf{F})^{\frac{1}{2}}$ of the polar decomposition of the deformation gradient $\mathbf{F} = \mathbf{R}\mathbf{S}$. Or more specifically, on the so called *Green-Lagrange strain* $\boldsymbol{\varepsilon}$,

$$\boldsymbol{\varepsilon} = \frac{1}{2}(\mathbf{F}^T \mathbf{F} - \mathbf{I}) = \frac{1}{2}(\mathbf{S}^2 - \mathbf{I}), \quad (2.3)$$

where \mathbf{I} is the second order identity tensor. This measure has the nice property of being null when the considered system is not deformed.

⁶This results from the dimensionality reduction of three-dimensional elasticity to plate theory (see Ciarlet (2021a))

⁷Note that this bending measure is not necessarily the same from one plate model to another.

⁸This can be easily seen by noticing that a deformation gradient \mathbf{F} can be obtained through a linear deformation $\boldsymbol{\varphi}(\mathbf{r}) = \mathbf{F}\mathbf{r}$ which keeps the system planar.




Figure 2.8 – To obtain the elastic force $d\mathbf{f}$ acting on the line element dl (—) at material point \mathbf{r} with normal \mathbf{n} , the formula $\mathbf{p} = b\mathbf{F}(\mathbf{r})\boldsymbol{\sigma}(\mathbf{r})\mathbf{n}$ can be used. Note that it is natural for this picture to be in two dimensions as we are dealing with two dimension elasticity in the context of shells.

Second Piola-Kirchhoff stress

In our work, we chose the membrane energy model based on observations obtained from experiments (see Chapter 3). Since those observations relate the deformation to a measured force, we need a way to derive forces from the present membrane energy ψ_m . Such a quantification of force can be done through the *second Piola-Kirchhoff stress tensor* $\boldsymbol{\sigma}$ which is a second order tensor given by⁹

$$\boldsymbol{\sigma} = \frac{1}{b} \frac{\partial \psi_m}{\partial \boldsymbol{\varepsilon}}, \quad (2.4)$$

where b is the thickness of the considered shell. This tensor quantifies forces acting on surface element. Concretely, the force \mathbf{p} acting on a line element dl with normal \mathbf{n} at material point \mathbf{r} , both given in the *reference state*, can be computed as

$$\mathbf{p} = b\mathbf{F}(\mathbf{r})\boldsymbol{\sigma}(\mathbf{r})\mathbf{n}. \quad (2.5)$$

Hence, a total force \mathbf{f} acting on a line l can be obtained by integrating this force over the line,

$$\mathbf{f} = \int_l b\mathbf{F}(\mathbf{r})\boldsymbol{\sigma}(\mathbf{r})\mathbf{n} dl. \quad (2.6)$$

The quantities involved in these relations are schematised in Figure 2.8.

Nonlinearity

We can note that the use of the Green-Lagrange strain $\boldsymbol{\varepsilon} = \frac{1}{2}(\mathbf{F}^T\mathbf{F} - \mathbf{I})$ introduces a geometric nonlinearity in the model. A non-linearity also comes from the second Piola-Kirchhoff stress tensor $\boldsymbol{\sigma}$ which is multiplied by the deformation gradient \mathbf{F} to obtain the force. Geometrically linear models which use the *Cauchy strain* $\boldsymbol{\varepsilon}^l = \frac{1}{2}(\mathbf{F} + \mathbf{F}^T)$ as a measure of deformation and the *Cauchy stress* $\boldsymbol{\sigma}^l = \frac{1}{b} \frac{\partial \psi_m}{\partial \boldsymbol{\varepsilon}^l}$ to derive the force. While the two measures highly simplify

⁹In mathematical derivations of elasticity theory, stress tensors is usually obtained as way to formulate forces in the context of continuum mechanic rather than from the energy (see Ciarlet (1988)). However, we prefer to present it through the relations that are central to the understanding of our work.

the equations of elasticity thanks to their linearity, they do not properly take into account rotation of the material. This can be seen in the Cauchy strain tensor which is not invariant by rotation of the deformation. And this issue is also apparent in the computation of force. A linear force density \mathbf{p} is computed as follows using the Cauchy stress tensor $\boldsymbol{\sigma}^l$

$$\mathbf{p} = b\boldsymbol{\sigma}^l(\mathbf{r})\mathbf{n}. \quad (2.7)$$

Here, compared to the expression which uses the second Piola-Kirchhoff stress $\boldsymbol{\sigma}$, there is no multiplication by the deformation gradient, which is important to orient the force in the right direction.

Since the vane deforms quite easily when acted upon manually, we cannot use the Cauchy strain which are not fit for handling large displacement. Hence, our model will be based on the Green-Lagrange strain and second Piola-Kirchhoff stress tensor. Those will be respectively named *strain* and *stress* in the rest of the document, as there will be no ambiguity with other type of stress and strain.

In anticipation of our analytical model which will be presented in Chapter 4, we note that this geometrical non-linearity will be the only source of non-linearity within our material. Indeed, our experiments show that the relationship between stress and strain is linear. This can be striking as a linear relationship between stress and strain are often only valid for small deformation. However, we will show that, despite the presence of large deformation within our experiments, this linear relation between stress and strain is in fact able to properly model the elastic behaviour of the vane.

Chapter 3

Lab experiments

As discussed in the previous chapter, we aim at an elastic membrane model for the vane which requires few parameters. Most of all, we want this model to be inspired by quantitative observations on the vane made through experiments. In this chapter, we describe the experiments that helped us design our vane model and also discuss the observations we could make from them. Note that these experiments are not exclusively used for the design of the model: the data gathered through them is also central to the validation of the model (see Chapter 5).

The work presented in this chapter has been divided between me and one of my supervisors Victor Romero-Graemeña. Victor focused on setting up the experiments and gathering the resulting data, and I focused on the design of the experimental scenarios and the analysis of the data.

3.1 Scenarios

As presented earlier, our membrane model takes the form of an energy density $\psi_m(\boldsymbol{\varepsilon})$, where $\boldsymbol{\varepsilon}$ is the strain tensor. Ergo, it would be natural to set up experiments whose observation tells us how the internal energy of the material varies with its deformation. However, to the best of our knowledge, it is not possible to directly measure energy. Hence, the model form needs to be inferred from the evolution of other quantities. In the case of elasticity, it is natural to go back to how they were originally modelled and look at how forces behave as the material is deformed¹.

The strain tensor, which quantifies local deformation, can be expressed in matrix form by choosing a basis,

$$\begin{bmatrix} \varepsilon_{xx} & \varepsilon_{xy} \\ \varepsilon_{xy} & \varepsilon_{yy} \end{bmatrix},$$

where the first basis vector is taken as the x-axis and the second one as the y-axis. The components of this matrix change with different kinds of deformation,

¹Actually, it is also impossible to measure forces directly. Instead, usual force sensor measure the deformation of small material in contact with the point at which we would like to know the force. Models for the small material use are well known, which allows one to derive the force from the deformation of this material. No such mechanism currently exists for measuring elastic potential energy.




Figure 3.1 – Each component of a matrix representation of the strain tensor is correlated with a specific deformation.

ε_{xx} with an extension along the x-axis, ε_{yy} with an extension along the y-axis, and ε_{xy} with a shearing (see Figure 3.1).

Hence, to analyse the behaviour of the feather vane with the purpose of building a model, we find it natural to use three experiments which would each excites one of these deformations². As the deformations depend on the choice of the basis, we have to fix one to properly choose what deformation we want to perform in our experiments. Due to the underlying geometry of the feather vane, it is natural to choose the x-axis to be along the barbs and the y-axis to be orthogonal such that they form a direct basis. This is the basis that will be used in the rest of the document.

More concretely, to study those deformations, we place rectangular vane samples between two clamps which are moved to enforce the wanted target displacement (see Figure 3.2). The force required to move the clamps is measured, giving us a force-displacement relation that we are able to study. More details concerning this protocol are given in the next section.

Considering those three specific deformations to study the membrane behaviour of an anisotropic material is nothing new. In the textile industry Kawabata and Niwa (1989) suggested an objective way to measure textile qualities hoping that it supersedes the more subjective methods used at the time. In the graphics community, this method has been used to measure the parameters of a Saint-Venant Kirchhoff membrane model for textiles (Volino et al. 2009).

In these scenarios, the vertical edges of the sample are not constrained, hence, the behaviour of the sample at the edge can be substantially different from the one at the centre of the sample. If such a phenomenon were to appear, it would hinder the easy analysis of the scenarios. Volino et al. (2009) mitigate this by considering samples that are far wider than they are high (see Figure 3.3). While this trick is easy to perform with cloth which can come in large size, it does not transfer easily to the feather. Cutting such samples from the feather vane is not really possible. Those samples need to be large enough

²If the feather vane was known to be isotropic, which it is not, we could have limited our study to the stretch in only one direction. While the direction orthogonal to the stretch would not be constrained, its deformation would still have to be measured to derive the full model. For example, it would be necessary to find the Poisson's ratio if the model was considered linear.




Figure 3.2 – Experimental scenarios each related to the deformation associated to the components ε_{xx} (left), ε_{yy} (centre) and ε_{xy} (right) of the matrix representation of the strain tensor. Those are the same scenario considered by Kawabata and Niwa (1989) and Volino et al. (2009) in their works.

to average away the individual behaviours of barbs, but we are also limited by the size of the vane itself. This prevents us to have a large difference between the width and height of the sample.

Thankfully, in the extension scenarios, the behaviour at the free edges of the sample is the same as at any point of the feather. However, the phenomenon does appear in the shearing scenario in the form of buckling (see Figure 3.4). We could have put in the effort to deal with the more complicated problem of in-plane elasticity due to the phenomenon appearing at edges. But, to the best of our knowledge, no post-processing effort could deal with the buckling which prevents us from using this scenario to study the membrane behaviour of the vane.

To avoid this buckling behaviour without preventing our study of the shearing deformation, we instead use another extension scenario with an orientation at a quarter turn from the x-axis. While this scenario still displays a non-uniform behaviour due to the free edges, it is exempt of buckling on the range of deformation considered.

In the rest of the document, the scenario in which an extension along x-axis is done will be called the *longitudinal scenario*, the one in which the extension is along y-axis, *transverse scenario* and the last one *tilted scenario*.




Figure 3.3 – In their experimental setup, Volino et al. (2009) used rectangular cloth samples that were much wider than high. This allows them to mitigate the effect of the cloth at its free edges.




Figure 3.4 – First (left) and last (right) state of the shearing scenario. In this scenario the sample vane buckles out of plane. This buckling is visible as a darker part of the sample at the top of the last state of the shearing scenario. This prevents the proper analysis of the membrane behaviour of the vane.




Figure 3.5 – To study shearing deformation within the feather vane, we consider an extension test oriented at a quarter turn from the x -axis. Unlike the previously presented shearing scenario, the extension test doesn't present any buckling in the range of deformation we are interested in.

3.2 Setup

In the previous section, we discussed the different experimental scenarios we perform to study the feather vane. In this section, we describe precisely how we set up those scenarios in the lab and what we chose to measure.

Experiment on a sample

When a sample is too large, the orientation of barbs can vary across its surface. The heterogeneity incurred by such a variation increases the complexity of the behaviour analysis. Also, if a sample is too small, we will not be able to focus our study on the overall behaviour of the vane. Hence, we take our samples as large as possible while keeping the barb orientation uniform. Also, in the tilted scenario, it is important to look at samples with a range of different aspect ratios between their two sides. The significance of this aspect ratio on the behaviour of the samples will be shown in Section 3.3.

Samples are cut out from the vane using a paper guillotine³. They 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 ε_{\max} . For the transverse scenario, the maximal strain is chosen to be $\varepsilon_{\max} = 0.4$ as fracture appears for greater strains. In the tilted scenario, strain at fracture is highly dependent on the aspect ratio of the sample. Hence, instead of choosing the maximal strain from the point of fracture, we arbitrarily use the same maximal strain used for the transverse scenario. In the longitudinal scenario, the sample slips out of the clamp early in the experiments due to high stiffness of the deformation. Therefore, we chose a maximal strain before skipping occurs, which we found to be around $\varepsilon_{\max} = 0.01$. In each step we displace the motor to a given position, stop, and make a force measurement, denoted by f , for 2 s at 5 kHz. Hence, one data point is the average of 10 000 measures, allowing us to highly reduce the uncertainty on the force measurement. In addition, the displacement, denoted by d , is taken from the readings of the motors (see Figure 3.6).

Beside the force and the displacement, we also measure the width, denoted by W , and height, denoted by H , as they are necessary to relate the measured force and displacement to stress and strain. In addition to those, we also measure the thickness b of each sample, since it has a substantial impact on the measure of forces, and we have seen it to vary across the vane. 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 micrometer, we perform thickness measurements after the traction test is done.

³The paper guillotine was chosen to produce the samples as it allow precise cut — through the juxtaposition of a ruler — without tearing the barbs apart during the process. We could have considered other obvious tools beside the paper guillotine such as a knife or a pair of scissors. But they were not appropriate for the task. A knife would have allowed the use of a ruler, but the shearing induced by the cut induces a separation of the barbs. And, while a pair of scissors would have avoided any damage, its lack of precision was problematic.




Figure 3.6 – **Top:** Real feather and example patches. **Bottom left:** Experimental setup for traction tests. **Bottom right:** Zoom on the traction experiment.




Figure 3.7 – During the traction of each sample, we measure force f and displacement d . Also, width W , height H and thickness of the sample are measured.



Figure 3.8 – We run our experimental scenarios with mute swan feathers (top) and guinea fowl feathers (bottom). The mute swan feathers are especially easy to work with due to their mostly uniform barb orientation. Each square shown under the feather has 1 cm sides

The setup and execution of each of experiment can take thirty minutes of manipulation on average. A great deal of precision and attention is required throughout the process to make sure that no sample breaks during the testing. This high cost of experimental data generation will drive us to produce a multistep validation protocol to avoid any additional in-lab work (see Chapter 5).

Origin of samples

While those precise measures allow us to properly analyse each sample individually, we also need to be careful of the origin of each of them to allow proper comparison between observations obtained from two different samples.

In fact, due to the biological nature of the feather vane, there can be substantial variation between measures on distinct samples. We noticed that this variation is especially high when looking at feathers of different species. This variability is also present to a lesser extent between samples taken at different places in the vane.

To simplify our work, we chose to focus our observations to the behaviour of feathers of only two bird species, the mute swan and the guinea fowl (see Figure 3.8). Our quantitative analysis mainly focuses on the mute swan as their feathers are the easiest to deal with among the cheap and commercially available feathers⁴. The observations of the guinea fowl feather have the purpose of a somewhat quick check that our qualitative observations on feather of swans at the bases of our model are also found in other bird species. Also, to analyse the origin of the variability between different region of the vane, we diversify the region from which the samples are cut out, while still making sure that each sample have a uniform orientation of barbs.

⁴Their feathers are large — around 25 cm, making them easier to manipulate. Also, the barb orientation is mostly uniform across the vane, making it easier to cut out uniform samples out of them.




Figure 3.9 – Measured force-displacement in load-unload cycles of increasing amplitude of a vane sample. After one cycle, the force is negative in the final state while it was null in the initial state. This hysteresis effect shows that there are dissipative effect in plays in the vane mechanics. We believe that those dissipative effect are due to the friction between barbules while they slide against each others.

Hysteresis

Instead of considering scenarios in which only one single traction is performed, as we have done in this PhD, load-unload cycles could have been performed. We first performed such cycles and observed hysteresis, showing that the deformation is actually not purely elastic (see Figure 3.9). We believe that this hysteresis is inherent to the frictional nature of the connection between barbules. Since we limit our work here to a non-dissipative model, we chose to only present pure traction experiments (without cycling) from which we can obtain elastic parameters.

3.3 Results and observations

Measure

As discussed in Section 2.3, our model does not work directly with forces and displacements but with stress and strain tensors. Hence, instead of looking at the relation between the measured force and displacement, we rather observe how a measure analogous to the stress tensor relates to one analogous to the strain tensor.

The measure we chose are the components of the strain and stress tensors along the axis aligned with the movement of the clamped, denoted here as y -axis, ε_{yy} and σ_{yy} , under the assumption of uniform deformation of the sample. Those measures are completely defined by the measured force f , size of the sample W , H , and b as well as the measured displacement d . In this context of

assumed uniformity, we will express the different quantity with a $\tilde{\cdot}$ to put an emphasis on the fact that they arise from an assumption not necessarily verified in our experiments. With these considerations, the deformation gradient $\tilde{\mathbf{F}}$ is expressed as,

$$\tilde{\mathbf{F}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 + \frac{d}{H} \end{bmatrix},$$

where we omitted the dependence on the material point as the quantity is uniform over the sample. Hence, using Equation (2.3), the strain tensor can be expressed as

$$\tilde{\boldsymbol{\varepsilon}} = \begin{bmatrix} \tilde{\varepsilon}_{xx} & \tilde{\varepsilon}_{xy} \\ \tilde{\varepsilon}_{xy} & \tilde{\varepsilon}_{yy} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 0 & 0 \\ 0 & (1 + \frac{d}{H})^2 - 1 \end{bmatrix}.$$

The stress $\tilde{\boldsymbol{\sigma}}$ can be linked to the measured force under the same consideration of uniformity by integrating along the clamped edge,

$$\begin{aligned} f &= b\hat{\mathbf{y}}^T \int_{\Gamma_t} \tilde{\mathbf{F}}\tilde{\boldsymbol{\sigma}}\mathbf{n} \, dr \\ &= bw \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 + \frac{d}{H} \end{bmatrix} \begin{bmatrix} \tilde{\sigma}_{xx} & \tilde{\sigma}_{xy} \\ \tilde{\sigma}_{xy} & \tilde{\sigma}_{yy} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ &= bw(1 + \frac{d}{H})\tilde{\sigma}_{yy} \end{aligned}$$

where Γ_t is the clamped edge of the sample, \mathbf{n} is the normal of this edge going outward of the sample, and $\hat{\mathbf{y}}$ is the unit vector along y-axis. Inverting this we obtain,

$$\tilde{\sigma}_{yy} = \frac{H}{bW(H+d)}f.$$

While the uniform longitudinal stress $\tilde{\sigma}_{yy}$ and the uniform longitudinal strain $\tilde{\varepsilon}_{yy}$ are not necessarily the actual stress and strain present in the sample, they have the nice property to be independent of the size and thickness of the sample, and they are the closest measures to the actual stress and strain we could find. Moreover, as discussed later in this section, it happens that the feather vane deforms uniformly in the longitudinal and transverse scenarios, making our chosen measures actually match the strain and stress present in the sample.

Variability

The data we gathered through our experimental scenarios are plotted in Figure 3.10. The first thing noticeable is the dispersion of the data. In both the transverse and longitudinal scenarios the slopes of plotted curves vary by $\pm 25\%$ around their middle value. Considering our choice of measure, it is unlikely that this variation comes from a difference in sample size. We believe that this dispersion comes from two other possible sources. Either from the intrinsic variation expected from biological materials⁵. Or from the hysteresis

⁵Unlike materials produced in the industry, where an emphasis is put on repeatability and uniformisation of processes, biological material can have some slight differences from one object to another. For example, knot in wood are not always found at the same places, and two feathers found at the same place on two different birds of the same species might not have the exact same geometry.

