

Two-Scale Topology Optimization with Microstructures

BO ZHU, MÉLINA SKOURAS, DESAI CHEN, WOJCIECH MATUSIK, MIT CSAIL




Fig. 1. Our two-scale topology optimization framework allows to optimize continuous material properties mapping to printable microstructures (*left*) to fabricate high-resolution functional objects (*middle*) and minimum compliant structures (*right*).

In this paper we present a novel two-scale framework to optimize the structure and the material distribution of an object given its functional specifications. Our approach utilizes multi-material microstructures as low-level building blocks of the object. We start by precomputing the material property gamut – the set of bulk material properties that can be achieved with all material microstructures of a given size. We represent the boundary of this material property gamut using a level set field. Next, we propose an efficient and general topology optimization algorithm that simultaneously computes an optimal object topology and spatially-varying material properties constrained by the precomputed gamut. Finally, we map the optimal spatially-varying material properties onto the microstructures with the corresponding properties in order to generate a high-resolution printable structure. We demonstrate the efficacy of our framework by designing, optimizing, and fabricating objects in different material property spaces on the level of a trillion voxels, i.e several orders of magnitude higher than what can be achieved with current systems.

CCS Concepts: •Computing methodologies → Physical simulation;

Additional Key Words and Phrases: microstructures, metamaterials, 3D printing, topology optimization

ACM Reference format:

Bo Zhu, Mélina Skouras, Desai Chen, Wojciech Matusik. 2017. Two-Scale Topology Optimization with Microstructures. *ACM Trans. Graph.* 36, 5, Article (July 2017), 17 pages.
DOI: 0000001.0000001_2

1 INTRODUCTION

Many engineering problems focus on the design of complex structures that needs to meet high level objectives such as the capability

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 ACM 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.

© 2017 ACM. 0730-0301/2017/7-ART \$15.00
DOI: 0000001.0000001_2

to support localized stresses, optimal tradeoffs between compliance and mass, minimal deformation under thermal changes, etc. One very popular approach to design such structures is topology optimization. Topology optimization generally refers to discretizing the object of interest into small elements and optimizing the material distribution over these elements in such a way that the functional goals are satisfied [Bendsøe and Sigmund 2004]. Traditionally, topology optimization focused on designs made of homogeneous materials and was concerned with macroscopic changes in the object geometry. With the advent of multi-material 3D printing techniques, it is now possible to play with materials at a much higher resolution, allowing to obtain much finer designs and, thus, improved functional performances. Unfortunately, standard techniques for topology optimization do not scale well and they cannot be run on objects with billions of voxels. This is because the number of variables to optimize increases linearly with the number of cells in the object. Since many current 3D printers have a resolution of 600DPI or more, a one billion voxel design occupies only a 1.67 inch cube.

One direction to handle this issue is to work with microstructures corresponding to blocks of voxels instead of individual voxels directly. Some recent works followed this direction and proposed to decouple macro structural design and micro material design [Coelho et al. 2008; Nakshatrala et al. 2013; Rodrigues et al. 2002]. However, these approaches remain computationally expensive and, in most cases, limited to the well-known minimal compliance problem. The second direction to reduce the problem complexity is to temporarily ignore the geometry of the microstructures and consider only their macroscopic physical behaviour. However, this introduces new difficulties as the space of material properties covered by all printable microstructures is much wider than the properties of the base materials. For example, microstructures made of alternating layers of soft and stiff isotropic materials exhibit an anisotropic behaviour as they are able to stretch more easily in one direction than in the others. This implies that not only the *ranges* but also

the *number* of physical parameters needed to describe the physical behaviour of these microstructures increases. Therefore, in order to work with the material properties of microstructures, one needs to solve two challenging problems: (i) computing the *gamut* – i.e the set – of the material properties achievable by all microstructures, (ii) efficiently optimizing the distribution of these high-dimensional material properties inside the layout of the object.

Most previous algorithms working in the material space focused on optimizing a single material property such as density or material stiffness, for which analytical formulas describing the property bounds exist [Allaire and Kohn 1993]. On the contrary, optimizing the structure and material distribution of an object in a high dimensional material property space remains an open problem. In this work, we propose a new computational framework for topology optimization with microstructures that supports design spaces of multiple dimensions. We start by computing the gamut of the material properties of the microstructures by alternating stochastic sampling and continuous optimization. This gives us a *discrete* representation of the set of achievable material properties, from which we can construct a *continuous* gamut representation using a level set field. We then reformulate the topology optimization problem in the continuous space of material properties and propose an efficient optimization scheme that finds the optimized distributions of multiple material properties simultaneously inside the gamut. Finally, in order to obtain fabricable designs, we map the optimal material properties back to discrete microstructures from our database.

Our general formulation can be applied to a large variety of problems. We demonstrate its efficacy by designing and optimizing objects in different material spaces using isotropic, cubic and orthotropic materials. We apply our algorithm to various design problems dealing with diverse functional objectives such as minimal compliance and target strain distribution (see Figure 1). Furthermore, our approach utilizes the high-resolution of current 3D printers by supporting designs with trillions of voxels. We fabricate several of our designs, thus, demonstrating the practicality of our approach.

The main contributions of our work can be summarized as follows:

- We present a fully automatic method for computing the space of material properties achievable by microstructures made of a given set of base materials.
- We propose a generic and efficient topology optimization algorithm capable of handling objects with a trillion voxels. The key of our approach is a reformulation of the problem to work directly on continuous variables representing the material properties of microstructures. This allows us to cast topology optimization as a reasonably sized constrained optimization problem that can be efficiently solved with state of the art solvers.
- We validate our method on a set of test cases and demonstrate its versatility by applying it to various design problems of practical interest.

2 RELATED WORK

Topology Optimization. Topology optimization is concerned with the search of the optimal distribution of one or more materials within a design domain in order to minimize some input objective function

while satisfying given constraints [Bendsøe and Sigmund 2004]. Initially applied to the structural design in engineering [Bendsøe 1989], topology optimization has been extended since then to a variety of problems including micromechanism design [Sigmund 1997], mass transfer [Challis and Guest 2009], metamaterial design [Cadman et al. 2013; Sigmund and Torquato 1996], multifunctional structure design [Yan et al. 2015], coupled structure-appearance optimization [Martinez et al. 2015]. Many algorithms have been proposed to numerically solve the optimization problem itself. We refer to the survey by Sigmund and Maute [2013] for a complete review. In the very popular SIMP (Solid Isotropic Materials with Penalization) method, the presence of material in a given cell is controlled by locally varying its density. A binary design is eventually achieved by penalizing intermediate values for these densities. In practice, this method works well for two-material designs (e.g., a material and a void), but generalizing this method to robustly handle higher dimensional material spaces remains challenging. Instead of considering only discrete structures, free material optimization [Haber et al. 1994; Ringertz 1993] optimizes structures made of continuous material distributions constrained by analytical bounds. Another class of methods rely on *homogenization*. They replace the material in each voxel of the object by a mixture of the base materials whose material properties can be analytically derived. While optimal microstructures are known for certain classes of problems (laminated composites in the case of the minimum compliance problem), this is not the case in the general setting, for which using a specific subclass of microstructures can lead to suboptimal results. In a sense, our work is a generalization of these approaches and aims to handle a wider range of materials for which theoretical bounds on the material properties are not known a priori.

Although they are largely used in engineering, standard methods for topology optimization suffer from a major drawback : the parametrization of the problem at the voxel level makes them extremely expensive and largely impedes their use on high resolutions models such as the ones generated by modern 3D printing hardware. High-performance GPU implementations with careful memory handling can be used to push the limits of what can be done (a couple of million variables in the implementation by Wu et al. [2016]), but such approaches rely on specificities of the minimum compliance problem and are difficult to generalize. To counteract the effects of the explosion of variables in finely discretized layouts, Rodrigues et al. [2002] alternatively proposed an interesting formulation where microstructure designs and macroscopic layouts using the effective properties of the underlying microstructures were hierarchically coupled and treated simultaneously. This initial work has been extended in multiple ways [Coelho et al. 2008; Nakshatrala et al. 2013; Xia and Breitkopf 2014; Yan et al. 2014]. Alexandersen and Lazarov [2015] proposed a fast simulation algorithm for optimizing complete macroscopic structures made of layered or periodic microstructures. However, these methods still need to handle variables defined at the microstructure level and therefore they remain relatively costly. The most related work is the method proposed by Xia et al.[2015b], which also relies on a database to speed up computations. However, their work specifically targets minimum compliance problems in the structural design which allows them to

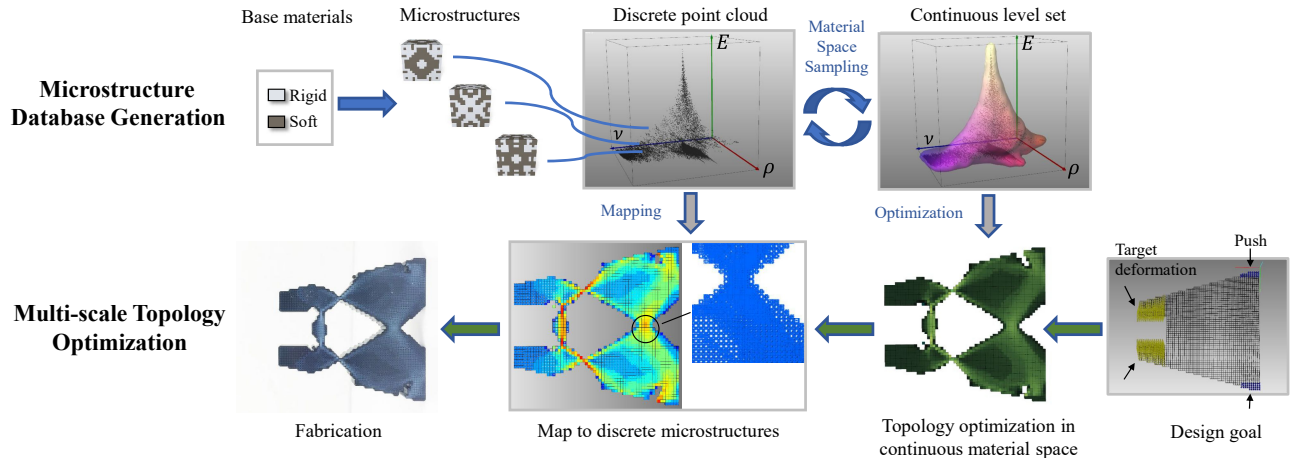


Fig. 2. Algorithm overview. We start by precomputing the gamut of material properties that can be achieved with all material microstructures of a given size. Next, we run our topology optimization algorithm that optimizes the material properties of the object within this gamut such as to minimize some functional objective. Finally, we map the optimal continuous material properties back to microstructures from our database to generate a printable object.

approximate the macroscale behaviour of the microstructures with a particular strain-based interpolating function.

