Inverse Elastic Cloth Design with Contact and Friction

Romain Casati, Gilles Daviet, Florence Bertails-Descoubes

Inverse Elastic Cloth Design with Contact and Friction

Romain Casati Gilles Daviet Florence Bertails-Descoubes

INRIA and Laboratoire Jean Kuntzmann (Grenoble University, CNRS), France

Abstract

Physically based cloth modeling is classically achieved through a trial and error process. The rest (undeformed) configuration of the cloth, often represented as a 2D pattern assembly, is edited geometrically and adjusted iteratively depending on the feedback provided by a static cloth simulator, which predicts the deformed 3D shape under gravity and contacts. Matching a reference 3D shape while keeping the time of the modeling process reasonable is thus difficult, unless the user possesses advanced skills in real cloth tailoring. In contrast, in this paper we investigate a new, inverse strategy for modeling realistic cloth intuitively. Our goal is to take as input a target (deformed) 3D shape, and to interpret this configuration automatically as a stable equilibrium of a cloth simulator, by retrieving the unknown rest shape. In the presence of gravity and frictional contact, such an inverse problem formulates as an ill-posed nonlinear system subject to nonsmooth constraints. To select and compute a plausible solution, we design an iterative two-step solving process. In a first step, contacts are reduced to frictionless bilateral constraints, and starting from an as-flat-as possible pose, a unique rest pose is retrieved using the adjoint method on a regularized energy. The second step modifies this rest pose so as to project bilateral forces onto the admissible Coulomb friction cone, for each contact. We show that our method converges well in most cases towards a plausible rest configuration, and demonstrate practical inversion results on various cloth geometries modeled by an artist.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling —Physically based modeling I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: cloth modeling, frictional contact, inverse problem, nonsmooth optimization

1 Introduction

Cloth simulation has been a very active research area over the last decades, yielding compelling animations of cloth in movies that are now hardly distinguishable from reality. In parallel, the variety of garments simulated, now ranging from simple cloaks to complex multi-layered dresses, has been considerably enriched in recent years. Such a diversity in shape and motion encourages artists to create personalized garments expressing not only aesthetics and social rank of a given character, but also some more personal traits (e.g., severity, shyness, spirit), which actively contributes to the overall story telling.

However, modeling a specific virtual garment still remains a difficult and tedious task for a designer, as current modeling tools offer very little user control over the final (simulated) shape of the cloth. On the one hand, a number of sophisticated tools allow to tailor garments as one would do in real: The user designs 2D cloth patterns, then sews them together in 3D, and finally adjusts the result iteratively based on the simulation of the "drape" — that is, the *stable equilibrium* pose of cloth under external forces [Volino et al. 2005; Marvelous Designer 2010; Umetani et al. 2011]. Though useful to design realistic patterns which may then be manufactured in real, such a trial and error process remains complex and unnatural when dressing a virtual character with a precise final shape in mind (e.g., the designer should anticipate the length of the seams so as to obtain folds of a given size located in a specific area). On the other hand, many different WYSIWYG approaches, ranging from sketch-based modeling [Turquin et al. 2007] to image-based capture [Bradley et al. 2008], have provided the user with intuitive and semi-automatic tools for digitizing the 3D shape of a specific garment, from an imaginary or real reference. Such geometric tools allow the user to focus only on the final shape of the cloth, thus providing a maximal level of control over geometric details. However, although the resulting geometry generally corresponds to an equilibrium configuration of the cloth under external forces, it cannot be interpreted physically in a simple way. If one naïvely plugs such a final pose as the rest (undeformed) state parameter of a cloth simulator, sagging will occur when simulation is activated, and all the modeling efforts and time will be lost.

Figure 1: *Breathing physics into a purely geometric garment directly modeled in 3D under gravity. From the 3D mesh of the cloth and the body mesh, our method automatically infers the rest pose of a cloth simulator such that a stable equilibrium of the cloth simulator under gravity and frictional contact closely matches the input geometry. Starting from this input pose, the cloth can then be animated in a plausible manner, for instance under wind.*

In this paper we explore a new strategy for modeling virtual cloth both realistically and intuitively, by combining the fine user control offered by geometric tools together with the predictive capabilities of simulation. Our idea is to devise an automatic *inverse modeling* process able to interpret a merely geometric drape pose as a stable equilibrium of a cloth mechanical simulator. Assuming material properties (mass, stiffness, friction coefficients) of the cloth are known, our algorithm takes as input a target (deformed) 3D surface, resulting from either geometric design or capture, and automatically retrieves a plausible rest shape for the cloth (non-

necessarily flat due to seams) as well as frictional contact forces at play. When simulated with such parameters under external forces, the cloth closely matches the input pose, and may be subsequently animated.

2 Related Work

Cloth modeling and simulation has a rich history in Computer Graphics and Computer-Aided Design. Within less than three decades, practicable solutions have been proposed to design extremely complex garments and to animate them with an impressive level of details. Such remarkable advances were seamlessly integrated into the work flow of the entertainment industry, yielding compelling visual effects in movies and arousing a growing interest in the fashion industry. Here we briefly discuss major approaches for cloth simulation before focusing on current techniques for modeling realistic garments. A recent survey on the topic, including advances in both Computer Graphics and Computer-Aided Design, can be found in [Liu et al. 2010].

Cloth Simulation Since the 90's, a large body of work in Computer Graphics contributed to the simulation of cloth dynamics. The work by Baraff and Witkin [1998] was the first to allow for the stable simulation of full-size garments in reasonable timings, thanks to a first-order implicit discretization of the dynamics. Subsequent works contributed on the robust treatment of contact and friction [Bridson et al. 2002], on better formulations for bending [Bridson et al. 2003; Grinspun et al. 2003], and on the limitation of stretching [Goldenthal et al. 2007; English and Bridson 2008; Thomaszewski et al. 2009]. Current challenges include the incorporation of real material properties (see, e.g., [Wang et al. 2011]) and the question of interactivity at high resolutions (see, e.g., [Kim et al. 2013; Hahn et al. 2014]).

Despite the profusion of simulation models, it is noteworthy that all of them share a similar *nodal* representation of cloth. That is, cloth is always modeled as a mesh whose vertices (3D positions) serve as degrees of freedom for the corresponding dynamical system.

Physics-based cloth tailoring The real process for making wearable garments involves the design of 2D patterns, i.e., flat patches which are then manufactured with real fabric and sewn together to create the final garment. Early cloth simulation methods mimicked this process in order to dress virtual characters before animating them [Carignan et al. 1992]. To improve the garment modeling process, Volino and colleagues [2005] have later imagined an intuitive and interactive design environment, where the user can simultaneously edit 2D patterns and visualize the 3D result under gravity and body contacts thanks to a fast draping simulator. Such an interactive physics-based tailoring process is now the standard work flow used in many commercial tools like the popular Marvelous designer software [2010]. It has also inspired further work in academia [Umetani et al. 2011], which has improved the 3D responsiveness to 2D pattern editing and thus greatly smoothed the creative flow.

Virtual tailoring is powerful in the sense that it allows experts in pattern design to create and experiment with complex garments without having to manufacture them. Moreover, final results may not only be used for making real garments, but also for dressing virtual avatars in a realistic way. This has motivated a number of VFX studios to adopt such a modeling strategy [Marvelous Designer 2010; CreativeBloq 2012]. An advanced usage however requires some specific skills in pattern sewing, making the technique pretty complex and unnatural for traditional 3D designers [CreativeBloq 2012]. Moreover, such a trial and error process proves unpractical for reproducing the exact look of a visual reference provided for instance as a painting, a photograph or a 3D capture of a real drape,

.

Geometry-based cloth modeling An interesting alternative in cloth modeling is to leverage 3D geometric tools for directly sculpting the final pose regardless of the underlying physical process at the origin of the cloth deformation. Two families of geometric methods are currently available: manual geometric design, and automatic geometric acquisition. On the one hand, advanced shape editing tools allow 3D designers to carve the drape of a garment directly around a virtual character (see, e.g., [Porumbescu et al. 2005]). For fairly simple garments, sketch-based interfaces may greatly improve the intuitiveness and speed of the process [Igarashi and Hughes 2003; Turquin et al. 2007]. Once a final garment has been created, some geometric transfer methods can be used for automatically adapting it to various character morphologies, while preserving the style [Brouet et al. 2012]. On the other hand, thanks to the considerable advance of image-based capture these latest years, it becomes now affordable to acquire precisely the full 3d geometry of static cloth with folds and wrinkles [White et al. 2007; Bradley et al. 2008], thus allowing for an automatic garment creation from a real source.

Geometry-based techniques are appealing because they provide the user with full control over the final 3D shape of the garment. Realistic but also imaginary garments can be created this way, leaving aside the flat patch-based structure of garments and focusing instead on the free-form 3D appearance. Yet, since the modeling process is done independently from physical simulation, the resulting geometry cannot be interpreted mechanically in a straightforward way. This implies that the modeled garment cannot be animated nor physically modified easily. A naïve approach consists in plugging the resulting shape as the rest (undeformed) configuration of a shell simulator. However, once the simulation gets started, the shape irremediably sags under gravity and may greatly diverge from the desired pose, thus ruining all prior modeling efforts (see Figure 2).