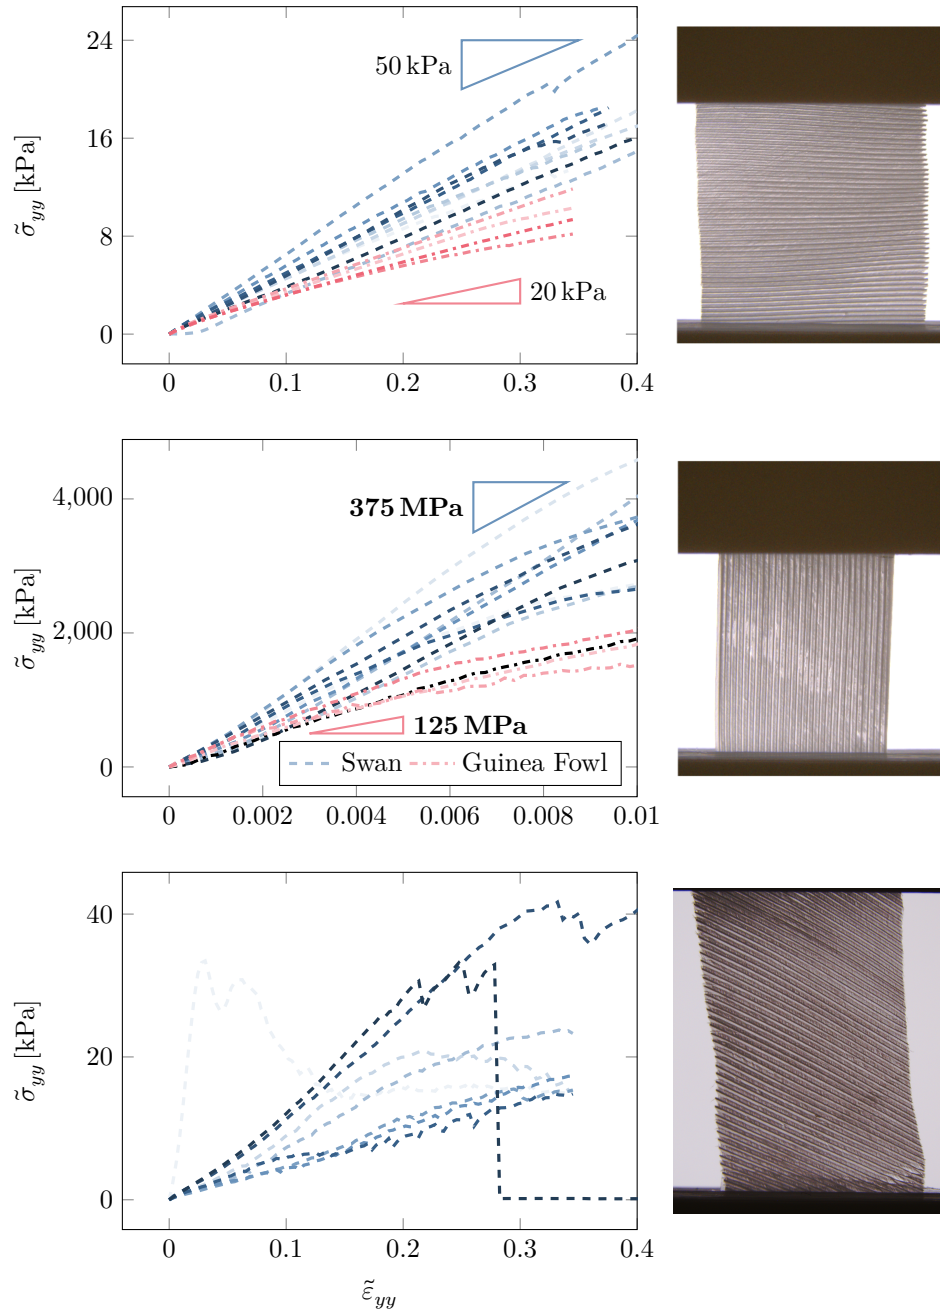


Figure 3.10 – Experimental stress-strain measurements for feather samples from swans (---) and guinea fowls (-.-.-) in the three experimental scenarios: transverse (top); longitudinal (middle); tilted (bottom). Pictures of one of the samples on swan’s feathers are shown on the right for each scenario. The geometry of the samples used here are listed in Table 3.1. In these measurements, one can notice the flagrant linear relation between stress and strain in the transverse and longitudinal scenario. This linear relation is not really visible in the tilted scenario, but might be explainable by the geometric non-linearity arising from the Green-Lagrange strain formulation. Another noticeable aspect of these measurements is the already mentioned strong anisotropy. The stiffness measured in both the transverse and longitudinal scenario differ by four orders of magnitude, which, as said in Chapter 1, is larger than similar materials which only have two orders of magnitude difference between their highest and lowest stiffness.

	W [mm]	H [mm]	b [μm]		W [mm]	H [mm]	b [μm]
---	15	14	120	---	9	18	73
---	8	14	100	---	7	15	133
---	12	11	120	---	9	23	60
---	10	11	110	---	8	13	120
---	10	13	95	---	9	16	129
---	12	11	140	---	4	12	116
---	11	13	120	---	10	11	80
---	10	15	100	---	7	19	125
---	17	14	112	---	10	10	117
---	10	13	125	---	6	6	100
----	11	12	120	----	7	10	145
----	16	12	150	----	2	6	130
----	8	10	110	----	5	11	140
----	11	12	120				

(a) Transverse scenario				(b) Longitudinal scenario			
	W [mm]	H [mm]	b [μm]		W [mm]	H [mm]	b [μm]
---	6	15	120				
---	3	14	110				
---	9	16	132				
---	12	16	155				
---	8	14	143				
---	16	16	130				
---	14	19	114				
---	13	18	103				

(c) Tilted scenario			
	W [mm]	H [mm]	b [μm]
---	6	15	120
---	3	14	110
---	9	16	132
---	12	16	155
---	8	14	143
---	16	16	130
---	14	19	114
---	13	18	103

Table 3.1 – Width W , height H , and thickness b of the samples whose stress-strain relation are measured in Figure 3.10. The width and height of samples used for the tilted scenario are especially important as the behaviour is highly influenced by the aspect ratio of the sample (see Figure 3.11). These geometrical measures are represented in Figure 3.7.

noticed in the previous section. Due to this non-elastic behaviour, a feather sample does not return to its initial configuration when relaxed. This introduces a variability in the possible initial state of the tested sample, and, hence, a variability in the observed measures.

In the tilted scenario, a stronger variability can be seen, however, unlike the longitudinal and transverse scenario, most of the variation comes from the aspect ratio of the tested sample. Here, the narrower a sample is the more barbs are not constrained by the clamps (see Figure 3.11). Those unconstrained barbs are free to move relative to each other, hence, instead of calling on their bending stiffness to deform as the clamped barbs do, they rather rely on the less stiff deformation of the barbules. This makes narrow samples much softer. The strong impact of the aspect ratio on the behaviour of samples in the tilted scenario will need to be properly taken into account in the validation process




Figure 3.11 – Pictures of samples used in the tilted scenario, each having a different aspect ratio in their rest state (top). When elongated (bottom), the behaviour of samples is highly dependent on their aspect ratio. The more elongated a sample is, the more barbules are not constrained by the clamps. Hence, more barbules are free to move, making the sample less rigid during the deformation.

(see Chapter 5).

Stiffness

Another noticeable aspect in the data is the stiffness observed in the longitudinal scenario. The stiffness there, mostly coming from the stiffness of barbs, is four order of magnitude higher than in the transverse scenario where the deformation is mostly supported by the barbules. Due to this stiffness we could not observe how samples behaved for large deformations as they were sliding out from the clamps. Thankfully, we should not need to look at larger deformation along the barb orientation, since other much softer deformations should occur before the material ever attain higher deformations.

Aside from preventing us to reach large displacements, this large stiffness ratio between transverse and longitudinal deformations is indicative of a strong anisotropy rarely seen in most materials studied in computer graphics and mechanics. This strong anisotropy will make our numerical modelling efforts quite arduous as will be explained in Section 4.3.




Figure 3.12 – Picture of sample used in the transverse scenario in its initial state (left) and final state (right). In this scenario, samples do not show any change in width when they are elongated, i.e. they do not show any Poisson effect. The Poisson’s ratio can be computed from the width W and height H at rest, at in the deformed state W' , H' : $\nu_y = -\frac{\tilde{\epsilon}_{xx}}{\tilde{\epsilon}_{yy}} = -\frac{(W'^2 - W^2)H^2}{(H'^2 - H^2)W^2}$. Here, since we have $W = W'$, the Poisson’s ratio is null $\nu_y = 0$.

Poisson’s effect

In the transverse scenario, the width of samples does not seem to vary as the material is deformed (see Figure 3.12) This is unusual as the width of most materials decrease as they are pulled to preserve their volume, a phenomenon commonly referred to as the *Poisson’s effect*⁶. Materials displaying such a behaviour are, among others, rubber, most metals, some textiles. The relation between length and width change is usually quantified by the Poisson’s ratio, denoted by ν_y , where x is the axis along which the material is stretched. This is computed from the ratio between the strain in the direction of extension $\tilde{\epsilon}_{yy}$ and the strain in the direction transverse of the extension $\tilde{\epsilon}_{xx}$.

$$\nu_y = -\frac{\tilde{\epsilon}_{xx}}{\tilde{\epsilon}_{yy}}. \quad (3.1)$$

This means that this ratio is positive for materials that tend to keep their volume as they deform — and taking the value 0.5 when exactly incompressible. In our case, the Poisson’s ratio seem to be null, which is striking as the underlying material, β -keratin is known to have a Poisson’s ratio of 0.4 (Khani et al. 2022).

We believe that the apparent absence of Poisson effect in the feather vane is due to the underlying structure of the vane and mostly to the stiffness of the barbs compared to that of the barbules. In the transverse scenario, the width of the sample is given by the length of the barbs. For the considered range of deformation in the transverse scenario, stresses are not strong enough to deform the barbs, hence the absence of width change.

⁶We wrote “most materials”, some do not present such behaviour, as is obviously seen here in the case of the feather vane. For another example of materials whose width does not decrease while they are stretched are the set of auxetic materials, instead their width increase, their volume becoming larger as they deform.

In the longitudinal scenario, the maximum displacement is too small to make any conclusion on the presence of Poisson effect without more precise measuring devices. However, since the deformation of the barbs should be negligible due to the range of stresses considered, we will still consider the Poisson effect to be null in this orientation in our modelling. Moreover, it will be seen in our modelling section (see Section 5.3) that the absence of Poisson effect in one direction implies an absence of Poisson effect in orthogonal direction for the sake of energy conservation.

Unlike the longitudinal and transverse scenario, the tilted scenario displays a strong Poisson's effect. This would be unexpected if we were dealing with an isotropic material, where the behaviour should be the same in all directions, however, such a Poisson's effect is expected from an anisotropic material. We will see in Chapter 5 that our model can properly capture this effect.

Overall behaviour

Let focus back on our original aim: building a model for the vane. The stress-strain relation observed in the context of the transverse scenario strikes as being linear. In light of this, we will be considering a linear elastic model — but geometrically non-linear due to the use of the Green-Lagrange strain — in the next section. This assumption might be strong at first glance as the data produced through the other scenarios are not quite linear.

First, in the longitudinal scenario, the stress slightly softens toward the end of the observed range of strain. While this could be problematic if we had to consider much larger deformations where such non-linear behaviour would have a substantial effect, extension of barbs will be kept small and in the regime of linear behaviour thanks to the high stiffness ratio as pointed out in the previous subsection.

Secondly, in the tilted scenario the stress seems to be subject to some substantial stiffening until the barbule separation comes into play and causes the sharp drop of stress. Here, this non-linear behaviour cannot be dismissed as easily as we have done for the longitudinal scenario where the range of deformation was much smaller. Nevertheless, as stipulated earlier in this section, the phenomenon shown in this scenario are more complex and their non-linear nature is not sufficient to dismiss our choice of linear relation between stress and strain. In fact, this non-linear behaviour could well be predicted by our suggested model thanks to the non-linear representation of the deformation through the Green-Lagrange strain as will be shown in Chapter 5.

For all these reasons, we chose to model the membrane behaviour of the feather vane through a linear relation between the stress and the strain. The formal description of this model will be discussed in the next chapter.

Mute swan and guinea fowl

In Figure 3.10 we showed the results for samples taken from feathers of both mute swans and guinea fowls. The guinea fowl samples are less stiff than the swan feathers, but their behaviour is similarly linear and strongly anisotropic. Hence, while our quantitative validation of our proposed membrane model, presented in Chapter 5, uses only on the data obtained from swan feathers, we

believe that it is highly probable that those validation results could be extended to other bird species.

Chapter 4

Macroscopic model for the membrane

4.1 Analytical Model

In the previous chapter, our observations suggested a linear relationship between stress $\boldsymbol{\sigma}$ and strain $\boldsymbol{\varepsilon}$, this means that our membrane model takes the form

$$\boldsymbol{\sigma} = \mathbf{E} : \boldsymbol{\varepsilon},$$

where \mathbf{E} is a fourth order tensor usually called the *elasticity tensor*.

Note that the suggested model is equivalent to an anisotropic Saint-Venant Kirchhoff model. This model is usually used to model the linear relation between stress and strain due to small deformations. However, in our case, the choice of the model does not arise from such considerations. In fact, as seen in the previous chapter, the considered deformation are quite large, going up to 40%. This model was chosen simply because it seems natural regarding our data.

A simple expression for the elasticity can be worked out if we express its component in the orthonormal basis whose first axis is along the barbs. Note that this basis is different from the one used in the description of the scenarios. To make this difference clear, we denote the axis along the barb by v-axis, and the one orthogonal to the barbs by w-axis. Using indices i, j, k and l that spans v and w , we have

$$\sigma_{ij} = \sum_{k,l} E_{ijkl} \varepsilon_{kl}.$$

Here, the elasticity tensor components E_{ijkl} must satisfy some well, known conditions for the model to be a proper elastic model.

Constraint due to general elasticity

Since both the stress tensor $\boldsymbol{\sigma}$ and the strain tensor $\boldsymbol{\varepsilon}$ are symmetric we must have,

$$E_{ijkl} = E_{ijlk} = E_{jikl}.$$

This allows us to use Voigt's notation, which is often used in the context of elasticity theory as it is simpler to work with. This notation consists in writing




Figure 4.1 – In our notations, we choose the v -axis to be along the barb and the w -axis to be orthogonal to the barbs. Everywhere, the vane is locally symmetric with respect to the v -axis (.....).

the independent component of the stress and strain tensor in vector form, and the components of the elasticity tensor in matrix form:

$$\begin{bmatrix} \sigma_{vv} \\ \sigma_{ww} \\ \sigma_{vw} \end{bmatrix} = \underbrace{\begin{bmatrix} E_{11} & E_{12} & E_{13} \\ E_{21} & E_{22} & E_{23} \\ E_{31} & E_{32} & E_{33} \end{bmatrix}}_{\mathbf{E}} \begin{bmatrix} \varepsilon_{vv} \\ \varepsilon_{ww} \\ 2\varepsilon_{vw} \end{bmatrix}.$$

Note that the shearing ε_{xy} is multiplied by two. This is to make sure that the double contraction of the stress and strain tensor is equal to the scalar product of their vector form, i.e. $\boldsymbol{\sigma} : \boldsymbol{\varepsilon} = [\sigma_{xx} \ \sigma_{yy} \ \sigma_{xy}] [\varepsilon_{xx} \ \varepsilon_{yy} \ 2\varepsilon_{xy}]^T$.

Here, to ensure energy preservation, the matrix form of the elasticity tensor needs to be symmetric.¹

$$\begin{bmatrix} \sigma_{vv} \\ \sigma_{ww} \\ \sigma_{vw} \end{bmatrix} = \begin{bmatrix} E_{11} & E_{12} & E_{13} \\ E_{12} & E_{22} & E_{23} \\ E_{13} & E_{23} & E_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_{vv} \\ \varepsilon_{ww} \\ 2\varepsilon_{vw} \end{bmatrix}.$$

Constraint due to the feather vane

This expression of a general membrane model linear in the strain has six parameters, but we can reduce them to three using the following geometric and mechanical considerations specific to the feather vane.

First, at any point on the vane of a feather, the vane is locally symmetric around the axis collinear to the barbs (See Figure 4.1). This symmetry, usually called *orthotropy*, removes some components from the elasticity tensor (Milton 2002, p. 37), namely $E_{13} = 0$ and $E_{23} = 0$.

Secondly, in Section 3.3 we have seen that the Poisson's ratio in the transverse scenario ν_w is very close to zero as we were unable to see any Poisson effect. Since, this Poisson's ratio can be expressed in terms of our model parameters as (see Appendix A.1 for the details of the derivation)

$$\nu_w = \frac{E_{12}}{E_{11}},$$

¹Otherwise, given any elastic object, it would be possible to formulate a series of deformations that cycles back to a rest configuration of the considered object while generating work, creating this way a free source of energy. If we require the stress to derive from a hyperelastic energy, which intrinsically conserve energy, we would naturally fall on the same condition for the elasticity tensor.

we know that $E_{11} \gg E_{12}$. This can give the intuition that E_{12} should be set to 0. However, to properly decide on this, the Poisson's ratio along the barb ν_v , i.e. in the longitudinal scenario have to be considered:

$$\nu_v = \frac{E_{12}}{E_{22}}.$$

Here, if the E_{12} is set to 0, this would imply the absence of Poisson's ratio in this direction. But, E_{12} could very well be of the same order as E_{22} , which would imply the presence of a Poisson effect when stretching along the barb. The presence of this Poisson effect is difficult to determine from our longitudinal scenario, since the tested deformation range is small (less than 1%). However, since deformations along the barbs should very small due to their apparent stiffness, we will consider that there is virtually no Poisson effect in this direction, meaning that $E_{22} \gg E_{12}$. With this, we set E_{12} to 0.

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

$$\begin{bmatrix} \sigma_{vv} \\ \sigma_{ww} \\ \sigma_{vw} \end{bmatrix} = \begin{bmatrix} E_{vv} & 0 & 0 \\ 0 & E_{ww} & 0 \\ 0 & 0 & E_{vw} \end{bmatrix} \begin{bmatrix} \varepsilon_{vv} \\ \varepsilon_{ww} \\ 2\varepsilon_{vw} \end{bmatrix}, \quad (4.1)$$

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_{vv} as the *longitudinal modulus*, E_{ww} as the *transverse modulus* and E_{vw} as the *shear modulus*.

Discussion

This elasticity model is close to the a linear orthotropic elastic materials that as already been used for cloth simulation (Volino et al. 2009). However, it presents some differences. First, our model is heterogeneous across the shell. Second, due to the extreme stiffness ratio between barbs 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.

4.2 Numerical model

As stipulated in the introduction, we aim to have a numerical model of the vane model. This would allow us to validate our membrane model and produce a feather numerical model with the purpose of helping animator in their work. The first step consists in building a numerical model of the different experimental scenarios.

To do so, we start from by presenting the associated continuous elasticity problem. Our experiments are done in a quasistatic manner: as said in the previous chapter, after each displacement the sample is kept in place for 2s while measurements are done. Hence, we consider a static elasticity problem rather than a dynamic one as it is easier to analyse². We denote the sample

²As discussed later in this section, the static problem is implemented numerically as a dynamic problem due to the limitation of our software. Hence, the use of the static case for the base problem might seem unnatural considering our numerical implementation. However, this static problem is much easier to handle in all our analytical analysis of the experiments.




Figure 4.2 – Schematic of the reference state (left) and deformed state of the scenarios considered in Section 3.1 with the notations associated to Equation 4.2.

reference configuration by $\Omega = [0, W] \times [0, H]$, the top clamped edge by $\Gamma_t = [0, W] \times \{H\}$, the bottom clamped edge by $\Gamma_b = [0, W] \times \{0\}$, and the side edges by $\Gamma_s = \{0, W\} \times [0, H]$. With these notations, the static elasticity boundary value problem for a given scenario can be written as

$$\begin{aligned}
 -\nabla \cdot \mathbf{F}\boldsymbol{\sigma}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Omega; \\
 \mathbf{u}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Gamma_b; \\
 \mathbf{u}(\mathbf{r}) &= d\hat{\mathbf{y}}, \forall \mathbf{r} \in \Gamma_t; \\
 \mathbf{F}\boldsymbol{\sigma}(\mathbf{r})\mathbf{n}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Gamma_s,
 \end{aligned} \tag{4.2}$$

where $\boldsymbol{\sigma}$ is the stress computed through our membrane model presented in the previous section, d represents the measured displacement in the experimental scenario and $\mathbf{n}(\mathbf{r})$ is the normal at material point \mathbf{r} (see Figure 4.2). In this boundary value problem, the first and last equations enforce the balance of forces inside the material, $\nabla \cdot \mathbf{F}\boldsymbol{\sigma}$ being the surface force density³. The two others equations are the constraint imposed by the clamps.

The standard numerical methods for solving this problem are the finite element method and its derivatives. We used two different software to solve this problem: ARCSim simulator specifically catered to cloth simulation and already well known to the people of the graphics community and us (Narain et al. 2012); and FENICS an open-source general purpose finite element solver (Baratta et al. 2023). All results shown in this thesis were produced with ARCSim. FENICS was used to cross-validate the results of ARCSim on simpler scenario (see Figure 4.3)

The static problem given in Equation (4.2) was straightforward to implement with FENICS, in which we used triangular linear elements. However, some care must be taken when implementing it in ARCSim. For the purpose of this section, ARCSim can be seen as a finite element method targeted at dynamic elasticity problem which uses triangular linear elements and a lumped mass

³The last equation consists in the balance of forces on the edges of the sample. As nothing is acting on those edges, the forces should be null, otherwise the sample edges would be in movement and, hence, not in static equilibrium.

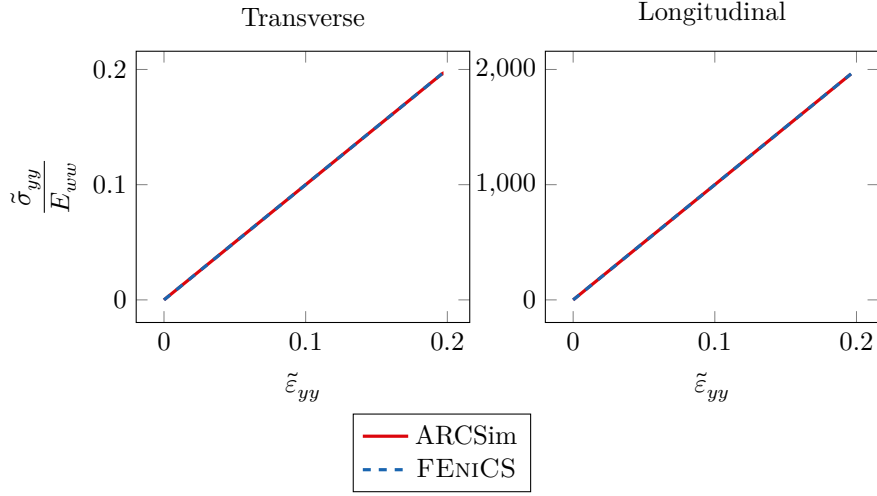


Figure 4.3 – We compare the measured stress $\tilde{\sigma}_{yy}$ nondimensionalised by the transverse modulus E_{ww} predicted by both ARCSim (—) and FENICS (---) in both the transverse (left) and longitudinal (right) scenarios. Simulations are run with a stiffness ratio $\alpha = \frac{E_{vv}}{E_{ww}}$ between the transverse and longitudinal modulus of 10^4 . Both ARCSim and FENICS predicts the same thing.

matrix⁴. Due to the dynamic nature of ARCSim, we cannot solve directly the static elasticity problem. Instead, we have to work from the dynamic equivalent of this boundary value problem,

$$\begin{aligned} \rho \ddot{\mathbf{u}}(\mathbf{r}) - \nabla \cdot \mathbf{F}\boldsymbol{\sigma}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Omega; \\ \mathbf{u}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Gamma_b; \\ \mathbf{u}(\mathbf{r}) &= d(t)\hat{\mathbf{y}}, \forall \mathbf{r} \in \Gamma_t; \end{aligned} \quad (4.3)$$

where $\ddot{\mathbf{u}}(\mathbf{r})$ is the acceleration at \mathbf{r} , ρ is the surface mass density, which is considered uniform over the sample, and the imposed displacement $d(t)$ now depends on the time. Note that the equation on the side edges is not present anymore as we are considering dynamic behaviour.

