

ERGOBOSS: Ergonomic Optimization of Body-Supporting Surfaces

Danyong Zhao, *Member, IEEE*, Yijing Li, *Member, IEEE*, Siddhartha Chaudhuri, Timothy Langlois, *Member, IEEE*, and Jernej Barbič, *Senior Member, IEEE*

Abstract—Humans routinely sit or lean against supporting surfaces and it is important to shape these surfaces to be comfortable and ergonomic. We give a method to design the geometric shape of rigid supporting surfaces to maximize the ergonomics of physically based contact between the surface and a deformable human. We model the soft deformable human using a layer of FEM deformable tissue surrounding a rigid core, with measured realistic elastic material properties, and large-deformation nonlinear analysis. We define a novel cost function to measure the ergonomics of contact between the human and the supporting surface. We give a stable and computationally efficient contact model that is differentiable with respect to the supporting surface shape. This makes it possible to optimize our ergonomic cost function using gradient-based optimizers. Our optimizer produces supporting surfaces superior to prior work on ergonomic shape design. Our examples include furniture, apparel and tools. We also validate our results by scanning a real human subject's foot and optimizing a shoe sole shape to maximize foot contact ergonomics. We 3D-print the optimized shoe sole, measure contact pressure using pressure sensors, and demonstrate that the real unoptimized and optimized pressure distributions qualitatively match those predicted by our simulation.

Index Terms—ergonomics, CAD, optimization, shape design, FEM, contact, 3D printing.

1 INTRODUCTION

Humans routinely interact with supporting surfaces – we sit on them (chairs), lie on them (beds and sofas), stand on them (footwear), and work all day with our hands on them (mice and keyboards). It is important for our health and well-being to shape these surfaces to be comfortable and ergonomic [1]. Traditionally, such contoured surfaces are shaped by craftspersons who draw upon generations of design knowledge to form shapes that are likely to be comfortable for a wide range of people, often with human testers in a feedback loop to refine the design. However, the traditional design process does not scale easily to *customized* support for individual people: it is expensive, time-consuming, requires extensive interaction and iteration with the target user, and is usually not based on fine-grained biomechanical modeling of the human body. At the same time, customized fabrication methods such as 3D printing have become more and more accessible to the general public [2]. Yet as long as customized *design* remains difficult, such fabrication methods remain grossly under-utilized.

A promising path to closing this gap is to develop computational methods that can *automatically* synthesize supporting surfaces whose contact ergonomics are personalized to a specific person. In the geometric shape design literature, shapes have been optimized for various objectives, typically aesthetics, visual plausibility and crude mechanical function [3]. A few methods have tried to explore human

interaction with the object, using non-physics-based data-driven models [4], [5], [6], [7], or a precomputed approximate pressure map for an inaccurate contact distribution (a negative image mold) that is used only as a fixed importance map for geometric support synthesis [8], [9].

In this paper, we propose an automatic, physics-based technique to optimize supporting surfaces for specific people, with known geometric (bodyshape) and elastic (tissue) properties. Our method is based on the computational contact mechanics between the human body and the supporting surface. While humans are obviously biomechanically complex with intricate muscle and fat geometry, we make our optimization tractable by modeling human body parts as consisting of a deformable layer (with measured material properties) surrounding a rigid core. The supporting surface is modeled as rigid. Based on ergonomics literature, we define a meaningful and tractable *ergonomic cost function* that scores the pressure distribution exerted by the support on the human. We estimate this contact pressure using large-deformation FEM deformable object elasticity. Importantly, and specifically for our optimization use-case, we design our cost function to be *differentiable* with respect to the parameters defining the geometry of the supporting surface. Hence, the latter can be efficiently optimized using gradient descent in order to minimize the cost.

We apply our pipeline to design optimal contact surfaces of chairs, shoe insoles, and computer mice. We use a custom acquisition pipeline to gather elasticity maps of the human body parts in contact with these surfaces, and use these material properties in our simulation. We validate our results against baselines, as well as in an experiment with an actual fabricated surface, to show that our pipeline is accurate and robust. In practice today, ergonomic surfaces are rarely designed completely using computer models; but instead

• Danyong Zhao, Yijing Li and Jernej Barbič are with the Department of Computer Science, University of Southern California, Los Angeles, CA, 90089.

Siddhartha Chaudhuri is with Adobe Research and IIT Bombay.

Timothy Langlois is with Adobe Research.

E-mail: danyongz@usc.edu, yijingl@usc.edu, sidch@adobe.com, tlangloi@adobe.com, jnb@usc.edu




Fig. 1. **Optimizing the shape of the supporting surface (chair) to maximize ergonomics of contact against a human.** We model the deformations of the human soft tissue using the Finite Element Method. Given any shape of the supporting surface, we calculate the contact pressures using physically based simulation, and score the ergonomics using an ergonomic score function that penalizes localized contact forces. Our gradient-based optimizer then optimizes the supporting surface chair shape to minimize the score.

real prototypes are manufactured, real people evaluate them by interacting with them, and then the design is iterated. Our work can shorten the design loop because it can provide a good initial design for this practical subsequent design process, and/or makes it possible to virtually prototype the design iterations.

Our technical contributions include a stable, easily computable, and differentiable ergonomic function to measure the ergonomics between a rigid object and a nonlinear elastic deformable object. We also give a phenomenological simulation model of the human subskin soft tissue that is sufficiently complex to be able to express realistic tissue deformations, yet not overly complex so that it can still be simulated efficiently and iterated in a design loop. The model consists of a rigid core surrounded by and coupled to a layer of soft nonlinear elastic tissue. We give a stable method to simulate the resulting coupled rigid-deformable human mechanical system in contact with an external object (a supporting surface), with spatially varying contact regions and pressures. Finally, we give an efficient method to optimize the shape of the object to maximize the ergonomic comfort of the contact (Figure 1). To the best of our knowledge, our work is the first to successfully incorporate an ergonomic objective, nonlinear large deformation FEM elasticity, and *distributed unilateral* contact into ergonomic shape design of man-made objects. *Distributed contact* is defined as contact with (potentially) several simultaneous contact sites, each distributed over a non-zero surface area [10]. Contact is *unilateral* if it can only push but not pull, i.e., it has same characteristics as real-world contact. Note that distributed unilateral contact is the most general and most realistic contact between two 3D objects, and is significantly harder to process and optimize than single-point contacts, bilateral attachments, or permanent conforming contact. Furthermore, we also contribute a new experimental method to measure real distributed contact pressures between two curved surfaces. We demonstrate how this can be reliably achieved using standard, inexpensive resistance-based pressure sensors. We use our method to experimentally validate that our optimized pressure distributions qualitatively match real-world pressures.

2 RELATED WORK

Shape design: There are various functional and aesthetic goals which shape design has attempted to address, and many methods explored to reach these goals. Topology optimization aims to add or remove material from a structure to optimize its properties. Often this is done by optimizing a dense grid of voxels to maximize the stiffness or robustness of a structure while minimizing its weight, in the presence of either static [11], [12] or stochastic [13], [14] external forces. Other methods optimize a level-set representation of the geometry [15]. While capable of producing intricate structures, topology optimization is computationally expensive, and is not generally used with nonlinear material models.

Other shape optimization methods keep the user in the loop and rely on precomputation to generate suitable exploration spaces. Shugrina et al. [16] adaptively sample the design space of a parametric object, mapping out which parts of the space satisfy constraints (such as manufacturability). Then, during runtime, users can quickly explore the design space, while ensuring they remain in the valid constraint region. Yumer et al. [17] allow high-level geometrical editing by relying on crowd-sourced comparisons of semantic attributes. This data is used to learn a mapping between semantic attributes and geometry. While capable of producing a wide variety of useful shapes, the amount of pre-computation that would be required to capture the fine variation between different humans is immense. Instead we rely on optimization based on human models.

