

Physical validation of simulators in Computer Graphics: A new framework dedicated to slender elastic structures and frictional contact

VICTOR ROMERO, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

MICKAËL LY, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

ABDULLAH HAROON RASHEED, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

RAPHAËL CHARRONDIÈRE, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

ARNAUD LAZARUS, Sorbonne Université, CNRS, Institut Jean Le Rond d'Alembert, UMR 7190, France

SÉBASTIEN NEUKIRCH, Sorbonne Université, CNRS, Institut Jean Le Rond d'Alembert, UMR 7190, France

FLORENCE BERTAILS-DESCOUBES, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France




Fig. 1. Four simple but rich physical protocols we introduce to Computer Graphics in order to verify the physical correctness of codes for thin elastic rods and plates, as well as dry frictional contact. For each test, our experimental data is shown side-by-side with a validated simulation output. From left to right: the **Cantilever** test (DISCRETE ELASTIC ROD simulation), the **Bend-Twist** test (SUPER-HELIX simulation), the **Lateral Buckling** test (FENICS SHELL simulation), and the **Stick-Slip** test (ARGUS simulation). While experimental and synthetic geometries indeed match, our validation protocols are actually much richer than plain geometric comparisons. Indeed, they are characterized by scaling laws which only depend on a few dimensionless parameters, making them ideal for benchmarking robustly a large diversity of codes across different physical regimes, without having to worry about scales or dimensions.

We introduce a selected set of protocols inspired from the Soft Matter Physics community in order to validate Computer Graphics simulators of slender elastic structures possibly subject to dry frictional contact. Although these simulators were primarily intended for feature film animation and visual effects, they are more and more used as virtual design tools for predicting the shape and deformation of real objects; hence the need for a careful, quantitative validation. Our tests, experimentally verified, are designed to evaluate carefully the predictability of these simulators on various aspects, such as bending elasticity, bend-twist coupling, and frictional contact. We have passed a number of popular codes of Computer Graphics through our benchmarks by defining a rigorous, consistent, and as fair as possible methodology. Our results show that while some popular simulators for

plates/shells and frictional contact fail even on the simplest scenarios, more recent ones, as well as well-known codes for rods, generally perform well and sometimes even better than some reference commercial tools of Mechanical Engineering. To make our validation protocols easily applicable to any simulator, we provide an extensive description of our methodology, and we shall distribute all the necessary model data to be compared against.

CCS Concepts: • **Computing methodologies** → **Animation; Physical simulation**;

Additional Key Words and Phrases: Slender elastic structures, dry frictional contact, code benchmarking, experimental validation

ACM Reference Format:

Victor Romero, Mickaël Ly, Abdullah Haroon Rasheed, Raphaël Charrondière, Arnaud Lazarus, Sébastien Neukirch, and Florence Bertails-Descoubes. 2021. Physical validation of simulators in Computer Graphics: A new framework dedicated to slender elastic structures and frictional contact. *ACM Trans. Graph.* 40, 4, Article 66 (August 2021), 19 pages. <https://doi.org/10.1145/3450626.3459931>

1 INTRODUCTION

In the last decades, a number of numerical simulators have been provided by the Computer Graphics (CG) community to compute the statics and dynamics of slender elastic structures, possibly subject to frictional contact. Initially, these codes were primarily intended for the movie industry where the important evaluation criteria are visual realism and artistic authoring. Such simulators have been successfully leveraged to animate, for instance, hair, fur and cloth of

Authors' addresses: Victor Romero, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France, victor.romero@inria.fr; Mickaël Ly, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France, mickael.ly@inria.fr; Abdullah Haroon Rasheed, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France, haroon.rasheed@inria.fr; Raphaël Charrondière, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France, raphael.charrondiere@inria.fr; Arnaud Lazarus, Sorbonne Université, CNRS, Institut Jean Le Rond d'Alembert, UMR 7190, 4 place Jussieu, case 162, F-75005, Paris, France, arnaud.lazarus@upmc.fr; Sébastien Neukirch, Sorbonne Université, CNRS, Institut Jean Le Rond d'Alembert, UMR 7190, 4 place Jussieu, case 162, F-75005, Paris, France, sebastien.neukirch@upmc.fr; Florence Bertails-Descoubes, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000, Grenoble, France, florence.descoubes@inria.fr.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2021/8-ART66 \$15.00

<https://doi.org/10.1145/3450626.3459931>

virtual characters [Baraff and Witkin 1998; Baraff et al. 2003; Daviet 2020; Daviet et al. 2011; Kaufman et al. 2014; McAdams et al. 2009].

Yet nowadays, the growing demand of the movie industry for mixing virtual and real shots in a seamless fashion constantly pushes forward the limits of simulation in terms of realism, up to some quantitative matching with reality [Museth 2020]. Besides, with the advance of computational fabrication, CG simulators are now more and more used for virtual prototyping tasks, where the ultimate goal is not to make images, but to build actual objects with desired properties. Among recent developments in this area, we can cite cloth material identification and prototyping [Bartle et al. 2016; Wang 2018; Yang et al. 2017], architectural design [Gavriil et al. 2020; Laccone et al. 2019; Panetta et al. 2019], fabrication of soft robots [Coevoet et al. 2019; Vanneste et al. 2020], and the exploration of new meta-materials [Guseinov et al. 2020; Martínez et al. 2019; Schumacher et al. 2018]. Furthermore, in various communities ranging from Computer Vision to Biomedicine, a number of works now study how to learn physics from simulated data in order to predict reality [Liang et al. 2019; Yang and Lin 2016]; for training, CG simulators are often preferred to commercial, specialized FEM softwares, because in general the former offer a free, lighter, and more user-friendly environment compared to the latter.

Common to all these new applications is the concern to *predict* reality in a *quantitative* way – this is in contrast with traditional image making which instead focuses on capturing the real world qualitatively. Unfortunately, tools to measure the physical correctness and accuracy of simulators remain scarce in the Computer Graphics community. As a result, there is *no guarantee* that current CG simulators offer the predictability level required by the target application. While not critical in movie making, wrong or unreliable predictions may yield serious consequences in some more sensitive domains like architectural design or healthcare.

Our goal in this paper is to improve the reliability of CG simulators so as to push forward their usage not only in hyper-realistic special effects, but also in material design. To this aim, we introduce a set of simple yet rich and compact physical protocols, originally designed to measure physical parameters in Soft Matter Physics, which can, in turn, be run easily to test the physical correctness of a model, and the validity of a code. Defining a simple yet robust evaluation metric, we test a number of popular codes from Computer Graphics and Computational Mechanics, and carefully analyze our results.

Our study is focused on static scenarios involving slender elastic structures (rods and plates) subject to clamped-free boundary conditions, gravity, and possibly natural curvature and frictional contact. The methodology we propose here is however general and could be extended to other situations. We hope our approach will be useful in the future to all the people who wish to design a new model or reimplement an existing one, by allowing them to check the physical validity of their model or/and of their code.

2 RELATED WORK

Unlike Mechanical Engineering which has a long tradition in physical and numerical validation, Computer Graphics is not used to evaluating its models quantitatively in a systematic manner. A first reason is that for movie and game making – the historical applications of Computer Graphics – many subjective criteria enter the

loop, such as visual perception, authoring, or aesthetics; all of them remaining extremely hard to assess. A second reason is that phenomena of interest have long differed between Mechanical Engineering and Computer Graphics. While Computer Graphics was from its early stage interested in large displacements of structures, buckling, tearing, entanglement, and dynamics, all these phenomena were long considered as undesirable behaviors in Mechanical Engineering, thus restraining the search for corresponding numerical models and appropriate validation tools.

Since these early times, both fields have considerably evolved. While many researchers from Computer Graphics have turned from pure visual effect to material design applications – the latter requiring more predictability than the former – a sub-community of Mechanical Engineering now considers so-called extreme mechanics [Krieger 2012] as a new fruitful direction for research [Reis 2015]. However, validation methods applicable to the simulation of complex (or extreme) phenomena, such as post-buckled rods/plates and frictional contact, remain scarce. We review here the main practices used in both fields.

Overly simple scenarios. For planar rods, elementary tests can be performed under the small deflection assumption, comparing code results against solutions to the linear Euler-Bernoulli beam equations. Various scenarios can be considered, such as the linear cantilever beam experiment (clamped-free or clamped-clamped boundary conditions) or the multiple point bending test (horizontal beam subject to multiple vertical pressure points), which all generate some simple analytical solutions (see e.g. [Timoshenko 1953]). For instance, Martin et al. [2010] compared their unified model for rods and shells to the linear cantilever test, and Panetta et al. [2019] validated their rod model implementation using the four point bending test, characterized by a constant bending strain between the two downwards pressure points. However, such tests are not valid anymore for large displacements of the structure, which is the case we are interested in. In this paper we considerably enrich the well-known linear cantilever scenario by considering nonlinear clamped rods and plates subject to possibly large gravitational forces (**Cantilever** test and **Lateral Buckling** test, respectively), as well as bending/twisting instabilities emerging for naturally curved structures (**Bend-Twist** test).

For frictional contact, a few authors have devised some simple tests like the falling drape over a plane [Li et al. 2018], which generates a constant velocity motion but only tests non-penetration as well as sliding friction. We go one step further by testing the sticking-sliding threshold (**Stick-Slip** test), independently of the material properties.

Qualitative only scenarios. To increase the complexity of validation scenarios, researchers in Computer Graphics soon departed from quantitative evaluation and turned instead to a qualitative assessment of their models. Progressively, a number of tests like cloth on rotating sphere [Bridson et al. 2002], the cloth funnel [Harmon et al. 2008], the reef knot [Harmon et al. 2009], or rod plectonemes [Spillmann and Teschner 2007], became popular and were adopted for subsequent benchmarking studies [Bergou et al. 2008, Sec. 9.2] [Li et al. 2018, 2020], even though these remained merely visual. Notable exceptions to this trend are a few studies mostly conducted in collaboration with physicists, which can be found on

rods [Bergou et al. 2008, Sec. 9.1], viscous threads [Bergou et al. 2010], and rigid body impacts [Smith et al. 2012]. In particular, Bergou et al. [2008] validate the DISCRETE ELASTIC ROD model on the Michell (or Zajac)’s buckling experiment, which depicts a characteristic 2D/3D instability in closed twisted filaments [Goriely 2006]. The **Bend-Twist** buckling test we introduce and leverage in this paper is similar in spirit [Miller et al. 2014], but has the advantage of extending the **Cantilever** experiment naturally as it features a very similar setting (gravity, same clamped-free boundary conditions) while adding a new degree of freedom to the system (naturally curly shape). This natural progression from one test to another makes our full validation suite consistent and practical to address as a whole.

Data-driven simulation & measurement protocols. Another common practice in Computer Graphics to verify the applicability of a code to material design is to feed it with real data. By adapting standard measurement protocols, such as for instance the ones provided by the ASTM International or the Kawabata Evaluation System (KES) [Kawabata and Niwa 1989] for cloth, or by creating new protocols suited for the model used [Wang et al. 2011], the aim is to fit model parameters to measurements so that the geometric output of the code matches the geometric output of the experiment. The underlying model does not need to be entirely physical but should possess enough parameters so that it can capture and reproduce the diverse phenomena at play in the experiments. This strategy has been widely used for the fast simulation of soft tissues [Bickel et al. 2009], the characterization of garments properties [Clyde et al. 2017; Miguel et al. 2012; Wang et al. 2011], or the design of inflatable structures [Skouras et al. 2012, 2014] and materials with desired properties [Bickel et al. 2010]. Of course, the next stage amounts to developing completely blind models through machine learning.

However, validation is much more than fitting: it requires setting physical parameters to their value *measured* in an *independent* fashion. Compared to validated physical models, fitted models are likely to possess a limited degree of predictability. Interestingly, we show in this paper that while our methodology is dedicated to the evaluation of physical numerical models, it can also be exploited to assess the potential of a fitted or partially fitted model.

Particular dimensional studies. In Mechanical Engineering, probably the most popular technique to validate a code quantitatively is to compare it against particular physical experiments spread across the literature. For instance, in the mechanics of plates and shells, well-defined scenarios have been defined over time [Sze et al. 2004], such as the sheared hemispherical shell or the pinched cylindrical shell. Such benchmark problems can be complex, but in general they rely on a particular, dimensional choice of parameters. As a result, although the test can be hard to set up in practice, it only assesses a particular setting of the code. Moreover, with these scenarios it can be tricky to compare various, disparate codes in a robust way.

Towards a compact quantitative validation. In contrast to particular dimensional benchmarks, each one of our tests includes a dimensionless *scaling law* which compactly represents several regimes of a complex physical experiment, hence describing a wide parameter range. Such a law can be verified against the output of any simulator able to reproduce the corresponding experimental setup, should the

simulator be dimensionalized or dimensionless. This verification actually assesses many facets of the simulator at once.

Our main inspiration originates from the work of Rasheed et al. [2020] in Computer Vision. In their study, they train a neural network on a cloth simulator in order to predict friction coefficients in real cloth motion. To make sure their network learns from realistic data, they first validate their simulator against the stick-slip scaling law recently proposed in Soft Matter Physics by Sano et al. [2017].

We build on Rasheed et al.’s validation idea and significantly extend it by proposing a set of four measurement protocols (among which is Sano et al.’s protocol) as new tests for code validation. We revisit in depth each one of these tests, reconstruct the corresponding scaling law accurately (adapting an existing one and devising a new one), and validate them experimentally.

3 CONTRIBUTIONS

We evaluate the *physical realism* of popular simulators in Computer Graphics for simulating thin elastic rods and plates, possibly subject to contact and dry friction. To this end, we carefully select and adapt four measurement protocols from the Soft Matter Physics community, never introduced to Computer Graphics so far (Section 4). These protocols, namely **Cantilever**, **Bend-Twist**, **Lateral Buckling**, and **Stick-Slip**, are meant to assess various physical aspects of a simulator, such as bending, bend-twist coupling, or frictional contact (see Figure 1). Note that the **Cantilever** test we introduce in this paper is a *dimensionless* test, meaning that it is able to characterize full deformation *regimes* (depending on a combination of parameters), and not just particular deformations based on specific parameters (e.g. a particular material). In this respect, the **Cantilever** test is thus much richer than the *dimensional* cantilever test usually leveraged in Computer Graphics for validation or parameter fitting [Miguel et al. 2012; Wang et al. 2011].

We thoroughly derive scaling laws associated to each protocol (the one for **Lateral Buckling** being new), and perform careful experiments to validate theoretical predictions. These scaling laws are made to be sufficiently simple and rich for revealing the capacity of a given numerical model to reproduce and predict reality. We then select a number of popular codes from Computer Graphics as well as two reference codes from Mechanical Engineering (Section 5) and define a thorough methodology for benchmarking (Section 6). We give a (positive or negative) score to each code evaluated on all applicable protocols. Through a careful analysis of our results, we show that our tests help discriminate the codes sharply (Section 7). We finally address limitations of our approach (Section 8), before concluding. Our protocols and data are made available to the community as supplementary materials, so that a thorough evaluation of additional codes can be performed easily by other researchers.

