



HAL
open science

Freely orientable microstructures for designing deformable 3D prints

Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez,
Pierre-Alexandre Hugron, Fabrice Neyret, Sylvain Lefebvre

► **To cite this version:**

Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, et al.. Freely orientable microstructures for designing deformable 3D prints. *ACM Transactions on Graphics*, 2020, 39 (6), pp.1-16. 10.1145/3414685.3417790 . hal-02524371v3

HAL Id: hal-02524371

<https://inria.hal.science/hal-02524371v3>

Submitted on 8 Sep 2020

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

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

Freely orientable microstructures for designing deformable 3D prints

THIBAUT TRICARD, Université de Lorraine, CNRS, Inria, LORIA

VINCENT TAVERNIER, Université Grenoble Alpes, CNRS, Laboratoire Jean Kuntzmann, Inria

CÉDRIC ZANNI, Université de Lorraine, CNRS, Inria, LORIA

JONÀS MARTÍNEZ, Université de Lorraine, CNRS, Inria, LORIA

PIERRE-ALEXANDRE HUGRON, Université de Lorraine, CNRS, Inria, LORIA

FABRICE NEYRET, Université Grenoble Alpes, CNRS, Laboratoire Jean Kuntzmann, Inria

SYLVAIN LEFEBVRE, Université de Lorraine, CNRS, Inria, LORIA




Fig. 1. Our technique synthesizes fiber-like microstructures that can be additively manufactured. This induces an elastic response that we call *rigid-transverse*: the fibrous microstructures are very rigid along their elongated axis while being comparatively very flexible in the (local) orthogonal plane. **Left:** Without any spatial gradation, our microstructures are formed by square diamond profiles defined in the v, w plane and extruded along the axis u . **Right:** 3D printed object embedding our synthesized microstructures. The designer directly manipulates the fiber orientations, controlling how the cube volume reshapes under large deformations. Here, the top face collapses while gripping the bar, as the fibers prevent stretch along their main direction (see orange dashed line). Note the more disorganized structures at the bottom, which cancels the directional effect in this area.

Nature offers a marvel of astonishing and rich deformation behaviors. Yet, most of the objects we fabricate are comparatively rather inexpressive, either rigid or exhibiting simple homogeneous deformations when interacted with.

We explore the synthesis and fabrication of novel microstructures that mimic the effects of having oriented rigid fibers in an otherwise flexible material: the result is extremely rigid along a transverse direction while being comparatively very flexible in the locally orthogonal plane. By allowing free gradation of the rigidity direction orientation within the object, the microstructures can be designed such that, under deformation, distances along fibers in the volume are preserved while others freely change. Through a simple painting tool, this allows a designer to influence the way the volume reshapes when deformed, and results in a wide range of novel possibilities. Many gradations are possible: local free orientation of the fibers; local control

of the overall material rigidity (structure density); local canceling of the effect of the fibers, obtaining a more isotropic material.

Our algorithm to synthesize the structures builds upon procedural texturing. It produces a cellular geometry that can be fabricated reliably despite 3D printing walls at a minimal thickness, allowing prints to be very flexible. The synthesis algorithm is efficient and scales to large volumes.

CCS Concepts: • **Computing methodologies;**

Additional Key Words and Phrases: procedural, textures, noise, phasor, Fabrication

ACM Reference Format:

Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, Fabrice Neyret, and Sylvain Lefebvre. 2020. Freely orientable microstructures for designing deformable 3D prints. *ACM Trans. Graph.* 39, 6, Article 211 (December 2020), 16 pages. <https://doi.org/10.1145/3414685.3417790>

1 INTRODUCTION

With the advent of additive manufacturing (AM), we can envision the design and fabrication of objects exhibiting a wide palette of deformation behaviors. The elastic response of a part can be significantly modified by carefully crafting a small scale architecture within its volume, a *metamaterial*. Yet, despite major advances, we are still relatively limited when it comes to obtaining large controlled deformations within an object, across different directions.

A promising approach is to exploit a *contrast* in the directional elastic response of a shape: being near non-stretchable along some

Authors' addresses: Thibault Tricard, Université de Lorraine, CNRS, Inria, LORIA; Vincent Tavernier, Université Grenoble Alpes, CNRS, Laboratoire Jean Kuntzmann, Inria; Cédric Zanni, Université de Lorraine, CNRS, Inria, LORIA; Jonàs Martínez, Université de Lorraine, CNRS, Inria, LORIA; Pierre-Alexandre Hugron, Université de Lorraine, CNRS, Inria, LORIA; Fabrice Neyret, Université Grenoble Alpes, CNRS, Laboratoire Jean Kuntzmann, Inria; Sylvain Lefebvre, Université de Lorraine, CNRS, Inria, LORIA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2020/12-ART211 \$15.00

<https://doi.org/10.1145/3414685.3417790>

directions while being comparatively very flexible along others. Under large deformations, the contrast induces internal curvature, resulting in highly non-homogeneous deformations that would be challenging to obtain otherwise. Recent examples include pneumatic plates [Siéfert et al. 2019] and the design of flextures [Shaw et al. 2019]. However, the studied structures are tailored to a specific case and offer limited design freedom. In this work we propose a freely orientable metamaterial with a high directional rigidity contrast. It affords designers unprecedented freedom, allowing to mimic the behavior of flexible fibrous materials (iniquitous in animal/plant tissue), and to trigger highly anisotropic, large deformations within a volume.

Numerically simulating under such large deformations is extremely challenging. Notably due to thin walls, complex stochastic geometry, and the difficulty to model the effect of layered fabrication. For large deformations, we thus resort to physical measurements on actual fabricated samples. This has the advantage to take into account the actual final product as opposed to an idealization of it. This is especially important as we print walls at the smallest possible thickness. The downside is that we cannot explore the material space or undertake inverse problems in a numerical fashion. However, our structures provide a fiber analogy allowing an intuitive approach for modeling.

The designer directly controls the local fiber orientations and structure sizing through a painting interface and can fabricate designs at low cost, on widely available filament 3D printers.

This idea is illustrated in Figure 1 where the fiber-like microstructures — produced by our method — makes the top of the deformed cube wrap around the bar being pressed on it. The yellow dashed line in Figure 1 (right) outlines a curve along which our structure resists stretch and preserves the curvilinear distance, as a rigid fiber would. We call the resulting material *transverse-rigid* as its maximal rigidity is along the selected transverse (fiber) direction and is much stiffer than its average response in the orthogonal plane.

Our approach allows to freely orient and grade such an elastic directional response throughout the volume of an object. The structures afford for contrast gradation between the most rigid and most flexible directions (see Figure 1, bottom part of the cube). This allows to locally dampen or cancel the directional contrast, progressively transitioning towards a structure that has the same elastic response in all directions.

Defining periodic structures having a transverse-rigid behavior in a fixed reference frame is relatively straightforward. But without the ability to freely orient the transverse direction and grade the rigidity contrast, the design possibilities are severely restricted. Making such controls possible, however, requires synthesizing a geometry with specific spatially varying characteristics (Section 3 and Figures 2 and 3). This leads to the conflicting requirements of having changes in local orientation and scale while preserving an overall smooth, locally regular structure.

Besides inducing the target elastic response, a key challenge is to ensure that the synthesized geometry can be fabricated with AM. For the large deformations we target, internal voids and thin walls are paramount, as the combination of both will produce the least dense and most flexible objects while being less prone to fracture than open-cell (truss) structures. On most technologies, this leads

to the major difficulty of avoiding any sort of support structure within the volume. 3D printing thin walls — at the minimal possible thickness — also implies enforcing a maximum overhang constraint: no wall should have an angle that exceeds a maximum with respect to the vertical direction of fabrication. Otherwise, the solidified curves forming the walls in each slice will not vertically bond to one another [Livesu et al. 2017].

The microstructure we synthesize is carefully crafted to achieve all the required properties. It prints reliably on even the most affordable AM process (filament printers), behaves like a transverse-rigid material, it can be freely spatially graded in terms of orientation and rigidity contrast, and it is efficient to synthesize, making it scalable (embarrassingly parallel algorithms with no global memory storage of the result).

To achieve this, we define the microstructures from a stochastic process introduced in the field of procedural texturing [Lague et al. 2010a]. The stochastic component provides direct control over the average distribution of geometric features, which translates into control over the average mechanical response. However, this comes with a specific challenge: such processes generate structures that can have a large number of localized discontinuities that negatively impact the structural properties. We propose a specialized filtering method that removes most of these issues from the final results.

We describe the geometry synthesis process in Section 3 and Section 4. To verify that the geometry we synthesize indeed produces a transverse-rigid oriented response, we fabricate a variety of samples covering the parameter space and submit each to mechanical testing. We report results in Section 5. We also fabricate a number of designs with interesting deformations, shown in Section 6 as well as in the accompanying video.

Contributions. In summary, our contributions are:

- (1) The definition, from stochastic processes, of a microstructure with high transverse-rigidity, which orientation, scale, and rigidity contrast can be freely graded.
- (2) A filtering approach to strongly reduce the occurrence of defects in the synthesized structures.
- (3) An evaluation of the obtained results in terms of manufacturability and mechanical properties.




Fig. 2. **Left:** Microstructure without gradations. The extruded diamond shape has a high rigidity in the direction of extrusion u and is comparatively very flexible in the two orthogonal directions v and w . **Right:** Illustration of orientation gradation along an arbitrary direction field. Conforming to changes in direction while maintaining a near-constant structure scale inevitably yields scattered singularities (red circles). Several types of singularities exist, the ones shown here are not problematic.