Traditionally, shape optimization methods operate with static constraints (e.g. Umetani et al. [18], Wang et al. [19]), although some methods consider dynamic properties of the object being optimized. There exist ways to design spinning [20], standing [21], flying [22] and floating objects [23]. Dynamics-Aware Coarsening [24] can speed up the design of compliant (large deformation) objects with dynamic motion, such as jumping. In our work, we focus on optimizing the human ergonomics of contact, as opposed to the object's dynamic or static properties. Several publications in computational fabrication design geometry or material properties under equilibrium constraints [25], [26], typically using the implicit function theorem, as in

our work. However, previous work has not incorporated unilateral contact, or proposed a methodology to validate such design by measuring real contact pressures on curved surfaces. Recently, Montes et al. [27] optimized the patterns of tight-fitting clothing so that the fit to the human is as ergonomic as possible. Because they modeled tight-fitting clothing, they could assume that the contact is permanent and conforming (essentially a bilateral attachment). In contrast, we model *distributed unilateral* contact between the human and the external object. As such, our method can optimize the ergonomics of contact between the human and arbitrarily shaped, not necessarily tight-fitting objects. Like in the real world, the specific contact sites and their sizes are not known a priori, but are determined by our ergonomic score optimizer together with the human tissue deformations and the elastostatic equilibrium of the human against contact forces and gravity. In the medical community, researchers have evaluated the human body contact ergonomics for various body parts, such as feet and buttocks [28], [29]. They modeled the individual musculo-skeletal anatomy [28], [29], which is difficult to generate due to the required medical imaging, organ templates, and anatomy transfer to the subject. Furthermore, unlike us, they did not aim to optimize the shape of external objects in contact with the human, but rather merely calculated the contact strains for given fixed external object shapes.

Furniture comfort evaluation: Considering how much time humans spend interacting with furniture, evaluating furniture comfort is important for designers. There are general ergonomic guidelines on how to choose and adjust chairs, e.g., from the US General Services Administration [30], but designing individual chairs to maximize comfort is still an active research area. Different computational measures of comfort have been used, such as joint angles [31] and interface pressure between the human and chair [32]. We choose to use interface pressure to quantify model comfort, as it correlates well with user comfort ratings [32], [33]. Savonnet et al. [34] provide an extensive review of various deformable body models (from twenty-seven different papers) that have been created to predict the comfort of chair models or the development of pressure sores; they conclude that “These models share a lack of validation and generally make little allowance for anthropometric diversity.” Generally, models aiming to predict pressure sores differentiate between internal soft tissue (i.e., fat, muscle, and skin were modeled separately), while models to predict seating comfort generally group the internal tissues together. Multiple material models have been used to model the soft tissue deformation, including Mooney-Rivlin, neo-Hookean, Ogden, and linear elasticity. We build our human model as a rigid-core (bones) surrounded by a deformable layer (all soft tissue), using the neo-Hookean material model with compression resistance. The depth of the soft tissue, and its material properties, are calibrated based on force measurements from a user, in a process similar to [35].

Layered models are useful for animation of characters [36], where inner layers can control large-scale behavior, and outer layers provide localized deformation due to contacts or other forces. Pauly and Pai [37] investigated techniques for modeling contact between quasi-rigid objects. Terzopoulos and Witkin [38] decompose deformation

into reference (rigid) components and displacement components. The reference component evolves according to rigid-body dynamics, allowing the model to handle large deformations with linear elasticity. This was extended to deformable reference shapes [39]. Implicit surfaces have been used for the outer layers [40]. A series of papers [41], [42] present a fully coupled elastodynamic system of a rigid core surrounded by a single layer of deformable tetrahedra. Recently, layered models were extended to articulated “cores” [43], [44]. However, prior work has not used layered models for shape optimization involving contact. In our problem, we only require the final converged state of the contact simulation, and its gradient. We present a layered elastostatic model, which we found to be faster and differentially more stable than elastodynamic models.

Ergonomic furniture: Previous studies have investigated the design of furniture based on ergonomics. Saul et al. [45] developed a system which extruded 2D sketches into 3D furniture, then physically simulated a rag doll (composed of rigid links) sitting on the shapes to check gross characteristics such as sizing and balance of the furniture. The results from this quick simulation can be used to adjust the furniture. This method is incapable of capturing fine differences between humans, or pressure differences along the body. Zheng et al. [46] use a rigid skeleton model to get rough body part lengths of users, and use ergonomics to infer contact configurations between parts (e.g., the model’s back should be in contact with the seat back). The method then adjusts the various rigid parts of the chair to support the model, and allows users to interactively edit the model. This method can quickly show various chairs adapted for the target pose, but cannot capture important interface pressure distributions. Perhaps most similar to our work is that of Leimer et al. [8]. This method uses a rough estimate of contact pressure, akin to what would happen if a body model was immersed in a negative image mold (so that every downward-facing vertex was supported). This pressure distribution is used to fit a Catmull-Clark subdivision surface, ensuring areas of high pressure are supported. While capable of automatically producing shapes supporting various body poses, pressure induced by a negative image mold can be quite different from that of real supporting surface and this is important for ergonomic design validation. In recent work, Leimer et al. [9] gave an optimization approach whereby desirable contact pressure distributions are computed by maximizing their smoothness and keeping the net forces and torques on the human skeleton links as close to equilibrium as possible. Although plausible, their approach relies heavily on geometric skinning-based deformation, which may be inaccurate. Unlike their work, we model human tissue elastic properties and simulate the human in contact against the supporting surface using computational contact mechanics of deformable objects. An important key difference to Leimer’s work is that they require knowing *a priori* what human vertices are in contact. However, real contact between a human and an external object is complex and spatially varying (distributed). Simply postulating known contact areas leads to substantial inaccuracies, as we experimentally show in Section 5. Our method determines the entire distribution of the contact pressures *automatically* using contact mechanics in each iteration of shape optimization,




Fig. 2. **Modeling and optimizing the contact between a human and a man-made object.** The man-made object is fixed in space during each contact simulation, whereas the human body is subjected to gravity. Our optimizer explores the space of man-made object shapes to maximize contact ergonomics.

which leads to significantly more realistic contact pressure distributions. We compare to both publications [8], [9] in Section 5.

3 SIMULATION MODEL

Our goal is to optimize the surface geometry of an external object (a supporting surface) to maximize the ergonomics of the contact against the human. The supporting surface is modeled as a triangle mesh. The human body is modeled as a FEM deformable layer of varying thickness, surrounding a rigid core that models the inner hard components such as bones. The local thickness and stiffness of the deformable layer, as well as its surface geometry, is measured via actual human studies, approved by the medical ethics committee (Institutional Review Board) at our institution. Mass is distributed across the body according to realistic estimates. The deformable layer is modeled as a tetrahedral mesh between the outer and inner surface of the body, and the volume inside the inner surface (the “core”) is modeled as a rigid body. While this simplifies human anatomy, such phenomenological models have been recently demonstrated to offer a good tradeoff between the modeling complexity and invested computational costs, for human body contact mechanics simulations [35].

Given a human body and a supporting surface, modeled as above, we define an *ergonomic cost function* (§ 4.1) that estimates how uncomfortable the contact between the human and the supporting surface is. Specifically, we follow the standard observation [32], [33] that high contact pressures on the body are uncomfortable, and define our cost function to penalize such high contact pressures. For each specific shape of the supporting surface, the spatially varying contact pressures are computed using an FEM simulation (§ 3.3) that couples the rigid and deformable components of the body model, and resolves the human contact against the supporting surface. In Section 4, we then optimize the shape of the supporting surface, by calculating the gradients of the ergonomic function with respect to the shape of the supporting surface.