Figure 2: *Naïve initialization of the rest pose by the target (left) yields sagging when statics is simulated (right), thus diverging from the target pose. In contrast, our inverse cloth design strategy allows to recover a plausible rest pose, under which the target pose matches a static configuration of the cloth simulator.*

Inverse cloth modeling We believe that a new, promising way for modeling cloth both realistically and intuitively is to investigate an *inverse modeling* strategy. Inverse problems have been studied in a huge amount of fields ranging from Mechanical Engineering to Meteorology, Electro-Magnetism and more recently Biology. The basic principle is to retrieve unknown parameters of a model

from the observation of the real object of interest. In Mechanical Engineering, inverse problems can be classified in two main categories [Beck and Woodbury 1998]: (a) inverse *measurement* problems, where material properties (e.g., mass, stiffness, internal damping) are looked for; and (b) inverse *design* problems, where the undeformed configuration of the object is sought for.

For cloth, all existing studies focus on case (a), that is on the estimation of material parameters while assuming the rest shape to be perfectly flat [Jojic and Huang 1997; Bhat et al. 2003; Bouman et al. 2013]. Such methods, useful to identify the type of fabric involved (knit, linen, satin, etc.), are applied to simple rectangular cloth patches hung under gravity, generally in the absence of contacts or with a coarse approximation of them [Jojic and Huang 1997]. Because they assume the rest pose to be flat, they cannot be used for inverting full (sewn) garments, unless the exact structure of the cloth assemblage (shape of flat panels, location of darts, etc.) can be provided [Hasler et al. 2007].

In contrast, in this paper our goal is to allow the user to animate any geometrically-designed garment in a both practical and consistent way, with minimum input. That is, the user may only provide the mesh of the garment and that of the underlying body, as well as material parameters that he/she can freely choose depending on the type of motion desired. As we deliberately disregard the internal structure of the garment and want to be able to take any 3D mesh as input (not necessarily resulting from an assembly of flat patches under gravity), we consider that the rest shape of the garment cannot be determined easily and should be searched for in the space of 3D (warped) surfaces. Thus, we rather draw our attention to inverse design methods (b), where the rest shape is left as unknown.

Inverse elastic design A few approaches in Computer Graphics have recently focused on the identification of the rest pose of elastic structures such as 2D fibers [Derouet-Jourdan et al. 2010], mass-spring systems [Twigg and Kačić-Alesić 2011], and 3D finite element models [Chen et al. 2014]. Yet, these approaches are either too specific, or in contrast too generic, to take properly into account the peculiar structure of cloth models . Furthermore, none of them incorporates frictional contact in the inversion process.

The method which is closest in spirit to ours is the inverse modeling strategy recently introduced by Derouet-Jourdan and colleagues [2013] in the case of hair, which extends [Derouet-Jourdan et al. 2010]. In this work, the authors show that it is possible to retrieve a plausible rest shape for each fiber such that at equilibrium, under gravity and frictional contact forces, the configuration of the fiber assembly matches an arbitrary set of clamped curves given as input. One key point which makes the inverse problem tractable is to rely upon curvature-based reduced models for individual fibers (e.g., super-helices [Bertails et al. 2006]), for which elastic forces depend linearly on the unknown rest configuration (natural curvature and twist). Inverse statics thus formulates as a linear equation w.r.t. the rest pose, subject to frictional constraints. Using some reasonable assumptions on the hair rest pose so as to select one solution out of the many mathematically admissible, the regularized problem boils down to a quadratic program under second-order conical constraints (SOCCQP). The latter can be solved robustly and efficiently using a solver initially devised for the forward hair dynamical problem [Daviet et al. 2011].

To the best of our knowledge, such an inverse design strategy has never been adopted for modeling garments. Because of the nodal formulation pertaining to cloth models, *strong nonlinearity* is introduced in the static equations compared to [Derouet-Jourdan et al. 2013], leading to some new and important computational challenges.

3 Our Contribution: Inverse cloth design

We propose the first method for interpreting an arbitrary surface geometry as the mechanical result of cloth draping subject to gravity and frictional contact. Our idea is to explore *inverse cloth design*, where the rest position of the garment is left as an unknown of the problem, and where (unknown) frictional contact forces may also be considered to account for the input pose. Material parameters (mass, stiffness, friction coefficient) are, in contrast, not considered as unknowns and may be freely chosen by the user.

We consider a nodal cloth model similar to Baraff and Witkin's [1998], but with a non-flat rest shape [Grinspun et al. 2003]. Our inverse static problem then formulates as an ill-posed nonlinear problem subject to conical constraints (Section 4). To tackle this challenging problem in a robust manner, we contribute on the following aspects:

- In the unconstrained case, we define *draping* as an implicit function Φ and apply the adjoint method on a modified function Φ_λ so as to better drive convergence (Section 5).
- We show that frictionless bilateral constraints can be included easily to the unconstrained case (Section 6).
- In the presence of unilateral frictional contact, we show how to formulate draping as a local constrained minimization problem, and generalize the adjoint method to this constrained case (Section 7).

Finally, we successfully apply our inversion method to various cloth input geometries ranging (Section 8) and discuss limitations and further improvements (Section 9).

4 Cloth Inverse Static Problem

Shell cloth model To free ourselves from the knowledge of the pattern structure of garments, we choose to model a garment in one single piece. As the rest position of a sewn garment may not be flat (for instance in the case of a tight skirt or a puff sleeve), we model cloth as a thin elastic *shell*. This way, both the input (deformed) pose and the unknown rest pose live in the same 3D space.

Another reason to consider a non-flat rest pose is that it allows us to accommodate with possible discrepancies of the input mesh. Imagine a tablecloth, whose rest position is presumably flat. If the input pose is generated externally, there is little chance that we retrieve a perfectly flat rest pose after inversion. If we want to satisfy both the equilibrium condition and this flatness constraint, we have to tolerate a slight alteration of the input shape. However, we believe that the exact shape of the rest pose needs not be set as a hard constraint, since the rest pose is not intended to be visualized in 3D applications. We prefer to match the *input pose* as closely as possible, while targeting an as-flat-as-possible rest pose.

In practice we use the thin elastic shell model from [Grinspun et al. 2003].

Notation We represent the cloth geometry as a triangular mesh made of n nodes. Vector $x_\mathrm{t} \in \mathbb{R}^{3n}$ collects the 3D input (deformed) positions of the nodes, while vector $\bar{x} \in \mathbb{R}^{3n}$ collects their unknown rest positions. If $F(x,\bar{x})$ is a vector-valued function defined on $\mathbb{R}^{3n} \times \mathbb{R}^{3n}$, then $\mathrm{D}_x F(x,\bar{x})$ and $\mathrm{D}_{\bar{x}} F(x,\bar{x})$ denote its Jacobian w.r.t. the first and second variable, respectively. For a single variable function $\Phi(\bar{x})$, we shall omit the subscript in the Jacobian, that is writing $\mathrm{D}\Phi(\bar{x})$. For a real-valued function $f(x,\bar{x})$, we shall use the gradient notation $\nabla_x f(x,\bar{x})$ and $\nabla_{\bar{x}} f(x,\bar{x})$ instead, with $\nabla_x f(x,\bar{x}) = (\mathrm{D}_x f(x,\bar{x}))^T \in \mathbb{R}^{3n}$. Scalar product between two vectors x and y will be denoted by $x^T y$.

4.1 Shell potential energy

For now, let us consider that only conservative forces act on our cloth model. Other forces such as frictional contact will be considered when formulating the statics in Section 4.2.

Conservative forces derive from the total potential energy $E_p(x,\bar{x})$ of our shell model, which can be decomposed as

$$E_p(x,\bar{x}) = E_{\text{int}}(x,\bar{x}) + E_{\text{g}}(x),$$

where $E_{\text{int}}(x,\bar{x})$ is the internal elastic energy of the shell, and $E_{\text{g}}(x)$ its gravitational energy. Note that other external energies, such as potential fields (e.g., for modeling penalty forces), could also be considered in this formulation.

Gravitational energy Let m_i be the mass of the i^{th} node, g the constant of gravity, and e_z the unitary upward vector. Gravitational energy reads

$$E_{\text{g}}(x) = \sum_{i \in \text{nodes}} m_i\, g\, x_i^T e_z.$$

Note that E_{g} is independent of the rest pose $\bar{x}$ of the cloth.

