Designing Volumetric Truss Structures for Computational Fabrication

Rahul Arora¹*

Alec Jacobson^{1†} Time ¹University of Toronto

Timothy R. Langlois^{2‡}

Karan Singh^{1§} ²Adobe Systems David I.W. Levin¹¶



Figure 1: We generate stress-aligned 3D trusses to fabricate objects which are structurally-sound and lightweight. These structures consist of families of smooth, continuous, curves tracing stress lines, and can be fabricated with multiple manufacturing techniques.

ABSTRACT

We present the first algorithm for designing volumetric Michell Trusses. Our method uses a parametrization approach to generate trusses made of structural elements aligned with the primary direction of an objects stress field. Such trusses exhibit high strength-toweight ratio while also being aesthetically pleasing. Unlike traditional approaches to structural optimization, our method produces trusses that can be edited as a post process but retain structural optimality. We also demonstrate the structural robustness of our designs via mechanical testing. Our algorithm permits an exciting combination of control and structural soundness which we believe serves as an important compliment to existing structural optimization tools and as a novel standalone design tool itself.

Index Terms: General and reference—Design; Computing methodologies—Physical simulation

1 INTRODUCTION

A primary objective of engineering is to develop the stiffest possible structure by using the least amount of material. The design of many structures in our everyday life such as bridges, bikes, and buildings such as stadia follow this principle. These structures often form *trusses* that also manifest an aesthetic appeal, and are therefore of interest to graphics and computational design communities (see [3]).

Automatically designing such minimal structures is challenging since material must be positioned optimally to retain structural soundness. Existing approaches for automatically designing lightweight structures typically formulate the problem as a strength maximization problem, subject to some constraint on the amount of material used in the structure. The two main frameworks for optimal material placement are Topology Optimization [1], which uses a voxel-based representation to explore the design space, and the Ground Structure Method [2] which works on a truss discretization.

Both of these methods have inherent limitations, rooted in the requirement of an overprescribed set of design variables (either voxels or bars) as initialization. At a high level, both the strategies involve optimizing for a sparse subset of the initial layout. Both methods are difficult to control, and efficiently editing the output after optimization is still an open problem.

In this paper, we take a different approach to generate aesthetically pleasing, light, and strong structures. Instead of starting with an overprescribed solution and sparsifying it, we use one-dimensional cylindrical structural elements to define a truss, and formulate its design as a fitting problem for these elements. Michell [4] laid the foundations for creating such trusses by proving that for a given material budget, all elements of the optimal (stiffest) truss must follow paths of maximum strain. Structures which fulfill this property are called *Michell Trusses*. Hence, by aligning the individual elements with the principal stress directions of an object's stress tensor field, a structurally sound design can be created without needing to fill the entire shape volume with material and later sparsifying it.

We present an algorithm to design Michell trusses inside arbitrary 3D domains. Rather than optimizing an initial guess, we treat truss optimization as a fitting problem, in the vein of recent hex-meshing approaches. Our method requires only a single solve of the static equilibrium equations to compute a continuous stress field. We then use a novel parametrization method to produce a graph of a prescribed resolution where each graph edge is as aligned as possible with the underlying stress tensor field. Our method avoids many of the difficulties of previous methods, its initialization is trivial, and requires no regularization terms to avoid high-frequency artifacts.

Importantly, our approach generates object-spanning material curves that are consistently labelled. This allows us to easily edit our Michell Truss after creation, laying the groundwork for controllable structural optimization. Implementing such customization with existing methods is a grueling task.

2 STRESS-ALIGNED TRUSS NETWORK GENERATION

The input to our method is a tetrahedral mesh, with Dirichlet and Neumann boundary conditions specifying fixed vertices and forces, respectively. We begin by performing standard linear finite element analysis to compute the Cauchy stress tensor field on the mesh. The crux of our algorithm lies in generating a stress-aligned 3D texture parametrization on the tet mesh. Integer isolines of the parametrization can then be extracted as a labelled graph embedded in 3D. Finally, graph elements are inflated to cylinders to get a truss structure.

2.1 Intermediate Frame Fields

Naïvely, a Cauchy Stress tensor field $\sigma(\mathbf{x})$ can be interpreted as a frame field by representing each tensor by its three eigenvectors.

