A practical aerodynamic model for dynamic textile manipulation in robotics

Franco Coltraro[a,b,*], Jaume Amorós[b,c], Carme Torras[a], Maria Alberich-Carramiñana[a,b]

[a] *Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain.*
[b] *Mathematics Department, Universitat Politècnica de Catalunya, Barcelona, Spain.*
[c] *Centre de Recerca Matemàtica, Barcelona, Spain*

Abstract

We study an aerodynamic model describing the interaction between cloth and air, with applications to dynamic textile manipulation by robots. After introducing the model, we investigate its theoretical implications by using an analytically solvable system: the damped pendulum. We deduce that aerodynamic forces in the model manifest themselves as a lifting force, more pronounced when the cloth transitions from rest to dynamic motion. The resulting aerodynamic model is simple, with no additional computational cost. The model is validated by comparing cloth simulations to real-world data as captured by a Motion Capture System: the results demonstrate errors of less than 1 cm even for size A2 textiles. Furthermore, we develop an a priori formula for estimating the parameters of the model for various textiles without further optimization. This formula allows us to present a challenging robotics application: a dynamic flattening motion is designed in simulation and then successfully executed by a robot without any fine-tuning or modification. The outcome, a smooth and rapid laying of the real textiles, demonstrates the minimal *sim-to-real* gap of our model even when aerodynamics plays a leading role.

Keywords: aerodynamics, cloth manipulation, simulation, robotics, sim-to-real gap

*Corresponding author
Email address: franco.coltraro@upc.edu (Franco Coltraro)

1. Introduction

The manipulation of cloth is currently a frontier problem in robotics [1]. It encompasses household and hospitality tasks, for which there is significant demand for automatization, such as folding and storing laundry, setting linens on beds and tables [2], as well as industrial tasks like wrapping, sealing with tape, and managing textile components in machinery ranging from sailboats to small unmanned aerial vehicles [3]. The problem owes its frontier status to the distinct complexity of cloth dynamics. When modeled as a surface, a piece of cloth has an infinite-dimensional space of possible states, whose topological analysis encompasses significant aspects of knot theory [4].

One possible path to overcome such difficulties for robotic applications is through the use of cloth simulators [5]. The strategy is in principle simple: given a cloth model (usually a set of ordinary differential equations) either use it for designing optimal control trajectories [6] or for generating data that then can be used to train any learning algorithm [7]. The main problem that is usually faced by this approach is the so called *sim-to-real gap*: often what is learned or designed in simulation does not translate to well to the real robot. There are many reasons why this happens, but the main culprit is usually the (lack of) realism of the chosen simulator [8].

1.1. Inextensibe cloth model

While the local structure of textiles suggests modeling them as elastic thin shells [9], this approach is limited due to the tendency of cloth to buckle (or wrinkle) under even minimal compression [10]. In previous works [11, 12], the authors developed a mechanical model treating cloth as an inextensible surface that bends freely (in [11]) and implemented a cloth simulator based on this model, incorporating friction, collisions, and self-collisions (in [12]). The model was shown to be realistic by comparing simulations of manipulation tasks, including shaking, spreading a piece of cloth flat on a table, and hitting a piece of cloth with a stick, with real-world recordings of these tasks.

This comparison between real-life robotic manipulation of cloth and its simulation demonstrated that **aerodynamic effects** are critical factors influencing cloth motion. Even at the modest speeds typical of human manipulation in domestic settings, incorporating aerodynamic forces into the mechanical model of cloth is necessary to achieve simulations with a margin of error of approximately 1 cm when comparing the simulated and actual positions of points of interest on the cloth.

1.2. A practical aerodynamic model for textiles

The science of aerodynamics applies fluid mechanics to model the interaction between a moving object and its surrounding atmosphere [13]. For textiles, this interaction becomes particularly complex as the motion of the garment's edges and the surrounding air turbulences form a coupled problem [14]. This coupling significantly increases the computational cost of simulations. Thus, simulating aerodynamic forces using fluid mechanics becomes unfeasible for its use in robotic cloth manipulation, e.g. in the development of efficient control strategies. Fast simulations are also crucial for generating the vast amounts of data needed to train learning algorithms, such as neural networks. Therefore, the use of a practical, yet realistic aerodynamic model for cloth simulation becomes of critical importance for robotic applications.

In this work, we analyze analytically and empirically the solution proposed by the authors for the first time in [11] to the dilemma of velocity versus accuracy in simulating aerodynamic effects on cloth motion. This is achieved with a simple computational modification: decoupling the inertial and gravitational masses of the cloth. Assigning a smaller gravitational mass to the garment turns out to be equivalent to assuming the presence of a lifting force, akin to the buoyant force that would be experienced if the cloth were immersed in a fluid. Our lifting force correction tries to reflect the effect of turbulences, similar to friction, that the motion of the cloth generates and exerts on itself. Our inclusion of aerodynamics in the mechanical model of cloth, as a correction to the gravitational mass, incurs no computational cost and can be applied with coarse meshings, enabling fast simulations. Since this correction factor represents the resistance of the surrounding air to the motion of the cloth, its optimal value is expected to depend on the specific structure and density of the cloth, the speed of its motion, and the size of the garment.

1.3. Contributions

In this work, we report for the first time our analysis of this dependence and identify default values for the mass correction factor. These default values, which serve as an approximation for the aerodynamic effects in the simulation of various robotic manipulation tasks, achieve a margin of error of 1 cm when comparing the position of key points in simulation with real recordings of the tasks. To the best of our knowledge, this is the first comprehensive study in literature that examines how size, speed, and material

properties affect the aerodynamics of textiles in such a way that is relevant for robotic applications. Moreover, we show that physically, the proposed solution is a practical model in which aerodynamic forces are represented by a lifting force, which becomes more pronounced as the cloth starts moving from a resting position. Finally, we give a non-trivial robotic application of having an accurate model that can predict the future behavior of cloth, proving that the *sim-to-real* gap of our cloth model with the aerodynamic correction is minimal, i.e. that what we plan in simulation can be executed by a real robot with essentially the same results. The original contributions of this work are:

1. A theoretical understanding of the gravitational mass correction in the aerodynamic model.
2. An extensive study of the aerodynamic effects on textiles, analyzing the dependency of the model on fabric type, size, and speed.
3. A formula for determining a-priori the values of the physical parameters in the model without any optimization.
4. A non-trivial robotic application demonstrating the minimal sim-to-real-gap of the aerodynamic model.

1.4. Organization

The structure of the article is as follows: after a discussion of prior work on the aerodynamics of cloth motion in §2, we present our inextensible cloth model in §3, including its aerodynamic-motivated correction to the gravitational mass factor. In §4, we discuss a theoretical justification for this simplified aerodynamic correction. §5 describes our experiments with various robotic cloth manipulation tasks, including the determination of the aerodynamic parameter through comparisons between the experiments and their simulations, as well as the margin of error introduced by using a default value of the parameter based solely on the type of cloth. In §6, we highlight the application of our model to a task where aerodynamics plays a decisive role: dynamically flattening a piece of cloth on a table with a robot. Finally, in §7, we summarize the conclusions drawn from our analysis.

2. Related work

The effect of aerodynamics on the motion of textiles is a relatively underdeveloped field of study. Although some simplified models exist, e.g. [15], modeling aerodynamic forces acting on a deformable object submerged in a fluid is highly challenging (see [14] and [16]), particularly near the boundaries of the cloth. This difficulty arises, in part, because non-trivial motions induce turbulence, and the non-rigid response of cloth to such turbulence is highly unpredictable [17]. Historically, this lack of research stems from the fact that the foundational problems of aerodynamics were focused on rigid or, at most, elastic bodies in an airflow (see [13, 18]). On the other hand, textile engineering has primarly studied the behavior of cloth during the manufacturing process, where it is subject to high tension and largely unaffected by airflow (see [10]).

The effect of wind, or fluid flows in general, on flags, sails, and membranes can be studied using elasticity theory, modeling the surfaces as thin shells ([19–21]; see also references reviewed in the last source). These studies primarily focus on 2-dimensional cross-sections of the flow due to the computational challenges posed by the coupling between the shape of the flag and the turbulence it generates in the surrounding air motion. Furthermore, cloth buckles immediately under even slight compression, a characteristic that limits the applicability of elasticity theory for its analysis.

Another approach has been the attempt to model explicitly the complex coupling between airflow and fabric motion using high-fidelity simulations, primarily based on immersed boundary (IB) methods and fluid–structure interaction (FSI) models. These methods have provided insights into cloth-like dynamics, by simulating flapping flags [22], parachutes [23], and flexible filaments [24]. IB methods avoid the need for deformable meshes by embedding the deformable object in a Cartesian grid, enforcing interaction forces through additional momentum terms [25]. This allows for flexible geometries but still demands significant computational resources, particularly for 3D simulations. Traditional FSI approaches, such as Arbitrary Lagrangian–Eulerian (ALE) frameworks, provide more accurate coupling but require dynamic remeshing, iterative solvers, and extensive memory use, further limiting their feasibility for fast computation. Overall both methods are computationally expensive, requiring high-resolution grids, small time steps, and often large-scale parallel computing, making them impractical

for robotics applications [22, 23]. Given these constraints, other researchers in the community are exploring reduced-order models and data-driven approaches to approximate cloth aerodynamics while maintaining computational efficiency [26]. Recent efforts, such as spectral submanifold reductions and proper orthogonal decomposition, aim to capture the dominant dynamics with significantly fewer degrees of freedom. While promising, these techniques are still under development and are often very case-specific.

Finally, in the context of cloth simulation in computer graphics, Ling, Damodaran and Gay [27] (see also ch. 7 of [28] for a more detailed description) conducted an in-depth analysis of the behavior of textiles in an airflow, going beyond the study of cross-sections and suitable for cloth simulation. Drawing from aerodynamics research, the authors in [27] model air as an incompressible, inviscid fluid governed by the Bernouilli equation. A piece of cloth in this fluid is treated as a moving boundary condition, which can only be studied numerically because of its deformability. Meshing the cloth enables the application of the Panel Method with ring vortexes of the flow (see [18]) to model the interaction between the airflow and the deformation of the garment. While this method has been successfully applied to simulate cloth motion, it has a significant drawback: its high computational cost. This is mainly due to the fact that it requires a fine meshing of the garment, which increases computation costs across all aspects of the cloth simulation. This high cost precludes its use for real-time simulation, and even for generating the vast amount of data required to train (deep) learning algorithms.

3. Inextensible cloth model and aerodynamics

In this section we give a self-contained presentation of the intextensible cloth model developed in [11] so that we can later discuss how to incorporate aerodynamics into the equations.