3.1 Modeling the human body

Figure 2 gives an overview of how we model the human body in contact with a man-made rigid object (the support-




Fig. 3. **Contact between the human body and a supporting surface.** Left: human sitting in a chair. Middle: foot contacting a shoe sole under human body gravity. Right: hand grasping a mouse.

ing surface). Our IRB-approved study acquired the human body shape by scanning real persons using the Artec Eva scanner [47]. For our chair example, we scanned the lower back, buttocks, torso and thigh. We approximated the sitting pose by asking the subject to position themselves on all fours, with pillows supporting the torso. For our foot example, we scanned the human foot. For our mouse example, we used a generic hand mesh purchased online [48]. Note that it is sufficient to scan and model the larger region of interest of the human body; it is not necessary to scan the entire human body.

We measured human body elastic stiffnesses (Young’s modulus) and deformable layer depths on contact regions using a technique similar to the system described in [35]. We captured 60 samples across the lower back, buttocks and thighs on the subject, in a pose and muscle activations approximately similar to sitting. Subjects were fully clothed wearing tight-fitting pants. This exposed as little privacy of the subject as possible. For each example, we measured the human in a pose similar to the one in the example, and with similar muscle activations. The measured Young’s moduli ranged from 5,000 Pa to 15,000 Pa. After that, we interpolated captured soft layer stiffnesses and depths to the whole mesh. We then used ElTopo [49] to push the surface mesh inward with the interpolated depths, generating the inner surface of our human model. Finally we tetrahedralized the volume between the outer and inner surface using TetGen [50] to create the tetrahedral mesh for the FEM simulation.

3.2 Modeling supporting surfaces

We model the chair by starting from a $0.5\text{m} \times 0.2\text{m} \times 0.5\text{m}$ cube, which we beveled into a round arc in the knee region to avoid a sharp edge contact below the knee. Although this “chair” shape is simplistic, it reasonably captures the essence of the contact between the human and a chair. Our neutral shoe sole is flat, and was obtained by resizing a commercial shoe sole model purchased online [48]. The neutral mouse model was created by modeling two non-concentric spheres, followed by a CSG union operation and a few iterations of Taubin smoothing [51]. This created a reasonable initial “mouse” shape, reminiscent of ergonomic vertical mice such as the Logitech MX Vertical, to serve as our baseline neutral model. All triangle meshes were sufficiently subdivided to provide enough degrees of freedom for shape deformation. Figure 3 gives our examples.

3.3 Nonlinear elastostatics coupled to the rigid core

The human is simulated via two-way coupling between a 6-DOF rigid object (the “rigid core”) and nonlinear large-deformation elasticity of the deformable layer, modeled as a tetrahedral mesh using FEM. The coupled rigid core + deformable system is subjected to gravity, whereas the rigid supporting surface is not simulated and is pinned in space. This causes the FEM mesh to be in distributed contact against the rigid supporting surface (Figure 2). Our model is designed so that the static contact equilibrium under gravity is differentiable with respect to changing the shape of the supporting surface.

The deformable layer is modeled using three-dimensional FEM nonlinear elasticity; we use invertible isotropic hyperelastic neo-Hookean energy [52] with additional energy terms added for volume preservation [53]. Let $M \in \mathbb{R}^{3n \times 3n}$ be the mass matrix, $u \in \mathbb{R}^{3n}$ the vector of mesh vertex displacements, and $f_{\text{int}} = f_{\text{int}}(u) \in \mathbb{R}^{3n}$ and $K = K(u) = df_{\text{int}}(u)/du \in \mathbb{R}^{3n \times 3n}$ the internal elastic forces and the tangent stiffness matrix, respectively. Assemble contact forces acting on mesh vertices into $f_{\text{contact}}(u) \in \mathbb{R}^{3n}$. Contact forces only apply to vertices on the exterior surface of the tet mesh. We furthermore divide the tet mesh vertices into 2 groups: vertices that are attached to the rigid core (by fixed constraints; the “Attached vertices”, denoted by A) and the rest of the vertices (the “Free vertices”, denoted by F). We split the forces and matrices correspondingly into the A and F components.

Consider an infinitesimally small displacement δu and infinitesimally small rigid body motion of the core $\delta q = [(\delta x)^T (\delta \theta)^T]^T$, where x is the position of the center of mass of the rigid core (in world coordinates), and $\delta \theta \in \mathbb{R}^3$ is an infinitesimally small rotation by angle $\|\delta \theta\|$ around axis $\delta \theta / \|\delta \theta\|$. Because the attached vertices are constrained to the rigid object, we have $\delta u^A = B \delta q$, where

$$B = \begin{bmatrix} I & -\widetilde{h}_1 \\ \vdots & \vdots \\ I & -\widetilde{h}_a \end{bmatrix}. \quad (1)$$

Here, h_i is the vector from the rigid core’s center of mass to the current position of attached vertex i , and integer a denotes the number of attached vertices. For a vector $x \in \mathbb{R}^3$, \widetilde{x} denotes the skew-symmetric matrix with the property $\widetilde{x}y = x \times y$ for all $y \in \mathbb{R}^3$.

We can now state our elastostatic equilibrium equations. Both the elastic object and the rigid object must be in equilibrium under gravity, contact to the supporting surface, and attachment forces f_{attach} between the rigid core and the deformable object:

$$K^{\text{FF}} \delta u^{\text{F}} + K^{\text{FA}} B \delta q + f_{\text{int}}^{\text{F}} + f_{\text{contact}}^{\text{F}} = M^{\text{F}} \hat{g}, \quad (2)$$

$$K^{\text{AF}} \delta u^{\text{F}} + K^{\text{AA}} B \delta q + f_{\text{int}}^{\text{A}} + f_{\text{contact}}^{\text{A}} + f_{\text{attach}} = M^{\text{A}} \hat{g}, \quad (3)$$

$$B^T f_{\text{attach}} + w = 0, \quad \text{for} \quad (4)$$

$$\hat{g} = [0 \ -g \ 0 \ 0 \ -g \ 0 \ \dots \ 0 \ -g \ 0]^T \in \mathbb{R}^{3n}, \quad \text{and} \quad (5)$$

$$w = [F^T \ \tau^T]^T + [0 \ -mg \ 0 \ 0 \ 0]^T. \quad (6)$$

Here, \hat{g} is the gravity acceleration vector (we use $g=9.81$), m is the mass of the rigid core, $w \in \mathbb{R}^6$ contains the external forces and torques acting on the rigid body (the “wrench”),

consisting of the external force $F \in \mathbb{R}^3$ and external torque (around the center of mass) $\tau \in \mathbb{R}^3$ acting on the rigid core, plus gravity. Note that the torque of gravity is zero. In our chair and mouse examples, there are no external forces except gravity. In the foot example, we model the force and torque of the rest of the human body onto the foot, as explained in Section 5.2. Equations 2, 3 and 4 establish the static equilibrium of the free vertices, attached vertices and the rigid core, respectively. By eliminating f_{attach} using Equation 3, we obtain the following linear system for an incremental change in the elastostatic equilibrium:

$$\begin{bmatrix} I & B^T \end{bmatrix} K \begin{bmatrix} I & B \end{bmatrix} \begin{bmatrix} \delta u^{\text{F}} \\ \delta q \end{bmatrix} = \begin{bmatrix} I & B^T \end{bmatrix} \left(M \hat{g} - f_{\text{int}} - f_{\text{contact}} \right) + \begin{bmatrix} 0 \\ w \end{bmatrix}, \quad (7)$$

$$\delta u = \begin{bmatrix} I & B \end{bmatrix} \delta u^{\text{F}}. \quad (8)$$

3.4 Computing the static equilibrium

We perform a Newton-Raphson iteration where we repeatedly solve for δu and δq , followed by a line search along the search direction $(\delta u, \delta q)$. We are solving a difficult nonlinear contact mechanics problem, and as such the presence of the line search is critical. Furthermore, it is critical to perform collision detection *inside each iteration of the line search* (below). We emphasize that doing so is *not* commonly done in computer graphics applications, and has taken substantial experimentation to discover. Without these algorithmic features, the shape design approach in this paper converges poorly; but with the approach as presented here, we were able to consistently produce stable and convergent results.