Internal elastic energy Given a deformed cloth with nodal positions x and a rest configuration $\bar{x}$, its internal energy reads [Grinspun et al. 2003]

$$E_{\text{int}}(x,\bar{x}) = E_{\text{membrane}}(x,\bar{x}) + E_{\text{bending}}(x,\bar{x})$$

with the membrane (stretch and shear) energy

$$E_{\text{membrane}}(x,\bar{x}) = \frac{k_e}{2} \sum_{\text{edges}} \frac{(e-\bar{e})^2}{\bar{e}} + \frac{k_A}{2} \sum_{\text{faces}} \frac{\left(A-\bar{A}\right)^2}{\bar{A}}, \quad (1)$$

and the bending energy

$$E_{\text{bending}} = \frac{k_\theta}{2} \sum_{\text{edges}} \frac{\bar{e}}{\bar{h}} \left(\theta - \bar{\theta}\right)^2. \quad (2)$$

The quantities e, A, and θ represent the deformed per-edge length, per-triangle area, and per-edge dihedral angle respectively, which can all be computed from the deformed configuration x. Their rest counterparts are denoted by the barred quantities $\bar{e}$, $\bar{A}$, and $\bar{\theta}$, which can similarly be computed from $\bar{x}$. .

Strong nonlinearity in the rest pose Looking at expressions (1) and (2), it is noteworthy that E_{int} can be written formally as

$$\begin{aligned} E_{\text{int}}(x,\bar{x}) &= \frac{1}{2}(\psi(x)-\psi(\bar{x}))^T \mathbb{K}(\bar{x})(\psi(x)-\psi(\bar{x})) \\ &= \frac{1}{2}\left\|\psi(x)-\psi(\bar{x})\right\|^2_{\mathbb{K}(\bar{x})}, \end{aligned} \quad (3)$$

where ψ is a vector-valued function whose components map x to either edge length e, triangle area A, or bending angle θ, and $\mathbb{K}(\bar{x})$ is a symmetric positive definite matrix collecting inverses of edge length or inverses of triangle area, multiplied by stiffnesses. As edge lengths, triangle areas, and bending angles are nonlinear w.r.t. positions, both $\psi(\bar{x})$ and $\mathbb{K}(\bar{x})$ depend *nonlinearly* on $\bar{x}$.

Because of such a nonlinearity, we anticipate here the fact that inverse statics will be a difficult problem, *even in the absence of contact and friction*. This will be indeed confirmed in next section.

To circumvent such a difficulty, one may think of considering a better-suited parameterization of the model. For instance, choosing $\bar{e}$, $\bar{A}$, and $\bar{\theta}$ as new unknowns of our problem, instead of $\bar{x}$. Unfortunately, those three quantities are not independent from each other, and thus do not form a consistent parameterization .

4.2 Inverse static problem

For the sake of clarity, let us introduce the function $F(x,\bar{x}) = \nabla_x E_p(x,\bar{x})$. Mechanically, F represents the sum of conservative forces of the cloth.

Our goal is to interpret the input pose x_{t} as an equilibrium configuration of the model. We would also like this equilibrium to be *stable*. That is, the cloth should go back to the x_{t} configuration when slightly moved apart from x_{t}.

Without contact In the absence of contact, $\bar{x}$ should thus satisfy

$$F(x_{\text{t}},\bar{x}) = 0. \quad (4)$$

Moreover, a stability condition may be written on the positiveness of $\text{D}_x F(x_{\text{t}},\bar{x}) = \nabla^2_x E_p(x_{\text{t}},\bar{x})$. Yet, such a condition is not convenient in practice because it involves the computation of eigen values of the Hessian $\nabla^2_x E_p(x_{\text{t}},\bar{x})$, which may be computer-expensive. Instead, we'll see in Section 5.2 how to guarantee stability of our equilibrium by formulating an appropriate *draping* function of the cloth. For now, we shall only consider the equilibrium condition (4).

Using Expression (3), Equation (4) reads

$$\left(\text{D}\psi(x_{\text{t}})\right)^T \mathbb{K}(\bar{x})\left(\psi(x_{\text{t}}) - \psi(\bar{x})\right) + \nabla_x E_{\text{g}}(x_{\text{t}}) = 0. \quad (5)$$

Equation (5) is, as expected, strongly nonlinear w.r.t. our unknown variable $\bar{x}$. It is not even linear in $\psi(\bar{x})$, because the matrix $\mathbb{K}$ also depends (nonlinearly) on $\bar{x}$. Moreover, the dimension of the problem scales with the cloth size, which often reaches thousands of particles for full-size garments. Inverse statics without contact thus already turns out to be a difficult, large and nonlinear problem.

This contrasts with the simplicity of inverse static problems generated from reduced fiber models [Derouet-Jourdan et al. 2010; Derouet-Jourdan et al. 2013]. In such cases, $\mathbb{K}$ is independent of $\bar{x}$ and $\psi(\bar{x})$ is invertible (and even equal to the identity), thus making it straightforward to retrieve $\bar{x}$ from Equation (5).

With frictional contact In the case when the target pose x_{t} is in contact with an external object (typically, the body of the character), we should not only rely on the rest pose $\bar{x}$ to account for the fact that x_{t} is an equilibrium, but also on (unknown) unilateral contact forces r. Moreover, *dry friction* has to be considered as it clearly helps cloth remain still . Our goal is thus to find a proper balance between the role of the rest pose and that of contact and friction, so that x_{t} closely matches a static pose of the cloth.

In the presence of frictional contact, cloth equilibrium reads

$$F(x_{\text{t}},\bar{x}) = f_c(x_{\text{t}})$$

where f_c collects all (non-conservative) frictional contact forces, which are a priori unknown. As in [Derouet-Jourdan et al. 2013] we use the exact (implicit) Signorini-Coulomb frictional contact model introduced to graphics in [Bertails-Descoubes et al. 2011]. Please refer to [Acary and Brogliato 2008] for a comprehensive analysis of this model.

Unlike conservative forces, frictional contact forces have no explicit expression, but they should satisfy some constraints. In the sequel we recall such constraints in the static case.

For the sake of simplicity, let us assume (C) that our mesh is sufficiently fine such that it is possible to consider contacts only at vertices. Then we can write $f_c = Hr$ with H a 3x3 block diagonal matrix where the i^{th} block is defined by

$$H_{(i,i)} = \begin{cases} \mathbf{0} & \text{if } x_{\text{t}}^i \text{ is not a contact point} \\ R_i & \text{otherwise,} \end{cases}$$

with R_i the rotation matrix that transforms the local (i.e., z-upward oriented) contact coordinates into world coordinates. In the static case, the cloth is necessarily sticking to the body (no relative velocity). Hence each local contact force r_i must satisfy $r_i \in K_{\mu^i}$, where μ^i is the coefficient of friction at contact i and K_{μ^i} the corresponding Coulomb friction cone .

Note that without assumption (C), H would not be block-diagonal anymore, but we would still have a linear relationship between r and f_c.

To sum up, the rest position $\bar{x}$ should satisfy the following problem,

$$\begin{cases} F(x_\mathrm{t},\bar{x}) &= H(x_\mathrm{t})r \\ r_i \in K_{\mu^i}. \end{cases} \tag{6}$$

Problem (6) turns out to be a nonlinear problem subject to a conical constraint. Moreover, two unknowns are now involved in the problem: the rest position $\bar{x}$, and the local frictional contact force r. To the best of our knowledge, such a challenging (both nonlinear and nonsmooth) inverse problem has never been tackled in the past.

5 Inversion without Constraints

Let us start examining the unconstrained case. This case will be generalized easily in Section 6 to the case when some vertices are fixed.

5.1 Failure case: Root-finding formulation

Using the equilibrium condition (4), the inverse static problem satisfies a root-finding problem,

$$\text{find } \bar{x} \text{ such that} \quad F(x_\mathrm{t},\bar{x}) = 0. \tag{7}$$

Relaxation of the equilibrium In the case when input data x_t is slightly noisy, this problem may however not have a solution. Instead, it is safer to consider the following minimization problem,

$$\min_{\bar{x}} \underbrace{\frac{1}{2}\|F(x_\mathrm{t},\bar{x})\|^2}_{g(\bar{x})}. \tag{8}$$

Because g is smooth and coercive on $\mathbb{R}^{3n}$, Problem (8) is guaranteed to possess a solution $\bar{x}^\star$, which is also solution to the original problem (7) if and only if the objective function $g(\bar{x})$ vanishes at $\bar{x}^\star$.

We have initially considered solving this root-finding formulation, but have quickly departed from this idea for several reasons listed below.

Drawbacks of the method First, even though (7) possesses a solution $\bar{x}^\star$, this only guarantees that $E_p(x_\mathrm{t},\bar{x}^\star)$ is a *critical* point of energy, *not* necessarily a minimum of energy. Energy maxima and saddle points would thus need to be explicitly discarded at some point.

Second, solving Problem (8) (and thus Problem (7)) remains difficult in practice. Indeed, even in elementary 2D cloth scenarios, function g is nonconvex and features multiple local minima. Unsurprisingly, classical minimization schemes fall into such local minima, and there is no obvious way to regularize the function so as to guide minimization towards a global minimum.