For this dynamic problem to behave quasistatically, we chose the surfacic mass density ρ such that elastic waves are able to propagate within the whole material in less time than a step size h . The elastic wave speed c for an isotropic material of Young's modulus E is given by

$$c = \sqrt{\frac{E}{\rho}}. \quad (4.4)$$

Hence, to be conservative we replace the Young's modulus E by the smallest modulus in our material and obtain the following constraint on the surface mass density ρ

$$\rho < \frac{h^2 E}{H^2}, \quad (4.5)$$

⁴ARCSim is not explicitly coded as such, however, in Section B.1 we prove the equivalence of the formulation.

where H is the height of the tested sample. In our simulations, we chose to set

$$\rho = \frac{h^2}{10H^2} \min\{E_{vw}, E_{vv}, E_{ww}\}. \quad (4.6)$$

In addition to making sure that elastic waves propagate fast enough, they also have to be dissipated. Thankfully, this is taken care by the numerical dissipation induced by the time integration used within the software.

In ARCSim, this dynamic problem is discretised in time using an implicit Euler method and discretised in space using linear triangular elements as presented earlier. The resulting non-linear problem is equivalent to a usual Galerkin's method with lumped mass matrix. Equations are given in Appendix B.1.

Previous versions of ARCSim were only solving a linearised version of this problem at each time-step. Since our problem is badly conditioned due to the high stiffness ratio, the solution provided by the linearised equation was too far from the proper solution of the non-linear problem. Hence, we replace the internal solver with the non-linear constrained optimisation problem solver IPOPT (Wächter and Biegler 2006). See Appendix B.1 for more details on the implementation.

4.3 Numerical locking

Unfortunately, one serious issue rapidly become apparent in our simulations. We identify the presence of *numerical locking*, i.e. the higher the stiffness ratio between the two material directions

$$\alpha = \frac{E_{vv}}{E_{ww}}, \quad (4.7)$$

the slower the convergence of the finite element model.

Numerical locking and how we deal with it is discussed in this section. In the next section, we discuss how we mitigated the ill-conditioning of our time-stepping problem.

Identifying Numerical Locking

Some modelling problem can sometimes result in solutions that are poorly discretised by textbook finite element methods. This is the case for incompressible fluid and elasticity, where solutions for displacement fields should have zero divergence (Brezzi and Fortin 1991). This also appears with some scenarios for thin shells, where only flexural deformations should appear and membrane deformations should be absent (Chapelle 1995). In both of those cases, conventional elements poorly discretise the set of solutions.

In practice, this results in a degradation of the convergence as the constraint on the problem increases (Chapelle 1995). For incompressible materials, this happens when the Poisson's ratio approaches 0.5. And, in the case of thin shells, this degradation occurs when the thickness vanishes.

Hence, we first identify the locking problem by analysing the convergence of the finite element model in the context of the tilted scenario as the stiffness ratio α increases (see --- in Figure 4.4). In those results, the clear degradation of

the convergence points toward a locking phenomenon, meaning that our choice of elements is inappropriate⁵.

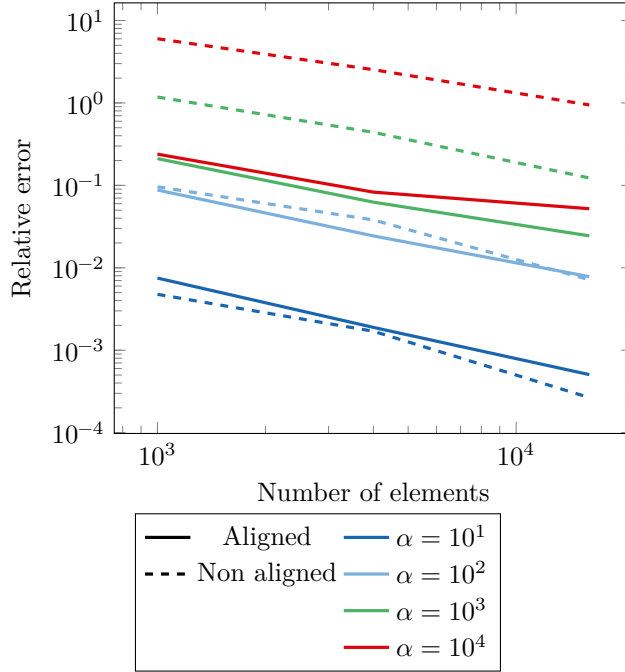


Figure 4.4 – Finite element convergence of the simulated tilted scenario with for different type of meshes and stiffness ratios α . Meshes are either aligned (—) or unaligned (---) with the orientation of barbs. The error is computed as the relative difference between the evaluated stress at 5% strain and the same stress for a simulation with aligned mesh and 128×10^3 elements. Its value is compared with the reference for the stiffness ratio $\alpha = 10^4$. As suggested by Chapelle’s methodology (Chapelle 1995), these plots clearly show a numerical locking phenomenon when an unaligned mesh is used: the convergence of the finite element method substantially deteriorates as the stiffness ratio is increased. This is shown by an upward shift in the curve in loglog scale, which means an increase in the prefactor in the actual relationship between relative error and number of elements. In contrast, the use of an aligned mesh drastically reduces this degradation of convergence.

In fact, the inability of the elements can be understood from the pictures of our in-lab tilted scenario (see Figure 4.5). Here, a discontinuity in the strain can be identified between the bound barbs and the shear band. This discontinuity implies a discontinuity in the derivative of the displacement, and is not properly discretised by our linear triangular elements mesh generated through a random Delaunay triangulation. Thankfully, in the case of linear elements, continuity of the derivative between elements is not enforced. Hence, we mitigate the locking phenomenon by aligning edges of elements with the discontinuity (see

⁵This 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 (J. S. Hale et al. 2018).





Figure 4.5 – Images of the first (left) and last (right) state of the tilted scenario. A discontinuity in the derivative of the displacement can be observed in the sharp transition between the lighter and darker region of the sample, marked by a red line here.

Figure 4.6). For the purpose of generalising this solution to scenarios that differ from the tilted scenario, we can notice that this discontinuity appears due to the high stiffness of barbs. Therefore, in all cases, elements should be aligned with barbs.

With this, degradation of the convergence with the increase of the stiffness ratio α is mitigated (see Figure 4.4 and Figure 4.7).



(a) Non aligned mesh

(b) Aligned mesh

Figure 4.6 – Examples of meshes used for the simulation of the vane samples. Those are meshes used for the tilted scenario. We either used mesh whose elements are not aligned with the barb direction (left) or mesh whose elements are aligned with those directions (right).

The locking phenomenon we show here is very similar to the one seen in the work of Yu et al. (2006), where, instead of having one really stiff direction, they have two very stiff directions and a low shearing stiffness. In this context, discontinuity in the strain also appears. Our solution of aligning elements with the discontinuity was taken from their work.

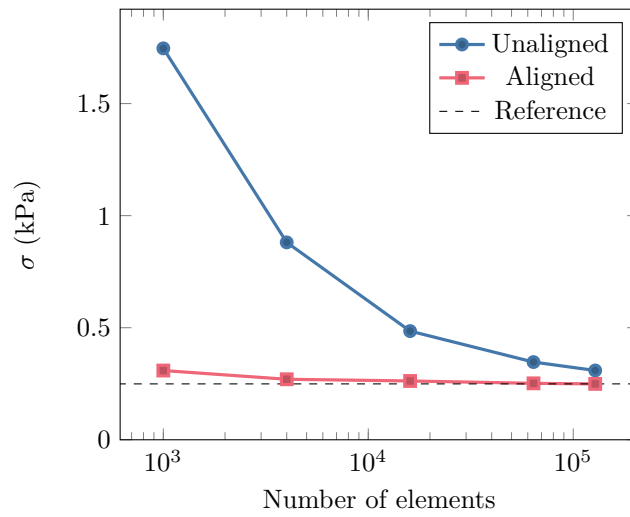


Figure 4.7 – Stress σ for a deformation of 5% in the tilted scenario, as a function of the number of elements. The stiffness ratio is fixed to our experimental value $\alpha = 10^4$. The convergence for simulations using an arbitrary mesh (\bullet) is considerably worse than an aligned mesh (\blacksquare). The reference value ($---$) is the one obtained with a finer aligned mesh made of 128×10^3 elements.

Chapter 5

Calibration and validation

Since it is impossible to mathematically *prove* that any model is able to predict anything¹, their predictive power is usually demonstrated through validation. Validation generally consists in showing the considered model is capable to predict something on which it was not build upon. Some well-known examples of such validation are the prediction of black holes by the theory of general relativity, or the prediction of the planet Uranus in the solar system through the gravity model of Newton. In both cases, the structure and parameters of the model were not chosen using the knowledge of the existence of either Uranus nor the black holes, and the ability of those model to predict the presence of those stellar objects is a demonstration of their predictive power.

The field of machine learning has picked up the same method to make sure that their models are predictive. Consider the example of a neural network which predicts the presence of some predefined object in pictures. In this case the neural network parameters are set from a given set of tagged images. Then, as is done for any validation, the predictive power of the trained neural network is tested against *another* set of tagged images. The importance of using a set of images different from the one used for the training of the parameter is easy to explain here. The parameters of the model are chosen such that the model is able to give the tag associated to the images of the training set, hence, if the training set was used to evaluate the predictivity of the model, the model would obviously pass with flying colours, making the validation rather useless.

Hence, to demonstrate how predictive our membrane model is, we first need to choose a set of experiments on which the parameters of our model, E_{vv} , E_{ww} , and E_{vw} , are to be calibrated. And in a second time, we show that our model is capable of predicting the elastic behaviour of the feather vane on an experiment different from the one used for the calibration of the parameters.

In this chapter, we first show in Section 5.1 that the three scenarios discussed in Chapter 3 are not sufficient to calibrate directly the parameters of our membrane model and validate them. In Section 5.2, we introduce an intermediate model which can be validated using those scenarios. In Section 5.3, this intermediate model is used to produce numerically the missing experiments necessary for the validation of our membrane model. The validation protocol

¹Indeed, a formal description of the modelled phenomenon would be required, however, there are no such description of natural phenomenons. One only has access to outcomes of the phenomenon.

Lab experiments

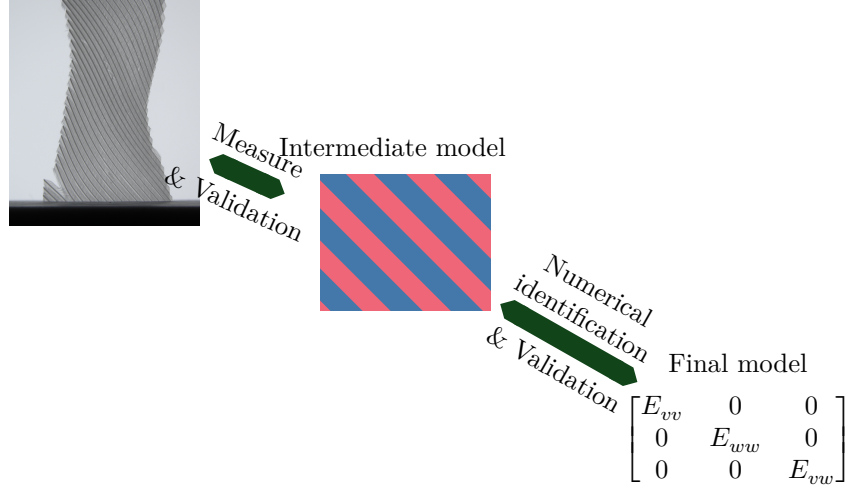


Figure 5.1 – Schematic summarising our validation protocol. We introduce an intermediate model which is validated against our experimental data. Then, this intermediate heterogeneous model is used to produce enough data to validate the final macroscopic model.

we follow is summed up in Figure 5.1.

5.1 Avoiding manual work

To calibrate the parameters of our model on the scenarios described in Chapter 3, we need to analyse how each of those parameters influence the prediction of our model for those scenarios. Then, those parameters are chosen such that the outcome of our model for the measured stress, denoted by $\tilde{\sigma}_{yy}^s$, matches as best as possible the experimental data gathered through execution of those scenarios in the lab.

We first consider the transverse scenario, where the barbs are orthogonal to the displacement. For a measured displacement d , the associated boundary value problem given in Equation (4.2) with our model gives the relationship for the measured stress of

$$\begin{aligned} \tilde{\sigma}_{yy}^s &= \frac{1}{2} E_{wv} \left(\left(\frac{d}{H} + 1 \right)^2 - 1 \right) \\ &= E_{wv} \tilde{\varepsilon}_{yy}, \end{aligned} \quad (5.1)$$

where the quantity marked by a $\tilde{\cdot}$ are measured under the assumption of uniformity instead of being the exact value seen in the sample, as was done in Chapter 3. Here, our model shows a linear relationship between the measured strain $\tilde{\varepsilon}_{yy}$ and the measured stress $\tilde{\sigma}_{yy}^s$, with coefficient E_{wv} . Hence, the parameter E_{wv} can be obtained through a simple linear regression on the experimental data. Equivalent results can also be obtained to measure the longitudinal modulus E_{vv} using the longitudinal scenario.




Figure 5.2 – The mesoscale 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 T and the width of barbs w_B , which allows us to define the barbs surface fraction as $\Phi_B = \frac{w_B}{T}$.

Therefore, we are left with only the shearing scenario, while we need at least two scenarios, one for the calibration of the last uncalibrated parameter, the shearing modulus E_{vw} , and another for the validation of the model.

Since we want to avoid additional lab experiments given their difficulty, we choose not to validate our membrane model directly on real experiments. Instead, we introduce in the next section an intermediary, so-called *mesoscale* numerical model, which is validated on a unique experimental tilted traction test, this mesoscale model is then meant to serve as a reference for the validation of our macroscale model through numerical experiments.

5.2 Mesoscale Model

Introduction of the model

The intermediate model we choose reproduces the structure of the vane at the barb scale, through a heterogeneous shell consisting of stiff bands representing the barbs, connected by soft bands representing the barbules (see Figure 5.2). For simplicity, the material of both bands are chosen to be isotropic linear. While this material choice is arbitrary, the validation of our model will demonstrate its appropriateness. Isotropic linear materials are parametrised by their Young’s modulus, which quantifies the resistance of the material to deformation, and their Poisson’s ratio, which quantifies how their volume changes as they are deformed. We have shown in Section 3.3 that the feather does not feature any Poisson effect: we reflect that in the model by setting the Poisson’s ratio of the material of both bands to zero. The two Young’s modulus, one for each type of bands, are left to be determined through the experimental data.

Due to the lower scale of this model compared to our membrane model, we will call it the *mesoscale model*. Also, to put an emphasis on the difference between the two models, the membrane model will now be referred as the *macroscale model*.

To validate this model, we have to go through the same process discussed in the previous section: we first calibrate the parameters, then validate the model with those calibrated parameters. In the mesoscale model, there are three parameters to calibrate: one geometric parameter emanating from the heterogeneous aspect of the model, the barbs surface fraction Φ_B , and two




Figure 5.3 – Surface coverage of barbs and barbules. Left: the segmented barbs (red) and barbules (green). Right: the original image of the patch region.

mechanical parameters, the Young's modulus of the barbs E_B and the one of the barbules E_b .

Calibration of the parameters of the model

Since the barb surface fraction is a geometric quantity, we avoid the setup of a complex experimental montage and simply measure its value from 50 zoomed images of different regions on a feather vane (Figure 5.3 (right)). Intensity thresholding is performed to segment the barbs from the barbules (green and red, respectively, in Figure 5.3 (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 barbs surface fraction $\Phi_B = 0.14$, with a standard deviation of 0.02.

The two remaining parameters, E_B and E_b , need to be set from the experimental data obtained in both the transverse and longitudinal scenarios.

In the longitudinal scenario, for a measured strain of $\tilde{\epsilon}_{yy}$, the associated boundary value problem given in Equation 4.2 with the mesoscale model gives the predicted measured stress of

$$\tilde{\sigma}_{yy}^s = [\Phi_B E_B + (1 - \Phi_B) E_b] \tilde{\epsilon}_{yy}. \quad (5.2)$$

Hence, denoting by Δ_{\parallel} the slope of the experimental stress-strain curves obtained through a linear regression, we have a first constraint on the mechanical parameters of the mesoscale model,

$$\Delta_{\parallel} = \Phi_B E_B + (1 - \Phi_B) E_b. \quad (5.3)$$

In the transverse scenario, the boundary value problem is harder to work with. Here, the predicted measured stress $\tilde{\sigma}_{yy}^s$ is the solution of a cubic polynomial from which we were initially unable to extract the missing constraint on the two mechanical parameters (see Appendix C.1). Instead, we look at the simpler boundary value problem where the highly stiff bands representing the barbs are considered to be rigid. In this case, the relationship between the measured stress $\tilde{\sigma}_{yy}^r$ and measured strain $\tilde{\epsilon}_{yy}$ is

$$\tilde{\sigma}_{yy}^r = E_b \frac{(\Phi_B - \tilde{\lambda}) \left[(\Phi_B - \tilde{\lambda})^2 - (\Phi_B - 1)^2 \right]}{2\tilde{\lambda} (\Phi_B - 1)^3}, \quad (5.4)$$

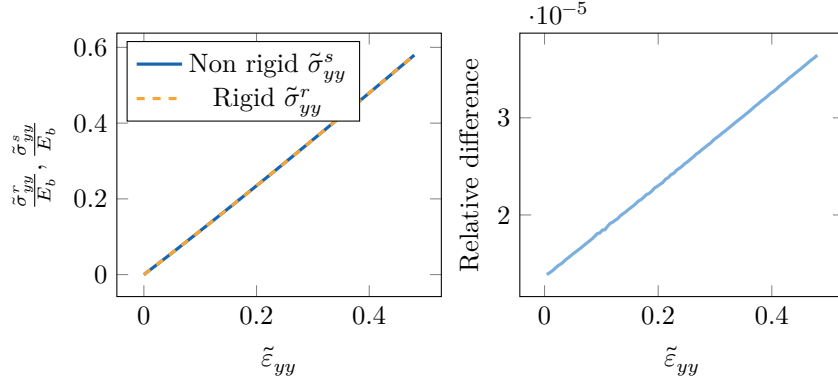


Figure 5.4 – Plot of the stress-strain relationship obtained through an analytical resolution of the transverse scenario with the mesoscale model (left). The plotted stress are nondimensionalised through a division by the Young’s modulus of the barbules E_b . Two different stresses are plotted: the stress obtained with a Young’s modulus for barbs that is four order of magnitude bigger than the Young’s modulus for barbules $E_B = 10^4 E_b$ (—); the stress obtained with the band representing the barbs considered to be rigid (---). The relative difference between the two stresses are considerably small (right).

where $\tilde{\lambda} = \frac{d}{H} = \sqrt{2\tilde{\epsilon}_{yy} + 1}$ is the measured stretching of the sample. As expected, we can see through a numerical comparison that the stress predicted with rigid bands $\tilde{\sigma}_{yy}^r$ is very close to the one predicted with stiff bands $\tilde{\sigma}_{yy}^s$. In fact, their relative difference is in the order of 10^{-5} for a barb surface fraction of $\Phi_B = 0.14$ and four orders of magnitude between the stiffnesses of barbs and barbules, $\frac{E_B}{E_b} = 10^4$ (see Figure 5.4). Moreover, the rigid problem results with an expression that is much easier to work with than the one obtained through the stiff problem.

Formally computing this expression through the computer algebra system SymPy (Meurer et al. 2017), we obtain the predicted slope Δ_{\perp} of the linear regression of the stress-strain relation in the transverse scenario as

$$\Delta_{\perp} = E_b \left[\frac{-15\Phi_B^2\lambda_{\max} - 25\Phi_B^2 + 18\Phi_B\lambda_{\max}^2 + 39\Phi_B\lambda_{\max} + 23\Phi_B}{5(\Phi_B - 1)^3(\lambda_{\max} + 1)^3} - \frac{1}{(\Phi_B - 1)^3} \right], \quad (5.5)$$

where λ_{\max} is the maximum stretch considered. With this second constraint on our mechanical parameters, we can calibrate both the Young’s modulus of the barbs E_B and the one of the barbules E_b .