Fabrication-oriented Optimization. The last decade has witnessed an increasing interest by the computer graphics community in the design of tools and algorithms targeting digital fabrication of physical artifacts. The range of media and applications addressed in previous literature is very diverse and we focus our discussion on systems targeting 3D printing. The problem of optimizing the material assignment for the individual voxels of an object in order to control its large scale behaviour has been studied in different contexts. Starting with optical properties, Hašan et al. [2010] and Dong et al. [2010] provided methods for printing objects with desired subsurface scattering properties. Stava et al. [2012] later considered stability of 3D printed objects, Zhou et al. [2013] explored structural strength while Chen et al. [2014] focused on rest shape optimization. Closer to our present work, frameworks for the design of objects with desired mechanical behaviours have been proposed by Bickel et al. [2010] and Skouras et al. [2013]. Like these works, our system allows to match given input deformations. However, while these previous systems assume a small set of available base materials and use these base materials in relatively coarse discretizations, our system combines the base materials into microstructures to expand the design possibilities. Also relevant is the tool presented by Xu et al. [2015] that allows to interactively design heterogeneous materials for elastic objects subject to prescribed displacements and forces, and the material optimization approach proposed by Panetta et al. [2015]. However, these methods may output materials that are not available in the real world for non-convex manifolds of material properties. By contrast, we guarantee that all the microstructures used are always realizable in such cases, which is one of the key contributions of our work. Lastly, in an effort to unify individual contributions when dealing with inverse modeling problems, Chen et al. [2013] proposed an abstraction mechanism to facilitate the development of goal-based methods. The output of most of

these systems is a per-voxel material composition, which cannot be efficiently represented using simple surface meshes. Vidimče et al. [2013] introduced a fabrication-specific language and a programming pipeline for a procedural material synthesis that lift this limitation.

Microstructures and Metamaterials. Microstructures can be defined as small scale assemblies made of one or several base materials, whose macroscale properties can be very different from those of the original materials. Many materials found in the nature are microstructures when observed at a sufficiently small scale. Microstructures can also be engineered so as to define composites with improved capabilities or even *metamaterials* with exceptional properties. For example, Lakes [1987] presented in 1987 the first man-made structure with negative Poisson’s ratio, i.e., a structure which transversally expands when it is axially stretched. The design of composites and metamaterials is an active research field inspiring myriads of works [Andreassen et al. 2014; Babaee et al. 2013; Cadman et al. 2013; Sigmund 1997; Sigmund and Torquato 1996; Wang et al. 2014]. While many of these works are concerned with the inverse modeling of specific microstructures or families of microstructures, the study of the *space* of properties that these microstructures can achieve as a whole has been investigated much less. Theoretical bounds have been derived without experimental validation [Lipton 1994; Milton and Cherkaev 1995; Ting and Chen 2005]. Taking into account additive manufacturing constraints, Schumacher et al. [2015] and Panetta et al. [2015] recently investigated the design of tileable and printable microstructures. In the first part of this paper, we further explore this line of research and focus on the generation and characterization of databases of microstructures with maximal material property coverage. In particular, we present a novel approach combining a probabilistic search and a continuous optimization that allows us to *fully automatically* explore the gamut of material properties that can be achieved by assembling given base materials.

3 OVERVIEW

Given as input a set of base materials, an object layout, and functional objectives, the goal of our system is to compute the material distribution inside the object in order to optimize these functional objectives. In our approach, we do not solve the problem directly, instead we work with microstructures made of the base materials and the space of physical material properties spanned by them. The complete pipeline of our system, illustrated in Figure 2, can be decomposed into three stages.

Material Space Precomputation. In the first stage, we estimate the gamut of material properties covered by all possible microstructures made by spatial arrangement of base materials. Since exhaustively computing the properties of all these microstructures is, in practice, intractable, we progressively increase the material space by alternating a stochastic search and a continuous optimization. The first step introduces discrete changes in the materials of the microstructures and allows emergence of new types of microstructures. The second step allows to locally push the material space boundaries by refining the microstructure shapes. After completing this stage, we obtain a discrete representation of the space of material properties and the mapping between these properties and the corresponding microstructures.

Gamut-based Continuous Topology Optimization. In the second stage, we construct a smooth continuous gamut representation of the material property space by using a level set field. We define our topology optimization problem directly in this space. Our approach minimizes the objective function over possible material parameters while asking for strict satisfaction of the physics constraints – typically, the static equilibrium – as well as the strict satisfaction of the physical parameter bounds. Taking advantage of our gamut representation as a level set, we formulate this last constraint as limiting the material properties to stay on the negative side of the level set. This guarantees that the material properties that we use in the optimization are always physically realizable.

Fabrication-oriented Microstructure Mapping. In the last stage, we generate a printable result by replacing each cell in the object layout with a microstructure whose material properties are the closest to the continuous material assignment resulting from the optimization. We also take into account the boundary similarity across adjacent cell interfaces to improve the connectivity between microstructures. This results in a complex, high-resolution, multi-material model with optimized functional specifications.

4 MECHANICS

In this section, we briefly introduce the background material for simulating deformable objects. We will use these concepts when computing the material properties of the microstructures (Section 5) and in the topology optimization algorithm (Section 6). We refer to the course by Sifakis and Barbic [2012] for a more comprehensive exposition.

4.1 Material Model

Most available materials for 3D printers are elastic materials. Assuming small deformations, we use linear elasticity to compute both

the mechanical behaviour of the entire object and the microstructures. In such a setting, the relation between the linear strain ϵ and Cauchy stress σ at every material point is given by

$$\sigma = \mathbf{C}\epsilon, \quad (1)$$

where \mathbf{C} , the so-called elasticity tensor, can be described by 21 parameters [Bonet and Wood 1997].

Working in such a high dimensional space is prohibitive and therefore we focus on materials having a certain numbers of symmetries, such as orthotropic materials for which the elasticity tensor is defined by 12 parameters (4 parameters in 2D), cubic materials defined by 3 parameters, and isotropic materials defined by 2 parameters. For example, the tensor for a 3D cubic structure can be written (using the Voigt notation) as

$$\mathbf{C} = \begin{pmatrix} (1-\nu)\hat{E} & \nu\hat{E} & \nu\hat{E} & & & \\ \nu\hat{E} & (1-\nu)\hat{E} & \nu\hat{E} & & & \\ \nu\hat{E} & \nu\hat{E} & (1-\nu)\hat{E} & & & \\ & & & \mu & & \\ & & & & \mu & \\ & & & & & \mu \end{pmatrix} \quad (2)$$

with $\hat{E} = E/((1-2\nu)(1+\nu))$ and where E is the Young's modulus of the material, ν its Poisson's ratio, and μ its shear modulus.

Alternatively, one can also use Lamé's parameters to define the tensor \mathbf{C} , which simplifies the derivation of the tensor with respect to the elastic parameters. In this case, the tensor has the form

$$\mathbf{C} = \begin{pmatrix} 2\mu+\lambda & \lambda & \lambda & & & \\ \lambda & 2\mu+\lambda & \lambda & & & \\ \lambda & \lambda & 2\mu+\lambda & & & \\ & & & \mu & & \\ & & & & \mu & \\ & & & & & \mu \end{pmatrix}. \quad (3)$$

The tensor for a 2D orthotropic structure can be written as

$$\mathbf{C} = c \begin{pmatrix} E_x & \nu_{yx}E_x \\ \nu_{xy}E_y & E_y \\ & & \mu/c \end{pmatrix} \quad (4)$$

with $c = 1/(1-\nu_{xy}\nu_{yx})$, and where E_x and E_y are the Young's moduli along the two principal axes, μ is the shear modulus, ν_{xy} is the Poisson's ratio corresponding to a contraction in the direction y when an extension is applied along the x axis, and ν_{yx} verifies $\nu_{yx}E_x = \nu_{xy}E_y$.

Letting \mathcal{R}^n denote the space of n material properties, we then write each point $\mathbf{p} \in \mathcal{R}^n$ as $\mathbf{p} = [\rho, \mathbf{e}]$, where ρ is the density of the material and \mathbf{e} are the other material parameters. Our gamut, i.e. the set $\mathcal{M} \subset \mathcal{R}^n$ of material properties corresponding to microstructures of a given resolution, is made of a finite number of points. However, by increasing the resolution of the microstructures this gamut gets denser and denser so that we assume that it can be approximated by the union of continuous n -dimensional manifolds and can be represented using a distance field.

4.2 Discretization

Following standard finite element methodology, we discretize the object in regular voxels and compute its deformed state when subject to external forces \mathbf{f}_{ext} using the well-known relation

$$\mathbf{K}\mathbf{u} = \mathbf{f}_{\text{ext}}, \quad (5)$$

where \mathbf{K} is the stiffness matrix of the system, and \mathbf{u} are the displacements at the nodes of the voxels.




Fig. 3. Level set gamuts for two dimensional cubic microstructures (*top row*) and three dimensional cubic microstructures (*bottom row*). The first column shows the projection of the sample points in the space parametrized by the density ρ , the normalized Young's modulus \hat{E} and the Poisson's ratio ν . The second to the fourth columns show three slices of the four dimensional level sets corresponding to different values for the shear modulus G .

Note that we use the same approach to simulate both the mechanical behaviour of the microstructures and the object macroscopic behaviour. However, we work at two different scales. To simulate the microstructures, we assume that each of its voxels is made of an homogeneous base material, whereas for determining the large scale behaviour of the object, we assume that each of its cells corresponds to a microstructure. The properties of the individual microstructures are determined from 6 harmonic displacements (or 2 displacements in 2D) using numerical coarsening as described by Kharevych et al. [2009].

We solve the static equilibrium Equation 5 using a fast multigrid solver based on the implementation by Dick et al. [2011].

5 MATERIAL SPACE EXPLORATION

The first step in our pipeline is to determine the space of physical properties that can be achieved when combining the base materials into microstructures of a predefined size.