Finally, in the (most probable) case when solving (8) numerically leads to some non-zero residual — which may happen either because (7) has no exact solution (g does not vanish at $\bar{x}^\star$) or because the minimization of (8) gets stuck into a local minimum — the computed solution $\bar{x}^\star_\mathrm{comp}$ turns out to be pretty useless in our case. Indeed, for such rest pose $\bar{x}^\star_\mathrm{comp}$, the target shape x_t is not an equilibrium configuration of the cloth . The only thing we know is that this point corresponds to a low slope of energy in the x-direction. Hence, setting $\bar{x}^\star_\mathrm{comp}$ as the rest pose of the cloth and starting the cloth dynamic simulator from the x_t state will irremediably result in some motion towards a true equilibrium configuration, which has no reason to be close to x_t.

In summary The main flaw of this root-finding formulation is to enforce the search for a vanishing point $F(x,\bar{x})$ *strictly along the* $x = x_\mathrm{t}$ *subspace*, even if it means finding a $\bar{x}^\star$ that does not guarantee equilibrium at $(x_\mathrm{t},\bar{x}^\star)$. Instead, we prefer to formulate the problem the other way round. That is, explore $E_p(x,\bar{x})$ on a slightly enlarged region close to but not restricted to the subspace $x = x_\mathrm{t}$, so as to find a pair $(x^\star,\bar{x}^\star)$ which *exactly satisfies* equilibrium, with $x^\star$ as close as possible to x_t .

5.2 Least-squares formulation

Instead of considering the equilibrium equation (7) corresponding to the balance of forces, let us model our full inverse problem — including the stability condition — using an energy-based formulation. For a $\bar{x}$ parameter to be found, the target pose x_t should match a local minimum of the function $x \mapsto E_p(x,\bar{x})$. That is, x_t should correspond to a *draped pose* of the cloth simulator initialized with $\bar{x}$ as the rest pose.

Implicit draping function In the case when $x \mapsto E_p(x,\bar{x})$ is a strictly convex function, then a unique draped pose $x^\star$ exists, corresponding to the (global) minimum of the energy. We can thus define a *draping function* on $\mathbb{R}^{3n}$ as $\Phi : \bar{x} \mapsto \mathrm{argmin}_x E_p(x,\bar{x})$.

In the general case however, the potential energy is not convex. It may feature several local minima, thus precluding the definition of such a unique mapping on $\mathbb{R}^{3n}$. This is illustrated in Figure 3.

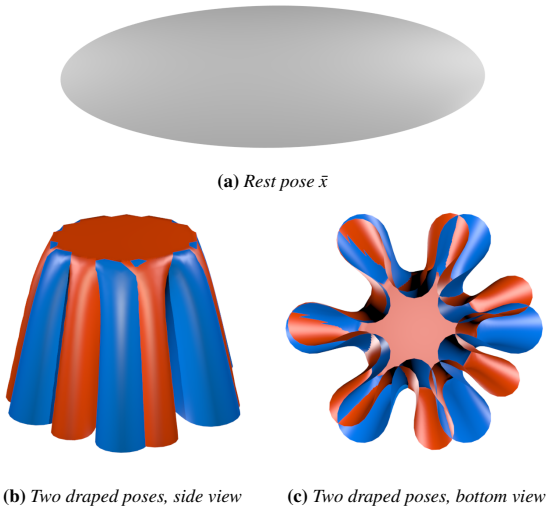

(a) *Rest pose* $\bar{x}$

(b) *Two draped poses, side view* **(c)** *Two draped poses, bottom view*

Figure 3: *For this tablecloth, at least two different local minima of potential energy (one in blue and one in orange) may be obtained from the same rest pose* $\bar{x}$ *(in light grey).*

Yet, in the nonconvex case a draping function Φ can still be defined *locally*. Let us consider a rest pose $\bar{x}^\star$ and a deformed pose $x^\star$ such that $x^\star$ is a strict local minimum of $x \mapsto E_p(x,\bar{x}^\star)$. We thus have $F(x^\star,\bar{x}^\star) = 0$ and $\mathrm{D}_x F(x^\star,\bar{x}^\star) \succ 0$. According to the implicit function theorem, there exists a unique mapping Φ from a neighborhood $\bar{V}$ of $\bar{x}^\star$ to a neighborhood V of $x^\star$ such that $x = \Phi(\bar{x})$ with $\bar{x} \in \bar{V}$

and $x \in V$. Intuitively, in our case computing $\Phi(\bar{x})$ would consist in minimizing the potential energy $x \mapsto E_p(x,\bar{x})$ in the neighborhood of $(x^\star, \bar{x}^\star)$. Moreover, the theorem says that $\mathrm{D}_x F(x^\star, \bar{x}^\star)$ is invertible in $V \times \bar{V}$ and that Φ is differentiable in $\bar{V}$, with an explicit expression for its Jacobian,

$$\mathrm{D}\Phi(\bar{x}) = -\left(\mathrm{D}_x F(\Phi(\bar{x}),\bar{x})\right)^{-1} \left(\mathrm{D}_{\bar{x}} F(\Phi(\bar{x}),\bar{x})\right). \tag{9}$$

Least-squares minimization Thanks to the existence of a local draping function Φ, our inverse problem literally reads

$$\text{find } \bar{x} \text{ such that } \quad x_\mathrm{t} = \Phi(\bar{x}), \tag{10}$$

and can be relaxed in the least squares sense as

$$\min_{\bar{x}} \underbrace{\frac{1}{2}\|\Phi(\bar{x}) - x_\mathrm{t}\|^2}_{J(\bar{x})}. \tag{11}$$

Similarly as before, Problem (11) is guaranteed to possess a solution $\bar{x}^\star$, which is solution to (10) if and only if the objective function $J(\bar{x})$ vanishes at $\bar{x}^\star$. However, in contrast to our former problem (8), if $J(\bar{x}^\star)$ does not exactly vanish we still get some very useful information. Indeed, solving (11) provides us with a deformed pose $\Phi(\bar{x}^\star)$ which is *guaranteed* to be a minimum configuration of the energy function $x \mapsto E_p(x,\bar{x})$ parameterized by $\bar{x}$. Moreover, this deformed pose is as close as possible to the target x_t.

Formulation (11) is common for inverse problems. To minimize J in practice, we use the limited-memory BFGS approach [Nocedal and Wright 2006, Section 9.1], which relies on the computation of ∇J. A classical strategy for efficiently computing ∇J is the adjoint method [Giles and Pierce 2000], which is briefly described in Section 5.3. The adjoint method has the special feature of relying only on the evaluation of Φ to compute ∇J, and not that of $\mathrm{D}\Phi$. However, in our case we are faced with an awkward situation, as the draping function Φ is defined *locally*. In Section 5.4 we shall explain how to evaluate Φ in a consistent way during subsequent steps of the minimization. Our full algorithm for minimizing J is given in Algorithm 1.

5.3 The adjoint method

To solve the minimization problem (11) robustly, it is desirable to compute the gradient ∇J accurately. Differentiating the objective function J gives

$$\nabla J = \mathrm{D}\Phi(\bar{x})^T \left(\Phi(\bar{x}) - x_\mathrm{t}\right). \tag{12}$$

Recall that an explicit expression for $\mathrm{D}\Phi(\bar{x})$ is provided by (9). However, it is dense in general and its computation requires a full matrix inversion.

Actually, we don't really need to compute $\mathrm{D}\Phi(\bar{x})$, but only ∇J. By replacing $\mathrm{D}\Phi(\bar{x})$ with its expression (9), the adjoint method consists in decomposing the computation of ∇J into two steps,

$$\left\{ \begin{array}{rcl} \left(\mathrm{D}_x F(\Phi(\bar{x}),\bar{x})\right)^T p & = & x_\mathrm{t} - \Phi(\bar{x}) \\ \nabla J(\bar{x}) & = & \left(\mathrm{D}_{\bar{x}} F(\Phi(\bar{x}),\bar{x})\right)^T p, \end{array} \right. \tag{13}$$

where p is called the *adjoint state*. Note that the adjoint method requires only one evaluation of Φ at $\bar{x}$ to compute $\nabla J(\bar{x})$.

5.4 Evaluation of the draping function

The draping problem consists in finding one local minimum of the potential energy, which amounts to solving the following problem,

$$\text{Given } \bar{x}, \text{ find } x \quad \text{s. t.} \quad F(x,\bar{x}) = 0 \text{ and } \mathrm{D}_x F(x,\bar{x}) \succ 0. \tag{14}$$

Naïve draping algorithm As already mentioned, this problem has no unique solution, and thus one cannot define a true draping function, at least globally. Still, it is easy to find an admissible x by minimizing the function $x \mapsto E_p(x,\bar{x})$ locally (using for instance the Newton method), starting from the initial guess $x_0 = x_\mathrm{t}$. By analogy with the draping function, let us call Φ this draping algorithm, which takes as input $\bar{x}$ and returns x.