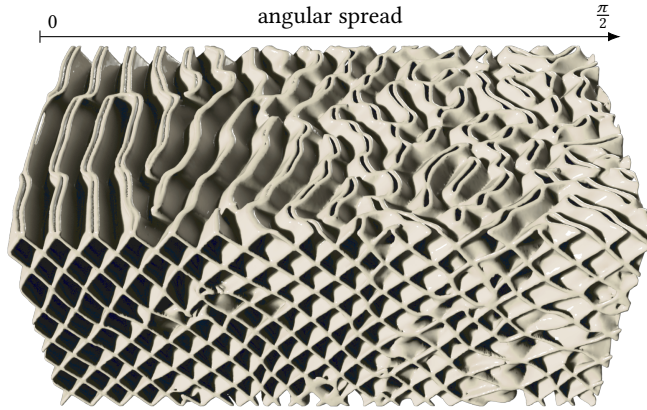


Fig. 3. *Rigidity contrast gradation*: the structure transitions from a high directional rigidity contrast (left) to a material with a similar response in all directions (right).

2 MICROSTRUCTURES FOR ADDITIVE MANUFACTURING

There is a large body of work considering how microstructures, or small-scale architectures, modify the macro-scale mechanical behavior of an object. We describe here those most related to our approach, and in particular those studied in the context of AM.

Throughout this discussion, and in the remainder of the paper, it is important to recall that the mechanical behavior of microstructures is to be understood as the *average behavior* of a sufficiently large volume filled with the microstructure [Gibson and Ashby 1999].

Microstructures are synthesized inside an object following a variety of approaches. A traditional way is to repeat carefully optimized tiles within a periodic grid [Panetta et al. 2017, 2015; Schumacher et al. 2015; Zhu et al. 2017]. The periodicity affords for efficient simulation of material properties [Michel et al. 1999] and slicing [Brennan-Craddock 2011]. Gradation is achieved by choosing different tiles in different locations, ensuring their boundaries match, either by construction [Panetta et al. 2015], optimization [Schumacher et al. 2015] or parametric interpolation [Fryazinov et al. 2013; Li et al. 2015]. Microstructures can also be defined through subdividing geometries [Kuipers et al. 2019; Medeiros e Sá et al. 2015; Wu et al. 2016], in which case density is directly controlled by the subdivision level.

While tile-based and subdivision approaches may allow for large deformations (locally placing tiles with anisotropic response), the underlying grid makes it very difficult to follow arbitrary orientations and arbitrary control fields. Indeed, the pre-computed tile set can only practically include a small set of orientations, given that all possible transition tiles would also have to be included.

Another approach is to generate microstructures from stochastic processes [Martínez et al. 2016, 2018, 2017]. The study of parameterized stochastic microstructures [Pham et al. 2019; Tarantino et al. 2019] has been gaining research focus in the last few years due to certain advantages it offers compared to periodic microstructures. First, stochastic microstructures allow for very progressive gradation and increased design freedom, as no grid or subdivision rule is constraining the structures [Siddhant et al. 2020]. Namely, the spatial gradation of a stochastic process is intrinsically encoded by its parameters. This is particularly relevant for anisotropic 3D

microstructures that are graded following an arbitrary 3D vector field. Recent works dealing with the 3D periodic case [Geoffroy-Donders et al. 2020] require additional processing to grade and conform the periodic geometry to a surface. By contrast, stochastic microstructures do not require any additional processing. Second, it has been reported that stochastic microstructures are more resilient to fabrication imperfections arising from symmetry-breaking issues, compared to periodic microstructures [Portela et al. 2020]. Finally, procedural stochastic microstructures afford for scalable and efficient slicing [Pasko et al. 2011; Vidimče et al. 2013].

When using microstructures, spatial gradations may be directly designed [Ion et al. 2016; nTopology 2019], or may be optimized for various objectives such as rigidity [Li et al. 2015; Zhu et al. 2017], target deformations [Panetta et al. 2015], or balance [Wu et al. 2016].

Microstructures dedicated to additive manufacturing typically take fabrication constraints into account. These include minimal geometry size, maximal overhang [Martínez et al. 2018; Wu et al. 2016], the existence of unsupported islands in slices [Panetta et al. 2015] as well as overlaps in deposition paths [Kuipers et al. 2019].

Freely orientable microstructures. Several works orient structures in a volume or along a surface, for instance, conformal truss structures [Wang et al. 2005], periodic structures following a texture mapping analogy [Chen 2007], or procedural synthesis with free orientation [Martínez et al. 2017].

It is also possible to directly optimize a microstructure within a part, performing high-resolution topology optimization [Aage et al. 2017; Wu et al. 2018] that allows complex oriented structures to emerge. A recent development consists in synthesizing microstructures following the output of a topology optimization process [Geoffroy-Donders et al. 2020; Groen et al. 2020; Pantz and Trabelsi 2008; Wu et al. 2019]. These techniques produce cellular (hexahedral) and lattice (truss) structures following an optimized direction field. This has direct links to the field of hex-meshing [Gao et al. 2017; Jakob et al. 2015; Ray et al. 2006].

To the best of our knowledge, none of the approaches we mentioned above allow for directional grading of a microstructure inducing large, inhomogeneous deformations within the object. Martínez et al. [2017] orient orthotropic microstructures within a volume. However, the open-cell beam structures can only be fabricated on SLA/SLS: the beams take arbitrary angles well beyond overhang constraints and beams print poorly on FDM machines due to the many small disconnected regions appearing in every slice. The structures also tend to be fragile under large deformations, as stress accumulates on joints: for the same density, closed-cell structures have in general higher fracture toughness than open-cell structures [Maiti et al. 1984]. The closed-cell structure introduced by Martínez et al. [2018] can be graded but is orientable only in the plane aligned with the printer bed, and is significantly less collapsible than our structures (see Section 5). The alternative closed-cell structure proposed by Kuipers et al. [2019] is very flexible and collapsible along the vertical axis; however, no gradations other than density are possible.




Fig. 4. Phasor noise with a sawtooth wave profile (inset). The main orientation of the kernels and the amount of angular spread is controlled by user-defined fields. The conflicting requirements of following orientation and preserving a constant wavelength yield two singularity types: curves — breaking the pattern (in red) — and points — creating branching waves (in green). The profile is distorted in those areas, introducing unwanted high-frequency content and discontinuities.

3 SYNTHESIZING TRANSVERSE-RIGID MICROSTRUCTURES

We now describe how we define and synthesize our microstructures. Note that despite guidelines (e.g. [Ashby 2006]) and in-depth studies of specific cases (e.g. [François et al. 2017; Schumacher et al. 2018]), it remains very challenging in general to precisely predict the link between a small-scale geometry and its macro-scale mechanical behavior. This is made even more difficult if taking into account large deformations and fabrication with layered AM processes.

In this section we define a geometry that we expect to result in the target mechanical behavior. We back up our prediction with detailed measurements in Section 5.3 and supplemental material, Section 6, as well as demonstration samples in Section 6.

We synthesize a cellular pattern, whose walls are meant to be printed using the minimal thickness of the target AM process. Using thin walls allows producing the highest number of cells per unit of length, which helps to lower the variance of the mechanical behavior in small volumes.

We do not directly model the walls of the structure, but instead, design a function that synthesizes a black and white cellular pattern whose boundaries will represent the walls (see Figure inset). So our procedure returns the color of the cell enclosing an evaluation point at a coordinate x . We later extract the wall geometry in each AM fabrication layer (Section 3.3).

Our structure is controlled through two user-defined fields F and Γ . F defines a frame field in space with three orthogonal vectors $u(x)$, $v(x)$, and $w(x)$ at each point x inside the object. u has unit scale while both v , w have a free, but equal scale. This scale controls the local cross-sectional size of the structure, which forms elongated cells aligned with u everywhere. Γ defines an angular spread of directions in the local u, v plane, with random rotations around w .




Fig. 6. Black and white cellular pattern.

This is an angle in radians, from 0 to 2π , which controls how much the cells orientations locally deviate from u .

The fields are created by the designer using a basic painting tool that we discuss in Section 6. In practice, the user only controls u and the scale of v and w : their directions are obtained as $v = u \times z$ and then $w = v \times u$.

In terms of physical properties, we seek to design a structure such that the transverse rigidity is everywhere aligned with the u axis defined by F . The local density is inversely proportional to the scale of v, w , and derived from the wall thickness. Note that ideally the solidified thickness in a 3D printed slice should adjust for the local slope of the wall, from a ratio of 1 (vertical) to $\sqrt{2}$ (max overhang of 45 degrees). Γ changes the rigidity contrast from the maximum at the given density towards a more isotropic result. Γ has no significant impact on density. Please refer to Section 5.3 for measurements validating these relationships on actual samples.

3.1 Design challenges

Consider a very flexible periodic 2D pattern in the xy -plane, typically a bending dominated structure [Ashby 2006]. By extruding it vertically along the z -axis, we produce a 3D pattern that remains flexible in the xy -plane but is very rigid along the z -axis (as vertical walls strongly resist extension and compression).

Thus, we can expect that by sweeping such a 2D pattern along a direction u in space, the alignment of the extruded walls with u would locally trigger a similarly aligned local transversal rigidity. To control the rigidity contrast, we have to modulate the microstructure geometry from a pattern that has mainly walls along u , to a pattern that has a more uniform distribution of wall orientations.