Caveat: What this study is not about. We do not evaluate performance nor user-friendliness of a code, even though some rough speed estimation of use and run are mentioned for merely indicative purposes. Neither do we evaluate the versatility of a code. Some codes are specific to one task (e.g. a static planar rod simulator), others might be more general (e.g. a dynamic cloth simulator). The range of phenomena which can be simulated by a code is not taken into account in the overall score attributed to a code.

4 SELECTION & DESIGN OF PHYSICAL PROTOCOLS

We present here the four protocols that will serve to construct our benchmarking study. All the details necessary to reproduce our experiments are available in our supplementary document (Section 3).

4.1 Scaling laws

Dimensional analysis [Buckingham 1914] and scaling laws are useful in Physics to understand the nature of phenomena. It is also a powerful tool for experimentalists who use it to measure unknown parameters. Yet we show here that scaling laws also have a high potential in terms of code validation. Consider for example the period $T_{pd} = 2\pi\sqrt{\ell/g}$ of the small-amplitude oscillations of a pendulum of length ℓ , where g is the gravity acceleration. If one were to test a simulator for the dynamics of such a pendulum, one could either choose several values of ℓ and g and compare the numerical period to the formula $T_{pd} = 2\pi\sqrt{\ell/g}$, or plot the numerical period as a function of ℓ/g and verify that it follows a square-root law. The second method has several advantages: (i) it illustrates that two different numerical tests should yield the same period provided the ratio ℓ/g is the same in both tests, (ii) it can detect some discrepancies in a simulator that would fail in a certain parameter range (small ℓ/g for example), and (iii) it allows one to compare simulators working with dimensional parameters (i.e. having actual units, for example meters for ℓ) to simulators working with dimensionless parameters (for example using the dimensionless time $\bar{t} = t/\sqrt{\ell/g}$ and hence finding the dimensionless period $\bar{T}_{pd} = 2\pi$). In conclusion, using scaling laws allows us to test, on a unique set-up, a complete family of parameter values with all points collapsing on a single curve, which we shall refer to as the *master curve* in this paper.

Building a set-up with a simulator is a long process and it is more efficient if this set-up can serve to test multiple configurations of parameters at once. Once the code for reproducing a given experiment has been written, it is relatively straightforward to launch a large number of simulations in batches to sample the parameter range automatically. Having the results collapsing on a master curve is then very rich in terms of information. The readability of the results is considerably higher than in the case of a particular dimensional study where one would compare, for each run, 2 simulations or 1 simulation and 1 experiment.

4.2 The Cantilever test

We discuss a 2D test dating back to the 30's [Bickley 1934], which is commonly used by the Soft Matter community for inferring mechanical parameters of rods and ribbons, see e.g. Duclaux [2006] or Fargette [2017]. The main advantage of this test is that, from purely geometrical considerations and measurements, one can estimate, with good accuracy, the elastic behaviour of the material. Hereafter we show that the same geometrical considerations can be harnessed either for intensive validation of physics-based simulators, or for calibration of non physics-based simulators.

The soft cantilever consists of a slender, naturally straight rod or ribbon, clamped horizontally, and deformed in 2D by the action of gravity, see Figure 2, left. For a rod of radius r , the parameters involved in the problem are the rod bending rigidity (EI), its material density (ρ), its length (L), the area $A = \pi r^2$ of its cross-section, and the gravity acceleration (g). Here E is the Young elastic modulus

and $I = \frac{\pi}{4}r^4$ the second moment of area of the cross-section of the rod. It turns out that one can construct a unique length scale for this problem, namely the gravito-bending length $L_{gb} = \sqrt[3]{EI/(\rho Ag)}$, which compares the resistance to bending with the gravitational force. Thus, for an elastic object deformed by gravity with length L , its equilibrium shape is effectively determined solely by the dimensionless gravito-bending parameter $\Gamma = (L/L_{gb})^3$, which reads

$$\Gamma = \frac{\rho Ag L^3}{EI}. \quad (1)$$

Note that in the case of ribbons, when the cross-section is a rectangle of width w and thickness h , with $h \ll w$, the Young modulus E becomes $E^* = E/(1-\nu^2)$, where ν is the Poisson ratio of the material [Shield 1992]. As a result, Γ is simply replaced with $\Gamma^* = (1-\nu^2)\Gamma$,

$$\Gamma^* = \frac{\rho A^* g L^3}{D w}, \quad (2)$$

where A^* is the area of the rectangular cross-section, $A^* = wh$, and D is the bending rigidity of the plate, $D = Eh^3/(12[1-\nu^2])$.

The parameter Γ gives an immediate notion of the expected behaviour for the deformation of an elastic object. Large values of Γ lead to large deformations, and conversely low values yield small deformations. However, the main advantage of this test is that one can quantitatively compare simulations with experimental and analytical results for a broad range of mechanical and geometrical parameters. For such comparison, here we chose the aspect ratio of the final shape, H/W (see Figure 2), and show that this aspect ratio is an accurate indicator of the quality of the simulation: indeed the mapping between Γ and H/W is unique (strictly monotonic), so there is no need to compare global equilibrium shapes or curvature distributions.

Master curve: solution of the planar elastica. The final shape for rods and ribbons is governed by the planar elastica equations (i.e. the Kirchhoff equations for 2D rods, see [Landau and Lifshitz 1959]). The equilibrium for the internal force \mathbf{F} and moment \mathbf{M} for an infinitesimal variation of the arc length s read $d\mathbf{F}/ds + \mathbf{f}_{ext} = \mathbf{0}$ and $d\mathbf{M}/ds + \mathbf{t} \times \mathbf{F} = \mathbf{0}$, with \mathbf{t} the unit tangent to the rod. Free-end boundary conditions reads $\mathbf{F}(s=L) = \mathbf{0} = \mathbf{M}(s=L)$. The constitutive relation between the bending moment and the curvature $d\theta/ds$ reads $M = EI d\theta/ds$, with θ the angle between the horizontal and the tangent, see Figure 2, left. When gravity is the only external force, $\mathbf{f}_{ext} = \rho Ag \mathbf{e}_y$, yielding $\mathbf{F} = \rho Ag(L-s)\mathbf{e}_y$ and $d\mathbf{M}/ds + \rho Ag(L-s)\cos\theta\mathbf{e}_z = \mathbf{0}$. Introducing dimensionless variables $\bar{s} = s/L$, $\bar{x} = x/L$, and $\bar{y} = y/L$, the boundary value problem to be solved finally reads

$$\frac{d^2\theta}{d\bar{s}^2} + \Gamma(1-\bar{s})\cos\theta = 0 \text{ with } \theta(0) = 0 \text{ and } \frac{d\theta}{d\bar{s}}(1) = 0 \quad (3a)$$

$$\frac{d\bar{x}}{d\bar{s}} = \cos\theta \text{ with } \bar{x}(0) = 0 \quad (3b)$$

$$\frac{d\bar{y}}{d\bar{s}} = \sin\theta \text{ with } \bar{y}(0) = 0 \quad (3c)$$

where the aforementioned parameter Γ arises naturally. For Γ values smaller than ~ 200 , this boundary value problem is easily solved with simple shooting techniques [Ascher et al. 1995]. However, as Γ

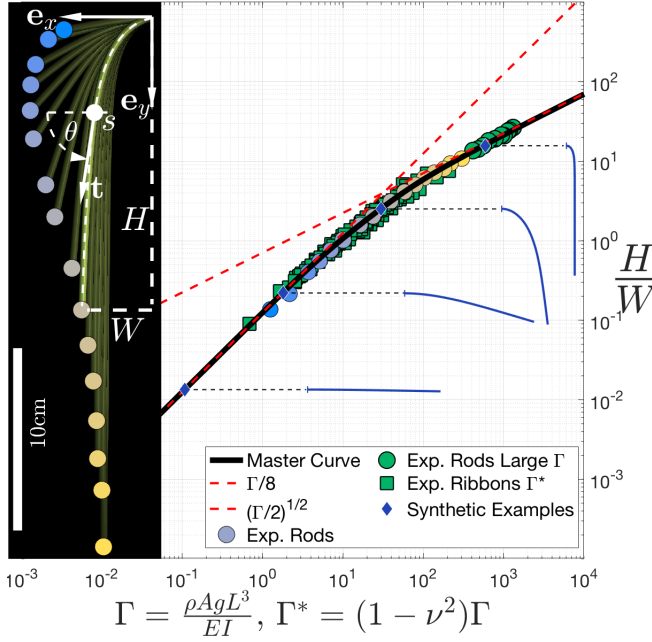


Fig. 2. **Master curve and experimental validation** for the Cantilever test. *Left*: snapshots of experimental rods under gravity, with 16 different Γ values. *Right*: our computed master curve (in black) representing the aspect ratio of the rod's shape as a function of Γ , together with our experimental results. After adjusting the value of the Young modulus (horizontal shift), our experimental data (circles for rods and squares for ribbons) are in perfect agreement with the master curve. Colored circles correspond to the experimental rods depicted on the left, the color providing a one-to-one correspondence between shapes and points of the graph. Four synthetic rods, simulated with DISCRETE ELASTIC ROD (see Section 5), are shown for illustration purposes.

gets larger the problem becomes stiff, and more powerful continuation methods are required. We use the collocation-based continuation package AUTO07p [Doedel et al. 1991] to compute the solution of (3) over 8 decades of Γ values, and we plot the aspect ratio $H/W = y(L)/x(L)$ as a black solid line in Figure 2, right. Furthermore, one can find analytical expressions for the aspect ratio for small and large Γ values, these are $H/W = \Gamma/8$ for $\Gamma \ll 1$ [Gere 2004], and $H/W = \sqrt{\Gamma/2}$ for $\Gamma \gg 1$, plotted as red dashed lines in Figure 2.

Experimental validation. In addition to the analytical and numerical solutions, we include experimental results. We fabricated a long rod by filling a straight PVC tube with vinylpolysiloxane. The rod was clamped at different lengths and the aspect ratio H/W was measured, see circles in Figure 2. We measured the cross-section radius $r \approx 2.0\text{mm}$ and the rod density $\rho \approx 1200\text{kg/m}^3$. The only uncertain parameter was the Young modulus, for which we only had an approximate range of values given by the manufacturer. Starting from an arbitrary value within this range, our measurements nicely collapsed onto a single curve having the exact same shape as the

master curve, albeit slightly shifted horizontally (E is inversely proportional to Γ). We could adjust the value as $E \approx 1.08\text{MPa}$ through a simple alignment of our measurements on top of the master curve¹.

We also present experiments on a ribbon cut from an acetate sheet of thickness of 0.17mm . The ribbon has dimension $3 \times 27\text{cm}^2$. It is clamped horizontally at different lengths and the aspect ratio H/W is reported as squares in Figure 2. The best alignment in this case was obtained for $E^* = E/(1 - \nu^2) \approx 4.19\text{GPa}$.

4.3 The Bend-Twist test

While planar equilibria of naturally straight rods are a good starting point for the mechanical validation and calibration of a simulator, most of the complex phenomenology observed for elastic rods arises from the coupling between bending and twist. A physical situation where this coupling is at play was described by Miller et al. [2014]. We consider a rod of length L with a circular cross-section, and we note A the area, and I the second moment of area, of the cross-section. The rod is made from an isotropic, uniform, elastic material of density ρ , but is cast in such a way that its natural shape is a twist-free circle of radius R . The rod is then clamped vertically at one extremity and is left sagging under the action of gravity, while its other extremity is free. Depending on the material and physical parameters, the suspended equilibrium shape can be planar (with no twist), or deformed into a 3D curly shape, see top panel in Figure 3.

Master curve: frontier between 2D and 3D configurations. The deformation of rods in 3D is governed by the Kirchhoff equations, see e.g. Audoly and Pomeau [2010]. We write the position $\mathbf{r} = (x, y, z)$, the tangent $\mathbf{t} = d\mathbf{r}/ds$, and the arc length $s \in [0, L]$. The local frame $\mathbf{t}(s), \mathbf{d}_1(s), \mathbf{d}_2(s)$ is used to follow the curvatures and twist deformations along the rod. We write τ the twist, $\kappa_{1,2}$ the two curvatures, and wrap these in a 3D vector $\kappa(s) = \{\tau, \kappa_1, \kappa_2\}^T$. In its natural, uniform configuration κ_0 , the rod has no twist and is curved in the plane orthogonal to \mathbf{d}_2 as $\kappa(s) = \kappa_0 = \{0, 0, 1/R\}^T$.

The local orthonormal frame follows the Darboux equation

$$d\mathcal{R}/ds = \mathcal{R}[\kappa]_{\times}, \quad (4)$$

where we use the matrix $\mathcal{R} = [\mathbf{t}|\mathbf{d}_1|\mathbf{d}_2]$ and the vector cross-product operator².

The twist and curvatures are found by integrating the balance of torques $\mathcal{K}_3 d\kappa/ds + [\kappa]_{\times} \mathcal{K}_3(\kappa - \kappa_0) = [\mathbf{e}_x]_{\times} \mathcal{R}^T \mathbf{T}(s)$ where $\mathbf{T}(s) = \rho Ag(s - L)\mathbf{e}_z$ is the internal force and $\mathcal{K}_3 = EI \text{diag}_3(1/(1 + \nu), 1, 1)$ the diagonal stiffness matrix. Introducing dimensionless variables

$$\bar{s} = \frac{s}{R}, \quad \varphi = \frac{L}{R}, \quad \bar{\kappa}(\bar{s}) = \kappa(s)R, \quad \bar{\kappa}_0 = \kappa_0 R, \quad \bar{\mathcal{K}}_3 = \mathcal{K}_3/EI, \quad (5)$$

we finally obtain the dimensionless equation

$$\bar{\mathcal{K}}_3 \frac{d\bar{\kappa}}{d\bar{s}} + [\bar{\kappa}]_{\times} \bar{\mathcal{K}}_3(\bar{\kappa} - \bar{\kappa}_0) = \frac{\Gamma}{\varphi^3}(\bar{s} - \varphi)[\mathbf{e}_x]_{\times} \mathcal{R}^T \mathbf{e}_z, \quad (6)$$

¹Actually, the quantity that we wish to measure accurately is the product EI . An adjustment on the **Cantilever** master curve proves more robust than using for instance a traction test, which would measure only the Young modulus E .

