Eurographics Symposium on Geometry Processing 2015
Mirela Ben-Chen and Ligang Liu
(Guest Editors)

Volume 34 (2015), Number 5

Reforming Shapes for Material-aware Fabrication

Yong-Liang Yang[1,2] Jun Wang[3,2] Niloy J. Mitra[4]

[1]University of Bath [2]KAUST [3]Baidu Research [4]University College London

Figure 1: *We present a data-driven framework to reform (i.e., reshape) an input multi-component mesh to simplify fabrication from a target built material. The input model (left) can be geometrically and topologically modified to a reformed shape (middle). Part conjoining specifications such as joint information can also be inferred for real fabrication (right).*

Abstract

As humans, we regularly associate shape of an object with its built material. In the context of geometric modeling, however, this inter-relation between form and material is rarely explored. In this work, we propose a novel data-driven reforming (i.e., reshaping) algorithm that adapts an input multi-component model for a target fabrication material. The algorithm adapts both the part geometry and the inter-part topology of the input shape to better align with material-aware fabrication requirements. As output, we produce the reshaped model along with respective part dimensions and inter-part junction specifications. We evaluate our algorithm on a range of man-made models and demonstrate a variety of model reshaping examples focusing only on metal and wooden materials.

1. Introduction

Geometric form of a physical object is strongly dictated by its built material. This is not surprising since materials differ as to how they can be shaped towards a geometric form. For example, it is desirable, both in terms of increased convenience and reduced waste, to cut wood into straight planks, while others like metal sheets or plywood can easily be given a curved profile without much additional cost. Thus built materials affect both the shape of the parts and how they can be interconnected. As humans, we often correlate geometric appearance of an object, even in absence of any texture or color information, with its fabrication material (see Fig. 2).

In geometric modeling, shape and physical material are rarely considered together. Traditionally, as graphics objects

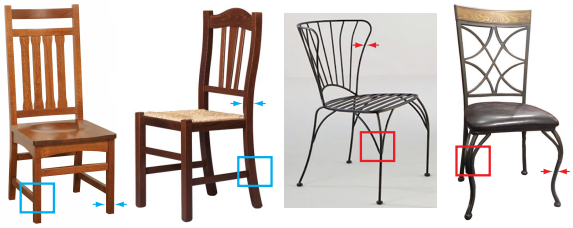

Figure 2: *Geometric properties such as thickness of individual parts and contact angles between parts are characteristics of the built materials. Blue box highlights near right angles between wooden parts, while red highlights more flexible angles between metal parts.*

are largely used in virtual environments, such an approach is entirely justified. Later, one can assign any (virtual) material texture to any geometric shape. Only recently, with growing interest in actual fabrication to bring virtual objects back to the real world, fabrication-aware modeling has gained popularity [BBO*10, LSH*10, HBA12, PWLSH13, SP13]. Such methods, however, focus on rationalization (i.e., close approximation) of a target shape by deriving an economic and feasible solution. Instead, we focus on reforming shapes by actually reshaping the target shape to simplify fabrication.

Given an input shape, we are interested in *how it will look if fabricated using a target built material*. Essentially, we investigate the implicit relation between shape and built material. For example, the chair model shown in Fig. 1-bottom left with thin parts connected at narrow angles would be difficult to fabricate in its original shape with wood. A reformed shape with thicker parts meeting together at near-orthogonal angles is much more suited for wood-based fabrication (Fig. 1-middle). We present a data-driven reshaping algorithm for this purpose that proposes a new shape along with necessary fabrication specifications (Fig. 1-right).

Starting from a set of objects with part materials assigned, in a preprocessing stage, we extract the correlation between shape of parts and their assigned raw material in the context of component-based man-made objects. Note that we do *not* assume the components to have semantic labels (chair leg, table top, etc.). In this paper, we only consider two types of materials: wood and metal. Based on the learned information, we adapt input models to suit fabrication from a target material. Since different materials favor different part pair contact profiles (i.e., contact angles), the proposed algorithm adapts the parts to better conform to preferred contact profiles. This involves both topological adaptations by breaking old and creating new contacts, and geometric adaptations by solving for appropriate new contact locations for corresponding part profiles. Finally, we use a standard FEM simulator to identify possible stress hotspots in the designed shape. If so, we use a simple parameter refinement search to arrive at a better shape for fabrication. Our system outputs both part dimensions and inter-part connection specifications to simplify physical fabrication of the reshaped objects. We evaluate our framework on a range of data sets and demonstrate its efficacy for novel material-driven object reshaping. In summary, our contributions include: (i) formulating a computational connection between geometric form and physical fabrication material; and (ii) introducing a novel data-driven algorithm to reshape input objects towards specified built material.

2. Related Work

Shape deformation and synthesis. Shape deformation deals with warping an input shape based on user provided positional specifications while trying to best maintain certain model properties. Such properties can preserve local geometric details (c.f., [BS08]); regularize deformations to be as-rigid-as-possible [SA07] or near-isometric [KMP07]; or conform to inter- and intra-part relations analyzed from the input models [GSMCO09]. In the context of content creation, there exist different methods to synthesize model variations starting from a collection of input shapes. For example, modeling can be performed by mixing-and-matching among model parts [FKS*04, KJS07, LVW*15]; sampling from a learned part-based probability distribution [CKGK11, KCKK12]; using an evolutionary algorithm to create model variations [XZCOC12]; exchanging compatible part substructures to create plausible model variations [ZCOM13]; or applying user-controlled group metamorphosis [San13]. More recently, algorithms have been proposed to minimally change input shapes such that the fabricated final shape is resilient and physically stable [SVB*12, UIM12, PWLSH13]. In contrast, we explore relations between material and form, and how they influence each other resulting in large form changes.

Fabrication-aware modeling. Rapid advances in accessible and economic fabrication possibilities have renewed interest in fabrication-aware modeling. Starting from an input shape and construction material, the goal of such methods is to best approximate the input shape while conforming to constraints arising due to fabrication methodology. Examples include incorporating curved folds to create freeform surfaces for folding single planar sheet of material [KFC*08]; rationalization of freeform surfaces with triangular, quad, cylindrical, and other primitive panels [EKS*10, FLHCO10, SS10]; designing cardboard chairs [SLMI11]; fabricate materials with target behavior [BBO*10]; generating fabricatable parts and connectors from an input wooden furniture model using a grammar-based method [LOMI11]; or adapting input models to facilitate construction from planar pieces [HBA12, SP13, CPMS14]. Schulz et al. [SSL*14] presented an interactive design system to create new fabricatable models based on predesigned templates in a database, while Koo et al. [KLY*14] focus on facilitating creation of works-like-prototypes from functional specifications. Unlike previous work, we aim at reforming shapes for material-aware fabrication. Our method can adapt a multi-component model to a different fabrication context (e.g., metal→wood, wood→metal) using pre-knowledge from a database. To the best of our knowledge, existing algorithms do not support this formulation.

