

Ergonomics-inspired Reshaping and Exploration of Collections of Models

Youyi Zheng

Han Liu

Julie Dorsey

Niloy J. Mitra

Abstract—This paper examines the following question: given a collection of man-made shapes, e.g., chairs, can we effectively explore and rank the shapes with respect to a given human body – in terms of how well a candidate shape fits the specified human body? Answering this question requires identifying *which* shapes are more suitable for a prescribed body, and *how* to alter the input geometry to better fit the shapes to a given human body. The problem links physical proportions of the human body and its interaction with object geometry, which is often expressed as ergonomics guidelines. We present an interactive system that allows users to explore shapes using different avatar poses, while, at the same time providing interactive previews of how to alter the shapes to fit the user-specified body and pose. We achieve this by first constructing a fuzzy shape-to-body map from the ergonomic guidelines to multi-contacts geometric constraints; and then, proposing a novel contact-preserving deformation paradigm to realize a reshaping to adapt the input shape. We evaluate our method on collections of models from different categories and validate the results through a user study.

Index Terms—geometric deformation, ergonomics, shape analysis, reshaping

1 INTRODUCTION

Humans come in various shapes and sizes. The field of ergonomics focuses on accommodating such human variations with design goals applied to functional objects. Various objects in everyday use are shaped and given form based on their intended use and target user. Guidelines on ergonomics (e.g., [1] and [2]) summarize years of such research, prototyping, product reviews, design experiences, and user feedback to provide recommendations for geometric shapes based on their target usage (see Figure 1). For example, Panero and Zelnik [1] summarize multiple design studies to prescribe good dimensions for typical household objects calibrated to human body size.

In computer graphics, object geometries are either acquired directly from, or modeled after real objects. Hence, such objects often inherit or mimic real world object specifications. In this work, we first investigate if one can classify and rank objects directly based on their target functions and associated ergonomic considerations specified in the form of target avatar poses. In contrast to typical classification strategies based on geometric descriptors, we present a novel view on *ergonomics-inspired* categorization, ordering, and exploration of unorganized shape collections.

Moreover, recent advances in digital fabrication have expanded the range of affordable, custom design possibilities. As geometric models can now be

easily fabricated, incorporating their target use is an important consideration. Since robust solutions exist to digitally capture high quality rigged avatars in a matter of minutes [3], as a second question, we investigate how to use such human avatars specified in intended-use poses to reshape the geometry of virtual objects. This leads to an *ergonomics-inspired* geometric reshaping of existing shapes. By *reshaping*, we refer to adapting both the part proportions and their relative arrangements so the new shape better conforms to the target usage. For example, Figure 2 shows results of classifying and reshaping a set of chair models based on different human avatar poses.




Fig. 1. Ergonomic guidelines linking human posture to geometry of a chair (left) and a bike (right).

- Y. Zheng and J. Dorsey are with Yale University.
E-mail: ckcbeebee@gmail.com; julie.dorsey@yale.edu
- H. Liu is with KAIST.
E-mail: han.liu@kaist.edu.sa
- N. Mitra is with University College London.
E-mail: n.mitra@cs.ucl.ac.uk




Fig. 2. Starting from (a) a collection of shapes and a few user-specified human shapes, (b) our method automatically deforms the models to fit the human shapes, and ranks the models using their deformation costs to produce a categorization of the input models with respect to their suitability to different human shapes (c).

We address both questions using a novel contact-based deformation paradigm. First, we map ergonomic guidelines to a set of contact specifications between the human avatar and input shape. Such guidelines often come in the form of multiple specifications, making it nearly impossible for lay users to satisfy by manual adjustment. For example, craftsmen [4] may make as many as twelve adjustments to a basic chair design (e.g., adjusting the headboard, lumbar support, or the angle between seat and back); while, naive users can struggle to appropriately adjust even three degrees of manual controls [5]. We propose algorithms to automate these steps. Starting from part-based input shapes, we first map the ergonomic specifications to a set of multi-contact constraints. We then, in absence of parameterized model templates, propose a novel reshaping algorithm to computationally adapt an input shape to simultaneously satisfy these constraints. We propose an iterative optimization that alternately deforms the input shape and revalidates the target contact specifications.

Our work makes the following contributions: first, we link human body and poses to shape deformation with the goal to facilitate personal customization; second, we provide a pose-driven shape exploration tool that subsequently enables a novel shape classification setup; third, we design a novel contact-based edit propagation algorithm to simultaneously accommodate ergonomics constraints for reshaping multiple objects according to a target human body. We evaluate our algorithm in the context of chairs, bikes, beds, and gym equipments and validate the results through a user study.

2 RELATED WORK

Our work is closely related to prior work on shape exploration, geometric classification, and research in human factors.

Shape exploration. A number of approaches have been described on shape exploration. While such tools allow users to quickly browse large data collections, a primary goal is to find effective embeddings where interactive exploration is made possible [6], [7], [8], [9]. These methods normally start from object geometry alone and extract commonalities among a family of shapes using mechanisms such as geometric descriptors or functional maps. Our method couples ergonomic guidelines with geometric deformation to measure inter- and intra- shape similarity.

Shape functional analysis. Our method is related to shape functionality analysis. Functionality is often related to shape semantics. While such cues are challenging to infer only from geometry, taking advantage of some prior knowledge would make the problem tractable. For example, previous methods leverage extrinsic shape structure (e.g., meaningful shape segmentation) to relate intra- component level interactions with functional partitions [10], [11], [12]. Notably, in a parallel work, Kim et al. [13] leverage pre-defined skeleton shapes to search for body-to-shape contacts and deformations and use such connections to reveal shape functionality. Our approach is complementary as we study the problem of shape-to-body deformation for reshaping an object, and subsequently use the deformation cues for geometric exploration and classification.

Geometry and human factor. In computer vision, geometry reasoning has been studied and used for human workspace reasoning [14] or indoor scene understanding [15]. These methods often use extensive training of existing image data to reason about an input image, and the process of human measurements is typically performed offline. Coupling geometric modeling with physical simulation has also been studied in recent research work of interactive chair modeling [16]. For an extensive discussion of history and

development of ergonomics in the context of chairs we recommend Cranz’s book [17], especially Chapters 3 and 5. As examples of linking geometric form with function, Shen et al. [18] utilize object structure for effective object assembly, while the AttriIt [19] system links semantic tags to object parts and their coupled variations.

Shape deformation. Shape deformation is a long-standing topic in geometry processing [20]. The shape deformation methods mainly fall into two classes: those that aim to preserve the shape’s local properties, such as curvature, differential coordinates [21], [22], [23], and those that aim to preserve global structures, such as symmetry and inter-part relations [24], [25], [26]. Li et al. [27] deform input man-made objects to make them amenable to stacking. In contrast, we focus on ergonomics guidelines for customizing man-made objects to target specific human body sizes and poses.

