

Designing Robotically-Constructed Metal Frame Structures

Z. Ma^{1,2}, A. Walzer¹, C. Schumacher², R. Rust¹, F. Gramazio¹, M. Kohler¹, M. Bächer²

¹ETH Zurich ²Disney Research



Figure 1: We present a computational method that aids with the design of structurally-sound metal frames, tailored for automated fabrication using existing robotic fabrication processes. By varying the density of bars, we enable transitions from narrower to wider regions (see chair and table designs, left), or the embedding of artistic targets (letters, right).

Abstract

We present a computational technique that aids with the design of structurally-sound metal frames, tailored for robotic fabrication using an existing process that integrates automated bar bending, welding, and cutting. Aligning frames with structurally-favorable orientations, and decomposing models into fabricable units, we make the fabrication process scale-invariant, and frames globally align in an aesthetically-pleasing and structurally-informed manner. Relying on standard analysis of frames, we then co-optimize the shape and topology of bars at the local unit level. At this level, we minimize combinations of functional and aesthetic objectives under strict fabrication constraints that model the assembly of discrete sets of bent bars. We demonstrate the capabilities of our global-to-local approach on four robotically-constructed examples.

CCS Concepts

• *Applied computing* → *Computer-aided design*; • *Computing methodologies* → *Shape modeling*; *Optimization algorithms*;

1. Introduction

Frame structures made of bent and welded steel bars are omnipresent, with applications in furniture design, concrete reinforcement, art, and architectural ornamentation most prominent (see Fig. 2 left). Because manual assembly is tedious and labor-intensive, robotic construction processes have been proposed. For example, Hack et al. [HWMF*17] propose a robotic system that seamlessly integrates CNC bar bending, cutting, and welding, requiring only limited intervention by a human (Fig. 2 middle).

While such processes achieve an unseen level of complexity when it comes to bar and weld connection counts, Hack et

al. [HWMF*17] assume the use of a regular grid structure, putting severe restrictions on the complexity of fabricable surfaces. As we illustrate in Fig. 2 right, a wider class of surfaces can be approximated if “inbetween” bars are added whenever neighboring bars branch too far apart, or removed whenever they become too close. Moreover, to approximate large and complex shapes like the head in Fig. 2 left, the surface has to be decomposed into fabricable units, and it is unclear how to place and orient bars at the local unit level to achieve an aesthetically pleasing and structurally-sound metal frame at the global scale (Fig. 2 right).

In this paper, we propose a fabrication-aware technique that aids with the design of robotically-constructed metal frame structures of



Figure 2: The 12-meter tall metal frame sculpture “Wonderland” by Spanish sculptor Jaime Plensa and the Diamond Chair design by Harry Bertoia (left). Robotically-constructed rebar for concrete applications (Mesh Mould [HLGK15], middle). Regular structures are insufficient to achieve high curvature detail, and it is unclear how to decompose complex shapes into fabricable units with globally aligned bar structures (right). Lower left image by Sandstein; license: Attribution-ShareAlike 3.0 Unported.

high geometric complexity. Given an input shape, we first compute an approximate stress field under an expected load case. This field is then used to guide the initialization of a globally-aligned pattern from which we generate an initial frame structure.

To assist the user with the decomposition of the model into fabricable units, we provide interactive feedback on where to best place cuts, taking structural and aesthetic considerations into account. Modeling fabricability constraints of the robotic construction process, we optimize each unit patch to fulfill structural and aesthetic requirements while remaining as close as possible to the desired target surface. To this end, we optimize the position of welded connections, and the inclusions of bar segments. We demonstrate the efficacy of our global-to-local technique on several demonstrations, targeting furniture design, or applications in art or architecture.

Succinctly, our technical contributions are

- an integrated design and robotic fabrication process to construct metal frames from input shapes of moderate curvature and complexity,
- a decomposition and global alignment to make the technique scale-invariant, and orient the bars in a structurally- and aesthetically-informed manner, and
- a co-optimization of the shape and topology of bars at the local unit level, maximize the strength-to-weight ratio and distance to the input under strict fabricability constraints.

While we target fabrication on a specific robotic process, we believe that our decomposition, global alignment, and local co-optimization has applications in related processes where a discrete set of entities are assembled into a large-scale structure.

2. Related Work

Computational Design Fabrication-oriented design has garnered an increasing level of interest in recent years. Most notable in

our context are techniques that aid with the design and fabrication of complex *wire structures*. Early work [MIG*14] introduced a method to create 3D printed wireframe structures for fast prototyping of shapes. Wu et al. [WPGM16] extended this approach to a 5DOF system, increasing the complexity of shapes that can be fabricated.

Apart from fast prototyping, wire structures have applications in shape abstraction. Lira et al. [LFZ18] proposed a hybrid meta-heuristic model to create such abstractions with as few wires as possible. On the basis of Chebyshev nets, Garg et al. [GSFD*14] introduced a method to design wire mesh structures. Focusing on aesthetic aspects of wire-based surface design, Zehnder et al. [ZCT16] described a technique for the design of visually-pleasing, structurally-sound ornamental curve networks. In contrast, we target the automated construction of structures made of thicker bars instead of wire, requiring a different modeling.

The work most akin to ours in graphics is Huang et al.’s Frame-Fab [HZH*16]. They assumed a volumetric frame structure as input, and their output is tailored for robotic thermoplastic extrusion. In contrast, we initialize frames from surface meshes, keep them close to the input with an MLS-based objective when optimizing their shape and topology, and target robotic assembly from steel bar, requiring a different modeling of fabricability constraints. While our fabrication process dictates a particular fabrication sequence, the sequencing of instructions could be co-optimized for other robotic processes [WPGM16, HZH*16].

Robotic Fabrication Robotic fabrication processes offer flexibility [DWW*18], hence are well-suited for automated construction at the architectural scale. The robotic assembly of brickwork [DSG*16], or the digitization of concrete [WLR*16] are exciting applications. Most closely related to our integrated design and fabrication of metal frames is work that targets the fabrication of reinforcement for concrete: Cortsen et al. [CRE*14] introduces a robotic system to create double-curved rebar meshes from a digital model. Hack et al. [HLGK15] present a technique to create 3D printed in-place formwork for concrete. They later extended their approach to enable the design of structures that serve as molds and reinforcement [HWMF*17].

While we rely on a similar fabrication process, we enable scale-invariance with our decomposition and global alignment, and the approximation of a wider class of shapes with a fine-grained modeling of fabrication constraints and the support for branching.

Truss and Frame Optimization The analysis of trusses and frames is well-understood, and has been a central topic in structural engineering [Log16]. The optimal design of such structures is an active area of research, with the seminal work on *Mitchell trusses* [Mic04] serving as source of inspiration for more recent work. Truss optimization [Ach99, Ach07], where combinations of shape, topology, and sizing parameters are optimized, has been addressed with a variety of methods [Sto16]. Targeting graphics, Smith et al. [SHOW02] propose a system for the synthesis of truss structures. Tang et al. [TSG*14] explore the use of polyhedral meshes for truss optimization, and target applications in architecture [TSG*14], and Jiang et al. [JTSW17] introduce a method to find material-efficient space structures.

Because forces *and* moments are transferred at welded connections, we rely on *frames* in our modeling. In contrast to standard synthesis and optimization of frames, robotic processes severely restrict the gamut of feasible structures. Hence, they require a different modeling. We believe that (1) our process-aware, structurally-informed initialization of globally-coherent frames (global step), and (2) our fabrication-aware co-optimization of strength-to-weight ratio, target matching, artistic, and regularity objectives (local step) are applicable to other robotic processes.