3. Overview

Given a query model and user-specified target materials for individual parts, our goal is to use the query model as reference, and generate a fabricatable model subject to prescribed material constraints (see Fig. 1). The main idea is to synthesize new geometric forms based on example models in a database. More specifically, in a preprocessing step, we perform shape analysis for all the database models and the input

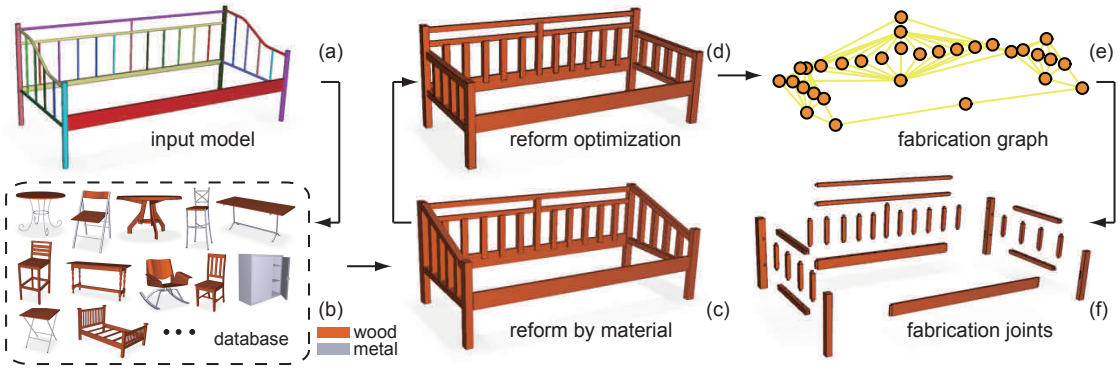

Figure 3: *Given an input model (a) and a database consisting of a set of models with labeled part material (b), our system can automatically reform the input model according to user-specified target material (wood in this case) by learning the material-geometry correlation from the database (c). Our system can also optimize the part configuration of the reformed model to facilitate fabrication in the target material context (d). The fabrication specifications (encoded in a fabrication graph, see Fig. 12) of the reformed model can also be inferred from our system (e) to facilitate fabrication (f).*

model, to extract geometric features relating to fabrication materials. In the reform process, we first extract new model parts from the database using the input model for guidance. This is based on optimizing the part similarity and pairwise similarity between the input model and the database models. However, due to material change, the new part configurations (e.g., the contact angle enclosed by neighboring parts) are revised in an optimization step, to facilitate fabrication in the target material context. We also infer fabrication specifications (e.g., woodworking joints) from the database, making the real fabrication feasible. The whole pipeline of our framework is illustrated in Fig. 3. The reformed part parameters are locally refined based on a FEM-based simulation testing, and if needed, are readjusted to ensure a durable design. This step often is optional as the data-driven designs mostly come with reasonable part parameter suggestions.

4. Algorithm

4.1. Preprocessing

Data preparation The input to our framework is a multi-component model $\mathcal{P}$ expressed as a set of parts $\{P_1, P_2, ..., P_N\}$. Each part is represented by a triangular mesh. If there is no pre-knowledge, a part segmentation based on triangle face connectivity is performed. The user may further divide and/or re-group parts if necessary. We support a database containing 152 multi-component models from the Google 3D Warehouse, which we manually tagged with fabrication material (wood/metal/other) for each part. The parts tagged with 'other' are ignored in the analysis like small parts that are not key elements of the model (e.g. screw, chair leg pad, etc.).

Part analysis For each model part, we compute an oriented bounding box (OBB) using principle component analysis of its mesh vertices. To overcome irregular triangulation, we further optimize the OBB by iteratively fixing one dimension and optimizing the other two dimensions using rotating calipers [Tou83]. We normalize each model by uniformly scaling according to the diagonal length of its OBB. We also find part thickness is closely related to fabrication material, and perform thickness estimation as follows. We first sample points (1000 in our experiments) on mesh surface. Starting from each sample point $\mathbf{s}_i$, we shoot a ray along its normal (i.e., the normal of the mesh face where the point resides), and find the first intersection point $\mathbf{t}_i$ that hits the part. The distance $d_i = \|\mathbf{s}_i - \mathbf{t}_i\|$ is defined as the thickness of sample point $\mathbf{s}_i$. We then perform a voting over all sample points. The distance value $d^\star$ that gets the most votes (within a margin of $\pm 10\%$) is defined as the part thickness t.

Contact analysis For a given model, we build a contact graph $G_c = (\mathcal{V}, \mathcal{E}_c)$ to encode the spatial relation among parts (see Fig. 4). Each node $v_i \in \mathcal{V}$ represents a single part. If the minimum distance between two parts, P_i and P_j, is less than a given threshold d_c, we add an edge $e_{ij} \in \mathcal{E}_c$ connecting the corresponding nodes in G_c. For two parts that are in contact, we also store a contact point $\mathbf{c}_{ij}$. We densely sample the two contact parts and extract nearby points (within d_c). The contact point is estimated by the barycenter of the nearby points.

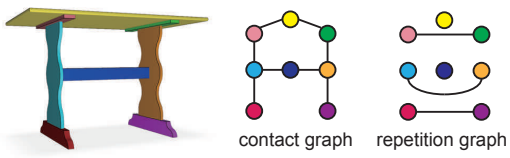

Figure 4: *A simple table model along with its contact graph and repetition graph.*

Repetition detection We also build a repetition graph G_r to encode the part repetitions (see Fig. 4). We detect congruent parts by aligning their OBB's and measuring root mean square distance. The repetition graph $G_r = (\mathcal{V}, \mathcal{E}_r)$ has the same nodes as G_c, while each edge $e_{ij} \in \mathcal{E}_r$ connects two congruent parts.

4.2. Evaluating Similarity

To formulate the connection between geometric form and fabrication material for shape reform, we correlate the query model with the example models using a set of similarity metrics. These metrics capture the similarity in terms of part geometry, part material, and the configuration of neighboring parts, which help to adapt to the new materials while respecting the original structure of the model.

Shape similarity We define shape similarity $\rho_{shp}(P_i, P_j)$ of two parts P_i and P_j based on part's OBB, area, and thickness. First, we measure the bounding box similarity as $\rho_{obb}(P_i, P_j) = \exp(-|\mathbf{b}_i - \mathbf{b}_j|^2/\sigma_{\mathbf{b}}^2)$, where $\mathbf{b}_i$ is a vector with 3 entries that represents P_i's size from its OBB. We sort the 3 entries and use $L1$ norm to measure the difference. Let $\sigma_{\mathbf{b}}$ be the standard deviation of bounding box difference that should be considered similar, so as other σ's defined in different contexts hereinafter. Second, suppose A_P is the area of part P and A_{obb} is the area of its OBB, we also compute an area ratio $r = A_P/A_{obb}$ of each part. The area ratio similarity is defined as $\rho_{area}(P_i, P_j) = \exp(-(r_i - r_j)^2/\sigma_r^2)$. We also define a thickness similarity as $\rho_{thick}(P_i, P_j) = \exp(-(t_i - t_j)^2/\sigma_t^2)$. Note that sophisticated geometric descriptors (e.g., shape distribution [OFCD02]) can alternatively be used, but may restrict suitable form changes when adapting to different materials.