We assume that the cloth S has been discretized into a triangular or quadrilateral mesh and the position of all its N vertices or nodes (denoted by $p_i(t) = (x_i(t), y_i(t), z_i(t)) \in \mathbb{R}^3$) at time $t \geq 0$ is given by $\boldsymbol{\varphi}(t) = (\mathbf{x}(t), \mathbf{y}(t), \mathbf{z}(t))^\intercal \in \mathbb{R}^{3N}$, where $\mathbf{x}(t), \mathbf{y}(t), \mathbf{z}(t)$ denote the x-coordinates (resp. y and z) of all the nodes of the discrete surface in $\mathbb{R}^N$ at time t.

Then, the inextensible cloth model without collisions consists in the following set of ordinary differential equations (ODEs):

$$\begin{cases} \rho \mathbf{M} \ddot{\boldsymbol{\varphi}} = \mathbf{f}_\rho - \kappa \mathbf{K} \boldsymbol{\varphi} - \mathbf{D} \dot{\boldsymbol{\varphi}} - \nabla \mathbf{C}(\boldsymbol{\varphi})^\intercal \boldsymbol{\lambda} \\ \mathbf{C}(\boldsymbol{\varphi}) = 0 \end{cases} \tag{1}$$

where the meaning of each term and parameter is:

1. $\rho > 0$ is the density of the cloth (which we are assuming to be homogeneous), $\mathbf{M}$ is the augmented mass matrix, $\mathbf{f}_\rho = -\rho \mathbf{M} \mathbf{g}$ is the force of gravity and $\mathbf{g} = (0, \ldots, 0 | 0, \ldots, 0 | g, \ldots, g)^\intercal$ where $g = 9.8 \mathrm{m/s}^2$,

2. the stiffness matrix (we are using the isometric bending model described in [29]) is $\mathbf{K} = \mathbf{L}^\intercal \mathbf{M} \mathbf{L}$ where $\mathbf{L}$ is an approximation of the point-wise Laplacian and $\kappa > 0$ is a bending constant,

3. we have added Rayleigh damping: $\mathbf{D} = \alpha \mathbf{M} + \beta \mathbf{K}$, where α and β are positive parameters dampening slow and fast oscillations respectively [30],

4. and finally $\boldsymbol{\lambda}(t)$ are the Lagrange multipliers ensuring *inextensibility* (to be described next) and other possible positional constraints, such as manipulating the textile by prescribing the position of its corners.

Besides further positional constraints introduced by manipulation, the smooth function $\mathbf{C} : \mathbb{R}^{3N} \to \mathbb{R}^{n_C}$ is responsible for modeling the *inextensibility* of the textile through the satisfaction of the equality constraints $\mathbf{C}(\boldsymbol{\varphi}) = \mathbf{0}$. Each constraint $C_i(\boldsymbol{\varphi}(t)) = 0,\ i = 1, \ldots, n_C$ is in fact a quadratic function of its argument and we have n_C of them, depending on the number of nodes N of the discretization.

Indeed, for each interior (non-boundary) node index $k \in \mathrm{Int}(S)$ of the discretized surface S, we have three constraints given by $E_k(t) - E_k(0) = 0$ or $F_k(t) - F_k(0) = 0$ or $G_k(t) - G_k(0) = 0$, where

$$\begin{aligned} E_k(t) &= \frac{1}{m_k} \sum_{i,j=1}^{N} \langle p_i(t), p_j(t) \rangle \cdot \int_S \mathcal{N}_k \partial_\xi \mathcal{N}_i \partial_\xi \mathcal{N}_j dA, \\ F_k(t) &= \frac{1}{m_k} \sum_{i,j=1}^{N} \langle p_i(t), p_j(t) \rangle \cdot \int_S \mathcal{N}_k \partial_\xi \mathcal{N}_i \partial_\eta \mathcal{N}_j dA, \\ G_k(t) &= \frac{1}{m_k} \sum_{i,j=1}^{N} \langle p_i(t), p_j(t) \rangle \cdot \int_S \mathcal{N}_k \partial_\eta \mathcal{N}_i \partial_\eta \mathcal{N}_j dA. \end{aligned} \tag{2}$$

Here $m_k > 0$ is one third the sum of the areas of all incident triangles (one fourth in the case of quadrilaterals) to the kth node p_k of the surface and $\mathcal{N}_i : S \to \mathbb{R}$ are piece-wise smooth continuous indicator functions (also called a Lagrange basis or hat functions) such that $\mathcal{N}_i(p_j) = \delta_{ij}$ is the Kronecker delta. See [11] for more details and an efficient algorithm to evaluate these constraints.

Remark 3.1 (Preservation of the metric of the surface)**.** Assume that $S \subset \mathbb{R}^3$ is a smooth surface moving through space, such that its motion is given by a family of surfaces $\{S_t\}_{t\geq 0}$ isometric to $S = S_0$. Then, it can be seen that Equations (2) are obtained by discretizing with *finite elements* the following system of partial differential equations:

$$\langle \varphi_\xi, \varphi_\xi \rangle(t) = E_0,\ \langle \varphi_\xi, \varphi_\eta \rangle(t) = F_0,\ \langle \varphi_\eta, \varphi_\eta \rangle(t) = G_0 \text{ for } t \geq 0, \tag{3}$$

where $\varphi^t : R \to S_t$ is a smooth parametrization of S_t, R is an open set of the plane and $(\xi, \eta) \in R$, $t \geq 0$ represents time, E_0, F_0, G_0 are the coefficients of the first fundamental form or metric of S that are constant in time (but not necessarily in R), and $\varphi_\xi(t) = \partial_\xi \varphi^t$ denotes partial differentiation with respect to the variable ξ.

3.1. Modeling of aerodynamics through virtual mass

Finally, as mentioned in the introduction, in order to model the aerodynamic effects of air resistance on cloth we allow inertial and gravitational mass to be different. Hence we write $\mathbf{f}_\delta = -\delta \mathbf{M}\mathbf{g}$, setting $\delta \leq \rho$ as a new parameter to be fitted. We call $\delta > 0$ *virtual mass*. Therefore the equations of motion become:

$$\begin{cases} \rho \mathbf{M} \ddot{\boldsymbol{\varphi}} = -\delta \mathbf{M}\mathbf{g} - \kappa \mathbf{K}\boldsymbol{\varphi} - (\alpha \mathbf{M} + \beta \mathbf{K})\dot{\boldsymbol{\varphi}} - \nabla \mathbf{C}(\boldsymbol{\varphi})^\intercal \boldsymbol{\lambda} \\ \mathbf{C}(\boldsymbol{\varphi}) = 0, \end{cases} \tag{4}$$

where the meaning of all parameters of the model is summarized in Table 1.

Remark 3.2 (Computational cost)**.** Observe that the introduction of the virtual mass parameter δ incurs in no extra computational cost, since it only modifies the gravitational force calculation. This means the simulations run as fast as the underlying cloth simulator itself – a significant advantage over fluid-coupled methods, where air is simulated independently.

Parameter	Meaning
ρ	Density (inertial mass)
δ	Virtual (gravitational) mass
κ	Bending/stiffness
α	Damping of slow oscillations
β	Damping of fast oscillations

Table 1: Physical parameters of the inextensible model and their meaning.

Remark 3.3 (Relationship with a lifting force)**.** Notice that since we assume that $\delta \leq \rho$, we have that

$$-\delta \mathbf{M}\mathbf{g} = -\rho \mathbf{M}\mathbf{g} + \gamma \mathbf{M}\mathbf{g}$$

where $\gamma = \rho - \delta \geq 0$. Hence, by allowing the inertial and gravitational masses to be different, we are actually introducing (a constant) force opposing gravity. This may be viewed as incorporating into our model the upwards component of the aerodynamic forces (such as lift and drag) to which an object moving through air is subject.

4. Theoretical consequences of introducing the virtual mass parameter

The goal of this section is to understand theoretically the consequences of the introduction of the virtual mass parameter $\delta > 0$. Notice that solving analytically system (4) is impossible even for the simplest motions and cloth topologies. For that reason we study a very well known and simpler system whose analytical solutions are easy to calculate: the *damped pendulum*. The goal of this section is to study what happens to the solutions of the damped pendulum when we introduce a virtual mass into its equations. The theoretical analysis using the damped pendulum model will demonstrate that reducing δ relative to the inertial mass leads to lower initial velocities, effectively capturing the impact of air resistance when the cloth transitions from rest to motion.

Consider the motion of a damped pendulum of length $l > 0$ and inertial mass ρ. Then (e.g. using Lagrangian mechanics) its equations of motion can be described by

$$\rho l^2 \ddot{\theta}(t) = -\delta g l \theta(t) - \beta \dot{\theta}(t) \tag{5}$$

where $\theta = \theta(t)$ is the angle that the pendulum makes w.r.t the equilibrium position, i.e. its Cartesian coordinates are given by $x = l\sin\theta$ and $y = -l\cos\theta$, $\beta > 0$ is a damping parameter and we have made the small angle approximation $\sin\theta \approx \theta$. Notice that we have already introduced the virtual mass parameter $0 < \delta < \rho$ into the equation. Then, manipulating the equation algebraically we obtain

$$\ddot{\theta} + 2b\dot{\theta} + \omega^2\theta = 0, \quad b = \tfrac{\beta}{2\rho l^2}, \quad \omega^2 = \tfrac{\delta g}{\rho l}. \tag{6}$$

As usual, to find the solutions we look for (possibly complex) functions of the form $\theta(t) = e^{rt}$. Substituting our *Ansatz* into (6) we find that $r = -b \pm \sqrt{b^2 - \omega^2}$. Thus, if $b \neq w$ the general solution is of the form

$$\theta(t) = c_1 e^{r_1 t} + c_2 e^{r_2 t},$$

where $r_1 = -b + \sqrt{b^2 - \omega^2}$, $r_2 = -b - \sqrt{b^2 - \omega^2}$ and c_1, c_2 are determined by the initial conditions $\theta(0) = \theta_0$ and $\dot{\theta}(0) = \dot{\theta}_0$. Notice that when $\omega > b$ we are actually computing the complex exponential, but we can always recover the real solutions by using Euler's formula.