Unfortunately, eventhough it is initialized with x_t, such a method is likely to return a local minimum *that is far from* x_t. In this case, the inverse method for solving (11) may get stuck at a wrong place and thus hardly converge to a global minimum. Furthermore, our procedure Φ is not continuous w.r.t. the $\bar{x}$ variable, as two close positions for $\bar{x}$ might lead to completely different local minima x. Hence, during inversion it may happen that two subsequent steps of the adjoint evaluate Φ at two unrelated places. In this case the computation of ∇J becomes meaningless .

Regularized draping algorithm Our goal is to build a draping procedure that remains *consistent* with our inversion algorithm. That is, if x_t is a local minimum of energy, the Φ procedure should be able to return it. To this aim, we penalize the energy to be minimized so as to avoid falling into a local minimum which is far from the target x_t. That is, we consider the new potential energy

$$E_p^\lambda(x,\bar{x}) = E_p(x,\bar{x}) + \frac{\lambda}{2}\|x - x_\mathrm{t}\|^2, \tag{15}$$

where $\lambda \geq 0$ is a regularization factor. Choosing a high value for λ helps "convexify" the potential energy around x_t. Of course it also leads to an energy that is quasi-independent of $\bar{x}$, thus far from the original one. Conversely, setting a low value restores the original energy, but we loose the benefit of penalization. Note that if E_p reaches a minimum at x_t for a given $\bar{x}$, then it is also a minimum for $E_p^\lambda, \forall \lambda \geq 0$. This encourages use to decrease the value of λ over successive calls of Φ during inversion, when the confidence in $\bar{x}$ increases (see Algorithm 1).

Our new procedure Φ^λ thus consists in minimizing $x \mapsto E_p^\lambda(x,\bar{x})$, starting from $x_0 = x_\mathrm{t}$. In practice we use the Newton-CG method [Nocedal and Wright 2006, Section 6.2], to perform the minimization, that is we solve at each step k the following linear system in Δ_{k+1},

$$\mathrm{D}_x F^\lambda(x_k,\bar{x})\,\Delta_{k+1} = -F^\lambda(x_k,\bar{x}) \tag{16}$$

where $\Delta_{k+1} = x_{k+1} - x_k$ and $F^\lambda(x,\bar{x}) = \nabla_x E_p^\lambda$. Compared to the pure Newton algorithm, the Newton-CG method has the advantage of always finding a descent direction Δ_{k+1}, in particular in the case when $\mathrm{D}_x F(x_k,\bar{x})$ is not positive-definite. To accelerate convergence, we perform an adaptive increment along Δ_{k+1} using a Wolfe linesearch. For most of our examples, we obtained convergence in a few iterations only. This good convergence rate — which contrasts with some earlier feedbacks on the Newton method [Volino and Magnenat-Thalmann 2007] — is explained by the fact that minimization starts from the target x_t, which is already fairly close to the solution.

5.5 As-flat-as-possible initial guess

Multiple admissible rest poses We have seen that from one given $\bar{x}$, multiple deformed configurations x may satisfy the draping conditions (14). Conversely, given an arbitrary shape x, multiple rest poses $\bar{x}^\star$ may satisfy the same conditions.

This non-uniqueness problem is not really an issue, as what we really want is to retrieve at least one rest pose $\bar{x}$ such that x_t satisfies the draping problem (14). That said, when we have the choice between several admissible rest poses $\bar{x}$, it is of course desirable to select the most plausible one.

As-flat-as-possible rest pose Because garments are designed from flat patterns, folds usually occur only because of the deformation under gravity. Thus, it is natural to target a rest pose $\bar{x}$ which is *as-flat-as-possible*. More precisely, since the topology of the garment is not necessarily that of a plane (in the case of a skirt or a sleeve, it has the topology of a cylinder), we would like to find the rest pose $\bar{x}$ which has the lowest possible Gaussian curvature.

Note that this choice has an influence on the realism of subsequent animation of the object, since $\bar{x}$ represent the shape that internal elasticity of the cloth wish to preserve during motion. Moreover, starting from a locally flat shape for the cloth would greatly help the texturing phase.

Flattened initial guess In practice however, it is difficult to penalize high curvatures of $\bar{x}$ in a manner that remains consistent with our inversion algorithm. Instead, we opt for a much simpler yet effective solution: we initialize our rest pose $\bar{x}$ with a flattened version of the target pose x_t, before starting our inversion algorithm.

To flatten the input configuration x_t, we minimize the internal energy $x \mapsto E_\mathrm{int}(x, \tilde{x}_\mathrm{t})$ with the rest pose $\tilde{x}_\mathrm{t}$ corresponding to x_t except that rest angles $\bar{\theta}$ are set to zero. Thus, minimization attempts to cancel edge dihedral angles of the target x_t while preserving edge lengths and face areas.

5.6 Full inversion algorithm

Let us define $J^{\lambda_k}(\bar{x}) = \frac{1}{2}\|\Phi^{\lambda_k}(\bar{x}) - x_\mathrm{t}\|^2$. Our full inversion method in the unconstrained case is summarized in Algorithm 1. To decrease λ at each step of the outer loop, in practice we chose $\alpha = 0.5$ in all our examples, except for the tablecloth for which progressive penalization was not used (one step with $\lambda = 5$ was sufficient).

Algorithm 1: Robust inversion algorithm without constraints.

Data: x_t, λ_0, $\alpha \in]0;1[$
Result: $\bar{x}$ that minimizes J
`// Warm start`
$\bar{x}_0 \leftarrow$ `Flatten`(x_t) ;
for $k = 0, 1, \cdots$ until $\lambda_k \leqslant \varepsilon_\lambda$ or $\|F(\Phi(\bar{x}_k), \bar{x}_k)\| \leqslant \varepsilon_\nabla$ **do**
  `// Find` $\bar{x}^{\lambda_k}$ `that minimizes` $J^{\lambda_k}(\bar{x})$
  $\bar{x}_0 \leftarrow \bar{x}_k$;
  for $i = 0, 1, \cdots$ until $\left\|\nabla J^{\lambda}(\bar{x}_i)\right\| \leqslant \varepsilon$ **do**
    `// Evaluate` Φ^{λ_k} `at` $\bar{x}_i$ `by draping`
    $\Phi^{\lambda_k}(\bar{x}_i) \leftarrow$`Newton-CG`$(x \to E_p^{\lambda_k}(x, \bar{x}_i), x_\mathrm{init} = x_\mathrm{t})$;
    `// Evaluate` ∇J^{λ_k} `at` $\bar{x}_i$ `by the adjoint`
    $\nabla J^{\lambda_k}(\bar{x}_i) \leftarrow$`Adjoint`$(\Phi^{\lambda_k}(\bar{x}_i))$;
    `// Compute L-BFGS descent direction`
    $d_i \leftarrow$`L-BFGS-Descent-Direction`$(\nabla J^{\lambda_k}(\bar{x}_i))$;
    `// Update` $\bar{x}$
    $\bar{x}_{i+1} \leftarrow \bar{x}_i + d_i$;
    `// Store best minimizer of` $J^{\lambda_k}(\bar{x})$
    $\bar{x}^{\lambda_k} \leftarrow \bar{x}_{i+1}$;
  end
  `// Update outer iterate`
  $\bar{x}_{k+1} \leftarrow \bar{x}^{\lambda_k}$;
  `// Decrease` λ
  $\lambda_{k+1} \leftarrow \alpha\lambda_k$;
end

6 Removing degrees of freedom

For some cases like a pinned piece of cloth, a flag, a tablecloth or a dress fitted at the waist, it may be useful to consider that predefined vertices of the cloth are fixed. In [Baraff and Witkin 1998], any vertex i may be prevented from moving in certain space directions thanks to a simple filtering process, equivalent to an orthogonal projection of the equations [Ascher and Boxerman 2003].

Let Π be the orthogonal projection matrix applied to the system (denoted by S in [Ascher and Boxerman 2003]). For instance, if vertex i is fixed, we have $\Pi_i = 0$, where Π_i is the i^{th} diagonal block of the matrix Π.

6.1 Modification of the inversion algorithm

To account for constrained vertices in the inversion method, it is sufficient to project draping onto the admissible space of positions, i.e., to replace (16) with

$$\Pi \mathrm{D}_x F(x_{k+1} - x_k) = -\Pi F(x_k).$$

As for the adjoint algorithm, it remains unchanged. .

Hence, inversion with fixed vertices does not add any complexity to the unconstrained case. Since vertices in contact are assumed to be fixed due to static friction, why not use this filtering method for dealing with frictional contact ? Actually, we show in the sequel that such a method is not equivalent to the one dealing with true frictional contact, because it can yield arbitrarily oriented forces.