Material similarity We simply define the material similarity between M_i and M_j ($M_i, M_j \in \{\text{wood, metal}\}$) as:

$$\rho_{mat}(M_i, M_j) = \begin{cases} 0, & \text{if } M_i \neq M_j \\ 1, & \text{if } M_i = M_j. \end{cases}$$

Spatial similarity Besides part similarity, we also define similarity for part pairs to correlate forms in a structural context. First, we consider the spatial relation between a pair of parts P_i and P_j. We use the Euclidean distance $d_{i,j}$ between part barycenters because it is independent of model orientations. The spatial similarity between two pairs of parts is defined as:

$$\rho_{pr}(d_{i,j}, d_{k,l}) = \exp(-\|d_{i,j} - d_{k,l}\|^2/\sigma_{pr}^2).$$

Contact angle similarity How to connect two parts and in what form they should contact largely depend on their fabrication material (see Fig. 2). An important observation is made for the contact angles between linear/curvilinear parts. Wooden parts usually contact in right angles, while metal parts can form flexible contact angles. This helps to better correlate material to object form. To estimate contact angle, we first identify whether the contact part is linear/curvilinear based on its OBB. If the OBB is elongated and the area ratio is near to 1, we use the maximal dimension as one side of the angle. Otherwise if the area ratio is small, we collect the sample points (those for detecting the contact point) in a local neighborhood centered at the contact point, and compute dominant principal direction to enclose the angle. In our implementation, the neighborhood size is 3 times the part thickness. If one of the two parts is not linear/curvilinear, we mark the contact angle as N/A. For simplicity, we do not separate acute and obtuse angles, i.e., $\alpha_{ij} \in [0, 90]^\circ$. The contact angle similarity is defined as:

$$\rho_{ca}(\alpha_{i,j}, \alpha_{k,l}) = \begin{cases} 1 & \alpha_{i,j}, \alpha_{k,l} \text{ N/A} \\ 0 & \alpha_{i,j} \in [0, 90]^\circ, \alpha_{k,l} \text{ N/A} \\ 0 & \alpha_{i,j} \text{ N/A}, \alpha_{k,l} \in [0, 90]^\circ \\ \exp(\frac{-(\alpha_{i,j} - \alpha_{k,l})^2}{\sigma_{ca}^2}) & \alpha_{i,j}, \alpha_{k,l} \in [0, 90]^\circ. \end{cases}$$

Please note that the contact angle similarity is 1 when two angles are N/A. This allows only counting other terms when both contact angles are not well defined.

In our experiments, plausible results are achieved with $\sigma_{\mathbf{b}} = 0.1, \sigma_r = 0.1, \sigma_t = 0.02, \sigma_{pr} = 0.2, \sigma_{ca} = 10$. More advanced weight learning strategy [LRFH13] can be used to further improve the results.

4.3. Reforming Shapes based on Target Material

Given a model $\mathcal{P}$ consisting of individual parts $\{P_1, P_2, ..., P_N\}$, the goal of shape reform is to adapt the geometry of the model when changing materials to $\{\bar{M}_1, \bar{M}_2, ..., \bar{M}_N\}$. The key idea is to use similar parts from the pre-tagged database $\{P_1^t, P_2^t, ..., P_K^t\}$ (with materials $\{M_1^t, M_2^t, ..., M_K^t\}$) to replace the model parts subject to material constraints. We formulate the shape reform problem as an optimization that jointly maximizes part similarity and pairwise similarity referring to the database.

Formulation For each part P_n with target material $\bar{M}_n$, we define a potential $\phi(P_n^t)$ to measure the probability of replacing P_n by P_n^t (with material M_n^t):

$$\bar{\phi}(P_n^t) = \rho_{mat}(\bar{M}_n, M_n^t)\tilde{\rho}_{shp}(P_n, P_n^t), \tag{1}$$

where $P_n^t \in \{P_1^t, P_2^t, ..., P_K^t\}$. Please note that here $\tilde{\rho}_{shp}(P_n, P_n^t) = \rho_{obb}(P_n, P_n^t)$, i.e., we only use OBB to measure the shape difference. Such relaxation allows suitable form changes to adapt to different materials. For example, a wooden board can be replaced by a curved metal part (see Fig. 6). If part compactness is required (e.g., cabinets), we further add ρ_{area} into $\tilde{\rho}_{shp}$.

For two parts P_i and P_j that are in contact, we define a pairwise potential $\bar{\psi}_c(P_i^t, P_j^t)$ to measure the probability of replacing P_i by P_i^t, and P_j by P_j^t. To do this, we compare with

all the contacting pair (P_u^t, P_v^t) in the training data set:

$$\bar{\psi}_c(P_i^t, P_j^t) = \sum_{(u,v)} \rho_{mat}(\bar{M}_i, M_u^t)\rho_{mat}(\bar{M}_j, M_v^t) \tilde{\rho}_{shp}(P_i^t, P_u^t)\tilde{\rho}_{shp}(P_j^t, P_v^t). \quad (2)$$

For two congruent parts P_i and P_j that are connected in the repetition graph, we define a pairwise potential $\psi_r(P_i^t, P_j^t)$ to encourage the same replacement part:

$$\bar{\psi}_r(P_i^t, P_j^t) = \delta_{ij}(P_i^t, P_j^t), \quad (3)$$

where $\delta_{ij}(P_i^t, P_j^t) = 1$ if and only if $P_i^t = P_j^t$, otherwise 0.

The overall potential of all part replacements $\mathcal{P}^t$ is defined as:

$$F(\mathcal{P}^t) = \prod_{P_n^t \in \mathcal{P}^t} \bar{\phi}(P_n^t) \prod_{(i,j)\in\mathcal{E}_c} \bar{\psi}_c(P_i^t, P_j^t)^{\alpha} \prod_{(i,j)\in\mathcal{E}_r} \bar{\psi}_r(P_i^t, P_j^t)^{\beta}, \quad (4)$$

where $\alpha = 0.1, \beta = 20$ are weighting parameters. We optimize the above multi-label assignment problem by loopy belief propagation. The implementation is based on the sum-product algorithm [KFL01]. In our implementation, we add a small ε for zero similarity values to avoid deterministic potentials and ensure convergence. Fig. 5 shows the effect of adding pairwise term. Note that the optimization is computationally prohibitive if we treat all parts in the database as replacement candidates. Instead, we use a subset by clustering geometrically similar parts as detailed in Sec. 5.

In general, one can randomly specify metal or wood for each part, which often leads to unrealistic results. A better solution is to perform material suggestion based on example models using a similar approach as in [JTRS12] (see Appendix A), and adapt model parts to different materials (i.e., metal→wood, wood→metal, see Fig. 6).

Restore contact After shape reform, we find for each part P_n a replacement part P_n^r. Then we scale P_n^r so that the dominant dimension of its OBB matches P_n. Note that we only scale the dominant dimension so that the thickness of P_n^r is not affected. We then perform an optimization that is similar to [KCKK12] to restore contacts on the reformed model.