At each Newton-Raphson iteration, we evaluate the nonlinear internal elastic forces $f_{\text{int}} = f_{\text{int}}(u)$ and matrix $K = K(u)$, and perform collision detection and contact resolution to calculate the collision forces f_{contact} (§3.5), and then solve Equation 7 for $(\delta u, \delta q)$. The line search finds $\eta_{\text{min}} \geq 0$ that minimizes the change of the energy E_{LS} away from the current configuration,

$$\begin{aligned} \Delta E_{\text{LS}}(\eta) &= E_{\text{elastic}}(u + \eta \delta u) - E_{\text{elastic}}(u) + \\ &E_{\text{contact}}(u + \eta \delta u) - E_{\text{contact}}(u) - \eta (\delta u)^T M \hat{g} - \eta (\delta q)^T w. \end{aligned} \quad (9)$$

Here, E_{elastic} is the nonlinear elastic energy of the FEM mesh, and E_{contact} is the collision energy (§3.5). For each η , we perform collision detection again and calculate $E_{\text{contact}}(u + \eta \delta u)$ with respect to the new contact configuration. The last two terms are the potential energy of the external forces and torques, and gravity, respectively. We use the Golden ratio algorithm [54] in our line search. Note that the $E_{\text{elastic}}(u)$ and $E_{\text{contact}}(u)$ terms do not depend on η and can be omitted. Finally, we update $u \rightarrow u + \eta_{\text{min}} \delta u$ and $q \rightarrow q + \eta_{\text{min}} \delta q$.

In the initial configuration, our human is typically positioned slightly above the supporting surface and therefore out of contact. To avoid singularities in the static equilibrium, we replace K with $K + \alpha M$ whenever there are no contacts; α is a tunable parameter; we use $\alpha = 10^3, 10^6, 10^4$ in the chair, shoe sole and mouse examples, respectively. In all our examples, this only occurs during the first few iterations while the human is “falling” under gravity, and is never needed again once the first contact appears.




Fig. 4. **Bidirectional contact model** is needed to correctly detect and penalize sharp features (which are very un-ergonomic). A uni-directional contact model of the top object misses the spike because vertex v is close to the bottom object's surface.

3.5 Contact

We resolve contacts using a custom-designed penalty method, by checking every vertex of the surface of the human tet mesh against the volume enclosed by the supporting surface's triangle mesh. We use the penalty method because it permits us to define contact forces that are differentiable (with respect to the collided point position) everywhere in the three-dimensional space except on a set of measure zero, and whose derivatives we can readily analytically compute. This enables us to define a differentiable ergonomic cost function (§4.1), calculate its gradient with respect to the shape of the supporting surface (§4.2) and then perform supporting surface shape optimization to maximize ergonomic comfort. Before we settled on our method, we tried modeling the volume enclosed by the supporting surface's triangle mesh using a signed distance field, and using it to define penalty forces [55]. We abandoned this approach because it was difficult to define the gradient of the contact force with respect to changing the vertex positions of the supporting surface and human meshes. We note that recent approaches demonstrated how to differentiate cloth in constraint-based contact [56]; however, these approaches were not demonstrated for a contact between solid objects whereby the control parameters consist of the entire shape (vertex positions).

We define our collision energy as follows. If human vertex a with world-coordinate position \mathbf{p}_a collides with the volume enclosed by the supporting surface, we find the closest point on the supporting surface; denote its position $\mathbf{p}_{\text{anchor}}$. Then, a contact spring is placed between \mathbf{p}_a and $\mathbf{p}_{\text{anchor}}$ with a collision energy of

$$E_{\text{contact}} = k_{\text{contact}} \|\mathbf{p}_a - \mathbf{p}_{\text{anchor}}\|^2, \quad (10)$$

where k_{contact} is the penalty stiffness coefficient. We adjusted k_{contact} by running a few static solves against the neutral shape of the supporting surface, until we obtained one whereby the maximum penetration depth in the static equilibrium is approximately 0.3mm. All our optimization results have a maximum penetration depth of approximately 0.3mm. This value is small and on the order of accuracy of optical surface scanners.

The penalty contact forces and their gradient can then be computed by the gradient and the hessian of E_{contact} . To accelerate the convergence of the static solver, it is very important to incorporate into the gradient the fact that $\mathbf{p}_{\text{anchor}}$ slides on the supporting surface as \mathbf{p}_a changes location. To this end, we derive and use three different gradients, depending on whether $\mathbf{p}_{\text{anchor}}$ is a vertex, on the interior

of an edge, or on the interior of a face of a triangle of the surface mesh of the supporting surface (Appendix A).

The non-differentiable set of measure zero consist of the boundaries of Voronoi regions of mesh vertices, edges and faces, and of the mesh itself. Although non-differentiable at those points, differentiability elsewhere greatly accelerated our optimization. The non-differentiability is dealt with by employing a line search, as explained in Section 3.4. We initially tried using a simpler contact formulation where we omitted the sliding terms from the gradient. Although this still gave us convergent optimization results, the sliding anchor method and contact gradients presented herein gave us an order of magnitude speedup in our optimizations.

4 OPTIMIZATION

We now describe how we optimize the shape of the supporting surface to maximize the ergonomics of contact with the human.

4.1 Ergonomic Cost Function

As suggested by prior work [32], [33], contact pressures and their spatial distribution is an objective way to measure comfort. How can a given pressure distribution be converted into a single, well-defined scalar cost function that both effectively measures ergonomics and can be efficiently computed and differentiated for shape optimization? The first idea may be to integrate the contact pressures, $\int_{\Gamma} p dS$, over the contact area Γ . However, this does not work because, for flat (or nearly flat) contacts, this integral approximately equals the magnitude of the gravity force of the human body, therefore having no differentiation power over different pressure distribution. After experimenting with different formulations, we arrived at our definition of the ergonomic cost function ζ ,

$$\zeta = \int_{\Gamma_{\text{human}}} p^3 dS + \int_{\Gamma_{\text{supporting surface}}} p^3 dS. \quad (11)$$

Observe that pressures p are non-negative scalars by definition, hence no norm or absolute value is needed on p . Note that we integrate both over the contact area on the human and that on the supporting surface. In a continuous setting these contact areas are the same, and the pressure distributions are also the same due to Newton's 3rd law. However, when Equation 11 is discretized over the contact areas on the human and supporting surface, the two integrals in the equation discretize differently. It is important to include both summation terms ("bi-directional contact"), otherwise one of the contact surfaces is permitted to become spiky without any penalization (Figure 4). Note that we use a uni-directional contact model (§3.5) in the static equilibrium calculations, but a bi-directional contact model in our ergonomic cost function (Equation 11). This is because the Hessian of E_{contact} (Equation 10) is symmetric positive-definite for uni-directional contact, but becomes indefinite under bi-directional contact, which decreases simulation stability (§3.4). In the remainder of this paper, we will use $f_{\text{contact}}^{\text{uni}}$ and $f_{\text{contact}}^{\text{bi}}$ when referring to the uni-directional contact forces on the human tet mesh vertices ($f_{\text{contact}}^{\text{uni}}$ from §3.5), and bi-directional contact forces on the union of human

tet mesh vertices and vertices of the triangle mesh of the supporting surface, respectively.