Computing the mechanical properties of microstructures, when arranged in periodic tilings, can be performed by probing the structure using a physical simulation. This approach, based on the homogenization theory, is a common practice and has been widely used in the past [Allaire 2012; Panetta et al. 2015; Schumacher et al. 2015]. However, while inferring the homogenized properties of individual microstructures is not particularly challenging, analyzing the space covered by *all* combinations of base materials is much more difficult due to the combinatorial explosion in the number of possible material arrangements. As an example, $16 \times 16 \times 16$ lattices made of only two materials corresponds to 2^{4096} microstructures: exhaustively probing all microstructures is clearly an impossible task. To address this issue, two possible avenues can be pursued: (i) we can try to sample the space of the microstructures, (ii) we can rely on the continuity between material parameters of the individual voxels and macroscopic properties of the microstructures in order to

generate new microstructures with desired properties. This second option is effective in reaching locally optimal values in the material property space. However, the function that maps the material assignment to material properties is nonlinear. In particular, very different microstructures can correspond to the same point in the material property space. Additionally, since the ratio of materials in each cell is bounded between zero and one, the continuous optimization converges slowly or stops moving when material distributions in many cells are at the lower or upper bound. Being able to jump out of a local optimum and discovering different variants is important in order to provide new exploration regions. We leverage these two approaches by combining them in a scheme that alternates between a stochastic search and a continuous optimization. We provide the technical details in the rest of this section.

5.1 Discrete Sampling of Microstructures

We aim at sampling the space of material assignments, i.e. microstructures, in such a way that we maximize the number of samples corresponding to microstructures whose material properties lie in the vicinity of the material gamut boundaries. We do not draw all samples at once but progressively enrich the database of microstructures as we refine our estimation of the material gamut boundaries (see Figure 6). This sampling strategy is motivated by the observation that a small change in the material assignment of a microstructure generally – but not always – translates to a small change of its material properties. By modifying microstructures located near the current boundaries of the material property gamut, we are likely to generate more structures in this area, some of which will lie outside of the current gamut.

Given a population of microstructures to evolve, we generate new samples from each microstructure by changing its material at random voxel locations. To rationalize computational resources, we want to avoid revisiting the same voxel twice. But we do not

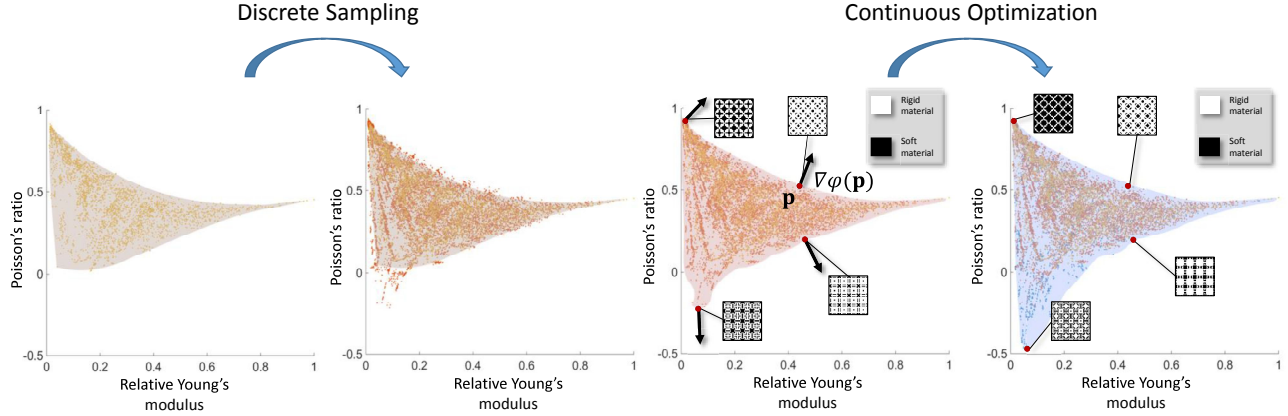


Fig. 4. One cycle of computing the microstructure gamut. Given a set of samples, we compute a signed distance function approximating the material gamut (*left*) and randomly perturb microstructures lying near the boundary to provide new seeds to the continuous algorithm (*middle left*). We then update the distance field and use the gradient of the signed distance function at the the boundary to define new target material points (*middle right*). These target material points are used in a continuous optimization that generates new samples (*right*).

want to privilege any particular order either. Ritchie et al. [2015] recently presented a Stochastically-Ordered Sequential Monte Carlo (SOSMC) method that provides a suitable approach. In SOSMC, a population of particles (here, our microstructures) corresponding to instances of a procedural program (here, the sequential assignment of materials to the voxels of the microstructures) are evolved so as to represent a desired distribution. During this process, the programs are executed in a random order and particles are regularly scored and reallocated in regions of high probability. In our particular settings, we use the scoring function

$$s(\mathbf{p}_i) = \frac{\Phi(\mathbf{p}_i)}{D(\mathbf{p}_i)} \times \frac{1}{D(\mathbf{p}_i)}, \quad (6)$$

where $\Phi(\mathbf{p}_i)$ is the signed distance of the material properties of particle i to the gamut boundary (see Section 5.3) and $D(\mathbf{p}_i)$ is the local sampling density at the location \mathbf{p}_i . We define the sample density as

$$D(\mathbf{p}_i) = \sum_{\mathbf{k}} \phi_{\mathbf{k}}(\mathbf{p}_i), \quad (7)$$

where $\phi_{\mathbf{k}}(\mathbf{p}) = \left(1 - \frac{\|\mathbf{p} - \mathbf{p}_{\mathbf{k}}\|_2^2}{h^2}\right)^4$ are locally-supported kernel functions that vanish beyond their support radius h , set to a tenth of the size of the lattice used for the continuous representation of the material gamut (see Section 5.3).

The first term in Equation 6 favors microstructures located near the gamut boundary. The normalization by D allows us to be less sensitive to the local microstructures density and to hit any location corresponding to the same level-set value with a more uniform probability. The second product is used to additionally privilege under-sampled areas.

Particles are resampled using systematic resampling scheme [Douc 2005] that is also used to initiate the population of particles. These particles are then evolved according to Algorithm 1. Additional details regarding the implementation of our algorithm are provided in the supplementary material.

Algorithm 1 Procedure for generating new microstructures

procedure GENMICROSTRUCTURE(input: microstructure M_i , output: microstructure M_o)

$M_o \leftarrow M_i$

while some voxels of M_o have not been visited **do**

while microstructure M_o is unchanged **do**

pick a random voxel v of M_o that has not been visited

assign a randomly chosen material to v

if M_o is manifold and $M_o \neq M_i$ **then**

accept the change

end if

end while

end while

end procedure

5.2 Continuous Optimization of Microstructures

The goal of the continuous optimization is to refine the geometry of the microstructures located at the boundary of the gamut in order to further expand along the normal directions (Figure 5).

We start continuous sampling by selecting a subset of microstructures lying on the boundary of the gamut as starting points for the continuous optimization. The discrete structures are mapped to continuous values close to 0.5. We used 0.5 ± 0.3 in our experiments. Doing so allows the topology optimization algorithms to move freely in the first steps and discover new structures. We show an example of reducing the Poisson's ratio of an initial structure in Figure 5. In the plot, the initial Poisson's ratio is close to 0.4 since the starting point is similar to a homogeneous block.

For each starting structure, we identify target material parameters using the gradient of the level set Φ at the initial discrete sample point \mathbf{p} (see Section 5.3) defined by $\mathbf{q} = \mathbf{p} + \nabla\Phi(\mathbf{p})$. We translate this target material parameters into an elasticity tensor \mathbf{C}_0 and density ρ_0 . Here ρ is the ratio of the two base materials in the microstructure.

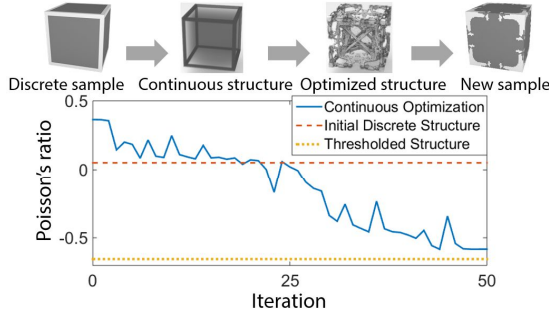


Fig. 5. The continuous sampling step uses topology optimization to expand the gamuts by refining existing structures. A good starting point is necessary for the optimization to find better solutions. We use discrete microstructures near the boundary to initialize topology optimization. At convergence, we threshold the values to obtain a new discrete sample.

Note that our problem formulation does not restrict us to a particular topology optimization algorithm or material distribution parametrization. We have experimented with two objective functions that worked equally well for our purposes. The first objective uses an energy based formulation [Xia and Breitkopf 2015a] to compute and optimize the elasticity tensor directly. At a high level, the optimization problem is

$$\arg \min_{\mathbf{x}} f(\mathbf{x}) = (\mathbf{C}(\mathbf{x}) - \mathbf{C}_0)^2 + w_\rho(\rho - \rho_0)^2, \quad \rho = \sum_i x_i, \quad (8)$$

where \mathbf{x} is the ratio of materials in each cell, and w_ρ controls the weighting between the displacement term and the density term. The authors of the method developed parameter heuristics to optimize for difficult cases such as negative Poisson's ratio structures. We naturally arrive at structures with negative Poisson's ratio without the parameter varying step in [Xia and Breitkopf 2015a] since our discrete samples allow us to explore a wide variety of initial designs.

The second objective is formulated using harmonic displacements [Kharevych et al. 2009; Schumacher et al. 2015] \mathbf{G} instead of the elasticity tensor directly. \mathbf{G} is a 6×6 symmetric matrix where each row corresponds to a strain in vector form. We use the target elasticity tensor \mathbf{C}_0 to compute the target harmonic displacements matrix \mathbf{G}_0 and minimize the objective function:

$$f(\mathbf{x}) = (\mathbf{G}(\mathbf{x}) - \mathbf{G}_0)^2 + w_\rho(\rho - \rho_0)^2. \quad (9)$$

This objective matches soft structures more accurately since entries of \mathbf{G} are inversely proportional to material stiffness.

Following the work by Andreassen et al. [2014], we use the method of moving asymptotes (MMA) [Svanberg 1987] to optimize the objectives using an implementation provided in the NLOPT package [Johnson 2014]. We run at most 50 iterations since it usually converges to a solution within 20-30 steps (Figure 5). MMA makes large jumps during the optimization while keeping track of the current best solution, thus causing the oscillation of the objective value. To force continuous material ratios towards discrete values, we experimented with the SIMP model with the exponent set to 3 and the Hashin-Shtrikman bound for isotropic materials described by Bendsoe and Sigmund [1999].

Either interpolation allows us to threshold the final continuous distribution and obtain a similar discrete sample. We tolerate small deviations introduced by the thresholding since our goal is to obtain a microstructure lying outside of the gamut rather than reaching a particular target. In practice, we observed that the material properties of the final discrete structures often did not change significantly after the thresholding step.