While our initialization and optimization are tailored to robotic construction processes, we share goals with work on material-minimizing forms and structures [KPWP17], and align bars with principal stress lines [PTP*15, PP18].

3. Overview

Provided with an input model (refer to Fig. 3 Target Surface), we first generate a continuous pattern where we take the bar density and global alignment into account (Global Pattern Initialization), then assist the user with the decomposition of inputs that are too large to fit in the workspace of the robotic fabrication process (Decomposition). After these global processing steps, we extract discrete frame structures from the continuous patterns for each unit (Local Unit Initialization), then co-optimize their shape and topology by parameterizing the location of welded connections, and assigning inclusion variables to a subset of the straight bar segments (Frame Optimization). After these local processing steps, all individual units are guaranteed to fulfill strict fabrication constraints, and are manufactured with the targeted robotic construction process. Manually welding the individual units into a monolithic structure, we then get a structurally-sound metal frame that approximates the input well (Fabricated Result).


3.1. Robotic Construction Process

Before we delve into technical details, it is important to gain a better understanding of the robotic construction process. For robotic construction, we rely on a custom tool head mounted on a 6-axis industrial robotic arm (see Fig. 4 top row) similar to the one described in Hack et al. [HWMF*17]. The tool head combines mechanical units to clamp and bend bars, and to insert, weld, and cut bars. We differentiate between *longer continuous* bars that are bent at a discrete set of locations (in magenta in Fig. 4), and *shorter straight* bars that connect pairs of continuous bars (in blue). Our metal frame units are built from left to right, or right to left as we illustrate in Fig. 4 bottom row and the accompanying video with a fabrication sequence. The following two atomic operations are repeatedly executed until a frame unit is built:

1. Utilizing the 6 DoFs of the robotic arm, the head is positioned and oriented along the leftmost or rightmost continuous bar, and the structure held fixed with the clamp unit while a discrete bend is introduced with the bending unit.
2. The arm is then positioned and oriented to connect the newly added bend with a location previously introduced on the neighboring continuous bar. To this end, when in position, the insertion unit is activated, and introduces a shorter bar segment that is then welded in place, and cut to length.

There are some additional considerations regarding how these two atomic operations can be combined. Because we can skip a discrete bending operation (operation 1), we can introduce shorter bar segments in the middle of a straight segment of a continuous bar (see inset, case 1).

Similarly, we can decide to not introduce a shorter bar segment (operation 2), hence fabricate frame structures with discrete bends that are not connected to neighboring continuous bars (case 2), or only connected on one side. Moreover, continu-



ous bars can end, or can be introduced during manufacturing (case 3). There is one exception: shorter segments can only be welded to continuous bars but not to one another, and not at the same location. Hence, we stagger frames by splitting such nodes, and introducing a small offset between the two shorter segments (staggering). Summarily, while our robotic fabrication process is specific, it enables the fabrication of frame structures of almost arbitrary network topology.

However, the mechanical units, and the dimensions and finite number of DoFs of head and arm, constrain the feasibility of the *shapes* of our frame structures. Bar segments have to observe minimal length constraints, and angle constraints keep us from introducing extreme curvature.

3.2. Fabrication-Aware Optimization

Parameterizing the inclusion of bar segments (topology), and the positions of discrete bending and welding locations (shape), we formulate a strength-to-weight ratio optimization on the unit-level (see Sec. 5; Frame Optimization), guaranteeing their fabricability with a set of constraints. We co-optimize

- the strength-to-weight ratio with an objective that measures the overall volume and compliance of the structure, together with a sparsity regularizer,
- a target matching objective that measures the distance between the frame to the user-specified input using feature-aware moving least squares (MLS),
- and regularity and artistic objectives that penalize irregularities, and enable the embedding of artistic targets in structures.

For the design of a globally coherent, structurally-sound, and visually-pleasing monolithic frame structure, a good initialization is key. As we discuss next, we adopt a recent stripe pattern approach [KCPS15].

4. Frame Initialization and Decomposition

To enable the fabrication of models too large for the reachable volume of the robot, or with locally too high curvature, we decompose our input into smaller units. In a naive approach, we could let the user decompose the model, then initialize each of the units with a regular frame. However, a frame structure consisting of units




Figure 3: Overview. Given a target surface as input (Target Surface), we first create an approximate frame structure with a continuous global parameterization (Global Pattern Initialization) that guides an interactive decomposition of the model into smaller units (Decomposition). For each unit, we then initialize a dense frame structure (Local Unit Initialization), and optimize the positions and topology of bars to meet weight, structural, fabricability, and aesthetic targets (Frame Optimization). Individual units are then constructed with a robotic process, and assembled to a monolithic structure (Fabricated Result).

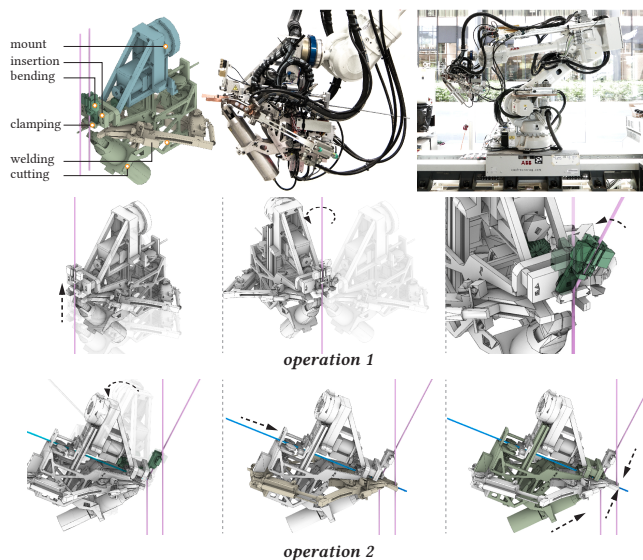


Figure 4: Fabrication Setup and Sequence. A close-up of our tool head (top, left), and the 6-axis robotic arm with mounted head (top, right). Mechanical clamping, bending, bar insertion, welding, and cutting units are used to implement an operation to (1) bend longer continuous bars and (2) insert, weld, and cut shorter straight bars.

with randomly-aligned frames would be (1) difficult to assemble, (2) visually-displeasing, and (3) cut into units across structurally-relevant stress lines.

To decompose models in a structurally-informed manner, we draw a parallel between frames and stripe patterns [KCPS15]: Using principal stresses to guide stripe alignment, and their magnitude to guide stripe density, we generate globally continuous patterns with branching points to maintain regular spacing. Moreover, we base the pattern generation on a stress field to generate stripes that align with global stress lines [PP18]. By aligning continuous bars with stress lines, and assisting the user with the informed decom-

position along these lines, we can significantly increase the structure's load-bearing capacity because stress-aligned bars absorb external loads with axial stresses instead of transverse stresses and moments.

4.1. Generation of Principal Stress-Aligned Stripe Patterns

Treating the model as a standard shell [PH14], and assigning a linear material, we analyze the target model under a user-specified load. For every point on the parameterized surface, we can then evaluate the Cauchy stress, and the principal stress direction and magnitude in the tangent plane. Because small stresses tend to be noisy for non-trivial and tessellated input, and the alignment of bars with principal stresses is most important in high-stress regions, we smooth a weighted version of the principal stress field. To this end, we feed the *Globally Optimal Direction Fields* approach by Köppl et al. [KCPS13] with the field of vectors whose magnitude is set to the squared principal stress, pointing in the direction of the principal stress.