Observe that our ergonomic cost function measures *discomfort*: meaning, higher values are less favorable ergonomically. Although it would be in principle “cleaner” to define a score function that is lower and higher when ergonomics is less and more favorable, respectively (say, by composing ζ with $x \mapsto -x$ or $x \mapsto 1/x$), doing so unnecessarily complicates the formulation without a clear gain. Hence, we adopted our “discomfort” measure ζ of Equation 11. Finally, the exponent 3 in Equation 11 could be replaced with some other suitable exponent such as, say 2 or 4 (see Figure 5), but, as argued, not 1. Higher exponents penalize localized pressures more aggressively, but are less stable due to discretization of contact.

Assemble vertex positions of the supporting surface into a vector $x \in \mathbb{R}^{3m}$. Our static equilibrium contact pressures are an implicit function of x , and therefore so is ζ . After discretizing Equation 11, we obtain

$$\zeta(x) = \sum_i s_i \bar{p}_i^3, \quad \bar{P} = (I - L)P, \quad (12)$$

$$P = P(y(x), x) \quad y(x) = [(u^F)^T(x), q^T(x)]^T, \quad (13)$$

where we have joined the two integrals of Equation 11 into one sum: i runs over each contact vertex on the human body mesh, and on the supporting surface (Figure 4). Quantity y contains the converged displacements and rigid body transformation of the human body in static equilibrium, s_i is the surface area associated with vertex i (either on human or supporting surface), and L is the 2×2 block-diagonal matrix whose upper-left and lower-right blocks are the umbrella Laplacian matrix of the human body and the supporting surface’s triangle mesh, respectively. Vector P contains contact pressures; its i -th component $p_i = f_{\text{coll},i}/s_i$ is the pressure computed at vertex i . Vector P depends on both y and x because the contact pressures depend on positions of the vertices of *both* the human and the supporting surface. In order to decrease mesh discretization artifacts, we include a Laplacian smoothing factor $I - L$ in the definition of ζ ; namely, vector $\bar{P} = (I - L)P$ contains the smoothed pressures; and \bar{p}_i is its component at vertex i .

To promote smoother optimized supporting surfaces, we also add a smoothing score $\psi(x)$ to the cost function. We initially formulated our optimization without this term, but we observed a small amount of noise, mostly in our mouse and foot examples when running from the negative initial guess. We define the smoothing score as the difference between the original shape x and the shape smoothed using Taubin smoothing [51],

$$\psi(x) = \|x - ((I - \mu L_{\text{tri},3})(I - \lambda L_{\text{tri},3}))^m x\|^2, \quad (14)$$

where $L_{\text{tri},3}$ is the three-dimensional umbrella Laplacian matrix of the supporting surface’s triangle mesh, m is the number of Taubin smoothing iterations, and λ and μ are Taubin smoothing parameters. The smoothing score only depends on the shape of the supporting surface; contacts do not affect this score. Finally, we define our combined objective function,

$$\phi(x) = \zeta(x) + \gamma \cdot \psi(x), \quad (15)$$




Fig. 5. **Sensitivity of our ergonomic cost to surface bumpiness.** We sit a human into a chair that has a bump. We then vary the spatial frequency (“spikiness”) of the bump and compute our ergonomic cost, under various cost exponents. For easier viewing, the scores are linearly re-scaled so that they are equal at the lowest frequency. It can be seen that all exponents are discriminative.

where γ is the strength of smoothing. Parameter γ was determined experimentally, by comparing $\zeta(x)$ and $\psi(x)$ on the initial guess, and setting κ so that $\zeta(x)$ and $\kappa\psi(x)$ have approximately equal magnitudes. We then explored several more values γ , computed by multiplication with 10^j , for a few integers $j \in \{-J, \dots, 0, \dots, J\}$. We used $J = 4$ in our examples. Further, we found that a small amount of additional post-process smoothing improved both the visual appearance and the ergonomic score in some scenarios (chair example in Section 5.1).

We minimize ϕ using nonlinear conjugate gradients combined with a line search. At each iteration, our optimization first computes the gradient of ϕ with respect to x , and projects it against previously explored search directions (as standard in nonlinear conjugate gradients [54]). This produces the conjugate gradient search direction d_i . We then perform a 1-dimensional line search

$$\min_{\kappa} \phi(x_i + \kappa d_i), \quad (16)$$

where x_i is the shape of the supporting surface at iteration i . The scalar $\kappa \geq 0$ is the step length in the line search. Note that ζ must be re-evaluated for each $x_i + \kappa d_i$, by computing the static equilibrium against the supporting surface with vertex positions $x_i + \kappa d_i$ (§3). To simplify the optimization, we use a pre-defined set of eight κ values, and test each one independently, which can be done by launching eight independent static equilibrium calculations in parallel. We set the eight κ values dynamically based on d_i , so that the maximum positional change of the supporting surface vertices equals $1\mu\text{m}$, $10\mu\text{m}$, 0.1mm , 0.2mm , 0.5mm , 0.7mm , 1mm and 2mm . Typically, the supporting surface shape changes little as the optimizer “probes” different values of κ . Therefore, in order to speed up convergence, we re-compute the static equilibrium starting from the values of u and q discovered while computing the previous iteration x_i .

4.2 Objective Function Gradient

The gradient of $\partial\psi/\partial x$ is easy to compute because ψ is a quadratic function of x . The gradient $\partial\zeta/\partial x$ can be computed using the chain rule,

$$\frac{\partial\zeta}{\partial x} = \frac{\partial\zeta}{\partial f_{\text{contact}}^{\text{bi}}} \left(\frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial x} + \frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial y} \frac{\partial y}{\partial x} \right). \quad (17)$$

The difficult term to compute in the equation above is $\frac{\partial y}{\partial x}$. We compute it using the implicit function theorem [25], [26], [27], but adapted here for our rigid core elastostatic formulation. When the system is in static equilibrium, the right-hand side of Equation 7 is zero. Let δx be an infinitesimal change of the supporting surface shape. Then, the contact force changes infinitesimally by $(\partial f_{\text{contact}}/\partial x)\delta x$, and the updated static equilibrium satisfies

$$\hat{K} \begin{bmatrix} \delta u^F \\ \delta q \end{bmatrix} = - \begin{bmatrix} I \\ B^T \end{bmatrix} \frac{\partial f_{\text{contact}}^{\text{uni}}}{\partial x} \delta x, \quad (18)$$

where \hat{K} is the system matrix of Equation 7. Finally, we combine Equations 17 and 18,

$$\frac{\partial\zeta}{\partial x} = \frac{\partial\zeta}{\partial f_{\text{contact}}^{\text{bi}}} \left(\frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial x} - \frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial y} \hat{K}^{-1} \begin{bmatrix} I \\ B^T \end{bmatrix} \frac{\partial f_{\text{contact}}^{\text{uni}}}{\partial x} \right). \quad (19)$$

In order to compute the gradient, we therefore only need to solve a single linear system

$$\hat{K}z = \left(\frac{\partial\zeta}{\partial f_{\text{contact}}^{\text{bi}}} \frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial y} \right)^T. \quad (20)$$