We are now facing an especially difficult task, as we seek to have strongly oriented regular patterns in some locations (locally an extruded 2D pattern), and progressively grade towards evenly distributed spatial structures. Meanwhile, the walls have to obey everywhere the overhang constraints of AM processes, as they will be printed at minimal thickness. We address these challenges next.

3.2 Stochastic structures

Stochastic procedural texturing methods offer rich controls over the patterns that can be synthesized. These are usually focused on producing seemingly random continuous patterns that obey a specific spectral content in terms of frequency, direction and bandwidth [Lagae et al. 2010b]. We note that a perfectly regular stripe pattern can be described in terms of spectral content. For instance, a perfect sine wave is a Dirac in the spectral domain. As procedural methods with increasingly precise spectral content are developed [Goldberg et al. 2008; Lagae et al. 2009], these offer an interesting alternative to parameterization methods for generating stripe patterns (e.g. [Jakob et al. 2015; Knöppel et al. 2015]).

We define our microstructures from a phasor noise [Tricard et al. 2019]; a stochastic pattern generator that can produce sharp, contrasted oscillating patterns (see Figure 4). The oscillation's frequency and orientation can be spatially graded at will, and can spatially evolve from a locally regular periodic oscillation towards a stochastic pattern having no clear preferred orientation.




Fig. 5. Overview of the microstructure geometry synthesis, without any spatial gradation for clarity. We define two orthogonal solid phasor noises p_v and p_w with a sawtooth profile. Each produces a cyclic linear ramp in $[0, 1]$. Combining both defines local u, v parametrizations in $[0, 1]^2$. These local parametrizations are then used to define diamond cells, from which the structure walls are extracted. AM toolpaths are directly generated from this structure definition.

The produced oscillations are less regular than those produced by parametrization methods. However, this trade-off is compensated by the ability to spatially grade the patterns, as well as the computational efficiency of procedural methods: the approach scales trivially, is supported by embarrassingly parallel algorithms, and does not require meshing the domain. We therefore use this methodology to define our 3D microstructure geometry and later focus on improving the quality of the synthesized oscillations (Section 4).

3.2.1 Phasor noise. A phasor noise is an instantaneous phase field ϕ which is defined from a complex version \mathcal{G} of Gabor noise:

$$\phi(\mathbf{x}) = \text{Arg}(\mathcal{G}(\mathbf{x})) \quad (1)$$

with:

$$\mathcal{G}(\mathbf{x}) = \sum_j e^{-b\|\mathbf{x}-\mathbf{x}_j\|^2} e^{2i\pi f_j \mathbf{d}_j \cdot (\mathbf{x}-\mathbf{x}_j) + i\phi_j} \quad (2)$$

This is a sparse convolution process [Lewis 1989]. A finite set of kernels located at random locations \mathbf{x}_j in a domain combine into the final signal. Each kernel is a product of two complex functions. The first is a clamped (99% of stddev) Gaussian of bandwidth b centered on the kernel, the second is an oscillator of frequency f_j , phase ϕ_j and direction \mathbf{d}_j . In our case, \mathbf{x} , \mathbf{x}_j and \mathbf{d}_j are 3D vectors, thus \mathcal{G} and ϕ define solid noises (see Figure 5, left).

Feeding the phase $\text{Arg}(\mathcal{G})$ into a periodic function produces oscillations having the profile of the function. The produced patterns can be freely graded in orientation and frequency (see Figure 4).

3.2.2 Local parametrizations. To generate a cellular pattern such as the one shown in Figure 6, we need to determine whether a point belongs to a black or a white cell. A convenient approach is to define a 2D parametrization in space (e.g. a set of texture coordinates) from which we can associate every 3D point with a 2D point in a base profile. The profile defines the cross-sectional shape of the cells, as illustrated in Figure 5. The choice of profile will impact the final mechanical response as well as determine whether the result is fabricable under overhang constraints. We discuss it in Section 3.2.3.




Fig. 7. The axis of a diamond cell can be oriented within the full unit-sphere, while being under a 45 degrees overhang constraint. **Left:** Square diamond shape extruded along its normal, Y . All walls have angles of exactly 45 degrees with respect to the horizontal plane. **Right:** Rotation of the diamond cell around its horizontal diagonal (X). This further increases the vertical cross-section wall angles (red lines), which remain below the constraint.

Producing such a parameterization in our context is challenging. It is akin to computing a globally periodic parameterization (e.g. [Ray et al. 2006]) – with the additional challenge of providing the ability to grade the structure’s regularity based on the field Γ . The method should also easily scale to large volumes and not require meshing the domain.

We propose to exploit the oscillation profile control of phasor noise. In particular, an oscillation with a *sawtooth* profile can be seen as a periodic local 1D parameterization which varies linearly from 0 to 1 in each period. By computing two orthogonal oscillating waves each with a sawtooth profile, we obtain local 2D parameterization patches (each wave provides one coordinate).

We orient two 3D phasor waves such that their directions are locally orthogonal to \mathbf{u} and respectively oscillate along \mathbf{v} and \mathbf{w} . We denote the waves respectively p_v and p_w . At a given point \mathbf{x} , the phasor sawtooth evaluates as:

$$p_v(\mathbf{x}) = \frac{\text{Arg}(\mathcal{G}_v(\mathbf{x}))}{2\pi} \quad (3)$$

where \mathcal{G}_v has the form given in Equation 2, with the kernel directions \mathbf{d}_j obtained as $\mathbf{v}(\mathbf{x}_j)$ (and similarly for p_w , using $\mathbf{w}(\mathbf{x}_j)$), see Figure 5 (left). The phasor sawtooth waves depicted in the remainder of the paper are planar slices of the 3D noises.

To generate cells with an increasing angular spread of wall directions we use Γ to perturb the kernel orientations in p_v . We add a rotation around w to the base orientation of the kernels, randomly sampling an angle in $[-\Gamma(\mathbf{x}), \Gamma(\mathbf{x})]$. When $\Gamma(\mathbf{x}_i) = \pi/2$ this results in oscillations in all directions (Figure 4 right), that still properly define parametric patches. Roughly speaking, in 3D the structures locally take the geometry of an hourglass around the w axis.

For the sake of conciseness, we move implementation details of the Phasor noise – which follow the state of the art – to the supplemental material (Section 4 therein).

While very effective, defining the local parameterizations with phasor noise comes with its own challenges: the phasor fields have local defects that impact the structure and its mechanical properties. We discuss these in Section 3.2.4.

3.2.3 Diamonds and overhangs. To allow for free orientation, the transverse axis of the generated microstructures should be able to cover any direction in the full unit-sphere S^2 . However, the structure also has to always verify the overhang constraints for correct fabrication.

Under this objective, a square diamond extrusion profile has desirable properties. Let us consider a 2D diamond shape in the x, z plane having one of its diagonal aligned with the vertical printing direction z . The extrusion of the diamond shape along y defines a cell with a diamond cross-section (see Figure 7, left). The wall angles with the horizontal plane are exactly 45 degrees, which typically enforces the maximum overhang angle constraint [Livesu et al. 2017]. Any such diamond cell remains printable for arbitrary rotations around the x then z axis (see Figure 7, right), covering the full sphere of directions.

We however have to cancel any *roll* around u from v and w in F , ensuring that v remains orthogonal to z . While the user has complete freedom over u , we also have to constrain Γ to avoid overhangs. Indeed, Γ adds random rotations to the kernels around w , which has a similar effect as adding roll around the diamond axes. The range of values $\Gamma(\mathbf{x})$ can take is thus restricted based on $u(\mathbf{x})$ (see supplemental for details). The recommendation is thus to keep u in the x, y plane where a large Γ is desired. In practice this is a minor inconvenience: a large Γ is used to cancel the directional effects, and therefore the choice of u is less important in such regions.

We obtain the black diamond cells (Figure 6) by thresholding the local parametrizations, using the following formula:

$$\left| \frac{1}{2} - p_v(\mathbf{x}) \right| + \left| \frac{1}{2} - p_w(\mathbf{x}) \right| < \frac{1}{2}$$

White cells are obtained by reversing the inequality.

3.2.4 Singularities. Using phasor noises as a basis to generate our microstructures comes with a specific challenge: the waves exhibit *singularities* where they diverge from an ideal regular oscillation. These defects occur where the complex Gabor field vanishes ($\mathcal{G}(\mathbf{x}) = 0$), but also where the phase field of the phasor noise varies abruptly, distorting the oscillations. While Neyret and Heitz [2016] proposed an FFT-based iterative filter to reduce singularities in the context of Gabor noise, it is not applicable in the presence of spatial gradations.

Note that in a non trivial orientation field a number of singularities have to exist [Bommes et al. 2013]. However, these should





Fig. 8. Phasor noise sawtooth oscillations, with a Gabor kernel (red) and its overlapping neighbors (green) shown. Constant kernel orientation (left) and smooth changes of orientation (right). **Top:** Raw result. **Bottom:** Our algorithm aligns kernel phases towards a more regular result.

ideally be points (in a slice) or curves (in 3D) where oscillations meet or split (see Figure 2, right). For the sake of clarity, since we show in figures 2D slices of 3D noises we will refer to the ideal case as *point singularities*.

Point singularities at the apex of branching oscillations (forks) do not endanger the overall properties of the structure: around the singular point, the produced shape still exhibits a diamond extrusion profile and prints reliably, with only a small hole at their base.

Unfortunately, we also observe *curve singularities* in the phasor noise (*surface singularities* in the 3D case); across such curves the local parameterization values abruptly change and oppose on both sides (see red circles in Figure 4). These are more problematic than point singularities as they induce a change from a black to a white diamond (and vice-versa) resulting in walls orthogonal to the local diamond extrusion direction – see red walls in the inset. Such walls violate the overhang constraint if the orientation is vertical, and make the structure stiffer overall.