Now we wish to study the effect that $\delta < \rho$ has on the speed $\dot{\theta}(t)$ for small $t > 0$ when we start from $\theta(0) = \theta_0 < 0$ and $\dot{\theta}(0) = 0$. Physically, this means releasing the pendulum from the 'left'. It is not difficult to see, that for these initial conditions to be satisfied, we need $c_1 = -\frac{\theta_0 r_2}{r_1 - r_2}$ and $c_2 = +\frac{\theta_0 r_1}{r_1 - r_2}$. Hence the velocity of the pendulum becomes:

$$\dot{\theta}(t) = \frac{r_1 r_2}{r_1 - r_2}\theta_0 \left(e^{r_2 t} - e^{r_1 t}\right). \tag{7}$$

Next, on the one hand we have that $\frac{r_1 r_2}{r_1 - r_2} = \frac{\omega^2}{2\sqrt{b^2 - \omega^2}}$ and on the other, using Taylor expansions for the exponentials we get

$$e^{r_2 t} - e^{r_1 t} \approx (r_2 - r_1)t + (r_2^2 - r_1^2)\tfrac{t^2}{2} = -2\sqrt{b^2 - \omega^2}t + 2b\sqrt{b^2 - \omega^2}t^2.$$

Therefore

$$\dot{\theta}(t) \approx \omega^2\theta_0(bt^2 - t) = \tfrac{\delta g}{\rho l}\theta_0(bt^2 - t), \text{ for } t \text{ small}. \tag{8}$$

Notice how the damping parameter $b = \frac{\beta}{2\rho l^2}$ only affects the terms of order 2, so it is negligible near $t = 0$. Thus, by taking $\delta < \rho$ we are actually making the initial velocities $\dot{\theta}(t)$ *smaller* than what they otherwise would be

if $\delta = \rho$. Notice that this effect can not be achieved by having a larger b. We can interpret this as adding air resistance, which (akin to static vs dynamic friction) is larger when the pendulum is not moving.

Remark 4.1 (Critically damped case $b = \omega$)**.** It can be checked that when $b = \omega$, the general solution of (6) is of the form $\theta(t) = (c_1 t + c_2)e^{-bt}$. Imposing $\theta(0) = \theta_0 < 0$ and $\dot{\theta}(0) = 0$, we get $c_2 = \omega$ and $c_1 = \theta_0\omega$. Thus,

$$\dot{\theta}(t) = -\omega^2\theta_0 t e^{-bt} \approx -\omega^2\theta_0 t(1 - bt) = \omega^2\theta_0(bt^2 - t),$$

so we arrive to the same result.

We summarize the previous discussion in the following proposition.

Proposition 1. Consider the equations of motion of a damped pendulum of length $l > 0$ and inertial mass ρ:

$$\rho l^2\ddot{\theta}(t) = -\delta g l\theta(t) - \beta\dot{\theta}(t),$$

where we have introduced the virtual mass $\delta < \rho$ and made the small angle approximation $\sin\theta \approx \theta$. If we assume the initial conditions $\theta(0) = \theta_0 < 0$ and $\dot{\theta}(0) = 0$, then denoting $b = \frac{\beta}{2\rho l^2}$ we have that

$$\dot{\theta}(t) \approx \tfrac{\delta g}{\rho l}\theta_0(bt^2 - t), \text{ for } t \text{ small.} \tag{9}$$

5. Validation of the aerodynamic model with real recordings

In this section, we seek to study how the speed and size of textiles affect their motion and the optimal value ot their physical parameters within the cloth model. We will see that the aerodynamics model can capture very accurately dynamic recordings of the textiles under different circumstances, but at the cost of making the parameters depend on size and speed of movement. Hence we will develop a predictive and *a priori* formula for the value of the cloth's physical parameters, based on the fabric, its size and an estimation of the speed at which it will move.

5.1. Cloth's materials and recording setting

We employ four cloth materials (see Figure 1) and two different sizes: DIN A3 (0.297 x 0.420 m with area 0.1247 m^2) and DIN A2 (0.42 x 0.594 m with area 0.2495 m^2). Before performing the experiments they were ironed

to remove all considerations of plasticity from the validation process. In the table of Figure 1 we can also see the density of all the fabrics and some typical examples of garments made from them.

Fabric	**Density** $(\mathbf{kg \cdot m^{-2}})$	**Examples**
Polyester	0.1042	Silk-like
Wool	0.1804	Formal suit
Denim	0.3046	Jeans
Stiff-cotton	0.3046	Sack

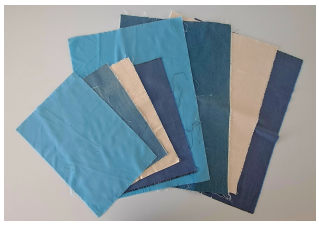

Figure 1: Density and examples of all the materials used in the experiments. In the picture we can see all the fabrics (size A3 and A2) employed. From left to right we have: polyester, denim, cotton and wool.

To record the motion of the textiles we use *Motion Capture Technology*, i.e. a system of cameras detects and tracks reflective markers that are hooked on the cloth. These markers, with a diameter of 3 mm and a weight of 0.013 g reflect infrared light, so the cameras are able to follow their motion through space. We use hardware and software from the manufacturer *NaturalPoint Inc*: five *Optitrack Flex 13* cameras surround the scene we wish to record and afterwards the recordings are processed with the software *Motive*. This combination of software and hardware offers sub-millimeter marker precision, in most applications less than 0.10 mm according to the manufacturers. For more details on the recording setting, see [31].

Remark 5.1 (Recorded coordinates of the fabrics)**.** For the A2 textiles we use 20 reflective markers, whereas for the A3 ones 12 are used. In both cases, the makers are placed equidistantly in order to obtain a faithful representation of the dynamics of the fabrics. Notice that the positions of the markers define a natural meshing of the recorded fabric. We denote the sequence of positions given by following the reflective markers with a frame every $\Delta t = 0.01$ seconds of the recorded fabric's nodes by $\{\boldsymbol{\phi}^0, \boldsymbol{\phi}^1, \dots, \boldsymbol{\phi}^m\}$. Thus, $\boldsymbol{\phi}^n \in \mathbb{R}^{12\times 3}$ (for A3) or $\boldsymbol{\phi}^n \in \mathbb{R}^{20\times 3}$ (for A2).

5.2. Cloth movements: shaking and twisting

The fabrics are recorded while being manipulated by a human and hence every movement will have its own unique variabilities. Every motion was

recorded twice on different days: **repetition I** with a special hanger and **repetition II** with bare hands. We record 2 types of motion:

1. Shaking: the cloth is held by its 2 upper corners and shaken back and forwards (approximately 15 to 20 cm) with respect to the x-axis (see Figure 2a).

2. Twisting: the cloth is also held by its 2 upper corners and rotated multiple times (approximately 30 degrees) with respect to the z-axis (see Figure 2b).

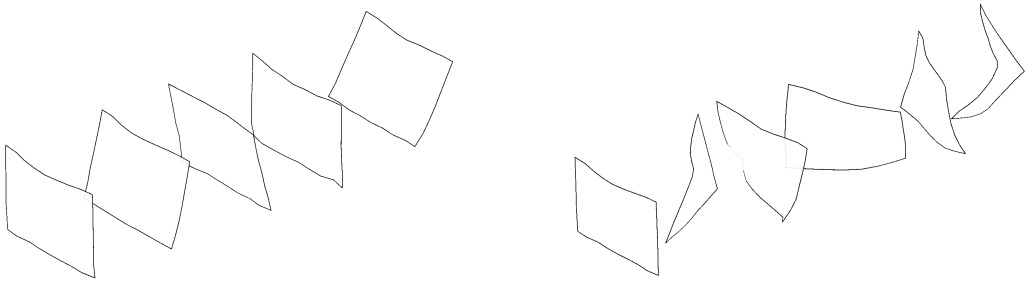

(a) Shaking motion (left to right): the cloth is shaken back and forwards.

(b) Twisting motion (left to right): the cloth is rotated with respect to the z-axis back and forth.

Figure 2: Schematic motion sequences for the shaking and twisting movements.

Each motion lasts approximately 15 seconds (with a frame every $\Delta t = 0.01$ seconds) and is performed at two different speeds: *slow* and *fast*. In Table 2 we can see some values of average speeds ($\mathrm{m} \cdot \mathrm{s}^{-1}$) comparing fast and slow motions for the twisting movement of the A2 textiles. We can of course observe some variability but overall the speeds are maintained pretty consistently.

Since we record the motion of eight (four size A3 and four size A2) textiles, the fabrics are shaken and twisted at two different speeds and we carry out two repetitions for each recording (hanger vs bare-hands), we end up with a total of 64 different recordings. For access to the raw recordings, see [31].

5.3. Parameter fitting methodology

Not all physical parameters are equally influential in the recorded motions described in this section (see [32]), therefore we will be adjusting only 2

Material	Slow I	Fast I	Slow II	Fast II
Polyester	0.081	0.250	0.090	0.247
Wool	0.093	0.228	0.085	0.256
Denim	0.092	0.308	0.090	0.292
Stiff-cotton	0.091	0.218	0.107	0.246

Table 2: Average velocities ($\mathrm{m}\cdot\mathrm{s}^{-1}$) for the twisting motion of the A2 textiles. We display the speeds for the first repetition (with a hanger) and for the second (with bare hands).

parameters: the damping parameter α and the virtual mass δ. All other physical parameters (e.g. bending) are set to 0 except for the cloth density ρ which is set to its corresponding value of Table 1. For a justification of this choice of parameters, see [11].

In order to compare recordings and simulations, we integrate numerically Equation (1) using the recorded trajectories of the two upper corners (i.e. the two markers placed at the top corners), and get a simulated sequence $\{\boldsymbol{\varphi}^0(\delta,\alpha), \boldsymbol{\varphi}^1(\delta,\alpha), \ldots, \boldsymbol{\varphi}^m(\delta,\alpha)\}$ (where $\boldsymbol{\varphi}^0 = \boldsymbol{\phi}^0$) for each value of the two parameters δ and α. For more details on how this numerical integration is performed, see [11]. Moreover, for the simulations we consider a refinement of the initial meshes given by the markers: for the A3 case we employ a 5×7 meshing and for the A2 case a 7×9 one. In order to obtain the optimal value of α and δ, we minimize the mean along time of the absolute error:

$$\bar{e}(\delta,\alpha) = \tfrac{1}{m}\sum_{n=1}^{m} e_n(\delta,\alpha) = \tfrac{1}{m}\sum_{n=1}^{m} \sqrt{||\boldsymbol{\varphi}^n(\delta,\alpha) - \boldsymbol{\phi}^n||^2_{\mathbf{M}}}, \tag{10}$$

where $||\cdot||_{\mathbf{M}}$ is the norm induced by the mass matrix $\mathbf{M}$, i.e. $||\mathbf{x}||^2_{\mathbf{M}} = \mathbf{x}^\intercal \cdot \mathbf{M} \cdot \mathbf{x}$. The use of the mass matrix $\mathbf{M}$ gives a greater weight to error in nodes limiting larger elements, and a smaller weight to error in nodes limiting smaller elements. From a mathematical viewpoint, we are estimating the integral of the error function e_n over the piece of cloth as its Riemann sum defined by the selected meshing.