Since our experimental data is made up of more than one stress-strain curve from which to extract a slope, we choose to measure the parameters by matching the slope of the linear regression of *all* experimental data points. See Figure 5.5 (—) for a representation of the linear stress-strain curve on which we fit the parameter. Here, it would have been more informative to consider the uncertainty on the calibrated parameters introduced by the spreading of the experimental stress-strain curve. However, we chose to focus on the average behaviour, since it should give a good preliminary idea of the usefulness of the solver. The study of the impact of this uncertainty is left for future work.

Table 5.1 – Mesoscale model measurements.

Φ_B	E_b [Pa]	E_B [Pa]
0.12	4.04×10^4	2.8×10^9
0.14	3.92×10^4	2.57×10^9
0.16	3.79×10^4	2.25×10^9

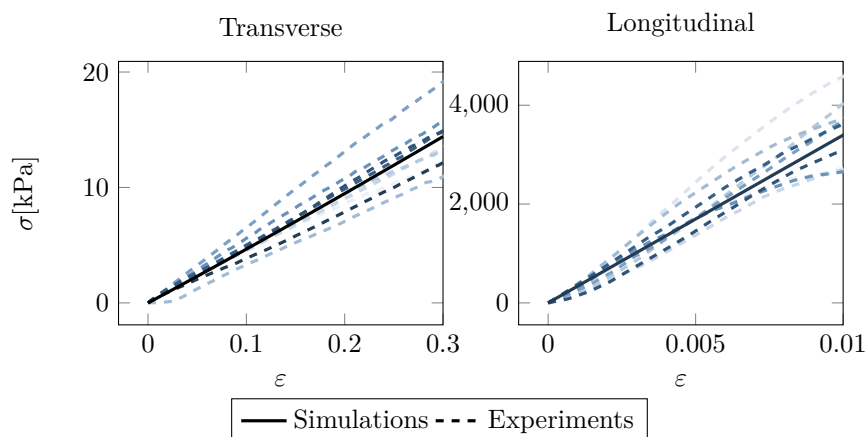


Figure 5.5 – Comparison of the mesoscale model behaviour with the behaviour observed in experiments. Parameters of the mesoscale model were chosen to minimise the distance to every data point of the experiments in both the transverse scenario (left) and the longitudinal scenario (right). The mesoscale model parameters resulting from our fitting procedure are listed in Table 5.1.

To take into account the substantial uncertainty on the barbs surface fraction $\Phi_B = 0.14 \pm 0.02$, we perform this procedure three times, with the surface fraction set to its average value and the values at a standard deviation from the average. The measured parameters are listed in Table 5.1.

Validation of the model

To validate our model, we compare its prediction with the experimental data obtained in the tilted scenario. In this scenario, the aspect ratio of the tested sample, defined as the height over the width $a = \frac{H}{W}$, has a substantial impact on the measured stress. Hence for each sample tested in-lab, we perform one simulation with the same aspect ratio as this sample. First, we focus on the prediction made with the barbs surface fraction set to its measured average value $\Phi_B = 0.14$ which is shown in Figure 5.6. Strikingly, this comparison shows that while the simulation manages to give a pretty good prediction of the measured stress of some samples, it overestimates the stress for most of the samples.

However, this difference between measured and predicted stress does not invalidate our mesoscale model. This difference comes from the presence of non-elastic events in the experiments which our pure elastic model was not built for. Pictures of those experiments show a rearranging of the barbules as

the sample deforms, allowing stress to dissipate and, this way, reaching lower values than the ones predicted (see Figure 5.7). Moreover, when none of these rearrangements can be seen in the pictures, the simulation manages to predict well the behaviour of the sample. From this, we conclude that our mesoscale model properly predicts the *elastic* behaviour of the feather vane which was our main purpose in this thesis.

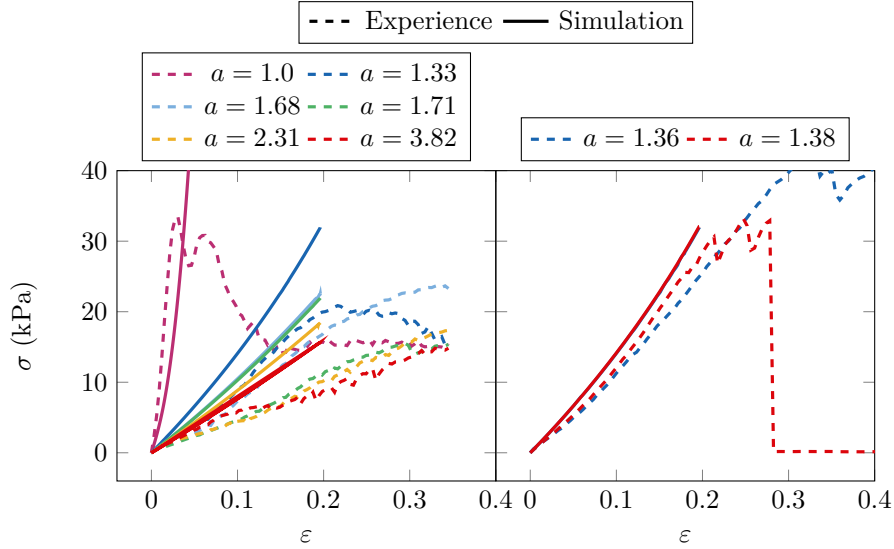


Figure 5.6 – Comparison of the mesoscale model with lab experiments on the tilted traction scenario. In some of those experiments (left), the barbules rearrange themselves (see Figure 5.7) allowing stress to dissipate and preventing our exclusively elastic model to properly predict the observed behaviour. Other experiments (right) do not have any of these barbules rearrangements. In these cases, our model properly predicts the measured stress-strain curves.




Figure 5.7 – Images of two tilted traction tests zoomed on the side of the tested sample. The sample shown in the subfigure 5.7a displays a non-elastic behaviour. In its rest state (left) the edge of the sample is smooth, while after extension (right) the edge has become ragged due to a rearrangement of the barbules. This rearrangement dissipates the bending energy of the barbules, allowing the samples to reach a lower energy state. The subfigure 5.7b presents a sample in which none of those rearrangements take place, as can be seen by the smooth sample edge even after extension.




Figure 5.8 – Comparison of the mesoscale model’s predictions against experimental data obtained on a sample with aspect ratio of 1.38. The simulation ran for the average and extreme measures of the barbs surface fraction. We can see that the predicted behaviour is not much impacted by the large uncertainty. This shows that our mesoscale model should be capable of providing a good estimation of the feather vane elastic behaviour for all potential barb surface fraction.

Influence of the barb surface fraction uncertainty

Until now, we have only looked at the average behaviour of the simulator for a barbs surface fraction of $\Phi_B = 0.14$. In this case, we have shown our mesoscale model to properly predict the elastic behaviour of the feather vane. However, due to the uncertainty on the barb surface fraction measurements, the model calibrated with a more precise measure might not be able to predict the observed behaviour. Thankfully, focusing on the sample with aspect ratio 1.38 we can see that this uncertainty has little impact on the predicted stress (see Figure 5.8).

This reinforces our claim that the mesoscale model is capable to estimate quantitatively the elastic behaviour of the feather vane. The modelling of non-elastic behaviour is left for future work.

5.3 Membrane model

With the mesoscale model validated against the experimental data, we can proceed to the validation of the macroscale model against it. Here, the calibration of the parameters of the macroscale model on the mesoscale model is a well-known process called *homogenisation*. This process consists in computing the parameters of a homogeneous model from a heterogeneous model such that the homogeneous model reproduces the average behaviour of the heterogeneous model.

In addition to the calibration problem being well-known, the mesoscale model itself has a lot in common with models for fibre reinforced matrices.




Figure 5.9 – Schematic of a fibre reinforced material taken from Maiti et al. (2022). Such materials are made of stiff fibre embedded in a suppler matrix. Fibre reinforced materials are quite close to our mesoscale model. In fact, taking a section of the material shown here would result exactly in our mesoscale model, i.e. in a material with bands of alternating properties.

The bands representing the barbules can be seen as the medium that is reinforced and the band representing the barbs as the fibres that are reinforcing this medium (see Figure 5.9). This type of material being crucial in the industry — think of reinforced concrete for example, their modelling, and most especially modelling through homogenisation, has been extensively studied (see for instance the introduction of Bleyer (2018) for a state of the art of these studies).

Unfortunately, due to brittleness of most reinforced material — meaning that they have a high stiffness compared to their resistance to fracture — studies have focused on the linear regime of these materials and never ventured in the study of large deformations. Hence, their work on homogenisation of the linear model is difficult to apply directly to our non-linear model.

Instead of performing the calibration through an analytical method, we proceed through a numerical method. As already shown in the first section of this chapter (see Section 5.1), the transverse modulus E_{ww} and the longitudinal modulus E_{vv} can be obtained from a linear regression of the experimental stress-strain curve produced respectively in the transverse and longitudinal scenarios. This only leaves the need to calibrate the shearing modulus E_{vw} .

We measure this parameter through bisection, performing macroscale model simulations with different values of E_{vw} to match the slope with that of the mesoscale model in the tilted scenario. As stated in the previous section, the aspect ratio has an influence on the predicted stress in the tilted scenario: here, the calibration of the shearing modulus was done with an aspect ratio fixed to 1.3. The rest of the validation protocol is done with the same fixed aspect ratio. The calibrated parameters are given in Table 5.2.

In Figure 5.11, we show that the macroscale model, whose parameters have been fit from the mesoscale model on the longitudinal, transverse and tilted traction tests, matches the same mesoscale model on traction scenarios with barbs oriented at various angles (see Figure 5.10). With this, our macroscale model is validated against the mesoscale model which was itself validated against experimental data. Hence, we claim that our macroscale model is capable of predicting the experimental elastic behaviour of the feather vane.

In addition to this quantitative validation, we also visually compare, on the

Table 5.2 – Macroscale model measurements. The macroscale measurements were obtained by fitting to the mesoscale model parameters.

Φ_B	E_{vv} [Pa]	E_{ww} [Pa]	E_{vw} [Pa]
0.12	3.60×10^8	4.80×10^4	2.65×10^4
0.14	3.60×10^8	4.80×10^4	2.78×10^4
0.16	3.60×10^8	4.80×10^4	2.88×10^4




Figure 5.10 – The mesoscale and macroscale model are compared on the tilted scenario with different barb orientation angles θ .




Figure 5.11 – Comparison between slopes of the stress-strain curves of the macroscale model and the mesoscale model in traction scenarios with barbs oriented at different angles. The slopes are nondimensionalised through scaling by the Young modulus of the barbules in the mesoscale model E_b . Comparison is done using the measured barb density $\Phi_B = 0.14$, as well as the two extreme values of the error range ± 0.02 .




Figure 5.12 – Tilted scenario performed on samples with different aspect ratios. This was performed experimentally (top), with the mesoscale model (middle) and with the macroscale model (bottom). For a given aspect ratio, each test displays shear bands of the same size and similar overall geometry.

tilted scenarios, the geometry of our macroscale model, mesoscale model, and real experiment (see Figure 5.12). 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.

5.4 Discussion

In this chapter, we have presented a parameter identification and validation protocol for our membrane model. This validation showed that our membrane model is able to properly predict the real elastic behaviour of the feather vane.

Two step protocol

The parameter identification protocol we suggest has some remote similarity with the work of Sperl, Sánchez-Banderas, et al. (2022) who are estimating parameter for a yarn level model of cloth from cloth level experiments. To do so, they first identify the parameter of a shell model from traction test on cloths samples. Then, they perform an inverse homogenisation procedure on the resulting calibrated shell model to obtain parameters for the yarn level model. Here, there are two main difference with our protocol. One is in the

process, where, instead of using a low scale calibrated model to calibrate a high scale model, they do the inverse: using a high scale calibrated model to calibrate a low scale model. The other is in the purpose of this two stage protocol. Their aim is to avoid the repeated simulations of the complex yarn level required for a straightforward fitting of the parameters. Both step of their protocol are more efficient than the naive approach as the shell model requires fewer degrees of freedom and the link between the shell model and the yarn level model can be done through homogenisation which also requires fewer degrees of freedom than a full simulation of a traction test. In comparison, our aim is to reduce the number of experiments that needs to be setup.

While we managed to reach this aim of reducing the number of experiments needed for calibration of our membrane model, we didn't really reduce the amount of work necessary, as we initially intended. In fact, the avoided effort necessary for the setup of more experiments was merely replaced by effort spent in the numerical modelling of the mesoscale model. Overall, the time spent on the two stage protocol probably is much greater than if we had performed the additional experiments, mostly due to workload induced by the implementation and debugging of the mesoscale model.

However, we believe that the introduced mesoscale model usefulness is not limited to our parameter identification protocol. First, the model being closer to the underlying structure of the feather vane, it is a first step toward understanding the link between this structure and the overall behaviour of the vane described by our membrane model. However, the path toward establishing such a link would probably be quite long (see Chapter 7). Secondly, the mesoscale model can be used to study other properties of the feather such as size-effect discussed in the conclusion of the thesis.

Compression

An interesting aspect of the validation protocol of the mesoscale model is that parameter identification didn't consider at all the behaviour in compression of the vane. Indeed, both the transverse and longitudinal scenarios which were used for the calibration exclusively induce extension within the material. Despite this restriction to non-compressive behaviour in the parameter identification, the mesoscale has no issue matching the elastic behaviour of the vane in the tilted scenario which is driven by compression of the barbules. This highlights that our model can capture a wide range of the vane membrane behaviour despite having so few parameters.

Lack of statistical method

A limitation of our protocol is its lack of rigour. Validation is made through a simple eyeball comparison and, while we roughly take into account the uncertainty on the measure of the barb surface fraction Φ_B , the impact of a lot of the uncertainty of other measures are not properly looked into. Even more, the flagrant variation between the stress-strain curve induced by the biological aspect of the studied material is only brushed upon. A careful statistical study of all those parameters could have provided an understanding of the qualities and limitation of our model. It might be fine to leave aside the uncertainty due to measures in the experimental protocol due to their low impact on the

model parameters. However, the impact of the biologically induced variations needs to be tackled in future works. Especially to show that our model is valid for most feathers instead of just the average feather.

To the best of our knowledge, such statistical studies are often not considered in the validation of models in both mechanics and computer graphics. This is somewhat understandable as, in mechanics, the materials are produced in very control environment, leaving little place for impactful variations, and in computer graphics, the appearance of the phenomenon is often the only aspect of interest, hence, slight variations would often go unnoticed.

Other fields dealing with populations, such as social study and biology, that are much more subject to internal variation, are likely more used to using statistical tool to analyse their hypothesis and model. And, hence, more prone to use those tool when tackling any problem.

This is why, coming from both mechanics and computer graphics, it was not natural to consider the use of statistical methods for the study of our material. As said earlier, this will need to be corrected in future works studying feathers.

Chapter 6

Model of a full feather

In this chapter, we present a preliminary modelling of the elasticity of a full feather which aims to be as visually similar to the behaviour of a real feather as possible. We believe that three ingredients are crucial to properly reproduce the movement of a feather:

- The modelling of the rachis and its interaction with the vane;
- A non-uniformity of the distribution of barbs orientations;
- The modelling of the flexural behaviour of the vane.

6.1 Rachis

Consider the thin rod geometry of the rachis, it would seem appropriate to model it through a rod model which would be coupled with the vanes modelled as two shells. We rather model the whole feather as one shell where the rachis is represented by a thin band of stiff linear isotropic band (see Figure 6.1). While this is probably less efficient than the first solution, it is easier to deal with, as our software ARCSim is capable to deal with heterogeneous shells but not thin rods.

In our workflow, the overall shape of the feather is given by a triangular mesh which describes its reference state. The rachis is given by its center-line described as a spline as well as two scalars, its width at the base and its width




Figure 6.1 – The rachis and the vane are represented using the same shell mesh. The rachis is modelled by a band of stiff isotropic material, in orange here.

at the end. Those scalar are usually chosen such that the rachis is wider at its base than at its end, as is the case in real feathers.

In the present work, we choose to keep the thickness of the rachis uniform across its representing band. This is a considerable approximation as the thickness of the rachis substantially changes from its base to its end. However, this approximation seems to be enough for the deformations considered during the thesis (see Section 6.6). The variation of the thickness might have to be taken into account in future works if more involved deformations or more precise prediction are required.

6.2 Non-uniformity of barb orientation

Modelling

The distribution of barb orientations is modelled by an angle field over the reference state of the feather $\theta : \Omega \rightarrow [0, \pi]$, which gives the angle $\theta(\mathbf{r})$ between the barb orientation and an arbitrarily chosen x-axis at point \mathbf{r} . The membrane model is parametrised by this vector field such that the transverse modulus properly penalises deformation orthogonal to the barb orientation. More formally, the membrane energy density $\psi_m(\mathbf{r})$ at material point \mathbf{r} is given by

$$\psi_m(\mathbf{r}) = b\varepsilon(\mathbf{r}) : \mathbf{E}(\theta(\mathbf{r})) : \varepsilon(\mathbf{r}), \quad (6.1)$$

where $\mathbf{E}(\theta(\mathbf{r}))$ is our membrane model tensor presented in Section 5.3 whose basis is changed by a rotation of angle $-\theta(\mathbf{r})$ (see Section A.2 for its detailed expression).

In the context of the finite element method, we chose to discretise the barbs angle by a piecewise constant field. We could not consider higher order discretisation as it would have prevented the implementation of the solutions to the different numerical issues discussed in Section 4.3. Moreover, this saves us the time-consuming implementation of higher order elements in ARCSim.

Meshing

The meshing algorithm consists in generating the edges we want to be present in the mesh and letting a triangulation software (triangle from Shewchuk (1996) here) produce a proper triangulation which contains those edges and vertices. Hence, we only describe the generation of those vertices and edges here. The description heavily relies on the schematics shown in Figure 6.2.

The input of the meshing is provided by the user. It consists in the boundaries of the feather as a mesh (—), the center-line of the rachis as a spline (—) and guides indicating the overall orientation of barbs as splines (—). The meshing algorithm takes also three input parameters, the width at the base w_b and the end w_e of the rachis, as well as the element target edge size e .

The meshing algorithm starts by generating the vertices (\bullet) of the rachis, placing them next to its center-line at a distance $\frac{w_b}{2}$ at the base and $\frac{w_e}{2}$ at the end of the rachis, decreasing the distance at a constant rate between the two tips (see Figure 6.2b). Vertices that do not fall into the provided boundary of the feather are removed (\circ). Edges (—) are placed between the kept vertices (see Figure 6.2c). Edges are also added between the boundary of the feather

and the first and last vertices, making sure that the region representing the rachis is separated from the vane.

Then, starting at each rachis vertex, we add the vertices of the vane. To do so, the two closest guides are interpolated to produce a new guide passing by the considered rachis vertex (see Figure 6.2d). The vane vertices are placed along this guide such that the distance between them respects the target edge length e . Vertices falling out of the boundary of the feather are naturally not added. Finally, vertices are linked by edges and the last vertex is linked to a vertex of the boundary.

Initialising barb orientation

We want each element to have at least one edge aligned with the barb orientation to mitigate numerical locking issues. Thankfully, the elements generated by the triangulation software should be incident to at least one of the edges whose presence was enforced. Since those edges were generated to be as aligned as possible with the user provided barb orientation, we simply assign the angle each of those edge makes with the x-axis to its incident faces (see Figure 6.3).