As an illustration of the scope of the problem, a structure of fixed orientation is about four times stiffer in v, w compared to u when these singularities are not dealt with (see details in supplemental material, Figure 9). This results in a significantly lower rigidity contrast, which is detrimental to our use cases.

We focus in Section 4 on the problem of filtering *curve singularities* from phasor noises.

3.3 Toolpath extraction

For the sake of completeness, before we proceed to filtering let us briefly describe how we extract the final structure for 3D printing. We seek to produce a single, as thin as possible wall between diamonds. Given a slice, we first draw an image of the sliced structure by evaluating whether each pixel belongs to a black or white diamond cell, as defined in Figure 5. This results in a binary image, from which we extract a solidification trajectory along each boundary between black and white pixels.

This is a robust process, which only drawback is the image discretization error. However, we can use a higher image resolution – smaller than half of the minimal thickness – since memory consumption is low: we work in a single slice and evaluate the pixel ‘color’ using a GPU pixel shader. The shader evaluates the phasor


noises with a sawtooth profile directly from the kernels. The kernels themselves use little memory, and the memory allocated for a slice is freed immediately after trajectory extraction.

4 FILTERING SINGULARITIES

The main drawback of relying on phasor noises is that the oscillations exhibit many singularities even in simple cases (F smooth overall and $\Gamma = 0$). One would instead expect a very regular result with smooth and progressive changes, and only a limited number of singularities where waves fork (see Figure 8).

We now describe how to strongly reduce the occurrence of singularities, while preserving the variability associated with changes of orientation, frequency, and angular spread.

Recall that our focus is on *curve* singularities, which have a detrimental effect on the synthesized geometry (see Section 3.2.4 and Figure 4). Such singularities most notably arise when two neighboring kernels with similar direction \mathbf{d}_j are out of phase — see Figure inset.



Our approach works at two fundamentally different levels. The stochastic process of the phasor noise cannot possibly produce phase coherency between kernels across long distances. Thus, we apply an iterated regularization process that operates on the phasor noise kernels and seeks to align the phases of neighboring kernels — removing inconsistencies — in order to produce longer coherent wavefronts (see Figure 8 left). It is described in Section 4.1.

A perfect phase alignment is generally impossible in the presence of variations in the frame field F . Therefore, even after phase alignment, some local singularities remain where neighboring kernels of different orientation or frequency combine. To further improve the result, we define a specialized filter that locally removes deviations from an idealized oscillation estimated in every sampling point. We describe the filter in Section 4.2.

4.1 Iterated phase alignment

To regularize the oscillations of the phasor noise, we first adjust the phase shifts of neighboring kernels, as illustrated in Figure 8. Interestingly, this problem appears in a related form in recent interactive quad meshing techniques [Evgeny and Harders 2019; Lichtenberg et al. 2018]. These operate on triangles and have one oriented oscillator (akin to a phasor kernel) per vertex. They attempt to globally align the phases of the per-vertex oscillators, so as to define continuous patterns across triangles by applying an iterative diffusion process. To achieve high performance, the phase at every vertex is iteratively updated, in parallel for all vertices, from the average of phases suggested by the neighbors. Note that the orientation and frequency of the oscillators remains fixed. Our iterated phase alignment for phasor noise bears some similarities but operates without a mesh on the initial kernels of the underlying (complex) Gabor noise.

At each iteration, the algorithm adjusts the phase of each kernel to be the phase of the average complex oscillator obtained from its neighbors. Aligning phases between kernels suppose they share a similar orientation. Thus, the average is weighted by the dot product of the orientation vectors of the kernel and its neighbor. The dot

Algorithm 1 Improving the phase of kernel k_j

```

a ← 0
for all  $k_i \in \mathcal{N}(k_j)$  do
    a ← a + max(0,  $\mathbf{d}_j \cdot \mathbf{d}_i$ ) ·  $e^{2i\pi f_i \mathbf{d}_i \cdot (\mathbf{x}_j - \mathbf{x}_i) + i\varphi_i}$ 
end for
 $\varphi_j \leftarrow \text{Arg}(a)$ 
    
```




Fig. 9. Some singularities in a phasor sawtooth example. Point singularities (a) represent forks and merges in the pattern. They correspond to rotations around a singular point in the phasor field. Curve singularities (b) introduce major discontinuities in the result as well as profile distortions in various directions. They are characterized by long curves with a high phase gradient. Phasor field images show $\text{Arg}(\varphi(\mathbf{x}))$ encoded as the hue.

product result is clamped to zero. The algorithm for one iteration is given in Algorithm 1, where $\mathcal{N}(k_j)$ is the set of kernels neighboring k_j (they belong to the same or neighboring cells in the Gabor noise evaluation grid [Lagae et al. 2009]).

We repeat Algorithm 1 for a fixed number of iterations (20 in our implementation), which we experimentally adjusted (please refer to supplemental material for details, Section 5). The algorithm trivially maps to an efficient GPU implementation, with one thread per kernel. The average time for one iteration on an NVIDIA GTX 1080 is 7ms for 551K kernels for a 100 mm × 100 mm × 100 mm cube.

4.2 Evaluation Filter

While phase alignment removes overall inconsistencies, the complex local interactions between disagreeing kernels still create local singularities. We propose a filter that regularizes the signal, filtering away discontinuities which locally distort the noise profile.

We start by examining the singularities and their properties in order to identify and classify them, and quantify their presence in the generated noise.

4.2.1 Identifying singularities. Let us first consider a simple case where the frequency and the orientations are constant in space, noted respectively f and \mathbf{d} . In this case, we can extract the *phasor field* of the phasor noise. It is obtained at each point \mathbf{x} by dividing the complex Gabor noise by a global oscillator ω_o :

$$\omega_o(\mathbf{x}) = e^{2i\pi f \mathbf{d} \cdot \mathbf{x}} \quad (4)$$

The phasor field is thus obtained as:

$$\varphi(\mathbf{x}) = \frac{\mathcal{G}(\mathbf{x})}{\omega_o(\mathbf{x})} \quad (5)$$




Fig. 10. Phasor noise filtering with increasing filter kernel sizes. **Left:** phasor field argument encoded as hue. **Right:** phasor sawtooth. **Inset:** filter. **Top to bottom:** increasing kernel size. Note how the phasor field becomes smoother and remaining singularities tend to become point singularities.




Fig. 11. Filtered phasor sawtooth (right) vs. phasor sawtooth with an equivalent noise bandwidth b (left). Our filter formulation results in a pattern that matches the input fields more closely and is less prone to leading to undesired angular spread.

It represents the phase shift between a specific phasor noise pattern and a hypothetical perfectly regular global oscillator (i.e. a simple sine wave). Examples of phasor fields are shown in Figure 9.




Fig. 12. Phasor field (right) obtained by division of the input noise (left) by the corresponding global oscillator. Note how there are strong oscillations along a preferred direction in the field, which is unexpected from a phasor field such as in Figure 10.




Fig. 13. Non-attenuated filter (middle) vs. attenuated filter (bottom). Without attenuation, the filter privileges the direction given by the input orientation field and discards other orientations. This erases the angular spread (red), which is a superposition of kernels of spread-out directions. The attenuated filter does not exhibit this behavior (green), and thus better preserves the properties of the original noise (top).

REMARK 1. *The smoother the phasor field is, the less the pattern deviates from an ideal oscillator. Thus the phasor field variations characterize deviations from the ideal signal (the singularities).*

Note that there is a notion of scale in this remark: deviations become problematic when they happen over a short distance compared to the oscillations wavelength. This wavelength is determined by the local frequency:

REMARK 2. *The frequency of the phasor noise defines the relative variation scale at which deviations can be considered as singularities.*

Indeed, variations that are smooth with respect to the local frequency are not problematic, as the phasor noise smoothly and progressively shifts around the ideal oscillator. However, where the phasor field varies too abruptly, the noise profile is locally compressed, which results in discontinuities. This is coherent with the visual examination of the fields in Figure 9.

Remarks 1 and 2 open the possibility to filter singularities by directly smoothing the phasor field (Section 4.2.2). However, this only holds for a constant orientation and frequency over the whole surface. A primary objective of our work is to allow arbitrary spatial gradations, therefore a more general point of view is needed.

We note that under the assumption that the control fields are smooth, they can be *locally* considered constant. We start by rewriting the global oscillator for an arbitrary origin point \mathbf{c} :

$$\omega_{\mathbf{c}}(\mathbf{x}) = e^{2i\pi f \mathbf{d} \cdot (\mathbf{x} - \mathbf{c})} \quad (6)$$




Fig. 14. Comparison of the contribution of the different parts of the method to the reduction of singularities. **Left:** from top to bottom: input orientation and scale, original phasor sawtooth wave, phasor sawtooth wave after iterated kernel phase alignment, phasor sawtooth wave after filtering, phasor sawtooth wave after iterated phase alignment and filtering. From left to right: various input fields and random seeds. **Right:** close-up of the effect of the iterated phase alignment, filtering, and both methods on the second example. Both approaches contribute to the overall improvement (see circled singularity).

The field $\frac{\mathcal{G}(\mathbf{x})}{\omega_{\mathbf{c}}(\mathbf{x})}$ is the same as $\varphi(\mathbf{x})$ with an additional constant phase shift $e^{2i\pi f \mathbf{d} \cdot (-\mathbf{c})}$. Thus, it has the same gradient, and gives the same result under Remark 1.