To generate an aligned stripe pattern, we then provide the resulting smoothed vector field (Fig. 5, Principal Stress Field) as input to the pattern generator proposed by Köppl et al. [KCPS15]. Besides a vector field, the generator takes a scalar field that controls the desired density of stripes as an additional input. For regions with zero principal stress, we set this second parameter to the lowest bar density supported by our fabrication process, assigning the maximum supported bar density to regions where the principal stresses peak. For stresses between the two extremes, we linearly interpolate the density value.

We then run the stripe pattern generator twice, once to align longer continuous bars with principal stresses, and once to align shorter straight bars with orthogonal directions, rotating the smoothed vector field by 90 degrees. The result is a pair of parameterizations that can be plugged into a symmetric periodic function like the cosine to visualize the stripes (Fig. 5, Stress-Aligned Stripe Pattern).




Figure 5: Frame Initialization & Decomposition To globally align frame structures, we generate stress-aligned stripe patterns (Stress-Aligned Stripe Pattern) where stripes align with a principal stress field (Principal Stress Field). Assisting the user with the decomposition of the model along stress lines (Decomposition, cuts in red), we initialize parameterizations (Local Parameterizations) to generate longer continuous bars, and shorter straight bar segments (Extracted Frame). Shorter bars are staggered for fabricability (Staggered Frame).

4.2. Structurally-Informed Decomposition

To assist users with the decomposition, we let them choose suitable locations on the input model with visualized stripe pattern, then trace out a cut by following the isolines in either direction until we reach boundaries or an already introduced cut (see red lines in Fig. 5). Interactively adding and removing cuts, users can quickly converge to a structurally-informed decomposition.

4.3. Local Frame Initialization

After decomposition, we initialize a frame structure for every fabricable unit. To achieve the best possible result with our topology and shape optimization, the initial frame structure should be as close as possible to fulfilling our fabricability constraints. Stripe patterns help with this task as stripes can be regularly spaced within the range of densities that our robotic construction supports. However, a hard constraint when initializing frames is that longer continuous bars cannot cross, nor overlap. The same holds for shorter straight bar segments.

To extract frames that meet these desiderata, we build on the routing approach proposed by Pereira et al. [PRM14]. Instead of routing in 3D, we solve for parameterizations that enable the extraction of continuous bars along stress lines from one boundary to its opposite, and shorter bars in orthogonal directions, respectively. To this end, we split the boundary of a unit patch into four segments, and assign the values 0 and 1 to two non-consecutive segments. We then seek a scalar field s defined at every point $\mathbf{x} \in \Omega$ on the patch surface, with isolines of constant value in the interval $[0, 1]$, that align with the stripe pattern, and start and end on the other two non-consecutive segments (Fig. 5 Local Parameterizations). Inspired by the energy formulation used for Poisson Image Editing [PRM14], we solve for a scalar field whose gradient aligns with the guidance field $\mathbf{g}(x)$, set to a scaled version of the stripe pattern field, rotated by 90 degrees

$$\min_{s(\mathbf{x})} \int_{\Omega} \|\nabla s(\mathbf{x}) - \mathbf{g}(\mathbf{x})\|^2 d\mathbf{x} \quad \text{s.t.} \quad s(\mathbf{x}) = \begin{cases} 0 & \text{for } \mathbf{x} \in \partial\Omega_0 \\ 1 & \text{for } \mathbf{x} \in \partial\Omega_1. \end{cases}$$

To find the unique minimizer of this energy, we solve the corresponding Poisson problem $\Delta s(\mathbf{x}) = \nabla \cdot \mathbf{g}(\mathbf{x})$ for the scalar field. We do this twice, alternating the boundary segments and rotating the guidance field by another 90 degrees.

To initialize a frame structure, we first trace continuous bars that are sufficiently far apart in the first parameterization, adding additional continuous bars that end and start at branching points if neighboring bars are not sufficiently dense. We then use the second parameterization to trace along orthogonal directions (Fig. 5 Extracted Frame), *staggering* the shorter straight bar segments because weld connections between shorter and longer bars cannot be introduced at the same location (Fig. 5 Staggered Frame).

While we have not experimented with multiple load cases, a smooth maximum of the principal stress fields could be used when initializing structures.

5. Frame Optimization

After initialization, our model is decomposed into units consisting of a dense set of bars, “zippering” well at the unit-unit boundaries. We achieve the latter property by aligning unit boundaries with global stress lines, and by using 0-1 boundary conditions for the generation of the two local parameterizations.

However, while bars are regularly spaced according to densities supported by the robotic process, there are various other fabrication constraints that limit the shapes of fabricable units. For example, the maximum bending angle, and potential head-structure collisions, limit the local curvature we can achieve. Moreover, bars can likely be removed without stresses in other bars exceeding their limits under the user-specified loads.

To optimize the shape and topology of our frames, we parameterize the locations of weld connections and bends, and assign continuous inclusion variables to the shorter straight bar segments that range between the maximum cross-sectional area (keep the bar) and zero (remove the bar). Co-optimizing the strength-to-weight ratio together with a target matching objective under strict fabricability constraint, we generate construction-ready unit frames that can be welded together to form a sound monolithic structure. Furthermore, with a regularizer favoring regularly-spaced bars, and an objective enabling the embedding of artistic targets, we provide the user with mechanisms to control artistic aspects of structures.

5.1. Representation and Analysis

Because the resulting structures consist of a network of straight bar segments \mathcal{E} , and deformations are expected to remain within the

small strain, elastic regime, standard frame modeling lends itself for analysis (see, e.g., [Log16]). For frames, in contrast to the more widely used truss modeling, forces *and* moments are transferred at connections \mathcal{V} .




Figure 6: Representation and Analysis Parameterized rest configuration on the left, and its deformed configuration on the right.

As we illustrate in Fig. 6, we represent nodes \mathbf{p}_i as 3D points that displace to locations $\mathbf{p}_i + \mathbf{u}_i$ when loads or moments are applied to the structure. To quantify local orientation changes at node i , we quantify rotations about the global coordinate axes with a set of three angles $\phi_i \in \mathbb{R}^3$. When defining inclusion variables, we differentiate between shorter straight segments \mathcal{E}_s , and segments \mathcal{E}_l that belong to longer continuous bars, only parameterizing the former with variables a_{ij} .

When optimizing the strength-to-weight ratio of a frame while keeping stresses within limits, it is important to be able to quantify how changes to shape and topology effect the compliance of the overall structure, and the stresses in individual bars. To this end, we minimize the total potential energy

$$f_{\text{sim}}(\mathbf{u}, \phi) = \sum_{(i,j) \in \mathcal{E}} E_{ij}^{\text{int}}(a_{ij}, \mathbf{p}_i, \mathbf{p}_j, \mathbf{u}_i, \mathbf{u}_j, \phi_i, \phi_j) - \sum_{i \in \mathcal{V}} E_i^{\text{ext}}(\mathbf{u}_i, \phi_i)$$

to first-order optimality. The internal energy E_{ij}^{int} stored in every edge (i, j) depends on the cross-sectional area, and the positions, displacements, and orientations of adjacent nodes (see App. A). The external energy at node i sums up the work $E_i^{\text{ext}} = \mathbf{f}_i^T \mathbf{u}_i + \mathbf{m}_i^T \phi_i$ that external forces \mathbf{f}_i and moments \mathbf{m}_i do on the structure. Because our analysis objective is quadratic in the deformation variables for a linearly elastic material, we can compute the optimal displacements \mathbf{u} and orientations ϕ by solving a linear system of equations [Log16].