As metrics to evaluate the fitting of the model, we use the previously introduced **absolute error**:

$$e_n(\delta,\alpha) = \sqrt{||\boldsymbol{\varphi}^n(\delta,\alpha) - \boldsymbol{\phi}^n||^2_{\mathbf{M}}}, \tag{11}$$

and the following **spatial standard deviation**:

$$s_n(\delta, \alpha) = \sqrt{\mathrm{Var}\left(||p_i^n(\delta, \alpha) - \hat{p}_i^n||_{\mathbb{R}^3}\right)}, \tag{12}$$

where $p_i^n(\delta, \alpha, \mu)$ (resp. $\hat{p}_i^n$) is the position at time $t_n = n\Delta t$ of the ith node of the simulated cloth (resp. recorded textile) and the variance is taken along all the nodes $i = 1, \ldots, N$ of the mesh. This metric gives us an idea of the spatial distribution of the errors on the mesh.

Remark 5.2. As a consequence of fast and abrupt movements, some of the markers in the recorded vector $\boldsymbol{\phi}^n$ disappear for small amounts of time, in those cases, they are simply excluded from the computation of the errors (no interpolation is performed). In the case that a marker k was missing for a given frame n, we just omit its corresponding coordinate from the computation of the norm $||\cdot||_{\mathbf{M}}$ for that given frame, i.e. we use a reduced diagonal sub-matrix of $\mathbf{M}$ by removing the kth row and column. On the other hand, since the simulated cloth has a finer resolution than the recording, we must only use the subsample of the simulated nodes that coincide with the recorded ones to compute the error metrics.

5.4. Optimization results

To find the optimal value of the parameters, we minimize the average on time of the absolute error (10) by performing a sweep search on the space $(\delta, \alpha) \in [0, \rho] \times [0, 4\rho]$. The lower and upper limits are selected based on physical (they have to be positive) and empirical (if they are too large they drag the clothes too much, as if they were underwater) considerations. We found this optimization method to be faster and more robust that other more sophisticated minimization algorithms used e.g. in [11]. This is likely due to the fact that the function we are trying to minimize is not completely smooth (see Figure 3) and hence has many local minima.

The optimal value of the parameters for repetition I (with a special hanger), together with the defined error metrics can be seen in Table 3.

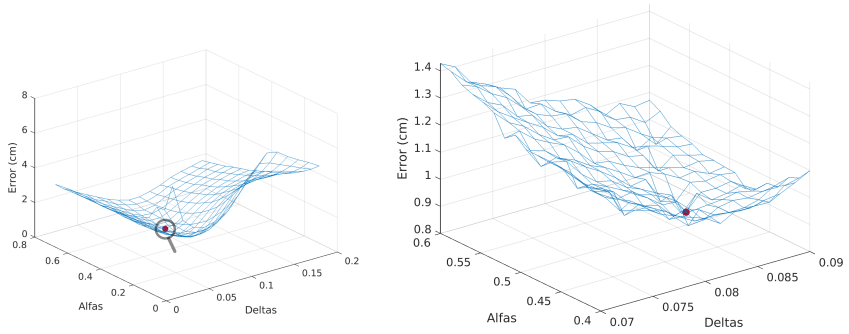

Figure 3: Surface plot of the error function $\bar{e}(\delta, \alpha)$ for the fast shaking motion of A2 wool (left) and a close-up near the detected minimum (right). Notice the presence of noise in the close-up. The maximum corresponds to the non-physical case of a very low value of δ, i.e. no gravity.

Material	Size	Motion	Speed	Error	STD	Alpha	Delta
"Stiff-cotton"	"A2"	"Shake"	"Slow"	0.433	0.661	0.336	0.16
"Stiff-cotton"	"A2"	"Shake"	"Fast"	0.516	0.77	0.657	0.181
"Stiff-cotton"	"A2"	"Twist"	"Slow"	0.324	0.533	0.496	0.222
"Stiff-cotton"	"A2"	"Twist"	"Fast"	0.562	0.918	0.817	0.284
"Stiff-cotton"	"A3"	"Shake"	"Slow"	0.268	0.711	0.336	0.201
"Stiff-cotton"	"A3"	"Shake"	"Fast"	0.336	0.798	0.817	0.222
"Stiff-cotton"	"A3"	"Twist"	"Slow"	0.267	0.711	0.336	0.243
"Stiff-cotton"	"A3"	"Twist"	"Fast"	0.315	0.823	0.898	0.263
"Wool"	"A2"	"Shake"	"Slow"	0.511	0.836	0.199	0.058
"Wool"	"A2"	"Shake"	"Fast"	0.991	1.485	0.437	0.082
"Wool"	"A2"	"Twist"	"Slow"	0.451	0.738	0.437	0.131
"Wool"	"A2"	"Twist"	"Fast"	0.844	1.314	0.532	0.095
"Wool"	"A3"	"Shake"	"Slow"	0.289	0.747	0.294	0.082
"Wool"	"A3"	"Shake"	"Fast"	0.447	1.083	0.579	0.107
"Wool"	"A3"	"Twist"	"Slow"	0.219	0.663	0.294	0.107
"Wool"	"A3"	"Twist"	"Fast"	0.361	0.919	0.579	0.107
"Polyester"	"A2"	"Shake"	"Slow"	0.887	1.187	0.17	0.026
"Polyester"	"A2"	"Shake"	"Fast"	1.36	1.819	0.307	0.041
"Polyester"	"A2"	"Twist"	"Slow"	0.444	0.685	0.197	0.041
"Polyester"	"A2"	"Twist"	"Fast"	1.152	1.617	0.307	0.048
"Polyester"	"A3"	"Shake"	"Slow"	0.463	0.903	0.142	0.033
"Polyester"	"A3"	"Shake"	"Fast"	0.471	0.99	0.28	0.041
"Polyester"	"A3"	"Twist"	"Slow"	0.243	0.689	0.115	0.048
"Polyester"	"A3"	"Twist"	"Fast"	0.359	0.854	0.307	0.055
"Denim"	"A2"	"Shake"	"Slow"	0.628	1.138	0.416	0.139
"Denim"	"A2"	"Shake"	"Fast"	0.841	1.252	0.898	0.201
"Denim"	"A2"	"Twist"	"Slow"	0.438	0.932	0.416	0.222
"Denim"	"A2"	"Twist"	"Fast"	0.905	1.35	0.978	0.263
"Denim"	"A3"	"Shake"	"Slow"	0.306	0.877	0.416	0.181
"Denim"	"A3"	"Shake"	"Fast"	0.334	0.959	0.737	0.201
"Denim"	"A3"	"Twist"	"Slow"	0.266	0.764	0.416	0.243
"Denim"	"A3"	"Twist"	"Fast"	0.288	0.804	0.737	0.305

Table 3: Summary of results of the parameter fitting for the first repetition (with hanger) of the recordings. For all the 32 recordings we display the characteristics of the recording (fabric, size, motion and speed), and the optimal value of the fitted parameters along with their associated absolute error and spatial standard deviation.

Moreover, the optimal value of the parameters for repetition II (with bare-hands), together with the defined error metrics can be seen in Table 4.

Material	Size	Motion	Speed	Error	STD	Alpha	Delta
"Stiff-cotton"	"A2"	"Shake"	"Slow"	0.446	0.682	0.416	0.139
"Stiff-cotton"	"A2"	"Shake"	"Fast"	0.561	0.898	0.737	0.181
"Stiff-cotton"	"A2"	"Twist"	"Slow"	0.351	0.569	0.496	0.222
"Stiff-cotton"	"A2"	"Twist"	"Fast"	0.681	1.098	0.898	0.305
"Stiff-cotton"	"A3"	"Shake"	"Slow"	0.256	0.651	0.336	0.181
"Stiff-cotton"	"A3"	"Shake"	"Fast"	0.331	0.805	0.737	0.222
"Stiff-cotton"	"A3"	"Twist"	"Slow"	0.282	0.716	0.496	0.284
"Stiff-cotton"	"A3"	"Twist"	"Fast"	0.329	0.838	0.978	0.305
"Wool"	"A2"	"Shake"	"Slow"	0.559	0.842	0.247	0.07
"Wool"	"A2"	"Shake"	"Fast"	0.964	1.456	0.532	0.095
"Wool"	"A2"	"Twist"	"Slow"	0.421	0.72	0.294	0.095
"Wool"	"A2"	"Twist"	"Fast"	0.958	1.408	0.627	0.107
"Wool"	"A3"	"Shake"	"Slow"	0.316	0.796	0.342	0.082
"Wool"	"A3"	"Shake"	"Fast"	0.403	1.028	0.532	0.107
"Wool"	"A3"	"Twist"	"Slow"	0.235	0.627	0.294	0.095
"Wool"	"A3"	"Twist"	"Fast"	0.468	1.076	0.627	0.144
"Polyester"	"A2"	"Shake"	"Slow"	1.025	1.327	0.142	0.026
"Polyester"	"A2"	"Shake"	"Fast"	1.378	1.761	0.28	0.041
"Polyester"	"A2"	"Twist"	"Slow"	0.52	0.811	0.17	0.033
"Polyester"	"A2"	"Twist"	"Fast"	1.22	1.696	0.252	0.041
"Polyester"	"A3"	"Shake"	"Slow"	0.467	1.227	0.17	0.033
"Polyester"	"A3"	"Shake"	"Fast"	0.575	1.247	0.28	0.041
"Polyester"	"A3"	"Twist"	"Slow"	0.543	1.429	0.17	0.033
"Polyester"	"A3"	"Twist"	"Fast"	0.528	1.166	0.417	0.055
"Denim"	"A2"	"Shake"	"Slow"	0.512	0.978	0.336	0.139
"Denim"	"A2"	"Shake"	"Fast"	0.739	1.083	0.817	0.181
"Denim"	"A2"	"Twist"	"Slow"	0.526	0.983	0.336	0.222
"Denim"	"A2"	"Twist"	"Fast"	0.876	1.318	0.978	0.243
"Denim"	"A3"	"Shake"	"Slow"	0.376	1.194	0.416	0.181
"Denim"	"A3"	"Shake"	"Fast"	0.368	1.15	0.737	0.201
"Denim"	"A3"	"Twist"	"Slow"	0.334	0.991	0.978	0.222
"Denim"	"A3"	"Twist"	"Fast"	0.375	1.121	0.496	0.222

Table 4: Summary of results of the parameter fitting for the second repetition (with bare-hands) of the recordings. For all the 32 recordings we display the characteristics of the recording (fabric, size, motion and speed), and the optimal value of the fitted parameters along with their associated absolute error and spatial standard deviation.