5.3 Continuous Representation of the Material Gamut

We represent the gamut of material properties using a signed distance field that is computed from the material points associated to the sampled microstructures (see Figure 4). First, we normalize each coordinate p_i of \mathbf{p} to constrain the scope of the level set to an n -dimensional unit cube. Then we compute the level set values on the cell centers of an n -dimensional Cartesian grid that encloses this unit cube. We draw inspiration from the methods for surface reconstruction used in particle fluid rendering [Ando et al. 2013; Bhattacharya et al. 2011; Zhu and Bridson 2005] and extend it to n dimensions. In this case, a signed distance field is generated from a set of points by evaluating an implicit distance function Φ at each point $\mathbf{p} \in \mathcal{M}$. We initialize the signed distance field using the implicit function $\Phi(\mathbf{p}) = \|\mathbf{p} - \bar{\mathbf{p}}\| - r$ from [Zhu and Bridson 2005] where $\|\cdot\|$ is the Euclidean distance between two points in \mathcal{M} , and $\bar{\mathbf{p}}$ is the average position of the neighboring points of \mathbf{p} within a range of $2r$. Note that the signed distance is initially defined only near the boundary of the gamut. In order to sample the distance on the entire domain, we propagate the 0-level set surface using the fast marching algorithm and solve an explicit mean curvature flow problem defined as $\partial\Phi/\partial t = \Delta\Phi$ [Osher and Fedkiw 2006].

Having a continuous representation of the gamut of materials achievable by the microstructures, we can now reformulate the topology optimization problem directly in the material space.

6 TOPOLOGY OPTIMIZATION

A classic topology optimization problem consists of optimizing the shape and structure of a given object defined by a prescribed domain in order to minimize some cost function. For example, the standard topology optimization minimizes the compliance of the object while satisfying the static equilibrium and the total weight constraint. Since the topology of the object is unknown *a priori*, a method of choice is to define the shape of the object through its material distribution and to locally work with material densities. To this end, the design layout is voxelized and a density variable is assigned to every cell of the discretized domain. By penalizing intermediate values for these densities, a binary distribution corresponding to the object's final layout can be eventually obtained.

In this work, we extend the traditional topology optimization algorithm in multiple ways. First, we do not compute a binary material distribution at the cell level as commonly done. Instead, we leverage our database of microstructures and ask for each cell to be filled with one of the microstructures. By doing so we change the topology of the object at a finer scale, i.e. within each cell. This is done by working with the macro-scale material properties of the microstructures instead of their geometry directly. The second difference is that our algorithm can be used with parametrizations of the material property space that are more complex than the single

density parameter per cell that is commonly used in the standard topology optimization algorithm. Indeed, in our generalized topology optimization problem, each cell c_i contains an n -dimensional material parameter $\mathbf{p}_i \in \mathcal{R}^n$. We use \mathbf{p} to denote the stacked vector of material parameters in all cells. Given a signed distance function $\Phi(\mathbf{p}_i)$ that defines the gamut, our new topology optimization problem is then written as

$$\begin{aligned} \min_{\mathbf{p}} : \quad & S(\mathbf{p}, \mathbf{u}) \\ \text{s.t.} : \quad & \mathcal{F}(\mathbf{p}, \mathbf{u}) = 0 \\ & \Phi(\mathbf{p}_i) \leq 0, \quad 1 \leq i \leq N_c \end{aligned} \quad (10)$$

where S is a real-valued objective function that depends on the material parameters and the displacement vector \mathbf{u} of the entire object at the elasticity equilibrium. The equality constraint $\mathcal{F} = 0$ requires \mathbf{u} to satisfy the elasticity equilibrium and the inequality constraint $\Phi \leq 0$ guarantees that the material properties of each cell stay inside the precomputed gamut.

In our examples, the material parameter \mathbf{p} consists of the density ρ and the elasticity parameters \mathbf{e} . We split our objective function into an elasticity term $C(\mathbf{e}, \mathbf{u})$ that controls the deformation behavior (see Section 6.1) and an optional density term $\mathcal{V}(\rho)$ that controls the overall mass of the object. The density term can be written as

$$\mathcal{V}(\rho) = \left(\sum_{i=1}^{N_c} \rho_i V_i - \hat{M} \right)^2, \quad (11)$$

where V_i is the cell volume and \hat{M} is the target overall mass. When one of the base material is void, the use of the density term allows to modify the topology of the object at a larger scale than the one of the microstructures, and thus to change the external shape of the object. In fact, even for multi-material designs involving base materials with similar mass densities, we noted that we could use the density term to encourage the presence of soft material in the structure. By removing the external cells entirely made of the soft material, we could then decrease the mass of the structure without significantly changing its mechanical behaviour. Alternatively, the density term can also be used to control other quantities related to the ratios of the different materials such as the cost of the object. For specific problems, we can also add spatially-varying weight control terms to Equation 11. For example, we can control the target weight of each individual cell by adding a local term $(\rho_i - \hat{\rho}_i)^2 V_i$.

Assuming static equilibrium, the elasticity constraint is written as

$$\mathcal{F}(\mathbf{e}, \mathbf{u}) = K(\mathbf{e})\mathbf{u} - \mathbf{f}_{\text{ext}} = \mathbf{0}, \quad (12)$$

where \mathbf{f}_{ext} are the external loads applied to the object.

The gamut constraint for a point \mathbf{p}_i in the material property space is described by an n -dimensional level set function $\Phi(\mathbf{p})$. We have $\Phi(\mathbf{p}_i) < 0$ for a point inside the gamut, $\Phi(\mathbf{p}_i) > 0$ for a point outside the gamut, and $\Phi(\mathbf{p}_i) = 0$ for a point on the boundary of the gamut. The value of Φ represents the n -dimension Euclidean distance to the level set boundary. The gradient of Φ are evaluated by a finite difference operation on the signed distance field.

We used a standard gradient-based numerical optimizer (Ipopt [Wächter and Biegler 2006] in our implementation) to solve Equation 10. We enforced the elasticity equilibrium constraint using

the adjoint method. The optimizer only needs to take the function values of S and Φ along with their gradients as input.

6.1 Elasticity Objectives

We used two different types of objective functions for the elasticity term in our topology optimization algorithm. These two types of objectives allowed us to design a wide range of objects.

Target Deformation. Our algorithm takes a vector of nodal target displacements and boundary conditions (external forces, fixed points, etc.) as input. Then, it automatically optimizes the material distribution over the object domain to achieve the desired linear deformation assuming a linear elastic behavior.

We define the deformation objective as

$$C_d(\mathbf{e}, \mathbf{u}) = (\mathbf{u} - \hat{\mathbf{u}})^T \mathbf{D}(\mathbf{u} - \hat{\mathbf{u}}), \quad (13)$$

where $\hat{\mathbf{u}}$ is the vector of the target displacements, \mathbf{D} is a diagonal matrix that determines the importance of each nodal displacement. We use \mathbf{D} to define the subset of nodes that we are interested in. For example, we can set most entries of \mathbf{D} to zero and focus on a portion of the domain (see Figure 13).

Minimum Compliance. We have experimented with the same objective as the one used in the standard topology optimization algorithm where the compliance C_c is defined as

$$C_c(\mathbf{e}, \mathbf{u}) = \mathbf{u}^T \mathbf{K}(\mathbf{e})\mathbf{u}. \quad (14)$$


In the commonly used SIMP algorithm, the stiffness matrix \mathbf{K}_i of each cell i depends on the artificial density value ρ_i through an analytical formula such as $\mathbf{K}_i = \rho_i^3 \mathbf{K}_0$ where \mathbf{K}_0 corresponds to the stiffness matrix of the base material. In contrast, the stiffness matrix in our objective function is directly computed from the material parameters of the material space and forced to correspond to a realizable material thanks to our gamut constraints.

Like in standard algorithms, we regularized the problem to avoid checkerboard solutions by applying a smoothing kernel on the material properties that favors smooth variations of the material parameters over the object layout. Our optimizer supports multiple objectives by linearly combining weighted objective functions.

7 MAPPING MATERIAL PROPERTIES TO MICROSTRUCTURES

After running the topology optimization algorithm, we generate a printable result by replacing each cell in the object lattice by a microstructure whose material properties match the optimal ones.

Material properties of the microstructures are computed using the homogenization theory which is more accurate with a smooth transition between the geometries of neighboring cells. While smoothness in the material parameters can be easily enforced, it does not imply topological similarity of nearby microstructures. For example, any translation of a given microstructure in a periodic tiling will result in a microstructure geometrically different but with






Fig. 6. Gamuts computed with our discrete-continuous sampling scheme for 2D cubic structures (*left*), 2D orthotropic structures (*second from left*), 3D cubic structures (*second from right*) and 3D cubic structures with 0.35 as Poisson's ratio (*right*). The plots show the results for the projection of the gamuts on the plane defined by the macroscale Young's modulus along the x axis (normalized by the Young's modulus of the stiffest base material) and the Poisson's ratio corresponding to a contraction along the y-direction when the material is stretched along the x-direction. The blue dots correspond to the generated samples, the orange dots correspond to the microstructures from Schumacher et al. [2015] and the yellow dots correspond to the microstructures from Panetta et al. [2015].

exactly the same mechanical properties. Fortunately, our database is very dense and multiple microstructures generally map to similar points in the material property space, offering several variants. To further increase the number of possibilities, we also incorporate an additional exemplar for each microstructure by translating it by half its size, which preserves its cubic or orthotropic symmetry without changing its properties. We then run a simple but effective algorithm that picks the microstructure exemplars that minimize the boundary material mismatch across adjacent cells. We quantify this mismatch by $\mathcal{I} = \sum_{i=1}^{N_c} \mathcal{I}_i$, where \mathcal{I}_i is the contribution associated to the cell i and corresponds to the number of boundary voxels filled with materials that are different from the ones of the voxels' immediate neighbours across the interfaces.

Our algorithm proceeds as follows:

- For each cell, we define a list of possible candidates by picking all the microstructures mapping to material points lying in the vicinity of the optimal material point and we randomly initialize the cell with one of the candidates.
- We compute the mismatch energy \mathcal{I}_i associated to each cell i and sort the cells according to their energy.
- We pick the first cell in the sorted list, i.e. the one with the highest energy and assign to it the microstructure candidate that decreases the energy the most. If we cannot decrease the cell energy, we move to the next cell in the list.
- We update the mismatch energies of all the impacted cells and we update the priority list.
- We repeat the last two steps until the mismatch energy \mathcal{I} cannot be decreased anymore.