6.3 Flexion modelling

As we chose to focus on the membrane model during the course of the thesis, we leave the design of a flexural model from experiments for future works. We instead use an ad hoc anisotropic model linear in the curvature. Bending forces are derived from an energy associated to each edge,

$$\frac{A_1 + A_2}{6} D(\beta) (\eta - \eta_0)^2, \quad (6.2)$$

where A_1 and A_2 are the area of the triangles incident to the edge, η and η_0 are respectively the curvature and rest curvature at the edge, and $D(\beta)$ is the bending modulus. The curvature is computed through the method described in Grinspun et al. (2003)

$$\eta = 9 \frac{e\chi}{A_1 + A_2}, \quad (6.3)$$

where e is the size of the edge and χ is the dihedral angle between the triangle that join at the edge. 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 parametrised by the longitudinal bending modulus D_{\parallel} (bending the barbs themselves) and the transverse bending modulus D_{\perp} (bending the vane keeping the barbs straight) (see Figure 6.4) The relationship between the bending modulus and the tilt angle is chosen such that $D(0) = D_{\perp}$ and $D(\frac{\pi}{2}) = D_{\parallel}$:

$$D(\beta) = \frac{D_{\perp} - D_{\parallel}}{2} \cos(2\beta) + \frac{D_{\perp} + D_{\parallel}}{2}. \quad (6.4)$$

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

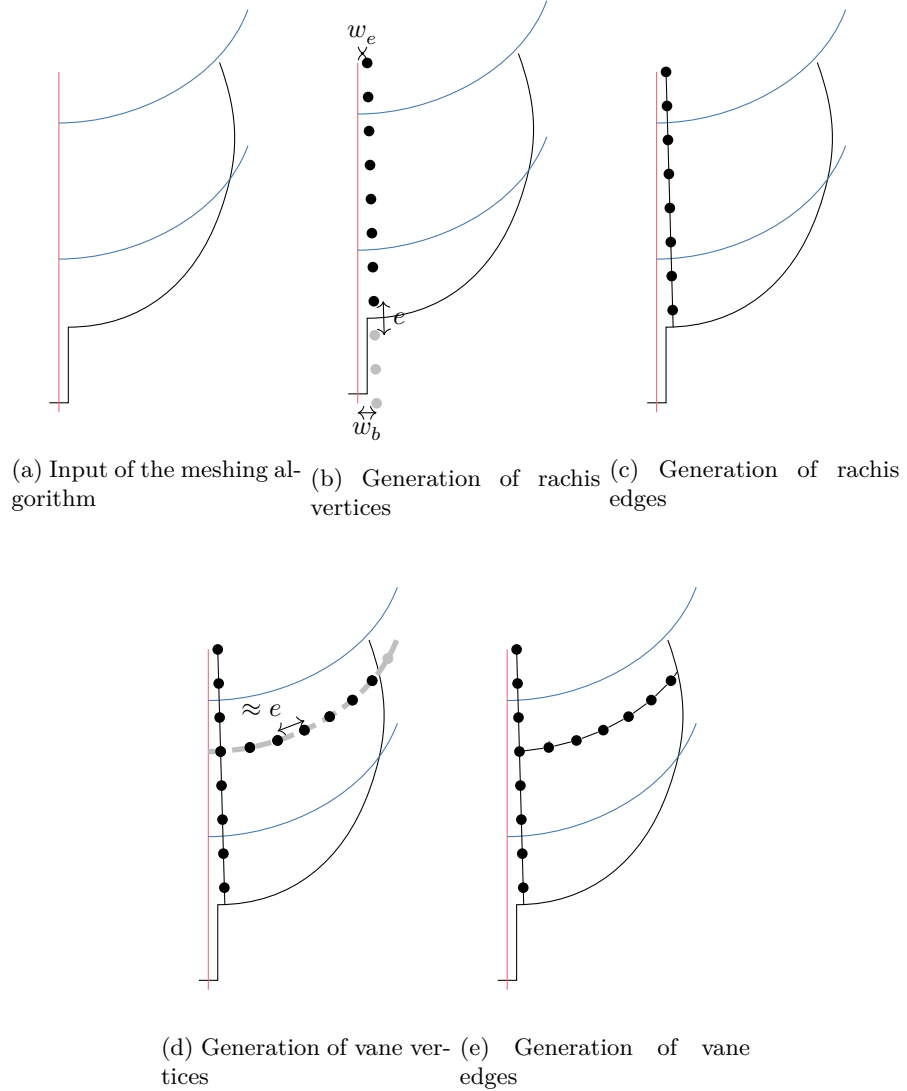


Figure 6.2 – Schematics displaying the different steps of our meshing algorithm. Here, only the left side of the feather is shown, barb orientation is given separately for the left and right side. Unlike the generation of the rachis, no symmetry with respect to the center-line is assumed when generating the vertices and edges of the vanes. These schematics are described within the main text.

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 through an area based average of both angles,

$$\beta = \frac{A_1\beta_1 + A_2\beta_2}{A_1 + A_2}. \quad (6.5)$$




Figure 6.3 – Schematics of the initialisation of the barb angle field θ . The barb angle field is chosen to be discretised as a piecewise constant field. Each element incident to edges that were placed along the barb orientation during the meshing (—) are given as barb angle $\theta(\mathbf{r})$ the angle of this incident edge.




Figure 6.4 – Schematics of two different bending direction. Here, the barbs are represented as black lines. One bending orientation is transverse bending (left) and is associated to the transverse bending modulus D_{\perp} , the other is longitudinal bending (right) and is associated to the longitudinal bending modulus D_{\parallel} .

6.4 Parameter calibration

Wanting to show that our model performs better than the isotropic cloth model used in the industry (Bowline and Kačić-Alesić 2011; Weber and Gornowicz 2009; Augello et al. 2019; Heckenberg et al. 2011; Haapaoja and Genzwürker 2019), we qualitatively compare the visuals produced by both our full feather model and a shell model close to the one used in the industry with a video taken from a real feather. To make the comparison as fair as possible, we do our best to calibrate the parameters of both models on a scenario separate from the one used for the comparison. The present section deals with this calibration. The final parameters used in the comparison presented in the next section are listed in Table 6.1.

Since the validation of our model is done through a qualitative visual comparison, we believe that it is sufficient to calibrate the parameters by choosing them manually such that the produced visuals qualitatively match a video of a




Figure 6.5 – Scenario where the rachis of a feather is held at its base and pulled at its tip. This scenario was performed with a real feather (top) and within the simulation using an isotropic membrane model (middle) and our anisotropic membrane model (bottom). This scenario was used to select the mechanical parameters of the rachis model and vane bending model.

real feather. Hence, aside from the membrane parameters which were already calibrated in the previous section, the rachis parameters and bending parameters of the vane are chosen by testing different sets of parameters until the simulation resembles to the naked eye our video of a real feather. The scenario used for the calibration consists in a feather where the base of its rachis is kept still while the tip is pulled away from its rest state (see Figure 6.5).

Our model

As discussed in Section 6.1, the rachis band has three parameters, its thickness, Young's modulus and Poisson's ratio. Its Poisson's ratio is arbitrarily chosen to be 0.49 to be as close as possible to incompressible, whereas we chose its Young's modulus by measuring it ourselves through a simple traction test on the rachis which resulted in the value of 34 MPa. While it would have been natural to directly measure the thickness of the rachis on the real feather, the resulting value was inducing an out of plane buckling of the rachis, which is absent from the video of the real feather (see Figure 6.5 (top)). Instead, we set the thickness to be as close as possible to the measured value, while preventing

buckling in the considered scenario.

For the vane, the two bending moduli of our flexural model need to be calibrated. We noticed that the farther apart they are, the more difficult it is for the solver to converge¹. Hence, we decided to choose the closest possible values such that buckling of the vane on the side away from the movement is prevented (this is the top side in the images shown in Figure 6.6) and wrinkling looks as close as possible to the one in the real feather. Finally, the target edge length in the meshing algorithm was chosen as high as possible to obtain good performances, while keeping the produced visual close to the real feather. The resulting visuals for the membrane model are shown in Figure 6.5.




Figure 6.6 – Using too small a longitudinal bending modulus can cause the vane to buckle during the deformation (left). Using too high a transverse bending modulus reduces the wave length of the wrinkles, making the deformation unrealistic (right).

Isotropic model

The simulation with which we compare our approach is exactly the same as our model with the exception of the membrane model being isotropic, which means that the stiffness modulus are all equals $E_{vv} = E_{vw} = E_{ww}$. Thanks to this, differences appearing between the two models should come from the difference between the membrane model as it is only difference between the two simulations. This allows us to highlight the contribution of our membrane model without being hindered by the flexural and rachis model.

For the purpose of fairness, the parameters of this model were chosen independently of the parameters of our proposed feather model and such that the visuals produced are also as close as possible to the video of the real feather. While this results in the same parameters for the rachis, the parameters of the vane are naturally different. First, the Poisson's ratio of the vane is set to 0 to obtain the same effect found in the real feather. Then the Young's modulus of the vane is chosen small enough for the vane to not buckle during the flexion, but as high as possible. With the Young modulus set, the bending parameters are picked using the same process used for the calibration of the bending parameter of our anisotropic membrane model. As for our membrane model, results produced with the isotropic model are displayed in Figure 6.5.

¹We believe that this slow convergence comes from the high ratio between the two bending modulus inducing a poor condition number for the hessian of the problem.

Parameter	Value
Feather length	0.2 m
Tolerance	10^{-7}
Number of Newton steps	50
Target element edge length	5×10^{-3} m
Isotropic longitudinal bending modulus D_{\parallel}	5.0×10^{-3} N m
Isotropic transversal bending modulus D_{\perp}	2.5×10^{-8} N m
Isotropic Young's modulus	4.80×10^4 Pa
Anisotropic longitudinal bending modulus D_{\parallel}	2.5×10^{-4} N m
Anisotropic transverse bending modulus D_{\perp}	10^{-7} N m
Barrier coefficient κ (see Section 6.5)	5.04×10^7 N m $^{-1}$
Surface mass density	2.6×10^{-2} kg m $^{-2}$
Rachis Young's modulus	34×10^6 Pa
Rachis thickness	1.2×10^{-2} m
Rachis Poisson's ratio	0.49
Rachis surface mass density	0.54 kg m $^{-2}$
Time-step	0.011 s

Table 6.1 – Parameters used for generating the figures involving the full feather in the main document. Parameters whose value can be inferred from the main document are not listed. If a notation was given in the main document for a parameter, we recall this notation in the table.




Figure 6.7 – When compression reaches a certain point (left), elements whose elasticity follows the St. Venant-Kirchhoff model collapse to allow other elements to stretch (right).

6.5 Energy penalty for bending

Unfortunately, the interaction between the bending and membrane parameter chosen in the previous section results in some instabilities in our simulation of the full feather when the vane is under large compressive strain (see Figure 6.7). This unphysical behaviour has already been studied in the context of a model close to ours, the St. Venant-Kirchhoff three-dimensional elasticity model (Sautter et al. 2022).

This non-physical behaviour under large strain is usually not an issue when working with shells thanks to the fact that buckling appears before those prob-

lematic strains are reached. However, the high longitudinal bending modulus chosen to prevent the buckling of the vane in the previous section also unphysically prevents buckling under large compression. To enforce buckling under large compression, we use the model of Kikuuwe et al. (2009) designed to prevent compression to zero volume.

$$\frac{1}{2}b\varepsilon : \mathbf{E}(\theta) : \varepsilon + b\frac{\kappa}{12}\left(\frac{1-J}{6}\right)^3, \quad (6.6)$$

where J is the determinant of the deformation gradient, and κ is a parameter of the model. This model adds the right term to the St. Venant-Kirchhoff model only under compression, it removes the instabilities and allows buckling if its parameter κ is chosen sufficiently high (see Figure 6.8).

The use of a log barrier instead of the cubic term has been considered. Such barrier diverges to infinity as λ approaches 0: this enforces buckling for any parameter chosen. However, the nonlinear solver becomes unable to converge on the time-stepping problem, hence our choice for the gentler cubic term.

6.6 Qualitative Validation

The scenario we chose for the purpose of the visual qualitative validation is the same as the one used for the parameter identification protocol aside that pulling is not exerted from the side of the rachis but from the side of the vane (see Figure 6.9). This scenario was performed with a real feather, our calibrated feather model and the calibrated isotropic feather model. Our model manages to reproduce the bending of the rachis and the wrinkles in the vane, while the isotropic model is unable to reproduce either and deforms unrealistically. The unrealistic behaviour of the isotropic model is due to its softness. This behaviour could be prevented by using a higher Young's modulus, but, as mentioned in the previous section, this would result in an unrealistic buckling of the vane. This shows the importance of the modelling anisotropic aspect of the vane for applications where the appearance of the deformation is critical.




Figure 6.8 – Total internal energy Ψ of a one dimensional truss of length L and thickness h . The endpoints of the truss are moved away and brought close to each other (bottom left). The distance between the endpoints relative to the distance at rest is denoted by λ . The displayed curves (top) are the internal total energy when no buckling occurs (bottom right top) for both the St. Venant-Kirchhoff model (—) $\Psi = \frac{1}{2}LhE(\lambda^2 - 1)^2$ and the Kikuuwe model (—) $\Psi = \frac{1}{2}LhE(\lambda^2 - 1)^2 + Lh\frac{\kappa}{12}\left(\frac{1-\lambda}{6}\right)^3$, as well as the bending energy for an isometric deformation (bottom right bottom) (—) $\Psi = \frac{1}{2}\int_0^L D\eta^2 ds$ where D is the bending modulus and η is the curvature. This last curve was computed through ARCSim’s finite element simulation. The coefficient of Kikuuwe energy was chosen such that $\kappa = 3 \times 10^3 E$ and the bending modulus was chosen such that $D = \frac{1}{3}hL^2E$. When the membrane behaviour is modeled with the St Venant-Kirchhoff and the bending modulus is large, buckling does not occur under compression as it would not reduce the internal energy of the truss. In contrast, with the Kikuuwe model, a sufficiently large coefficient κ induces a membrane energy that is higher than the bending energy under compression, allowing buckling to emerge.




Figure 6.9 – 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.

Chapter 7

Conclusion

Through this thesis, we presented an elastic membrane model for the feather vane which was validated against experiments. This model was numerically implemented using linear triangular finite elements, which were aligned with the stiff barbs' direction to avoid locking. Moreover, we qualitatively showed that our membrane model could be used to model full feathers for the purpose of animation.

Link to the underlying structure

Knowing how the shape of the barbules and barbs impacts the phenomenon emerging at the vane scale, quantified by our membrane model, would be especially useful for engineering purposes. This would allow one to understand what is important to produce a material that would have some of the useful mechanics properties of feathers.

For most materials, this link can be established through the derivation of the parameters of a high scale material model in function of parameters of a lower scale model through either homogenisation or dimensionality reduction. A good example is the one dimensional Kirchhoff thin rod model. The link between the bending parameters appearing in this model and the lower scale geometry is well known thanks to its derivation through dimensionality reduction: the bending parameter D scales with the fourth power of the rod's radius r , i.e. $D = Cr^4$, where C is a constant. Those types of relation are especially important when one wants to design some structures with specific properties.

Unfortunately, our phenomenological modelling of the feather, i.e. based on observations rather than through derivation from existing model, prevents us to obtain an understanding on how the different aspects of the microstructure impact the behaviour at the vane scale.

To establish the link between the structure of the feather and its behaviour, one would first need to understand the behaviour of the structure in details through proper experimental observations. Unfortunately, observing this scale of the feather is hindered by the micrometer size of the barbules which would require precise equipment to analyse. Despite this apparent difficulty, F. Zhang et al. (2018a) have qualitatively investigated the mechanism of barbules' locking, laying this way the path to a barbules scale model that could then be linked to our higher scale model.

However, even if a barbule scale model of the feather vane were to be designed, deriving the link between this lower scale model and the overall behaviour of the vane would be difficult. In fact, consider the more than half a century old Drucker-Prager yield criterion (Drucker and Prager 1952) used to model the behaviour of granular medium when modelled as a continuum. While the behaviour of such granular a medium is straightforward to model at lower scale through rigid body mechanics and friction, the parameters of the criterion have yet to be linked to the shape and friction coefficient of the underlying grains. Hence, it is not clear how much work would be required to establish a connection between the barbules scale and the vane scale for the feather vane.

Size effect

In our model, deformation is measured by the Green-Lagrange strain which is dimensionless. This means that the behaviour of a considered system doesn't depend on its overall size. However, this is only true when the underlying structure is sufficiently small compared to the overall size of the system. Otherwise, the size of the heterogeneity at the lower scale has an impact on the overall behaviour (Bažant 2000).

We observed this phenomenon, called *size effect*, through our mesoscale model which emerged due to the bending of barbs (see Appendix A.3). This is unsurprising as this was observed in fibre reinforced material whose structure and behaviour are quite similar to the feather vane (Bleyer 2018). The impact of this phenomenon is probably small in the context of the deformation we considered in this work, since we managed to predict quite well the elastic behaviour as shown by our validation protocol. Taking the size effect into consideration might result in a small gain in precision for those deformations. But for other deformations, properly modelling their effect through higher order theory of elasticity, such as, among other, strain-gradient and coupled stress might be crucial (Ghiba et al. 2017; Lam et al. 2003). This phenomenon will need to be studied at some point to make advances on the modelling of feathers, especially if mechanical applications are to be considered.

Perspective

As mentioned in the introduction of this thesis, there is currently not much work on the mechanics of feather, hence, there is still a lot to work on to understand the details of the mechanical behaviour of the feather and design an appropriate model for it. Naturally, every phenomenon we have left aside during the course of the thesis would need to be investigated. This means looking into the non-elastic behaviour of the feather mentioned in Section 5.2 which is quite important considering its impact on measured stresses in our traction scenario. This could probably be modelled using a simple plasticity model. However, like our elastic model, its numerical implementation might encounter some difficulty due to the strong anisotropy. Obviously, flexion, fracture, and healing of the feather vane will have to be tackled.

As mentioned in Section 5.4, our analysis of the elastic membrane behaviour of the feather seriously lacks a consideration for the observed variance in the stiffness parameter. A proper study of its impact is of importance if we want

to properly support the claim that our model is capable to model the elastic behaviour of any feather.

Moreover, it would be interesting to advance the knowledge of size-effect within the feather as well as on the establishment of a link between the geometry of the structure and the macroscale parameters as mentioned in the two previous subsections.

Appendix A

Elasticity Theory

A.1 Poisson's ratio

In this section, we show that for an orthotropic material modelled by a linear stress-strain relation

$$\begin{bmatrix} \sigma_{vv} \\ \sigma_{ww} \\ \sigma_{vw} \end{bmatrix} = \begin{bmatrix} E_{11} & E_{12} & 0 \\ E_{12} & E_{22} & 0 \\ 0 & 0 & E_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_{vv} \\ \varepsilon_{ww} \\ 2\varepsilon_{vw} \end{bmatrix}, \quad (\text{A.1})$$

the Poisson's ratio ν_w along the w-axis is given by

$$\nu_w = \frac{E_{12}}{E_{11}}. \quad (\text{A.2})$$

To do so, we cannot consider a scenario different from the one considered in our experiments (see Section 3.1). In such scenarios, an orthotropic material doesn't have a uniform strain due to the constraint on the displacement along the v-axis (see Figure A.1) This means that there is an ambiguity on the strains ε_{vv} and ε_{ww} present in the definition of the Poisson's ratio $\nu_w = -\frac{\varepsilon_{vv}}{\varepsilon_{ww}}$.

Instead, we consider a scenario in which the material is stretched along its w-axis while its deformation along the v-axis is nowhere constrained. This




Figure A.1 – Images of the initial state (left) and last state (right) of a traction test where the sides of the tested sample are completely constrained. In such a traction test, if the tested sample's material is subject to a Poisson's effect, the resulting deformation will be non-uniform.



Figure A.2 – Example of an experimental setup where the displacement transverse to the direction of elongation is not constrained. Here, knitted patch is attached to a bar in a way that allow loops to move along the bar. This image was taken from Poincloux (2018).

scenario can be modelled by a boundary value problem quite similar to the one shown in Section 4.2

$$\begin{aligned}
 -\nabla \cdot \mathbf{F}\boldsymbol{\sigma}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Omega; \\
 u_w(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Gamma_b; \\
 u_w(\mathbf{r}) &= d\hat{\mathbf{y}}, \forall \mathbf{r} \in \Gamma_t; \\
 \mathbf{F}\boldsymbol{\sigma}(\mathbf{r})\mathbf{n}(\mathbf{r}) &= \mathbf{0}, \forall \mathbf{r} \in \Gamma_s,
 \end{aligned} \tag{A.3}$$

where, rather of constraining the whole displacement, only the displacement along w-axis is constrained.

Thanks to the relaxation of the constraint on the displacement, the solution to this problem is uniform. Meaning that there is no ambiguity when we talk about either the strain along the v-axis ε_{vv} and the one along the w-axis ε_{ww} .

Due to the absence of constraint on the side of the material, we have

$$\mathbf{F}\boldsymbol{\sigma}\hat{\mathbf{v}} = 0 \tag{A.4}$$

where the normal \mathbf{n} as been expressed in the referential basis.

Finally, from this we derive

$$\begin{aligned}
 \mathbf{F}\boldsymbol{\sigma}\hat{\mathbf{v}} &= 0 \\
 \implies \boldsymbol{\sigma}\hat{\mathbf{v}} &= 0 \\
 \implies \sigma_{vv} &= 0 \\
 \implies E_{11}\varepsilon_{vv} + E_{12}\varepsilon_{ww} &= 0 \\
 \implies \nu_w &= -\frac{\varepsilon_{vv}}{\varepsilon_{ww}} = \frac{E_{11}}{E_{12}}.
 \end{aligned}$$

A.2 Elasticity tensor change of basis

As explained in Section 5.3, the diagonal shape of the elasticity tensor only appear for an appropriate choice of basis, namely, the orthonormal basis formed of the axis aligned with the barbs. Another choice of the basis would result in a dense elasticity tensor. This can be seen in the following formulas which gives the values of the components of the elasticity tensor for a basis rotated

by an angle of θ from our suggested one.

$$\begin{aligned}
E_{11} &= E_{vv} \cos^4(\theta) + \frac{E_{vw} \sin^2(2\theta)}{2} + E_{ww} \sin^4(\theta) \\
E_{22} &= E_{vv} \sin^4(\theta) + \frac{E_{vw} \sin^2(2\theta)}{2} + E_{ww} \cos^4(\theta) \\
E_{33} &= \frac{E_{vv} \sin^2(2\theta)}{4} + \frac{E_{vw} \cos^2(2\theta)}{2} + \frac{E_{ww} \sin^2(2\theta)}{4} \\
E_{12} &= (E_{vv} - 2E_{vw} + E_{ww}) \sin^2(\theta) \cos^2(\theta) \\
E_{13} &= E_{vv} \sin(\theta) \cos^3(\theta) - \frac{E_{vw} \sin(4\theta)}{4} - E_{ww} \sin^3(\theta) \cos(\theta) \\
E_{23} &= E_{vv} \sin^3(\theta) \cos(\theta) + \frac{E_{vw} \sin(4\theta)}{4} - E_{ww} \sin(\theta) \cos^3(\theta)
\end{aligned} \tag{A.5}$$

Notice that, since the matrix is dense for another basis, the Poisson's ratio can be nonzero. This explains the presence of a Poisson's effect in the transverse scenario.