To get a better grasp of the results, in Table 5 we can see the mean absolute error and the mean standard deviation averaged over: material, size, type of movement and speed for the first repetition of recordings. From the table, we can deduce that for the cloth model, the most challenging material to modelize is polyester. This is somehow to be expected because

its silk-like properties make the aerodynamic effects of air more prevalent.

Repetition I	$\bar{e}$ **(cm)**	$\bar{s}$ **(cm)**
Polyester	0.67	1.09
Wool	0.51	0.97
Denim	0.50	1.00
Stiff-cotton	0.38	0.74
A2	0.71	1.08
A3	0.33	0.83
Shake	0.57	1.01
Twist	0.47	0.89
Slow	0.40	0.80
Fast	0.63	1.11
Global	0.51	0.96

Table 5: Mean absolute error $\bar{e}$ and the mean standard deviation $\bar{s}$ with the optimal value of the parameters averaged over: fabric's material, size (A2 or A3), type of movement and speed for the first repetition of the recordings.

On the other hand, the errors are generally larger for the bigger textiles, this is again reasonable since they have double the area. The fast motions have larger errors, this is likely due to the fact that in that case aerodynamics are harder to model. Furthermore, we can see that both the shaking and twisting motions have comparable errors.

Finally, to get a sense of the goodness of the fitting for particular recordings, in Figure 4, 5 and the **supplementary video 1**, where we can see a visual comparison between two specific recordings of repetition I and their simulation with the optimal value of the parameters.

5.5. A priori forecast of α and δ

As we can see in Tables 3 and 4 , the values of α and δ that minimize the absolute error $\bar{e}$ vary substantially with respect to material, size and speed. This is fine if we only wish to validate the expressive capabilities of the model, but it becomes a problem if we want to use the model to forecast cloth behavior, e.g. in robotic applications.

Thus, we would like to find an *a priori* formula that we can use to forecast the value of these parameters without minimizing $\bar{e}(\alpha, \delta)$. Apart from

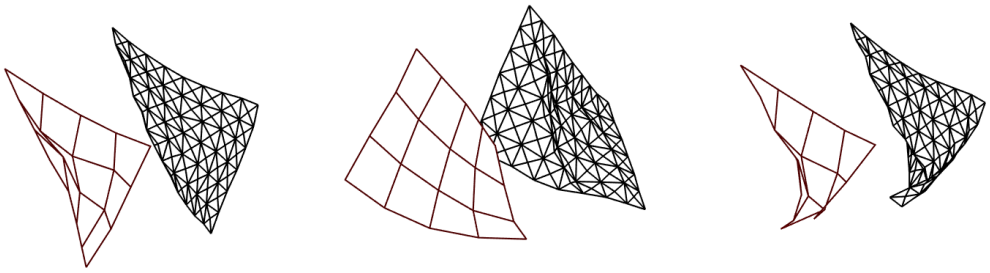

Figure 4: Three frames comparing the recorded fast twisting of A2 polyester (left) with its inextensible simulation (right). The error at the three depicted frames from left to right is 1.82, 1.73 and 1.36 cm respectively; being the average error of the whole simulation 1.15 cm.

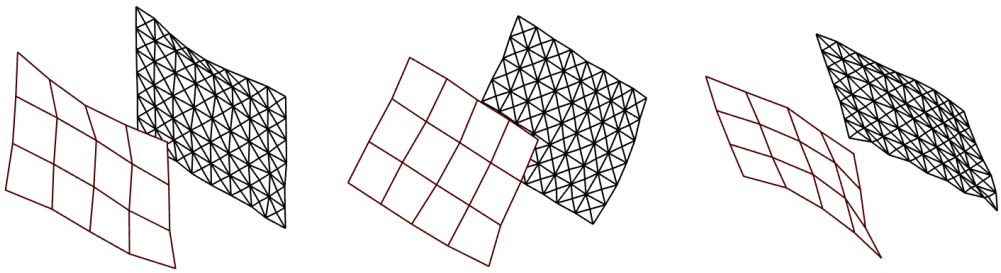

Figure 5: Three frames comparing the recorded fast shaking of A2 denim (left) with its inextensible simulation (right). The error at the three depicted frames from left to right is 0.63, 0.92 and 1.001 cm respectively; being the average error of the whole simulation 0.84 cm.

intrinsic properties of the textile (the density ρ and its size), this formula will also depend on the speed of motion of the cloth. We look then for formulas of the form:

$$\hat{\delta} = \delta_0 + \delta_1 S + \delta_2 V + \delta_3 \rho, \quad \hat{\alpha} = \alpha_0 + \alpha_1 S + \alpha_2 V + \alpha_3 \rho, \tag{13}$$

where S is a measure of size and V a measure of velocity. We have found that for our purposes using as S a normalized area of the cloth (1 for size A3 and 2 for A2) and as $[V] = \mathrm{m}^2 \cdot \mathrm{s}^{-2}$ the average (in time and over all the nodes) of 50% of the highest squared velocities gives the best results.

Now, in order to find the optimal values of $\{\delta_k, \alpha_k\}$ we minimize the

function given by:

$$\mathcal{R}(\delta_0, \ldots, \delta_3, \alpha_0, \ldots, \alpha_3) = \frac{1}{32} \sum_{r=1}^{32} \bar{e}(\hat{\delta}, \hat{\alpha})_r, \tag{14}$$

where $\bar{e}(\hat{\delta}, \hat{\alpha})_r$ corresponds to the absolute error (10) of the r-recording of either repetition I or II obtained using the values $\hat{\delta}, \hat{\alpha}$ given by Equation (13). Thus, in order to evaluate $\mathcal{R}$ for either repetition I or II, we need to perform 32 simulations. To minimize this function we employ this time the *Nelder-Mead Simplex Method*, which is a derivative-free algorithm.

We denote by $\mathcal{R}^*$ function (14) evaluated at the optimal values of $\{\delta_k, \alpha_k\}$. For comparison let $\mathcal{E}^*$ be the mean of the optimal errors found in the previous subsection (which will be by definition smaller), i.e. the mean of the errors displayed in Table 3 or 4.

	Repetition I	**Repetition II**
δ_0	-0.0223	-0.0359
δ_1	-0.0178	-0.0117
δ_2	0.0714	0.0780
δ_3	0.7664	0.7890
α_0	0.2082	0.2155
α_1	-0.1481	-0.1711
α_2	1.1804	1.4410
α_3	1.7440	1.9387
$\mathcal{R}^*$	0.578 cm	0.646 cm
$\mathcal{E}^*$	0.516 cm	0.560 cm

Table 6: Optimal values of the parameters δ_k, α_k for both repetition I (with hanger) and II (with bare hands) obtained by minimizing the function $\mathcal{R}$.

In Table 6 we can see the optimal values of the parameters $\{\delta_k, \alpha_k\}$ along with the mean absolute errors averaged over all recordings for both repetition I (with hanger) and II (with bare hands). Notice that the error $\mathcal{R}^*$ is comparable to $\mathcal{E}^*$ and hence the fitting is quite accurate. Moreover, both sets of parameters are estimated independently for the two repetitions and have very similar values and the same signs, which shows that they are significant and have a consistent meaning.

5.6. Prediction of the behavior of cloth in an unseen scenario

In this subsection we test the previously found a priori formulas with *out of sample* (i.e. unseen) recordings. This means that the recordings we will use have not been used to estimate the formula coefficients, i.e. we do not perform any optimization and just compute the value of the cloth parameters with the formulas

$$\begin{aligned} \hat{\delta} &= -0.0223 - 0.0178S + 0.0714V + 0.7664\rho, \\ \hat{\alpha} &= +0.2082 - 0.1481S + 1.1804V + 1.7440\rho, \end{aligned} \tag{15}$$

where, as before, ρ is the density of the cloth, S is a normalized area measure (in this scenario it will be always 2), V is the average (in time and over all the nodes) of 50% of the highest squared velocities. We have used the coefficients of the first repetition in Table 6.

Now we describe briefly the new recordings. For more details and how to simulate this scenario with our cloth model, see [12].

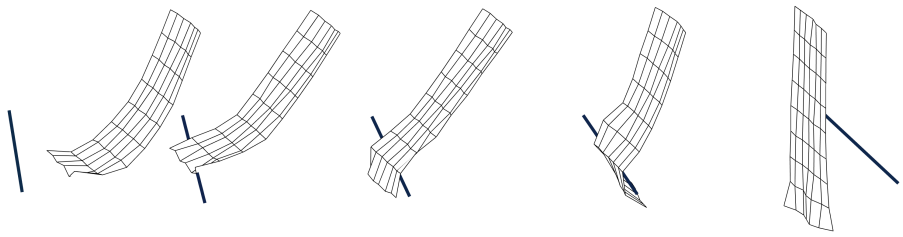

Figure 6: the cloth is held by its two upper corners and then is hit repeatedly with a long thin stick.

Hitting recordings: the A2 cloths were grasped by two corners, and held suspended in the air with the long sides perpendicular to the floor. Then, they were hit repeatedly with a long thin stick. Each textile is hit four times at various locations of the cloth with varied strengths and speeds as shown in Figure 6. The recordings lasted between 12 and 18 seconds. The stick has a length of 75 cm and a diameter of 1.5 cm. Apart from the 20 markers already attached to the A2 cloths, two new markers are put at both ends of the stick to record its trajectory. In total, we have 4 recordings, one for each material.

In Table 7 we show a comparison of the estimated parameters $\hat{\delta}, \hat{\alpha}$ using formulas (15), with the optimal ones δ^*, α^* taken from [12]. Their respective absolute errors are also displayed.

Material	$\hat{\delta}$	δ^*	$\hat{\alpha}$	α^*	$\hat{e}$ **(cm)**	e^* **(cm)**
Polyester	0.032	0.025	0.26	0.28	1.59	1.44
Wool	0.099	0.078	0.54	0.50	1.96	1.39
Denim	0.198	0.127	0.82	0.67	1.76	0.98
Stiff-cotton	0.203	0.165	0.90	0.89	1.34	1.07

Table 7: Comparison of the parameters $\hat{\delta}, \hat{\alpha}$ obtained with the formulas (15) with the optimal ones δ^*, α^* (taken from [12]) along with their respective mean absolute errors.

We can see that, in general, the formula gives a very reasonable estimation of the parameters (especially for α), but it always overestimates δ and this causes the errors $\hat{e}$ to be somehow larger. Still the results are very accurate considering how challenging this scenario is and that the parameters are computed with an *a priori* equation without any optimization at all.