3 ERGONOMICS GUIDELINES

Decades of meticulous user studies along with information gleaned from direct observations are available in the form of qualitative ergonomics guidelines. In the following, we focus on chairs for a clear description. Although there is still some debate regarding the relative importance of the various guidelines, we summarize here those commonly described across the different reference works we consulted [28], [17], [29]. For a more general reference, please see [1].

- *Chair seats should have correct height* to allow both feet to be fully supported (e.g., by the ground). A chair that is too high creates undue pressure at the knee/thigh; while if it is too short, forces the knee to be higher than the hip sockets.
- *Width and depth of chair seats* should conform to the user’s dimensions. Specifically, while the width is dictated by the avatar’s waistline, the depth is dictated by the length of the avatar’s thigh bones.
- *Flat uncontoured seats* are preferred to discourage a slouched or C-shaped posture.
- *Lumbar support* by providing low- or mid-back support can help hold good posture and prevent pain to the spine and neck.
- *Head support*, if provided, can help ease stress for the neck muscles and provide support for seating over extended periods and should be positioned based on the length of the user’s upper body.
- *Arm rests*, if provided, lend support for reading, typing, painting, and similar activities.

A seemingly obvious solution is to design adjustable chairs. Besides being more expensive, such chairs, paradoxically, encourage worse user postures. For example, as observed in an early anthropometric study [5], with more than two dimensions to (manually) adjust, a person regularly forgets the previous

(comfortable) setting among the large space of possible adjustments. Thus, with increased freedom, the user ends up adjusting their own posture to fit an inappropriately dimensioned object. Such considerations also extend to other functional objects, however, they are somewhat less systematically studied in the literature we consulted.

4 ALGORITHM OVERVIEW

We now describe our framework. Our system operates in three steps. First, given a user-specified skeleton body, a set of geometric constraints are extracted from ergonomics guidelines (Section 3 and 5.3) by analyzing contact regions of body and shape geometry. Second, the geometric constraints are integrated into a contact-preserving deformation algorithm to reshape the input model to fit the skeleton body. Finally, the shape-to-body deformation costs are analyzed, and used for ranking and clustering the shapes.




Fig. 3. Modifying human body’s shape attributes (e.g., leg length, body length) and poses leads to different styles of chairs. Our method provides a quick view of how a shape can be altered to accommodate a specified avatar pose that abstracts both the human shape and desired configuration.

Our reshaping algorithm takes a geometric model with multiple components as input. We represent the model with a spatial relation graph where graph nodes denote shape components and graph edges their spatial relations. Given a specified human skeleton, we first map a set of ergonomic guidelines as contact constraints between the human skeleton and the shape. For example, in a sitting pose, a contact constraint between the hip of the skeleton and the top face of the chair seat is established. Such a constraint requires deforming the underlying seat to a specific

size, width, and height according to the skeleton pose and attributes. We formulate the reshaping as an optimization process and design an edit propagation algorithm to deform the shape according to these constraints in the regards that the shape structure is preserved. When the skeleton pose and attributes are modified, the contact constraints continue to induce new edits to the shape. This allows the user to interactively design new shapes for a human with particular characteristics (e.g., design a beach chair for a kid). Figures 3 and 6 show that different human shapes and poses lead to different reshaped chairs.

Once a model is deformed, we examine the deformation cost of the model compared to its original configuration. Such measurements of the deformation cost enable us to effectively rank and classify objects based on human poses (e.g., bench chair versus dining chair). Classification is then done by comparing the shape-to-body deformation costs embedded in a low-dimensional subspace.

5 THE APPROACH

5.1 Shape representation

The input to our algorithm is a family of man-made objects that interact with our human body (e.g., chairs, bed, bikes, etc.). We assume that the collection of shapes belong to the same category and are pre-aligned with consistent orientations [30]. Further, each shape is composed of multiple components that are tagged (such as seat, back, arm, legs). Note that one semantic part (e.g., a chair back) might contain

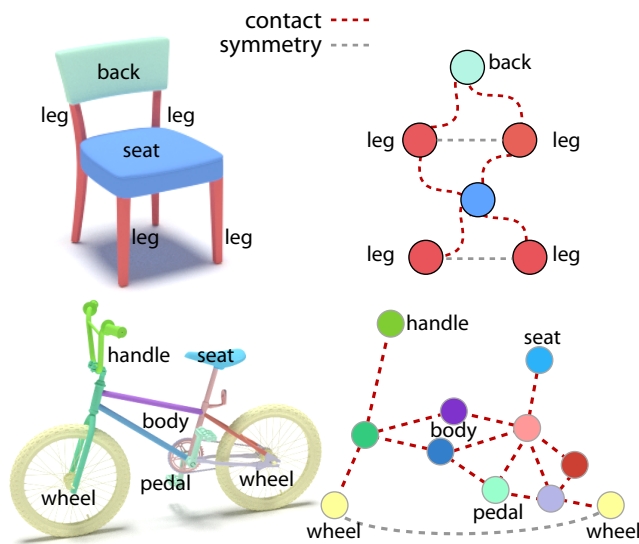


Fig. 4. We represent a chair and a bike model with their respective spatial graphs, where nodes denote individual shape components and the edges denote the connection between components (in our implementation, contact and symmetry relations are used).

multiple components [11]. We refer to existing co-segmentation methods [31], [32] for references.

We encode each shape with a *spatial relation graph* [30]. Each node in the graph corresponds to a component, while each edge denotes a contact or symmetry relation (see Figure 4). We fit each component with a primitive (in our system cuboid and cylinder, computed via PCA); we call such a primitive a *proxy*, which is later used to guide the deformation of the underlying component [26]. We compute the contact information between adjacent components by considering their point-to-point distances. In practice, two components are in contact if their closest point-to-point distance is smaller than 1% of the diagonal length of their smaller bounding boxes (see also [33]).

To account for preservation of shape structure, we need to preserve both the spatial relations among the components (or proxies) and the individual component characteristics. We show in the subsequent sections how such a representation enables a simple and robust contact-preserving deformation paradigm.

5.2 User-specific avatar shapes

Our system exposes to the user a predefined human shape with pose, which is represented as a rigid skeleton (Figure 5). We use a tree structure representation of the skeleton whose root node lies at the chest and each skeleton edge is enclosed with an ellipsoid denoting a body part. Each skeleton edge is also associated with attributes such as length, width, and thickness. As shown in Figure 5, there are a total of 20 nodes and 19 bones in a skeleton.