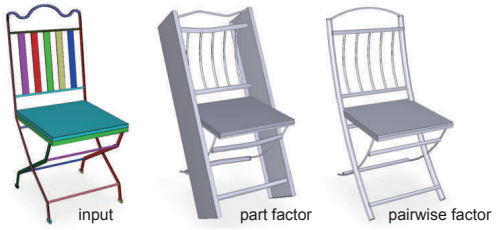

Figure 5: *Pairwise factor in the shape reform optimization is based on comparing neighboring configuration in a database. This helps to suggest right geometric configuration between parts.*

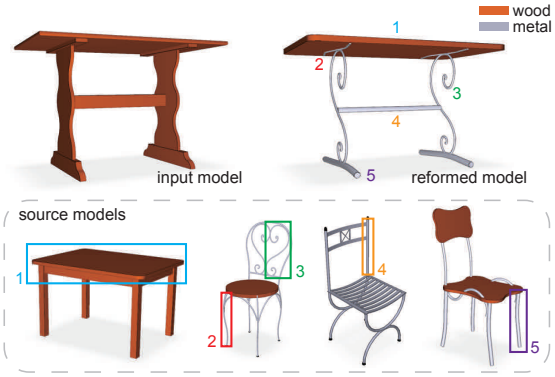

Figure 6: *A wooden table is reformed to have a metal structure based on example models in the database. Replacement parts are suggested using a data-driven optimization.*

The difference is that the contact points of the replacement parts are computed by projecting the original contact point $\mathbf{c}_{ij}$ to the relevant parts. Also, during the optimization, we only allow translation in the optimization to avoid thickness change due to scaling, which causes artifacts in the reformed results. Only translating parts can solve most cases except some rare hanging parts. Additional scaling of dominant dimension is performed as a post-processing. Fig. 7 shows the replacement parts before and after contact optimization.

4.4. Optimizing for Material-aware Part Configurations

Although we have considered pair-wise contact similarity in the material-aware shape reform stage, the spatial relation of neighboring parts still largely depends on the initial configuration. However, as mentioned before, how the parts should be manufactured and assembled is restricted by the fabrication material. Therefore, after reform, the updated model may have non-optimized structures which are not suitable for fabrication using the target material . Here we propose an example based optimization which can further improve the fabrication feasibility of the reformed model.

We use contact angle (see Sec. 4.2) to measure the fabrication feasibility. From the examples in the database, we build up histograms of contact angles between wooden parts and

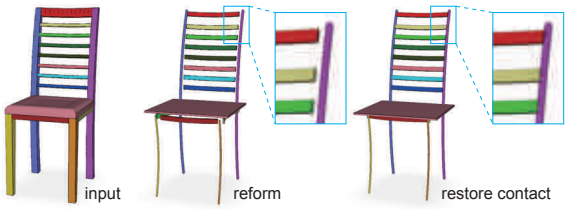

Figure 7: *We perform a contact optimization to jointly connect replacement parts.*

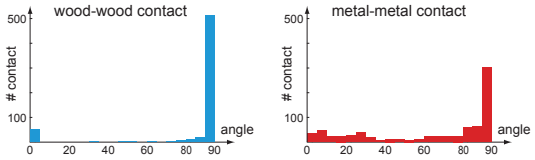

Figure 8: *We build up histograms of contact angles for wooden parts and metal parts respectively, to infer fabrication feasibility and guide the structure optimization.*

between metal parts respectively (see Fig. 8). This captures the possibility/difficulty of constructing different contact angles in practice. Then from the reformed model, we compute contact angle between parts and measure its feasibility from the histogram. If the feasibility is low (i.e., the number of hits in the histogram is less than 5), we mark the contact and specify its target angle from the nearby feasible angles in the histogram. Given all the infeasible contacts denoted by edge set $\mathcal{E}_{angle} \subset \mathcal{E}_c$ in the contact graph, each infeasible contact $e_{ij} \in \mathcal{E}_{angle}$ has two incident parts P_i and P_j, a contact angle θ_{ij}, and a target angle θ^*_{ij}. We would like to perform an optimization to relocate the contacting parts so that all contact angles become feasible:

$$\min \sum_{e_{ij} \in \mathcal{E}_{angle}} \delta_{ij} \cdot [\theta_{ij} - \theta^*_{ij}]^2, \text{ s.t. } \sum_i \delta_{ij} \geq 2 \quad (\mathbf{c}_{ik} \neq \mathbf{c}_{il}). \tag{5}$$

In the above equation, δ_{ij} is a binary variable with $\delta_{ij} = 1$ indicating P_i and P_j are in contact, otherwise 0; $\mathbf{c}_{i,j}$ is the contact point of P_i and P_j. The constraints ensure that there is no hanging part. We add an extra contact point for part who touches the ground (suppose the up-right orientation of the model is along z-axis).

Ideally, we should have $\delta_{ij} = 1 \ \ \forall e_{ij} \in \mathcal{E}_{angle}$ in the optimization. However, in practice it is not feasible to converge to all target angles. One simple case is shown in Fig. 9. This is mainly because metal allows much more contact freedom than wood. Instead, we relax the contacts by allowing $\delta_{ij} = 0$. On the other hand, to keep contact between P_i and P_j, we cannot simply set $\delta_{ij} = 1$, since δ_{ij} depends on the spatial relations of P_i and P_j. Instead, we ensure contact by adding geometric constraints that only allow P_i to slide on P_j (i.e., changing contact point), or rotate while keeping contact

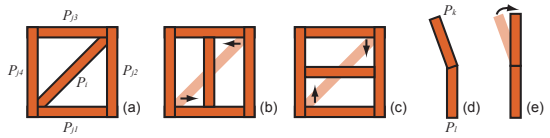

Figure 9: *To optimize the contact feasibility, we allow part to slide/rotate on other parts. (a) P_i cannot be orthogonal to all other four parts; (b), (c) by relaxing contact and allowing P_i to slide, two feasible configurations can be achieved; (d) P_k is allowed to rotate while fixing the contact point with P_l; (e) feasible configuration after rotation.*

point with P_j. Whether a part should slide or rotate depends on its relation with neighboring part. More specifically, rotate if the target angle is 0, otherwise slide.

To simplify the computation, we abstract each (elongated) part P_i by a line segment $s_i = (\mathbf{v}^s_i, \mathbf{v}^t_i)$ (using two end points along dominant dimension of its OBB). Given all the contact constraints, i.e., sliding pairs $\{e_{ij}\} \subset \mathcal{E}_{angle}$ (P_i slides on P_j), rotating pairs $\{e_{kl}\} \subset \mathcal{E}_{angle}$ (P_k rotates while contacting P_l), the formulation in Eqn. 5 can be interpreted as:

$$\begin{aligned}
\min & \sum_{e_{ij}} [\theta(s_i, s_j) - \theta^*_{ij}]^2 + \sum_{e_{kl}} [\theta(s_k, s_l) - \theta^*_{kl}]^2 \\
& + w_l \sum_{e_{kl}} [(\mathbf{v}^s_k - \mathbf{v}^t_k)^2 - (\hat{\mathbf{v}}^s_k - \hat{\mathbf{v}}^t_k)^2]^2 \\
& + w_r \sum_{(u,v)} \sum_{(m,n)} Repel(s_m, s_n), \\
& e_{mu}, e_{mv}, e_{nu}, e_{nv} \in \{e_{ij}\} \\
\text{s.t. } & \mathbf{v}^*_i = t_{ij}\mathbf{v}^s_j + (1 - t_{ij})\mathbf{v}^t_j, \quad t_{ij} \in [0,1], \quad \mathbf{v}^*_k = \mathbf{c}_{kl},
\end{aligned} \tag{6}$$

where $* \in \{s,t\}$ indicating which end of the part should slide/rotate, can be inferred from the initial configuration. We also add length preserving term for rotating parts. If P_m and P_n both slide on P_u and P_v, a repulsion term $Repel(s_m, s_n) = \exp[-(\mathbf{v}^s_m - \mathbf{v}^s_n)^2/\sigma^2] \cdot \exp[-(\mathbf{v}^t_m - \mathbf{v}^t_n)^2/\sigma^2]$ is added to avoid part overlap.