Then, substituting f and \mathbf{d} for $F(\mathbf{c})$ and $D(\mathbf{c})$ — the approximated parameters around \mathbf{c} — in $\omega_{\mathbf{c}}$, we obtain a local definition of the phasor field:

$$\varphi_{\mathbf{c}}(\boldsymbol{\delta}) = \frac{\mathcal{G}(\mathbf{c} + \boldsymbol{\delta})}{\omega_{\mathbf{c}}(\mathbf{c} + \boldsymbol{\delta})} \quad (7)$$

$\varphi_{\mathbf{c}}(\boldsymbol{\delta})$ provides a good approximation of the phase shift relative to an ideal oscillator at any offset $\boldsymbol{\delta}$ around \mathbf{c} , as long as $\boldsymbol{\delta}$ remains small compared to the spatial variations of the control fields.

This allows us to evaluate functions on the locally-defined phasor field in a neighborhood centered around \mathbf{c} . We use this property to define a measure for singularities and then express the evaluation filter.

4.2.2 Filter. From Equation 5 we see that the initial complex Gabor noise can be reconstructed from the phasor field, multiplying it back by the oscillator. This opens the possibility to directly manipulate the phasor field, and in particular, apply a smoothing filter that will remove abrupt variations. Following Remark 1, doing so will reduce the extent of singularities in the phasor noise.

Given g a Gaussian kernel of bandwidth a :

$$g(\mathbf{x}) = e^{-a\|\mathbf{x}\|^2} \quad (8)$$

The filtered complex Gabor noise is obtained as:

$$\mathcal{F}(\mathbf{x}) = (g \circledast \varphi_{\mathbf{x}})(\mathbf{x}) \cdot \omega_{\mathbf{x}}(\mathbf{x}) \quad (9)$$

Note that the convolution (\circledast) is applied to the variable δ , around point \mathbf{x} , hence filtering around the local phasor field.

Since $\omega_{\mathbf{x}}(\mathbf{x}) = 1$ (see Equation (6)), the filter simply becomes:

$$\mathcal{F}(\mathbf{x}) = (g \circledast \varphi_{\mathbf{x}})(\mathbf{x}) \quad (10)$$

An example of filtering is given in Figure 10.

Making the filter practical. At this stage we have defined a filter that improves the result. Unfortunately it is also computationally expensive due to the convolution in Equation 10. Used as is, obtaining the convolution result would require at each evaluation point \mathbf{x} a dense local sampling of both $\varphi_{\mathbf{x}}$ and the Gaussian filter. As we operate in 3D, this would add a very significant computational and/or memory overhead.

Fortunately, the convolution-based noise formulation has desirable properties with respect to Gaussian products that we can use to our advantage. As was done for texture filtering in the original Gabor noise [Lagae et al. 2009] we can analytically integrate the filtering operation directly into the kernels.

At a given evaluation point, applying the filter to the kernels we obtain new filtered kernels \mathcal{K}_j :

$$\mathcal{K}_j(\mathbf{x}) = e^{-\frac{ab}{a+b}\|\mathbf{x}-\mathbf{x}_j\|^2 - \frac{\pi^2\|\Delta_{j\mathbf{x}}\|^2}{a+b}} e^{2i(\pi f_j \mathbf{d}_j \cdot (\mathbf{x}-\mathbf{x}_j) + \frac{a}{a+b} \Delta_{j\mathbf{x}} \cdot (\mathbf{x}-\mathbf{x}_j)) + i\varphi_j} \quad (11)$$

where

$$\Delta_{j\mathbf{x}} = f_j \mathbf{d}_j - f_{\mathbf{x}} \mathbf{d}_{\mathbf{x}} \quad (12)$$

$f_{\mathbf{x}}$ and $\mathbf{d}_{\mathbf{x}}$ are the frequency and orientation sampled at the evaluation point \mathbf{x} . Recall that a and b are respectively the bandwidth of

the filter and the phasor noise. The full derivation of this equation is given in the supplemental material, Section 4.

This leads to a formulation with only a small overhead. Indeed, using the kernels defined in Equation (11) in place of the original ones *directly gives the filtered result*. Thus, evaluating the filtered pattern remains a simple sum of kernels, as with the original phasor noise. The kernel expression is only slightly more complex, while we entirely eliminate the sampling that was required to implement the convolution in Equation 10. The only remaining overhead comes from the larger footprints of the filtered kernels. This requires fetching a larger neighborhood compared to the non-filtered version. However, this only adds a limited, constant cost to the evaluation.

The $\Delta_{j\mathbf{x}}$ term can be interpreted as a weight that weakens the contribution from kernels whose frequency and orientation are too different from the ones given by the input field. Interestingly, when this term goes towards zero the filter becomes equivalent to decreasing the bandwidth of the noise (larger kernels). The importance of this term is revealed in Figure 11, where our filter is compared to simply reducing the noise bandwidth. In areas of quick orientation changes, decreasing the bandwidth leads to mixing orientations (the pattern locally results from summing kernels having different orientations, see bottom right part of the first image in Figure 11). The filter, on the contrary, preserves the input orientation everywhere.

4.2.3 Filtering and angular spread. The angular spread Γ locally perturbs the kernel orientations, resulting in a pattern locally mixing different orientations. It is important to not filter out these variations in the final result, as they are user-specified gradations (control of the rigidity contrast).

By construction, the filter we propose assumes only a single direction is present at any given point in the noise. This is not true of a phasor noise with angular spread, as illustrated in Figure 12. Thus if we would just extract the phasor field from such a noise and convolve it with a Gaussian filter, the extra oscillations introduced by the angular spread would be removed, see Figure 13, top.

To avoid this, we introduce an attenuation factor $A(\mathbf{x}) \in [0, 1]$ that permits a smooth transition between the original filter definition and a Gaussian filter with no preferred orientation. This is done by substituting $\Delta_{j\mathbf{x}}$ for $\Delta'_{j\mathbf{x}}$ in Equation (11):

$$\Delta'_{j\mathbf{x}} = A(\mathbf{x})\Delta_{j\mathbf{x}} \quad (13)$$

The factor A is expected to vary smoothly so as to not introduce discontinuities in the result. We derive it from the local angular spread of the phasor noise, please refer to the supplemental material, Section 4.3.

4.3 Results of singularity reduction

Figure 14 shows a variety of qualitative examples using the sawtooth profile, and for each, the effect of both approaches separately, and then combined. In all cases, we observe a significant reduction in the number of singularities by the regularization, as well as a much more regular oscillation pattern after the evaluation filter is applied.

Applying only the filter results in a high number of curve discontinuities which cannot be solved without aligning phases, but applying only the iterated phase alignment results in remaining local distortions. Using both approaches gives the best results.




Fig. 15. Two slabs of toolpaths visualized in our printer emulation tool. A red color indicates an overhang violation, a blue color an overlap. Sliced with 0.2 mm layers for a 0.4 mm nozzle. Structure scale was chosen to obtain a cube filled at 25% volume density. **Left:** $\Gamma = 0$. **Right:** $\Gamma = \frac{\pi}{2}$.

We provide in the supplemental material a quantitative analysis of both iterated alignment and filter by deriving a *singularity energy* from the local ideal oscillator definition, see Section 3.

5 ANALYSIS

We perform our analysis for the case of fused-filament 3D printing, the most widespread and inexpensive additive manufacturing technique. For the sake of conciseness, we report here our main conclusions and include additional data in the supplemental material, Section 6.

All results are 3D printed using flexible *NinjaFlex 85A* filament on a *Creativity CR10* modified to replace the initial Bowden extruder by a direct extrusion mechanism (The E3D Titan Aero); this increases reliability by ensuring the gear pushing the (flexible) filament is as close as possible to the nozzle.

5.1 Fabricability

While our structures nearly enforce geometrical overhang and overlap constraints everywhere, we cannot provide *strict* guarantees. We thus perform additional analysis of the generated deposition trajectories (Gcode).

We measure overhang violation and overlaps by rasterizing the paths in high resolution images (0.04 mm per pixel). Figure 16 top, reports statistics for a cube before singularity filtering. After filtering (Figure 16 bottom) the number of deviations is strongly reduced. Figure 15 visualizes two slabs of toolpaths with color-coded deviations, revealing that they are spatially scattered. Thus, their effect is negligible during fabrication, which is confirmed experimentally as the structures fabricate reliably on all our 3D printers.

5.2 Fidelity to control fields

We now consider how well the synthesized geometries match the user-provided target fields. First, we measure how well the scale of \mathbf{v}, \mathbf{w} in F correlates with the final structure density. Second, we verify that the structure orientation matches the user specifications (\mathbf{u} in F).

For analysis, we convert the produced solidification trajectories (G-code) into a high resolution voxel model. We then analyze the data to extract local density and local orientation measures.




Fig. 16. 35 mm cube sliced at 0.2 mm thickness for 0.4 mm nozzle, filled at 25% density, with $\Gamma = 0$. **Left:** histogram of the length of trajectory segments exceeding the overhang constraint (less than 40% of the extrusion disc supported from below). **Right:** Same for overlap statistics. A segment is counted when the overlap along it exceeds 10% of the nozzle diameter. Graphs are clamped after the first 95% of segment counts for readability. Filtering strongly reduces the occurrence of deviations, which are small and scattered, having no impact during fabrication.

The density ρ_d is estimated at each point by filtering the voxel model using a 3D Gaussian filter. The filter standard deviation is chosen to cover the width of two diamond cells and measure the local average behavior. We then measure its correlation with respect to the scale of \mathbf{v}, \mathbf{w} in F .