Fig. 5. Human avatar used in our system consists of a skeleton whose bones represent body parts. The user can modify the shape and poses of the avatar by simple mouse operations. The middle and right column show two representative poses used in our system.

We tag each bone and each node with semantic attributes. For example, the skeleton bone which corresponds to the body part is tagged as “*body-bone*.” The user can use these semantic tags to alter the length, width, and thickness of individual body parts in a small panel (see also in accompanying video).

Besides the pre-specified poses, we allow the user to design their own poses and shape attributes. The user can drag a skeleton node to move it in 3D, the




Fig. 6. In our system, we design four types of avatar poses (sitting, lying, bench sitting, bar sitting) for chairs. Here, each pose induces a modification to the original chair (shown in blue) to be fit to the skeleton shape.

screen mouse movement is mapped as translation in the corresponding skeletal bone plane (defined by its consecutive skeleton bones, multiplied by the current viewpoint transformation). We restrict all transformations to be rigid to preserve the body rigidity. The user can also move the entire body by translating the root node, again with mouse dragging.

5.3 Mapping poses to shape constraints

In this section, we explain how to map the ergonomic guidelines into geometric constraints between the user-specified human shape and a man-made model. For clear exposition, we will primarily focus on chairs, given their widespread use. Other examples (e.g., beds, bikes, etc.) are similarly handled.

We first allow the user to indicate a target pose by either loading a pre-defined pose skeleton or to design her/his own shape (by modifying an existing one). Next, our system provides restrictions to a specified pose based on ergonomic guidances, i.e., the upper leg and lower leg should be roughly orthogonal in a normal sitting pose, while being roughly parallel when seated on a beach chair. We do not allow extreme sitting poses such as hands over head, hips on legs, etc.

We identify a set of contact regions between the body and the chair according to the ergonomic guidelines. For example, the hip and the back part of the legs should contact the top face of the seat, the lower arm will be in touch with chair arms around the center regions, the chair back should be approximately supporting the neck, etc. Please see the appendix for a complete list of ergonomic constraints between human body and various objects used in this paper. All

these types of ergonomic constraints are pre-specified in our system and automatically included when a new pose is created by the user.

The ergonomic relations between the skeleton and the chair are computationally converted into geometric constraints. In particular, for chairs we derive the following types of geometric constraints based on the user-specified poses. The constraints and their specifications are:

- *seat width*: $w_{seat} = (\sum_j^n w_{hip}^j) \times 1.2$, where n is the number of bodies and w_{hip}^j is the width of the hip of body j ;
- *seat height*: $h_{seat} = h_{hip}$, where h_{hip} is the height of the hip (Figure 7);
- *seat length*: $l_{seat} = l_{u_leg}$, where l_{u_leg} the length of the upper legs;
- *arm height*: $h_{arm} = h_{l_arm}$, where h_{l_arm} is the height of the lower arms (Figure 7 right);
- *back length*: $l_{back} = h_{b_top} - h_{seat}$ is the top height of the body minus the height of the seat (Figure 7 left), and
- *seat back angle*: $\theta = angle(u_leg, body) - 20^\circ$, $angle(u_leg, body)$ is the angle between the upper legs and the body spine (Figure 7 left).

These geometric constraints are directly discretized in our system. For example, the contact relation between the hip and the chair seat will lead to a constraint in the height of the seat top face to be a specific value h . Other types of constraints, e.g., the angle between the body and the leg, will lead to a sliding angle constraint between the seat and the back, and so on. For those types of constraints, we allow for a 5 – 15% range of sliding among the exact values to account for the subsequent deformation stiffness. For example, the seat width is allowed to be $1.1 \times - 1.2 \times$ the width of the hips. Figure 7 shows an illustrative mapping of ergonomics to geometric constraints.

Since our aim is not dealing with precise physical




Fig. 7. Mapping of ergonomic guidelines to geometric constraints. For chairs, we derive four types of constraints, namely, height, width, length, and angle, as shown here in the figure. The dotted arrows indicate the specifications from the skeleton shape/pose.




Fig. 8. The contact-preserving deformation paradigm. Starting from a grouped set of geometric constraints (e.g., height, length, angle, etc.), our propagation method iteratively deforms the shape components to fulfill the sets of constraints, while preserving the shape’s structure and connectivity (contacts).

fabrication but instead the deformation and exploration of shapes based on the human ergonomics guidelines, we do not enforce any hard constraints. Once the geometric constraints are derived, we attach them to the model components for subsequent contact-preserving deformation, as explained next. Note that these constraints are edit constraints that are used for shape deformation, while conforming to the subsequent structure optimization.

5.4 Contact-preserving deformation

Given the set of edit constraints derived from the human body and the input model, we now deform the model to meet the constraints while preserving the underlying shape structure. Shape deformation methods, either implicit or explicit, aim to optimize a given shape to fulfill the edit constraints while preserving the shape’s intrinsic properties, or more importantly, the structure. Recently, advanced techniques have been exploited [25], [26], which either use identifying shape features (e.g., feature curves) or high-level deformation delegators (e.g., controllers) to induce the underlying structural optimization, which typically involves complicated strategies to transfer the edits progressively or hierarchically. These methods often proceed one edit at a time, thus it is not clear whether their methods could handle cases when multiple edit constraints are in dependence.

We introduce a simpler and more efficient contact-based edit propagation mechanism to address the problem. The key observation of our approach is that, shape structure is largely encoded in the spatial relations among components [30], while the spatial relations are often linked in the form of spatial connectors, i.e., contacts. Thus, preserving the contact and symmetrical relations can in turn preserve spatial shape structure. Please note that if parametric models are available (e.g., telescopic joints for bike seat into the frame, etc.), then they can be used instead of contact-based deformation model. However, we found such parametric models to be much less prevalent in repositories.

The algorithm: The input to our algorithm is a set of components, denoted as $\{P_1, P_2, \dots, P_n\}$ (sorted by the decreasing number of associated geometric constraints), and groups of constraints $\{C_1, C_2, \dots, C_k\}$ distinguished by names, i.e., heights, length, width, angle, etc. Our goal is to deform the components with respect to the constraints while maintaining all the intra-shape contact relations. A simple approach is to apply these constraints at an initial step and then optimize the remaining components. However, since the constraints can be dependent (e.g., the angle between the back and seat depends on the position of the seat), simply fixing one component according to a set of constraints may lead to a dead-lock. Also, the contact relations may not be explicitly expressed as parameters of transformations, which makes a global optimization difficult to formulate.

Thus, as in iWires [25], we design an algorithm

Data: Input Model $M := \{P_1, \dots, P_n\}$ and ergonomic-inspired constraint groups $C := \{C_1, \dots, C_K\}$ obtained from specified skeleton pose.