²For a 3D vector $\mathbf{u} = \{u_0, u_1, u_2\}^T$ the cross-product operator is represented by the skew symmetric matrix

$$[\mathbf{u}]_{\times} = \begin{bmatrix} 0 & -u_2 & u_1 \\ u_2 & 0 & -u_0 \\ -u_1 & u_0 & 0 \end{bmatrix}, \text{ such that } \mathbf{u} \times \mathbf{v} = [\mathbf{u}]_{\times} \mathbf{v}.$$




Fig. 3. **Master curve** ($\nu = 0.5$) and **experimental validation** for the **Bend-Twist** test. We use the following code for background color: orange when the rod shape is 2D, and turquoise when it is 3D. *Top*: our experimental rods under gravity, with different values of Γ and φ . *Bottom*: phase diagram where our numerical master curve (black) separates the 2D (orange) and 3D (turquoise) regimes. Our experimental points, at $\sqrt{\Gamma}/\varphi \approx 0.43$, with L from 2.5 cm to 14.6 cm, are drawn with large colored circles, the color labels identifying the experimental snapshots displayed on the top panel. We also plot (purple triangles) the numerically-found frontier points of Miller et al. [2014], together with their experimental points (small circles). All kinds of data fairly agree with each other. Four synthetic rods (A – D), simulated with SUPER-HELIX (see Section 5), are shown for illustration purposes.

with Γ defined in Section 4.2. Equations (4) and (6) are subject to the boundary conditions $x(0) = y(0) = z(0) = 0$ (clamping point located at the origin), $\mathcal{R}(0) = [-\mathbf{e}_z | -\mathbf{e}_x | \mathbf{e}_y]$ (vertical downwards clamping), and $\bar{\kappa}(\varphi) = (0, 0, 1)$ (circular natural shape). Both are integrated in the interval $\bar{s} \in [0, \varphi]$.

From Equation (6), we see that this test involves two dimensionless parameters: the gravito-bending parameter Γ , and the curvature parameter φ . Note that $\varphi/(2\pi)$ corresponds to the number of turns the rod makes in its natural shape³. The full analysis of this equation is out of reach for this paper. However the key observation is that despite the fact it is always possible to find a planar solution (twist $\tau(s) = 0, \forall s$), depending on the values of the parameters Γ and φ , it

³In the **Bend-Twist** protocol, rod self-contact is not accounted for.

might not be stable. There is indeed a specific region in the phase diagram where the planar solution is unstable and is then replaced by a 3D solution. The parameter range, see Figure 3, then has two well-defined regions. Planar solutions (in orange), are observed for low and high values of $\sqrt{\Gamma}/\varphi$ and, in contrast, 3D solutions (in turquoise) prevail for rod with large $\varphi/(2\pi)$, but are restricted to $\sqrt{\Gamma}/\varphi \lesssim 0.62$. Above this limit, the rod remains planar whatever its length and natural curvature, and for long length cases, the deformed shape is mainly straight with a hook near the free-end to comply with boundary conditions. The frontier between 2D and 3D solutions is characterized by a pitchfork bifurcation and we use AUTO07p (with $\nu = 0.5$) to numerically follow this bifurcation point and plot the transition curve in black in Figure 3.

Experimental validation. Curly rods were fabricated by injecting vinylpolysiloxane in a PVC tube coiled around a rigid cylinder and letting the elastomer cure for 24 hours [Miller et al. 2014]. We obtained samples with a circular cross-section of radius ≈ 1.8 mm, a natural curvature $1/R \approx 57.02 \text{ m}^{-1}$, and a small natural twist $\tau \approx 2.65 \text{ m}^{-1}$. This small intrinsic twist is due to the manufacturing process as soon as rods have $L > 2\pi R$, but it only slightly affects the results. We used the tabulated value of vinylpolysiloxane Poisson's ratio, $\nu = 0.5$, and we reused the Young's modulus value, $E \approx 1.08$ MPa, measured in Section 4.2. The density $\rho \approx 1267 \text{ kg/m}^3$ was obtained by simply weighting our samples.

We took snapshots of the rod clamped at different lengths, see top panel in Figure 3, and reported experimental parameters with large circles in the phase diagram, see bottom panel in Figure 3. Discrimination between 2D and 3D configurations was done visually⁴. We found a good agreement between our experimental measurements and the numerical 2D/3D transition. We complemented the diagram with experimental points (small circles) and numerically-computed frontier points (triangles) from Miller et al. [2014].

4.4 The Lateral Buckling test

We now turn to a test involving an elastic structure with a width w commensurated to its length L , namely an elastic plate. The thickness h of the plate is small compared to both L and w . As in the **Bend-Twist** test, we let the plate hang and sag under its own weight and wait for a (pitchfork) bifurcation to occur. In its reference configuration, the plate lies in the (x, z) plane with $0 \leq X \leq L$ and $-w \leq Z \leq 0$. Gravity is oriented along the $-\mathbf{e}_z$ direction. The plate is clamped along its $X = 0$ edge while the other three edges are free. The clamped orientation is vertical, along the z axis. In the deformed configuration, the plate can either stay in the (x, z) plane or deflect laterally in the $\pm \mathbf{e}_y$ direction and adopt a 3D shape. This lateral buckling typically occurs when gravity overcomes the elastic rigidity of the structure.

Lateral buckling instabilities are an important mode of failure in civil engineering where long steel beams carry important shear loads, and the history of lateral buckling analysis goes back to the

⁴The sharp reader may notice that some experimental rods (Figure 3 top right) are labeled as 2D whereas they slightly look 3D. This 3D shift is due to our residual natural twist stemming from fabrication constraints, and not to any buckling of the rod. We accounted for this residual when labeling the experimental rods as 2D or 3D (see also our supplementary document, Section 3).

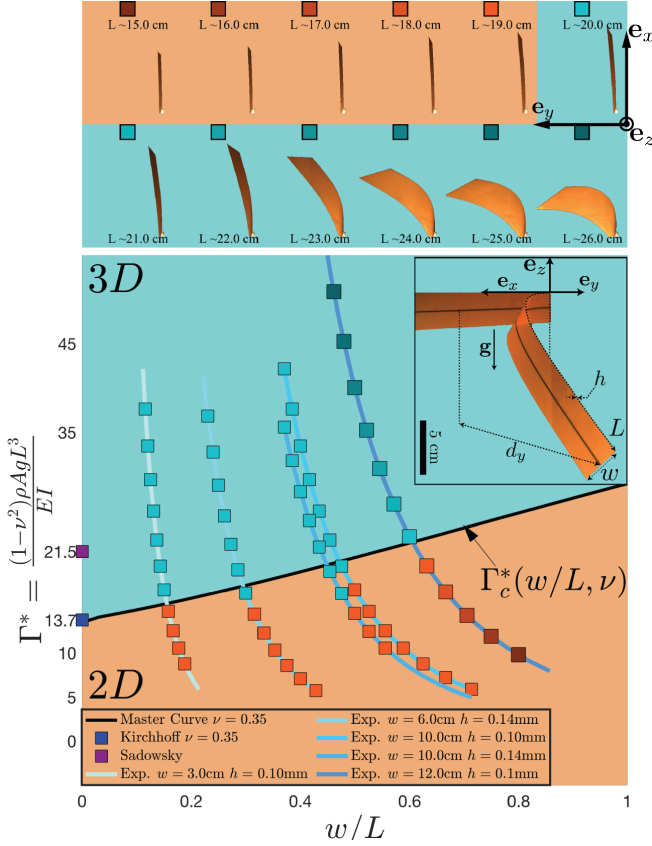


Fig. 4. **Master curve** ($\nu = 0.35$) and **experimental validation** for the **Lateral Buckling** test. *Top*: our experimental plates under gravity, with varying aspect ratio w/L . The background (originally black) has been colored to indicate whether the plate lies in 2D (orange) or has buckled in 3D (turquoise). *Bottom*: phase diagram where the computed master curve (in black) separates the 2D regime (in orange) from the 3D buckled regime (in turquoise). Five series of experiments are presented. Each series has fixed w and h values, and consists in increasing the length L of the suspended part of the plate. The first four experiments have their squares color-coded according to the 2D/3D nature of the equilibrium shape. For the uppermost experiment, the equilibrium shapes are shown in the top panel, and the color coding of the squares is used for correspondence.

19th century [Michell 1899]. Still, we are not aware of any study carefully analyzing and building up a collapsing curve (with only two dimensionless parameters) for this test. This is what we propose to achieve here.

As in previous sections we nondimensionalize, i.e. we rescale, lengths with L and forces with Dw/L^2 , where D is the bending rigidity of the plate, $D = Eh^3/(12[1 - \nu^2])$ (see Section 4.2). Once these two parameters are set away, we are left with the following free parameters: the Poisson ratio ν , the width w/L , the thickness h/L , and the total weight $\Gamma^* = Mg/(Dw/L^2)$, where the total mass of the plate is noted $M = \rho hwL$. Recall from Section 4.2 that Γ^* differs from $\Gamma = Mg/(EI/L^2)$ by a scaling factor, $\Gamma^* = (1 - \nu^2)\Gamma$, with $Dw = E^*I$.

Master curve: Our test experiment consists in holding the plate clamped as explained above, and increasing the parameter Γ^* while keeping the other parameters fixed. For $\Gamma^* = 0$ the plate is in its reference configuration. As Γ^* increases, the plate only suffers in-plane shear and extension, but stays in the (x, z) plane. Only when $\Gamma^* > \Gamma_c^*$ does the plate buckle in the third direction and normal curvature sets in. The critical Γ_c^* value depends on the other free parameters $\Gamma_c^* = \Gamma_c^*(w/L, h/L, \nu)$. To generate an interesting master curve, our idea was to plot Γ_c^* as a function of w/L , as shown in Figure 4, bottom panel. We work with thin plates $h/L \ll 1$, $h/w \ll 1$, that is in the regime where only little *transverse* (i.e. through the thickness) shear deformation is present. Consequently h is not a parameter anymore and we assume we have reached the limit where $\Gamma_c^* = \Gamma_c^*(w/L, \nu)$. To be on the safe side numerically, we nevertheless use the S8R thick shell finite-elements setup of ©ABAQUS to determine Γ_c^* . We set $\nu = 0.35$ (corresponding to polyester, which we use in experiments) and compute the threshold curve with the buckling eigenvalue analysis module of Abaqus using a high resolution mesh, i.e. for each w/L , we compute the critical Γ_c^* for which the total stiffness matrix, composed of the matrix of the base state and the tangent matrix, becomes singular. The results are plotted in Figure 4, using $h/L = 6 \times 10^{-4}$. We see in Figure 4 that $\Gamma_c^*(w/L = 0)$ corresponds to the value evaluated with the Kirchhoff elastic rod theory $\Gamma_c^*(\text{kirchhoff}) = 18.2(1 - \nu)\sqrt{1 + \nu}$ [Michell 1899], and that, remarkably, Γ_c^* increases approximately linearly (with slope ≈ 14.5) with w/L in the range $0 < w < L$ that we considered. Interestingly, we note that this simple linear dependence, extremely practical to evaluate numerical codes, actually matches the observations previously made in a lateral buckling problem with a different loading [Reissner 1995].

To evaluate the dependence of $\Gamma_c^*(w/L, \nu)$ w.r.t. the Poisson ratio ν , we computed the curve for several ν values and found them to be approximately parallel, each one emerging from the Kirchhoff $w = 0$ limit. We conclude that an approximated formula for Γ_c^* may be written as $\Gamma_c^*(w/L, \nu) \approx 18.2(1 - \nu)\sqrt{1 + \nu} + 14.5 \frac{w}{L}$.

Experimental validation. We perform experiments with naturally flat Poly-Styrene plates with thickness $h = 0.10$ or 0.14 mm. We measure densities to be $\rho_{0.10} = 1410$ and $\rho_{0.14} = 1260$ kg/m³. Using the **Cantilever** test, we measure the Young modulus to be $E_{0.10}/(1 - \nu^2) \approx 7.75$ GPa and $E_{0.14}/(1 - \nu^2) \approx 4.19$ GPa. For the Poisson ratio we use the value $\nu = 0.35$, tabulated at polymerdatabase.com. We then have $L_{gb,0.10}^* = 7.8$ cm and $L_{gb,0.14}^* = 8.2$ cm.

We use five samples of widths cut in the range $3 < w < 12$ cm, and having either $h = 0.10$ or 0.14 mm. For each sample, the plate is clamped vertically, with a suspended length L emerging from the clamp. Each experiment consists in increasing L from $L = 14$ to $L = 27$ cm. For each length value, we place a square on the phase diagram of Fig. 4 whose color indicates the 2D/3D nature of the equilibrium shape. Writing $y = \Gamma^*$ and $x = w/L$, each experiment follow a hyperbola of equation $y = (w/L_{gb})^3 1/x^3$.

Experimentally this is a challenging system as the **Lateral Buckling** instability is very sensitive to both the verticality of the clamp and the flatness of the plate. Despite these difficulties, we retrieve a fairly good agreement with the numerical prediction, as shown in the phase diagram in Figure 4.

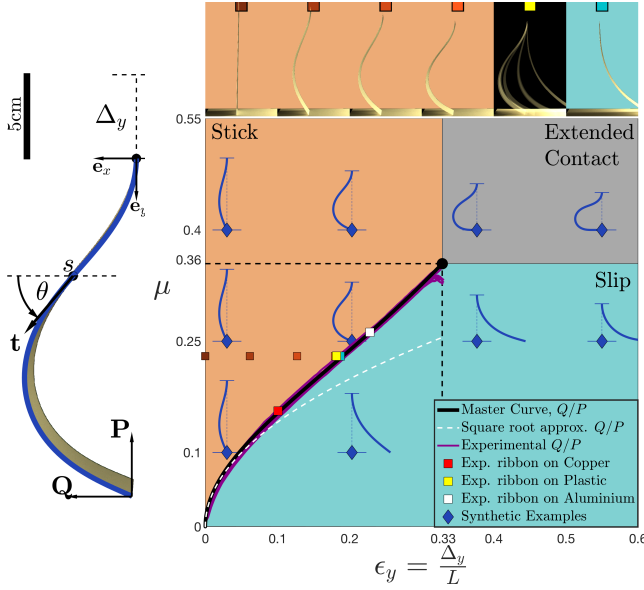


Fig. 5. **Master curve and experimental validation** for the **Stick-Slip** test. *Left*: Experimental snapshot of a **Stick-Slip** test with a high friction coefficient, for $\epsilon_y = 0.2$. The superposition of the numerical planar elastica solution (blue line), for the same value $\epsilon_y = 0.2$, illustrates the very good agreement between experiments and numerics. *Top-Right*: our experimental ribbons under frictional contact with the ground, with vertical deflection ϵ_y . The background (originally black) has been colored to indicate whether the rod is sticking (orange) or slipping (turquoise after slippage, black during the dynamic transition). *Bottom-Right*: phase diagram where the master curve (in black) separates the sticking regime (in orange) from the slipping regime (in turquoise). The extended contact region (in gray) is not used in our validation protocol. The (numerical) black and (experimental) purple curves are plots of the ratio Q/P as a function of ϵ_y . Ten synthetic rods (in dark blue), simulated with SUPER-HELIX 2D coupled with So-BOGUS (see Section 5), are depicted on the phase diagram for illustration purposes.

4.5 The Stick-Slip test

Up to now we have dealt with suspended elastic objects deformed by the action of gravity. However, in many applications in Computer Graphics and Computer Vision, elastic slender objects are supposed to interact with their surroundings. Thus, testing contact and frictional solvers is of foremost importance for assessing realistic scenarios.