The orientation \mathbf{o}_d is extracted by casting rays from each voxel: the longest rays follow the direction of the generated diamond cells, and should match the input direction \mathbf{i}_d . The distribution of $\arccos(\mathbf{i}_d \cdot \mathbf{o}_d)$ represents how much spread around the input direction there is in the resulting model.

Figure 17 reveals the close correlation between the user specified fields and the final, 3D printed geometry. The analysis is performed on our most complex example, described in Section 6 and shown 3D printed in Figure 29.

5.3 Mechanical behavior

Simulating the structures under large deformations would be extremely challenging and prone to large approximations due to thin walls, complex stochastic geometry, as well as buckling and contacts. We instead measure the true experimental mechanical behavior of our structures by performing physical measurements on fabricated samples, using a uniaxial compression testing machine *Instron 3345*.

Figure 18 shows measurements of the elastic response along the main axes for varying values of Γ . For $\Gamma = 0$, we clearly see the desired rigidity contrast, with a \mathbf{u} axis significantly stiffer than \mathbf{v}, \mathbf{w} across the full range of deformation – including beyond buckling, which is an extreme scenario along the \mathbf{u} axis for our use cases. As Γ increases we see how the contrast diminishes until we obtain

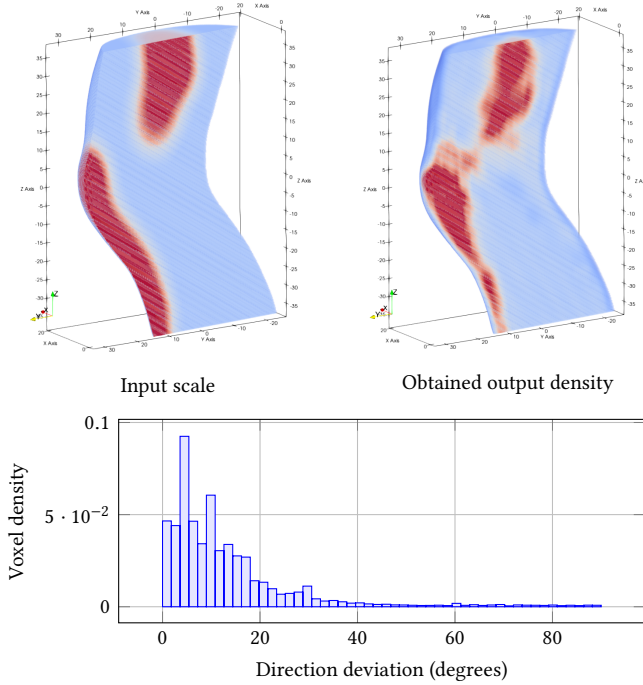


Fig. 17. Fidelity of the synthesized geometry to the control fields, for the 3D printed knee example (Figure 29). Although the output density suffers from filtering artifacts at the geometry boundary, the rest of the model correlates strongly with the input scaling field (Pearson correlation coefficient: 0.814). The extracted orientation also follows the input orientation closely as the bottom histogram shows, with most structures within 20 degrees of the input direction.

a quasi-equivalent response along u, v and w . $\Gamma = \frac{\pi}{2}$ makes the material roughly equivalent along its three main axes. We provide further analysis in supplemental material, Section 6 (for varying densities in particular).

We note that the diamond shape induces a stiffer response along the $v + w$ and $v - w$ axes compared to any other direction in the v, w plane. Nevertheless this behavior has little impact when designing parts, especially as the orientation gradations and the stochasticity of the synthesis process break exact alignments.

Finally, we have compared our method with that of [Martínez et al. 2018] as implemented in the software IceSL¹. The parameters of each method are adjusted in order to obtain transversal rigidity. The results are shown in Table 1. The experiments reveal that our method obtains a higher contrast in transversal rigidity, which is reflected in a greater collapsibility in the transverse axis (see Figure 19).

6 PRINTED EXAMPLES

In this section we design and fabricate several examples with our method, demonstrating the variety of deformation behaviors that can be achieved by manipulating the control fields (orientation, scale, angular spread).

¹<https://icesl.loria.fr/>

Table 1. We compare our approach having $\gamma = 0$ with [Martínez et al. 2018] having $\mu = 0.85$ (see article for details of the effect of parameter μ) to produce two comparable structures exhibiting transverse rigidity. The specimens are cubes of 30 mm of side with similar density. We perform compression tests to evaluate the directional Young’s modulus (MPa) in each orthogonal direction. Our method creates a structure that is more compliant on x an y , stiffer along z , thus manifesting a greater rigidity contrast (see Figure 19).

Axis	[Martínez et al. 2018]	Ours
x	0.76 MPa	0.12 MPa
y	0.60 MPa	0.08 MPa
z	1.175 MPa	2.78 MPa

User painting tool. The fields are stored and manipulated as low resolution 3D textures (64^3), which are directly painted by users through a dedicated interface shown in Figure 20. We describe here our simple approach for this, but of course many other options would be possible.

The user selects a brush radius, a scale factor (for v, w), a value for Γ , and paints onto the shape. The stroke direction is used as the vector u . The brush updates the volume field within a radius around its center. To paint inside the volume, the user can interactively move a cut plane through the part.

Direct visual feedback is obtained by rendering the phasor local parameterizations along the surface or cuts (without singularity correction). This gives an accurate indication of the structures that will be locally synthesized.

Usage guidelines. We found the analogy with long rigid fibers to be effective when designing the fields. Thanks to the very high rigidity contrast of our structures, a designer can reason in only two terms: directions that will be rigid (as if long rigid fibers were embedded in the print) and directions that will be flexible. Under deformation, the fibers resist stretch and pull/push other parts of the volume alongside them. These other parts then easily collapse along flexible directions. These effects cascade throughout the volume.

In addition to these base principles, some typical patterns exist that can be reused in many parts. For instance, the cylindrical collapse of Figure 1 generates a gripper/engulfing behavior also visible in Figures 26, 22. Radial and tangential patterns trigger (respectively) collapse or absorption effects as in Figure 24. Changes to Γ cancel out the contrast achieving different effects along a shape, as in Figure 25. These behaviors can combine towards more complex effects such as in Figures 23, 29, 27, 28. We next detail each of these examples.

Examples. Our first example, shown in Figure 1, is a cube filled with a complex structure. The direction field rotates such that the transverse rigidity remains parallel to an axis, located in the middle of the *top* face. Some angular spread is added to the opposite (bottom) face, making it more rigid in all directions. Pressing on the top face collapses the cube while wrapping the top around the pressing object, an unexpected gripping behavior for the unsuspecting observer. Pressing on the bottom face produces nothing special, the cube reacts as if made of a homogeneous material. Such gripping behaviors are useful for soft robotics.

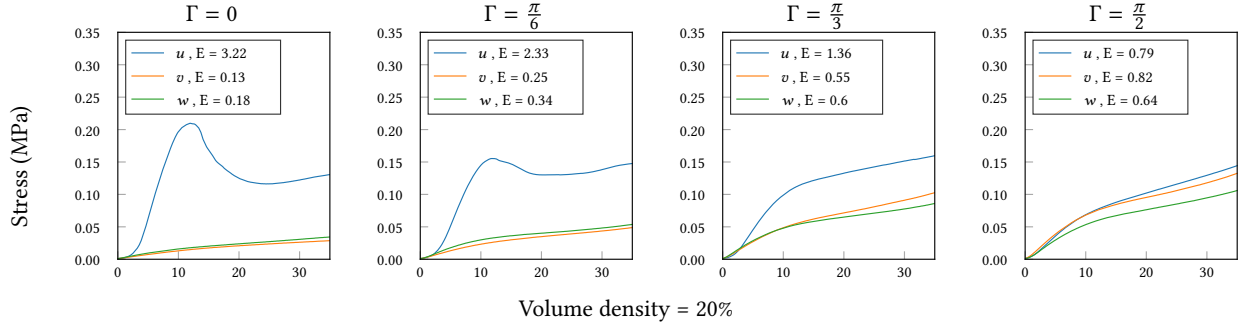


Fig. 18. Compression test of different specimens (3D printed cubes). Scaling chosen to obtain 20% volume density. Showing the varying amount of angular spread Γ , from 0 to quasi-isotropy ($\pi/2$). For each specimen we measure the stress-strain response [Jones 1975] in the three orthogonal directions u , v , w denoted with different colors. The Young's modulus E is in Mpa. The printing direction is along the axis w . Please refer to the text for discussion.




Fig. 19. Visual comparison of the compliance of our approach (right) compared to [Martinez et al. 2018] (left) when a 20N load is applied along the x axis.

Figure 24 shows a thick circular plate, where one face has rigidity oriented radially and the other in a tangential manner. This reveals how the direction field alone strongly influences the deformation. In Figure 24 (bottom) we see that pressing the center has no global effect since the structures flex and absorb the distortion. In Figure 24 (top) the radial arrangement of rigidity makes the entire plate bend under pressure. This is easily explained by the orientation of the structures that prevent distances to change within the volume along specific directions. Note that this very different behavior occurs on both sides of the same object.

Figure 23 is another circular plate where a tangential spiral displacement occurs under radial compression, thanks to the specific arrangement produced by the direction field. This example could also act as a shock-absorbing wheel for a robot or rover.

Figure 27 is a toy mechanism inspired by Figure 1 right. The fields are overall the same, but an internal geometry allows a column to raise from the opposite face when pressing on the sides. This can serve as an actuator or a lateral pressure sensor (measuring the displacement of the column).