5.7. Discussion of the results

Our previous work [11] already indicated that the inextensible model together with the virtual mass parameter was capable of reproducing faithfully the shaking motion of A3 textiles. This work establishes more firmly this proposition: we have evaluated a varied, more exhaustive set of recordings than in our prior work. We have enlarged the set of motions by adding a new twisting movement and a larger cloth size (A2). This has resulted in a significant resource for validating any aerodynamic model for cloth motion: two sets of 32 recordings (each cloth being recorded two times on different days: with a special hanger or bare hands) each lasting approximately 15 seconds. We have estimated the optimal values of the physical parameters and the model has achieved very low mean errors (less than 1 cm) and standard deviations (see Tables 3, 4 and 5), even for the A2 textiles and fast motions.

Notice that our experimental results validate the introduction of the virtual mass parameter δ: for each textile and motion, the optimal value of δ identified from the motion capture data is lower than the corresponding ρ for the particular fabric, indicating that a lift force is indeed opposing gravity, exactly as the theory predicts.

Finally, we have found a predictive formula in order to obtain *a priori* estimates for the values of the parameters α and δ of the model. This formula depends on the density of the textile, its size and more importantly its speed. The formula was found to yield values for the parameters which in turn still give rise to very low absolute errors. As an immediate application, we made use of this *a priori* formula to test its accuracy with an unseen collision scenario. This was done with very satisfactory results (see Table 7) by using (15) to compute an estimate of the physical parameters of the clothes without performing any optimization at all.

6. Dynamic flattening of the A2 cloths with a robot

In this section we give a non-trivial robotic application (a dynamic flattening task) of having an accurate model that can predict the future behavior of cloth. We also show that the *sim-to-real* gap of our cloth model is small, meaning that what we plan in simulation can be executed by a real robot with essentially the same results. This success indicates that the theoretical model not only aligns with the experimental data but also effectively bridges the *sim-to-real* gap, confirming the validity of our aerodynamic correction. All real and simulated videos described in this section are shown in the **supplementary video 2** to this manuscript.

The task we consider is that of laying the A2 clothes flat onto a table in a dynamic fashion. The fabrics are grasped by their two upper corners, and held suspended in the air with their short sides perpendicular to the table. In order to lay them down dynamically, we designed a trajectory –see Figure 7, we only control the 2 upper corners– which has the following parts:

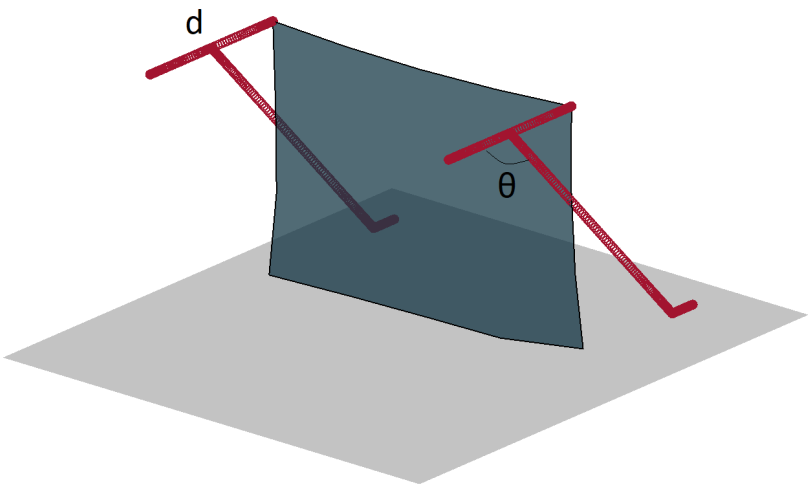

Figure 7: Trajectory used for the dynamic flattening of the textiles. The two upper corners are controlled in the following way: they are moved back in a line a distance d and then forward a distance $d/2$. Afterwards, they are moved downwards diagonally until they reach the table with a certain angle θ. Finally, the 2 corners are moved forward parallel to the table 5 cm.

1. The corners are moved back in a straight line (parallel to the table) a certain distance d.

2. Then, the corners are moved forward in a straight line (again parallel to the table), this time a distance $d/2$.

3. Next, the corners are moved downward until the table is reached in a diagonal straight line with a certain angle θ with respect to the table plane.

4. Finally, the 2 corners are moved forward parallel to the table 5 cm.

We remark that the two trajectories for the upper corners are identical up to translation.

Remark 6.1 (Parametrization of the trajectories)**.** We parametrize all straight segments with equations of the form:

$$\gamma(t) = q_0 + \tau(t)\vec{w}, \quad \tau(t) = -\tfrac{4}{3}\tfrac{v}{T^2}t^3 + 2\tfrac{v}{T}t^2, \tag{16}$$

where $\vec{w}$ is a unit vector, v is parameter controlling the maximum velocity at which we wish to move, $T = \frac{3}{2}\frac{d}{v}$ and d is the length of the parametrized segment. Notice that the $\mathcal{C}^\infty$ parametrization (16) has the following two desirable properties: (i) $||\gamma(T)-\gamma(0)|| = d$ since $\gamma(0) = q_0$ and $\gamma(T) = q_0+d\vec{w}$ and (ii) $||\gamma'(t)|| = \tau'(t),\ t \in [0,T]$ is a parabola with maximum value v at $T/2$ such that $||\gamma'(0)|| = ||\gamma'(T)|| = 0$. This parametrization allows easy control of the length d and velocity v at which we go along the segment, while at the same time being robot-friendly since it is smooth and starts and ends with zero velocity.

To simulate the collision between the cloths and the table we follow the method described in [12]. Moreover, the real table is covered with a tablecloth so that the friction between the textiles and the table is very high. Then, in order to execute the trajectories shown in Figure 7 with the simulator, we only need an initial mesh for each textile (cotton, denim, polyester and wool) and the value of its physical parameters shown in Table 1. The initial mesh is simply taken from the first frame of the recordings described in the previous section, whereas for the physical parameters of the cloths we use the a priori formulas (15). Thus, ρ is set to its corresponding value according to Figure 1, $S = 2$ (since we are only dealing with the A2 textiles) and V is computed using the full planned trajectory of the 2 upper corners (i.e. we compute it as the average in time and between the two corners of 50% of their highest squared velocities). For each of the 4 textiles we select reasonable but different values of d (around 35cm) and θ (around 135 degrees) that result in successful flattening motions in simulation. Each motion lasts 3 seconds and is indeed very dynamic, with peak velocities of $0.9\ \mathrm{m} \cdot \mathrm{s}^{-1}$.

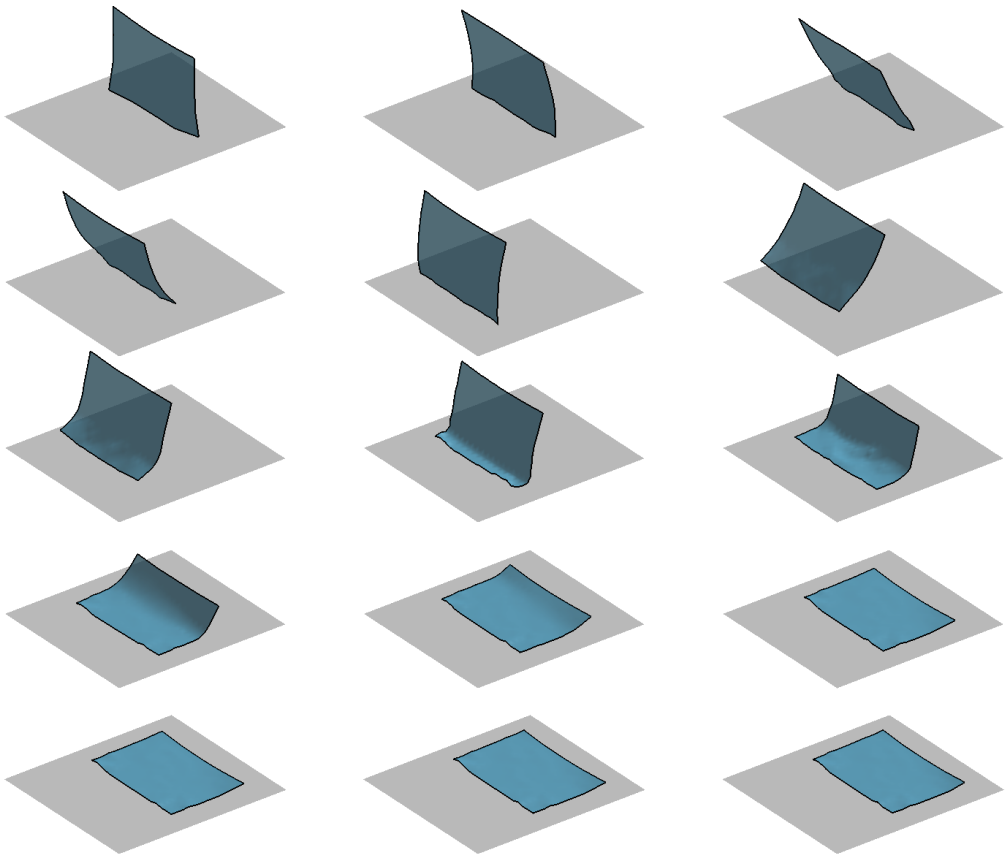

Figure 8: Equally spaced frames of the simulation for the denim sample using the trajectory shown in Figure 7. For the physical parameters of the cloth we use the a priori formulas (15).

In Figure 8 we show 15 equally spaced frames of the simulation for the denim sample using the trajectory shown in Figure 7. As we can see in the Figure (and in the supplementary video), the motion results in a successful flattening of the simulated textile.

Figure 9: Equally spaced frames taken from a video recording of the WAM robotic arm performing a dynamic laying of the real A2 denim textile onto the table. The trajectory of the end-effector is exactly the same of that used in the simulation. Comparing visually with the simulated frames of Figure 8 we can observe that the qualitative state of the cloth is very similar in both the real and simulated frames.

Next, we reproduce the exact same trajectory of the two upper corners with a robotic arm and the real textile. In order to do so, we use a Barrett WAM robot arm with a hanger (see Figure 9). We simply average the planned trajectory of the 2 upper corners and then we compute the inverse kinematics so that the end-effector of the WAM follows precisely this averaged curve. In order to have the same initial conditions as in the simulated environment, we must only set the robotic arm so that the textile is at the same height with respect to the table as in the simulation. The result (shown in 15 equally spaced frames in Figure 9) is also a successful and smooth dynamic laying of the real A2 denim textile onto the table. Notice that even though we did not perform any fine-tuning or optimization whatsoever, if we compare the simulated frames with the real ones, the qualitative state of the cloth is very similar in both the real and simulated frames.