Rasheed et al. [2020] recently introduced the work of Sano et al. [2017] as a way to evaluate the ARGUS code [Li et al. 2018] that they used for generating cloth training data. Inspired by their approach, we similarly leverage the Sano’s stick-slip criterion for precisely evaluating the realism of frictional contact solvers. Unlike them however, we re-build the mathematical formulation of the master curve from scratch and carefully verify it experimentally, with a more elaborated experimental set-up compared to [Sano et al. 2017]. Furthermore, we test multiple codes on this master curve, including plate and 2D rods simulators coupled with various frictional contact solvers.

In this so-called **Stick-Slip** test, a clamped elastic strip, of length L and rigid enough to neglect gravity effects ($\Gamma \ll 1$), is pushed quasi-statically against a solid substrate, see Figure 5. The strip and the substrate interact through normal (P) and tangential (Q) forces at the contact. As the vertical displacement (Δy) increases forces do as well.

Three well-defined regimes are observed in this test depending on the friction coefficient (μ) and the vertical strain ($\epsilon_y = \Delta y/L$): stick, slip, and extended contact (orange, turquoise, and gray regions respectively in Figure 5). In the stick phase, contact is localized at the lower end of the elastic strip and the contact point does not move. In the slip phase, friction is no longer able to hold the end of the strip which slips along the horizontal plane. Finally, in the extended contact regime the contact interface spreads to a larger portion of the elastic strip. These three regions are separated by (i) a *Master curve* (computed below), (ii) the line $\epsilon_y = \epsilon_{y,c} \approx 0.33$, and (iii) the line $\mu = \mu_c \approx 0.36$. These three border curves meet at the triple point $(\epsilon_{y,c}, \mu_c)$. The phase diagram of Figure 5 does not depend on any mechanical or geometrical property of the system. Therefore, one can compare frictional contact solvers over a large range of mechanical scenarios and test their ability to reproduce the phase diagram.

Master curve: transition between stick and slip regimes. To compute the **Stick-Slip** master curve, which partitions the stick and slip regions, we compute the equilibrium of a clamped planar elastica as the parameter ϵ_y is varied. The equilibrium solution gives the forces $\mathbf{P} = -P\mathbf{e}_y$ and $\mathbf{Q} = Q\mathbf{e}_x$ and friction is accounted for with the Amontons-Coulomb law. More precisely, once contact is made, the strip buckles as soon as the vertical force $P > 20.2 EI/L^2$. At buckling the horizontal force is $Q = 0$ and no friction is required for the equilibrium to hold. Nevertheless as ϵ_y is increased the ratio Q/P grows, and slip occurs as soon as Q/P reaches the Coulomb coefficient μ . The transition (master) curve is then reached when the clamped-pinned equilibrium is such that $Q/P = \mu$. We therefore compute the ratio Q/P at equilibrium, as function of ϵ_y and plot it in the phase diagram, see the black curve in Figure 5. Using the dimensionless arc length $\bar{s} = s/L$, equilibrium is found as the solution of the following boundary value problem,

$$\frac{d^2\theta}{d\bar{s}^2} = \frac{PL^2}{EI} \sin \theta + \frac{QL^2}{EI} \cos \theta \text{ with } \theta(0) = \pi/2, \frac{d\theta}{d\bar{s}}(1) = 0 \quad (7a)$$

$$\frac{d\bar{x}}{d\bar{s}} = \sin \theta \quad \text{with } \bar{x}(0) = 0, \bar{x}(1) = 0 \quad (7b)$$

$$\frac{d\bar{y}}{d\bar{s}} = \cos \theta \quad \text{with } \bar{y}(0) = 0, \bar{y}(1) = 1 - \epsilon_y^* \quad (7c)$$

where we note ϵ_y^* the value of ϵ_y on the transition curve.

Note that in Eq. (7a), P and Q only appear divided by the scaling force EI/L^2 and will therefore going to be proportional to it. Hence, the ratio $\mu = Q/P$ will be independent of the mechanical (EI) and geometrical (L) properties of the elastic strip.

One can analytically solve Eq. (7) up to a second order in θ by writing $\sin \theta \approx \theta$ and $\cos \theta \approx 1 - \theta^2$. With this weakly nonlinear approximation, the solution for the master curve reads

$$\frac{Q}{P} \approx 0.445\sqrt{\epsilon_y}. \quad (8)$$

This solution (dashed white line in Figure 5) is not accurate enough to serve in the present test. A detailed, semi-analytical, solution of this problem can be found, see e.g. [Mikata 2006], but involves elliptic functions and a root-finding step. Here, without any loss of accuracy, we simply solve system (7) numerically and plot Q/P as a function of ϵ_y , see the black curve in Figure 5. The triple point $(\epsilon_{y,c}, \mu_c) \simeq (0.33, 0.36)$ is reached when the equilibrium solution is such that $\theta(s = L) = \pi$.

Experimental validation. We quasi-statically compress a 260 μm thick, naturally flat, elastic strip of Poly-Styrene against a rigid substrate. The strip is clamped at the top end to a Futek load cell which measures the load P . The load cell itself is attached to a Thorlabs LTM300 motorised stage used to control the vertical displacement $\Delta_y = \epsilon_y L$. The substrate is mounted on two horizontal free-to-roll cylinders, to minimise resistance to horizontal displacement. A second Futek load cell holds the substrate horizontally and measures the load Q . An experiment starts at $\epsilon_y = 0$ and consists in increasing the vertical displacement with a series of 1mm steps, each followed by a wait-in-position step of 2 seconds where 2000 force measurements are performed and averaged, for both P and Q . At each displacement step we take a snapshot of the strip to compare its shape with numerical computations (see e.g. the image on Figure 5). We use different sample lengths ($L = 20.5$ to 28.5cm) and widths ($w = 3$ to $w = 4.5\text{cm}$) to explore different L^2/EI values. We stop experiments at vertical strain $\epsilon_y = 0.33$, and make sure no material plasticity was encountered during the test.

Depending on the frictional characteristics of the substrate we obtain two qualitatively different responses of the system.

On the one hand, for smooth substrates (copper, plastic, and aluminium), we observe the stick to slip transition for $\epsilon_y = \epsilon_y^*$. At this transition point, the force measurements yield a ratio Q/P which is in good agreement with the numerical **Stick-Slip** master curve, see square markers in Figure 5. Note that the value of Q/P at this point is a way to evaluate the Coulomb friction coefficient μ of the material. On the other hand, for rough substrates (e.g. sandpaper with grit size 180) the elastic strip first sticks then transitions into the *extended contact regime*. In this case, plotting the experimentally measured ratio Q/P as a function of ϵ_y yields the purple curve in Figure 5, which is in good agreement with the numerical master curve. For such substrates we cannot infer the friction coefficient μ from the experimental data, we only know that $\mu > \mu_c \simeq 0.36$.

5 SELECTION OF EXISTING CODES

We have selected a number of numerical models from Computer Graphics, which are well-known and widely used in the community. As a baseline for comparison, we also picked two reference codes from Mechanical Engineering, among which the widely used commercial software ©ABAQUS. The full list of the 14 codes used (together with some variants) is provided in Table 1, left column.

Of course, our selection is not exhaustive: it is rather meant to gather some representative examples. But we hope that the methodology we describe in this paper will encourage others to validate their own numerical model, which will help the community enrich these results.

5.1 Implementation

As far as possible, to limit implementation bias, we have striven to use the *original code* from the authors. This was possible for all but four of the academic numerical models we have selected, either by using the free implementation distributed online by the authors (SUPER-CLOTHOID, LIBSHELL, ARCSIM, SO-BOGUS, ARGUS, FENICSSHELL), or upon direct request to the authors (SUPER-HELIX, DISCRETE ELASTIC ROD, SUPER-RIBBON, PROJECTIVE FRICTION). The four exceptions comprise PROJECTIVE DYNAMICS, DISCRETE SHELL, BRIDSON-HARMON, and VISCOUS FRICTION. However, these four algorithms are simple enough so that we could reimplement them quite easily with a high level of confidence. Finally, for the commercial software ©ABAQUS, we had to use a license paid by our lab, and did not have access to the source code.

It is noteworthy that the codes we have tested are very disparate in terms of implementation, parameter input, and general behavior. In the supplementary document (Section 6) we summarize the main characteristics of all the codes we have tested.

5.2 Thin elastic rods

A staple of rod models in Computer Graphics is the DISCRETE ELASTIC ROD model [Bergou et al. 2010, 2008], which has been used for various applications ranging from hair simulation [Daviet 2020; Kaufman et al. 2014] to the design of meta-materials [Schumacher et al. 2018] and grid shells [Panetta et al. 2019]. The DISCRETE ELASTIC ROD (dynamic) model relies on a nodal formulation of the rod centerline tightly coupled to a material frame via some material twist around the discrete Bishop frame. While the first paper rigidly fixes the material frame as the one minimizing the potential elastic energy of the rod (assuming a vanishing cross-section), the second one considers in contrast some twist inertia and updates the material frame in time through parallel transport. The latter model has the advantage of yielding a sparse-block stiffness matrix at first order, allowing for fast implicit integration. The implementation that we have used was provided by the authors; it is based on this second method, which offers more stability than the first one.

We have also considered the family of curvature-based models, namely the SUPER-HELIX [Bertails et al. 2006] and the SUPER-CLOTHOID [Casati and Bertails-Descoubes 2013] models. In contrast to position-based formulations, these models rely on a curvature-based parameterization, allowing the formulation of the kinematics of the rod (i.e. bending and twisting modes) in a fully reduced way, without the need for inextensibility constraints. In the SUPER-HELIX model, each element is characterized by uniform material curvatures and twist and takes the form of a circular helix (or circular arc in 2D). In our test, we have used the 2D and 3D implementations provided by the authors (dynamic codes). Super-helices were successfully used in Computer Graphics for simulating the forward and inverse dynamics of human hair [Daviet et al. 2011; Derouet-Jourdan et al. 2013; Hu et al. 2017], and their reduced parameterization has been recently harnessed in the context of soft robotics [Boyer et al. 2020]. Besides, the approach was extended to the SUPER-CLOTHOID model, which features material curvatures and twist linearly varying along the elements. For our tests we took a static version (provided by the authors) of the reference (dynamic) implementation available online [Casati and Bertails-Descoubes 2013].

We initially intended to test further popular rod models from Computer Graphics [Hadap 2006; Pai 2002; Spillmann and Teschner 2007], but unlike for plates/shells, we had trouble finding reference implementations. As rod models are often tricky to implement, we found it too risky to use a non-reference implementation which may not fairly reflect the original published method. Hence, for the sake of fairness, we refrained from testing further rod models. However, our result table is meant to be dynamic and filled in by others with their own models and implementations.

5.3 Thin elastic plates, shells and ribbons

Plates and shells. For plate and shell computations, Computer Graphics codes compete in attempting to formulate simple yet accurate membrane and bending formulations, as well as their coupling. For in-plane deformations, linear elasticity has been proven to be computationally cheap and yet relies on strong theoretical bases ensuring its accuracy and showing its limits. However, for out-of-plane deformations, to the best of our knowledge, no model has been validated yet.

A widespread bending energy for plates or shells in Computer Graphics has been simultaneously proposed by Grinspun et al. [2003] and Bridson et al. [2003]. Although the former derived it from a discretization of the square of the difference of the mean curvatures and the latter from heuristics, both yield a similar (up to a scaling coefficient) hinge energy based on the dihedral angle⁵ between faces at their shared edges. The simplicity of this energy made it a popular way to measure the bending in Computer Graphics applications [Coros et al. 2012; Skouras et al. 2012; Wang et al. 2011]. Bridson et al. [2003]’s model is part of the ARCSIM implementation⁶ [Narain et al. 2013, 2012] and, introducing a slight modification of the bending coefficient (from $\frac{1}{4}$ to 3), we could readily switch to Grinspun et al. [2003]’s bending model. This modified code is denoted in the following as DISCRETE SHELL + ARCSIM. Note that for ARCSIM, we used the more recent ARGUS implementation [Li et al. 2018] for the sake of simplicity (without frictional contact, ARGUS boils down to ARCSIM). By default we set the adaptive option switched on, otherwise we refer to ARCSIM NON ADAPTIVE.

Another, more recent approach, comes from the discrete geometry community [Chen et al. 2018; Weischedel 2012]. By replacing metrics such as the fundamental forms used in the continuous description of shells by discrete equivalents, new numerical shell models were derived. The open-source library LIBSHELL [Chen et al. 2018] implements a discrete shell model based on Koiter’s thin shell energy following this principle. The code possesses a few variants, characterized by the choice of elastic behavior – Saint Venant-Kirchoff (SrV-K) or Neo-Hookean (NH) laws – and the choice of discretization scheme employed for the computation of the normals – one “first order” MIDEAGEAVERAGE and two “second order” MIDEAGESIN and MIDEGETAN schemes [Weischedel 2012]. Note that this overall approach departs from traditional finite-elements formulations which discretize quantities over functional spaces. By incorporating the [Grinspun et al. 2003]’s bending model into LIBSHELL, we

were able to test yet another variant of this approach, referred to as DISCRETE SHELL (+ LIBSHELL). We expect this variant to be close to DISCRETE SHELL + ARCSIM as, in principle, both models use the same in-plane and out-of-plane energies. Anticipating our results (Section 7), we however noticed some discrepancies between them.

Ribbons. Lying in between rods and plates, ribbons have recently regained some interest in the Soft Matter Physics community [Fosdick and Fried 2015]. Based on the Sadowsky-Wunderlich theory [Sadowsky 1929; Wunderlich 1962], and obtained through dimension reduction, equations for ribbons have been reformulated as enriched Kirchhoff rod equations [Dias and Audoly 2015]. Inspired by the SUPER-CLOTHOID model, an implementation of the Sadowsky-Wunderlich ribbon model has recently been proposed, based on elements with a linear normal curvature and a quadratic twist [Charondière et al. 2020]. This work has been published in Mechanical Engineering, but since it is deeply related to the super-model formulation introduced in Computer Graphics, we have included it in our study. Moreover, since the model lies in-between rods and plates, we found it particularly interesting to assess its relevance w.r.t. our rod and plate tests. We have used the implementation borrowed from the authors, that will be referred to as SUPER-RIBBON in the following.

5.4 Frictional contact

Perhaps one of the most popular model for frictional contact in Computer Graphics is the Bridson-Harmon algorithm [Bridson et al. 2002; Harmon et al. 2008], which explicitly resolves non-penetration and Coulomb constraints by successive filters applied to the object velocity. Such a model is known for yielding satisfying qualitative results (e.g. avoiding any interpenetration artifact) at a cheap cost.