4.3 Deformable Supporting Surface

Our method can also be adapted in optimizing a deformable supporting surface. To do this, we first add all DOFs of vertices on the deformable supporting surface to the system variable u^F . Since the contact force depends on the deformed positions of the vertices rather than their rest positions, we need to modify Equation 13 to

$$P = P(y(x)) \quad y(x) = [(u^F)^T(x), q^T(x)]^T. \quad (21)$$

The gradient of the ergonomic score becomes

$$\frac{\partial\zeta}{\partial x} = \frac{\partial\zeta}{\partial f_{\text{contact}}^{\text{bi}}} \frac{\partial f_{\text{contact}}^{\text{bi}}}{\partial y} \frac{\partial y}{\partial x} \quad (22)$$

Similarly to rigid supporting surfaces, when the system is in static equilibrium, the right-hand side of Equation 7 is zero, and we get

$$\hat{K} \begin{bmatrix} \delta u^F \\ \delta q \end{bmatrix} = - \begin{bmatrix} I \\ B^T \end{bmatrix} \frac{\partial f_{\text{int}}}{\partial x} \delta x. \quad (23)$$

We therefore need to solve a linear system of equations equal in form to that of Equation 20.

5 RESULTS

We give three examples demonstrating ergonomic contact between a human and a rigid supporting surface (Figure 3). Our optimization is able to reduce the ergonomic score by large percentages (Table 1), resulting in pressure distributions whereby the ‘‘hot-spot’’ pressures are reduced 2-3x, and their carrying load evenly distributed across the entire

Supporting surface	Chair	Shoe sole	Mouse
Initial guess shape	Flat plane	Flat plane	Negative image
Initial guess score	4.99×10^{11}	2.99×10^{13}	5.48×10^7
Optimized score	1.39×10^{11}	5.36×10^{12}	1.37×10^7
Percentage reduced	72.07%	82.10%	74.96%
Smoothing strength	1×10^{13}	1×10^{19}	1×10^{13}
Optimization time	4 hours	19 hours	13 hours
#Human vtxs	10,231	14,060	77,749
#Supp. surface vtxs	26,891	5,894	13,614
#Contact vtxs	1,285	575	2,171

TABLE 1

Ergonomic scores and optimization timings. Human (deformable), supporting surface (rigid) and contact vertex counts refer to the tet mesh, triangle mesh and the static equilibrium, respectively. The human exterior surface is the tet mesh’s exterior surface.

contact surface. All examples were computed on a 3.00 GHz Intel Xeon i7 CPU E5-2687W v4 processor. We used 8 threads for optimization as we attempt 8 step lengths (κ) in one line search (Equation 16).

5.1 Human sitting in a chair

We designed an ergonomic chair for a specific person. We scanned and simulated the subject’s lower back, buttocks and upper thighs region. The other parts of the body are modeled as rigid because they do not deform in our scenario. We acquired the depths and Young’s moduli at different body locations, and created the rigid core by offsetting the external shape using the spatially varying depths [49]. The rigid core can rotate arbitrarily. Our chair is generated from a smoothed flat block as initial guess. The human sits in the central region of the chair, and it is not our intention to permit horizontal sliding. Therefore, we lock the two directions of horizontal translation, and only permit the rigid core to translate along the gravity direction. Figure 6, top, shows that our optimization can greatly reduce the ergonomic score (by 72.1%). The output pressure distribution is also clearly better on the optimized chair. Note that we perform similar locking of translations in all of our examples. Doing so prevents sliding and therefore we did not need to employ friction forces in our contact model or the shape optimization, which makes the optimization simpler, faster and more tractable. In this particular example, we found that adding a small amount of additional Taubin smoothing ($\mu = -0.53, \lambda = 0.5$) to the final optimized chair smoothed out minor numerical irregularities during optimization and further reduced the ergonomic score (from 3.62×10^{11} to 1.39×10^{11}). All of our chair examples have this smoothing applied before computing pressures and ergonomic score. This postprocessing was not significant in the other examples (mouse, shoe) and was omitted.

The above result was obtained starting from a flat chair as an initial guess. Another potential initial guess is the ‘‘negative image’’ chair obtained by subtracting the human body from a flat chair using CSG boolean subtraction. We create the negative image by penetrating the human into the chair by 5cm. To remove the sharp edges caused by CSG, we use Laplacian smoothing. Figure 6, bottom, shows that the negative image chair, when used directly as the chair shape, has a bad ergonomic score because of the high pressures along the edges. Although we smoothed the edges, the

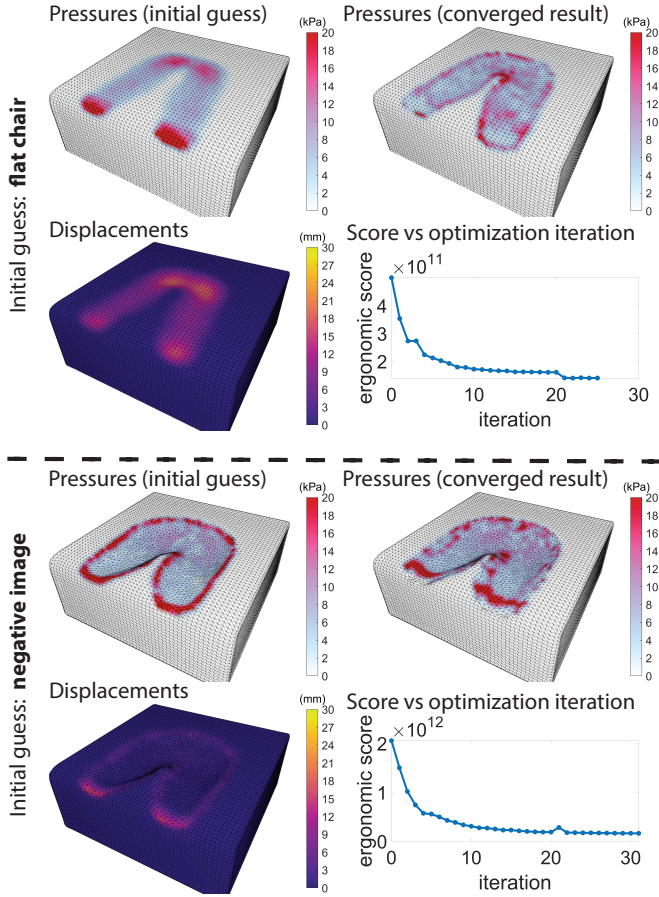


Fig. 6. **Optimized pressure maps and ergonomic scores (chair).** “Displacements” gives the magnitude of vertex displacements between the initial guess and the converged results. When using the flat chair initial guess (top), the ergonomic score was reduced from 4.99×10^{11} to 1.39×10^{11} (72.1%); the maximum penetration depth of the optimized chair was 0.15mm. When using the negative image initial guess (bottom), the ergonomic score was reduced from 2.05×10^{12} to 1.58×10^{11} (92.3%); the maximum penetration depth of the optimized chair was 0.16mm.

human body will still generate high pressures along the edges because of the deformation. When starting from the “negative image” initial guess, our optimization method can reduce the pressure substantially (92.3%), however the end result is inferior to optimization that started from a flat chair (Figures 6). We also tried other penetrations: 1cm, 2cm, 3cm, 4cm, 6cm; they all produced inferior results. The 6cm was unstable because the negative image was too deep, creating a concave chair.