In Figure 10 we display the final state of the cloths for all 4 materials and 3 independent trials performed with the robot. As we can see in the figure, all trials achieve a successful flattening of the textiles (see also the supplementary video for one recording with each real fabric) with the exception of the polyester sample. This fabric presents (different, somehow small) wrinkles in all 3 trials.

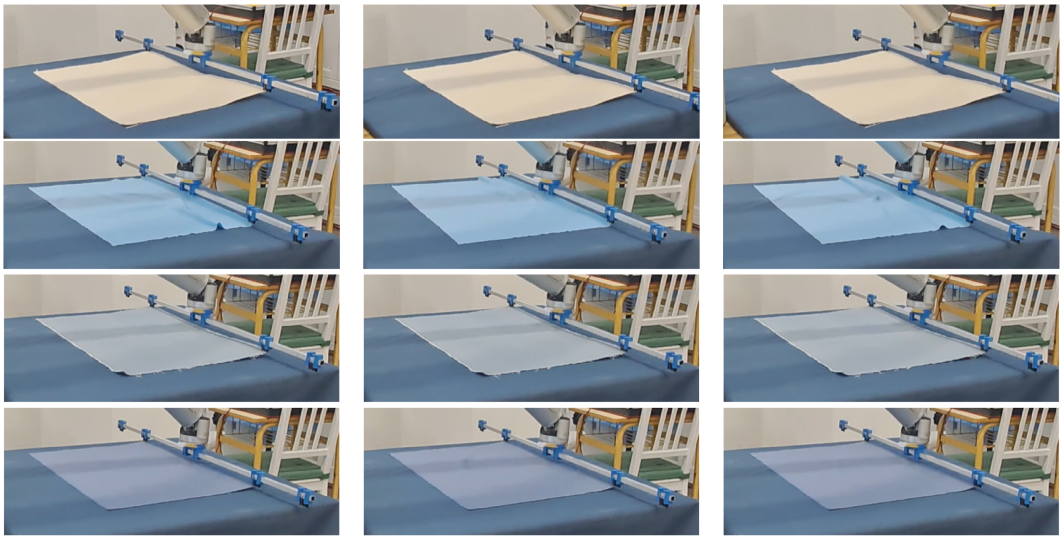

Figure 10: Final state of the cloths for all 4 materials and 3 independent trials performed with the robot. All trials (with the exception of polyester) achieve a successful full flattening of the textiles without any wrinkles.

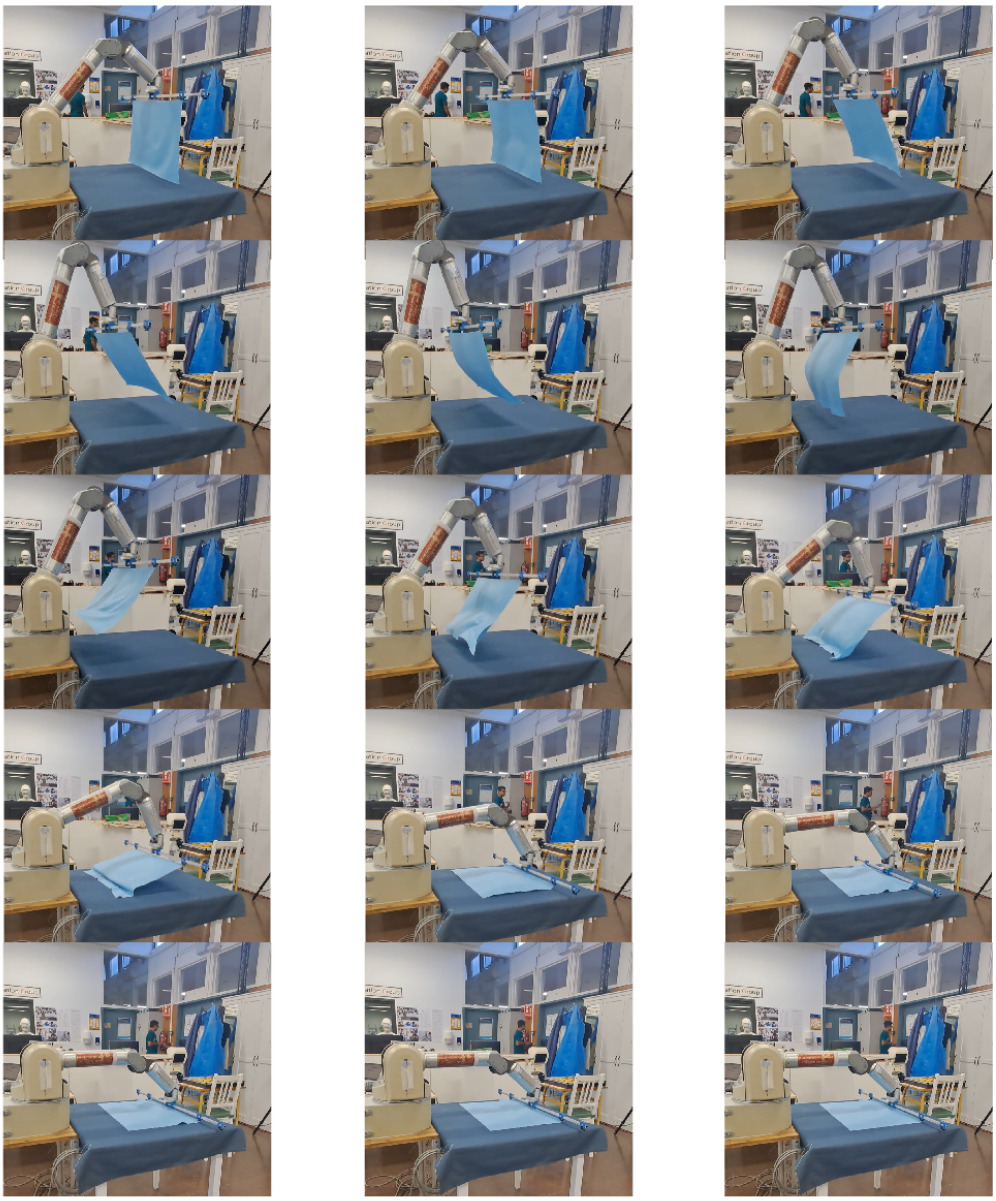

Figure 11: Equally spaced frames taken from a video recording of the WAM robotic arm performing a dynamic laying of the real A2 polyester textile onto the table. The material is so light (25 grams) that the motions at these speeds are already turbulent near the boundaries (see the 8th and 9th frame). Moreover, when the fabric approaches the table, a chamber of air is created between the cloth and the table which causes the cloth to 'spring' forwards.

In order to understand what is happening, in Figure 11 we display 15 equally spaced frames of the whole motion with the real polyester textile. Our best guess is that the polyester material is so light (the A2 sample weights 25 grams) that the motions at these speeds are already turbulent (see the 8th and 9th frame of Figure 11). Moreover, when the fabric approaches the table, a chamber of air is created between the cloth and the table which then causes the cloth to 'jump' forwards (this can be very clearly seen in the video).

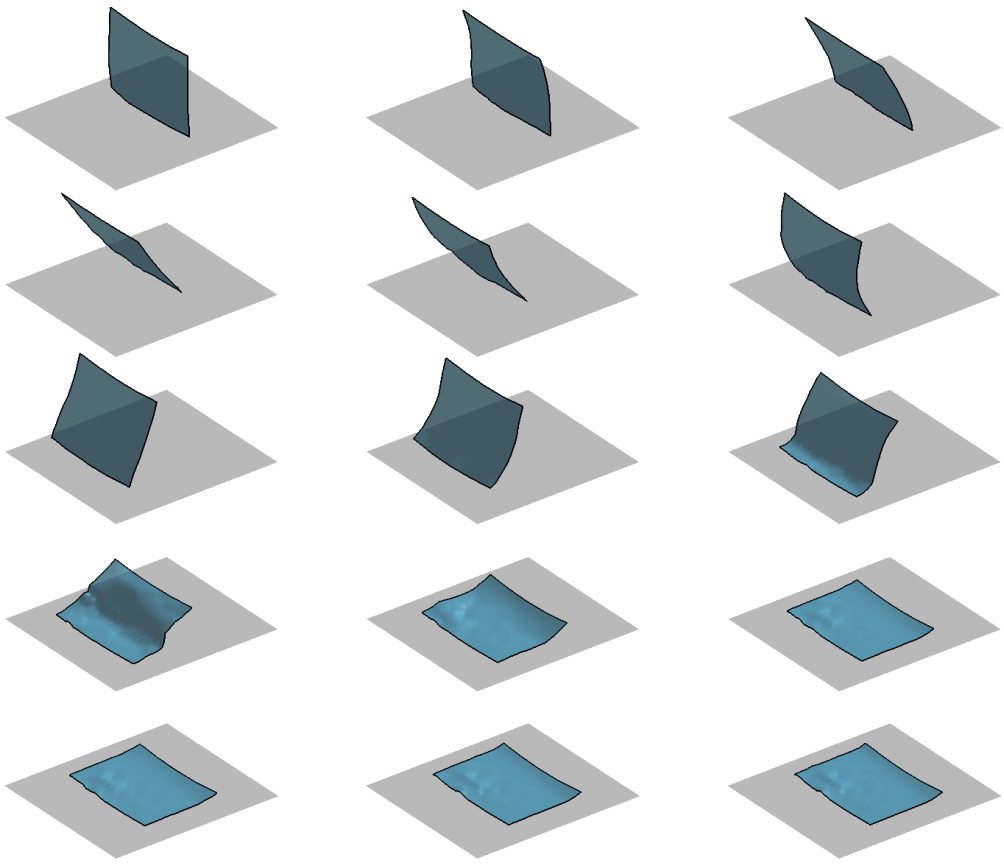

Figure 12: Equally spaced frames of the simulation for the polyester sample using the trajectory shown in Figure 7. Although the states of the real polyester sample are followed quite faithfully (see Figure 12), all the fine aerodynamic details are not being simulated with the same level of detail and hence the final result of the simulation differs from that of the real fabric.

Finally, in Figure 12 we display 15 frames of the corresponding simulation. There we can see that, although the states of the polyester are followed quite faithfully, all the fine aerodynamic details are not being simulated properly and hence the final result differs. This was also somehow expectable since polyester was the hardest material to simulate with the cloth model (see Table 5).