Result: Reformed model M .

```

1. while  $i < K$  do
   $C_i = \{c_1, c_2, \dots, c_l\} \in C$ ;
  a. while  $j < l$  do
    i. apply  $c_j \rightarrow P_r \in M$ ;
    ii.  $\Phi \leftarrow P_r$ ;
  end
  while  $\Phi \neq M$  do
    i. Find  $P_m \in \Phi$  which has the maximum
       number of contacts and whose neighbors
       are not all in  $\Phi$ ;
    ii. Find a neighbor  $P_q \in M/\Phi$  of  $P_m$  which
        has the maximum number of deformed
        contacts;
    iii. Deform( $P_q$ ) with its deformed contacts;
    iii.  $\Phi \leftarrow P_q$ ;
  end

```

Algorithm 1: Contact-based deformation propagation.



Fig. 9. Different skeleton shapes and poses lead to different candidate ranking results. Here only 6 chairs (out of 45) are shown for illustration purposes. Please see accompanying video for a complete list.

that transfers the edits in a local, progressive manner. We first organize the constraint into semantic groups as $G = \{C_1, C_2, \dots, C_k\}$, where $C_j = \{c_j^1, c_j^2, \dots, c_j^l\}$ is a set of constraints that have the same type (i.e., heights, weights, lengths, etc.). We then greedily process them to enforce one type of constraint at a time, starting from $C_i = \{c_i^1, c_i^2, \dots, c_i^l\}$ (for example, the heights constraints). For each $c_i^l \in C_i$, we compute the transformation T_i^l for the corresponding component in order to align it to meet c_i^l (e.g., to lift a seat to a certain height h , the transformation is computed as a translation that maps the proxy center o_i to a new position, such that the height of the proxy's top face meets the height h). Width and length constraints are similarly treated except for those constraints which involve multiple components such as angle, wherein we rotate and deform the proxies in accordance to the skeleton pose. Once all the c_i^l -s $\in C_i$ have been handled, we propagate deformation to the remaining (i.e., unedited) components based on contact and symmetry relations.

Edit propagation: Figure 8 illustrates the contact-based propagation. When we deform a component P_i , all its contact points are transformed accordingly (Figure 8(b)). Let us denote the set of already treated components as $\Phi = \{P_k, \dots, P_l\}$ and the set of transformed contacts as $\Theta = \{c_p, \dots, c_q\}$. The propagation treats one component at a time. Each time we look for one component that has the largest number of contacts, say P_m . We look in the neighbors of P_m based on the relation graph and find the neighbor that has not been treated and has the largest number of transformed contacts $\in \Theta$, denoted as χ . Then, χ is selected as the next component and the process continues (see in Figure 8(c)). If all neighbors of P_m have been treated, we proceed to the next P'_m with the largest number of contacts. In Figure 8(c), the chair seat is selected as P_m , since it has the largest number of contacts, and the chair back component is selected as χ , since it has

the largest number of deformed contacts in Θ .




Fig. 10. We quantify deformation cost in terms of volumetric variations in three dimensions (as the velocity of a valid transformation). In this example, it is more expensive to deform an office chair into a bench chair.

To deform a component, given its deformed set of contacts: $\{s_1, s_2, \dots\} \rightarrow \{s'_1, s'_2, \dots\}$, we find the best transformation matrix $T_{4 \times 4}^*$ such that the following energy is minimized:

$$T^* := \arg \min_T \sum_i \|T(s_i) - s'_i\|^2. \quad (1)$$

We solve the minimization using least-squares. The propagation continues until all components are treated. Figure 8 illustrates a simple 2D example of edit propagation. To reduce unnecessary deformations and preserve the original properties of individual components, we adjust T as in the method of [26], i.e., for a cylindrical shape, we retain its cylindrical property during deformation by enforcing uniform scales along its two non-principal axes.

Once C_i is enforced and the transformations are propagated, we proceed to the next constraints group $C_{i+1} \in G$. Note that the constraints are applied one at a time, it may happen that when a C_i is applied, it violates a previous applied constraints C_j , ($j < i$). To address this issue, we enforce the deformation as follows: each time when a component P_i is to be

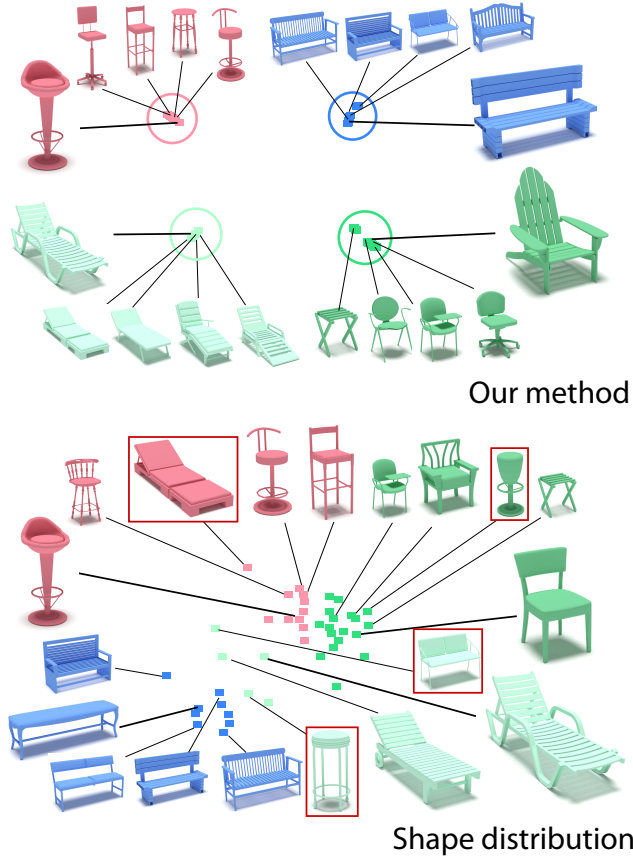


Fig. 11. Embedding of chair models into 2D using multi-dimensional scaling (MDS). Only representative shapes are shown. Note that different styles get clustered differently (indicated by different colors). Note the standard shape distribution produces erroneous results (such as the green bar stool erroneously assigned to the cluster of reclining chairs.)

treated, we enforce all the previous constraints that have been applied to P_i in C_j -s.

Algorithm 1 is an overview of the contact-preserving deformation pipeline. Unlike the propagation methods used in [25], [26] which process groups of elements (wires, controller feature curves) progressively, we delegate the propagation by enforcing the contact relations among components to preserve shape structure. To avoid conflicts, as in [25], we do not revisit any component once it is propagated. The typical number of contacts involved in a model is typically much smaller than the number of feature curves, which makes our method simpler and faster. As shown in the accompanying video, our system supports deforming of dozens of shapes instantly. Please note that by contacts we refer here to the intra-shape contacts, i.e., the contacts among shape components. The body-to-shape contact relations are pre-mapped as geometric constraints (Section 5.3) which serve as edit constraints to the deformation system.