In our implementation, we enumerate all the slide/rotate possibilities, specify the corresponding constraints, and perform quasi Newton method (L-BFGS) to solve Eqn. 6 (with $w_l = 1, w_r = 0.1, \sigma = 0.05$). The solution with minimal objective function value is selected as the optimal result. The

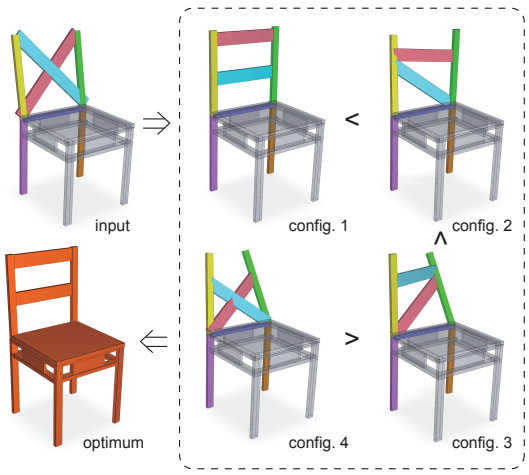

Figure 10: *Material-specific optimization on a reformed chair model. We use contact angle histogram to detect non-optimal contact angles. The incident parts are shown in dark colors. We optimize the angle differences between parts under different contact constraints, resulting in different configurations. We sort the configurations based on contact angles and select the optimum one.*

model is updated accordingly by aligning parts to the corresponding line segments.

To initialize the optimization, we fix the parts that have no contact part at one of its end, or have symmetric parts with no angle problem. We also restrict two ends of a part to slide on non-contacting parts. This allows the optimization to run faster, and explore more interesting variations with topology and geometry change (see Fig. 10).

4.5. Generating Fabrication Specifications

So far we have generated the geometry of individual parts and optimized their configurations according to target materials. In practice, how to assemble multiple parts to form a real functional object is a non-trivial task [Hyl08]. It not only depends on the materials of neighboring parts (e.g., metal parts are connected by welding, wooden parts are connected by specific wooden joints), but also their geometry and spatial correlation (see Fig. 11). In this section, we would like to further investigate the fabrication specifications in terms of how neighboring parts are conjoined. The goal is to infer joint type based on part geometry, and further refine its geometry (for wooden parts) to ensure the overall assembly.

Fabrication type inference The actual fabrication of a multi-component object requires different conjoining methods to assemble neighboring parts. Fig. 11 shows several representative joint types that are resolved in our framework. These joint types are classified into three categories (i.e., wood-wood, wood-metal and metal-metal) based on the material configuration of the neighboring parts. For each category, different fabrication techniques can further be applied to assemble parts with different characteristics.

Similar to shape reform, we also perform a data-driven approach to infer different joint types for the reformed model. Namely, we derive joint type by comparing with example models with pre-tagged joints. Suppose P_i and P_j are the incident parts of the query joint, the similarity function of assigning the same type of joint as P_u^t and P_v^t is expressed as:

$$\phi_{i,j;u,v} = \rho_{mat}(M_i, M_u^t)\rho_{mat}(M_j, M_v^t) \tilde{\rho}_{shp}(P_i, P_u^t)\tilde{\rho}_{shp}(P_j, P_v^t)\rho_{pr}(d_{i,j}, d_{u,v})\rho_{oa}(\mathbf{a}_{i,j}, \mathbf{a}_{u,v}), \quad (7)$$

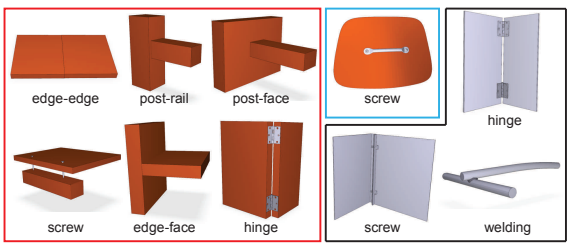

Figure 11: *Representative joint types of wood-wood (red box), wood-metal (blue box) and metal-metal (black box).*

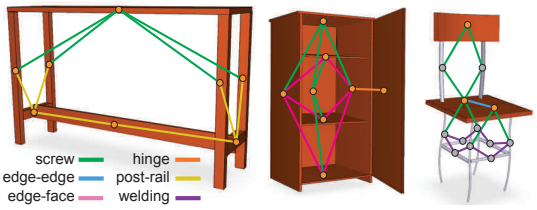

Figure 12: *Different fabrication joint types (shown as graph edges) are inferred using a data-driven approach.*

where $\mathbf{a}_{i,j}$ is a 9D vector that measuring angles ($\in [0,90]^\circ$) between individual principal directions of two OBB's. Specifically, suppose the principle directions of P_i is $\{X_{i0}, X_{i1}, X_{i2}\}$ (in order), then $\mathbf{a}_{i,j}$ is the vector representation of the 3×3 matrix $\{\angle(X_{ik}, X_{jl})\}_{k,l=0,1,2}$ and $\rho_{oa}(\mathbf{a}_{i,j}, \mathbf{a}_{u,v}) = \exp(-|\mathbf{a}_{i,j} - \mathbf{a}_{u,v}|^2/\sigma_{oa}^2)$, where $\sigma_{oa} = 360^\circ$ in our tests. The definition of the other terms can be found in Sec. 4.2.

For joint assignment, we classify all pre-tagged joints into multiple clusters based on the specified joint types. The type of the query joint is inferred by its k nearest neighbors (k=3 in our implementation) among the pre-tagged models from different clusters (see Fig. 12). Please note that our joint assignment is entirely geometry-based and does not rely on any semantic information of the models/parts (up-right/front orientation, part labels, etc.). For functional parts, this may lead to ambiguities. For example, a closed cabinet door can not be separated from cabinet back. Then the user is expected to clarify such ambiguities. After the assignment, for neighboring parts connected by edge-edge, post-rail, post-face, edge-face joints, we further verify whether the part is a tenon or mortise part (tenon tongue or mortise hole will be generated accordingly, see Fig. 14) based on pairwise relations (e.g., the contact location) for later processing.

Fabrication-aware part refinement Based on the inferred joint types, we further seek to create tangible joint shapes by refining neighboring wooden parts. However, the dimensions of the neighboring parts may not be compatible after shape reform (see inset), making the actual fabrication unrealizable. We thus perform a proxy-based deformation [ZFCO*11] to resize neighboring parts and ensure the right configuration. We use part's OBB as proxy and solve for the optimal dimensions of the OBB (the centroid is kept fixed). We optimize the closeness to the original OBB's subject to compatibility constraints (e.g., the tenon part should lie in the mortise part along major axis) on neighboring OBB's. This results in a quadratic function (in terms of OBB dimensions) with linear constraints, which can

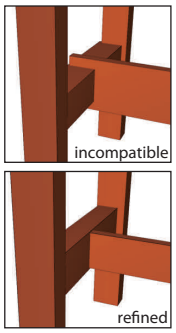

be explicitly solved by Quadratic Programming. After optimization, all the related parts are updated (scaled) according to the OBB's. Please note that for part refinement and the subsequent joint formation, we only focus on wooden parts, since metal parts can be easily conjoined in real fabrication, where additional nails, hinges, filler material are used.