Figure 26 is an example design of a simple chair. We oriented the structure to follow the seat surface and be less dense near it.

Figure 28 is a pair of pliers synthesized with our method. The jaws and the hand grips are denser whereas the joint part is softer and uses a radial field to guide the deformation.

Our last and example is an illustration of a knee-like ‘joint-less’ articulation, shown in Figure 29 and which principle is more clearly isolated in Figure 21. A flexible front plate is created using a low rigidity contrast. Behind, two main directions are used in a way that 1. allows the volume to collapse where we want the fold to occur and 2. resists lateral bending. We also added a region of low contrast to mimic the bone in the upper and lower leg. This inner core is easily felt when pressing on the sides and gives vertical rigidity to the overall part.

To the best of our knowledge, no other available approach can produce results with such deformation behaviors under large displacements. Freedom of gradation in orientation, density (through scale), and rigidity contrast (through Γ) interact to provide very effective control over the volume reshaping under external loads. Our examples are only scratching the surface of potential applications in soft robotics, soft mechanisms, prosthetics, and orthoses design, as well as shock or vibrations absorbing pads.

7 LIMITATIONS

Our structures are stiffer along directions $v + w$, $v - w$ while they would ideally be equivalently flexible for all directions in the v , w plane. On less demanding fabrication technologies, using a hexagonal profile would reduce this effect.

Our current implementation does not adapt the deposited thickness with respect to the local wall slope, which implies that density varies with slope. While we did not observe detrimental effects in our results (where walls are often everywhere slanted), a varying width deposition would provide a more accurate result [Kuipers et al. 2020].

Finally, having the ability to perform simulation – even under crude approximations – would help explore more subtle designs and effects, and help the adoption of the approach for engineering applications. This is however an orthogonal, albeit important and difficult problem.

8 CONCLUSION

By allowing to design where and how a volume compresses or extends when manipulated, our method enables unprecedented control over 3D printed objects. We made careful trade-offs to obtain a balance between reliability of fabrication – including on the most




Fig. 20. User interface for designing a frame field. From **left to right** : The original field with only one direction, The field with another direction painted along the **red** line, a section view of the field (section in **green**), the field with one more direction painted along the **blue** line on the **green** section.




Fig. 21. Principle of a knee-like flexure. A flexible front plate provides the overall strength. The V-shaped transverse-rigid structures behind resist lateral bending, while freely collapsing at the rear to let the front flex in the unique allowed direction (at comparable stress levels).




Fig. 22. A structured plate. The orientations produce a strong collapsing/gripping effect when pressing on a specific location.




Fig. 23. A wheel design with a spiraling direction field. **Left**: The structure at rest, with black lines drawn on it to better outline the deformation. **Right**: The wheel under radial compression. Note the spiral motion occurring around the center, thanks to the spiraling fiber structures.

affordable filament printers — and freedom of gradation of the mechanical properties. Experimental tests confirm that a high-contrast




Fig. 24. **Top**: Thick disc with radial fiber pattern. Pressing in the center lifts the sides up. **Bottom**: With a circular direction field, the central pressure is absorbed and has no effect on the sides.




Fig. 25. Plate with a progressive gradation of Γ along its length. One end bends easily across its width while the other is overall more rigid. If pressed on a specific point it bends out of plane, due to how the deformations are absorbed on one side and transferred to the other.

transverse rigidity is achieved over a wide range of deformations. The ability to print walls at the thinnest thickness with a single material allows to produce very flexible objects using off-the-shelf materials.

Our technique is defined from a stochastic approach. This affords for freedom of gradation of the rigid axis orientation, as well as spatial gradation in density and rigidity contrast. Thanks to its favorable computational properties, our technique allows rapid, efficient extraction of toolpaths and is highly scalable.




Fig. 26. Seat printed with our method, a lower frequency is used on the top of the structure to make it softer




Fig. 27. Example of a soft mechanism created with our method. Lateral pressure on the sides of the object induces a deformation that makes the embedded cylinder move up, lifting a small object.




Fig. 28. Pliers printed with our method. The direction control induces the movement. Scaling control is used to rigidify the hand grips and the jaws, only allowing bending in the joint part.

In terms of applications, an important feature of our method is that it works on low-end 3D printers using inexpensive materials. This will allow enthusiasts, makers, and students in engineering to experiment with novel designs. We encourage this by releasing the technique in a freely available design/3D printing software.

Code availability. We provide the full implementation of a 2D version, available as a Shadertoy <https://www.shadertoy.com/view/WtjfwW>.





Fig. 29. Knee 3D printed using our diamond structures. (a) Input model: in blue the direction field D , and in green the angular spread field. (b) shows the generated toolpaths. (c) and (e) show the printed model at rest (d): Knee bended. (f): Stress is applied across the section. The less dense regions collapse, while the denser parts remain the same.

ACKNOWLEDGMENTS

This work was supported partly by the french PIA project « Lorraine Université d'Excellence » (ANR-15-IDEX-04-LUE), by the ANRMuF-Fin (ANR-17-CE10-0002), and the project IMPRIMA(ANR-18-CE46-0004). We thank Salim Perchy and Pierre Bedell for their help in the integration of the method in IceSL and for their help with 3D printers.

REFERENCES

Niels Aage, Erik Andreassen, Boyan S Lazarov, and Ole Sigmund. 2017. Giga-voxel computational morphogenesis for structural design. *Nature* 550, 7674 (2017), 84. <https://doi.org/10.1038/nature23911>

M.F Ashby. 2006. The properties of foams and lattices. *Philosophical Transactions of the Royal Society A* 364, 1838 (2006), 15–30. <https://doi.org/10.1016/b978-0-08-096668-7.00028-0>

David Bommes, Bruno Lévy, Nico Pietroni, Enrico Puppo, Claudio Silva, Marco Tarini, and Denis Zorin. 2013. Quad-Mesh Generation and Processing: A Survey. *Comput. Graph. Forum* 32, 6 (2013), 51–76. <https://doi.org/10.1111/cgf.12014>

James Brennan-Craddock. 2011. *The Investigation of a Method to Generate Conformal Lattice Structures for Additive Manufacturing*. Ph.D. Dissertation. Loughborough Univ.

Yong Chen. 2007. 3D Texture Mapping for Rapid Manufacturing. *Computer-Aided Design and Applications* 4, 6 (2007), 761–771. <https://doi.org/10.1080/16864360.2007.10738509>