8 RESULTS

We first analyzed our microstructure sampling algorithm for 2D and 3D microstructure gamuts. Then we used these precomputed gamuts and we designed and optimized a wide variety of objects with our topology optimization algorithm.

8.1 Microstructure Sampling

We evaluated our method on two- and three-dimensional microstructures made of one or two materials. For the 2D case we considered patterns with cubic and orthotropic mechanical behaviors that

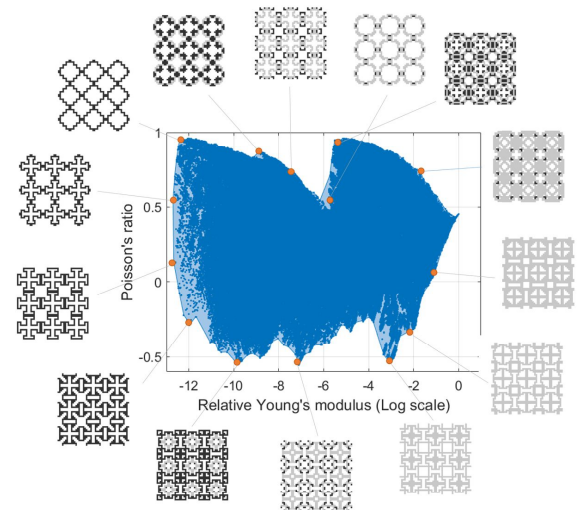


Fig. 7. Gamut corresponding to 2D cubic microstructures made of two materials and void. The Young's modulus of the microstructures is plotted using a logarithmic scale. We show above some examples of microstructures lying near the estimated boundary of the gamut, i.e. with extreme material properties. The dark color corresponds to the softer material, while the light grey color is used for the stiffer material.

can be described with 4 parameters (3 elasticity parameters and density) and 5 parameters (4 elasticity parameters and density) respectively. In 3D we computed the gamut corresponding to cubic structures with 4 parameters. In all cases, the size of the lattice for the microstructures was set to 16 in every dimension. We used isotropic base materials whose Young's modulus differed by a factor of 1000 and having 0.48 as Poisson's ratio. We initially computed the databases for two-material microstructures, but also adapted these databases for microstructures made of a void and a stiff material. In the later case, we replaced the softer material by void, filter out all the microstructures with disconnected components and, in the 3D case, filled the enclosed voids and recomputed the homogenized properties. We provide a comparison between the initial and postprocessed databases in the supplementary material.




Fig. 8. Resulting material property distributions when running our topology optimization algorithm in the orthotropic (*top left*), cubic (*top right*), isotropic (*middle left*) and an analytically defined gamut $E \geq \rho^3 E_0$ (*middle right*), with the material property space dimensions ranging from five to two. We compare our algorithm with the standard SIMP method with power index $p = 1$ (*bottom left*) and $p = 3$ (*bottom right*). For these figures, we computed the color of each cell by mapping every base vector of the normalized parameter space to a color range and linearly interpolating the colors associated to each of the parameters. In this example, the left side of the cantilever is fixed while a discretized, linearly varying, distributed force is applied to the bottom side (see red arrows in the top left picture).

The resulting postprocessed gamuts are also depicted in Figure 6. Our databases contain 274k, 388k and 88k 2D cubic, 2D orthotropic and 3D cubic microstructures respectively and took from 15 hours to 93 hours to compute, which correspond to 68, 19 and 5 sampling cycles, respectively.

We first compared our results to the ones obtained by Schumacher et al. [2015] and observed a significant increase in the coverage of the material space, even for 2D microstructures where we used a coarser discretization. This comforts us with the idea that topology optimization only, while helpful to locally improve the microstructure geometries, is suboptimal when one aims to discover the entire gamut of physical properties. The diversity of the microstructures that we obtained is also much richer, thus providing a larger set of options for the practical use of microstructures. Note that they employed some regularization to avoid thin features. For 16^3 microstructures, we found regularization unnecessary since they are manifold and have a minimal feature size of $1/16$ of the lattice size, which is the same order of magnitude as the thinnest parts of Schumacher’s microstructures. For completeness, we also compared our database of 3D microstructures to the one of Panetta et al. [2015] at 16^3 and 64^3 grid resolutions (Figure 6, *right*). Our initial database was computed with 0.48 as Poisson’s ratio and is shown in the supplementary material. For this comparison, we then recomputed the material properties of the microstructures using the same

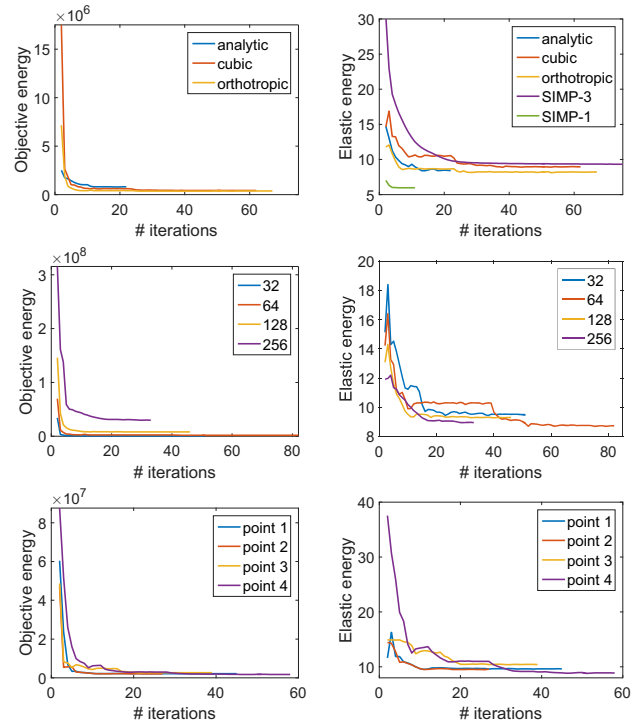


Fig. 9. Convergence tests. Variation of the objective energy (*left*) and the elastic energy (*right*) of a beam being optimized for minimum compliance as the optimization progresses. The convergence plots correspond to the beam of Figure 8 when optimized using different material spaces (*top*), different resolutions for the beam lattice (*middle*) when using cubic microstructures, and different initial material properties for the cubic microstructures (*bottom*).

Poisson’s ratio as Panetta’s, i.e 0.35, which affects the extremal values of the obtained gamut. For the 64^3 microstructures, we used morphological operations in the discrete step and sensitivity filtering with a radius of 3 voxels in the continuous step to limit the minimum feature size to $1/32$ of the lattice size [Sigmund 2007]. Note that this comparison is provided for reference only since our microstructures are cubic while Panetta’s are isotropic (a subset of cubic). Furthermore, they target a different 3D printing technology with self-supporting constraints not imposed here. Finally, we also obtained a dense sampling in the interior of the space, as a result of the randomness inherent to our approach. This reduces the need of running costly optimization in these areas and occurs even if we do not explicitly enforce any sampling there.

We also experimented with three-material 2D cubic microstructures (two solid materials with Young’s moduli differing by a factor of 1000 and with 0.48 as Poisson’s ratio, plus a void material). The resulting database contains about 800k microstructures that can potentially be printed. The corresponding gamut and some examples of the generated microstructures are shown in Figure 7.

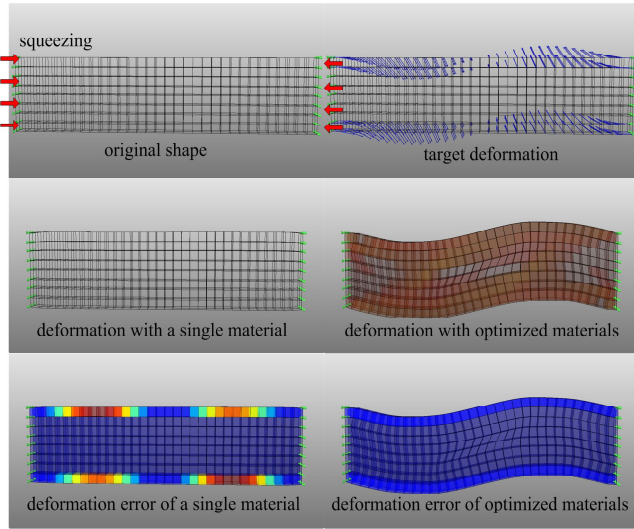


Fig. 10. Optimizing a beam to make it bend when it is squeezed. A beam with optimized material properties can take the desired shape (right) whereas a beam with homogeneous material properties can only axially deform under small deformation (left). In this example, the vertices on the vertical sides are fixed in their target positions, while the other vertices are free to move. Target displacements are set on the nodes of the cells of the two horizontal boundary layers. The color plot for the bottom beams shows the deformation error of each cell as defined by Equation 13.

8.2 Topology Optimization

We tested our topology optimization algorithm on a number of simple test cases and large scale examples. Detailed analysis and discussion of the results is provided below.

Impact of the Material Space. We evaluated the impact of the chosen material space on a 2D cantilever beam with optimized minimum compliance. We tested our topology optimization algorithm on isotropic, cubic and orthotropic gamuts as well as with the virtual materials used in the traditional SIMP approach and for which the stiffness of the material $E = \rho^p E_0$, $p \geq 1$ is a function of the density ρ of the cell and the stiffness of the base material E_0 . We also tested our algorithm on an analytical gamut with allowed stiffnesses E defined by $E \leq \rho^3 E_0$. The results are shown in Figure 8. It can be noted that, as the dimension of the material space increases, the final energy of the system decreases. This is to be expected since higher dimensional space means larger gamuts. Thus, when using cubic materials, the minimum compliance objective function reaches 3% lower energy than the standard SIMP method with power index 3. This difference reaches 11% when we use orthotropic materials. It is worth noting that the lowest elastic energy is achieved when we use the traditional SIMP method with $p = 1$ (as shown in Figure 9). However, this solution does not correspond to a realizable structure since some of the optimized materials do not correspond to any microstructure.