^{*}e-mail: arorar@cs.toronto.edu

[†]e-mail: jacobson@cs.toronto.edu

[‡]e-mail: tlangloi@adobe.com

[§]e-mail: karan@dgp.toronto.edu

[¶]e-mail: diwlevin@cs.toronto.edu

Because each $\sigma(\mathbf{x})$ is a Hermitian matrix, its eigenvectors are guaranteed to form an orthogonal basis. However, such a frame field is almost certain to be non-smooth as the direction of each eigenvector can be arbitrarily flipped or interchanged. Our key observation is that the tensor itself is a useful, symmetry agnostic frame field representation and we leverage this notion to fit a frame field.

We define a "good" fit between a frame and a stress tensor as one where the first axis of the frame is aligned with the primary eigenvector of the stress tensor and the other two axes are aligned with the second and third eigenvectors (though it does not matter which aligns with which). Let the stress tensor of the i^{th} tet be eigendecomposed as $\sigma^i = Q^i \Lambda^i (Q^i)^{-1}$. We define the following frame-tensor matching function:

$$E_{data}^{i}(\mathbf{r}_{2},\mathbf{r}_{3}) = \|\mathbf{r}_{2}^{T}S^{i}\mathbf{r}_{2}\|_{F} + \|\mathbf{r}_{3}^{T}S^{i}\mathbf{r}_{3}\|_{F},$$
(1)

where $S^i = Q^i (\Lambda^i)^{1/2}$ and $\mathbf{r}_j \in \mathbb{R}^3$ is the *j*th direction vector of our frame. This cost function has a set of identical minima at every frame alignment which satisfies our criteria.

Next we need a method for disambiguating the local minima in Equation 1. Typically this is done combinatorially, but here we follow the approach Solomon et al. [5] and instead use a smoothness energy to produce a well-fitted, consistently aligned frame field. While we borrow their Laplacian smoothing term, we avoid their frame field representation as it requires an extra projection step. Instead we represent a frame at the centroid of a tetrahedron using rotation matrices, parameterized via the matrix exponential, $R^i = \exp(\sum_{j=1}^4 [\omega^j]) \in \mathbb{R}^{3\times 3}$. Here, $\omega^j \in \mathbb{R}^3$ are angular velocity vectors at the vertices of each tetrahedron and the $[\cdot]$ operator computes the cross product matrix. Combining the data term with Laplacian smoothing gives us the weighted optimization

$$\boldsymbol{\omega}^* = \underset{\boldsymbol{\omega}}{\operatorname{argmin}} \sum_{i=1}^{N} E_{data}^i(\boldsymbol{\omega}) + \alpha \frac{1}{2} \boldsymbol{\omega}^T L \boldsymbol{\omega}$$
(2)

where *N* is the number of tetrahedra in our finite element mesh and α is a scalar weight. Here, we make an observation that the only purpose of the smoothness energy is to help us choose an appropriate local minima to descend into. Therefore, our final fitting algorithm is Augmented Lagrangian-esque in that we repeatedly minimize Eq. 2 with increasingly smaller α until the cost stops decreasing.

2.2 Parametrization Computation

We use our smooth, data-aligned frame field to compute a stressaligned parametrization from which we will create our Michell Truss. We define $\Omega \in \mathbb{R}^3$ as the world space that our object occupies and $\mathbf{u} \in \mathbb{R}^3$ as a volumetric texture domain. We chose our structural members to lie along the coordinate lines of \mathbf{u} and seek to find a parametrization $\mathbf{u} = \phi(\mathbf{x}) : \Omega \to \mathbb{R}^3$ that aligns these coordinate lines with our frame field. Formally we seek a $\phi(\mathbf{x})$ such that

$$\frac{\partial \phi}{\partial \mathbf{x}} \mathbf{r}_{\mathbf{i}} = \mathbf{e}_i \quad \forall i \in \{1, 2, 3\}$$
(3)

at the center of each tetrahedron in our mesh. Here, \mathbf{e}_i is the column vector representing the *i*th standard basis vector of \mathbb{R}^3 .