Structural validation After part refinement, we perform standard design-simulation strategy to validate the functionality of the reformed model. In the current solution, we directly employ structural analysis on the whole model by merging neighboring parts with the same material. After adding boundary constraints (i.e., fixed constraints, boundary load), we use the finite element method in [COM15] to compute the stress map over the entire model. If the imposed stress is beyond the yielding stress of the specified material, we further increase the thickness of the weak parts and refine the other parts in the reformed model. The coupled simulation and refinement is iteratively performed (1-2 steps in our test) to generate physically stable shapes (see Fig. 13). In our test, the yielding stress is 8.6MPa for wood and 200MPa for metal. Empirically, we found this mode is more used as a sanity check since the data-driven part suggestions are often good to start with.

Forming joints With the right joint type and part geometry, we form joint shapes by sculpting corresponding wooden parts to ensure the assembly after fabrication. This is based on several simple CSG operations. First, we use mortise part to subtract tenon part and detect the contact face on the tenon part. Second, we construct the tenon geometry by scaling the contact face and extending to the mortise part. Different scaling factors are applied to generate different physical linkages, such as blind tenon for post-rail, half-blind tenon for edge-edge, through tenon for edge-face. Finally, the tenon/mortise part is sculpted by adding/subtracting the tenon geometry (see Fig. 14).

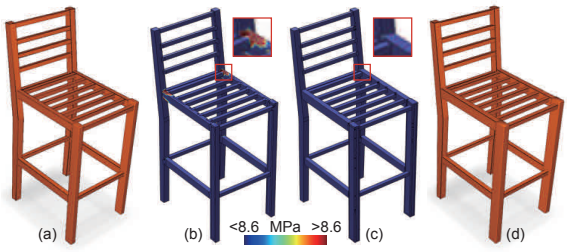

Figure 13: *Coupled structural validation and shape refinement ensures physically stable reformed shapes. (a) The input model for structural validation. (b) Structural analysis of the input model. The region in which the stress is beyond the yielding stress is shown in red. (c) By enlarging the thickness of the weak parts and part refinement, a physically stable model is obtained as in (d).*

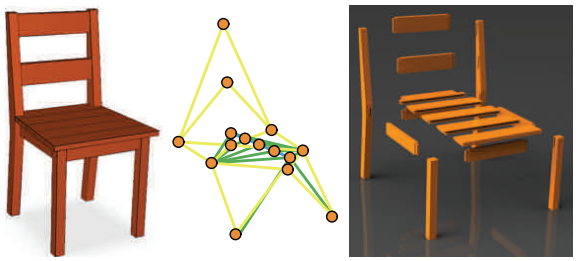

Figure 14: *Fabrication joints can be created based on the right joint type and part geometry. Please refer to Fig. 12 for explanation of edge-color convention.*

5. Evaluation

In this section, we evaluate the proposed framework. First, we test our algorithm using different settings on a large number of query models. Then we explore several reform variations. Further, we show the conducted user study and the statistics of our framework. All results are generated with the whole database unless mentioned otherwise. The database models are mixed of chairs (83 models, 1304 parts), tables (32 models, 350 parts), beds (19 models, 588 parts) and cabinets (18 models, 221 parts).

Results We test our material-aware reform algorithm on a number of models (see Fig. 15). We present reform results for the same input models to have only wooden/metal parts. To highlight the generality of our framework, the material context of the input is not inferred here.

The algorithm produced plausible reformed models adapted to the target material constraints. For example, in 'to metal' cases, straight bar become curved arc, flat boards become planar snakes, the thickness of the reformed parts is small (opposite effects can be observed in 'to wood' case). All these patterns conform to the correlations between built material and geometric forms (see Fig. 2). If the target material happens to be the underlying material of the input model, the resultant model usually has a comparable structure with the input model, while exhibiting feasible geometric variations. Note that as marked by the blue box (also see other results in Fig. 1, Fig. 3), our framework can also optimize the inter-part topology of the reformed model, to facilitate fabrication in the target material context. For example, part connections are broken at the chair back and leg, so that wooden parts can be optimized to meet at near 90°. On the other hand, we also find some challenging case as shown in red box. This is not only caused by inappropriate modeling/segmentation of the input table legs, but also the complicated spatial configurations of the output. A small amount of user interactions may further be required for such cases.

Besides replacing and optimizing individual parts, fabrication joints on the reformed models are also inferred from the database (see Fig. 15). Our data-driven approach can resolve

Figure 15: *Material-aware shape reform on a number of models with target material set to wood and metal, respectively. The fabrication joint inference result is shown next to the reformed model. Ambiguous joints are highlighted by dash lines.*

Figure 16: *Given an input model (left), multiple reform results can be generated from different example dataset.*

most of the joint types. Ambiguous joints are mainly caused by the lack of shape semantics. For example, for metal bed, bed top and bed body need to be screwed to ease assembly/disassembly.

User control One option for shape reform is to use material suggestion (see Appendix A) to infer the current material context, and adapt model parts to different materials. Statistical numbers are discussed later. The rest reform results in the paper, including mixed target materials, are generated referring to the current material context.

Given a single input model and the target material of each part, material-specific reform variations can be achieved by exploring different subspaces in the database. Fig. 16 shows several reformed metal chairs from the same wooden chair. Each reformed model is generated by randomly selecting 20 chairs in the database as a new example dataset.

Fig. 17 shows a shape reform result across different type of models in the database. Note that the bed board cannot find a compatible part if only search in a chair database. Stretching a chair seat to be a bed board results in large distortion. On the other hand, forms with the same material share common features across different type of models. For example, compatible metal parts can be retrieved from a chair dataset to form a plausible metal bed frame.

User study We conducted a user study to verify if our material-aware reform results coincide with the human experience of correlating shape and built material. We showed 35 material-aware shape reform results (the test models are randomly picked from our database, see supplemental material) to 30 computer science and EE students with varied

Figure 17: *Given a bed with all wooden parts (left), we can reform for mixed materials, e.g., the bed frame adapts to metal and the bed board adapts to wood. The reformed result from the whole database (middle). The reformed result from only the chair dataset (right).*

Figure 18: *It is not easy to recognize material from geometric shape alone if material characteristics cannot be inferred from individual parts and part configurations (e.g., contact angles). Left: Metal parts of the model are mainly straight and most of the contact angles are 90 degrees. It accords with typical wooden chairs. Right: Wooden parts of the table are thick, while the overall structure being similar to a typical metal table.*

background. For each reformed model, the users were asked to infer its major built material from the geometry alone.

The average hit rate was 91.5% indicating that human expectation generally agreed with target material and geometric form correlation. Fig. 18 shows two typical failure cases where users' choice did not match the specific material.