In optimizations, we then seek to minimize a design objective over a constraint manifold spanned by an equilibrium constraint, lower and upper bounds on cross-sectional areas, fabricability constraints, and bounds on allowable stresses

$$\min_{\mathbf{a}, \mathbf{p}} f_{\text{design}}(\mathbf{a}, \mathbf{p}, \mathbf{u}, \phi) \quad \text{s.t.} \quad \begin{aligned} & \nabla_{(\mathbf{u}, \phi)} f_{\text{sim}}(\mathbf{u}, \phi) = 0, \\ & 0 \leq a_{ij} \leq a_{\text{max}}, \quad (i, j) \in \mathcal{E}_s, \\ & \text{(fab. constraints, Sec. 5.3),} \\ & \text{(stress bounds, Sec. 5.4).} \end{aligned}$$

Holding external loads and moments fixed, the structure's response changes if we alter the design parameters. Hence, the displacements and local orientations *implicitly* depend on the cross-sectional areas \mathbf{a} and nodal displacements \mathbf{p} . We will further discuss this dependency towards the end of this section.

While we can solve this minimization on the monolithic structure for smaller input, the number of variables and constraints make

numerical optimizations prohibitively expensive for larger targets. To make optimizations scalable, we first analyze the monolithic structure under user-specified loads and moments, then extract the forces and moments along the unit patch boundaries that preserve the equilibrium state at the individual unit-level. We then solve instances of the above design optimization on the unit level, setting the local “external” forces and moments to the extracted set.

While not necessary for the load cases of moderate magnitude we considered in our demonstrations, we could alternate between the global analysis and local design optimizations, solving the latter in parallel. These global-local strategies, while heuristic in nature, have proven effective in various applications in graphics [SA07, BML*14]. Introducing per-load-case displacements and orientations, our design optimization could further be used to determine optimal frame structures under a multitude of load cases.

Before we discuss fabrication constraints and stress bounds, we will discuss a set of objectives that enable the generation of structurally-sound and aesthetically-pleasing frames at the furniture and architectural scale.




Figure 7: Design Objectives Our target matching objective (left) keeps frames close to the user-specified input, while our regularity objective penalizes differences between the initial and the current rest configuration (middle, left), and favors smoothness for segments belonging to continuous bars (middle, right). To embed artistic targets, we move connector points towards attraction directions that we compute from a user-specified scalar field (right).

5.2. Design Objectives

Our design objectives balance structural, target matching, and artistic targets.

Strength-to-Weight Ratio To remove bars that are not needed, we minimize a weighted sum of the overall volume and compliance of the structure

$$f_{\text{s-to-w}} = \frac{1}{V_0} \sum_{(i,j) \in \mathcal{E}} V_{ij} + \alpha \frac{1}{C_0} \sum_{i \in \mathcal{V}} E_i^{\text{ext}}(\mathbf{u}_i, \phi_i). \quad (1)$$

The volume V_{ij} of segment (i, j) is either set to a_{ij} or a_{max} times its length $\|\mathbf{p}_i - \mathbf{p}_j\|$, dependent on whether the bar's cross-sectional area is parameterized or not. To normalize the two terms, we divide them by the volume V_0 and compliance C_0 of the initial structure, respectively. Their relative importance is controlled by the factor α . To favor sparse solutions, we add the approximate L0-regularizer [STC*13]

$$\frac{1}{|\mathcal{E}_s|} \sum_{(i,j) \in \mathcal{E}_s} (a_{ij}^2)^{\frac{1}{\gamma}} \quad (2)$$

to our objective. We set the parameter $\gamma > 2$ to 3 for all our demonstrations, and normalize the term, dividing it by the number of edges with parameterized cross-sectional areas.

Target Matching While the initial frame structure approximates the target surface well, the optimization may move nodes away to fulfill fabrication, or other constraints. To keep frames close to the target, we therefore penalize the signed distance of connections to an implicit representation of the input mesh [ÖGG09]. To this end, we minimize the signed difference between nodes i and neighboring surface points \mathbf{x}_j with normals \mathbf{n}_j (compare with Fig. 7 left)

$$f_{\text{target}} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \frac{1}{2} \left(\frac{\sum_j \mathbf{n}_j^T (\mathbf{p}_i - \mathbf{x}_j) w_{ij}(\mathbf{p}_i)}{\sum_j w_{ij}(\mathbf{p}_i)} \right)^2, \quad (3)$$

with weights w_{ij} decaying with the distance between pairs i and j , and across sharp corners or edges [ÖGG09]. The latter enables the preservation of sharp features.

Regularity To penalize irregularity, we seek to minimize differences between the initial and current rest length of segments, and the initial and current angle between adjacent pairs of edges from the sets $\mathcal{E}_s \times \mathcal{E}_l$ and $\mathcal{E}_l \times \mathcal{E}_l$ (see Fig. 7 middle). Defining the oriented edge $\mathbf{e}_{ij} = \mathbf{p}_j - \mathbf{p}_i$, we penalize differences in rest length with per-edge terms $\frac{1}{2} (\|\mathbf{e}_{ij}\| - \|\mathbf{e}_{ij}^0\|)^2$. To penalize differences in rest angles $\beta_{ijk} = \angle(\mathbf{e}_{ij}, \mathbf{e}_{ik})$ between pairs of edges, we minimize terms of the form $\sin^2(\beta_{ijk} - \beta_{ijk}^0)$. Note that the use of the squared sine has the advantage that we can express the term with dot- and cross-products between the two oriented edges.

While we could use the sine-penalty for both types of adjacent edges, we favor a term that penalizes deviations from straightness, and discourages proximity to bending angle limits for pairs of segments belonging to a continuous bar. This choice tends to help with circumnavigating local minima, and positively affects fabrication time also. To this end, we use the term $(1 + \cos(\beta_{ijk}))^\tau$ to measure the deviation of the angle $\beta_{ijk} = \angle(\mathbf{e}_{ij}, \mathbf{e}_{ik})$ from 180 degrees, setting the penalty parameter τ to 5 for all our demonstrations.

Summing up the individual penalty terms, normalizing the three penalty types by dividing by the number of edges, or adjacent edge pairs in the sets $\mathcal{E}_s \times \mathcal{E}_l$ and $\mathcal{E}_l \times \mathcal{E}_l$, and weighing them, we define our regularity objective f_{reg} .

Embedding Artistic Targets To enable the embedding of artistic targets, we let users define scalar fields $s(\mathbf{x})$ on target surfaces, where surface points \mathbf{x}_j with positive values $s(\mathbf{x}_j) > 0$ attract connectors, and those with negative values $s(\mathbf{x}_j) < 0$ repel them. To define a direction a connector i should move toward, we evaluate the scalar field for surface points in a neighborhood of the initial position \mathbf{p}_i^0 (see Fig. 7 right), computing an average attraction direction

$$\mathbf{d}_i = \sum_j s(\mathbf{x}_j) w(\|\mathbf{p}_i^0 - \mathbf{x}_j\|) \frac{\mathbf{p}_i^0 - \mathbf{x}_j}{\|\mathbf{p}_i^0 - \mathbf{x}_j\|}, \quad (4)$$

where we give attraction and repulsion forces for points that are closer \mathbf{p}_i^0 a higher weight with a decay function w . We then measure

if connectors move toward these directions by minimizing

$$f_{\text{art}} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \mathbf{d}_i^T (\mathbf{p}_i - \mathbf{p}_i^0). \quad (5)$$

While we keep the attraction directions constant, we could update them by exchanging the initial connector position in \mathbf{d}_i with its current position. This is in particular important if we expect frame structures to move significantly. In our implementation, we chose the decay function $w(\|\mathbf{p}_i^0 - \mathbf{x}_j\|) = \frac{1}{1 + \|\mathbf{p}_i^0 - \mathbf{x}_j\|}$. While this is a good choice if the attraction directions remain constant, a function with a finite neighborhood is desirable if we made the attraction directions depend on the current positions \mathbf{p}_i .