6.2 Bilateral constraints vs. unilateral contact

Filtering the system of equations by Π is equivalent to the adding of some frictionless bilateral constraints to the vertices. More precisely, each vertex i is subject to the linear constraint $I_3 - \Pi_i$, where I_3 is the 3×3 identity matrix. If b_i is the corresponding force (i.e., , the Lagrange multiplier) applied onto vertex i, we have

$$(F)_i(\Phi(\bar{x}), \bar{x}) = (I_3 - \Pi_i)\, b_i,$$

where $(F)_i(\Phi(\bar{x}), \bar{x})$ is the i^{th} block of size 3 of F evaluated at point $\Phi(\bar{x}), \bar{x}$. Thus, for a fixed vertex i, we have direct access to its applied force $b_i = (F)_i(\Phi(\bar{x}), \bar{x})$.

Figure 4a displays these bilateral forces (black arrows) for a simple 2D case. In this example, our target equilibrium shape x_t is depicted in red. It corresponds to a curve, slightly curly at both ends, lying on a circular object. To invert this pose, we started from the flattened shape represented as the blue dotted line, and fixed all contacting point during inversion. Inversion retrieves the solid blue configuration $\bar{x}^\star$, straight in the middle and curly at both ends. Note that the resulting equilibrium pose $\Phi(\bar{x}^\star)$ (in green) is identical to the target x_t.

At a first glance, this rest pose looks consistent with the input data. Indeed, only points outside the contact zone have been displaced from the initial flat guess. However, it is noteworthy that the "gluing" forces, represented with black arrows, cannot be interpreted as consistent contact forces. Some of them are penetrating the circular obstacle while others, correctly pointing outwards the obstacle, share the same orientation upon the circle as if the friction coefficient was varying along the surface.

In contrast, our inversion with frictional contact, displayed in Figure 4b, automatically yields forces belonging to the local Coulomb friction cone (see next section). Because this method does not tolerate forces to have any arbitrary orientation like in the previous case, it retrieves a more plausible rest pose $\bar{x}^\star$. In particular, it was able to reverse the stretching applying to the contacting zone of the

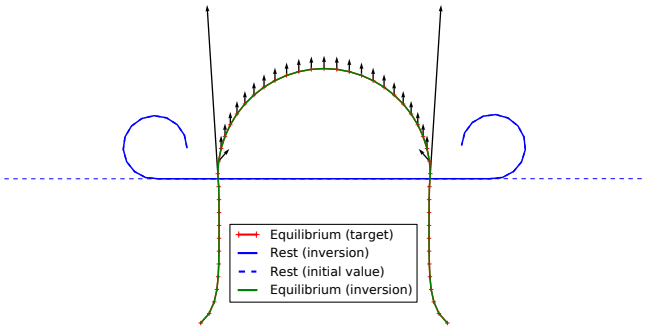

(a) *Inversion with contact modeled by fixing positions.*

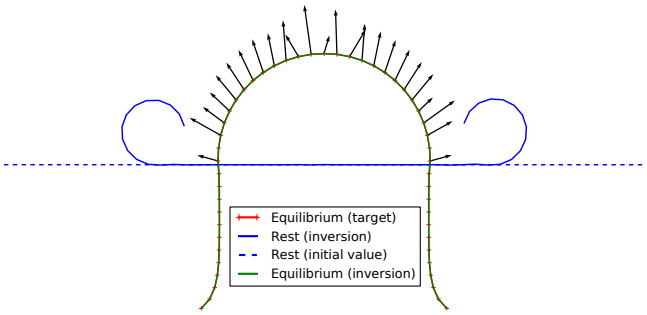

(b) *Inversion with unilateral contact and dry friction.*

Figure 4: *Comparison of the yielded contact forces in the bilateral case and in the unilateral case with dry friction.*

target, and to retrieve a slightly more compressed rest pose than the filtering method.

The filtering method is thus limited to constraints that should still hold in the rest configuration. Yet, it can still be used to generate a good warm start for inverting the system with contact and friction (see next section).

7 Including Frictional Contact

Starting from the equilibrium condition (6), our inverse static problem formulates as

$$\text{find } (\bar{x}, r) \text{ such that } \begin{cases} F(x_\mathrm{t}, \bar{x}) &= H(x_\mathrm{t}) r \\ r_i \in K_{\mu^i}. \end{cases} \tag{17}$$

Obviously this problem has no unique solution, but as previously, we show that an initial as-flat-as-possible warm start greatly helps select a plausible rest pose $\bar{x}$.

Similarly as in Section 5, we do not want to solve strictly Problem (17), but rather relax the problem using a least-squares formulation. This raises two major challenges:

How to compute draping ? It is noteworthy that the balance of forces under frictional contact (17) involves two dual variables: (a) the primal (position) x variable, and (b) the dual (force) r variable. As is, such a mixed formulation is not adapted for computing a static pose x given a rest pose $\bar{x}$. In Section 7.1 we show how to formulate draping as a constrained minimization problem on the primal variable x only, using the fact that contact point locations are known from the target pose x_t.

How to generalize the adjoint method ? Recall that the adjoint method is defined starting from an implicit equation $F(x,\bar{x}) = 0$ where F is assumed to be differentiable. When unilateral contact with friction is considered, we come up with a new implicit equation $F_c(x,\bar{x})$ accompanied with a conical constraint (see Problem (17). This new format thus does not fit in with the adjoint framework straightforwardly. In Section 7.3 we introduce a nonsmooth projection to convert our draping problem into a zero-finding formulation, and show that the adjoint keeps on working well in this nonsmooth case.

7.1 Towards a least-squares formulation

Primal vs. dual formulation The static formulation (17) is pretty inconvenient, as the constraint is expressed on the dual variable r. Moreover, E_p is nonlinear with respect to x, making it difficult to express $\bar{x}$ as a function of r. Note that this contrasts with reduced fiber models that feature linear elastic forces, which allows for a straightforward elimination of the primal variable $\bar{x}$ in the system [Derouet-Jourdan et al. 2013].

One solution would be to linearize F locally, so as to formulate a local problem depending only on the dual variable r, as done for instance in [Kaufman et al. 2014]. However, such a method would be computationally expensive in our case, as it would yield a dense Delassus operator.

Instead, by using some tools of convex analysis, we demonstrate here that a much simpler formulation of (17) can be derived, which only depends on the *primal* variable x.

A new constrained minimization problem Let us assume for now that $x \mapsto E_p(x,\bar{x})$ is a convex function. Let C be a convex set of $\mathbb{R}^{3n}$. From convex analysis [Hiriart-Urruty and Lemaréchal 1993, Theorem 2.1.4] we have the following result,

$$\min_{x \in C} E_p \iff -F(x) \in \mathcal{N}_C(x), \tag{18}$$

where $\mathcal{N}_C(x)$ is the normal cone of C at x, defined as

$$\mathcal{N}_C(x) = \{d \in \mathbb{R}^{3n}, d^T (c-x) \leq 0 \quad \forall c \in C\}.$$

In the particular case when the convex set C is a cone $K_{\frac{1}{\mu}}$ of aperture $\frac{1}{\mu}$, the normal cone boils down to

$$\mathcal{N}_{K_{\frac{1}{\mu}}}(x) = -K_\mu \cap x^\perp \quad \text{where } x^\perp = \{y \in \mathbb{R}^{3n}, y^T x = 0\}.$$

Thus, the equivalence (18) can be expressed as

$$\min_{x \in K_{\frac{1}{\mu}}} E_p(x) \iff K_\mu \ni F(x) \perp x \in K_{\frac{1}{\mu}}. \tag{19}$$

This condition is illustrated in .

Let us assume our target pose x_t has one contact located at the i^{th} node $x_{\mathrm{t},i}$. From (17) we get

$$(F(x_t, \bar{x}))_i \in R_i K_{\mu^i}, \tag{20}$$

where $R_i K_{\mu^i}$ is still a cone of aperture μ^i, corresponding to a rotated cone K_{μ^i} .

Besides, let us consider the following minimization problem

$$\min_{x \in C} E_p(x,\bar{x}) \quad \text{with } C = \left\{ x \in \mathbb{R}^{3n}, H^T(x - x_\mathrm{t}) \in K_{\frac{1}{\mu}} \right\}, \tag{21}$$

where $K_{\frac{1}{\mu}} = K_{\frac{1}{\mu^1}} \times \ldots \times K_{\frac{1}{\mu^n}}$. We have

$$\begin{aligned} x \in C &\iff H^T(x - x_\mathrm{t}) \in K_{\frac{1}{\mu}} \\ &\iff R_i^T (x_i - x_{\mathrm{t},i}) \in K_{\frac{1}{\mu^i}} \quad \forall i \text{ in contact} \\ &\iff (x_i - x_{\mathrm{t},i}) \in R_i K_{\frac{1}{\mu^i}} \quad \forall i \text{ in contact.} \end{aligned}$$

Thus, using Equivalence (19), it follows that Problem (21) is equivalent to solving

$$\begin{cases} \forall i \text{ in contact}, & R_i K_{\mu^i} \ni (F(x,\bar{x}))_i \perp R_i(x_i - x_{\mathrm{t},i}) \in K_{\frac{1}{\mu^i}} \\ \forall i \text{ not in contact}, & (F(x,\bar{x}))_i = 0. \end{cases}$$

For contact-free points, the initial problem (17) is automatically satisfied when $x_\mathrm{t} = x$. For points in contact, when $x = x_\mathrm{t}$, the right-hand side of the above complementarity condition vanishes, leading exactly to (20) and thus, again, to our initial problem (17). Very interestingly, this means that $(\bar{x}, r)$ is solution to (17) if and only if x_t is solution to (21).