5.2 Human standing on a shoe sole

We purchased the triangle meshes of a human foot and a flat shoe sole from [48]. We tetrahedralized the human foot using our method described in Section 3.1. To simulate a real human standing on the shoe sole, we model the presence of the upper human body using an external gravity force equaling to half of the body mass. Note that the human body’s center of mass is typically not on the vertical line through the rigid core’s center of mass; we model this effect by adding an external torque applied to the rigid core (Figure 7). By tweaking this torque, we can simulate the effect of




Fig. 7. **The rigid core model of the foot.** Left: The outer layer and the inner layer of the foot. The inner layer forms a rigid core of the foot. Right: The deformable tetrahedra form the outer layer.

the human body leaning forwards or backwards. Although our system does not incorporate motor control, setting a specific value of this torque (either positive or negative) will shift the contact pressure distribution on the shoe sole to compensate for the leaning body. This effectively simulates a stably standing human, and is possible as long as the external torque is consistent with a body center of mass being within the shoe sole surface area. Similar to the chair example, we locked the two directions of horizontal foot translation. In addition, because we are not interested in yaw and roll motion of the foot, we only allow the rigid core of the foot to rotate around the pitch axis.

Similar to the chair model, we also use a flat shoe sole and a negative image as our initial guesses (Figure 8). Our results demonstrate that the optimizer works well on the flat shoe sole (82.1% ergonomic score reduction). When starting from the negative image, the score reduction is smaller (29.0%). Overall, the best score was obtained starting from a flat shoe sole.

5.3 Human grasping a mouse

In this example, we explored how to create an ergonomic shape of the mouse as it contacts the human hand. The mouse neutral shape was modeled using CSG operations on spheres, followed by Taubin smoothing. We rigged the human hand using linear blend skinning, and posed it so that it sits on the mouse in a natural configuration. Although we do not articulate the human hand in our system, or model the biomechanics of the upper arm, our model does permit calculating the contact pressures experienced by the hand due to gravity, as it rests on top of the mouse. We locked all hand rotation axes and only permit the hand to translate up and down. Our optimizer can greatly reduce the ergonomic score for both the neutral mouse (93.5%) and negative image (75.0%) initial guesses. The negative image method produced a better ergonomic score than the neutral mouse (Figure 9).

5.4 Human sitting in a deformable chair

We also demonstrate that our method can be adapted to designing the rest shape of a deformable supporting surface. We used the same scanned human model as in the rigid

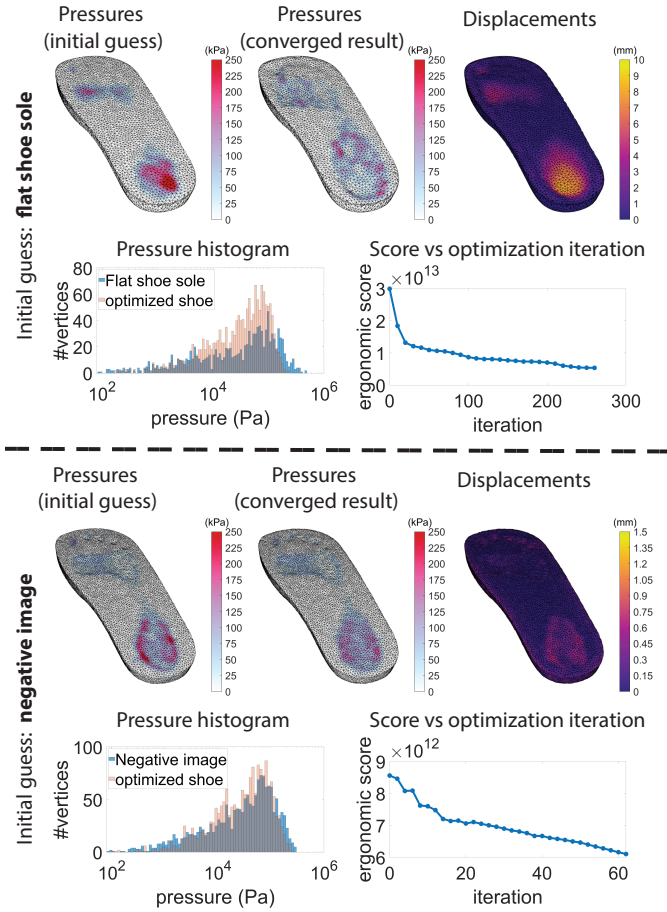


Fig. 8. Optimized pressure maps and ergonomic scores (shoe sole). “Displacements” gives the magnitude of vertex displacements between the initial guess and the converged results. When using the flat shoe sole initial guess (top), the ergonomic score was reduced from 2.99×10^{13} to 5.36×10^{12} (82.1%); maximum penetration depth of the optimized shoe sole was 0.30mm. When using the negative image initial guess (bottom), the ergonomic score was reduced from 8.56×10^{12} to 6.08×10^{12} (29.0%); maximum penetration depth of the optimized shoe sole was 0.52mm. In both results, it can be seen that the pressure histogram shifts to the left (i.e., to smaller pressures). The bins at smaller pressures also became taller as the optimized shoe sole has a larger number of contacts, each carrying a smaller contact pressure. The grey color in the pressure histogram denotes the overlap of the blue color (negative image) and the pink color (optimized shoe).

chair example, and we also used a smoothed flat box as the initial guess. Like in the rigid case, we locked the two directions of horizontal translations. The top part of Figure 10 shows that our optimization decreases the ergonomic score of a deformable chair (50kPa Young’s Modulus) by 63.1%. The output pressure distribution is also clearly improved. The vertex displacements between the initial guess and the converged result are smaller than in the rigid chair example because the deformable material of the chair already accommodates some ergonomic “molding” to the human body. Because the rest shape of the chair becomes less important for ergonomic design (as it exhibits smaller variability), the number of iterations in our optimization becomes smaller.

We also modified the Young’s Modulus of the deformable chair to 20kPa (2.5x softer) and 125kPa (2.5x stiffer), to investigate how the chair material properties affect the results (middle and bottom parts of Figure 10). We

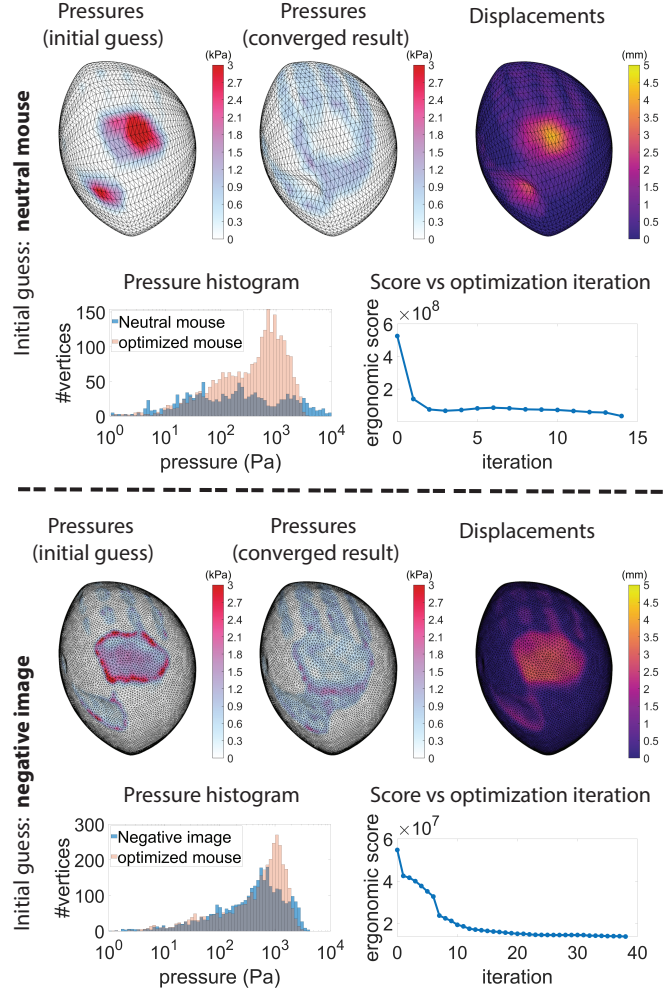


Fig. 9. Optimized pressure map and ergonomic scores (mouse). “Displacements” gives the magnitude of vertex displacements between the initial guess and the converged results. When using the neutral mouse shape (top) as the initial guess, the ergonomic score was reduced from 5.25×10^8 to 3.34×10^7 (93.5%); maximum penetration depth of the optimized mouse was 0.25mm. When using the negative image shape (bottom) as the initial guess, the ergonomic score was reduced from 5.48×10^7 to 1.37×10^7 (75.0%); maximum penetration depth of the optimized mouse was 0.13mm. Similarly to the shoe sole example, the pressure histograms shifted to the left (i.e., to smaller pressures) and became taller, due to larger number of contacts each carrying a smaller contact pressure; very high pressures disappeared (note: logarithmic pressure scale).