In contrast, implicit constraint-based methods that were later introduced in Computer Graphics focus on physical accuracy and on the accurate capture of the Coulomb threshold. We have selected the SO-BOGUS solver [Daviat et al. 2011], where the contact forces are treated as unknowns and are tied to the velocities through complementarity constraints that are resolved implicitly. This solver has been successfully used to simulate large assemblies of hair [Kaufman et al. 2014] and cloth [Li et al. 2018]. The SO-BOGUS code is freely available as a standalone library, and for our tests we coupled it with the SUPER-HELIX 2D model following the work of Daviat et al. [2011]. We also use its coupling with the cloth simulator ARCSIM, also known as the ARGUS code, which is freely distributed by their authors [Li et al. 2018]. Additionally, we consider the much faster PROJECTIVE FRICTION algorithm, which yields results qualitatively similar to ARGUS [Ly et al. 2020]. We used the original code provided by the authors.

Finally, although the VISCOUS FRICTION model – which simply sets friction forces proportional to the object velocity – is not adapted to solid friction, it is quite often used for its simplicity. We thus decided to include it in our study to emphasize on its differences with dry frictional contact models.

5.5 Reference codes of Mechanical Engineering

Originally released in 1978, ©ABAQUS [Dassault-Systems 2005] is a widely used commercial software designed for the numerical modelling of solid and fluid mechanics as well as electromagnetic

⁵Bridson et al. [2003] rely on the sine of the dihedral angle whereas Grinspun et al. [2003] use the angle itself, which is equivalent in the limit of small angles.

⁶We used version 0.2.1 of ARCSIM, recommended by their authors over version 0.3.1 when shells and fracture are not needed.

problems. It is based on finite-elements analysis and supports various types of elements (including high-order elements for rods, plates and shells), both in static and dynamic settings. The user may choose between a graphical user interface and Python scripting for using the software. ©ABAQUS's robustness and capability to treat nonlinear problems involving contact, plasticity, etc., as well as explicit and implicit dynamics, have made it one of the reference tools in Mechanical Engineering, both in industry and academia. We therefore use it as a reference in our validation work, taking advantage of its many ready-to-use models, and despite the fact that most of its routines are closed.

In contrast FENICS is an open-source computing platform which comprises a C++ core engine interfaced with Python [Habera et al. 2018]. FENICS codes are easily parallelized and run on laptops, desktops, or clusters. The FENICS environment allows users to quickly enter various types of partial differential equations through the Unified Form Language (UFL) with which users convert mathematical models into finite-elements formulation in a straightforward way. FENICSSHELL is a set of libraries for solving linear and nonlinear plate and shell models [Hale et al. 2018]. Shear-locking is dealt with using several techniques, among which MITC (Mixed Interpolation of Tensorial Components) or PSRI (Partial Selective Reduced Integration, which we have used). As with UFL the user only has to write down the elastic energy, the physical model is easily changed and moreover several Newton and linear algebra solvers are accessible (including for exemple PETSC). The actual implementations for the **Cantilever** and **Lateral Buckling** tests are based on an example provided online [Brunetti et al. 2018].

6 METHODOLOGY FOR BENCHMARKING

Comparing a number of codes thoroughly is a particularly challenging task. Common and unconscious biases include treating codes heterogeneously, favoring one's code in place of others', or failing to give sufficient details for reliable replication by others. To avoid or at least limit these pitfalls, we have established a rigorous methodology which we strive to make reproducible.

6.1 As consistent as possible

Benchmarking codes requires a common methodology even though the set of tested codes may be fairly diverse. In our case for instance, some codes may be static while others dynamic, some may use dimensionless parameters while others dimensional data. To compare codes consistently, we should be able to free ourselves from peculiarities and offer clear and non-ambiguous guidelines that can be followed equally for each code to be tested.

To this aim, we advocate a (possibly iterative) two-step strategy:

- (1) First, prepare the simulation scenario, that is find a combination of *physical* parameters of the code C able to reproduce the test \mathcal{T} in the right physical regime;
- (2) Second, find, by continuation, the *optimal numerical* parameters of the code C able to "agree with the test".

After defining the meaning of "agreeing with the test" in terms of a test error, we formulate our final rule for assessing a code on a given test.

Test error. Quantifying the proximity of the results produced by a code C to the expected ones on a given test \mathcal{T} can be achieved by defining an error $\mathcal{E}(C, \mathcal{T})$, which measures a distance between the two. In practice, we compare the reference *master curve* of the test \mathcal{T} , over its full range, to the one generated by the code C . The code will be said to be in good agreement with the test result if the output of the code is indeed a continuous curve at our scale of observation, and if it matches the reference curve visually (i.e. $\mathcal{E}(C, \mathcal{T}) < \varepsilon_{\mathcal{T}}$ with a tolerance $\varepsilon_{\mathcal{T}} \approx 5\%$ of the plot size, in the sense of the Hausdorff distance).

Consistent solver calibration. For each test, we perform a large number of simulations by sampling a wide parameter range. From one parameter bound to another one (e.g. from a small value of Γ to a large one), the complexity of the simulations (e.g. small deformation vs. large deformation) may vary a lot, hence calling for different tunings of the discretization and numerical solver (number of elements, timestep, tolerance, etc.). However, to keep benchmarks as simple as possible, instead of building complex adaptive numerical strategies, we preferred to find out a common set of numerical parameters allowing for proper convergence of the solver over the *entire* parameter range. In other words, for a given code, the results reported stem from the same numerical configuration, and only the *physical* parameters vary.

Success versus failure. To simplify code evaluation, we have chosen a binary measure of success. A code C is assessed as successful (**OK**) on test \mathcal{T} if there exists a set of numerical parameters such that, for any refinement of these parameters and any combination of physical parameters p which lead to the same regime (or dimensionless data) $\mathcal{R}_{\mathcal{T}}$ of the test, the code output agrees with the master curve. Otherwise, it is considered as failing (**KO**) on test \mathcal{T} . Mathematically, this reads

$$\mathcal{T}(C) = \text{OK} \iff \exists n\text{Elts}^*, \bar{dt}^*, \bar{tol}^* \text{ such that} \\ \forall n \geq n\text{Elts}^*, \forall dt \leq \bar{dt}^*, \forall tol \leq \bar{tol}^*, \forall p \in \mathcal{R}_{\mathcal{T}}, \quad \mathcal{E}(C, \mathcal{T}) < \varepsilon_{\mathcal{T}}. (9)$$

As a result, a code C will be marked as **KO** on test \mathcal{T} if we fail to find a set of physical and numerical parameters passing the test \mathcal{T} , or, if we just find an arbitrary and isolated set of parameters which works, without any possible refinement.

Our success rule is thus highly demanding, since it requires, at the same time, a correct physical (dimensionless) behavior of the code ($\forall p \in \mathcal{R}_{\mathcal{T}}$), the numerical convergence in space and time of the numerical model ($\forall n \geq n\text{Elts}^*, \forall dt \leq \bar{dt}^*, \forall tol \leq \bar{tol}^*$), and the convergence towards the physical target ($\mathcal{E}(C, \mathcal{T}) < \varepsilon_{\mathcal{T}}$).

Note however that our rule does not discriminate between "fast" and "slow" codes, in terms of convergence. A code may be qualified of **OK** even though it may require a huge mesh resolution or/and an extremely small time step to reach the target, making such a configuration impractical. To compare code convergence speed (and thus their practicability), we hence adjoin the information of the numerical values ($n\text{Elts}^*, \bar{dt}^*, \bar{tol}^*$) to the test status (**OK**) in case of success.

Anticipating on the result section, the reader may take a look at Table 1, which summarizes all our test results.

6.2 As fair as possible

Dealing *equally* with some codes that are yours and some that have been written by others is difficult, for two main reasons. First, having a perfect knowledge of the model, its various options, and the underlying implementation details, of course helps understand how to tune things and quickly get the most out of the code. This process might be much slower and may eventually not converge to the optimal behavior when using a code that one does not master. Second, it is unconsciously always tempting to spend less efforts on others' code to make it work than on one's code, because of a lack of neutrality (even though one's honesty is unquestionable).

In the two cases mentioned above, the risk is clearly to underestimate the potential of others' code or method⁷. To mitigate this risk, we have adopted the principles listed below.

A set of well-experimented programmers. The authors of this paper who took in charge the benchmarking of Computer Graphics codes are only senior PhD students, postdocs or permanent researchers with good programming skills and a minimum experience of three years with complex programming in Computer Graphics.

Understanding of the underlying numerical model. Fortunately, most of the numerical models we have been using are well documented in open publications. Moreover, except for ©ABAQUS, we have access to all source codes, which has allowed us to check the conformity between the written model and its implementation. From the publications, we could also verify the limitations of the proposed algorithm, which generally matched our own observations. Overall, we were able, in most cases, to make sure that our observations were consistent (or at least, not contradictory) with the claims made in the original papers. The ©ABAQUS case was the most problematic one, because it is a closed code and its optimal usage may require specific training. To limit bias regarding our usage of this software, two authors of this paper were trained on ©ABAQUS for several weeks. Additionally, we have consulted several colleagues of Mechanical Engineering conducting their research with ©ABAQUS, who could share their expertise with us.

Overall, we estimate we did our best efforts to get the most of each software in our hands.

Systematic investigation in case of failure. After trying in vain parameter refinement or various tricks, when it comes that a code fails to a given test, we strive to analyze the origin of such a failure. And, if possible, we attempt to provide a (documented) fix to the code so that it may eventually pass the test. In our study, some of our fixes helped in succeeding a test, while others did improve results but not to the point of a success.

Future evaluations. Finally, despite all these precautions, there is theoretically always a small probability that we may have misleadingly categorized a code as a failure on a given test, because we may have failed to test all possible alternatives. In this case, we may have unwillingly produced false negatives. In contrast, given the high requirement level of each test, we believe that false positives are extremely unlikely to arise. Hence, successful codes can reasonably

be considered as *certified* by our study, under their precise condition of usage.

To mitigate the possible limitation on false negatives, we count on future benchmarking made by the community to refine and enrich our results. This is why reproducibility of our methodology is of foremost importance.

6.3 As reproducible as possible

Our goal is twofold. First, allow others to reproduce our benchmarking results (at least for publicly available codes). Second, allow others to validate rigorously their own models or codes, following our exact procedure. Our supplementary material "Recipe Manual for Validation" should give all the necessary algorithms and information to meet these two objectives. We also plan to deliver data regarding the reference master curves to compare against.

7 BENCHMARKING RESULTS

The general results of our benchmarking study are summarized in Table 1. As mentioned sooner, **OK** is the success score, while **KO** denotes a failure. A dash (–) indicates that the test is not applicable to the corresponding code, while a cross (×) shows that we have not tried to test the code because of expected failure. From these results one can already make a few important observations.

First, the most well-known rod models of Computer Graphics, namely DISCRETE ELASTIC ROD and SUPER-HELIX (and its higher-order counterpart SUPER-CLOTHOID), perfectly pass the two rod tests **Cantilever** and **Bend-Twist**. This is not surprising, as these models have benefited from physical insights from their construction stage. Furthermore, these codes have already been cross-validated [Casati and Bertails-Descoubes 2013] and used in other communities like Soft Matter Physics [Audoly and Pomeau 2010; Jawed et al. 2014].

In contrast, the popular adaptive model ARCSIM for plates fails on **Cantilever**, even though it is a mere 2D test. We carefully analyze such failure and attempt to provide a fix to the code.

Then, the only Computer Graphics code able to cope with the demanding **Lateral Buckling** test is LIBSHELL, although its success depends on the exact choice of the bending energy formulation. Note that FENICSHELL successfully passes the test as well.

Regarding frictional contact, most Computer Graphics solvers fail on the **Stick-Slip** test, mainly due to a lack of accuracy at convergence. The only successful Computer Graphics code is So-BOGUS, provided a low tolerance of the solver is used.

Finally, the commercial package ©ABAQUS, widely used in mechanics, performs moderately well. While it easily passes **Cantilever** and **Stick-Slip**, it surprisingly has a hard time on **Bend-Twist**, which is in contrast perfectly passed by the three Computer Graphics codes for rods that we have tested. It similarly fails on **Lateral Buckling**, where both LIBSHELL and FENICSHELL succeed.

Our results and observations are further detailed and commented on in the remainder of this section. Please also watch our accompanying video for an illustration of our benchmarking study. Many further information such as convergence plots or detailed experiments and results are provided in supplementary material. Moreover, our second supplementary material, called "Recipe Manual for Validation", explains all the rules, criteria and tips that we have been using to conduct the study thoroughly.

⁷Given the high level of requirement of each proposed test, we believe, in contrast, that overestimating results of one's code is impossible albeit being of bad faith. See also our discussion about false negatives and false positives in the end of this section.

Table 1. General result summary of all the benchmarks we have performed.

Tested Code	Cantilever	Bend-Twist	Lateral Buckling	Stick-Slip
Rod				
DISCRETE ELASTIC ROD [Bergou et al. 2010]	OK (300 elts)	OK (200 elts)	–	–
SUPER-HELIX [Bertails et al. 2006]	OK (50 elts)	OK (30 elts)	–	–
SUPER-CLOTHOID [Casati and Bertails-Descoubes 2013]	OK (20 elts)	OK (25 elts)	–	–
Ribbon				
SUPER-RIBBON [Charrondière et al. 2020]	OK (20 elts)	–	KO	–
Plate				
LIBSHELL [Chen et al. 2018]	OK (Res 0)	–	OK (Res +)	–
DISCRETE SHELL (+ LIBSHELL) [Grinspun et al. 2003]	OK (Res +)	–	KO	–
ARCSIM [Narain et al. 2012]	KO	–	×	–
DISCRETE SHELL + ARCSIM (tentative fix of ARCSIM)	KO	–	×	–
PROJECTIVE DYNAMICS [Bouaziz et al. 2014] (fit)	KO	–	×	–
Contact & friction				
VISCOUS FRICTION (+ SUPER-HELIX 2D)	–	–	–	KO
So-BOGUS [Daviet et al. 2011] (+ SUPER-HELIX 2D)	–	–	–	OK (dt=0.5 ms, tol = 10 ⁻¹³ N)
ARGUS (≈ ARCSIM + So-BOGUS) [Li et al. 2018]	–	–	–	KO
ARGUS NON ADAPTIVE (fix of ARGUS)	–	–	–	OK (dt=0.5 ms, tol = 10 ⁻¹³ N)
BRIDSON-HARMON [Bridson et al. 2002; Harmon et al. 2008] (+ ARCSIM)	–	–	–	KO
PROJECTIVE FRICTION [Ly et al. 2020]	–	–	–	KO
Reference codes in Mechanical Engineering				
FENICSHELL [Hale et al. 2018]	OK (Res 0, P _{2,3} elts)	–	OK (Res 0, P _{2,3} elts)	–
© ABAQUS	OK (200 P ₂ elts)	KO	KO	OK (dt=9 μs)


7.1 Results on the Cantilever test

The **Cantilever** test stands for an interesting 2D protocol to validate the bending formulation of all kinds of models – rods, ribbons, and plates – whether they are inextensible or not. Indeed, even though the master curve has been derived in the mere inextensible case, it remains valid in the presence of stretching within the range of Γ used, $\Gamma < 10^4$. When simulating a simple finite-difference rod model allowing for both bending and stretching, we were able to check that the in-plane elongation remains indeed negligible ($\approx 1\%$ at $\Gamma \approx 10^4$). In the codes allowing for stretching and shearing, we made sure to take a negligible thickness $h \ll w$ so that bending is the preferred deformation mode.