A.3 Size effect

As explained in Chapter 7, due to the absence of dimensions in the Green-Lagrange strain, the model is agnostic to the size of the considered system. Meaning that two systems, one where every distance is twice as big as in the other, would behave exactly the same. However, this is usually not true for real systems of small enough size. The heterogeneity of the underlying structure will impact the overall behaviour when the scale of the system gets near the scale of the heterogeneity, this phenomenon is commonly called a *size effect* (Bažant 2000). This phenomenon is also present in the feather vane.

To display this size effect we will consider the longitudinal scenario, but, instead extension, we will be looking at the behaviour of a sample under compression (see Figure A.3). Unfortunately, we didn't try this scenario in the laboratory. However, the impact of the size effect can still be looked into through the mesoscale model in which the heterogeneity is present (see Figure A.4). This simulation is run for different ratio of sample size over inter barb distance $N = \frac{W}{T}$. The results are shown in Figure A.5.

Here, it is clear that the behaviour is dependent on the size of the sample: the smaller the size, the higher the stiffness, hence, we are clearly in the presence of a size effect. The origin of the difference of behaviour between the different size is due to the barb bands being in flexion rather than in compression during the scenario. To understand this, it is easier to consider an equivalent setup where the size of the sample is fixed, but the width of the bands is changed. Since the bending modulus D of a band scales with the cube of its width $D = Cw_B^3$, where C is a constant, the forces at the end of each band also has a cubic dependence on the width of the bands. Then, when normalising by the width of the band to obtain the measured stress, we obtain a quadratic dependence on the size of the bands. This dependence is shown in Figure A.6. For a large enough sample, the impact of the bending of the barbs will become negligible which will make the size effect unnoticeable.




Figure A.3 – Schematics of the first (left) and last (right) step of scenario considered studying the size effect phenomenon within the vane of feathers. The scenario is the same as the longitudinal scenario, but in compression instead of extension.




Figure A.4 – Images of the first (left) and last (right) step of a simulation of the scenario used to display the size effect using the mesoscale model. Here the ratio between sample size over inter barb distance is $N = 25$.




Figure A.5 – Observed relation between non-dimentionalised measured stress $\frac{\tilde{\sigma}_{yy}}{E_B}$ and measured strain $\tilde{\varepsilon}_{yy}$ in the scenario used to display size effect. Here the mesoscale model with a ratio between the stiffnesses $\frac{E_B}{E_b}$ of 6.2×10^4 was used to produce those results. These observations show a clear dependence of the behaviour on the size of the sample.




Figure A.6 – Non-dimentionalised measured stress $\frac{\tilde{\sigma}_{yy}}{E_B}$ at -3% strain $\tilde{\varepsilon}_{yy}$ in function of the non-dimentionalised barbs width N^{-1} (—) Is also plotted a regression of this relation (—). As can be seen, the stress is almost a quadratic function of the width of the barbs. The dependence is probably not entirely quadratic due to the impact of the barbule bands.

Appendix B

ARCSim

B.1 ARCSim implementation

In this section, we will discuss how ARCSim and our modification to its source code solves the following continuous dynamic elasticity problem

$$\begin{aligned}\rho\ddot{\mathbf{u}}(\mathbf{r}) - \nabla \cdot \mathbf{F}\boldsymbol{\sigma}(\mathbf{F}(\mathbf{r})) &= \mathbf{0}, \forall \mathbf{r} \in \Omega; \\ \mathbf{u}(\mathbf{r}) &= \mathbf{g}(\mathbf{r}), \forall \mathbf{r} \in \Gamma,\end{aligned}\tag{B.1}$$

where Γ is the boundary of the reference state Ω , and $\mathbf{g} : \Gamma \rightarrow \mathbb{R}^3$ is a user defined dirichlet boundary condition.

Discretised problem

Similar to a textbook finite element method, ARCSim discretises the displacement in space through linear triangular elements and in time through an implicit Euler's method. In what follows, we chose to use the same notations for the discretised value and the continuous values, this should not introduce any ambiguity since they are never compared to each other. With this, the discretisation of the displacement will be represented by a vector (u_i) . Each of its scalar component u_i gives the displacement in one direction at one node of the chosen mesh. We denote the interpolant associated to the component u_i by $\boldsymbol{\varphi}_i : \Omega \rightarrow \mathbb{R}^3$, meaning that the discretised displacement field \mathbf{u} is given by $\mathbf{u} = \sum_i u_i \boldsymbol{\varphi}_i$. The material position of nodes on the reference state boundary are denoted by Γ . Also, we denote the time derivative of the displacement, the velocity, by \mathbf{v} . Quantities at time step t_n will be marked by a superscript n , meaning that the displacement at time step t_n is denoted by \mathbf{u}^n . Finally, the difference between two time step, the step size, will be denoted by h .

With these notations, ARCSim solves the following non-linear problem at each time step,

$$\begin{aligned}\frac{1}{h}\mathbf{M}(\mathbf{v}^{n+1} - \mathbf{v}^n) &= \mathbf{f}(\mathbf{u}^{n+1}) + \mathbf{C}\boldsymbol{\mu}; \\ \mathbf{u}^{n+1} &= \mathbf{u}^n + h\mathbf{v}^{n+1}; \\ \mathbf{u}^{n+1}(\mathbf{r}) &= \mathbf{g}(\mathbf{r}), \forall \mathbf{r} \in \Gamma,\end{aligned}\tag{B.2}$$

where \mathbf{M} is a lumped mass matrix, \mathbf{C} is the transpose of the jacobian of the boundary constraints, $\boldsymbol{\mu}$ are the associated Lagrange's multipliers, and the

nodal elastic forces $\mathbf{f}(\mathbf{u}^{n+1})$ are given by

$$\mathbf{f}(\mathbf{u}^{n+1}) = -\frac{\partial}{\partial \mathbf{u}^{n+1}} \int_{\Omega} \psi_m(\mathbf{u}^{n+1}; \mathbf{r}) \, d\mathbf{r}. \quad (\text{B.3})$$

The intuition behind these equations comes from a direct application of Newton's second law of motion, the product of the mass and acceleration of a body is equal to the sum of forces acting on this body. Here, the force is given by the gradient of the elastic energy.

This system of equations looks different from the one resulting from a textbook finite element method which usually use Galerkin's method. We prove in the following subsection that the system of equation used by ARCSim is actually equivalent to the one obtained through Galerkin's method.

ARCSim as a finite element method

For the proof, we will leave aside the boundary constraints and the integration of displacement to focus solely on the part that differ between ARCSim and usual Galerkin's method. With these removed, the system of equation solved by ARCSim can be written

$$\frac{1}{h} \rho A_i (v_i^{n+1} - v_i^n) = - \int_{\Omega} \frac{\partial \psi_m^{n+1}}{\partial u_i^{n+1}} \, d\mathbf{r}, \forall i, \quad (\text{B.4})$$

where A_i is a third of the total area of elements incident to the node associated to the displacement component u_i .

In comparison, the system of equation resulting from Galerkin's method gives

$$\frac{1}{h} \rho A_i (v_i^{n+1} - v_i^{n+1}) = - \int_{\Omega} \mathbf{F}^{n+1} \boldsymbol{\sigma}^{n+1} : \nabla \boldsymbol{\varphi}_i \, d\mathbf{r}, \forall i. \quad (\text{B.5})$$

To prove that these two system of equations are the same, we have to show that for any displacement component u_i , we have

$$\frac{\partial \psi_m}{\partial u_i} = \mathbf{F} \boldsymbol{\sigma} : \nabla \boldsymbol{\varphi}_i, \quad (\text{B.6})$$

or, equivalently, consider the differential of the considered displacement component du_i and show that the differential of the membrane energy surface density $d\psi_m$ is

$$d\psi_m = \mathbf{F} \boldsymbol{\sigma} : \nabla \boldsymbol{\varphi}_i \, du_i \quad (\text{B.7})$$

First, since the derivative of the membrane surface energy density ψ_m with respect to the Green-Lagrange strain $\boldsymbol{\varepsilon}$ gives the second Piola-Kirchhoff stress $\boldsymbol{\sigma}$, we have,

$$\begin{aligned} d\psi_m &= \frac{\partial \psi_m}{\partial \boldsymbol{\varepsilon}} : d\boldsymbol{\varepsilon} \\ &= \boldsymbol{\sigma} : d\boldsymbol{\varepsilon}. \end{aligned} \quad (\text{B.8})$$

Then, working from the definition of the Green-Lagrange strain $\boldsymbol{\varepsilon}$, we get

$$d\boldsymbol{\varepsilon} = \frac{1}{2} (\mathbf{F}^T d\mathbf{F} + d\mathbf{F}^T \mathbf{F}). \quad (\text{B.9})$$

Plugging this back into the differential of the membrane energy density, we get

$$d\psi_m = \frac{1}{2} \boldsymbol{\sigma} : (\mathbf{F}^T d\mathbf{F} + d\mathbf{F}^T \mathbf{F}). \quad (\text{B.10})$$

Since the second Piola-Kirchhoff strain is symmetric we have for any any matrix \mathbf{A}

$$\boldsymbol{\sigma} : \mathbf{A} = \boldsymbol{\sigma} : \mathbf{A}^T. \quad (\text{B.11})$$

Hence, we can simplify the expression for the differential of the membrane energy density,

$$d\psi_m = \boldsymbol{\sigma} : (\mathbf{F}^T d\mathbf{F}). \quad (\text{B.12})$$

Now, we work out the expression of the differential of the deformation gradient. We start by looking at the expression of the deformation gradient,

$$\begin{aligned} \mathbf{F} &= \frac{\partial \mathbf{u}}{\partial \mathbf{r}} \\ &= \frac{\partial}{\partial \mathbf{r}} \sum_j u_j \boldsymbol{\varphi}_j \\ &= \sum_j u_j \nabla \boldsymbol{\varphi}_j^T. \end{aligned} \quad (\text{B.13})$$

Hence, the differential of the deformation gradient is given by

$$d\mathbf{F} = \nabla \boldsymbol{\varphi}_i^T du_i. \quad (\text{B.14})$$

With this, we finally have the differential of the membrane energy expressed in function of the differential of the considered displacement component:

$$\begin{aligned} d\psi_m &= \boldsymbol{\sigma} : (\mathbf{F}^T \nabla \boldsymbol{\varphi}_i^T du_i) \\ &= \text{Tr}(\boldsymbol{\sigma} \nabla \boldsymbol{\varphi}_i \mathbf{F}) du_i \\ &= \text{Tr}(\mathbf{F} \boldsymbol{\sigma} \nabla \boldsymbol{\varphi}_i) du_i \\ &= \mathbf{F} \boldsymbol{\sigma} : \nabla \boldsymbol{\varphi}_i du_i, \end{aligned} \quad (\text{B.15})$$

Where we expressed the double contraction of two matrices using the trace and used its property to obtain the final expression. This ends our proof that ARCSim is equivalent to Galerkin's method using a lumped mass matrix.

Solving the non-linear system

As mentioned in Section 4.2, previous version of ARCSim only solved one linearisation of the non-linear time-stepping problem which we found to be insufficient for our purpose. Instead, we used the non-linear constrained optimisation problem solver IPOPT. To do so, we rephrase the discretised problem as the following optimisation problem

$$\min_{\substack{\mathbf{v}^{n+1} \\ \mathbf{v}^{n+1} \in \mathcal{C}}} (\mathbf{v}^{n+1} - \mathbf{v}^n)^T \mathbf{M} (\mathbf{v}^{n+1} - \mathbf{v}^n) + \int_{\Omega} \psi_m(\mathbf{u}^n + h\mathbf{v}^{n+1}; \mathbf{r}) d\mathbf{r}, \quad (\text{B.16})$$

where \mathcal{C} is the set of velocity that satisfy the boundary conditions.

Appendix C

Experimental scenarios

C.1 Intermediate model and transverse scenario

In Section 5.2, we derive the following expression for the slope of the linear regression of stress strain relationship obtained through the mesoscale model in the context of the transverse scenario

$$\Delta_{\perp} = E_b \left[\frac{-15\Phi_B^2 \lambda_{\max} - 25\Phi_B^2 + 18\Phi_B \lambda_{\max}^2 + 39\Phi_B \lambda_{\max} + 23\Phi_B}{5(\Phi_B - 1)^3 (\lambda_{\max} + 1)^3} - \frac{1}{(\Phi_B - 1)^3} \right]. \quad (\text{C.1})$$

In the present section, we detail the derivation of this relation.

First we consider the transverse scenario where the bands of both the barbules and of the barbs are deformable. Since both bands have a Poisson's ratio of zero, there is no Poisson's effect and the displacement is uniform per bands. Moreover, there is no displacement along the x-axis, hence, the problem is one dimensional. We denote by λ the stretch of the whole patch, which a parameter of the problem here, and respectively by λ_B and λ_b the stretch of the barb and barbule bands, which are degrees of freedom (see Figure C.1). Those two last stretch are linked to the overall stretch through the following relation

$$\lambda = \Phi_B \lambda_B + (1 - \Phi_B) \lambda_b. \quad (\text{C.2})$$

With these notations, we can write the force balance between the two bands,

$$\frac{1}{2} E_B \lambda_B (\lambda_B^2 - 1) = \frac{1}{2} E_b \lambda_b (\lambda_b^2 - 1). \quad (\text{C.3})$$

Using the computational algebra system SymPy (Meurer et al. 2017) equations we can determine our two unknowns λ_B and λ_b . This expression for λ_b can be found below in Equation (C.4) where common subexpressions have been collected to fit on the page. Unfortunately the resulting expression is difficult to analytically link to a linear regression of the experimental stress-strain curve to obtain a value for the Young's modulus. However, a stress-strain relationship can still be plotted numerically for a chosen set of parameters (see Figure 5.4).




Figure C.1 – Schematic of the transverse scenario modelled using the intermediate model. The overall stretch ratio λ and the stretch ratio of the barbs λ_B and of the barbules λ_b can be computed from the rest heights H, H_B, H_b and the deformed heights H', H'_B, H'_b : $\lambda = \frac{H'}{H}, \lambda_B = \frac{H'_B}{H_B}, \lambda_b = \frac{H'_b}{H_b}$.

$$\begin{aligned}
x_0 &= 3E_B; \\
x_1 &= E_B\Phi_B^2; \\
x_2 &= 3x_1; \\
x_3 &= -6E_B\Phi_B\lambda + \lambda x_0 + \lambda x_2; \\
x_4 &= \Phi_B^3; \\
x_5 &= -E_Bx_4 + E_b x_4; \\
x_6 &= 3E_B\Phi_B - E_B - x_2 - x_5; \\
x_7 &= \frac{1}{x_6}; \\
x_8 &= x_7(E_B\lambda^3 - \lambda x_1); \\
x_9 &= \frac{x_3^3}{x_6^3}; \\
x_{10} &= \frac{1}{x_6^2}; \\
x_{11} &= \lambda^2; \\
x_{12} &= \Phi_B x_0 x_{11} - x_0 x_{11} + x_1 + x_5; \\
x_{13} &= x_{10} x_{12} x_3; \\
x_{14} &= x_{10} x_3^2 - 3x_{12} x_7; \\
x_{15} &= \sqrt[3]{-\frac{9x_{13}}{2} + \frac{27x_8}{2} + x_9 + \frac{\sqrt{-4x_{14}^3 + (-9x_{13} + 27x_8 + 2x_9)^2}}{2}}; \\
\lambda_b &= -\frac{x_{14}}{3x_{15}} - \frac{x_{15}}{3} - \frac{x_3 x_7}{3}.
\end{aligned} \tag{C.4}$$

This plotted stress-strain relation is compared with the stress-strain relationship obtained when the band of the barb is considered rigid, i.e. $\lambda_B = 1$. The expression of this stress strain relationship is given in Equation (5.4) in

Section 5.2. To obtain this equation, we start from the expression of the force f acting on the top clamp

$$f = W \frac{1}{2} E_b \lambda_b (\lambda_b^2 - 1) \quad (\text{C.5})$$

From this force, we derive the stress $\tilde{\sigma}_{yy}^r$ of the overall material

$$\begin{aligned} \tilde{\sigma}_{yy}^r &= \frac{1}{W\lambda} f \\ &= \frac{E_b \lambda_b (\lambda_b^2 - 1)}{2\lambda} \end{aligned} \quad (\text{C.6})$$

We can then replace λ using Equation (C.2) to obtain Equation (5.4).

In Section 5.2, we showed numerically that the expression of the stress is pretty much the same in both the rigid and non-rigid case. Hence, we will derive the analytical expression of the slope of the linear regression Δ_{\perp} from the simpler expression derived through the rigid assumption. To find this expression, we simply express the linear regression problem as a mean square minimisation:

$$\min_{\Delta_{\perp}} \int_0^{\varepsilon_{\max}} (\tilde{\sigma}_{yy}^r(\tilde{\varepsilon}_{yy}) - \Delta_{\perp} \tilde{\varepsilon}_{yy})^2 d\tilde{\varepsilon}_{yy}. \quad (\text{C.7})$$

And we solve it using, once again, SymPy which results in the expression given in Equation (C.1).

C.2 Equipment

In this section, we provide a detailed description of the used devices and give the precision for the measuring captors in our experimental setup.

Quasi-static traction experiments are carried out by means of a motorized translational stage Zaber X-LRT0500BL-E08C, which allows us to control the pulling position to an accuracy of $\pm 1 \mu\text{m}$. Simultaneously, we measure the traction force with a load cell. Depending on the orientation of the barbs, forces may vary widely. For hundreds of milli-Newtons, we use a Futek LBS-200 with a maximum load of 5 N and an accuracy of $\pm 25 \text{ mN}$. For unit Newtons, we measure with a force gauge ME-systems KD-80s with a maximum load of 100 N and an accuracy of $\pm 0.5 \text{ N}$. The signals from the load sensors are acquired by a Ni DAQ at 16 bits and sampled at high frequencies to produce statistically significant measurements. Finally, each run of a traction experiment is imaged at the patch scale with a camera mounted on a microscope Leica z16.

Appendix D

French overall summary

D.1 Introduction

Les plumes (ou plumes de contour) sont des plumes qui parent les ailes et queues des oiseaux. En raison de leur structure multi-échelle hiérarchique, elles ont des propriétés hors du commun. Cette structure est constituée d'une tige centrale nommée *rachis*, de cette tige sortent des centaines d'autres appelées *barbes*, et sur chacune d'elles se trouve une centaine de branches nommées *barbules*. Ces barbules ont une structure qui diffère en fonction de leur orientation par rapport à la barbe. Celles dirigées vers la pointe extérieure du rachis ont des crochets à leur extrémité, celles dirigées vers la base du rachis ont la forme d'une gouttière dans laquelle les crochets peuvent venir se loger (voir Figure D.1) Cette particularité permet aux barbes de s'accrocher entre elles, formant ainsi une surface le long du rachis dénommée *étendard*. Cette structure permet à l'étendard de dissiper les chocs en se fracturant de façon réversible. L'oiseau peut ensuite reformer l'étendard abîmé en accrochant les barbules entre elles à travers le passage de son bec dans ses plumes (F. ZHANG et al. 2018a).




FIGURE D.1 – Schéma de la structure hiérarchique d'une plume. De chaque côté de la tige centrale, nommée *rachis*, se trouvent les *barbes*. Comme le rachis, les barbes sont bordées d'autres tiges, *barbules*. Dans le cas spécifique des plumes, les barbules pointant vers le haut de la plume ont des crochets qui leur permettent de s'accrocher à d'autres barbules. Les barbes attachées entre elles par leurs barbules forment *l'étendard*.

En plus d'avoir ce comportement inélastique particulier, l'étendard se comporte comme une coque fortement anisotrope, similaire aux coques renforcées

par des tiges. En raison de la grande rigidité des barbes comparée à celle des barbules, l'étendard est plus rigide de plusieurs ordres de grandeurs dans la direction des barbes que dans celle orthogonale aux barbes.

Nous considérons que ces propriétés hors du commun de la plume en font un objet d'étude intéressant pour augmenter les connaissances concernant l'ornithologie, permettre la création future de matériaux par biomimétisme, ainsi que produire des outils aidant à la création d'images animés par ordinateur de plumes. À notre connaissance, aucune modélisation de ce comportement de l'étendard n'existe à ce jour. Les études de la mécanique des plumes se sont concentrées sur des aspects différents. L'interaction entre deux barbes due à l'imbrication des barbules a été étudiée à plusieurs reprises (KOVALEV et al. 2014; CHEN et al. 2016; F. ZHANG et al. 2018a). Certaines études ont aussi analysé le comportement mécanique de la plume complète, rachis compris (BHAR et al. 2019; ZHAO et al. 2020b). Ainsi, nous proposons un modèle de membrane caractérisant la forte anisotropie de la plume.

La modélisation numérique de matériaux fortement anisotropes voire quasiment inextensibles dans une direction ou deux directions est confrontée à un phénomène de verrouillage lorsque des éléments finis standards sont utilisés (WRIGGERS et al. 2018; YU et al. 2006). C'est-à-dire que la vitesse de convergence de ces méthodes est considérablement réduite, ou que leur convergence vers la solution du problème continue est impossible (BREZZI et FORTIN 1991). Plusieurs travaux ont proposé des solutions pour limiter ces effets. Par exemple, dans le cadre de matériaux ayant deux directions quasi inextensibles, le verrouillage est soit résolu par un alignement des éléments avec ces directions, soit par une intégration réduite de la partie raide de l'énergie élastique, ou encore par l'introduction d'éléments mixtes (YU et al. 2006; TEN THIJE et AKKERMAN 2008). Pour les cas où l'extension dans une direction seulement est raide, par exemple pour les matériaux renforcés par des fibres, les solutions proposées se reposent sur l'utilisation d'éléments finis mixtes (WRIGGERS et al. 2018; SCHRÖDER et al. 2016).