6 APPLICATIONS

User Interface. Figure 12 shows an overview of our user interface that consists of two panels:

- an interaction panel, which
 - allows the user to modify the human shape (represented as a skeleton) – by modify we mean making changes to both the shape attributes and the pose;
 - displays the original and deformed shapes (in blue and green respectively);
- a preview panel, which displays the ranked shapes according to the current human shape.

The user can use the preview panel to browse the shapes. Once the user clicks a preview shape in the bottom panel, the shapes displayed in the top panel are updated to show the new deformed shape which fits the human body.

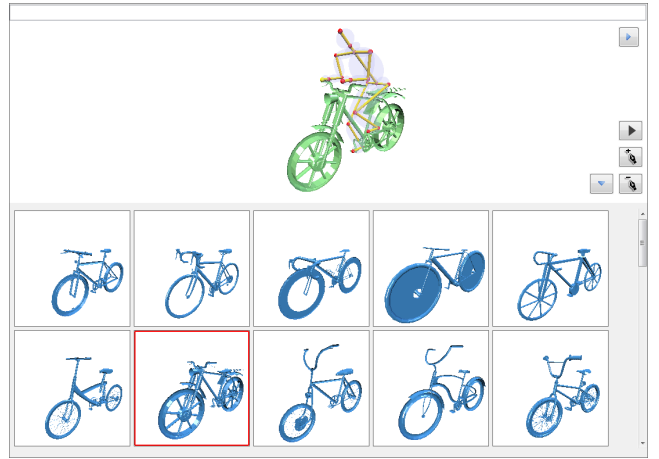


Fig. 12. The user interface of our system consists of two panels: a top interaction panel where users can modify the human shape and get immediate feedback of a deformed shape in accordance to the human shape, and a bottom preview panel to present the ranked models according to the suitability of the models to the current avatar shape and pose.

The user can modify the pose of the skeleton by editing its nodes. This is performed by simple mouse dragging. The user can also modify the geometry by changing the semantic attributes of the skeleton, such as leg length, body width, hip width, etc. When the user edits the human shape, the system provides immediate feedback with deformed shapes shown in the preview windows.

In our current implementation, we provide the user with default human skeleton shape and poses. For example, we include avatar poses for chairs such as normal sitting, bench sitting, beach lying, and bar sitting, etc. (Figures 3 and 5).

Pose-driven geometric ranking. The deformation involved in reshaping a given model indicates how much effort it is required to make a reformed model for a given skeleton body. This enables an intuitive



Fig. 13. Clustering of a set of 45 chairs using 4 annotated poses: lying, bar sitting, bench sitting, and normal sitting. The color models show the clustering results. Sample deformation results are shown in the bottom row, black chairs denote original shape. Not all the models are shown.

algorithm for geometric selection. Specifically, we measure the deformation cost for each deformed chair based on the induced transformation.

Since a common semantic part may contain multiple components, a direct comparison of deformation energy can lead to incorrect results. As we recall that the edits are propagated through all the components, thus, to measure the deformation cost of a chair model M , a simple method might be to add up transformations from all components. However, since a common semantic part (e.g., chair back) can contain different number of components and this number varies across models, thus we compute the average deformation cost for each *semantic part* and compare them across chairs.

We measure for each semantic part p_i how much deformation is induced during the propagation by computing the scale variation of its axis-aligned bounding box. Denote the bounding box of p_i before and after deformation as B_i and B'_i respectively, the deformation cost for p_i is measured as (see Figure 10):

$$e_{p_i} := \Pi_{j=x,y,z} |1 + |\Delta s_j|| + \Pi_{j=x,y,z} |1 + |\Delta t_j||. \quad (2)$$

Here $\Delta s_j, \Delta t_j$ are the scaling and translational differences in each dimension (i.e., x, y, z) of B_i and B'_i . The

total deformation cost for M is then defined as:

$$E_M = \frac{1}{N} \sum_{p_i \in G} e_{p_i}, \quad (3)$$

where G is the set of semantic parts (i.e., seat, back, base, etc.) and $N = |G|$. Figure 9 shows a simple ranking of chairs according to the computed deformation costs, which are ranked in accordance with the skeleton shapes in the first column.

Pose-based object classification. A direct application of our method is pose-based object classification. For example, given a collection of chairs along with a set of annotated skeleton poses, we can classify the chairs as sitting chairs or bench chairs, etc.

We provided different skeleton poses for each family of objects, e.g., normal sitting (chair), reclining (bed), racing (bike), bench press (gym), etc. We compute the deformation costs $E_{M_i}^k$ for each chair M_i according to a particular skeleton k . For a model M , its deformation costs form a vector

$$V_M = (E_M^1, E_M^2, \dots, E_M^k), \quad (4)$$

where k is the total number of skeleton poses. We then compute the pairwise similarity between each pair of models (defined as their L2 norm distance) and embed the resultant similarity matrix into 2D using multi-dimensional scaling (MDS). Figure 11 shows the