6.1. Discussion of the results

In this section we have given a non-trivial robotic application of many of the ideas presented in this work. Using the cloth model described in Section 3, we were able to design successful dynamic trajectories to flatten the four real A2 textiles shown in Figure 1. The trajectories were first benchmarked in simulation (see Figure 7) and then executed without any fine-tuning with a real robot. In order to do so, we used the aerodynamic formulas (15), which allowed us to estimate the values of the physical parameters $\alpha > 0$ and $\delta > 0$ of each individual fabric without performing any optimization; just by using its density, its size and an estimation of the speed at which it was going to move. The *sim-to-real* transfer of our simulator for this scenario was proven to be very small, since what happened in simulation also took place in real life (see Figure 10). The main exception to this was the polyester sample: being the lightest textile it was the most affected by aerodynamic forces. For this material our cloth model was not able to properly simulate all the fine aerodynamic details needed to flatten the fabric 100% without any wrinkles (compare Figure 11 with Figure 12). Still, the final result is quite satisfactory (see Figure 11) taking into consideration how challenging the motion of this textile can be for any cloth simulator since aerodynamics and friction play leading roles.

7. Conclusions

In this work we have analyzed a simple and practical method to account for the effect of aerodynamic forces on the motion of a piece of cloth at low speeds, i.e. speeds at which air can be approximated as an incompressible, inviscid fluid. This method, justified by the study of a dampened pendulum (see Proposition 1), consists in replacing the cloth density ρ with a correction factor $\delta \leq \rho$. This correction is applied to compute the gravitational force acting on the cloth, rescaling its weight by the factor $\delta/\rho \leq 1$. This aerodynamic correction factor, δ, can be applied to any mechanical model of cloth

to account for aerodynamic forces in the simulation of cloth motion. It has no additional computational cost if the mechanical model already involves calculating the weight of the cloth. Hence, it is an ideal candidate to be used for real-time simulation and control, making it well suited for robotics applications where efficiency is key. The parameter δ reflects the friction between the moving cloth and the surrounding air. As a result, it depends on the local structure of the cloth, such as its constituent fabric and density, as well as two global characteristics: the size of the cloth and its speed relative to its surrounding air.

We have applied this aerodynamic correction factor to the inextensible cloth model developed in [11, 12]. To validate the aerodynamics model and the value of the parameter δ, two different tasks involving cloth motion are recorded with a Motion Capture System. Each task is performed multiple times with cloth pieces varying in fabric type, density, size and speed of motion, totaling 64 recordings of about 15 seconds each. For each motion, the positions of key points on the cloth are recorded throughout the task's execution. The resulting dataset is then compared to simulations of the same tasks using the inextensible cloth model incorporating the aerodynamic correction parameter δ. By minimizing the error between the simulated and real positions of the cloth, the optimal value of δ is identified. This best-fit results in a mean absolute error of less than 1 cm in the comparison between the simulated and the actual positions of the cloth's key points. Moreover, we analyze for the first time how the resulting parameter δ depends on fabric type, garment size, and speed of motion.

As a final validation of the proposed aerodynamic correction and the inextensible cloth model, we used the model to simulate a task in which aerodynamic and friction forces play a dominant role in controlling the motion of cloth: dynamically flattening size A2 textiles onto a table. The simulations are performed using the default parameter values determined in the previous tasks, with no fine-tuning. The resulting simulation is then used to directly control a robotic arm, which successfully executes the flattening task with the real textiles. The final validation also serves as a demonstration of the suitability of the inextensible cloth model proposed in [11, 12], combined with the simple aerodynamic correction δ, for robotic cloth manipulation. This includes its use as a real-time control mechanism and as a generator of simulations for training learning algorithms. The authors plan to extend the application of the model after further exploration of parameter values across

a broader range of cloth types.

8. Acknowledgements

F. Coltraro was partially supported by the ClothIRI (CSIC 202350E080) project and the RobIRI 2021-SGR-00514 AGAUR project. He is staff hired under the Generation D initiative, promoted by Red.es, an organisation attached to the Ministry for Digital Transformation and the Civil Service, for the attraction and retention of talent through grants and training contracts, financed by the Recovery, Transformation and Resilience Plan through the European Union's Next Generation funds. Jaume Amorós and Maria Alberich–Carramiñana have been partially supported by the projects PID2019-103849GB-I00 and PID2023-146936NB-I00 financed by the Spanish State Agency MICIU/AEI/10.13039/501100011033 and by ERDF/EU, and by the GEOMVAP 2021-SGR-00603 AGAUR project. This research work was partially funded by the European Commission – NextGenerationEU, through MomentumCSIC Programme: Develop Your Digital Talent.

References

[1] A. Longhini, Y. Wang, I. Garcia-Camacho, D. Blanco Mulero, M. Moletta, M. Welle, G. Alenyà, H. Yin, Z. Erickson, D. Held, J. Borràs, D. Kragic, Unfolding the literature: A review of robotic cloth manipulation, Annual Review of Control, Robotics, and Autonomous Systems (2024).

[2] I. Garcia-Camacho, M. Lippi, M. C. Welle, H. Yin, R. Antonova, A. Varava, J. Borràs, C. Torras, A. Marino, G. Alenyà, D. Kragic, Benchmarking bimanual cloth manipulation, IEEE Robotics and Automation Letters 5 (2020) 1111–1118.

[3] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, Y. Mezouar, Robotic manipulation and sensing of deformable objects in domestic and industrial applications: a survey, The International Journal of Robotics Research 37 (2018) 688–716.

[4] L. Younes, Shapes and Diffeomorphisms, Applied mathematical sciences, Springer, 2019.

[5] H. Yin, A. Varava, D. Kragic, Modeling, learning, perception, and control methods for deformable object manipulation, Science Robotics 6 (2021) eabd8803.

[6] Y. Li, Y. Yue, D. Xu, E. Grinspun, P. Allen, Folding deformable objects using predictive simulation and trajectory optimization, Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (2015) 6000–6006.

[7] A. Longhini, M. C. Welle, Z. Erickson, D. Kragic, Adafold: Adapting folding trajectories of cloths via feedback-loop manipulation, IEEE Robotics and Automation Letters (2024).

[8] D. Blanco-Mulero, O. Barbany, G. Alcan, A. Colomé, C. Torras, V. Kyrki, Benchmarking the sim-to-real gap in cloth manipulation, IEEE Robotics and Automation Letters 9 (2024) 2981–2988.

[9] E. Grinspun, A. Hirani, M. Desbrun, P. Schröder, Discrete shells, 2003.

[10] J. Hearle, P. Grosberg, S. Backer, Structural mechanics of fibers, yarns, and fabrics, volume 1, Wiley-Interscience, 1969.

[11] F. Coltraro, J. Amorós, M. Alberich-Carramiñana, C. Torras, An inextensible model for the robotic manipulation of textiles, Applied Mathematical Modelling 101 (2022) 832–858.

[12] F. Coltraro, J. Amorós, M. Alberich-Carramiñana, C. Torras, A novel collision model for inextensible textiles and its experimental validation, Applied Mathematical Modelling 128 (2024) 287–308.

[13] J. D. Anderson, Fundamentals of aerodynamics, 5th ed., McGraw-Hill, 2011.

[14] X. Liu, L. Liu, An immersed transitional interface finite element method for fluid interacting with rigid/deformable solid, Engineering Applications of Computational Fluid Mechanics 13 (2019) 337–358.

[15] T. Mao, S. Xia, Z. Wang, Evaluating simplified air force models for cloth simulation, in: 11th IEEE International Conference on Computer-Aided Design and Computer Graphics, 2009, pp. 81–86.

[16] M. Ghalandari, S. Bornassi, S. Shamshirband, A. Mosavi, K. W. Chau, Investigation of submerged structures' flexibility on sloshing frequency using a boundary element method and finite element analysis, Engineering Applications of Computational Fluid Mechanics 13 (2019) 519–528.

[17] M. Shelley, J. Zhang, Flapping and bending bodies interacting with fluid flows, Annual Review of Fluid Mechanics 43 (2011) 449–465.

[18] J. Katz, A. Plotkin, Low-speed aerodynamics, Cambridge University Press, 2001.

[19] A. Fitt, M. Pope, The unsteady motion of two-dimensional flags with bending stiffness, Journal of Engineering Mathematics 40 (2001) 227–248.

[20] M. Argentina, L. Mahadevan, Fluid-flow-induced flutter of a flag, Proceedings of the National Academy of Sciences 102 (2005) 1829–1834.

[21] S. Alben, Dynamics of flags over wide ranges of mass and bending stiffness, Physical Review Fluids 7 (2022) 013903.

[22] W.-X. Huang, H. J. Sung, Three-dimensional simulation of a flapping flag in a uniform flow, Journal of Fluid Mechanics 653 (2010) 301–336.

[23] H. Zhu, Q. Sun, J. Tao, P. Tan, Z. Chen, M. Dehmer, G. Xie, Fluid-structure interaction simulation and accurate dynamic modeling of parachute warhead system based on impact point prediction, IEEE Access 9 (2021) 104418–104428.

[24] W.-X. Huang, S. J. Shin, H. J. Sung, Simulation of flexible filaments in a uniform flow by the immersed boundary method, Journal of Computational Physics 226 (2007) 2206–2228.

[25] R. Mittal, G. Iaccarino, Immersed boundary methods, Annual Review of Fluid Mechanics 37 (2005) 239–261.

[26] Z. Xu, B. Kaszás, M. Cenedese, G. Berti, F. Coletti, G. Haller, Data-driven modelling of the regular and chaotic dynamics of an inverted flag from experiments, Journal of Fluid Mechanics 987 (2024) R7.

[27] L. Ling, M. Damodaran, R. K. Gay, Aerodynamic force models for animating cloth motion in air flow, The Visual Computer 12 (1996) 84–104.

[28] D. House, D. E. Breen, Cloth modeling and animation, CRC Press, 2000.

[29] M. Bergou, M. Wardetzky, D. Harmon, D. Zorin, E. Grinspun, A quadratic bending model for inextensible surfaces, in: Proceedings of the Fourth Eurographics Symposium on Geometry Processing, SGP '06, Eurographics Association, 2006, pp. 227–230.

[30] O. C. Zienkiewicz, R. L. Taylor, J. Z. Zhu, The Finite Element Method: Its Basis and Fundamentals, Sixth Edition, Butterworth-Heinemann, 2005.

[31] F. Coltraro, J. Borràs, M. Alberich-Carramiñana, C. Torras, Tracking cloth deformation: A novel dataset for closing the sim-to-real gap for robotic cloth manipulation learning, The International Journal of Robotics Research (2025) 02783649251317617.

[32] R. Aranyó Llach, Validación del modelo de tela del proyecto Clothilde mediante la simulación y el cálculo numérico, Universitat Politècnica de Catalunya, Degree Thesis, 2020.