Figure 8: Fabrication Constraints To avoid head-frame collisions during robotic assembly, we limit the bending angle and bending plane for segments that belong to longer continuous bars (left), and the insertion angle for shorter bars that connect two neighboring continuous bars (right).

5.3. Fabrication Constraints

While our design objectives trade off functional and aesthetic targets, the robotic fabrication process poses hard limits: The angle and plane for a discrete bending operation, the length of bar segments, and the angle for an insertion operation of a shorter segment are constrained by head-structure collisions, and the limited degrees of freedom of the head w.r.t. the continuous bar it is moving along.

Bending Angle and Plane The bending tool limits the maximum angle, and we constrain the angle between adjacent bar segments from the set $\mathcal{E}_l \times \mathcal{E}_l$ to be smaller or equal to the maximum bending angle α_{bend} with constraints of the form (see Fig. 8 left)

$$\angle(-\mathbf{e}_{il}, \mathbf{e}_{ik}) \leq \alpha_{\text{bend}}. \quad (6)$$

The bending tool bends continuous bars in a plane perpendicular to the direction the insertion tool is extruding shorter straight bar segments toward. To avoid structure-head collisions when introducing a bend at location i , we favor if the normal of the bending plane lies in, or close to the local tangent plane at i . Approximating the tangent plane at i with the plane spanned by the already fabricated segment \mathbf{e}_{il} , and the yet-to-be-inserted, shorter segment \mathbf{e}_{ij} , we constrain the dot product between the normal of the bending plane, $\mathbf{e}_{il} \times \mathbf{e}_{ik}$, and the normal of the tangent plane, $\mathbf{e}_{il} \times \mathbf{e}_{ij}$, to not deviate too far from orthogonality

$$(\mathbf{e}_{ij} \times \mathbf{e}_{il}) \cdot (\mathbf{e}_{ij} \times \mathbf{e}_{ik}) \leq \epsilon_{\text{plane}}. \quad (7)$$

Bar Length The discrete bars that are inserted between continuous bars are welded using a clamp whose size limits the length that these bars can have. Additionally, the tool head has to move along the continuous bar, which introduces a minimal distance requirement between adjacent continuous bars. Together, these constraints introduce a lower and upper bound, l_{\min}^s and l_{\max}^s , on the length of shorter, straight bar segments:

$$l_{\min}^s \leq \|\mathbf{e}_{ij}\| \leq l_{\max}^s, \quad (i, j) \in \mathcal{E}_s. \quad (8)$$

Similarly, the size of the tool head imposes a minimal distance constraint between consecutive bends along continuous bars

$$l_{\min}^l \leq \|\mathbf{e}_{ij}\|, \quad (i, j) \in \mathcal{E}_l \quad (9)$$

where l_{\min}^l is the lower length bound.

Insertion Angle To insert a shorter bar segment, the tool head has to tilt forward or backward (Fig. 8 right), and the amount the head can tilt before colliding with the structure is limiting the two angles between the three oriented bars connecting i to j , k , and l

$$\angle(\mathbf{e}_{ij}, \mathbf{e}_{il}) \leq \alpha_{\text{back}} \quad \text{and} \quad \angle(\mathbf{e}_{ij}, \mathbf{e}_{ik}) \leq \alpha_{\text{forw}}. \quad (10)$$

5.4. Stress Bounds

Apart from ensuring fabricability, we have to safeguard against structural instabilities. To this end, we keep axial stresses

$$\sigma_{ij} = E\epsilon_{ij} = E \frac{\|(\mathbf{p}_j + \mathbf{u}_j) - (\mathbf{p}_i + \mathbf{u}_i)\| - \|\mathbf{e}_{ij}\|}{\|\mathbf{e}_{ij}\|}. \quad (11)$$

in segments (i, j) from yielding and buckling with constraints

$$|\sigma_{ij}| < \sigma_{\text{yield}}, \quad \text{and} \quad \sigma_{ij} > -\frac{1}{a_{\max}} \frac{\pi^2 EI}{(K\|\mathbf{e}_{ij}\|)^2}, \quad (12)$$

respectively [Log16]. We set the column effective length factor K in Euler's critical load criterion to 1.2, Young's modulus E to the tabulated value for the steel we use, and the area moment of inertia to $I = \frac{1}{4\pi} a_{\max}^2$. We deliberately choose a_{\max} instead of a_{ij} in the criterion. While this choice overestimates the buckling resistance for intermediate states, it guarantees that the constraints are fulfilled if we set non-zero cross-sectional areas to a_{\max} after optimizations.

5.5. Implementation Detail

We solve our optimization problem with KNITRO's implementation of the interior point method with BFGS [BNW06]. We treat all constraints explicitly, with the exception of the equilibrium constraint, which we implicitly enforce by updating the deformation variables \mathbf{u} and ϕ whenever the design variables change. This requires us to determine the derivatives of the deformed configuration w.r.t. design variables, which we can compute by applying the *implicit function theorem*.

6. Fabrication

While the procedure of the fabrication has been described in Sec. 3.1, we hereafter provide technical specifications.

Robotic Setup Our setup consists of a 6-axis ABB IRB 4600 robotic arm with a 40kg payload, a reach of 2.55m, and a position repeatability of less than 0.01 mm (Fig. 4). The arm is mounted on an ABB IRBT 2005 linear track and equipped with a process-specific, custom-built end effector [KHD*17] that is able to

- plastically cold-form continuous rebar with a diameter of up to 6 mm (Grade B500A or B500B) through a bending mechanism,
- pneumatically cut continuous rebar into straight segments of appropriate length,
- hydraulically clamp a pair of welding electrodes that are actuated by a four-bar linkage, onto pairs of bars, reaching a maximum of 6kN pressure, and
- cross-weld shorter bar segments to continuous bars with an industrial-grade resistance welding setup with a peak output current of 21 kA.

Constraint Bounds We experimentally determined bounds for our fabrication constraints: During insertion of shorter bar segments, the tool head can tilt backward by 8° ($\alpha_{\text{back}} = 98^\circ$), and forward by 15° ($\alpha_{\text{forw}} = 105^\circ$). The minimum and maximum length of the bar segments are $l_{\min}^s = 28\text{mm}$ and $l_{\max}^s = 105\text{mm}$, respectively, while the minimum distance between nodes (the minimum length of each segment on the continuous bar) is $l_{\min}^l = 15\text{mm}$. The bending tool can bend a bar by up to $\alpha_{\text{bend}} = 60^\circ$ in either direction, taking the spring-back angle of 12° into account. The bending plane constraint is $\epsilon_{\text{plane}} = \cos(20^\circ)$.

Steel Bar For fabrication, we use European standard BS EN 10080: B500A rebar with a diameter of 4.5mm and a material strength of $\sigma_c = 500\text{MPa}$. The Young's modulus of the material is $E = 210\text{GPa}$, and its shear modulus is $G = 150\text{GPa}$. For the continuous bars in our *Letter Wall Elements*, we used 6mm B500B rebar with the same material properties.