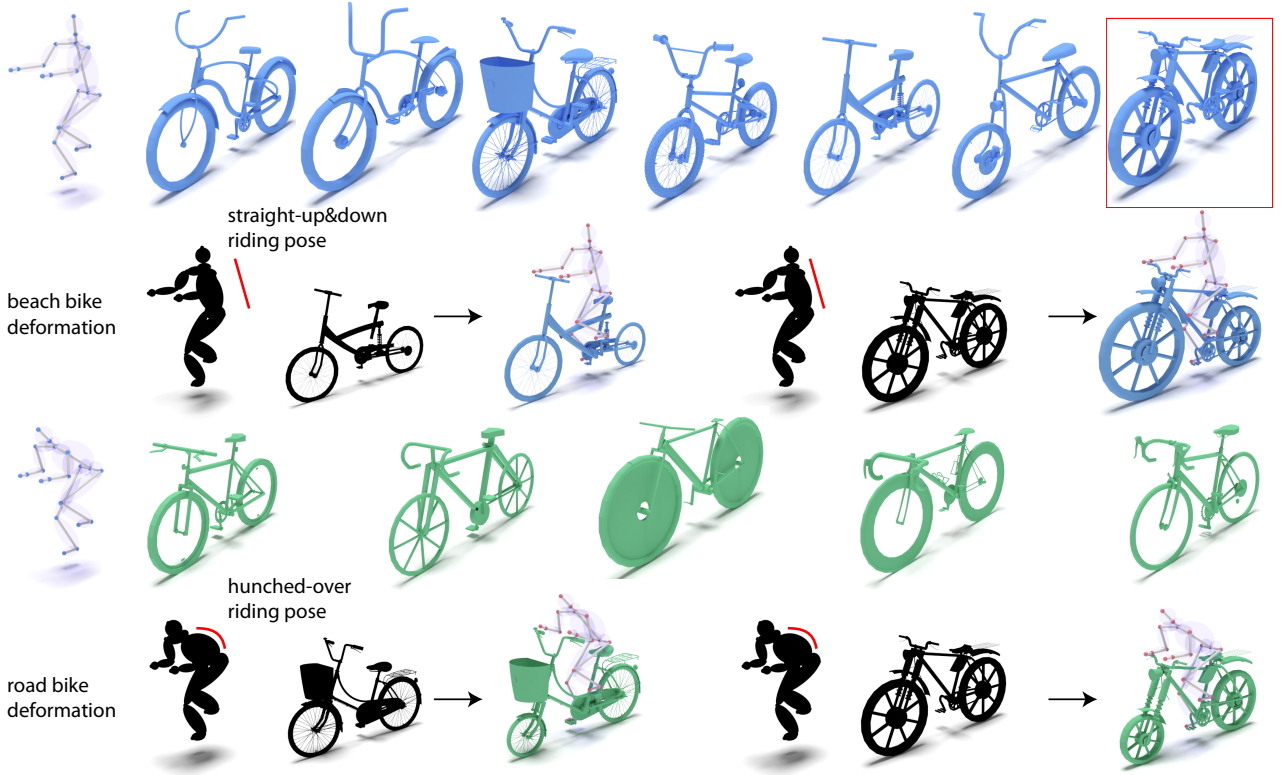


Fig. 14. Clustering of a set of 12 bikes using 2 annotated poses: upright bike and road bike riding poses. Cluster membership resulting from the classification is indicated by corresponding colors. Not all the models are shown.

embedding and clustering results for the chairs (with k-means). To further evaluate our method, we show the MDS results of chairs in Figure 11, using shape distribution [34] signatures for the models. For their results, we used the k-means algorithm in the 64-dimensional descriptor space of shape distribution. Note that different styles of chairs get more clearly separated with our pose-based embedding method.

Classification results: We evaluate our algorithm on 4 collections of man-made models, namely chairs, bikes, beds, and gym equipments. We include a data set of 45 chairs, 15 beds, 12 bikes, and 15 gym equipments, each of which typically consists of multiple styles. Similarly, for each style, we extract a set of geometric constraints and attach them to objects during deformation (see the supplemental file).

In Figures 13, 14, and 16, we show corresponding clustering results of chairs, bikes, beds, and gym equipments along with some of the sampled deformations. Our results clearly show that the ergonomic guidelines can largely help the process of functional classification, which is otherwise challenging to achieve by traditional approaches.

Human-centric content co-retrieval. The human workspace is an environment where objects are often correlated and serve for particular activities. Our framework can be used for human-centric content co-retrieval and co-placement. This is implemented as follows. Given a specified annotated human body

shape, we can retrieve the best suitable objects and deform them to suit a particular human shape based on pose and ergonomics guidelines.

In Figure 15, the user annotates an office sitting pose, and she then retrieves an office sitting chair, and an office desk along with a monitor which is placed on the table. The size, placements and orientation of the three objects are automatically determined by ergonomic guidelines (Figure 7). In this example, the avatar links the three objects with ergonomics, resulting in the modifications of sizes, placements, and contact relations among the three objects.

We annotate the spatial relation among individual




Fig. 15. With our framework, we can bring multiple objects into a coherent workplace. In this example, the size, position and arrangements of the chair, table and monitor are determined by the ergonomic guidelines derived with the human skeleton.

objects (e.g., the monitor should be placed on the table) and obtain the ergonomic constraints simultaneously between the human body and the three objects. The placements come into place at first, followed by sizes and orientations (see Figure 1). Since the three objects share spatial contact constraints with the human shape while the monitor and the table also share a placement constraint, we perform the deformation of the three objects sequentially in a spanning tree order to resolve conflicts. Please note that in this setup, we can also navigate through multiple collections of shapes (in this example, chairs, tables, and monitors). This allows the user to browse over best suited configurations for the current human shape – when the user selects one of the objects, the other two objects will get automatically re-ranked and updated in accordance with the human shape.

7 VALIDATION

User study: We evaluated our method with a user study. The main purpose of the user study is two-fold. First, we wanted to validate the fidelity of the deformed shape generated by our algorithm. Second, we wanted to evaluate how our selective exploration agrees with human perception. We designed the user study with two experiments. In the first experiment, we fixed 9 human avatar shapes, 4 for chair, 2 for bike, 3 for bed. For each human skeletal shape, we deformed the corresponding shape set by the ergonomics constraints derived. We then designed an interface which randomly shows one model (either deformed or original) at a time and we asked the user to select a best matching human avatar presented in a side view. This experiment helps validate whether the generated shapes and classification results agree with a desired human shape.

In the second experiment, we asked the user to do suitability selection in a finer granularity. In particular, for each participant, we randomly select 3-4 models from each collection. Given a human avatar, users were asked to rank the selected models (in original shape) according to the fitness of the model to the given human avatar. This process is repeated multiple times for all the prescribed avatars, and finally we compare the ranking results collected from the participants to those generated by our algorithm.

A total of 74 users participated in our experiments. In the first experiment, we found people to be fairly good at correctly identifying which models are more suitable to which human shapes when the functionalities of given models are visibly different. The statistics are encoded as confusion matrix as shown in Table 1. Our algorithm achieves an accuracy of over 85.4% for most model type (i.e., normal chairs, racing bikes, etc.). Interestingly, under some extreme deformation (e.g., deforming a beach chair to a bar stool), people

TABLE 1

Confusion matrices for chair, bike, and bed classification.

	bar stool	bench	beach chair	chair
ine bar stool	73	3	4	0
bench	9	72	8	22
beach chair	1	2	70	4
chair	0	6	8	106

	king bed	recline bed	single bed
ine king bed	43	3	7
recline bed	4	83	12
single bed	0	2	31

	road bike	racing bike
ine road bike	63	5
racing bike	8	37

often got confused in making judgements whether from its appearance or actual usability (Table 1 chair).