Draping with unilateral contact and friction In the general (nonconvex) case, we still benefit from the following sufficient condition: if x_t is solution to (21), then $(\bar{x}, r)$ is solution to (17). We thus transform the solving of (17) into the searching of $\bar{x}$ such that x_t matches a local minimum of Problem (21). Locally solving (21) thus defines a new "draping" procedure $\Phi(\bar{x})$ which takes into account frictional contact. Similarly as in the unconstrained case, we can thus imagine to derive a least-squares formulation for finding a $\bar{x}$ which allows x_t to match $\Phi(\bar{x})$. Still, two questions remain: (a) How to evaluate $\Phi(\bar{x})$ in practice ? and (b) How to modify the adjoint method so that our new draping function can fit in ? These two issues are examined in the following sections.

7.2 Evaluation of the draping function

Due to the simple nature of the constraints appearing in (21), the draping problem can be solved with a simple algorithm such as a (conjugate) projected gradient. Indeed, it suffices to remark that the orthogonal projection of x on the constraints is explicitely given by

$$\Pi_C\left(x^i\right) = \begin{cases} x^i \text{ if } x_\mathrm{t}^i \text{ is not a contact point} \\ x_\mathrm{t}^i + \Pi_{K_{\frac{1}{\mu^i}}}\left(R_i^\top (x^i - x_\mathrm{t}^i)\right) \text{otherwise,} \end{cases}$$

where $\Pi_{K_{\frac{1}{\mu^i}}}(\cdot)$ is the orthogonal projection on cone $K_{\frac{1}{\mu}}$.

The draping procedure is provided in Algorithm 2.

Algorithm 2: Projected gradient for evaluating the draping with contact and friction.

Data: $\bar{x}$, H, $K_{\frac{1}{\mu}}$, $\alpha \in [0,1]$
Result: x that satisfies Equation (17)
$x_0 \leftarrow x_\mathrm{t}$;
for $k = 0, 1, \cdots$ **do**
  $d_k \leftarrow F(x_k, \bar{x})$;
  if $\| x_k - \Pi_C(x_k - d_k) \|$ *sufficiently small* **then**
    return x_k ;
  end
  $\rho \leftarrow 1$;
  repeat
    $x_{k+1} \leftarrow \Pi_C(x_k - \rho d_k)$;
    $\rho \leftarrow \alpha\rho$;
  until $E_p(x_{k+1}, x_\mathrm{t}) < E_p(x_k, x_\mathrm{t})$;
end

7.3 Fitting in with the adjoint method

In order to use the formalism of the adjoint method, we can first rewrite Equation (17) using the orthogonal projection operator $\Pi_{H,K_\mu}\left(x^i\right) := R_i \Pi_{K_{\frac{1}{\mu^i}}}\left(R_i^\top x^i\right)$.

Equation (17) is then equivalent to the root-finding problem

$$h(x,\bar{x}) := (I - \Pi_{H,K_\mu}(\cdot)) \circ F(x,\bar{x}) = 0. \tag{22}$$

The adjoint algorithm can be now written in a similar fashion than without contacts. That is, for a given value of $\bar{x}$,

1. Compute $x^\star = \Phi(\bar{x})$;
2. Compute the adjoint state p as $\left[\frac{\partial h}{\partial x}(x^\star, \bar{x})\right]^\top p = x_\mathrm{t} - x^\star$;
3. Compute the gradient of J as $\nabla J(\bar{x}) = \left[\frac{\partial h}{\partial \bar{x}}(x^\star, \bar{x})\right]^\top p$.

Similarly as in the bilateral case, it can be shown that the adjoint method does not need to be modified. . Furthermore, in practice we have not observed any issue related to the fact that h is locally nonsmooth.

8 Results

8.1 Framework

We have tested our inversion method on three main 3D geometries:

- A tablecloth, generated from simulation. Contacts (onto a centered round table) are modeled by fixing positions, using the method of Section 6;
- A puff sleeve, modeled by an artist. Contacts (localized at both extremities of the sleeve) are also modeled by fixing positions;
- A skirt, modeled by an artist. Unilateral contacts with dry friction are considered between the garment and the body, using the method of Section 7.

Tablecloth:

Skirt: 5208 vertices and 755 contact points

Puff sleeve: 2045 vertices and 102 contact points

For each of these geometries, we evaluated the correctness of our method in two ways:

- First numerically, by evaluating the closeness of the equilibrium to the target pose. These results are compiled on Figure 5.
- Then visually: we put each equilibirum in a forward cloth simulator to ensure that (a) when the simulation starts, the cloth does not sags, and (b) when we add a new external force, such as wind, the cloth behaves in plausible manner. These simulations are shown in the accompanying video.

For the puff sleeve and the skirt, we were not able to match the target perfectly. However, in those cases we were still able to match the intended look of the cloth, and found a much better starting point for the simulation than originally available.

Performance While devising the fastest algorithm was not our primary goal, we are aware that being too slow could significantly limit the applicability of our method. The geometry which took the most time to invert was the skirt, which features about 5200 vertices, with a wall-clock time of just under three hours.

9 Discussion

Modeling cloth as a thin elastic shell In the case when cloth is modeled manually using a 3D free form interface, there is no guarantee that the result satisfies realistic cloth constraints. Typically,

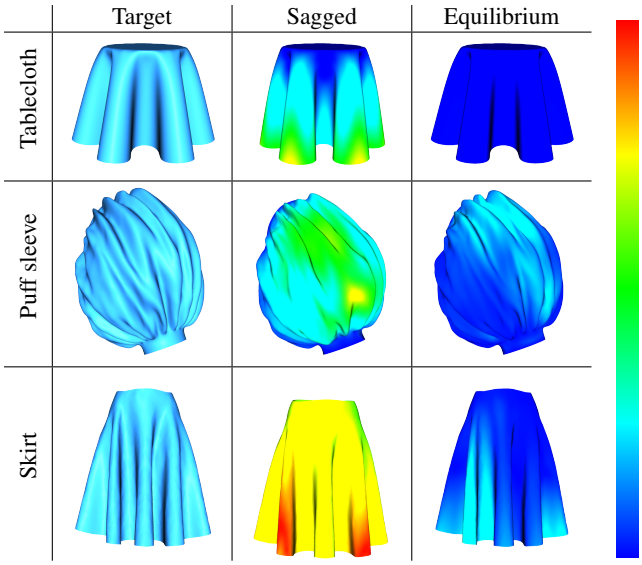

Figure 5: *Evaluation of our inversion method on our 3 examples. Left column shows the target pose; middle column compares the sagged (naïve) pose against the target, and right column compares our generated equilibrium $\Phi(\bar{x}^\star)$ against the target. Blue color indicates a very low error , whereas red indicates an error of the order of 10% of a typical scene length.*

some cloth parts may fail to preserve a low Gaussian curvature, leading not only to troubles in texturing but also to an unnatural overall appearance [CreativeBloq 2012] .

Advantages In comparison with the traditional work flow of pattern design, where the user continuously transitions between pattern editing and draping simulation, our modeling process is very different, as it is split into two successive phases: (a) the geometric modeling phase, which is fully interactive or automatic, and (b) our automatic (and fast) conversion into a physics-based cloth model. Unlike physics-based tailoring, our approach is geometry centered and physics is transparent to the user: geometry is freely explored by the user, while physics should obey (and not the reverse). Any shape is feasible, but of course physical realism may be altered (cloth may be made with shell patterns).

Limitations For the sake of simplicity, we assumed here the garment material properties (mass, stiffness) to be known. This simplifying assumption appears reasonable to us in a first step. Indeed, whereas a designer cannot easily guess the rest pose of a deformed configuration, she/he generally has some idea about the kind of material (silk, linen) she/he would like to simulate. Moreover, as suggested by some studies in fiber design [Derouet-Jourdan et al. 2010] and as we show here, for a well-chosen rest shape and regardless of the material properties, any given configuration may be interpreted as a static pose. This means that even if the designer chooses a "wrong" material, our method will still provide a rest pose which allows to interpret the input as a static pose. But of course, in such a case the rest pose may be pretty unrealistic . .

In some cases (like the puff sleeve) where the target is mainly explained by bending forces, our algorithm has more difficulty to converge to an equilibrium that is very close to the target. In practice we have noticed that the success of our convergence is highly dependent upon the formulation of the bending energy. We would like to investigate this issue in the future.

Still, our method produces results which greatly improve naïve approaches consisting in (a) initializing $\bar{x}$ with x_t, or (b) initializing $\bar{x}$ with a flattened version of x_t (see Figure 5).