Matching Quality. We evaluated for different examples the matching quality of the target deformation optimization. For the first test, we forced a beam to take an ‘S’ shape when undergoing tensile

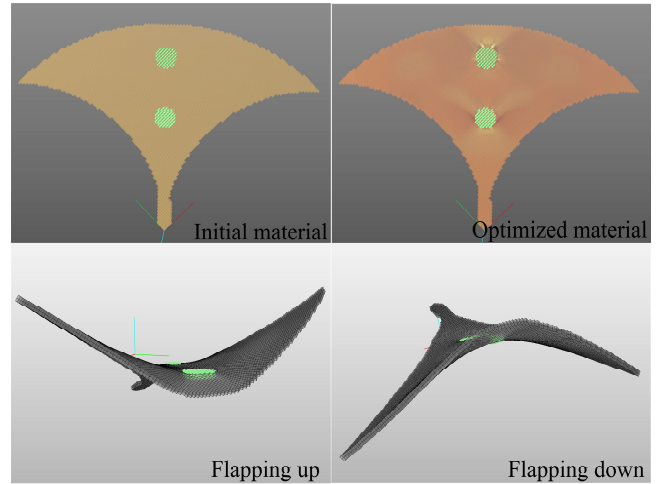


Fig. 11. Designing a soft ray. The wings of the ray are asked to flap up and down when vertices on its spine contract and expand. Constrained vertices are colored in green. The deformations achieved with the optimized materials are displayed on the bottom row.

forces (Figure 10). In order to avoid overfitting, we applied target displacements on the vertices of the boundary cells only. As depicted in the figure, the use of microstructures largely improves the global shape of the beam, which closely matches the target deformed shape. This becomes even more striking when compared to the behavior of a beam made of a homogeneous material.

We also validated our algorithm by designing a soft ray whose wings can flap using a compliant mechanism (see Figure 11 and accompanying video). Boundary conditions are applied on two circular areas located along the spine of the ray. Each disk has one degree of freedom for deformation, namely contracting or expanding along the disk normals. This mechanism resembles the one of many hydraulics-driven soft robots. We define two target deformation objectives corresponding to the flapping of the wings up and down, when alternatively contracting and expanding the two disks’ boundaries. By running our multi-objective topology optimization framework, we can compute an optimized material design that can achieve both deformation modes when the corresponding boundary conditions are exerted.

Convergence and Robustness. We evaluated the convergence rate of our topology optimization both on the minimum compliance problem and with the target deformation objective. For the minimum compliance problem, we used the same loading as the one of Figure 8. The corresponding results are shown in Figure 9 where we plot both the deformation energy of the structure as defined in Equation 14 and the original objective of the problem 10 that also includes the volume term defined by Equation 11. For all these examples, the algorithm converged after a couple of dozen iterations, irrespectively of the lattice resolution, i.e. the number of variables and the number of non-linear constraints. This demonstrates the scalability of our algorithm. We also tested the robustness of our algorithm by starting with different initial conditions. In this case, we initialized the material parameters of each cell with a random

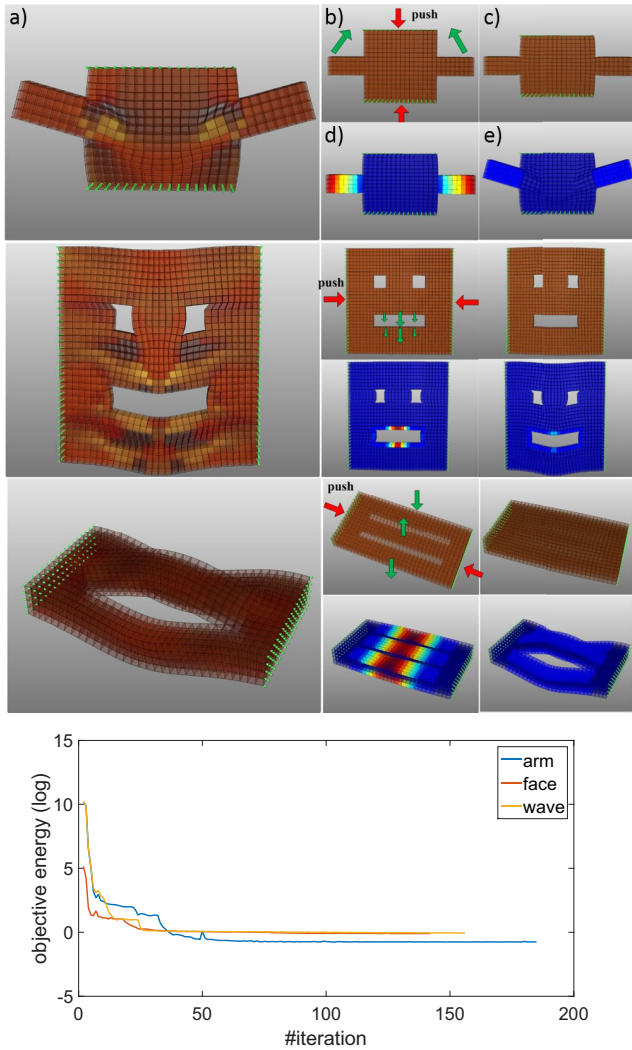


Fig. 12. Target deformation examples. The pictures on the left (a) show the deformed optimized result when the boundary conditions illustrated by the pictures (b) are applied. The orange meshes (c) correspond to the simulated deformations when a homogeneous material is used. The blue-to-red meshes indicate the relative deformation error of the unoptimized (d) and optimized (e) structures. Convergence rates corresponding to the optimization of these three examples are reported on the bottom figure.

material point projected onto the boundary of the gamut. Similar to other topology optimization schemes, we have no guarantee that we reach the global minimum of the function, and indeed, our algorithm sometimes converges to different solutions. However we note that these different solutions have a similar final objective value and are therefore equally good.

For the evaluation of the target deformation optimization, we ran our topology optimization algorithm on scenarios similar to the ones from Panetta et al. [2015] where different extruded structures are asked to deform into prescribed shapes when being compressed between two plates (see Figure 12). As shown in the figure, our

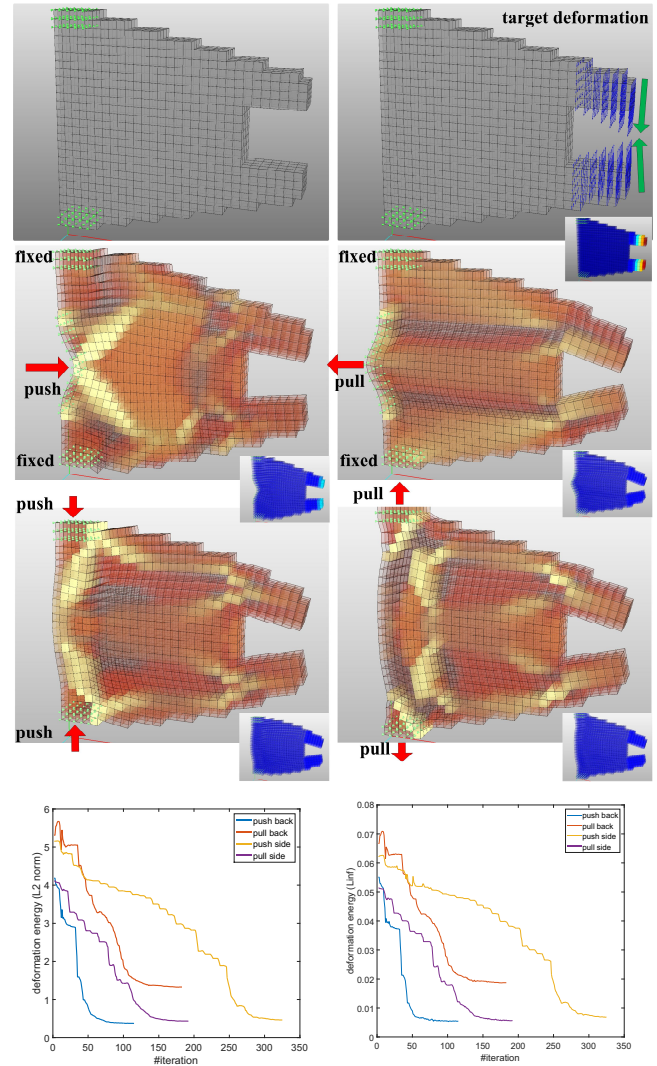


Fig. 13. Designing functional grippers. The top row shows the initial shape of the gripper (left), and the target deformation for the tip (right). The green dots correspond to the fixed vertices while the blue arrows correspond to the target displacements. The two middle rows correspond to the optimized results obtained for the specified boundary conditions. The inset pictures color-code the initial and final deformation error for the different examples. The convergence plots in the bottom row depict the change in the sum of the deformation errors corresponding to all the cells (left) and the value of the maximal cell error contribution (right) as the optimization progresses.

optimization algorithm successfully converges to the specified deformation behavior in less than 100 iterations for all the examples. We also tested the convergence rate when optimizing for functional mechanisms. To this end, we designed several grippers that can grasp objects by moving their tips when external forces are applied to their extremities. We experimented with four sets of boundary conditions, namely, pulling and pushing the back of the gripper horizontally, and compressing and stretching the extremities of the




Fig. 14. Comparison of simulated beams with homogenized material properties (*inset pictures*) to the ones using full microstructures (*large pictures*).

gripper vertically. As shown in Figure 13, these different settings lead to different material structures. The deformation errors of all the four designs converge to a low level after a couple of hundreds of iterations.




Fig. 15. Comparison of simulated cubes with different material patterns modeled by homogenized cells (*inset pictures*) and full resolution microstructures (*large pictures*).

Accuracy. We evaluated the accuracy of our algorithm on several optimized structures by comparing the deformation obtained when using the optimized homogenized material properties for each cell to the one obtained by a high resolution simulation in which every cell is replaced by its mapped microstructure. We first evaluated the accuracy on a deformable bar, one side of which was rigidly attached while the other was subject to different sets of external conditions (see Figure 14). We used a $8 \times 2 \times 2$ lattice to represent the bar with homogenized cells, which translates into a $128 \times 32 \times 32$ grid for the full resolution mesh. Similar stretching, bending and shearing behaviors were obtained for both sets of models. From a quantitative point of view, the differences amount to 5-10% in terms of average vertex displacement and 9%-33% in terms of elastic deformation energy (see Table 1). We further evaluated the effects of material patterns by running a similar comparison on a cube made of periodic layers of similar microstructures and with random assignments of microstructures (Figure 15). As reported in Table 1,




Fig. 16. Comparison of a simulated gripper with homogenized material properties (*inset picture, left*) to the one using full microstructures (*main picture, left*). The figures on the right show the vector field of vertex displacement of the two models. The blue-to-red colors represent the magnitudes of the displacements.

we show that the ratio between the magnitudes of the average vertex displacement differences is between 4% and 7%, and the elastic energy difference is between 10% and 19%. Finally, we also compared the behaviors of one of the grippers (Figure 16). The original optimized gripper is made of 3k elements while the high resolution version is made of 4M voxels. Overall, the two models exhibit similar global deformation behaviors, in particular in the tip area. Some differences can be observed on the left side of the gripper for which the high-resolution model exhibits a lower effective material stiffness than its homogenized counterpart. With the same displacement boundary conditions applied, the high-resolution model deforms about 25% more than the homogenized model.