7. Results

We have used our method to optimize and fabricate a Stool and Table design (Figs. 1,9,11), and two wall elements with engraved letters (Figs.13), illustrating use in furniture design and architectural ornamentation. In addition, we demonstrate our technique on several in-simulation examples, evaluating it under changes to load cases, target surfaces, and artistic targets, besides an analysis of our decomposition strategy and the direction-dependence of our optimized results. We report optimization parameters in Tab. 1.

Stool and Table Furniture, such as stools and tables, must withstand extensive live loads from a person sitting on them, or loads of objects placed on them. For both models, we keep the ground nodes fixed, and apply a uniformly distributed load of 1500N at the top, pointing in the direction of gravity (Figs. 9,11). Note that both models are significantly narrower in the mid-region, and widen towards the two ends. Hence, a direct fabrication is not possible, and a decomposition into at least three units is necessary (Figs. 10,12). To bridge the gap between the longer continuous bars in regions of varying curvature, our initialization inserts additional shorter continuous bars. For our Table with a higher variation in curvature, up to three shorter (continuous) bars are added between the continuous bars that run along the full height of the model. During shape

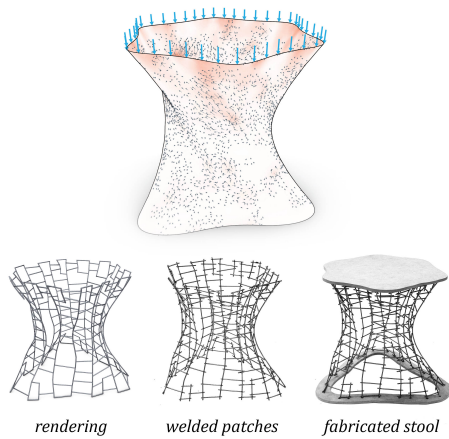


Figure 9: We apply a top load of 1500N to our Stool example, resulting in regions of high stress at the top of the model (top). Our technique generates a structurally-sound frame model (bottom left) that is fabricable with our robotic fabrication approach (bottom middle, right).




Figure 10: During the design process, the Stool example is split into three units (center). For each unit, we visualize the stress distribution before (left) and after (right) optimization.

and topology optimizations of the individual units, shorter bar segments that are not necessary for structural integrity, are removed. To make our fabricated models functional, we add wooden seats or top surfaces, and footing (Fig.1). As reported in Tab. 1, our structural optimization increases the collapse load of the structures.

Considering principal curvature, our Stool and Table examples consist of points of elliptical (both principal curvatures have same sign), hyperbolic (opposite sign), parabolic (one set to zero), and planar type (both set to zero). Summarily, if we decompose input models into units that consist of points at which the two principal curvatures are (1) bound from above and below with a constant of moderately high value, and (2) vary sufficiently slowly along the shortest path between any pairs of points, we can expect our global-to-local strategy to produce the desired fabricable output.




Figure 11: The Table example is subject to a 1500N top load (top). During initialization, up to three shorter continuous bars are added to bridge the gaps due to the extreme changes in circumference along the model's height.




Figure 12: Stress visualization for the three unit patches (Table) show an improvement after structural optimization (compare left to right), resulting in a higher collapse load (see Tab. 1).

Letter Wall Elements Beyond designing fabricable and functional output, we support the embedding of artistic input. To this end, we apply a user-specified *attraction field* $s(\mathbf{x})$ (see Sec. 5.2) to the structure. The optimization will then solve for a frame structure that balances structural properties and adherence to the artistic target. We demonstrate this approach on curved wall elements with embedded letters (Fig. 13). These 1.5 m tall wall elements are intended to withstand a moderate vertical load of 300N. The initialization creates a very regular structure that is ideal to absorb these forces. Through an attraction field defined by a texture, the optimization then allows for the letters *S* and *I* to appear. Notice that we only optimize the node positions of the structure, and leave the topology unchanged to avoid visual distraction. These letter examples demonstrate for models of increasing curvature, how well




Figure 13: Using shape optimization only, we embed artistic targets into Wall Elements. Starting from a regular frame (left), the optimization arrives at a solution that “embeds” the letters S and I through variations in densities and orientations of bars (center). The fabricated models maintain this visual appeal (right).

Table 1: We report the size of initialization grids (without additional continuous bars), the objective weights, optimization timings for our fabricated examples, and collapse loads before and after optimization. The timings were collected on a 4.0 GHz Intel Core i7-6700K quad-core processor with 32 GB of RAM.

| Model | Patch | Stool | | | Table | | | Letters | |
|-------------|------------------|-------|------|------|-------|------|------|---------|------|
| | | 1 | 2 | 3 | 1 | 2 | 3 | S | I |
| Grid | contin. | 5 | 5 | 5 | 4 | 4 | 4 | 15 | 15 |
| | discrete | 10 | 10 | 10 | 17 | 17 | 17 | 19 | 21 |
| | w_{s-to-w} | 2 | 2 | 2.2 | 2 | 2 | 2 | 2 | 2 |
| | γ | 1.5 | 1 | 0.9 | 1.5 | 1.5 | 1.5 | 1.5 | 1.5 |
| | w_{sparse} | 2.5 | 3.4 | 3 | 4 | 4 | 4.4 | - | - |
| Weights | w_{target} | 5 | 5 | 5 | 5 | 5 | 5 | 10 | 10 |
| | $w_{reg.length}$ | 0.01 | 0.07 | 0.01 | 0.04 | 0.03 | 0.03 | 0.03 | 0.03 |
| | $w_{reg.sine}$ | - | - | - | 10 | 10 | 10 | - | - |
| | $w_{reg.cosine}$ | 110 | 150 | 160 | 260 | 200 | 200 | 400 | 400 |
| | w_{art} | - | - | - | - | - | - | 1300 | 1200 |
| Opt Runtime | | 2m | 4m | 2m | 11m | 6m | 19m | 67m | 119m |
| Collapse | initial | 4510N | | | 5105N | | | 75N | 365N |
| Load | opt. | 4695N | | | 5150N | | | 305N | 410N |

connections can move significantly along, and in close proximity to the target surface, without getting trapped in local minima.

Parameter Modifications Changes to loading scenarios can have a significant effect on the resulting structure. In Fig. 14, we show an optimized wall element under three different loading scenarios: top load, shear load, and cantilever load. Only for the top load case, the continuous bars align well with stress lines, leaving shorter bar segments room to align with an artistic target. For the other two

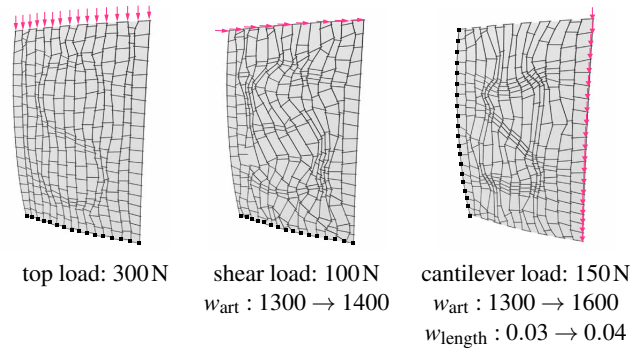


Figure 14: Sensitivity to Load Cases Changes to the loading scenario of our top-loaded wall element S (left) result in significantly different optimization results. Under a shear (middle) or cantilever load case (right), the structure’s load carrying capacity is significantly lower, and the artistic intent can no longer be fully expressed.