In the second experiment, we used Kendall Tau rank correlation coefficient [35] to compare the rankings from users and our algorithm. Kendall Tau value is a measurement of association between two random variables x, y , defined as:


$$\tau = \frac{n_c - n_d}{\frac{1}{2} \times n \times (n - 1)}, \quad (5)$$

where n_c is the number of concordant pairs $(x_i, y_j), i = 1, \dots, n, j = 1, \dots, n$, while n_d discordant pairs, and n is the number of observations for each variable x and y . The coefficient is in the range of $[-1, 1]$ with higher values indicating better agreement and -1 denoting perfect agreement. In our test, the Kendall Tau value was 0.717 in the second experiment. The score decreases when the models are within the same style such as the chair models shown in the same rows of Figure 9. This is expected as human judgements purely based on perception could be imprecise without being able to physically operate the models. The user study indicates that our pose-induced shape classification is mostly consistent with those of the human users.

Timing. Our algorithm runs at interactive rates. The edit propagation is linear in terms of the number of components. For the collection of 45 chairs we tested, it took less than 2s to process them all (see the accompanying video). All experiments were done on a desktop with Intel i5-4430 processor (3.2GHZ) and 8GB memory.

Limitations. Our method has some limitations. First, although our method uses ergonomic guidelines to deform the shapes to fit a given human skeleton, we can not precisely account for ergonomics, which is a more general and complicated problem. That means, our method can not be directly used for reshaping an object into a production-ready object. This remains a future direction to investigate. Second, in our experiments, we practically specified parameters deduced from ergonomic guidelines and map them to geometric constraints for different object categories (we used [1] for reference parameters). Although we

BEDS



GYM Equipments




Fig. 16. Clustering of a set of 15 beds using 3 annotated poses: single bed, double bed, and reclining (top) and of a set of 15 gym equipments using 3 annotated poses: bench press, riding/running, and level lift. Cluster membership resulting from the classification is indicated by corresponding colors. Not all the models are shown.

allow for partial stiffness between the human body and the underlying shape, the constraints are mostly empirically defined. An automatic method to learn such sets of parameters from object categories would be an interesting future direction. Finally, our method assumes the input is comprised of multiple, consistent and meaningful components (obtained via co-segmentation). If this assumption fails, the user needs to manually segment and tag the input shapes.

8 CONCLUSION

We introduced an algorithm to couple ergonomic guidelines with geometric reshaping, exploration and classification. The essential component of our algorithm is a pose-driven contact-preserving deformation paradigm that quickly adapts multiple shapes to a given human shape while allowing for fast preview of deformed shapes. Instead of browsing the shapes by abstracting them into a common space (e.g., using templates [6]), we allow the user to use a human skeleton to explore and classify shape collections by customizing the attributes and poses of the avatar.

In the future, we plan to further consider ergonomics for geometry analysis, for example, automatic learning of geometric parameters from ergonomics guidelines for categories of objects, ergonomic-driven content co-placement, etc. We would also like to examine physical simulators for the analysis of fidelity/stability of geometric objects and also their interaction with soft materials (e.g., leather).

ACKNOWLEDGEMENT

We would like to thank the anonymous reviewers for their constructive comments and Melinos Averkiou for proofreading the final version of the paper. We thank Moos Hueting and all the volunteers for their help with the user study. This work was supported in part by the ERC Starting Grant SmartGeometry (StG-2013- 335373) and the National Science Foundation (1302267).

REFERENCES

- [1] J. Panero and M. Zelnik, "Human dimension and interior space: A source book of design reference standards," Watson-Guptill, Ten Speed Press.
- [2] G. Salvendy, *Handbook of Human Factors and Ergonomics*. John Wiley & Sons, Inc., 2006.
- [3] J. Tong, J. Zhou, L. Liu, Z. Pan, and H. Yan, "Scanning 3d full human bodies using kinects," *Proc. IEEE Virtual Reality*, vol. 18, no. 4, pp. 643–650, 2012.
- [4] R. Erickson, "Furniture by robert erickson: chairs that fit," brochure, Nevada city, 1992.
- [5] F. I. R. Association, "Anthropometric data: Limitation in use," Jour. Information Library, 1961.
- [6] M. Ovsjanikov, W. Li, L. Guibas, and N. J. Mitra, "Exploration of continuous variability in collections of 3d shapes," *ACM TOG (SIGGRAPH)*, vol. 30, no. 4, pp. 33:1–33:10, 2011.
- [7] V. G. Kim, W. Li, N. J. Mitra, S. Chaudhuri, S. DiVerdi, and T. Funkhouser, "Learning part-based templates from large collections of 3d shapes," *ACM TOG (SIGGRAPH)*, vol. 32, no. 4, pp. 54:1–54:11, 2013.
- [8] S.-S. Huang, A. Shamir, C.-H. Shen, H. Zhang, A. Sheffer, S.-M. Hu, and D. Cohen-Or, "Qualitative organization of collections of shapes via quartet analysis," *ACM TOG (SIGGRAPH)*, vol. 32, no. 4, 2013.
- [9] M. Averkiou, V. Kim, Y. Zheng, and N. J. Mitra, "Shapesynth: Parameterizing model collections for coupled shape exploration and synthesis," *CGF (EUROGRAPHICS)*, 2014.
- [10] H. Laga, M. Mortara, and M. Spagnuolo, "Geometry and context for semantic correspondence and functionality recognition in manmade 3d shapes," *ACM TOG*, 2013.
- [11] Y. Zheng, D. Cohen-Or, M. Averkiou, and N. J. Mitra, "Recurring part arrangements in shape collections," *CGF (EUROGRAPHICS)*, vol. 33, p. to appear, 2014.
- [12] N. Fish*, M. Averkiou*, O. van Kaick, O. Sorkine-Hornung, D. Cohen-Or, and N. J. Mitra, "Meta-representation of shape families," *ACM TOG (SIGGRAPH)*, 2014, * joint first authors.
- [13] V. G. Kim, S. Chaudhuri, L. Guibas, and T. Funkhouser, "Shape2Pose: Human-centric shape analysis," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, Aug. 2014.
- [14] D. F. Fouhey, V. Delaitre, A. Gupta, A. A. Efros, I. Laptev, and J. Sivic, "People watching: Human actions as a cue for single-view geometry," in *ECCV*, 2012.
- [15] A. Gupta, S. Satkin, A. A. Efros, and M. Hebert, "From 3d scene geometry to human workspace," in *Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [16] G. Saul, M. Lau, J. Mitani, and T. Igarashi, "Sketchchair: An all-in-one chair design system for end users," in *Proc. Tangible, Embedded, and Embodied Interaction*, 2011, pp. 73–80.
- [17] G. Cranz, *The Chair: Rethinking Culture, Body, and Design*. W. W. Norton & Company, 1988.
- [18] C.-H. Shen, H. Fu, K. Chen, and S.-M. Hu, "Structure recovery by part assembly," *ACM TOG (SIGGRAPH Asia)*, vol. 31, no. 6, pp. 180:1–180:11, 2012.
- [19] S. Chaudhuri, E. Kalogerakis, S. Giguere, , and T. Funkhouser, "Attribl: Content creation with semantic attributes," *ACM Symposium on User Interface Software and Technology (UIST)*, Oct. 2013.
- [20] M. Botsch and O. Sorkine, "On linear variational surface deformation methods," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 213–230, 2008.
- [21] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel, "Laplacian surface editing," in *Proc. SGP*, 2004, pp. 175–184.
- [22] Y. Lipman, O. Sorkine, D. Levin, and D. Cohen-Or, "Linear rotation-invariant coordinates for meshes," in *ACM TOG (SIGGRAPH)*, 2005, pp. 479–487.
- [23] Y. Lipman, D. Levin, and D. Cohen-Or, "Green coordinates," *ACM TOG (SIGGRAPH)*, vol. 27, no. 3, pp. 78:1–78:10, Aug. 2008.
- [24] V. Kraevoy, A. Sheffer, A. Shamir, and D. Cohen-Or, "Non-homogeneous resizing of complex models," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 111:1–111:9, Dec. 2008.
- [25] R. Gal, O. Sorkine, N. J. Mitra, and D. Cohen-Or, "iwires: an analyze-and-edit approach to shape manipulation," *ACM TOG (SIGGRAPH)*, vol. 28, no. 3, pp. 33:1–33:10, 2009.
- [26] Y. Zheng, H. Fu, D. Cohen-Or, O. K.-C. Au, and C.-L. Tai, "Component-wise controllers for structure-preserving shape manipulation," *CGF (EUROGRAPHICS)*, vol. 30, no. 2, pp. 563–572, 2011.
- [27] H. Li, I. Alhashim, H. Zhang, A. Shamir, and D. Cohen-Or, "Stackabilization," *ACM TOG (SIGGRAPH Asia)*, vol. 31, no. 6, 2012.
- [28] A. C. Mandal, *The Sitting Position, Its Anatomy and Problems*. Daphne Publishing, 1984.
- [29] G. Gordon, "Design a chair that fits like a glove," Fine Woodworking, 1992.
- [30] Y. Zheng, D. Cohen-Or, and N. J. Mitra, "Smart Variations: Functional Substructures for Part Compatibility," *CGF (EUROGRAPHICS)*, vol. 32, no. 2pt2, pp. 195–204, 2013.
- [31] Q. Huang, V. Koltun, and L. Guibas, "Joint shape segmentation with linear programming," *ACM TOG (SIGGRAPH Asia)*, vol. 30, no. 6, pp. 125:1–125:12, 2011. [Online]. Available: <http://doi.acm.org/10.1145/2070781.2024159>
- [32] Y. Wang, S. Asafi, O. van Kaick, H. Zhang, D. Cohen-Or, and B. Chen, "Active co-analysis of a set of shapes," *ACM TOG (SIGGRAPH Asia)*, vol. 31, no. 6, pp. 165:1–165:10, 2012.