Statistics Cross validation of the material suggestion is tested on different training data sets. We randomly sample different subsets (with ratio 0.2, 0.4, 0.6, 0.8) from the database as training data and use the rest as testing data. For each ratio, we run 10 times and select the top 4 material suggestion accuracy, then compute the average. The statistics is shown in Fig. 19 (left). We also perform fabrication type inference on all the models in the database. The query model itself is leaving out in the training set during the inference process. The inference accuracy numbers of individual categories are summarized in Fig. 19 (right).

Performance Our experimental platform is with a 2.66 GHz Intel Xeon X5550 CPU. The training of the database is computed offline. Since we only have two candidate materials (metal/wood), it only takes 4.6 seconds to assign materials for a model with 27 parts. For material-aware reform, if we

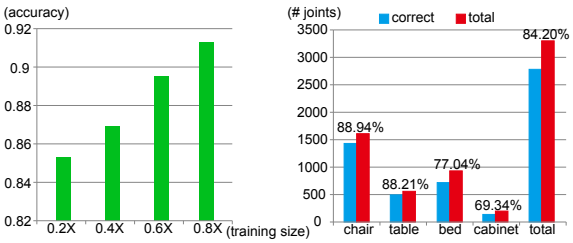

Figure 19: *Left: the accuracy of predicting fabrication materials using different proportions of the data base as training set. Right: the accuracy statistics of predicting fabrication joint types.*

consider all parts in the database to be candidates, the optimization is very expensive. We handle this problem in a preprocessing step. First, we filter out congruent parts (keeping only one instance). Then we cluster the rest parts into 80 groups by k-means clustering using feature vector of each part as its OBB dimension, area ratio and thickness (5D in total). For each cluster, the part with the smallest distance to the cluster center is selected as a candidate part. The reform based on clustered candidate parts can be computed efficiently, it takes 15.3s to reform the same model with 27 parts. The fabrication inference takes 3.9s for 55 contacts of the same model. The computational time of material suggestion and material-aware reform mainly depends on the number of parts of the input model, while for fabrication inference the dominant factor is the number of contacts.

6. Conclusion

We presented a data-driven algorithm that reforms a component-based input shape such that the reformed shape is better suited for fabrication using the target build material. We formulated this as an optimization that not only selects appropriate parts from the database, scales and positions them appropriately, but searches over non-trivial topological changes to ensure that the reformed shape conform to material-specific angle distributions. Finally, part connection types are inferred and necessary geometric modifications are suggested. We validated the algorithm on various models for wood and metal constructions, and evaluated the results using a user study that the classification results and reformed shapes match human perception of material just based on geometric shape.

Our reform algorithm is based on a part replacement strategy. Since different materials result in different part shaping abilities, if one part cannot find a geometrically similar part when adapting to a different target material, implausible result would be generated. Also, in this work we only focused on wood and metal materials. In the future, it would be interesting to extend the framework to also handle plastic and molded sheets. However, the challenge then would be to obtain initial parts since the input geometric meshes do not necessarily conform to material specific partitioning. Besides, our algorithm does not assume access to part labels (table leg, chair seat, etc.). As a result parts can potentially undergo large deformations, say a chair leg can get stretched to become a bar for the bed frame. In reality, however, materials have limits on maximum dimensions. It is desirable for the algorithm to take this into account when making part suggestions. Finally, it is worth mentioning that if the inferred material context (by material suggestion) of the input model matches the target material constraints, we could skip Sec. 4.3~4.4, and apply the rest algorithms to generate result.

Acknowledgements. We thank the reviewers for their comments and suggestions for improving the paper; Cristina Amati for help with the chair fabrication. This work was supported in part by University of Bath startup fund BA-FS6SYY, KAUST baseline funding, the ERC Starting Grant SmartGeometry (StG-2013-335373), and Marie Curie CIG.

Appendix A: Material Suggestion

Given a model $\mathcal{P}$ consisting of individual parts $\{P_1, P_2, ..., P_N\}$, in an optional initialization step, fabrication material $M_n \in \{\text{metal, wood}\}$ is assigned for each part $P_n, (1 \leq n \leq N)$. This helps to adapt model parts to different materials in the shape reform stage and is available to the user as an optional suggestion.

Formulation For each part P_n, we define a potential $\phi(M_n)$ to measure the probability of assigning material M_n, by comparing with all the parts in the training set $\{P_1^t, P_2^t, ..., P_K^t\}$ with materials $\{M_1^t, M_2^t, ..., M_K^t\}$:

$$\phi(M_n) = \sum_{k=1}^{K} \rho_{mat}(M_n, M_k^t)\rho_{shp}(P_n, P_k^t), \tag{8}$$

where $\rho_{shp}(P_n, P_k^t)$ is a combination of ρ_{obb}, ρ_{area} and ρ_{thick}.

For two parts P_i and P_j that are in contact, we define a pairwise potential $\psi_c(M_i, M_j)$ to measure the probability of assigning material M_i to P_i, and M_j to P_j, by comparing with all the contacting pairs (P_u^t, P_v^t) in the training data set:

$$\begin{aligned}\psi_c(M_i, M_j) = \sum_{(u,v)} & \rho_{mat}(M_i, M_u^t)\rho_{mat}(M_j, M_v^t) \\ & \rho_{shp}(P_i, P_u^t)\rho_{shp}(P_j, P_v^t) \\ & \rho_{pr}(d_{i,j}, d_{u,v})\rho_{ca}(\alpha_{i,j}, \alpha_{u,v}).\end{aligned} \tag{9}$$

For two congruent parts P_i and P_j, we define a pairwise potential $\psi_r(M_i, M_j)$ to encourage the same material:

$$\psi_r(M_i, M_j) = \rho_{mat}(M_i, M_j). \tag{10}$$

For the whole model, the potential of the material suggestion M is defined as:

$$F(M) = \prod_{M_n \in M} \phi(M_n) \prod_{(i,j)\in\mathcal{E}_c} \psi_c(M_i, M_j)^{\alpha} \prod_{(i,j)\in\mathcal{E}_r} \psi_r(M_i, M_j)^{\beta}, \tag{11}$$

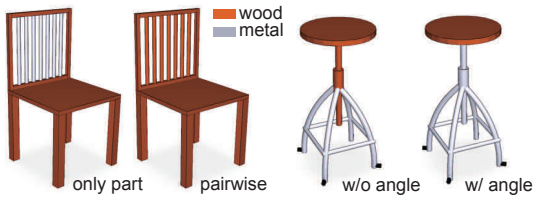

Figure 20: *(Left) Pairwise factor, which is based on comparing neighboring configuration in a database, helps to suggest right material configuration between parts. (Right) Additionally, inter-part angle similarity can help to make more appropriate material assignment to parts that are easier to fabricate.*

where $\alpha = 0.1, \beta = 20$ are weighting parameters. The optimal material suggestion is solved by loopy belief propagation. Fig. 20 shows the effect of pairwise factor and angle similarity.