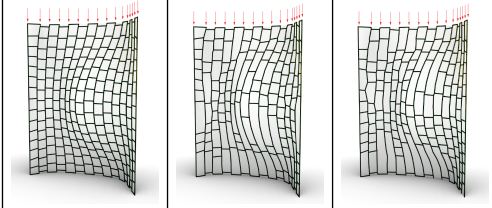
Figure 15: Higher-Curvature Target Surface and Artistic Input Modifications. Replacing the target surface for the S with a more curved surface changes the resulting frame structures (left), while the artistic intent can still be emphasized by the optimization. Similarly, the artistic target can be exchanged, for example by using a G texture as input (right).

load cases, our method favors the structural integrity of the result—defined through constraints—at the expense of the artist’s intent.

In Fig. 15, we investigate the influence of changes to target shapes, or artistic targets, on optimization results: We transfer the letter S to a target shape of higher curvature, and show that the artistic intent is preserved. Conversely, we show that another artistic target (letter G) can be embedded in the same target shape, with only slight modifications to optimization parameters.

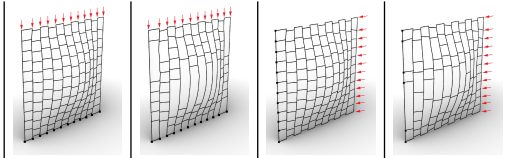
Decomposition Analysis As we demonstrate with a curved, uniformly loaded wall element with a bulge in its center in Tab. 2, our decomposition strategy results in a very substantial decrease in optimization runtime (10m vs. 1h31m) at the cost of higher average displacements and stresses. Note that both optimized models are fabricable.

Table 2: Decomposition Analysis We report the weight, average stress in bars, average displacement of nodes, as well as the run-time of a global optimization, and a subdivision into 2×2 units.



| Model | initial | no decomp. | 2×2 decomp. |
|--------------|---------|------------|----------------------|
| Weight | 8.4 kg | 7.2 kg | 7.2 kg |
| Avg. Stress | 7.1 MPa | 6.4 MPa | 9.1 MPa |
| Avg. Displ. | 1.61 mm | 1.50 mm | 2.21 mm |
| Fabricable | no | yes | yes |
| Opt. Runtime | - | 1h 31m | 10m |

Table 3: Direction-Dependence In this analysis, we report the weight and maximum stress for a wall element optimized under either vertical or horizontal loading.



| Load | vertical | | horizontal | |
|-------------|----------|-----------|------------|-----------|
| Model | initial | optimized | initial | optimized |
| Weight | 6.5 kg | 5.3 kg | 6.5 kg | 5.8 kg |
| Max. Stress | 199 MPa | 167 MPa | 224 MPa | 161 MPa |

Direction-Dependence To study the direction-dependence of the load carrying capacity of structures (Tab. 3), we uniformly load a wall element with a bulge in its center from the top (vertical, with nodes at the base fixed), and the side (horizontal, with boundary nodes on the opposite edge fixed). As expected, the performance of the structure is best if long continuous bars align with the direction of the load case (vertical). At the cost of additional weight, the optimization can compensate the performance gap in structural strength (horizontal).

8. Conclusions

We presented a computational technique that enables the *scale-invariant* fabrication of structurally-sound and aesthetically-pleasing metal frame structures at the furniture and architectural scale. While global optimizations for models with high bar counts are infeasible, we resolve this bottleneck with a global-to-local strategy: taking structural and fabricability considerations into account when initializing frame structures (global step), we start local unit optimizations from sensible initial guesses, requiring only local refinements (local step).

8.1. Limitations and Future Directions

While our global-to-local approach scales well with the number of bars, it introduces inconsistencies in forces and moments along unit boundaries. While we can easily check for the structural stability of the final optimization result, and have yet to identify a scenario where this assumption is responsible for non-fabricable output, the exploration of optimization strategies that alternate between a global analysis and local design optimizations is an exciting future direction.

In terms of fabrication, the complexity of metal frames is restricted by the fabrication constraints imposed by our robotic fabrication process. Especially the bounds on the insertion and bending plane angles turn out to be limiting factors. Because we explicitly model these constraints, we can utilize our tightly integrated design and fabrication technique to identify the ones that are most restrictive, and propose targeted updates to the fabrication hardware.

With our fabrication constraints, we only avoid local robot-structure collisions. We assume our input to be self-intersection free, and of moderate curvature and complexity. If pairs of surface regions come too close in a model, it is best to decompose it into units where such pairs are part of different units. In the future, we plan to extend our technique to handle global collisions.

Finally, our processing starts with a shell simulation of the target surface to extract stress lines. While this approximate analysis provides us with a sufficient estimate for frame initialization, our computational technique could benefit from a continuous shell model that homogenizes the behavior of the inherently discrete metal frame structures [KMOD09].

Appendix A: Analysis

The internal energy stored in a deformed bar segment is

$$E_{ij}^{int} = \frac{1}{2} \left[(\mathbf{u}_i - \mathbf{u}_j)^T \mathbf{A}_{ij} (\mathbf{u}_i - \mathbf{u}_j) + (\mathbf{u}_i - \mathbf{u}_j)^T \mathbf{B}_{ij} (\phi_i + \phi_j) + (\phi_i - \phi_j) \mathbf{C}_{ij} (\phi_i - \phi_j) + \phi_i \mathbf{D}_{ij} \phi_j \right]$$

where the elemental stiffness matrices [Log16]

$$\mathbf{A}_{ij} = \mathbf{R}_{ij}^T \begin{bmatrix} \frac{Ea_{ij}}{l_{ij}} & 0 & 0 \\ 0 & \frac{3Ea_{ij}^2}{\pi l_{ij}^3} & 0 \\ 0 & 0 & \frac{3Ea_{ij}^2}{\pi l_{ij}^3} \end{bmatrix} \mathbf{R}_{ij}, \quad \mathbf{B}_{ij} = 2\mathbf{R}_{ij}^T \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & \frac{3Ea_{ij}^2}{2\pi l_{ij}^2} \\ 0 & -\frac{3Ea_{ij}^2}{2\pi l_{ij}^2} & 0 \end{bmatrix} \mathbf{R}_{ij},$$

$$\mathbf{C}_{ij} = \mathbf{R}_{ij}^T \begin{bmatrix} \frac{Ga_{ij}^2}{2\pi l_{ij}} & 0 & 0 \\ 0 & \frac{Ea_{ij}^2}{\pi l_{ij}} & 0 \\ 0 & 0 & \frac{Ea_{ij}^2}{\pi l_{ij}} \end{bmatrix} \mathbf{R}_{ij}, \quad \mathbf{D}_{ij} = \mathbf{R}_{ij}^T \begin{bmatrix} 0 & 0 & 0 \\ 0 & \frac{3Ea_{ij}^2}{\pi l_{ij}} & 0 \\ 0 & 0 & \frac{3Ea_{ij}^2}{\pi l_{ij}} \end{bmatrix} \mathbf{R}_{ij}$$

depend on the constant rotation matrix \mathbf{R}_{ij} that rotates the vector $\mathbf{e}_{ij} = \mathbf{p}_i - \mathbf{p}_j$ to the global x-axis $[0, 0, 1]^T$, the rest length $l_{ij} = \|\mathbf{e}_{ij}\|$ of the bar, its cross sectional area a_{ij} , and the Young's and shear moduli E and G , respectively.