- [33] E. Kalogerakis, S. Chaudhuri, D. Koller, and V. Koltun, "A probabilistic model for component-based shape synthesis," *ACM TOG (SIGGRAPH)*, vol. 31, no. 4, pp. 55:1–55:11, 2012.
- [34] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," *ACM TOG*, vol. 21, no. 4, pp. 807–832, 2002.
- [35] M. G. KENDALL, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1/2, 1938.

APPENDIX

Mapping ergonomic constraints. Similar to the map of ergonomic guidelines to geometric constraints for chairs in section 6.3, we map human shapes to the bed, bike and gym models.

For beds:


- *seat width*: $w_{seat} = (\sum_j^n w_{body}^j) \times 1.5$, $w_{body} = w_{hip} + (w_{arm}) \times 2$, where n is the number of bodies and w_{body}^j is the width of the body j , defined by the sum of width of the hip and two arms of the body;
- *seat height*: $h_{seat} = w_{l_leg}$, where w_{l_leg} is the length of the lower leg, as considering the height of common beds;
- *seat length*: $l_{seat} = (l_{body}) \times 1.2$, where l_{body} the length of the body, defined by the sum of length of the head, neck, spine, upper leg, and lower leg;
- *seat back angle*: $\theta = angle(u_leg, body) - 20^\circ$, $angle(u_leg, body)$ is the angle between the upper legs and the body spine, as default, the angle is 90° , in the case of a lying pose with parallel legs and the body spine.

For bikes:

- *seat width*: $w_{seat} = w_{hip} \times 0.5$, where w_{hip} is the width of the hip;
- *seat height*: $h_{seat} = h_{hip}$, where h_{hip} is the height of the hip;
- *handle position*: $p_{handle} = (p_{l_hand} + p_{r_hand}) \times 0.5$, where p_{handle} is the center of top face of the proxy of bikes' handles, and p_{l_hand}, p_{r_hand} are the center of left and right hand of body respectively;
- *pedal height*: $h_{pedal} = h_{u_foot}$, where h_{u_foot} is the height of the higher foot.

For gym equipments:

- *seat width*: $w_{seat} = w_{hip}$, where w_{hip} is the width of the hip;
- *seat height*: $h_{seat} = h_{hip}$, where h_{hip} is the height of the hip;
- *arm position*: $p_{arm} = (p_{l_hand} + p_{r_hand}) \times 0.5$, where p_{arm} is the center of the proxy of arms, and p_{l_hand}, p_{r_hand} are the center of left and right hand of body respectively;
- *seat back angle*: $\theta = angle(u_leg, body) - 20^\circ$, $angle(u_leg, body)$ is the angle between the upper legs and the body spine.



Youyi Zheng is currently an assistant Professor at the School of Information Science and Technology, ShanghaiTech University. He obtained his PhD from the Department of Computer Science and Engineering at Hong Kong University of Science & Technology, and his M.Sc. and B.Sc. degrees from the Department of Mathematics, Zhejiang University. His research interests include geometric modeling, imaging, and human-computer interaction.



Han Liu is a Ph.D. candidate supervised by Niloy J. Mitra in Visual Computing Center at King Abdullah University of Science and Technology. She received her master's degree in Computer Software and Theory from Xiamen University in 2010. Her current research work focuses on geometric modeling and layout optimization.



Julie Dorsey is a Professor of Computer Science at Yale University, where she teaches computer graphics. She came to Yale in 2002 from MIT, where she held tenured appointments in both the Department of Electrical Engineering and Computer Science (EECS) and the School of Architecture. She received undergraduate degrees in architecture and graduate degrees in computer science from Cornell University.



Niloy J. Mitra is a Professor of Computer Science at University College London (UCL). His research focuses on algorithmic issues in shape analysis and geometry processing. He received the ACM SIGGRAPH Significant New Researcher award in 2013 and the BCS Roger Needham award in 2015.