References

[BBO*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. *ACM TOG (Siggraph) 29*, 3 (2010), 63:1–63:10. 2

[BS08] BOTSCH M., SORKINE O.: On linear variational surface deformation methods. *IEEE Trans. on Vis. and Comp. Graphics 14*, 1 (2008), 213–230. 2

[CKGK11] CHAUDHURI S., KALOGERAKIS E., GUIBAS L., KOLTUN V.: Probabilistic reasoning for assembly-based 3d modeling. *ACM TOG (Siggraph) 30*, 4 (2011), 35:1–35:10. 2

[COM15] COMSOL: http://www.uk.comsol.com/. 8

[CPMS14] CIGNONI P., PIETRONI N., MALOMO L., SCOPIGNO R.: Field-aligned mesh joinery. *ACM TOG 33*, 1 (2014), 11:1–11:12. 2

[EKS*10] EIGENSATZ M., KILIAN M., SCHIFTNER A., MITRA N. J., POTTMANN H., PAULY M.: Paneling architectural freeform surfaces. *ACM TOG (Siggraph) 29*, 4 (2010), 45:1–45:10. 2

[FKS*04] FUNKHOUSER T., KAZHDAN M., SHILANE P., MIN P., KIEFER W., TAL A., RUSINKIEWICZ S., DOBKIN D.: Modeling by example. *ACM TOG (Siggraph) 23*, 3 (2004), 652–663. 2

[FLHCO10] FU C.-W., LAI C.-F., HE Y., COHEN-OR D.: K-set tilable surfaces. *ACM TOG (Siggraph) 29*, 4 (2010), 44:1–44:6. 2

[GSMCO09] GAL R., SORKINE O., MITRA N. J., COHEN-OR D.: iwires: An analyze-and-edit approach to shape manipulation. *ACM TOG (Siggraph) 28*, 3 (2009), #33, 1–10. 2

[HBA12] HILDEBRAND K., BICKEL B., ALEXA M.: crdbrd: Shape fabrication by sliding planar slices. *Computer Graphics Forum (Proc. EUROGRAPHICS) 31*, 2 (2012), 583–592. 2

[Hyl08] HYLTON B.: *Illustrated Cabinetmaking: How to Design and Construct Furniture That Works*. Fox Chapel Publishing, 2008. 7

[JTRS12] JAIN A., THORMÄHLEN T., RITSCHEL T., SEIDEL H.-P.: Material memex: Automatic material suggestions for 3d objects. *ACM TOG 31*, 6 (2012), 143:1–143:8. 5

[KCKK12] KALOGERAKIS E., CHAUDHURI S., KOLLER D., KOLTUN V.: A probabilistic model for component-based shape synthesis. *ACM TOG (Siggraph) 31*, 4 (July 2012), 55:1–55:11. 2, 5

[KFC*08] KILIAN M., FLÖRY S., CHEN Z., MITRA N. J., SHEFFER A., POTTMANN H.: Curved folding. *ACM TOG (Siggraph) 27*, 3 (2008), #75, 1–9. 2

[KFL01] KSCHISCHANG F. R., FREY B. J., LOELIGER H. A.: Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theor. 47*, 2 (2001), 498–519. 5

[KJS07] KREAVOY V., JULIUS D., SHEFFER A.: Model composition from interchangeable components. In *Pacific Graphics Proceedings* (2007), pp. 129–138. 2

[KLY*14] KOO B., LI W., YAO J., AGRAWALA M., MITRA N. J.: Creating works-like prototypes of mechanical objects. 2

[KMP07] KILIAN M., MITRA N. J., POTTMANN H.: Geometric modeling in shape space. *ACM TOG (Siggraph) 26*, 3 (2007), #64, 1–8. 2

[LOMI11] LAU M., OHGAWARA A., MITANI J., IGARASHI T.: Converting 3d furniture models to fabricatable parts and connectors. *ACM TOG (Siggraph) 30*, 4 (2011), 85:1–85:6. 2

[LRFH13] LIN S., RITCHIE D., FISHER M., HANRAHAN P.: Probabilistic color-by-numbers: Suggesting pattern colorizations using factor graphs. *ACM TOG (Siggraph) 32*, 4 (2013), 37:1–37:12. 4

[LSH*10] LI X.-Y., SHEN C.-H., HUANG S.-S., JU T., HU S.-M.: Popup: Automatic paper architectures from 3d models. *ACM TOG (Siggraph) 29*, 4 (2010), 111:1–111:9. 2

[LVW*15] LIU H., VIMONT U., WAND M., CANI M.-P., HAHMANN S., ROHMER D., MITRA N. J.: Replaceable substructures for efficient part-based modeling. *Computer Graphics Forum (Proc. EUROGRAPHICS)* (2015). 2

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM TOG 21*, 4 (2002), 807–832. 4

[PWLSH13] PRÉVOST R., WHITING E., LEFEBVRE S., SORKINE-HORNUNG O.: Make it stand: Balancing shapes for 3d fabrication. *ACM TOG (Siggraph) 32*, 4 (2013), 81:1–81:10. 2

[SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Proc. of Symp. of Geometry Processing* (2007), pp. 109–116. 2

[San13] Morphological shape generation through user-controlled group metamorphosis. *Computers & Graphics 37*, 6 (2013), 620 – 627. 2

[SLMI11] SAUL G., LAU M., MITANI J., IGARASHI T.: Sketchchair: An all-in-one chair design system for end users. In *Proc. TEI* (2011), pp. 73–80. 2

[SP13] SCHWARTZBURG Y., PAULY M.: Fabrication-aware design with intersecting planar pieces. *Computer Graphics Forum 32*, 2pt3 (2013), 317–326. 2

[SS10] SINGH M., SCHAEFER S.: Triangle surfaces with discrete equivalence classes. *ACM TOG (Siggraph) 29*, 4 (2010), 46:1–46:7. 2

[SSL*14] SCHULZ A., SHAMIR A., LEVIN D. I. W., SITTHI-AMORN P., MATUSIK W.: Design and fabrication by example. *ACM TOG (Siggraph) 33*, 4 (2014), 62:1–62:11. 2

[SVB*12] STAVA O., VANEK J., BENES B., CARR N., MĚCH R.: Stress relief: Improving structural strength of 3d printable objects. *ACM TOG (Siggraph) 31*, 4 (2012), 48:1–48:11. 2

[Tou83] TOUSSAINT G.: Solving geometric problems with the rotating calipers. In *Proceedings of IEEE Melecon* (1983), pp. A10.02/1–4. 3

[UIM12] UMETANI N., IGARASHI T., MITRA N. J.: Guided exploration of physically valid shapes for furniture design. *ACM TOG (Siggraph) 31*, 4 (2012), 86:1–86:11. 2

[XZCOC12] XU K., ZHANG H., COHEN-OR D., CHEN B.: Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM TOG (Siggraph) 31*, 4 (2012), 57:1–57:10. 2

[ZCOM13] ZHENG Y., COHEN-OR D., MITRA N. J.: Smart variations: Functional substructures for part compatibility. *Computer Graphics Forum (Proc. EUROGRAPHICS) 32*, 2pt2 (2013), 195–204. 2

[ZFCO*11] ZHENG Y., FU H., COHEN-OR D., AU O. K.-C., TAI C.-L.: Component-wise controllers for structure-preserving shape manipulation. *Computer Graphics Forum (Proc. EUROGRAPHICS) 30*, 2 (2011), 563–572. 7