Even though it remains difficult to compare finely two models discretized with different approaches (e.g., finite-elements against models based on discrete differential geometry), we made sure we were using similar mesh resolutions for all the plate models that we have tested. For the **Cantilever** test we have typically used two kinds of meshes: a Res 0 mesh, made of 120 discretization nodes along the length, and a Res + mesh with a double-size resolution.

Figure 6 presents the results for all the tested codes. Except ARCSIM (whose results are shifted) and PROJECTIVE DYNAMICS (whose results are too scarce), all the codes are in excellent agreement with the master curve. However, for some bending plate models like DISCRETE SHELL, an increased resolution (Res +) was needed to cover the range $10^3 < \Gamma < 10^4$ properly. Computation time ranges between ≈ 1 h (for most models) and a few hours (for nodal models requiring a high number of elements or a high-resolution mesh).

Bridson et al. vs Grinspun et al.'s bending energy. Naturally, since ARCSIM's bending energy [Bridson et al. 2003] only differs from that of DISCRETE SHELL [Grinspun et al. 2003] by a multiplicative factor, both cannot pass the test. Our test thus allows to discriminate between the two factors, leading to the conclusion that Grinspun et al.'s factor is the correct one (see Figure 7, top).


 Fig. 6. Results on the **Cantilever** test.

However, we found the Grinspun et al. [2003]'s discrete energy to converge much more slowly compared to other mesh-based approaches like LIBSHELL or FENICSHELL. In practice, we had to use a twice as more resolute mesh (Res +) for DISCRETE SHELL (+ LIBSHELL) to observe a good agreement with the master curve.

Interestingly, Grinspun et al. [2006] report that DISCRETE SHELL does not converge in the general case, albeit in the case of triangular meshes with equilateral triangles. Our meshes were close to this ideal case, which probably explains our observation of convergence.

Investigations on ARCSIM. After fixing ARCSIM with the correct factor (DISCRETE SHELL + ARCSIM), we observe a reduction of the

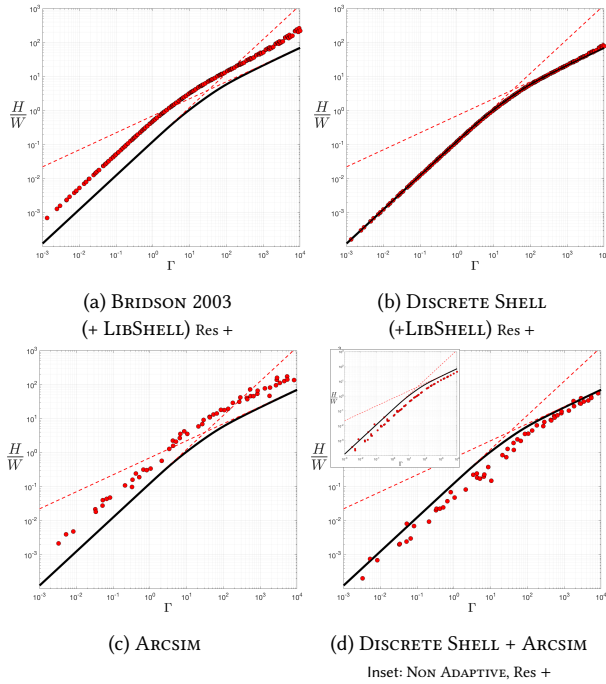


Fig. 7. Attempting to fix ARCSIM results on **Cantilever** by using the DISCRETE SHELL bending energy. Top: Comparing Bridson et al.'s bending formula (a) to Grinspun et al.'s one (b), in combination with LIBSHELL's membrane energy. Bottom: original ARCSIM results (c) and tentative fix (d) by using Grinspun et al.'s bending formulation (d). Increasing resolution (inset) unfortunately does not help improve results.

shift w.r.t the master curve. Nevertheless, results lie slightly below the curve and remain excessively scattered, as illustrated in Figure 7, bottom. Testing the code in its non-adaptive version and using an increased mesh resolution unfortunately did not help, as shown in the inset of Figure 7d. This result is all the more surprising as when using the LIBSHELL membrane part (instead of the ARCSIM membrane part), we easily check that replacing Bridson's factor with Grinspun's factor nicely removes the shift (Figure 7, top). Investigating further such discrepancies in ARCSIM was out of the scope of the present study, and we leave it for future work.

High Γ values. The region $\Gamma > 10^3$ is numerically challenging for the codes as the expected shapes are rods/ribbons heavily bent at the clamp and then almost vertical. Curvature-based models such as SUPER-HELIX or SUPER-RIBBON can easily capture this phenomenon with few elements (50 and 20 elements, respectively). However, nodal models like DISCRETE ELASTIC ROD, LIBSHELL or DISCRETE SHELL + ARCSIM require meshes with a fairly good resolution around the clamp (the latter required meshes twice as resolute as the former). For plates, without a sufficient resolution, the resulting mesh is unable to capture the smoothness and the results start to scatter (see supplementary document, for example Figure 27h).

A way to circumvent this issue would be to use adapted meshes, by subdividing more around the clamp. Additionally, for all models including rods, this would have the advantage to reduce the time spent in performing the test. However, for the sake of simplicity,

we decided in the present study to use uniform meshes for all the tested codes (except ARCSIM in its original version).

Successful $\text{\textcircled{A}}$ ABAQUS. The $\text{\textcircled{A}}$ ABAQUS software easily passes the test. Note that 200 2D Timoshenko beam elements with quadratic interpolation functions have been used. Based on a Newton solver with automatic stabilization, 200 static equilibria were found within approximately 1 hour (no warm-start, only changing the length of the rod every independent computation).

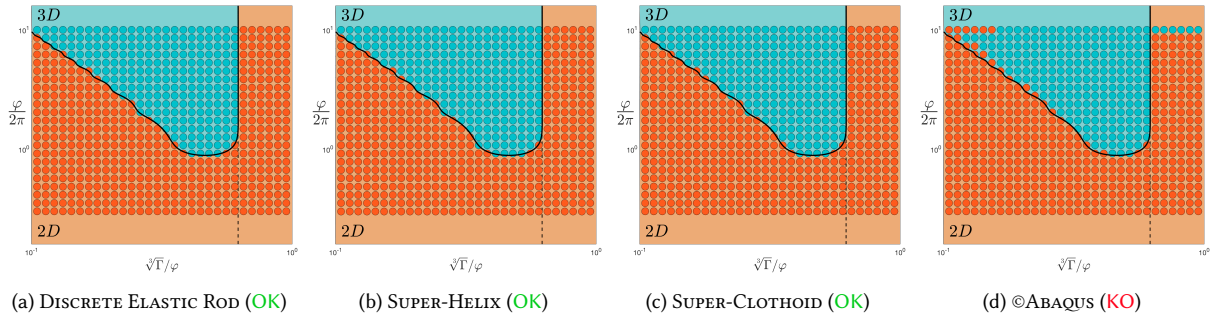
Tentative fitting of PROJECTIVE DYNAMICS. Finally, the PROJECTIVE DYNAMICS method, which relies on arbitrary (non-physical) weights in front of its three energy terms (gravity G , stretch S , bend B), provides a good example of calibration offered by our method. We can indeed fit at least its bending weight against the **Cantilever** master curve. We take $h \ll w$, for each vertex $G = 1/3A\phi hq$ where A is the area adjacent to the vertex, $S = Eh$, and $B = \alpha Eh^3/(12 - (1 - \nu^2))$. We found that for our implementation of PROJECTIVE DYNAMICS, $\alpha = 0.25$ gives the best agreement to the master curve. However, the output values yielded by the PROJECTIVE DYNAMICS code are scarce as the regime of this experiment is not favorable to the convergence of the method. In practice, we did not find any appropriate solver parameters allowing for a sampling as dense as for the other codes. Indeed, in this experiment the stretching coefficient S is much greater than the bending coefficient B , leading to a poor convergence of the local/global scheme, which is even worse at low Γ 's where G becomes much smaller than B . Nonetheless, the code produces fairly good results in the range $0.1 < \Gamma < 500$ beyond which it crosses the master curve and departs from it.

7.2 Results on the Bend-Twist test

The **Bend-Twist** test allows to challenge twist formulation of 3D rod models and displays a known numerical limit between 2D and 3D deformed configurations as explained in Section 4.3. Our three codes for rods, namely DISCRETE ELASTIC ROD, SUPER-HELIX, and SUPER-CLOTHOID, as well as $\text{\textcircled{A}}$ ABAQUS, were tested on this protocol.

Figure 8 shows that all the codes except $\text{\textcircled{A}}$ ABAQUS perform well. However, as the protocol is much more demanding than the **Cantilever** test, the simulation setup needs to be carefully designed.

Successful graphics rod models. The three rod models from Computer Graphics yield perfect results, provided enough elements are used for highly curved rods (200 elements for the nodal model DISCRETE ELASTIC ROD, 30 and 20 elements for the curvature-based models SUPER-HELIX and SUPER-CLOTHOID, respectively). The interested reader can see convergence plots in the supplementary material. Yet, to yield successful results, important details in the setup deserve to be emphasized. First, to generate a 3D configuration, it is generally needed to start from an initially 3D rod (otherwise there is a risk that the rod remains in an unstable 2D state). If one starts from an initially 2D curved rod, one idea is to apply a small transverse perturbation at the beginning of the simulation. Second, for this test it is crucial to make sure that dynamic models (DISCRETE ELASTIC ROD and SUPER-HELIX) have perfectly stabilized after each simulation. Otherwise there is a high risk to label wrongly a 2D configuration as a 3D one. For large ϕ and Γ (long and curly rods, forming a hook), such a stabilization can take a long time. We

Fig. 8. Results of the **Bend-Twist** test.

have spent many efforts attempting to reduce this time by damping the system, or using warm-starting (for a given point of the phase diagram, start from the final configuration of its left neighbor on the same line). Still, the time needed for DISCRETE ELASTIC ROD and SUPER-HELIX ranges from 6 hours to one day, while it was just about one hour for the static SUPER-CLOTHOID code.

Difficulties for ©ABAQUS. Similarly to SUPER-CLOTHOID, we have used ©ABAQUS with a static version relying on the Newton method. But as SUPER-CLOTHOID converges to the correct equilibria on the whole set of parameters with 20 elements, ©ABAQUS gives wrong results close to the 2D/3D frontier and at large natural curvature $\varphi/2\pi$. Increasing the number of elements (up to 785 for the largest $\varphi/2\pi$ value), adding some stabilization parameters to the Newton algorithm, and using previously computed deformed configurations as warm starts, all greatly improve a phase diagram that was initially much worse than the one presented in Figure 8d. But after days of trying, we have not been able to obtain entirely satisfactory results.

7.3 Results on the **Lateral Buckling** test

The **Lateral Buckling** test can be seen as an extension of the **Cantilever** test to elastic plates. Although both tests are meant to evaluate the bending formulation, the former is much more demanding than the latter. Indeed, **Cantilever** is a 2D test, whereas **Lateral Buckling** incorporates 3D effects through the width and is subject to a pitchfork instability (transverse buckling). A code not working perfectly well on **Cantilever** has little chance to pass the **Lateral Buckling** test. Hence, we consider as eligible for the **Lateral Buckling** test only the plate/shell codes that successfully passed the **Cantilever** test (OK), that is, SUPER-RIBBON, LIBSHELL, DISCRETE SHELL (+ LIBSHELL), FENICSSHELL, and ©ABAQUS. Each code took at most a few hours for performing the whole test.

Figure 9 summarizes our results. Only LIBSHELL and FENICSSHELL manage to pass this challenging test. As in the **Cantilever** test, we have used several kinds of mesh resolutions (yet different from those used for **Cantilever**): Res 0 (respectively Res +, Res ++) typically comprising respectively 50 (resp. 75, 100) points along the length and $50 \times w/L$ (resp. $75 \times w/L$, $100 \times w/L$) along the width.

Inaccuracy of DISCRETE SHELL. Investigating further the failure of DISCRETE SHELL, we found out that its bending formulation does not converge fast enough to yield a threshold in the range $0 < \Gamma^* < 40$ displayed in Figure 9. The threshold actually shows up at $\Gamma^* > 50$,

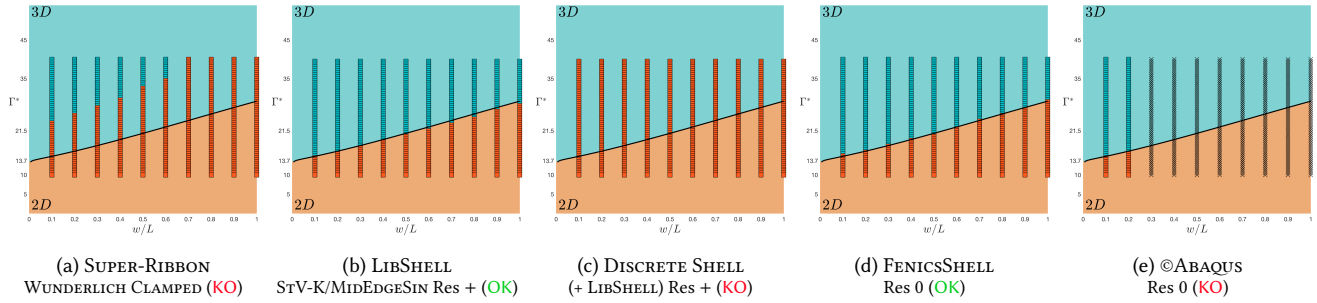

as shown in Figure 10a. The evolution of the threshold is linear, but with a steeper slope compared to the theory. We tried to improve results by increasing again the resolution of the mesh (Res +, 75 points along the length). However the threshold curve remained equally shifted from the theoretical master curve (Figure 10b), showing that the method would not converge to the right solution.

Failure of ©ABAQUS. Apart from two values of $w/L = 0.1$ and 0.2 for which the transition between 2D and 3D equilibrium shapes are correctly predicted by ©ABAQUS, the latter returns no result for $w/L > 0.2$. The reason of this failure lies in the incapacity of the Newton solver to converge to the starting point of the demanding **Lateral Buckling** protocol. Although an equilibrium is always found when the plate is slightly tilted at $\Gamma^* = 40$, when one strives to remove this rotation angle, the Newton solver keeps diverging whatever the stabilization options and the mesh resolution for $w/L > 0.2$, so that the **Lateral Buckling** experiment cannot even start.

Limitation of the Wunderlich ribbon model. Based on a one dimensional formulation, the SUPER-RIBBON code nevertheless exhibits a correct linear dependence of Γ^* with respect to w/L , although it displays a clear offset from the theoretical curve. We found out that this offset was caused by the developability constraint of the Wunderlich ribbon model, which turns out to be unrealistic in the **Lateral Buckling** test. Indeed, due to the orientation of the clamp, in-plane extension and shear are bound to happen. To overcome this limitation while still benefiting from a reduced formulation, it would be interesting to develop a stretchable ribbon model.