Fig. 17. Simulation of a cube made of a periodic arrangement of a single microstructure at different resolutions. Inset pictures correspond to the model with homogenized material properties, while main pictures correspond to the full resolution simulations.

The differences observed between the homogenized model behaviour and the full resolution simulation can be explained by two major factors: (i) numeral stiffness when using larger elements which tends to make the homogenized mesh slightly stiffer in particular when bending deformation arises, (ii) violation of the periodicity assumption when replacing each cell by a single microstructure. This issue can be reduced by replacing each cell by a *tiling* of microstructures. This was verified on a cube made of a periodic arrangement of a single microstructure (see Figure 17). And indeed, as we increase the resolution of the simulation grid, the error between the homogenized model and the full resolution version decreases and converges to similar values.

Orthotropic materials. We tested the behavior of our algorithm in a 5-dimensional space by using the gamut of 2D orthotropic microstructures depicted in Figure 6 (*middle*). To this end, we used

Table 1. Error statistics (SI units). The size of one microstructure is set to $1 \times 1 \times 1$.

Example	Mean displacement	Mean displacement difference	Elastic energy homogenized/full resolution	
Beam 1	6.47×10^{-3}	6.04×10^{-4}	6.85×10^{-5}	6.17×10^{-5}
Beam 2	6.47×10^{-3}	6.04×10^{-4}	1.63×10^{-5}	1.08×10^{-5}
Beam 3	5.07×10^{-3}	4.97×10^{-4}	2.38×10^{-4}	2.07×10^{-4}
Beam 4	8.78×10^{-3}	4.45×10^{-4}	3.33×10^{-4}	2.30×10^{-4}
Cube 1	3.62×10^{-3}	2.86×10^{-4}	3.64×10^{-3}	3.20×10^{-3}
Cube 2	4.35×10^{-3}	1.94×10^{-4}	6.82×10^{-3}	5.94×10^{-3}
Cube 3	5.42×10^{-3}	4.22×10^{-4}	7.81×10^{-3}	6.32×10^{-3}
Gripper	1.32×10^{-2}	6.90×10^{-3}	8.67×10^{-3}	5.70×10^{-3}
Cube 4	5.22×10^{-3}	4.89×10^{-4}	2.08×10^{-2}	1.63×10^{-2}
Cube 5	5.21×10^{-3}	2.17×10^{-4}	1.89×10^{-2}	1.63×10^{-2}
Cube 6	5.21×10^{-3}	1.32×10^{-4}	1.82×10^{-2}	1.63×10^{-2}

a regular lattice whose vertices on the left side were fixed and we applied parallel forces on the vertices of the opposite side. The goal in the test was to minimize the compliance of the structure. As can be seen in Figure 18 and in the accompanying video, we experimented with different force directions. Unsurprisingly, when a single cell is considered, the microstructure that we obtain has a structure that is aligned with the direction of the forces (see Figure 18, *left*). For a higher resolution lattice this is no longer true and the resulting overall structure becomes less intuitive (see Figure 18, *right*). Note that the resulting material distribution varies smoothly. By considering various alternatives for each material point, our tiling algorithm is able to map the material properties to microstructures which are well connected.

8.3 3D-Printed Designs

Leveraging our two-scale approach, we used our topology optimization algorithm to generate a wide variety of high resolution models that we 3D-printed. We used a Stratasys Objet Connex 500 and the two base materials *Vero Clear* and *Tango Black Plus* and used the database containing the three-dimensional cubic microstructures. The sizes and computation times of the resulting models are outlined in Table 2.

Since Ipopt performs a line search at each gradient step, one single step may correspond to multiple simulations. We show the average time required for taking a step in the last column of Table 2. For these large scale examples, Ipopt takes two hundred iterations in average to find a local minimum. Since our problem is formulated as a very general constrained continuous optimization, it is independent of the optimization package that is used and its speed could potentially be further improved by using alternative minimizers. We found Ipopt to be a good choice for its capability to efficiently handle a large number of inequality constraints, which is not the case of other popular minimizers used in topology optimization such as the method of moving asymptotes (MMA).

Our algorithm is mainly directed towards engineering applications and targets the design of objects undergoing small deformations. In the following examples, we sometimes intentionally exaggerated the target displacements (and scaled the external forces accordingly) for better visualization, which does not change the output of the algorithm with a linear material model.

Table 2. Statistics on the 3D-printed models. The last row uses the database of 64^3 microstructures.

Example	Grid Size	# Voxels	Time per FEM Solve [s]	Time per Step [s]
Beam	$96 \times 24 \times 4$	38M	0.7	5
Flexure	$32 \times 32 \times 16$	67M	1	12
Gripper	$64 \times 32 \times 8$	67M	1.7	10
Bunny	$32 \times 32 \times 32$	134M	0.6	4
Bridge	$128 \times 64 \times 32$	1074M	27	81
Bridge 2	$320 \times 160 \times 80$	1074G	1.3k	-

Beams with controlled deformation behaviour. We started by designing a 3D hollowed beam with a desired deformed shape. The beam was stretched by moving vertices on two opposite sides. Our topology optimization algorithm was run using a target deformation objective. The resulting optimized material properties and the 3D-printed structure are depicted in Figure 19.

Multi-Objective Flexure Design. We tested our algorithm on a multi-target deformation setting by optimizing the structure of a flexure mount with two different target shapes (see Figure 20). Here, our goal is to design a flexure that resists vertical loads while remaining compliant to horizontal loads. We assume that the object mounted on the flexure is connected to the flexure using a cylindrical connector that transmits the forces to the flexure via the connecting area. In the first scenario, vertical forces are applied to the points of the cylindrical area and we ask the flexure to stay as close as possible to its rest configuration. In the second scenario, horizontal forces are applied to the points of the cylinder and we ask the flexure shape to match the shape shown in the Figure 20.

Gripper. We verified the functionality of our grippers by fabricating two of them. For these results we ran the optimization on high resolution meshes of the version that grasps the object when the extremities of the gripper are pressed (see Figure 21). By changing the parameter controlling the ratio of the soft material, different designs based on different mechanisms can be achieved. When more soft material is used the gripper achieves its target deformation thanks to out-of-plane bending, while for stiffer designs, the grasping motion is achieved via in-plane deformation.

Minimal Compliance Examples. To demonstrate the scalability of our algorithm, we ran our topology optimization algorithm on a Stanford bunny (Figure 22) made of more than 100 million voxels and subject to two load case scenarios. We also designed two bridges of increasing resolutions. The first bridge was optimized using a lattice of half a million cells which corresponds to 1 billion voxels (Figure 23). For the second bridge, we used the database of 64^3 microstructures and a layout made of 4 million cells, which amounts to 1 trillion voxels. We initialized the topology optimization by running the algorithm on a lower resolution grid with 1.4 million elements and used the resulting parameters as initial material parameter values for the higher resolution optimization.




Fig. 18. Optimizing the orthotropic material parameters of a single cell (*left*) and a 32×32 lattice of cells (*right*) subject to directional forces. The vertices on the left side of the layout are fixed while forces are applied on the right vertices as depicted by the red arrows. Our simple but effective tiling algorithm allows to nicely transition between microstructures of smoothly material properties (*right, top*).




Fig. 19. An optimized hollow beam with target deformation. The left figure shows the target deformation and optimized material distribution. The right figure shows the 3D-printed structure and the achieved deformation.




Fig. 20. Optimizing a flexure mount. The flexure is connected to an object thanks to a cylindrical connector. We leave space for this connector by keeping a cylindrical area of the design layout empty of material. The material distribution of the flexure is optimized for two sets of external forces applied to the cylindrical area. Under vertical load, the flexure should stay close to the rest shape while under horizontal load, the flexure should deform according to the inset figure.

9 CONCLUSION

We have presented a computational framework for two-scale topology optimization. Our approach can efficiently optimize high resolution models that can be fabricated using multi-material 3D printing. Our first insight is to use a precomputation process to efficiently




Fig. 21. 3D-printed functional grippers. By setting different target ratios of the rigid material, different designs can be obtained. When more soft material is used the grasping behaviour of the gripper is obtained via out-of-plane bending (*top*), whereas more rigid material is used, the gripper deformation remains planar (*bottom*).

sample the space of microstructures and their corresponding material properties in order to define a continuous material property gamut. Our second insight is to use this gamut as a constraint in a generalized topology optimization framework to assign spatially-varying material properties throughout the optimized object. Finally, the volume with assigned material properties can be converted to a 3D model with corresponding spatially-varying microstructures.




Fig. 22. Stanford bunny optimized for two loading cases. Two sets of external forces are applied to the back and chest of the bunny as indicated by the arrows, and the color indicates the material distribution (left). Material parameters are mapped to microstructures to obtain an object that can be actually printed (right).

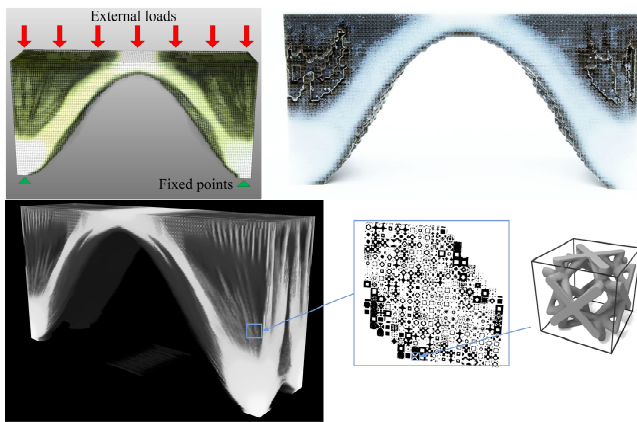


Fig. 23. Optimizing a bridge. The initial layout corresponds to a $128 \times 64 \times 32$ regular grid. We apply uniform loads on the upper plane deck. We compute the material parameters and set cells with extremely low stiffness to void (top left). We look up the microstructures and 3D print the bridge (top right). We scale the problem to 1 trillion voxels by using a lattice of 4 million elements where each element corresponds to a 64^3 microstructure (bottom left). We show a $20 \times 20 \times 1$ patch on the bridge with filled microstructures and a single microstructure with 64^3 voxels on the patch (bottom right).

We demonstrated the effectiveness of our approach on multiple examples and showed improvements over traditional binary topology optimization schemes.