Notre modèle de plume présentant une grande raideur dans une direction (celle des barbes), nous avons choisi de le résoudre numériquement à travers des éléments finis linéaires alignés selon la direction de plus grande raideur.

Pour résumer, les contributions qui seront présentées sont :

- Un modèle de membrane pour l'étendard des plumes ;
- Une comparaison de ce modèle avec des tests de traction sur des échantillons réels de plumes de cygne ;
- Une analyse du phénomène de verrouillage qui émerge de la forte anisotropie présente dans notre modèle. Cette étude confirme les analyses précédentes de l'état de l'art.

D.2 Méthode expérimentale

Afin de modéliser les étendards de plumes et de valider le modèle résultant, nous avons analysé plusieurs échantillons de plumes de cygne à travers des tests de traction. Comme la plume est clairement anisotrope due à la présence des barbes, les échantillons sont découpés à différents angles θ entre barbes et

direction d'extension. Nous avons choisi les directions transverse et longitudinale pour leur simplicité d'analyse et une direction en biais à 45° dont le but est d'activer les déformations de cisaillement (voir Figure D.2). Nous avons testé 10 échantillons différents pour la direction transverse, 9 échantillons pour la direction longitudinale et 8 échantillons pour la direction en biais.




FIGURE D.2 – Haut : Plume de cygne et exemples de découpes tracées en pointillé pour chaque orientation d'échantillon testé. Bas gauche : Photo du montage expérimental. Bas droite : Gros plan sur le test de traction.

De ces tests de traction, nous tirons des mesures de contraintes de Green-Lagrange ε et du second tenseur de déformation de Piola-Kirchhoff σ (voir Figure D.3). On peut noter que la raideur du test longitudinal est quatre ordres de grandeur plus grande que celle des autres tests, montrant bien le comportement fortement anisotrope de l'étendard. De plus, dans le test transverse le comportement semble linéaire malgré les grandes déformations, cela nous amène à choisir une relation linéaire entre déformation et contrainte pour notre modèle. Les résultats du scénario en biais présentent de grandes différences : cette disparité est discutée dans Annexe D.5.

D.3 Modèle

Afin de décrire notre modèle de membrane pour l'étendard, nous fixons le référentiel de l'espace de référence de sorte que l'axe x soit colinéaire aux barbes et que la base soit orthonormée. Dans la section précédente, nous avons identifié une relation linéaire entre contrainte et déformation dans les résultats expérimentaux. Ainsi, en prenant les mêmes notations que dans la section précédente, cette relation linéaire entre contrainte et déformation en notation de Voigt s'écrit

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix} = \underbrace{\begin{bmatrix} E_{11} & E_{12} & E_{13} \\ E_{12} & E_{22} & E_{23} \\ E_{13} & E_{23} & E_{33} \end{bmatrix}}_{\mathbf{E}} \begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{xy} \end{bmatrix}.$$

Ici, le tenseur d'élasticité \mathbf{E} est symétrique pour des raisons de conservation d'énergie. Les paramètres de ce tenseur peuvent être réduits en utilisant des arguments géométriques. L'étendard a un axe de symétrie local en tous points

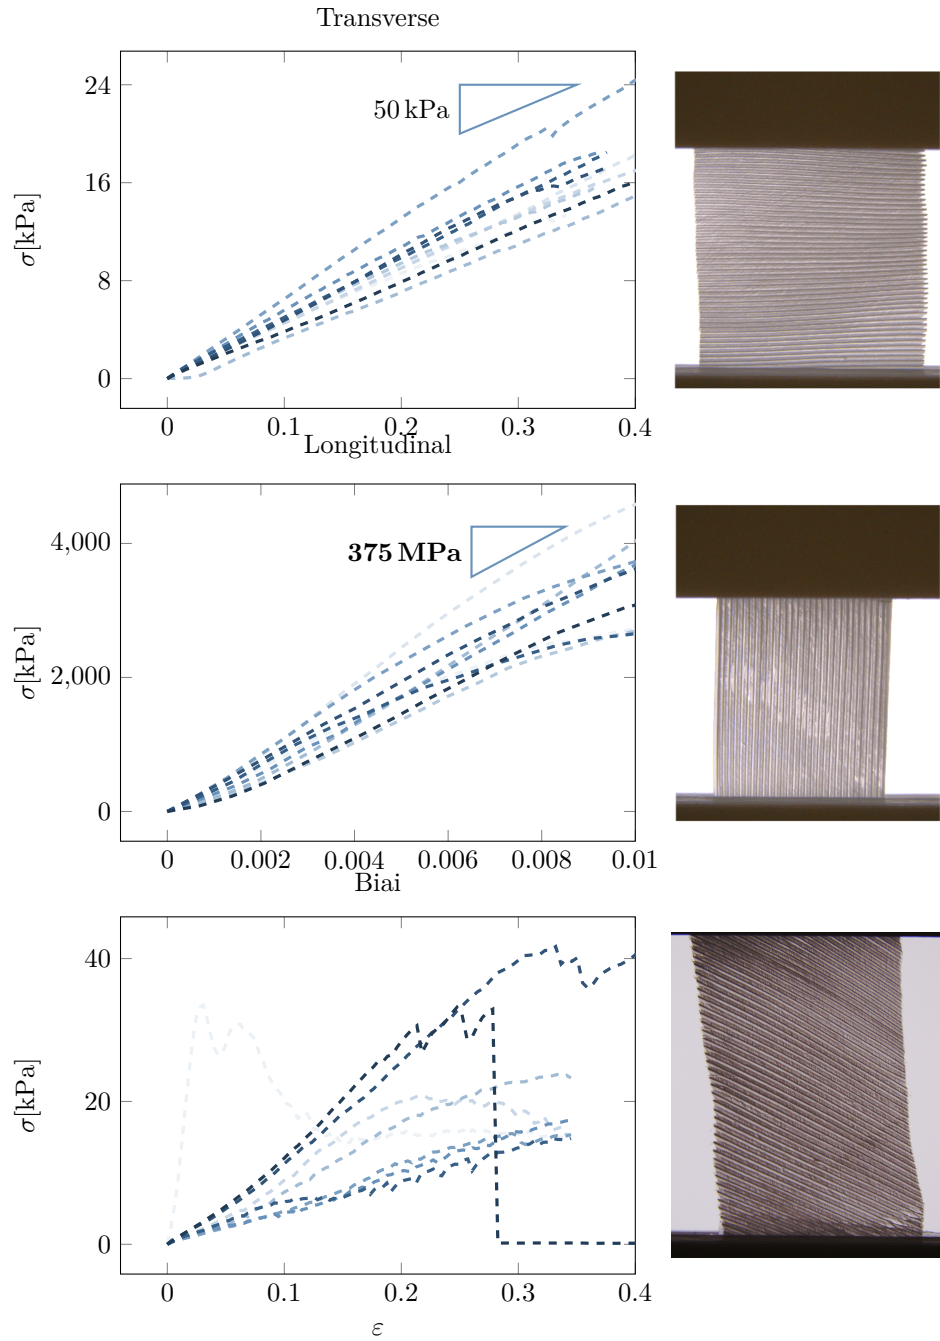


FIGURE D.3 – Résultats expérimentaux pour chaque orientation. Les dimensions associées aux échantillons testé sont listées dans Table D.1. Pour la direction transverse, les pentes des courbes varient entre 37 kPa et 61 kPa, et entre 263 MPa et 457 MPa pour la direction longitudinale.

	W [mm]	H [mm]	b [μm]		W [mm]	H [mm]	b [μm]
---	15	14	120	---	9	18	73
---	8	14	100	---	7	15	133
---	12	11	120	---	9	23	60
---	10	11	110	---	8	13	120
---	10	13	95	---	9	16	129
---	12	11	140	---	4	12	116
---	11	13	120	---	10	11	80
---	10	15	100	---	7	19	125
---	17	14	112	---	10	10	117
---	10	13	125				

(a) Transverse		(b) Longitudinal	
	W [mm]	H [mm]	b [μm]
---	6	15	120
---	3	14	110
---	9	16	132
---	12	16	155
---	8	14	143
---	16	16	130
---	14	19	114
---	13	18	103

(c) Biais

TABLE D.1 – Largeur W , hauteur H et épaisseur b des échantillons testées pour chaque orientations.

aligné sur les barbes (voir Figure D.4), ce qui induit un modèle orthotrope, c'est-à-dire $E_{13} = E_{23} = 0$. De plus, nous avons pu voir que les échantillons ne subissent aucun changement de largeur dans les tests transversaux et longitudinaux, ainsi il n'y a aucun phénomène de Poisson et on a $E_{12} = 0$. Avec cela, le modèle est réduit à la relation

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} E_{xx} & 0 & 0 \\ 0 & E_{yy} & 0 \\ 0 & 0 & E_{xy} \end{bmatrix} \begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{xy} \end{bmatrix},$$

où les paramètres du tenseur d'élasticité ont été renommés pour mettre en valeur les quantités qu'ils relient. Aussi, nous dénoterons parfois E_{xx} par *module longitudinal*, E_{yy} par *module transverse* et E_{xy} par *module de cisaillement*. Grâce aux résultats expérimentaux on peut déduire que le rapport entre le module longitudinal et le module transverse $\frac{E_{xx}}{E_{yy}}$ est de l'ordre de 10^4 , ce qui indique bien l'aspect fortement anisotrope des pennes. Dans la suite du résumé, nous appelons le rapport $\frac{E_{xx}}{E_{yy}}$ *rapport de raideur*.




FIGURE D.4 – En tous points, l'étendard admet un axe de symétrie colinéaire aux barbes (-----).




FIGURE D.5 – Maillage utilisé pour une simulation du scénario expérimental en biais. L'alignement des éléments sur l'orientation des barbes réduit considérablement le phénomène de verrouillage numérique.

D.4 Verrouillage

Nous implémentons notre modèle numériquement grâce au simulateur libre de plaques minces élastiques qui nous est le plus familier ARCSim (PFAFF et al. 2014) que nous modifions pour utiliser IPOPT comme solveur non-linéaire (WÄCHTER et BIEGLER 2006). Ce simulateur ne supporte actuellement que des éléments triangulaires linéaires, l'impact du type d'éléments utilisé est gardé pour des travaux futurs. La modélisation de notre modèle dans le cas fortement anisotrope présente un désagrément. Cette ennuie est la présence d'un phénomène de verrouillage numérique, c'est-à-dire que plus le rapport de raideur est élevé, plus la vitesse de convergence du modèle d'élément fini est faible, voire la convergence devient impossible.

Pour atténuer le phénomène de verrouillage, nous nous inspirons des travaux de YU et al. (2006) et TEN THIJE et AKKERMAN (2008), dans lesquels ils et elles considèrent un matériau ayant deux directions de très grande raideur comparée au module de cisaillement associé à ces directions. Une simulation naïve d'un tel matériau présente un phénomène de verrouillage. Une de leurs solutions pour pallier cela est d'aligner les éléments selon les directions de forte raideur. Nous la reproduisons dans notre contexte en faisant en sorte que chaque élément triangulaire ait au moins une arête alignée selon la direction des barbes (voir Figure D.5). Comme attendu, cet alignement limite considérablement le phénomène de verrouillage numérique (voir Figure D.6).

D.5 Validation

Notre modèle a trois paramètres indépendants E_{xx} , E_{yy} et E_{xy} . Les deux premiers paramètres peuvent être mesurés à travers les tests de traction transverses et longitudinaux. Le troisième E_{xy} nécessite un scénario en biais. Ainsi, la calibration seule du modèle demanderait d'utiliser nos trois scénarios de laboratoire, ce qui nous contraindrait à produire une nouvelle expérience afin de valider le modèle.

Pour éviter l'exécution laborieuse de nouvelles expériences en laboratoires,




FIGURE D.6 – Contrainte σ pour une déformation de 5 % en fonction du nombre d'éléments dans une simulation du scénario en biais. Dans ces simulations le rapport de raideur est fixé) $\frac{E_{xx}}{E_{yy}} = 10^4$. Les courbes de convergence sont données pour une simulation utilisant un maillage arbitraire (—●—) et pour une simulation utilisant un maillage aligné (—■—). La valeur de référence affichée (---) est celle obtenue avec la simulation sans contrainte utilisant le maillage aligné constituée de 128×10^3 éléments. La simulation utilisant un maillage non-aligné a une vitesse de convergence plus faible que la simulation utilisant un maillage aligné. Cela met en valeur le phénomène de verrouillage dans le cadre de cette simulation et montre que l'utilisation d'un maillage aligné pallie ce phénomène.

nous choisissons de ne pas valider notre modèle directement sur les expériences. Au lieu de cela, nous introduisons un modèle de membrane plus basse échelle pouvant être validé avec seulement nos trois scénarios. Dans le reste du document, nous appellerons ce modèle intermédiaire *modèle méso* et notre modèle d'étendard présenté plus tôt *modèle macro*. Ainsi après avoir validé le modèle méso, nous l'utilisons pour produire en simulation les scénarios nous manquant pour la validation du modèle macro.

Le modèle méso pour l'étendard se trouve à l'échelle des barbes. C'est un modèle de membrane hétérogène fait de bandes raides représentant les barbes connectées ensemble par des bandes souples représentant les barbules. Les matériaux de ces bandes sont décrits par le modèle de Saint Venant-Kirchhoff isotrope dont le coefficient de Poisson est nul. Ainsi, le modèle a seulement deux paramètres physiques, le module de Young des barbes E_B et le module de Young des barbules E_b (voir Figure D.7).

Ce modèle n'ayant que deux paramètres physiques, seulement deux des scénarios sont nécessaires pour les calibrer. Les autres paramètres étant géométriques, ils peuvent être mesurés sur les images des échantillons prises au début de chaque expérience. Les mesures des modules E_B et E_b sont données dans le tableau D.2.

Les prédictions de la simulation sont comparées avec le scénario en biais.

E_b [Pa]	E_B [Pa]	E_{xx} [Pa]	E_{yy} [Pa]	E_{xy} [Pa]
3.92×10^4	2.57×10^9	3.60×10^8	4.80×10^4	2.27×10^4

TABLE D.2 – Mesure des paramètres physiques du modèle méso et macro.

Dans ce scénario le comportement dépend fortement de l'aspect de forme de l'échantillon utilisé. Pour rendre compte de cette dépendance, des échantillons ayant des rapport de forme variant de 1.0 à 3.82 ont été utilisés (voir Figure D.8). Les calculs en simulation prédisent une contrainte légèrement plus élevée que la mesure expérimentale. Cela peut être expliqué en partie par des événements non-élastiques de réarrangement des barbules ayant lieu dans ces expériences (voir Figure D.9). Notre modèle étant complètement élastique, il ne peut rendre compte de ces événements dissipateurs d'énergie. Dans certains scénarios en biais exécutés en laboratoire, aucun réarrangement des barbules n'a pu être identifié visuellement : pour ces scénarios la simulation prédit un comportement bien plus proche de l'expérience.




FIGURE D.7 – Schéma du modèle intermédiaire *méso* utilisé pour la validation de notre modèle d'étendard. Ce modèle de membrane est formé d'une alternance de bandes élastiques raides et isotropes représentant les barbes (■) et de bandes élastiques souples et isotropes représentant les barbules (■) accrochant les barbes entre eux. Le modèle a quatre paramètres : deux paramètres physiques, le module d'Young des bandes représentant les barbes E_B , et celui des bandes représentant les barbules E_b ; deux paramètres géométriques, la largeur des barbes w_B et la distance entre deux barbes centre à centre T .

Le modèle méso étant validé, nous considérons qu'il peut être utilisé pour valider le modèle macro. Nous reproduisons en simulation avec le modèle méso les tests de traction sur un échantillon de rapport de forme 1.38 avec différentes orientations de barbes allant de 0° à 0.53° . Les paramètres E_{yy} , E_{xx} et E_{xy} du modèle macro sont calibrés sur les orientations 0° et 45° (voir Table D.2). Les résultats de la validation montrent un bon accord entre le modèle méso et macro (voir Figure D.10).

D.6 Conclusion

Dans ce résumé étendu, nous avons décrit notre montage expérimental qui nous a permis de suggérer un modèle que nous pensons être approprié pour l'étendard des pennes. Nous implémentons ce modèle grâce à des éléments finis linéaires. Pour pallier le phénomène de verrouillage numérique et de mauvais conditionnement dus à la forte anisotropie de notre modèle, nous proposons d'aligner les éléments le long des barbes et de remplacer la partie raide de

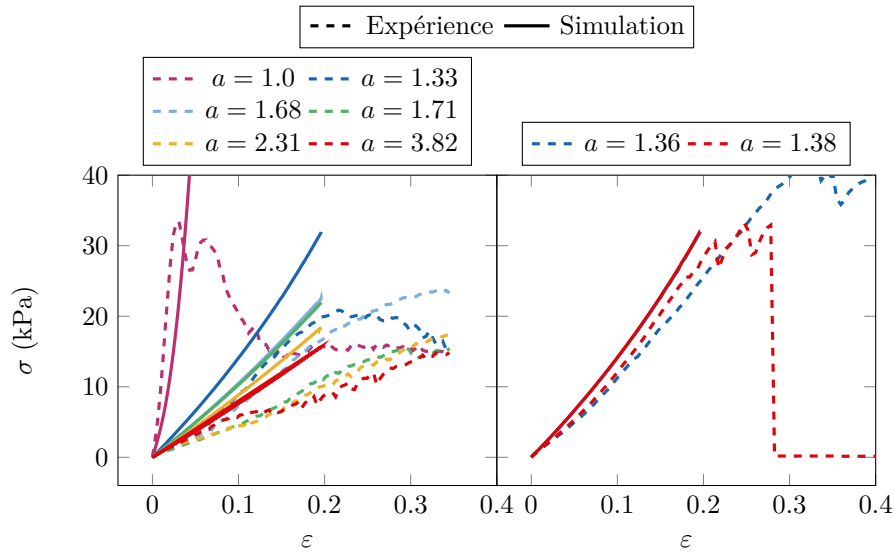


FIGURE D.8 – Comparaison de la simulation et de l'expérience sur le scénario en biais. Dans certaines expériences les barbules se réarrangent, dissipant ainsi l'énergie (gauche) (voir Figure D.9). Notre modèle d'étendard étant strictement élastique, il prédit une contrainte légèrement plus haute que celle observée dans l'expérience. D'autres expériences ne présentent aucun réarrangement (droite), en conséquence la simulation fournit de meilleures prédictions.



(a) Comportement non élastique

(b) Comportement élastique

FIGURE D.9 – Photographies de deux tests de traction. L'expérience montrée dans la sous-figure D.9a présente un comportement non-élastique ; à l'état au repos (gauche) les barbules du bord sont alignées, après extension (droite) ces barbules se réarrangent. Ce réarrangement des barbules permet aux barbes de glisser sans induire de cisaillement important et ainsi de réduire la contrainte sur le matériau. La sous-figure D.9b montre une expérience où ce comportement non élastique n'est pas présent.




FIGURE D.10 – Comparaison entre les pentes des courbes de contraintes-déformations du modèle méso et du modèle macro dans des scénarios de traction sur un échantillon de rapport de forme 1.38 avec des orientations de barbes à différents angles. Les pentes sont adimensionnées par le module de Young des barbules dans le modèle méso E_b .

l'énergie élastique par une contrainte d'inextensibilité. De plus, nous avons validé notre modèle indirectement, en utilisant un modèle intermédiaire plus bas niveau et validé expérimentalement, et ce pour réduire le nombre de scénarios différents à réaliser en laboratoire.

Annexe E

French chapter summary

E.1 Introduction

Les plumes d'oiseaux sont formées d'une structure hiérarchique de tige partant d'un rachis central, duquel partent des barbes, qui elles-mêmes portent des barbules les accrochant entre elles. Cette structure donne des propriétés mécaniques particulières à la plume, en particulier une forte anisotropie dirigée par l'orientation des barbes. À notre connaissance, très peu de travaux ont étudié le comportement élastique de la plume sous l'angle de cette forte anisotropie. Dans le cadre de cette thèse, nous proposons un modèle homogène à l'échelle de la plume pour l'élasticité non-flexurale produite à partir de résultats d'expériences faites sur des échantillons de plumes. Une implémentation numérique ainsi qu'un protocole de mesure est suggéré pour le modèle. Cette implémentation et ce protocole sont ensuite utilisés pour valider le modèle. De plus, nous suggérons une modélisation préliminaire d'une plume complète à des fins d'applications dans le domaine de l'informatique graphique.