Discriminating LIBSHELL variants. As in the **Cantilever** test, we note that the elastic law has little impact here, hence we evaluate LIBSHELL using only the Saint-Venant Kirchhoff (StV-K) material. All the variants of LIBSHELL pass the test, however it appears that the discretization schemes for the discrete normals computation are discriminated by this test. Indeed, the two "second-order" models MIDGESIN and MIDGETAN pass at the resolution Res + while the "first-order" model MIDEDGEAVERAGE requires an even increased resolution Res ++ (see supplementary document, Figure 13).

Successful FENICSSHELL. It is noteworthy that FENICSSHELL successfully passes the test using default resolution meshes, in comparison with LIBSHELL which requires higher resolution meshes. This can probably be explained by the fact that FENICSSHELL relies on

Fig. 9. Results of the **Lateral Buckling** test.Fig. 10. Detailed results of the **Lateral Buckling** test for DISCRETE SHELL and our two different mesh resolutions up to $\Gamma^* = 80$.

high-order elements ($P_{2,3}$), with faster convergence compared to LIBSHELL.

Complementary rotation experiment. The interested reader can refer to our supplementary document (Section 4) for an additional test based on this experiment. In this complementary test, the clamped orientation is tilted by an angle α from the vertical direction, and we record the continuous value of the lateral displacement d_y as α is varied. This rotation test has the advantage of yielding a smooth signal in place of the present binary output of the **Lateral Buckling** test, thus being less demanding compared to our retained test. In particular, it can be leveraged to analyze failures on the **Lateral Buckling** test.

7.4 Results on the **Stick-Slip** test

Finally, the **Stick-Slip** test assesses the ability of a code to properly capture the sticking-slipping threshold of Coulomb's friction. Figure 11 gathers our results obtained on our six different codes for a planar rod (or a plate) subject to frictional contact.

Successful So-BOGUS. We first notice that the **Stick-Slip** test is particularly demanding, as among our tested Computer Graphics solvers, only the nonsmooth implicit So-BOGUS solver is able to pass it, whether being coupled to a planar rod model like SUPER-HELIX 2D or to a plate model like ARCSIM – incorporated within ARGUS.

Yet, it must be noted that for successful agreement with the master curve, one has to carefully set the parameters of the simulation, such as the solver tolerance and the timestep. We set the timestep to

0.5 ms and the solver tolerance to 10^{-13} N. Furthermore, it is necessary to ensure the quasi-static behavior of the dynamic simulators when replicating the experiment. A detailed study on convergence of So-BOGUS utilizing the SUPER-HELIX 2D model is provided in the supplementary material, Figure 16.

No adaptivity for ARGUS. While the bending formulation of SUPER-HELIX 2D was validated on **Cantilever**, it is not the case for ARCSIM: we have seen in Section 7.1 that ARCSIM's bending energy (relying on Bridson et al. [2003]'s model) had to be fixed by a constant factor. Moreover, despite this fix, the agreement to the master curve remained unsatisfactory. Interestingly, succeeding in the **Cantilever** test is not necessary for ARGUS to pass the **Stick-Slip** test, which demonstrates the independence of the latter with respect to the relative scaling of the bending and gravitational energies. This way we retrieve the results previously obtained by Rasheed et al. [2020]. However, although it generates fair results, the original adaptive version of ARGUS does not properly match the master curve, presumably due to small impacts during adaptive refinement: we had to disable adaptivity in the ARGUS code to pass the **Stick-Slip** test (see Figures 11c and 11d), as was also done by Rasheed et al. [2020].

Insufficient accuracy for explicit solvers. Explicit frictional contact solvers from Computer Graphics, such as BRIDSON-HARMON and PROJECTIVE FRICTION, suffer from a lack of accuracy and are not able to pass the **Stick-Slip** test, even when using a small timestep and/or a large number of iterations. For both solvers we test a timestep down to 0.1 ms, and observe no improvement over refinement. For PROJECTIVE FRICTION we test the number of local/global iterations up to 500. Note that unlike VISCOUS FRICTION, both methods do capture a stick-slip threshold, but they yield excessive sticking due to their projective nature. This is especially prominent in the BRIDSON-HARMON method, which systematically overestimates the slipping threshold. Our study thus confirms the observations previously made by Li et al. [2018] when comparing ARGUS to BRIDSON-HARMON. In contrast, PROJECTIVE FRICTION manages to generate a correct threshold value for moderate friction coefficients (up to $\mu = 0.25$). Yet, in accordance with Ly et al. [2020]'s observations, the model reaches a convergence plateau which binds it to a mere sticking behavior for higher μ values.

Successful results for ©ABAQUS. Finally, ©ABAQUS performs very well on the **Stick-Slip** test. We work in a 2D-space and use 40 Timoshenko isoparametric beam elements with quadratic interpolation (B22) to model Sano's rod. The rod motion is simulated thanks to

an explicit dynamic analysis with large Rayleigh internal damping ($\alpha = 10$) to damp the lower frequencies that could appear. Contacts between the bottom surface and the beam nodes are enforced with a kinematic contact algorithm (predictor-corrector algorithm that modifies accelerations to obtain a corrected configuration in which the contact constraints are enforced). The contact tangential behavior is a basic Coulomb friction model with static friction μ . Sticking between the rod and the bottom surface is ensured by a penalty method. The minimal increment being 9.223×10^{-6} s to ensure stability of the explicit scheme for a total time of 60 s for one **Stick-Slip** test, it takes over 6.5×10^6 increments for a CPU time of approximately 1 hour in total.




Fig. 11. Results of the **Stick-Slip** test.

8 LIMITATIONS AND FUTURE WORK

In our study we have provided a rich set of protocols to validate physically any numerical simulator for rods, plates, or frictional contact. Still, our framework does not cover all possible validation facets and could be further developed to handle more complex cases.

Validation versus verification. As recalled by Museth [2020], *validation* queries the ability of a simulator to solve the correct physics, whereas *verification* deals with the correct numerical solving of a given model [Pham 1999]. In that sense, our study focuses on the physical *validation* of numerical codes, not on their numerical verification. Indeed, numerical verification is only meaningful once one knows exactly the (continuous) physical model from which the code derives, as well as the physical validity of this model. While this is often the case in Mechanical Engineering, which clearly decouples the use of well-established physical models from their subsequent finite-element discretization, this is not necessarily the case in Computer Graphics, where simulators may be directly written as new numerical models integrating physics and numerics altogether. Our validation methodology then allows us to treat equally simulators deriving from a continuous physical model from that directly written on a discrete form, making it applicable to any numerical model.

However, we have seen that numerical verification, such as the proof of numerical convergence, often turns out to be a *prerequisite* to physical validation. For instance, a suitable choice of numerical parameters (time step, number of elements) can only be done if one has the guarantee that results will improve over refinement. Convergence studies are not always conducted in Computer Graphics papers, and we encourage the authors of new models to provide them systematically in the future.

While convergence is indirectly assessed by our protocols, other relevant numerical properties are however not evaluated. Typically, important features that we miss are the *objectivity* of the model – i.e. its invariance w.r.t. rigid motion [Crisfield and Jelenić 1998] – and *non-locking* – i.e. the ability of the model to continue converging properly upon some parameter refinement, typically when the thickness of the structure tends towards zero [Arnold and Brezzi 1997]. We plan to investigate these issues in the future.

Towards more complex physical protocols. Although richer than the one commonly used, our protocols remain too elementary to evaluate the full richness of Computer Graphics simulators. To augment their validation potential, one first step would be to consider *dynamical* scenarios, such as vibrating structures [Virgin 2007]. More complex contacting scenarios, such as 2D and 3D confinement of buckled slender structures [Roman and Pocheau 2002], could also be considered. An extension of our methodology to the validation of thin shells would also be valuable, but finding relevant master curves, which exploit characteristic instabilities depending on a limited set of dimensionless parameters, remains a challenging task.

Finally, to be able to deal with fully integrated phenomena that are common to Computer Graphics scenarios (large fiber assemblies, cloth with frictional contact, etc.), one should develop some innovative experimental setups. We believe such a large step forward will only be possible through a strong cooperation between experts

from the three communities Computer Graphics, Mechanical Engineering, and Soft Matter Physics. We hope our paper will foster more research in this fascinating pluridisciplinary area.

9 CONCLUSION

We have proposed a new framework to validate quantitatively the physical realism of numerical simulators for rods, plates and frictional contact. Our methodology relies on the introduction of four physical protocols inspired from experimental measurement setups of Soft Matter Physics. As these protocols exhibit macroscopic instabilities which depend on very few dimensionless physical parameters, they are characterized by a clear master curve against which any code can be robustly compared. This way, we can easily determine whether a code successfully passes each validation test, and thus evaluate its level of physical realism. We have tested a number of representative codes from Computer Graphics, as well as two reference codes from Mechanical Engineering. Our results demonstrate that even though some popular codes fail, some may be improved upon slight modification. Furthermore, remarkably a few Computer Graphics codes do succeed on all tests and even perform better than reference models of Mechanical Engineering. We hope that our study will help the Computer Graphics community better assess and diffuse their models in the future, and that the set of proposed protocols will be enriched by a joined work across the communities.

ACKNOWLEDGMENTS

We would like to thank Benoît Roman and José Bico for their help with the **Cantilever** experiment, Mélina Skoura and Thibaut Métivet for their careful proofreading of our paper submission, Francesco Dal Corso for providing us with useful bibliographic references, as well as the anonymous reviewers for their constructive comments. This work was supported in part by the ERC grant GEM (StG-2014-639139).

REFERENCES

- D. Arnold and F. Brezzi. 1997. Locking free finite element for shells. *Math. Comput.* 66 (Jan. 1997), 1–14. <https://doi.org/10.1090/S0025-5718-97-00785-0>
- U. Ascher, R. Mattheij, and R. Russell. 1995. *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*. SIAM.
- B. Audoly and Y. Pomeau. 2010. *Elasticity and Geometry: From Hair Curls to the Non-linear Response of Shells*. Oxford University Press.
- D. Baraff and A. Witkin. 1998. Large Steps in Cloth Simulation. In *Computer Graphics Proceedings*. 43–54. <http://www.cs.cmu.edu/~baraff/papers/sig98.pdf>
- D. Baraff, A. Witkin, and M. Kass. 2003. Untangling Cloth. *ACM Transactions on Graphics* 22, 3 (2003), 862–870.
- A. Bartle, A. Sheffer, V. Kim, D. Kaufman, N. Vining, and F. Berthouzoz. 2016. Physics-driven Pattern Adjustment for Direct 3D Garment Editing. *ACM Trans. Graph.* 35, 4, Article 50 (July 2016), 11 pages.
- M. Bergou, B. Audoly, E. Vouga, M. Wardetzky, and E. Grinspun. 2010. Discrete Viscous Threads. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'10)* 29, 4 (2010). <http://www.cs.columbia.edu/cg/threads>
- M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun. 2008. Discrete elastic rods. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'08)* 27, 3 (2008), 1–12. <https://doi.org/10.1145/1360612.1360662>
- F. Bertails, B. Audoly, M.-P. Cani, B. Querleux, F. Leroy, and J.-L. Lévêque. 2006. Super-Helices for Predicting the Dynamics of Natural Hair. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'06)* 25 (2006), 1180–1187. Issue 3. <https://doi.org/10.1145/1141911.1142012>
- B. Bickel, M. Bäcker, M. Otaduy, H. Richard Lee, H. Pfister, M. Gross, and W. Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph.* 29, 4, Article 63 (July 2010), 10 pages. <https://doi.org/10.1145/1778765.1778800>
- B. Bickel, M. Bäcker, M. Otaduy, W. Matusik, H. Pfister, and M. Gross. 2009. Capture and Modeling of Non-Linear Heterogeneous Soft Tissue. *ACM Trans. Graph.* 28, 3, Article 89 (July 2009), 9 pages. <https://doi.org/10.1145/1531326.1531395>
- W.G. Bickley. 1934. The heavy elastica. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 17, 113 (1934), 603–622. <https://doi.org/10.1080/14786443409462419>
- S. Bouaziz, S. Martin, T. Liu, L. Kavan, and M. Pauly. 2014. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Trans. Graph.* 33, 4, Article 154 (July 2014), 11 pages.
- F. Boyer, V. Lebastard, F. Candelier, and F. Renda. 2020. Dynamics of Continuum and Soft Robots: A Strain Parameterization Based Approach. *IEEE Transactions on Robotics* (2020), 1–17. <https://doi.org/10.1109/TRO.2020.3036618>
- R. Bridson, R. Fedkiw, and R. Anderson. 2002. Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph.* 21, 3 (2002), 594–603. <http://www.cs.ubc.ca/~rbridson/docs/cloth2002.pdf>
- R. Bridson, S. Marino, and R. Fedkiw. 2003. Simulation of Clothing with Folds and Wrinkles. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '03)*. 28–36. <http://dl.acm.org/citation.cfm?id=846276.846281>
- M. Brunetti, J. S. Hale, and C. Maurini. 2018. Fenics-Shells demos. https://fenics-shells.readthedocs.io/en/latest/demo/nonlinear-naghd-cylindrical/demo_nonlinear_naghd-cylindrical.py.html
- E. Buckingham. 1914. On physically similar systems; illustrations of the use of dimensional equations. *Physical Review* 4, 4 (1914), 345–376.
- R. Casati and F. Bertails-Descoubes. 2013. Super Space Clothoids. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'13)* 32, 4, Article 48 (July 2013), 12 pages. <https://doi.org/10.1145/2461912.2461962> <http://www.inria.fr/bipol/people/casati/research.html#ssc>
- R. Charrondière, F. Bertails-Descoubes, S. Neukirch, and V. Romero. 2020. Numerical modeling of inextensible elastic ribbons with curvature-based elements. *Computer Methods in Applied Mechanics and Engineering* 364 (2020), 112922. <https://doi.org/10.1016/j.cma.2020.112922>
- H.-Y. Chen, A. Sastry, W. van Rees, and E. Vouga. 2018. Physical Simulation of Environmentally Induced Thin Shell Deformation. *ACM Trans. Graph.* 37, 4, Article 146 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201395>
- D. Clyde, J. Teran, and R. Tamstorf. 2017. Modeling and Data-Driven Parameter Estimation for Woven Fabrics. In *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA '17)*. Association for Computing Machinery, New York, NY, USA, Article 17, 11 pages. <https://doi.org/10.1145/3099564.3099577>
- E. Coevoet, A. Escande, and C. Duriez. 2019. Soft robots locomotion and manipulation control using FEM simulation and quadratic programming. In *RoboSoft 2019 - IEEE International Conference on Soft Robotics*. Seoul, South Korea. <https://hal.inria.fr/hal-02079151>
- S. Coros, S. Martin, B. Thomaszewski, C. Schumacher, R. Sumner, and M. Gross. 2012. Deformable Objects Alive! *ACM Trans. Graph.* 31, 4, Article 69 (July 2012), 9 pages. <https://doi.org/10.1145/2185520.2185565>
- M. A. Crisfield and G. Jelenić. 1998. Objectivity of strain measures in the geometrically exact three-dimensional beam theory and its finite-element implementation. *Proc. Royal Society of London, Series A* 455, 1983 (1998), 1125–1147.
- Dassault-Systems. 2005. Simulia Abaqus. <https://www.3ds.com/products-services/simulia/products/abaqus>
- G. Daviet. 2020. Simple and Scalable Frictional Contacts for Thin Nodal Objects. *ACM Trans. Graph.* 39, 4, Article 61 (July 2020), 16 pages. <https://doi.org/10.1145/3386569.3392439>
- G. Daviet, F. Bertails-Descoubes, and L. Boissieux. 2011. A Hybrid Iterative Solver for Robustly Capturing Coulomb Friction in Hair Dynamics. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 1–12.
- A. Derouet-Jourdan, F. Bertails-Descoubes, G. Daviet, and J. Thollot. 2013. Inverse Dynamic Hair Modeling with Frictional Contact. *ACM Trans. Graph.* 32, 6, Article 159 (Nov. 2013), 10 pages. <https://doi.org/10.1145/2508363.2508398>
- M. Dias and B. Audoly. 2015. “Wunderlich, Meet Kirchhoff”: A General and Unified Description of Elastic Ribbons and Thin Rods. *Journal of Elasticity* 119, 1 (Apr 2015), 49–66. <https://doi.org/10.1007/s10659-014-9487-0>
- E. Doedel, H. Keller, and J.-P. Kernevez. 1991. Numerical Analysis and Control of Bifurcation Problems (I) Bifurcation in Finite Dimensions. *International Journal of Bifurcation and Chaos* 1, 3 (1991), 493–520.
- V. Duclaux. 2006. *Pulmonary occlusions, eyelid entropion and aneurysm : a physical insight in physiology*. Ph.D. Dissertation. Université de Provence - Aix-Marseille I. <https://tel.archives-ouvertes.fr/tel-00130610>
- A. Fargette. 2017. *Soft Interfaces: from elastocapillary snap-through to droplet dynamics on elastomers*. Ph.D. Dissertation. Université Pierre et Marie Curie - Paris VI.
- R. Fosdick and E. Fried (Eds.). 2015. *The Mechanics of Ribbons and Moebius Bands*. Springer.
- K. Gavril, R. Guseinov, J. Pérez, D. Pellis, P. Henderson, F. Rist, H. Pottmann, and B. Bickel. 2020. Computational Design of Cold Bent Glass Façades. *ACM Transactions on Graphics (SIGGRAPH Asia 2020)* 39, 6, Article 208 (Dec 2020), 16 pages.