- Zuenko Evgeny and Matthias Harders. 2019. Wrinkles, Folds, Creases, Buckles: Small-Scale Surface Deformations as Periodic Functions on 3D Meshes. *IEEE Trans. Vis. Comput. Graph.* PP (05 2019), 1–1. <https://doi.org/10.1109/tvcg.2019.2914676>
- Marc L.M. François, Letian Chen, and Michel Coret. 2017. Elasticity and symmetry of triangular lattice materials. *Int. J. Solids Struct.* 129, Supplement C (2017), 18–27. <https://doi.org/10.1016/j.ijsolstr.2017.09.019>
- Oleg Fryazinov, Turlif Vilbrandt, and Alexander Pasko. 2013. Multi-Scale Space-Variant FRep Cellular Structures. *Computer-Aided Design* 45, 1 (2013), 26–34. <https://doi.org/10.1016/j.cad.2011.09.007>
- Xifeng Gao, Wenzel Jakob, Marco Tarini, and Daniele Panozzo. 2017. Robust Hex-Dominant Mesh Generation Using Field-Guided Polyhedral Agglomeration. *ACM Transactions on Graphics* 36, 4 (2017), 1–13. <https://doi.org/10.1145/3072959.3073676>
- Perle Geoffroy-Donders, Grégoire Allaire, and Olivier Pantz. 2020. 3-d topology optimization of modulated and oriented periodic microstructures by the homogenization method. *J. Comput. Phys.* 401 (2020), 108994. <https://doi.org/10.1016/j.jcp.2019.108994>
- Lorna J Gibson and Michael F Ashby. 1999. *Cellular solids: structure and properties*. Cambridge university press.
- Alexander Goldberg, Matthias Zwicker, and Frédo Durand. 2008. Anisotropic Noise. *ACM Transactions on Graphics* 27, 3 (Aug. 2008), 1–8. <https://doi.org/10.1145/1399504.1360653>
- Jeroen P Groen, Florian C Stutz, Niels Aage, Jakob A Bærentzen, and Ole Sigmund. 2020. De-homogenization of optimal multi-scale 3D topologies. *Computer Methods in Applied Mechanics and Engineering* 364 (2020), 112979.
- Alexandra Ion, Johannes Frohnhofer, Ludwig Wall, Robert Kovacs, Mirela Alistar, Jack Lindsay, Pedro Lopes, Hsiang-Ting Chen, and Patrick Baudisch. 2016. Metamaterial Mechanisms. In *UIST '16*. 529–539.
- Wenzel Jakob, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Instant Field-Aligned Meshes. *ACM Transactions on Graphics* 34, 6 (Nov. 2015), 1–15. <https://doi.org/10.1145/2816795.2818078>
- Robert M Jones. 1975. *Mechanics of Composite Materials*. Vol. 193. Scripta Book Company Washington, DC.
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2015. Stripe Patterns on Surfaces. *ACM Trans Graph.* 34, 4 (2015). <https://doi.org/10.1145/2767000>
- Tim Kuipers, Eugeni L. Doubrovski, Jun Wu, and Charlie C. L. Wang. 2020. A framework for adaptive width control of dense contour-parallel toolpaths in fused deposition modeling. [arXiv:cs.GR/2004.13497](https://arxiv.org/abs/2004.13497)
- Tim Kuipers, Jun Wu, and Charlie C.L. Wang. 2019. CrossFill: Foam Structures with Graded Density for Continuous Material Extrusion. *Computer-Aided Design* 114 (2019), 37 – 50. <https://doi.org/10.1016/j.cad.2019.05.003>
- Ares Lagae, Sylvain Lefebvre, Rob Cook, Tony Derose, George Drettakis, David S. Ebert, J.P. Lewis, Ken Perlin, and Matthias Zwicker. 2010a. A Survey of Procedural Noise Functions. *Computer Graphics Forum* 29, 8 (2010), 2579–2600.
- Ares Lagae, Sylvain Lefebvre, Rob Cook, Tony DeRose, George Drettakis, D. S. Ebert, J. P. Lewis, Ken Perlin, and Matthias Zwicker. 2010b. A Survey of Procedural Noise Functions. *Computer Graphics Forum* 29, 8 (2010), 2579–2600. <https://doi.org/10.1111/j.1467-8659.2010.01827.x>
- Ares Lagae, Sylvain Lefebvre, George Drettakis, and Philip Dutré. 2009. Procedural Noise using Sparse Gabor Convolution. *ACM Transactions on Graphics* 28, 3 (2009), 54–64. <https://doi.org/10.1145/1576246.1531360>
- J. P. Lewis. 1989. Algorithms for Solid Noise Synthesis. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques*. 263–270. <https://doi.org/10.1145/74333.74360>
- Dawei Li, Ning Dai, Xiaotong Jiang, and Xiaosheng Chen. 2015. Interior Structural Optimization Based on the Density-Variable Shape Modeling of 3D Printed Objects. *Int. J. Adv. Manuf. Technol.* 83, 9 (2015), 1627–1635. <https://doi.org/10.1007/s00170-015-7704-z>
- Nils Lichtenberg, Noeska Smit, Christian Hansen, and Kai Lawonn. 2018. Real-time field aligned stripe patterns. *Computers & Graphics* 74 (Aug. 2018), 137–149. <https://doi.org/10.1016/j.cag.2018.04.008>
- Marco Livesu, Stefano Ellero, Jonàs Martínez, Sylvain Lefebvre, and Marco Attene. 2017. From 3D Models to 3D Prints: An Overview of the Processing Pipeline. *Computer Graphics Forum* 36 (2017), 537–564. <https://doi.org/10.1111/cgf.13147>
- S.K. Maiti, M.F. Ashby, and L.J. Gibson. 1984. Fracture toughness of brittle cellular solids. *Scripta Metallurgica* 18, 3 (1984), 213 – 217. [https://doi.org/10.1016/0036-9748\(84\)90510-6](https://doi.org/10.1016/0036-9748(84)90510-6)
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi Foams for Additive Manufacturing. *ACM Transactions on Graphics* 35, 4 (2016), 44:1–44:12. <https://doi.org/10.1145/2897824.2925922>
- Jonàs Martínez, Samuel Hornus, Haichuan Song, and Sylvain Lefebvre. 2018. Polyhedral Voronoi diagrams for additive manufacturing. *ACM Transactions on Graphics* 37, 4 (2018), 15. <https://doi.org/10.1145/3197517.3201343>
- Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. 2017. Orthotropic k-nearest foams for additive manufacturing. *ACM Transactions on Graphics* 36, 4 (2017), 1–12. <https://doi.org/10.1145/3072959.3073638>
- Asla Medeiros e Sá, Vinicius Moreira Mello, Karina Rodriguez Echavarría, and Derek Covill. 2015. Adaptive voids. *The Visual Computer* 31, 6 (2015), 799–808.
- Jean-Claude Michel, Hervé Moulinec, and P Suquet. 1999. Effective properties of composite materials with periodic microstructure: a computational approach. *Computer methods in applied mechanics and engineering* 172, 1-4 (1999), 109–143. [https://doi.org/10.1016/S0045-7825\(98\)00227-8](https://doi.org/10.1016/S0045-7825(98)00227-8)
- Fabrice Neyret and Eric Heitz. 2016. *Understanding and controlling contrast oscillations in stochastic texture algorithms using Spectrum of Variance*. Research Report. LJK / Grenoble University - INRIA. 8 pages.
- nTopology. 2019. nTop design platform. <https://ntopology.com>.
- Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-case Stress Relief for Microstructures. *ACM Transactions on Graphics* 36, 4 (2017), 122:1–122:16. <https://doi.org/10.1145/3072959.3073649>
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Transactions on Graphics* 34, 4 (2015), 135:1–135:12. <https://doi.org/10.1145/2766937>
- Olivier Pantz and Karim Trabelsi. 2008. A Post-Treatment of the Homogenization Method for Shape Optimization. *SIAM Journal on Control and Optimization* 47, 3 (2008), 1380–1398. <https://doi.org/10.1137/070688900>
- Alexander Pasko, Oleg Fryazinov, Turlif Vilbrandt, Pierre-Alain Fayolle, and Valery Adzhiev. 2011. Procedural function-based modelling of volumetric microstructures. *Graphical Models* 73, 5 (2011), 165–181. <https://doi.org/10.1016/j.gmod.2011.03.001>
- Minh-Son Pham, Chen Liu, Iain Todd, and Jedsada Lertthanasarn. 2019. Damage-tolerant architected materials inspired by crystal microstructure. *Nature* 565, 7739 (2019), 305. <https://doi.org/10.1038/s41586-019-0968-y>
- Carlos M Portela, A Vidyasagar, Sebastian Krödel, Tamara Weissenbach, Daryl W Yee, Julia R Greer, and Dennis M Kochmann. 2020. Extreme mechanical resilience of self-assembled nanolabyrinthine materials. *Proceedings of the National Academy of Sciences* 117, 11 (2020), 5686–5693. <https://doi.org/10.1073/pnas.1916817117>
- Nicolas Ray, Wan Chiu Li, Bruno Lévy, Alla Sheffer, and Pierre Alliez. 2006. Periodic Global Parameterization. *ACM Trans. Graph.* 25 (2006), 1460–1485. <https://doi.org/10.1145/1183287.1183297>
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Transactions on Graphics* 34, 4 (2015), 136:1–136:13. <https://doi.org/10.1145/2766926>
- Christian Schumacher, Steve Marschner, Markus Cross, and Bernhard Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Transactions on Graphics* 37, 4, Article 148 (July 2018), 15 pages. <https://doi.org/10.1145/3197517.3201278>
- Lucas A. Shaw, Frederick Sun, Carlos M. Portela, Rodolfo I. Barranco, Julia R. Greer, and Jonathan B. Hopkins. 2019. Computationally efficient design of directionally compliant metamaterials. *Nature Communications* 10 (2019). <https://doi.org/10.1038/s41467-018-08049-1>
- Kumar Siddhant, Stephanie Tan, Zheng Li, and Dennis M Kochmann. 2020. Inverse-designed spinodoid metamaterials. *npj Computational Materials* 6, 1 (2020). <https://doi.org/10.1038/s41524-020-0341-6>
- Emmanuel Siefert, Etienne Reyssat, José Bico, and Benoit Roman. 2019. Bio-inspired pneumatic shape-morphing elastomers. *Nature Materials* 18 (01 2019). <https://doi.org/10.1016/j.cirp.2020.03.006>
- M.G. Tarantino, O. Zerhouni, and K. Danas. 2019. Random 3D-printed isotropic composites with high volume fraction of pore-like polydisperse inclusions and near-optimal elastic stiffness. *Acta Materialia* 175 (2019), 331 – 340. <https://doi.org/10.1016/j.actamat.2019.06.020>
- Thibault Tricard, Semyon Efremov, Cédric Zanni, Fabrice Neyret, Jonàs Martínez, and Sylvain Lefebvre. 2019. Procedural Phasor Noise. *ACM Transactions on Graphics* 38, 4 (July 2019), 57:1–13. <https://doi.org/10.1145/3306346.3322990>
- Kiril Vidimčec, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik. 2013. OpenFab: A Programmable Pipeline for Multi-material Fabrication. *ACM Transactions on Graphics* 32, 4 (2013), 136:1–136:12. <https://doi.org/10.1145/2461912.2461993>
- Hongqing Wang, Yong Chen, and David W Rosen. 2005. A hybrid geometric modeling method for large scale conformal cellular structures. In *ASME 2005 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. 421–427.
- Jun Wu, Niels Aage, Rüdiger Westermann, and Ole Sigmund. 2018. Infill Optimization for Additive Manufacturing—Approaching Bone-Like Porous Structures. *IEEE Trans. Vis. Comput. Graph.* 24, 2 (2018), 1127–1140. <https://doi.org/10.1109/tvcg.2017.2655523>
- Jun Wu, Charlie C.L. Wang, Xiaoting Zhang, and Rüdiger Westermann. 2016. Self-supporting rhombic infill structures for additive manufacturing. *Computer-Aided Design* 80 (2016), 32–42. <https://doi.org/10.1016/j.cad.2016.07.006>
- Jun Wu, Weiming Wang, and Xifeng Gao. 2019. Design and Optimization of Conforming Lattice Structures. *IEEE Trans. Vis. Comput. Graph.* (2019), 1–1. <https://doi.org/10.1109/tvcg.2019.2938946>
- Bo Zhu, Mélima Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-Scale Topology Optimization with Microstructures. *ACM Transactions on Graphics* 36, 5 (2017), 1. <https://doi.org/10.1145/3072959.3095815>