E.2 Méthode de modélisation

Plusieurs aspects sont à prendre en compte pour la modélisation de notre modèle. Premièrement, il nous a fallu choisir l'échelle de modélisation. Nous modélisons la plume à l'échelle de la plume, c'est-à-dire en tant que surface, car cela nous permet d'utiliser toute la théorie d'élasticité déjà existante. De plus, cela nous évite les difficultés apportées par la modélisation à une échelle inférieure, i.e. un nombre de degrés de liberté prohibitif et un besoin de matériel particulièrement précis pour le protocole de calibration. Deuxièmement, nous déterminons le modèle depuis des résultats expérimentaux en nous basant sur notre intuition. Nous avons préféré cela aux techniques d'homogénéisation et de réduction de dimension, car elle nécessite un modèle existant et une connaissance précise de la géométrie que nous n'avons pas. Nous aurions pu aussi considérer les méthodes basées sur les données, mais nous voulions nous éviter la production importante de données que la plupart de ces méthodes nécessitent. Finalement, la modélisation du comportement élastique de surfaces est souvent séparée en deux aspects, l'aspect flexural et l'aspect membranaire. Dans la thèse, nous nous concentrons exclusivement sur le comportement membranaire, laissant la modélisation de la flexion pour des travaux futurs.

E.3 Expériences

Notre montage expérimental consiste en des tests de tractions sur des échantillons rectangulaires de plumes. Ces tests sont exécutés pour différentes orientations des barbes pour proprement révéler le caractère anisotrope de la plume. De ces expériences est mise en exergue une relation linéaire entre le tenseur de déformation de Green-Lagrange et le second tenseur de contrainte de Piola-Kirchhoff. De plus, les vidéos de ces expériences montrent une absence d'effet de Poisson pour les orientations de traction longitudinale et transverse aux barbes.

E.4 Modélisation

Considérant les résultats des expériences, le comportement membranaire de la plume est modélisé par une relation linéaire entre les deux tenseurs susmentionnés. Par des considérations géométriques, nous montrons que ce modèle n'a que trois paramètres, le module longitudinal E_{xx} quantifiant la raideur longitudinale aux barbes, le module transverse E_{yy} et le module de cisaillement E_{xy} . Pour implémenter numériquement ce modèle, nous utilisons un modèle d'éléments finis linéaires et triangulaires. Dû à la forte anisotropie présente dans la plume, un problème de verrouillage est présent dans notre première modélisation numérique qui considère des maillages arbitraires à cause de discontinuités dans la dérivée du champ de déplacement. Nous réduisons ce verrouillage numérique en alignant les éléments sur cette discontinuité, c'est-à-dire avec les barbes.

E.5 Calibration and validation

Pour valider notre modèle, il nous faut d'abord produire un protocole de calibration de ses paramètres. Malheureusement, les expériences considérées ne sont pas suffisantes pour mesurer directement les trois paramètres de notre modèle ainsi que de le valider. Au lieu de produire plus d'expériences, nous considérons un modèle intermédiaire qui grâce à une représentation plus basse échelle de la plume est capable d'être calibré et validé à partir des observations expérimentales que nous avons à notre disposition. Ainsi, ce modèle intermédiaire validé, nous l'utilisons pour produire numériquement plus d'observations expérimentales qui nous permettent de calibrer et valider notre modèle.

E.6 Plume complète

Pour mettre en valeur l'utilité de notre modèle dans le cadre de l'informatique graphique, nous suggérons une modélisation préliminaire d'une plume complète utilisant notre modèle précédemment validé. Ici, la plume est modélisée par une surface comprenant le rachis et l'étendard. Comme notre modèle d'étendard ne considère que le comportement membranaire, nous introduisons un modèle de flexion anisotrope à deux paramètres. Aussi, pour aligner les éléments de notre modèle sur l'orientation des barbes, nous devons introduire un algorithme de maillage. Nous procédons ensuite à une validation qualitative de ce modèle à

travers une comparaison visuelle avec une plume réelle et un modèle similaire utilisant un modèle de membrane isotrope pour l'étendard.

Bibliography

- [1] Stuart S. Antman. *Nonlinear Problems of Elasticity*. Springer-Verlag, 2005. ISBN: 0387208801. DOI: 10.1007/0-387-27649-1. URL: <http://dx.doi.org/10.1007/0-387-27649-1>.
- [2] Basile Audoly and Claire Lestringant. “An energy approach to asymptotic, higher-order, linear homogenization”. In: *Journal of Theoretical, Computational and Applied Mechanics* (2023). DOI: 10.46298/jtcam.11414. URL: <https://hal.science/hal-04112136>.
- [3] Basile Audoly and Yves Pomeau. *Elasticity and geometry: from hair curls to the non-linear response of shells*. Oxford University Press, 2010.
- [4] Nicholas Augello et al. “Hummingbird: Dreamworks Feather System”. In: *ACM SIGGRAPH Talks*. 2019, 60:1–60:2.
- [5] Igor A. Baratta et al. *DOLFINx: The next generation FEniCS problem solving environment*. 2023. DOI: 10.5281/zenodo.10447665.
- [6] Zdeněk P. Bažant. “Size effect”. In: *International Journal of Solids and Structures* 37 (Jan. 2000), pp. 69–80. ISSN: 0020-7683. DOI: 10.1016/S0020-7683(99)00077-3. URL: [http://dx.doi.org/10.1016/S0020-7683\(99\)00077-3](http://dx.doi.org/10.1016/S0020-7683(99)00077-3).
- [7] Miklós Bergou et al. “Discrete Elastic Rods”. In: *ACM Trans. Graph.* 27.3 (Aug. 2008), 63:1–63:12.
- [8] Florence Bertails et al. “Super-helices for predicting the dynamics of natural hair”. In: *ACM Trans. Graph.* 25.3 (July 2006), pp. 1180–1187. ISSN: 0730-0301. DOI: 10.1145/1141911.1142012. URL: <https://doi.org/10.1145/1141911.1142012>.
- [9] Kinjal Bhar et al. “How localized force spreads on elastic contour feathers”. In: *Journal of The Royal Society Interface* 16.160 (2019), p. 20190267.
- [10] Jeremy Bleyer. “Multiphase continuum models for fiber-reinforced materials”. In: *Journal of the Mechanics and Physics of Solids* 121 (Dec. 2018), pp. 198–233. ISSN: 0022-5096. DOI: 10.1016/j.jmps.2018.07.018. URL: <http://dx.doi.org/10.1016/j.jmps.2018.07.018>.
- [11] R Bonser and P Purslow. “The Young’s modulus of feather keratin”. In: *Journal of Experimental Biology* 198.4 (1995), pp. 1029–1033.
- [12] Stephen D. Bowline and Zoran Kačić-Alesić. “Dynamic, Penetration-free Feathers in Rango”. In: *ACM SIGGRAPH Talks*. 2011, 35:1–35:1.
- [13] Franco Brezzi and Michel Fortin. *Mixed and Hybrid Finite Element Methods*. 1991.

- [14] D. Caillerie and E. Sanchez-Palencia. “Elastic thin shells: asymptotic theory in the anisotropic and heterogeneous cases”. In: *Mathematical Models and Methods in Applied Sciences* 05.04 (June 1995), pp. 473–496. ISSN: 1793-6314. DOI: 10.1142/S0218202595000280. URL: <http://dx.doi.org/10.1142/S0218202595000280>.
- [15] Dominique Chapelle. *Etude numérique du verrouillage de quelques méthodes d’éléments finis pour les coques*. Research Report RR-2740. Projet MODULEF. INRIA, 1995. URL: <https://inria.hal.science/inria-00073953>.
- [16] Qiang Chen et al. “An analytical hierarchical model explaining the robustness and flaw-tolerance of the interlocking barb-barbule structure of bird feathers”. In: *EPL (Europhysics Letters)* 116.2 (2016), p. 24001.
- [17] Philippe G. Ciarlet. *Mathematical Elasticity Volume I: Three-Dimensional Elasticity*. Studies in Mathematics and Its Applications 20. Academic Press, Elsevier, 1988. ISBN: 9780444702593; 0444702598.
- [18] Philippe G. Ciarlet. *Mathematical Elasticity: Theory of Plates*. Society for Industrial and Applied Mathematics, Jan. 2021. ISBN: 9781611976809. DOI: 10.1137/1.9781611976809. URL: <http://dx.doi.org/10.1137/1.9781611976809>.
- [19] Philippe G. Ciarlet. *Mathematical Elasticity: Theory of Shells*. Society for Industrial and Applied Mathematics, Jan. 2021. ISBN: 9781611976823. DOI: 10.1137/1.9781611976823. URL: <http://dx.doi.org/10.1137/1.9781611976823>.
- [20] Gabriel Cirio et al. “Yarn-Level Cloth Simulation with Sliding Persistent Contacts”. In: *IEEE Transactions on Visualization and Computer Graphics* 23 (2017), pp. 1152–1162. URL: <https://api.semanticscholar.org/CorpusID:9745874>.
- [21] D. C. Drucker and W. Prager. “Soil mechanics and plastic analysis or limit design”. In: *Quarterly of Applied Mathematics* 10.2 (1952), pp. 157–165. ISSN: 1552-4485. DOI: 10.1090/qam/48291. URL: <http://dx.doi.org/10.1090/QAM/48291>.
- [22] Tracy Ducasse. Oct. 16, 2004. URL: <https://www.flickr.com/photos/59089068@N00/1050687> (visited on 06/18/2024).
- [23] Moritz Flaschel. “Automated Discovery of Material Models in Continuum Solid Mechanics”. PhD. ETH Zürich, 2023.
- [24] Moritz Flaschel et al. “Unsupervised discovery of interpretable hyperelastic constitutive laws”. In: *Computer Methods in Applied Mechanics and Engineering* 381 (Aug. 2021), p. 113852. ISSN: 0045-7825. DOI: 10.1016/j.cma.2021.113852. URL: <http://dx.doi.org/10.1016/j.cma.2021.113852>.
- [25] J. Ghaboussi et al. “Knowledge-Based Modeling of Material Behavior with Neural Networks”. In: *Journal of Engineering Mechanics* 117.1 (1991), pp. 132–153. DOI: 10.1061/(ASCE)0733-9399(1991)117:1(132). eprint: <https://ascelibrary.org/doi/pdf/10.1061/%28ASCE%290733-9399%281991%29117%3A1%28132%29>. URL: <https://ascelibrary.org/doi/abs/10.1061/%28ASCE%290733-9399%281991%29117%3A1%28132%29>.

- [26] Ionel-Dumitrel Ghiba et al. “A variant of the linear isotropic indeterminate couple-stress model with symmetric local force-stress, symmetric nonlocal force-stress, symmetric couple-stresses and orthogonal boundary conditions”. In: *Mathematics and Mechanics of Solids* 22.6 (2017), pp. 1221–1266.
- [27] Eitan Grinspun et al. “Discrete Shells”. In: SCA '03 (2003), pp. 62–67.
- [28] Rasmus Haapaoja and Christoph Genzwürker. “Mesh-driven Generation and Animation of Groomed Feathers”. In: *ACM SIGGRAPH Talks*. 2019, 61:1–61:2.
- [29] Wouter Hagens. Apr. 19, 2021. URL: https://en.wikipedia.org/wiki/Down_feather#/media/File:Donsvlokje.jpg (visited on 06/18/2024).
- [30] J. S. Hale et al. “Simple and extensible plate and shell finite element models through automatic code generation tools”. In: *Computers & Structures* 209 (2018), pp. 163–181.
- [31] Jack S Hale et al. “FENICS-SHELLS: an open-source library for simulating thin structures”. working paper or preprint. Apr. 2018. URL: <https://hal.sorbonne-universite.fr/hal-01763370>.
- [32] Daniel Heckenberg et al. “Quill: Birds of a Feather Tool”. In: *ACM SIGGRAPH Talks*. 2011, 34:1–34:1.
- [33] Robert Hook. *De Potentia Restitutiva, or of Spring*. 1678.
- [34] Jonathan M. Kaldor et al. “Simulating knitted cloth at the yarn level”. In: *ACM Trans. Graph.* 27.3 (Aug. 2008), pp. 1–9. ISSN: 0730-0301. DOI: 10.1145/1360612.1360664. URL: <https://doi.org/10.1145/1360612.1360664>.
- [35] S. Kawabata and Masako Niwa. “Fabric Performance in Clothing and Clothing Manufacture”. In: *The Journal of The Textile Institute* 80.1 (Jan. 1989), pp. 19–50. ISSN: 1754-2340. DOI: 10.1080/00405008908659184. URL: <http://dx.doi.org/10.1080/00405008908659184>.
- [36] Marzieh Khani et al. “Water uptake by gecko -keratin and the influence of relative humidity on its mechanical and volumetric properties”. In: *Journal of The Royal Society Interface* 19.194 (Sept. 2022). ISSN: 1742-5662. DOI: 10.1098/rsif.2022.0372. URL: <http://dx.doi.org/10.1098/rsif.2022.0372>.
- [37] Ryo Kikuuwe et al. “An edge-based computationally efficient formulation of Saint Venant-Kirchhoff tetrahedral finite elements”. In: *ACM Trans. Graph.* 28.1 (Feb. 2009). ISSN: 0730-0301. DOI: 10.1145/1477926.1477934.
- [38] Alexander Kovalev et al. “Unzipping bird feathers”. In: *Journal of The Royal Society Interface* 11.92 (2014).
- [39] Olga Kozachenko. Oct. 14, 2020. URL: <https://unsplash.com/photos/white-and-gray-knit-textile-o9dtfshlJ60> (visited on 09/11/2024).
- [40] David CC Lam et al. “Experiments and theory in strain gradient elasticity”. In: *Journal of the Mechanics and Physics of Solids* 51.8 (2003), pp. 1477–1508.

- [41] G.D. Macleod. “Mechanical Properties of Contour Feathers”. In: *Journal of Experimental Biology* 87.1 (1980), pp. 65–72.
- [42] Saptarshi Maiti et al. “Sustainable Fiber-Reinforced Composites: A Review”. In: *Advanced Sustainable Systems* 6.11 (Sept. 2022). ISSN: 2366-7486. DOI: 10.1002/adsu.202200258. URL: <http://dx.doi.org/10.1002/adsu.202200258>.
- [43] Benny Mazur. Original image was slightly cropped. Oct. 25, 2008. URL: <https://www.flickr.com/photos/44545509@N00/2977350640> (visited on 09/11/2024).
- [44] Aaron Meurer et al. “SymPy: symbolic computing in Python”. In: *PeerJ Computer Science* 3 (Jan. 2017), e103. ISSN: 2376-5992. DOI: 10.7717/peerj-cs.103. URL: <https://doi.org/10.7717/peerj-cs.103>.
- [45] K. Miled et al. “Effective elastic properties of porous materials: Homogenization schemes vs experimental data”. In: *Mechanics Research Communications* 38.2 (2011), pp. 131–135. ISSN: 0093-6413. DOI: <https://doi.org/10.1016/j.mechrescom.2011.01.009>. URL: <https://www.sciencedirect.com/science/article/pii/S0093641311000115>.
- [46] G.W. Milton. *The theory of composites*. Cambridge monographs on applied and computational mathematics 6. Cambridge University Press, 2002. ISBN: 051104092X; 9780511040924; 0521781256; 9780521781251.
- [47] Sayed Abolfazl Mirdehghan. “1 - Fibrous polymeric composites”. In: *Engineered Polymeric Fibrous Materials*. Ed. by Masoud Latifi. The Textile Institute Book Series. Woodhead Publishing, 2021, pp. 1–58. ISBN: 978-0-12-824381-7. DOI: <https://doi.org/10.1016/B978-0-12-824381-7.00012-3>. URL: <https://www.sciencedirect.com/science/article/pii/B9780128243817000123>.
- [48] Wereld Museum, ed. 1907. URL: <https://hdl.handle.net/20.500.11840/128122> (visited on 09/11/2024).
- [49] Wereld Museum, ed. Jan. 1, 1972. URL: <https://hdl.handle.net/20.500.11840/128332> (visited on 09/11/2024).
- [50] Rahul Narain et al. “Adaptive anisotropic remeshing for cloth simulation”. In: *ACM Transactions on Graphics* 31.6 (Nov. 2012), pp. 1–10. ISSN: 1557-7368. DOI: 10.1145/2366145.2366171.
- [51] Tobias Pfaff et al. “Adaptive tearing and cracking of thin sheets”. In: *ACM Transactions on Graphics (TOG)* 33 (2014), pp. 1–9.
- [52] Samuel Poincloux. “Élasticité et treplements du tricot”. PhD. Sorbonne Université, 2018.
- [53] P.P. Purslow and J.F. V. Vincent. “Mechanical Properties of Primary Feathers From the Pigeon”. In: *Journal of Experimental Biology* 72.1 (Feb. 1978), pp. 251–260. ISSN: 0022-0949. DOI: 10.1242/jeb.72.1.251. eprint: <https://journals.biologists.com/jeb/article-pdf/72/1/251/1343731/251.pdf>.
- [54] Klaus Bernd Sautter et al. “Limitations of the St. Venant–Kirchhoff material model in large strain regimes”. In: *International Journal of Non-Linear Mechanics* 147 (Dec. 2022), p. 104207. ISSN: 0020-7462. DOI: 10.1016/j.ijnonlinmec.2022.104207.

- [55] Jörg Schröder et al. “A novel mixed finite element for finite anisotropic elasticity; the SKA-element Simplified Kinematics for Anisotropy”. In: *Computer Methods in Applied Mechanics and Engineering* 310 (2016), pp. 475–494. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2016.06.029>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782516306478>.
- [56] Jonathan Richard Shewchuk. “Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator”. In: *Applied Computational Geometry Towards Geometric Engineering*. Ed. by Ming C. Lin and Dinesh Manocha. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 203–222. ISBN: 978-3-540-70680-9.
- [57] Georg Sperl, Rahul Narain, et al. “Homogenized Yarn-Level Cloth”. In: *ACM Transactions on Graphics (TOG)* 39.4 (2020).
- [58] Georg Sperl, Rosa M. Sánchez-Banderas, et al. “Estimation of Yarn-Level Simulation Models for Production Fabrics”. In: *ACM Transactions on Graphics (TOG)* 41.4 (2022).
- [59] Theodore Sussman and Klaus-Jürgen Bathe. “A model of incompressible isotropic hyperelastic material behavior using spline interpolations of tension-compression test data”. In: *Communications in Numerical Methods in Engineering* 25 (2009), pp. 53–63. URL: <https://api.semanticscholar.org/CorpusID:17681230>.
- [60] R.H.W. ten Thije and R. Akkerman. “Solutions to intra-ply shear locking in finite element analyses of fibre reinforced materials”. In: *Composites Part A: Applied Science and Manufacturing* 39.7 (2008), pp. 1167–1176. ISSN: 1359-835X. DOI: <https://doi.org/10.1016/j.compositesa.2008.03.014>.
- [61] The Engineering Toolbox. *Young’s Modulus, Tensile Strength and Yield Strength Values for some Materials*. 2003. URL: https://www.engineeringtoolbox.com/young-modulus-d_417.html.
- [62] Pascal Volino et al. “A simple approach to nonlinear tensile stiffness for accurate cloth simulation”. In: *ACM Trans. Graph.* 28 (2009), 105:1–105:16.
- [63] Andreas Wächter and Lorenz T. Biegler. “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming”. In: *Mathematical Programming* 106 (2006), pp. 25–57.
- [64] Andrew J. Weber and Galen Gornowicz. “Collision-free Construction of Animated Feathers Using Implicit Constraint Surfaces”. In: *ACM Trans. Graph.* 28.2 (May 2009), 12:1–12:8.
- [65] Mingjiang Zhan; Richard P. Wool. “Mechanical properties of chicken feather fibers”. In: *Polymer Composites* 32 (6 2011).
- [66] Peter Wriggers et al. “Finite and Virtual Element Formulations for Large Strain Anisotropic Material with Inextensive Fibers”. In: (2018). URL: <https://api.semanticscholar.org/CorpusID:126421967>.
- [67] Xiaobo Yu et al. “Intra-ply shear locking in finite element analyses of woven fabric forming processes”. In: *Composites Part A: Applied Science and Manufacturing* 37.5 (2006), pp. 790–803. ISSN: 1359-835X. DOI: <https://doi.org/10.1016/j.compositesa.2005.04.024>.

- [68] F. Zhang et al. “Repairable cascaded slide-lock system endows bird feathers with tear-resistance and superdurability”. In: *Proceedings of the National Academy of Sciences* 115 (2018), pp. 10046–10051.
- [69] F. Zhang et al. *Repairable cascaded slide-lock system endows bird feathers with tear-resistance and superdurability — Movie S3*. 2018. URL: https://www.pnas.org/doi/suppl/10.1073/pnas.1808293115/suppl_file/pnas.1808293115.sm03.mp4 (visited on 09/25/2024).
- [70] Zhan Zhang et al. “Computational Design of Flexible Planar Microstructures”. In: *ACM Transactions on Graphics* 42.6 (Dec. 2023), 185:1–16. DOI: 10.1145/3618396. URL: <https://inria.hal.science/hal-04425030>.
- [71] Jing-Shan Zhao et al. “Shaking the wings and preening feathers with the beak help a bird to recover its ruffled feather vane”. In: *Materials & Design* 187 (2020), p. 108410. ISSN: 0264-1275. DOI: <https://doi.org/10.1016/j.matdes.2019.108410>.
- [72] Jing-Shan Zhao et al. “Shaking the wings and preening feathers with the beak help a bird to recover its ruffled feather vane”. In: *Materials & Design* 187 (2020), p. 108410. ISSN: 0264-1275. DOI: <https://doi.org/10.1016/j.matdes.2019.108410>. URL: <https://www.sciencedirect.com/science/article/pii/S0264127519308482>.