Future work So far we have not considered self-contacting for inversion. Adding self-contacts would not bring any difficulty to our theoretical framework. However, we would have to investigate a faster draping function so as to be able to scale up the number of contacts (currently limited to a few thousands).

Another interesting venue of our approach would be to generate realistic patterns from our retrieved rest pose. Indeed, current methods compute such patterns by flattening the deformed shape (the target), thus neglecting deformations due to stretching and shearing [Brouet et al. 2012]. With our approach we could generate more realistic patterns, and thus fully resolve the 3D→2D computation. In that sense our approach could be served to enrich current pattern design methods such as [Umetani et al. 2011]

References

ACARY, V., AND BROGLIATO, B. 2008. *Numerical methods for nonsmooth dynamical systems*, vol. 35 of *Lecture Notes in Computational and Applied Mechanics*. Springer.

ASCHER, U., AND BOXERMAN, E. 2003. On the modified conjugate gradient method in cloth simulation. *The Visual Computer 19*, 7-8, 526–531.

BARAFF, D., AND WITKIN, A. 1998. Large steps in cloth simulation. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'98)*, 43–54.

BECK, J., AND WOODBURY, K. 1998. Inverse problems and parameter estimation: integration of measurements and analysis. *Measurement Science and Technology 9*, 6, 839.

BERTAILS-DESCOUBES, F., CADOUX, F., DAVIET, G., AND ACARY, V. 2011. A nonsmooth Newton solver for capturing exact Coulomb friction in fiber assemblies. *ACM Transactions on Graphics 30* (February), 6:1–6:14.

BERTAILS, F., AUDOLY, B., CANI, M.-P., QUERLEUX, B., LEROY, F., AND LÉVÊQUE, J.-L. 2006. Super-helices for predicting the dynamics of natural hair. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'06) 25*, 1180–1187.

BHAT, K. S., TWIGG, C. D., HODGINS, J. K., KHOSLA, P. K., POPOVIĆ, Z., AND SEITZ, S. M. 2003. Estimating cloth simulation parameters from video. In *Proc. ACM SIGGRAPH/EG Symp. Comp. Animation*, SCA '03, 37–51.

BOUMAN, K. L., XIAO, B., BATTAGLIA, P., AND FREEMAN, W. T. 2013. Estimating the material properties of fabric from video. In *International Conference on Computer Vision (ICCV'13)*.

BRADLEY, D., BOUBEKEUR, T., AND HEIDRICH, W. 2008. Accurate multiview reconstruction using robust binocular stereo and surface meshing. In *Computer Vision and Pattern Recognition (CVPR'08)*.

BRIDSON, R., FEDKIW, R., AND ANDERSON, R. 2002. Robust treatment of collisions, contact and friction for cloth animation. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'02 21*, 3, 594–603.

BRIDSON, R., MARINO, S., AND FEDKIW, R. 2003. Simulation of clothing with folds and wrinkles. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'03)*, ACM-EG SCA, 28–36.

BROUET, R., SHEFFER, A., BOISSIEUX, L., AND CANI, M.-P. 2012. Design preserving garment transfer. *ACM Transactions on Graphics 31*, 4, 36:1–36:11.

CARIGNAN, M., YANG, Y., MAGNENAT-THALMANN, N., AND THALMANN, D. 1992. Dressing animated synthetic actors with complex deformable clothes. *SIGGRAPH Comput. Graph. 26*, 2 (July), 99–104.

CHEN, X., ZHENG, C., XU, W., AND ZHOU, K. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Trans. Graph. 33*, 4 (July), 95:1–95:11.

CREATIVEBLOQ, 2012. The making of Pixar Brave. http://www.creativebloq.com/making-pixars-brave-8123080, August.

DAVIET, G., BERTAILS-DESCOUBES, F., AND BOISSIEUX, L. 2011. A hybrid iterative solver for robustly capturing Coulomb friction in hair dynamics. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia'11) 30*, 139:1–139:12.

DEROUET-JOURDAN, A., BERTAILS-DESCOUBES, F., AND THOLLOT, J. 2010. Stable inverse dynamic curves. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia'10) 29* (December), 137:1–137:10.

DEROUET-JOURDAN, A., BERTAILS-DESCOUBES, F., DAVIET, G., AND THOLLOT, J. 2013. Inverse dynamic hair modeling with frictional contact. *ACM Trans. Graph. 32*, 6 (Nov.), 159:1–159:10.

ENGLISH, E., AND BRIDSON, R. 2008. Animating developable surfaces using nonconforming elements. *ACM Transactions on Graphics (Proc. of SIGGRAPH'08) 27*, 3, 1–5.

GILES, M., AND PIERCE, N. 2000. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion 65*, 3-4, 393–415.

GOLDENTHAL, R., HARMON, D., FATTAL, R., BERCOVIER, M., AND GRINSPUN, E. 2007. Efficient simulation of inextensible cloth. In *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'07*, ACM, New York, NY, USA, SIGGRAPH '07.

GRINSPUN, E., HIRANI, A., DESBRUN, M., AND SCHRÖDER, P. 2003. Discrete Shells. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'03)*, ACM-EG SCA, 62–67.

HAHN, F., THOMASZEWSKI, B., COROS, S., SUMNER, R., COLE, F., MEYER, M., DEROSE, T., AND GROSS, M. 2014. Subspace clothing simulation using adaptive bases. *ACM Trans. Graph. 33*, 4 (July), 105:1–105:9.

HASLER, N., ROSENHAHN, B., AND SEIDEL, H.-P. 2007. Reverse engineering garments. In *Computer Vision/Computer Graphics Collaboration Techniques*, A. Gagalowicz and W. Philips, Eds., vol. 4418 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 200–211.

HIRIART-URRUTY, J.-B., AND LEMARÉCHAL, C. 1993. *Convex Analysis and Minimization Algorithms I: Fundamentals*. Grundlehren Text Editions. Springer.

IGARASHI, T., AND HUGHES, J. 2003. Clothing manipulation. *ACM Trans. Graph. 22*, 3 (July), 697–697.

JOJIC, N., AND HUANG, T. 1997. Estimating cloth draping parameters from range data. In *In International Workshop on Synthetic-Natural Hybrid Coding and 3-D Imaging*, 73–76.

KAUFMAN, D., TAMSTORF, R., SMITH, B., AUBRY, J.-M., AND GRINSPUN, E. 2014. Adaptive nonlinearity for collisions in complex rod assemblies. *ACM Trans. on Graphics (SIGGRAPH 2014)*.

KIM, D., KOH, W., NARAIN, R., FATAHALIAN, K., TREUILLE, A., AND O'BRIEN, J. 2013. Near-exhaustive precomputation of secondary cloth effects. *ACM Transactions on Graphics 32*, 4 (July), 87:1–7. Proceedings of ACM SIGGRAPH 2013, Anaheim.

LIU, Y.-J., ZHANG, D.-L., AND YUEN, M. 2010. A survey on {CAD} methods in 3d garment design. *Computers in Industry 61*, 6, 576 – 593. Soft Products Development.

MARVELOUS DESIGNER, 2010. Marvelous Designer. http://www.marvelousdesigner.com.

NOCEDAL, J., AND WRIGHT, S. 2006. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer.

PORUMBESCU, S. D., BUDGE, B., FENG, L., AND JOY, K. I. 2005. Shell maps. *ACM Trans. Graph. 24*, 3 (July), 626–633.

THOMASZEWSKI, B., PABST, S., AND STRASSER, W. 2009. Continuum-based strain limiting. *Computer Graphics Forum (Proc. Eurographics'09) 28*, 2 (apr).

TURQUIN, E., WITHER, J., BOISSIEUX, L., CANI, M.-P., AND HUGHES, J. 2007. A sketch-based interface for clothing virtual characters. *IEEE Comput. Graph. Appl. 27*, 1 (Jan.), 72–81.

TWIGG, C., AND KAČIĆ-ALESIĆ, Z. 2011. Optimization for sag-free simulations. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'11)*, ACM-EG SCA, 225–236.

UMETANI, N., KAUFMAN, D., IGARASHI, T., AND GRINSPUN, E. 2011. Sensitive couture for interactive garment editing and modeling. *ACM Transactions on Graphics (SIGGRAPH 2011) 30*, 4.

VOLINO, P., AND MAGNENAT-THALMANN, N. 2007. Stop-and-go cloth draping. *The Visual Computer* (August), 669–677.

VOLINO, P., CORDIER, F., AND MAGNENAT-THALMANN, N. 2005. From early virtual garment simulation to interactive fashion design. *Computer-Aided Design Journal (CAD journal) 37* (March), 593–608.

WANG, H., RAMAMOORTHI, R., AND O'BRIEN, J. 2011. Data-driven elastic models for cloth: modeling and measurement. *ACM Transactions on Graphics (SIGGRAPH 2011) 30*, 4 (Aug.), 71:1–71:12.

WHITE, R., CRANE, K., AND FORSYTH, D. 2007. Capturing and animating occluded cloth. *ACM Transactions on Graphics 26*, 3 (July).