References

[Ach99] ACHTZIGER W.: Local stability of trusses in the context of topology optimization Part II: A numerical approach. *Structural Optimization* 17, 4 (dec 1999), 247–258. 2

- [Ach07] ACHTZIGER W.: On simultaneous optimization of truss geometry and topology. *Structural and Multidisciplinary Optimization* 33, 4 (2007). 2
- [BML*14] BOUAZIZ S., MARTIN S., LIU T., KAVAN L., PAULY M.: Projective dynamics: Fusing constraint projections for fast simulation. *ACM Trans. Graph.* 33, 4 (July 2014). 6
- [BNW06] BYRD R. H., NOCEDAL J., WALTZ R. A.: KNITRO: An integrated package for nonlinear optimization. In *Large Scale Nonlinear Optimization* (2006), Springer Verlag, pp. 35–59. 8
- [CRE*14] CORTSEN J., RYTZ J. A., ELLEKILDE L.-P., SØLVASON D., PETERSEN H. G.: Automated fabrication of double curved reinforcement structures for unique concrete buildings. *Robotics and Autonomous Systems* 62, 10 (2014). 2
- [DSG*16] DÖRFLER K., SANDY T., GIFTHALER M., GRAMAZIO F., KOHLER M., BUCHLI J.: *Mobile Robotic Brickwork*. 2016. 2
- [DWW*18] DAI C., WANG C. C. L., WU C., LEFEBVRE S., FANG G., LIU Y.-J.: Support-free volume printing by multi-axis motion. *ACM Trans. Graph.* 37, 4 (2018). 2
- [GSFD*14] GARG A., SAGEMAN-FURNAS A. O., DENG B., YUE Y., GRINSPUN E., PAULY M., WARDETZKY M.: Wire Mesh Design. *ACM Transactions on Graphics* 33, 4 (jul 2014), 66:1–66:12. 2
- [HLGK15] HACK N., LAUER W. V., GRAMAZIO F., KOHLER M.: Mesh mould: Robotically fabricated metal meshes as concrete formwork and reinforcement. In *Proceedings of the 11th International Symposium on Ferrocement and 3rd ICTRC International Conference on Textile Reinforced Concrete* (Bagnex, 2015), Brameshuber W., (Ed.), RILEM Publications SARL, pp. 347–359. 2
- [HWMF*17] HACK N., WANGLER T., MATA-FALCÓN J., DÖRFLER K., KUMAR N., WALZER A. N., GRASER K., REITER L., RICHNER H., BUCHLI J., KAUFMANN W., FLATT R. J., GRAMAZIO F., KOHLER M.: Mesh mould: an on site, robotically fabricated, functional formwork. In *Proceedings of the Second Concrete Innovation Conference (2nd CIC)* (2017). 1, 2, 3
- [HZH*16] HUANG Y., ZHANG J., HU X., SONG G., LIU Z., YU L., LIU L.: FrameFab: Robotic Fabrication of Frame Shapes. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 224:1–224:11. 2
- [JTSW17] JIANG C., TANG C., SEIDEL H.-P., WONKA P.: Design and volume optimization of space structures. *ACM Trans. Graph.* 36, 4 (July 2017), 159:1–159:14. 2
- [KCPS13] KNÖPPEL F., CRANE K., PINKALL U., SCHRÖDER P.: Globally optimal direction fields. *ACM Trans. Graph.* 32, 4 (July 2013), 59:1–59:10. 4
- [KCPS15] KNÖPPEL F., CRANE K., PINKALL U., SCHRÖDER P.: Stripe patterns on surfaces. *ACM Trans. Graph.* 34, 4 (2015), 39:1–39:11. 3, 4
- [KHD*17] KUMAR N., HACK N., DÖRFLER K., WALZER A. N., REY G. J., GRAMAZIO F., KOHLER M. D., BUCHLI J.: Design, development and experimental assessment of a robotic end-effector for non-standard concrete applications. In *2017 IEEE International Conference on Robotics and Automation (ICRA)* (May 2017), pp. 1707–1713. 8
- [KMOD09] KHAREVYCH L., MULLEN P., OWHADI H., DESBRUN M.: Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3 (2009). 11
- [KPWP17] KILIAN M., PELLIS D., WALLNER J., POTTMANN H.: Material-minimizing forms and structures. *ACM Trans. Graph.* 36, 6 (2017). 3
- [LFZ18] LIRA W., FU C.-W., ZHANG H.: Fabricable eulerian wires for 3d shape abstraction. *ACM Trans. Graph.* 37, 6 (Dec. 2018), 240:1–240:13. 2
- [Log16] LOGAN D. L.: *A First Course in the Finite Element Method*, 6th ed. Cengage Learning, 2016. 2, 6, 8, 11
- [Mic04] MICHELL A. G.: The limits of economy of material in frame-structures. *Philosophical Magazine* 8, 47 (1904). 2
- [MIG*14] MUELLER S., IM S., GUREVICH S., TEIBRICH A., PFISTERER L., GUIMBRETIERE F., BAUDISCH P.: Wireprint: 3d printed previews for fast prototyping. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2014), UIST '14, ACM, pp. 273–280. 2
- [ÖGG09] ÖZTIRELI A. C., GUENNEBAUD G., GROSS M.: Feature preserving point set surfaces based on non-linear kernel regression. *Computer Graphics Forum* 28, 2 (2009), 493–501. 7
- [PH14] PREISINGER C., HEIMRATH M.: Karamba—a toolkit for parametric structural design. *Structural Engineering International* 24, 2 (2014), 217–221. 4
- [PP18] PELLIS D., POTTMANN H.: Aligning principal stress and curvature directions. *Advances in Architectural Geometry. Klein Publishing Ltd* (2018), 34–53. 3, 4
- [PRM14] PEREIRA T., RUSINKIEWICZ S., MATUSIK W.: Computational light routing: 3d printed optical fibers for sensing and display. *ACM Trans. Graph.* 33, 3 (2014), 24:1–24:13. 5
- [PTP*15] PIETRONI N., TONELLI D., PUPPO E., FROLI M., SCOPIGNO R., CIGNONI P.: Statics aware grid shells. *Comput. Graph. Forum* 34, 2 (2015). 3
- [SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2007), SGP '07, Eurographics Association, pp. 109–116. 6
- [SHOW02] SMITH J., HODGINS J., OPPENHEIM I., WITKIN A.: Creating models of truss structures with optimization. *ACM Trans. Graph.* 21, 3 (2002), 295–301. 2
- [STC*13] SKOURAS M., THOMASZEWSKI B., COROS S., BICKEL B., GROSS M.: Computational design of actuated deformable characters. *ACM Trans. Graph.* 32, 4 (July 2013), 82:1–82:10. 6
- [Sto16] STOLPE M.: Truss optimization with discrete design variables: a critical review. *Structural and Multidisciplinary Optimization* 53, 2 (2016), 349–374. 2
- [TSG*14] TANG C., SUN X., GOMES A., WALLNER J., POTTMANN H.: Form-finding with polyhedral meshes made simple. *ACM Trans. Graph.* 33, 4 (2014). 2
- [WLR*16] WANGLER T., LLORET E., REITER L., HACK N., GRAMAZIO F., KOHLER M., BERNHARD M., DILLENBURGER B., BUCHLI J., ROUSSEL N., FLATT R.: Digital concrete: Opportunities and challenges. *RILEM Technical Letters* (2016). 2
- [WPGM16] WU R., PENG H., GUIMBRETIERE F., MARSCHNER S.: Printing arbitrary meshes with a 5dof wireframe printer. *ACM Trans. Graph.* 35, 4 (July 2016), 101:1–101:9. 2
- [ZCT16] ZEHNDER J., COROS S., THOMASZEWSKI B.: Designing structurally-sound ornamental curve networks. *ACM Trans. Graph.* 35, 4 (July 2016), 99:1–99:10. 2