Limitations and Future Work. First, while our sampling method outperforms current approaches in terms of the material space coverage and the approach converges to stable gamuts, we do not provide any theoretical guarantees that the gamut space cannot be further expanded. Second, we would like to investigate microstructures with additional properties, e.g., electrical or magnetic properties, their combined property gamuts, and the different applications they enable. Finally, our framework builds upon linear elasticity and optimizes the material distribution of objects subject to small deformations only. While this is enough for many engineering applications, extending our algorithm to the nonlinear regime would be an interesting direction for future work.

ACKNOWLEDGEMENTS

This research is supported in part by the Defense Advanced Research Projects Agency (DARPA) and Space and Naval Warfare Systems Center Pacific (SSC Pacific) under Contract No 66001-15-C-4030.

REFERENCES

- Joe Alexandersen and Boyan S. Lazarov. 2015. Topology optimisation of manufacturable microstructural details without length scale separation using a spectral coarse basis preconditioner. *Computer Methods in Applied Mechanics & Engineering* 290 (2015), 156–182.
- Grégoire Allaire. 2012. *Shape optimization by the homogenization method*. Vol. 146. Springer Science & Business Media.
- G. Allaire and R.V. Kohn. 1993. Optimal bounds on the effective behavior of a mixture of two well-ordered elastic materials. 51, 4 (1993), 643–674.
- Ryoichi Ando, Nils Thürey, and Chris Wojtan. 2013. Highly adaptive liquid simulations on tetrahedral meshes. *ACM Trans. Graph. (TOG)* 32, 4 (2013), 103:1–103:10.
- Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. 2014. Design of manufacturable 3D extremal elastic microstructure. *Mechanics of Materials* 69, 1 (2014), 1–10.
- Sahab Babaee, Jongmin Shim, James C. Weaver, Elizabeth R. Chen, Nikita Patel, and Katia Bertoldi. 2013. 3D Soft Metamaterials with Negative Poisson's Ratio. *Advanced Materials* 25, 36 (2013), 5044–5049.
- Martin P Bendsøe. 1989. Optimal shape design as a material distribution problem. *Structural optimization* 1, 4 (1989), 193–202.
- Martin P Bendsøe and Ole Sigmund. 1999. Material interpolation schemes in topology optimization. *Archive of applied mechanics* 69, 9–10 (1999), 635–654.
- Martin Philip Bendsøe and Ole Sigmund. 2004. *Topology Optimization: Theory, Methods, and Applications*. Springer Science & Business Media.
- Haimasree Bhattacharya, Yue Gao, and Adam Bargteil. 2011. A level-set method for skinning animated particle data. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. ACM, 17–24.
- Bernd Bickel, Moritz Bäcker, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010), 63:1–63:10.
- J. Bonet and R. D. Wood. 1997. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press.
- Joseph E Cadman, Shiwei Zhou, Yuhang Chen, and Qing Li. 2013. On design of multi-functional microstructural materials. *Journal of Materials Science* 48, 1 (2013), 51–66.
- Vivien J Challis and James K Guest. 2009. Level set topology optimization of fluids in Stokes flow. *International journal for numerical methods in engineering* 79, 10 (2009), 1284–1308.
- Desai Chen, David I. W. Levin, Piotr Didyk, Pitchaya Sitthi-Amorn, and Wojciech Matusik. 2013. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 135:1–135:10.
- Xiang Chen, Changxi Zheng, Weiwei Xu, and Kun Zhou. 2014. An Asymptotic Numerical Method for Inverse Elastic Shape Design. *ACM Trans. Graph. (Proc. SIGGRAPH)* 33, 4 (2014), 95:1–95:11.
- PG Coelho, PR Fernandes, JM Guedes, and HC Rodrigues. 2008. A hierarchical model for concurrent material and topology optimisation of three-dimensional structures. *Structural and Multidisciplinary Optimization* 35, 2 (2008), 107–115.
- Christian Dick, Joachim Georgii, and Rüdiger Westermann. 2011. A real-time multigrid finite hexahedra method for elasticity simulation using CUDA. *Simulation Modelling Practice and Theory* 19, 2 (2011), 801–816.
- Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. 2010. Fabricating Spatially-Varying Subsurface Scattering. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010), 62:1–62:10.
- Randal Douc. 2005. Comparison of resampling schemes for particle filtering. In *4th International Symposium on Image and Signal Processing and Analysis (ISPA)*. 64–69.
- RB Haber, P Pedersen, and JE Taylor. 1994. An analytical model to predict optimal material properties in the context of optimal structural design. *Urbana* 51 (1994), 61801.
- Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. 2010. Physical Reproduction of Materials with Specified Subsurface Scattering. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010), 61:1–61:10.
- Steven G Johnson. 2014. The Nlopt nonlinear-optimization package. (2014). <http://ab-initio.mit.edu/nlopt>
- Lily Kharevych, Patrick Mullen, Houman Owghadi, and Mathieu Desbrun. 2009. Numerical Coarsening of Inhomogeneous Elastic Materials. *ACM Trans. Graph. (Proc. SIGGRAPH)* 28, 3 (2009), 51:1–51:8.
- Roderic Lakes. 1987. Foam structures with a negative Poisson's ratio. *Science* 235, 4792 (1987), 1038–1040.
- Robert Lipton. 1994. Optimal Bounds on Effective Elastic Tensors for Orthotropic Composites. *Proceedings of the Royal Society A* 444, 1921 (1994), 399–410.

- Jonàs Martínez, Jérémie Dumas, Sylvain Lefebvre, and Li-Yi Wei. 2015. Structure and Appearance Optimization for Controllable Shape Design. *ACM Trans. Graph.* 34, 6 (2015), 229:1–229:11.
- Graeme W. Milton and Andrej V. Cherkaev. 1995. Which elasticity tensors are realizable? *Journal of Engineering Materials and Technology* 117 (1995), 483.
- PB Nakshatrala, DA Tortorelli, and KB Nakshatrala. 2013. Nonlinear structural design using multiscale topology optimization. *Computer Methods in Applied Mechanics and Engineering* 261 (2013), 167–176.
- Stanley Osher and Ronald Fedkiw. 2006. *Level set methods and dynamic implicit surfaces*. Vol. 153. Springer Science & Business Media.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4 (2015), 135:1–135:12.
- U. T. Ringertz. 1993. On finding the optimal distribution of material properties. *Structural & Multidisciplinary Optimization* 5, 4 (1993), 265–267.
- Daniel Ritchie, Ben Mildenhall, Noah D. Goodman, and Pat Hanrahan. 2015. Controlling Procedural Modeling Programs with Stochastically-ordered Sequential Monte Carlo. *ACM Trans. Graph. (Proc. SIGGRAPH)* 34, 4 (2015), 105:1–105:11.
- H Rodrigues, Jose M Guedes, and MP Bendsøe. 2002. Hierarchical optimization of material and structure. *Structural and Multidisciplinary Optimization* 24, 1 (2002), 1–10.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Trans. Graph.* 34, 4 (2015), 136.
- Eftychios Sifakis and Jernej Barbic. 2012. FEM Simulation of 3D Deformable Solids: A Practitioner’s Guide to Theory, Discretization and Model Reduction. In *ACM SIGGRAPH 2012 Courses*.
- Ole Sigmund. 1997. On the Design of Compliant Mechanisms Using Topology Optimization. *Mechanics of Structures and Machines* 25, 4 (1997), 493–524.
- Ole Sigmund. 2007. Morphology-based black and white filters for topology optimization. *Structural and Multidisciplinary Optimization* 33, 4 (2007), 401–424.
- Ole Sigmund and Kurt Maute. 2013. Topology optimization approaches. *Structural and Multidisciplinary Optimization* 48, 6 (2013), 1031–1055.
- Ole Sigmund and Salvatore Torquato. 1996. Composites with extremal thermal expansion coefficients. *Applied Physics Letters* 69, 21 (1996), 3203–3205.
- Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. 2013. Computational Design of Actuated Deformable Characters. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 82:1–82:10.
- Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. 2012. Stress relief: improving structural strength of 3D printable objects. *ACM Trans. Graph. (Proc. SIGGRAPH)* 31, 4 (2012), 48:1–48:11.
- Krister Svanberg. 1987. The method of moving asymptotes—a new method for structural optimization. *International journal for numerical methods in engineering* 24, 2 (1987), 359–373.
- T. C. T. Ting and Tungyang Chen. 2005. Poisson’s ratio for anisotropic elastic materials can have no bounds. *Quarterly Journal of Mechanics & Applied Mathematics* 58, 1 (2005), 73–82.
- Kiril Vidimčec, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik. 2013. OpenFab: A Programmable Pipeline for Multi-material Fabrication. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 136:1–136:12.
- A. Wächter and L. T. Biegler. 2006. On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming. *Mathematical Programming* 106, 1 (2006), 25–57.
- F. Wang, O. Sigmund, and J.S. Jensen. 2014. Design of materials with prescribed nonlinear properties. *Journal of the Mechanics and Physics of Solids* 69 (2014), 156–174.
- Jun Wu, Christian Dick, and Rüdiger Westermann. 2016. A System for High-Resolution Topology Optimization. *IEEE Trans. on Visualization and Computer Graphics* 22, 3 (2016), 1195–1208.
- Liang Xia and Piotr Breitkopf. 2014. Concurrent topology optimization design of material and structure within nonlinear multiscale analysis framework. *Computer Methods in Applied Mechanics and Engineering* 278 (2014), 524–542.
- Liang Xia and Piotr Breitkopf. 2015a. Design of materials using topology optimization and energy-based homogenization approach in Matlab. *Structural and Multidisciplinary Optimization* 52, 6 (2015), 1229–1241.
- Liang Xia and Piotr Breitkopf. 2015b. Multiscale structural topology optimization with an approximate constitutive model for local material microstructure. *Computer Methods in Applied Mechanics and Engineering* 286 (2015), 147–167.
- Hongyi Xu, Yijing Li, Yong Chen, and Jernej Barbivc. 2015. Interactive material design using model reduction. *ACM Trans. Graph.* 34, 2 (2015), 18:1–18:14.
- Xiaolei Yan, Xiaodong Huang, Guangyong Sun, and Yi Min Xie. 2015. Two-scale optimal design of structures with thermal insulation materials. *Composite Structures* 120 (2015), 358–365.
- X Yan, X Huang, Y Zha, and YM Xie. 2014. Concurrent topology optimization of structures and their composite microstructures. *Computers & Structures* 133 (2014), 103–110.
- Qingnan Zhou, Julian Panetta, and Denis Zorin. 2013. Worst-case Structural Analysis. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 137:1–137:12.
- Yongning Zhu and Robert Bridson. 2005. Animating sand as a fluid. *ACM Trans. Graph.* 24, 3 (2005), 965–972.