- J. M. Gere. 2004. *Mechanics of Materials* (6th ed.). Thomson-Brooks/Cole.
- Alain Goriely. 2006. Twisted Elastic Rings and the Rediscoveries of Michell's Instability. *Journal of Elasticity* 84 (September 2006), 281–299. <https://doi.org/10.1007/s10659-006-9055-3>
- E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin. 2006. Computing discrete shape operators on general meshes. *Computer Graphics Forum* 25, 3 (2006), 547–556. <https://doi.org/10.1111/j.1467-8659.2006.00974.x>
- E. Grinspun, A. Hirani, M. Desbrun, and P. Schröder. 2003. Discrete Shells. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'03)*. ACM-EG SCA, 62–67. <http://www.multires.caltech.edu/pubs/ds.pdf>
- R. Guseinov, C. McMahon, J. Pérez, C. Daraio, and B. Bickel. 2020. Programming temporal morphing of self-actuated shells. *Nature Communications* 11 (Jan. 2020). <https://doi.org/10.1038/s41467-019-14015-2>
- M. Habera, J. S. Hale, A. Logg, C. Richardson, J. Ring, M. E. Rognes, N. Sime, and G. N. Wells. 2018. The Fenics Project. <https://fenicsproject.org>
- S. Hadap. 2006. Oriented Strands - Dynamics of Stiff Multi-Body System. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation*. <https://doi.org/10.2312/SCA/SCA06/091-100>
- J. S. Hale, M. Brunetti, S. Bordas, and C. Maurini. 2018. Simple and extensible plate and shell finite element models through automatic code generation tools. *Computers & Structures* 209 (2018), 163–181.
- D. Harmon, E. Vouga, B. Smith, R. Tamstorf, and E. Grinspun. 2009. Asynchronous Contact Mechanics. *ACM Trans. Graph.* 28, 3, Article 87 (July 2009), 12 pages. <https://doi.org/10.1145/1531326.1531393>
- D. Harmon, E. Vouga, R. Tamstorf, and E. Grinspun. 2008. Robust Treatment of Simultaneous Collisions. *ACM Trans. Graph.* 27, 3, Article 23 (Aug. 2008), 4 pages. <https://doi.org/10.1145/1360612.1360622>
- D. Hinz and E. Fried. 2015. Translation of Michael Sadowsky's Paper 'An Elementary Proof for the Existence of a Developable Möbius Band and the Attribution of the Geometric Problem to a Variational Problem'. *Journal of Elasticity* 119, 1 (2015), 3–6.
- L. Hu, D. Bradley, H. Li, and T. Beeler. 2017. Simulation-Ready Hair Capture. *Computer Graphics Forum* (2017). <https://doi.org/10.1111/cgf.13126>
- M. Jawed, F. Da, J. Joo, E. Grinspun, and P. Reis. 2014. Coiling of elastic rods on rigid substrates. *Proceedings of the National Academy of Sciences* 111, 41 (2014), 14663–14668. <https://doi.org/10.1073/pnas.1409118111> http://www.cs.columbia.edu/cg/elastic_coiling
- D. Kaufman, R. Tamstorf, B. Smith, J.-M. Aubry, and E. Grinspun. 2014. Adaptive Nonlinearity for Collisions in Complex Rod Assemblies. *ACM Trans. Graph.* 33, 4, Article 123 (July 2014), 12 pages.
- S. Kawabata and Masako Niwa. 1989. Fabric Performance in Clothing and Clothing Manufacture. *The Journal of The Textile Institute* 80, 1 (1989), 19–50. <https://doi.org/10.1080/00405008908659184>
- K. Krieger. 2012. Extreme mechanics: buckling down. *Nature* 488, 7410 (Aug. 2012), 146–147.
- F. Laccone, L. Malomo, J. Pérez, N. Pietroni, F. Ponchio, B. Bickel, and P. Cignoni. 2019. FlexMaps Pavilion: a twisted arc made of mesostructured flat flexible panels. In *FORM and FORCE, IASS Symposium 2019, Structural Membranes 2019*. CIMNE, 498–504. <http://vcg.isti.cnr.it/Publications/2019/LMPPBC19>
- L. D. Landau and E. M. Lifshitz. 1959. *Theory of elasticity*. Pergamon London.
- J. Li, G. Daviet, R. Narain, F. Bertails-Descoubes, M. Overby, G. Brown, and L. Boissieux. 2018. An Implicit Frictional Contact Solver for Adaptive Cloth Simulation. *ACM Trans. Graph.* 37, 4, Article 52 (Aug. 2018), 15 pages.
- M. Li, D. Kaufman, and C. Jiang. 2020. Codimensional Incremental Potential Contact. <https://arxiv.org/pdf/2012.04457.pdf>
- J. Liang, M.C. Lin, and V. Koltun. 2019. Differentiable Cloth Simulation for Inverse Problems. In *Conference on Neural Information Processing Systems*.
- M. Ly, J. Jouve, L. Boissieux, and F. Bertails-Descoubes. 2020. Projective Dynamics with Dry Frictional Contact. *ACM Transactions on Graphics* 39, 4 (2020), 1–8. <https://doi.org/10.1145/3386569.3392396>
- S. Martin, P. Kaufmann, Mario Botsch, E. Grinspun, and M. Gross. 2010. Unified Simulation of Elastic Rods, Shells, and Solids. *ACM Trans. Graph.* 29, 4, Article 39 (July 2010), 10 pages. <https://doi.org/10.1145/1778765.1778776>
- J. Martínez, M. Skouras, C. Schumacher, S. Hornus, S. Lefebvre, and B. Thomaszewski. 2019. Star-Shaped Metrics for Mechanical Metamaterial Design. *ACM Transactions on Graphics* 38, 4 (July 2019), Article No. 82:1–13. <https://doi.org/10.1145/3306346.3322989>
- A. McAdams, A. Selle, K. Ward, E. Sifakis, and J. Teran. 2009. Detail preserving continuum simulation of straight hair. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'09)* 28, 3 (2009), 1–6. <https://doi.org/10.1145/1531326.1531368>
- A. G. M. Michell. 1899. Elastic stability of long beams under transverse forces. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 48, 292 (1899), 298–309.
- E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. Otaduy, and S. Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Computer Graphics Forum* 31, 2 (may 2012). <https://doi.org/10.1111/j.1467-8659.2012.03031.x>
- Y. Mikata. 2006. Complete solution of elastica for a clamped-hinged beam, and its applications to a carbon nanotube. *Acta mechanica* 190, 1–4 (Oct. 2006), 133–150.
- J. Miller, A. Lazarus, B. Audoly, and P. Reis. 2014. Shapes of a Suspended Curly Hair. *Physical Review Letters* 112, 6 (2014).
- K. Museth. 2020. Physics simulations: Is it Hollywood magic or rocket science. <http://computeranimation.org/2020/program.html#keynote1>
- R. Narain, T. Pfaff, and J. O'Brien. 2013. Folding and Crumpling Adaptive Sheets. *ACM Trans. Graph.* 32, 4 (2013), 51.
- R. Narain, A. Samii, and J. O'Brien. 2012. Adaptive Anisotropic Remeshing for Cloth Simulation. *ACM Trans. Graph.* 31, 6, Article 152 (Nov. 2012), 10 pages.
- D. K. Pai. 2002. STRANDS: Interactive Simulation of Thin Solids using Cosserat Models. *Computer Graphics Forum* 21, 3 (2002), 347–352. <https://doi.org/10.1111/1467-8659.00594>
- J. Panetta, M. Konaković-Luković, F. Isvoranu, E. Bouleau, and M. Pauly. 2019. X-Shells: A New Class of Deployable Beam Structures. *ACM Trans. Graph.* 38, 4, Article 83 (July 2019), 15 pages. <https://doi.org/10.1145/3306346.3323040>
- H. Pham. 1999. *Software reliability*. Springer.
- A.-H. Rasheed, V. Romero, F. Bertails-Descoubes, S. Wuhler, J.-S. Franco, and A. Lazarus. 2020. Learning to Measure the Static Friction Coefficient in Cloth Contact. In *CVPR 2020 - IEEE Conference on Computer Vision and Pattern Recognition*. Seattle, United States, 1–10. <https://hal.inria.fr/hal-02511646>
- P. M. Reis. 2015. A Perspective on the Revival of Structural (In)Stability With Novel Opportunities for Function: From Bucklephobia to Buckliphilia. *Journal of Applied Mechanics* 82, 11 (Sept. 2015). <https://doi.org/10.1115/1.4031456> 111001.
- E. Reissner. 1995. The problem of lateral buckling of cantilever plates. *ZAMM - Journal of Applied Mathematics and Mechanics* 8 (1995), 615–621.
- B. Roman and A. Pocheau. 2002. Postbuckling of bilaterally constrained rectangular thin plates. *Journal of the Mechanics and Physics of Solids* 50, 1 (2002), 2379–2401.
- M. Sadowsky. 1929. Die Differentialgleichungen des MÖBIUSSchen Bandes. *Jahresbericht der Deutschen Mathematiker-Vereinigung* (1929), 49–51. , translated in [Hinz and Fried 2015].
- T. Sano, T. Yamaguchi, and H. Wada. 2017. Slip Morphology of Elastic Strips on Frictional Rigid Substrates. *Physical Review Letters* 118, 17 (2017), 178001–5.
- C. Schumacher, S. Marschner, M. Gross, and B. Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (July 2018), 15 pages. <https://doi.org/10.1145/3197517.3201278>
- R. T. Shield. 1992. Bending of a beam or wide strip. *Quarterly Journal of Mechanics and Applied Mathematics* 45, 4 (1992), 567–573.
- M. Skouras, B. Thomaszewski, B. Bickel, and M. Gross. 2012. Computational Design of Rubber Balloons. *Comput. Graphics Forum (Proc. Eurographics)* (2012).
- M. Skouras, B. Thomaszewski, P. Kaufmann, A. Garg, B. Bickel, E. Grinspun, and M. Gross. 2014. Designing Inflatable Structures. *ACM Trans. Graph.* 33, 4, Article 63 (July 2014), 10 pages. <https://doi.org/10.1145/2601097.2601166>
- B. Smith, D. Kaufman, E. Vouga, R. Tamstorf, and E. Grinspun. 2012. Reflections on Simultaneous Impact. *ACM Trans. Graph.* 31, 4, Article 106 (July 2012), 12 pages. <https://doi.org/10.1145/2185520.2185602>
- J. Spillmann and M. Teschner. 2007. CoRDe: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'07)*. ACM-EG SCA, 63–72.
- K. Sze, X. Liu, and S. Lo. 2004. Popular benchmark problems for geometric nonlinear analysis of shells. *Finite Elements in Analysis and Design* 40, 11 (2004), 1551 – 1569. <https://doi.org/10.1016/j.finel.2003.11.001>
- S. Timoshenko. 1953. *History of strength of materials*. The McGraw-Hill Book Company, Inc.
- F. Vanneste, O. Gouy, J. Martinez, S. Lefebvre, H. Delingette, and C. Duriez. 2020. Anisotropic soft robots based on 3D printed meso-structured materials: design, modeling by homogenization and simulation. *IEEE Robotics and Automation Letters* 5, 2 (Jan. 2020), 2380–2386. <https://doi.org/10.1109/LRA.2020.2969926>
- L. Virgin. 2007. *Vibration of axially loaded structures*. Cambridge University Press.
- H. Wang. 2018. Rule-Free Sewing Pattern Adjustment with Precision and Efficiency. *ACM Trans. Graph.* 37, 4, Article 53 (July 2018), 13 pages.
- H. Wang, R. Ramamoorthi, and J. O'Brien. 2011. Data-driven elastic models for cloth: modeling and measurement. *ACM Trans. Graph.* 30, 4, Article 71 (Aug. 2011), 12 pages.
- C. Weischedel. 2012. A discrete geometric view on shear-deformable shell models.
- W. Wunderlich. 1962. Über ein abwickelbares Möbiusband. *Monatshefte für Mathematik* 66, 3 (June 1962), 276–289. <https://doi.org/10.1007/BF01299052>
- S. Yang, J. Liang, and M.C. Lin. 2017. Learning-based cloth material recovery from video. In *Proceedings of the IEEE International Conference on Computer Vision*. 4383–4393.
- S. Yang and M.C. Lin. 2016. MaterialCloning: Acquiring elasticity parameters from images for medical applications. *IEEE Transactions on Visualization and Computer Graphics* 22, 9 (2016), 2122–2135.