This can be restated as a linear system of equations by constructing the discrete directional gradient operator for each tet:

$$G^{i}(\mathbf{v}) = \left[\mathbf{v}_{x}^{i} * G_{x}^{i} + \mathbf{v}_{y}^{i} * G_{y}^{i} + \mathbf{v}_{z}^{i} * G_{z}^{i}\right],$$
(4)

where G_x , G_y and G_z are the discrete gradient operators of our tetrahedral mesh, $\mathbf{v} \in \mathbb{R}^3$ is the direction in which the derivative is to be measured (at the centroid of a tetrahedron) and *i* indexes our tetrahedra. We can assemble these local directional derivative operators into global matrices to produce the global operator $G(\mathbf{v})$.

We proceed by constructing three directional derivative operators, one for each frame director

$$\mathbf{G}_i = G(\mathbf{r}_i) \quad \forall i \in \{1, 2, 3\}.$$
(5)

We frame this problem as a weighted quadratic minimization, with a scalar weight β balancing between two orthogonal notions—uniform spacing of structural members and alignment with coordinate lines:

$$\phi^* = \underset{\phi}{\operatorname{argmin}} \beta \left\| \begin{bmatrix} \mathbf{G}_1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{G}_3 \end{bmatrix} \phi - \begin{bmatrix} \mathbf{1} \\ \mathbf{1} \\ \mathbf{1} \end{bmatrix} \right\|_F^2 + \left\| \begin{bmatrix} \mathbf{0} & \mathbf{G}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{G}_1 \\ \mathbf{G}_2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{G}_2 \\ \mathbf{G}_3 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_3 & \mathbf{0} \end{bmatrix} \phi - \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \right\|_F^2$$

Finally, before extracting the truss by tracing integer isolines of the parametrization, ϕ^* is uniformly scaled to get the desired density.

3 RESULTS AND CONCLUSION

We tested our method on a variety of shapes and fabricated the resulting trusses using various manufacturing processes. 2D models were fabricated using a laser cutter, while 3D models were fabricated using a two different additive manufacturing techniques—Fused Deposition Modelling (FDM, Fig. 1c) with soluble support material, and Selective Laser Sintering (SLS, Fig. 1a). Fabricating with such diverse manufacturing processes is untenable using existing approaches. Further, we plan to expand to pipe bending and dowel rod based construction in the future (Fig. 1b).

We performed mechanical tests to experimentally support our theoretical claim that stress-aligned trusses are strong. Our ABS plastic bridge (Fig. 1c), optimized for compression from the top, weighs 140 grams and is able to withstand the weight of an adult human weighing approximately 93 kg (205 lbs). We also utilized laser cutting to build an optimized bike frame for a wooden kids' bike (Fig. 1d) using 1/4'' Baltic Birch plywood. The bike was tested with a 5-yr old weighing 21 kg (46 lbs) and no failures occurred.

User Control. Unlike existing approaches for topology optimization and truss optimization, the labelled end-to-end curves produced by our method make our results amenable to user control and modification. Currently, we have implemented density selection—a user can simply select a subset of the labelled curves post-optimization—and vertex snapping, allowing a user to improve visual quality of the results by snapping the integer parameter grid to a set of specified vertices (for example, to sharp corners).

Conclusion. We have presented the first algorithm to build volumetric Michell trusses. We evaluated our algorithm by fabricating a variety of structures and performing mechanical tests. We believe that the labelled curves generated by our method also open up avenues for user-controlled structural optimization, and we have demonstrated some initial applications. In the future, we plan to explore a variety of manufacturing methods and structural requirements which can benefit using the user control afforded by our method.

REFERENCES

- M. P. Bendsøe and O. Sigmund. *Topology Optimization*. World Scientific, 2009.
- [2] R. M. Freund. Truss design and convex optimization. *Massachusetts Institute of Technology*, 2004.
- [3] C. Jiang, C. Tang, H.-P. Seidel, and P. Wonka. Design and volume optimization of space structures. ACM Trans. Graph., 36(4), July 2017.
- [4] A. G. M. Michell. LVIII. The limits of economy of material in framestructures. *The London, Edinburgh, and Dublin Philosophical Magazine* and Journal of Science, 8(47), 1904.
- [5] J. Solomon, A. Vaxman, and D. Bommes. Boundary element octahedral fields in volumes. ACM Trans. Graph., 36(3), May 2017.