observed that (as expected) the softer chair results in a lower (i.e., better) ergonomic score. Also, the rest shape of the soft chair changed less from its initial shape, as the chair’s greater deformability partially enables a good ergonomic outcome on its own. More rigid chairs, however, need a greater change in their rest shape to decrease the ergonomic score.

5.5 Smoothing Strength

In order to evaluate the effect of smoothing (γ) values to our optimization, we optimized the mouse with different γ values, 0, 10^{11} , 10^{13} , 10^{15} and 10^{17} . Each γ gives a different optimized shape (Figure 11) and ergonomic score (Table 2). The smoothing strength of 10^{13} gives the best optimized ergonomic score.

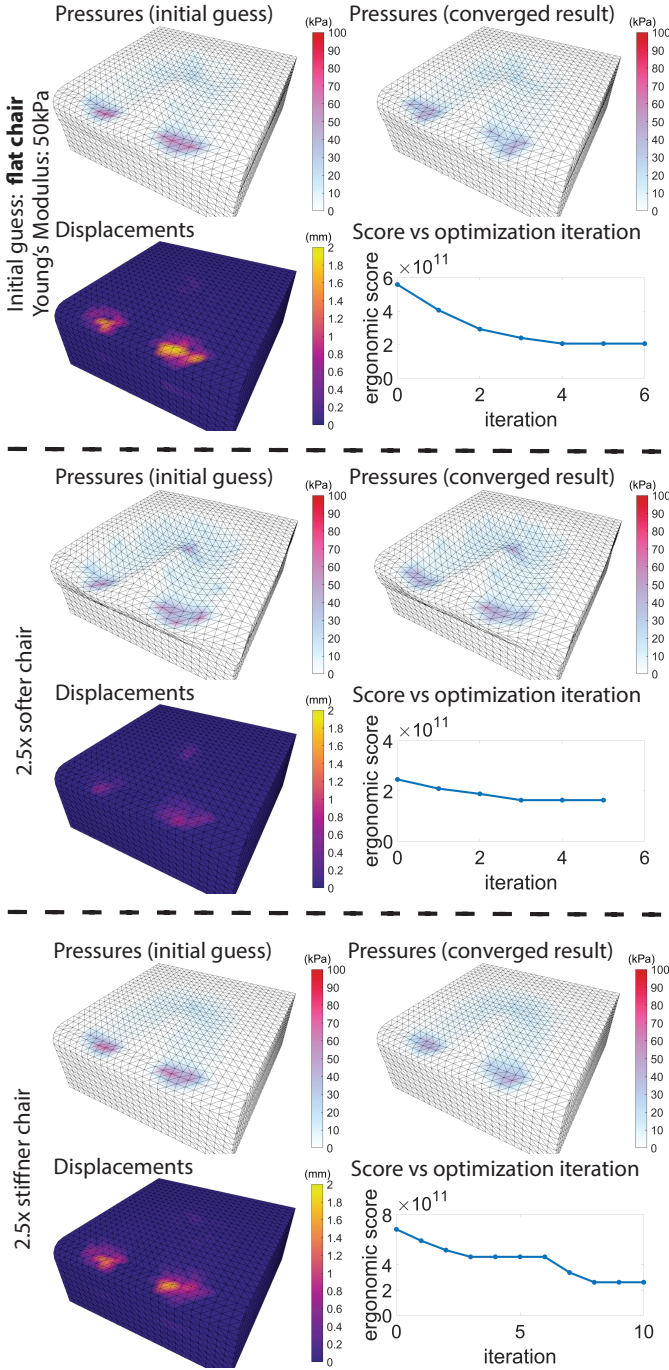


Fig. 10. **Optimized pressure maps and ergonomic scores (deformable chair).** “Pressures” gives the pressure field on the deformed chair when the simulation reaches the static equilibrium. “Displacements” gives the magnitude of vertex displacements between the rest shape of the initial guess and the rest shape of the converged result. All experiments use the flat chair initial guess. Top: under a Young’s modulus of 50kPa, the ergonomic score was reduced from 5.59×10^{11} to 2.06×10^{11} (63.1%); maximum optimized penetration depth was 0.20mm. Middle: softer chair (25kPa). The ergonomic score was reduced from 2.46×10^{11} to 1.63×10^{11} (33.6%). Bottom: stiffer chair (125kPa). The ergonomic score was reduced from 6.83×10^{11} to 2.62×10^{11} (61.7%).

6 PHYSICAL VALIDATION OF OUR METHOD

We validated our simulation results using a physical experiment where we compared pressure distributions measured

in the real-world to those predicted by our simulation method, both for the initial guess supporting surface and the optimized supporting surface. In doing so, we designed a new inexpensive experimental method to measure real distributed contact pressures between two curved surfaces. Doing so is challenging because one has to somehow insert measuring equipment between the two curved surfaces, but the presence of the sensors itself potentially changes the distributed contacts, and thus interferes with the contact measurements. Initially, we attempted to model the pressures using thin pressure measurement film (paper) from FujiFilm [57]. However, this did not work well because the optimized shoe is a general surface with non-zero Gaussian curvature (i.e., is not a developable surface). Consequently, when we placed the film in between the foot and the optimized shoe, the film warped, which interfered with the pressure measurements. Instead, we then used electrical resistance-based pressure sensors from Tekscan [58]. These sensors have a shallow profile and minimally interfere with pressures. By designing proper electrical circuits as described below, we were able to reliably measure the contact pressure distribution on the entire shoe for less than \$500.

First, we used the Alja-Safe material [59] to make a plaster model (Figure 12) of a real human foot, and scanned it with a 3D scanner (Artec Spider [60]). We then ran ergonomic shoe sole optimization against the real foot shape, as described in our paper. The ergonomic score was reduced by 82.8%, and the pressure map was visibly improved (Figure 14). After the optimization, we 3D-printed the optimized rigid shoe sole. We then asked the same subject to stand on a flat rigid surface, and on the optimized shoe sole. We used the Tekscan [58] A101 force sensor to measure real pressures at 78 sample locations on the shoe sole (Figure 12). This sensor can measure force in the range from 0N to 44N, and its diameter is 3.81 mm, which is suitable for our experiments. The Tekscan force sensor changes its resistance when a force is applied based on a known linear profile provided in manufacturer’s datasheet. We used the manufacturer-provided dedicated circuit to convert the force into a voltage signal and acquired the voltages into the computer using a digital acquisition card. We used a thin electrical strap provided by Tekscan to electrically connect the two pins on the force sensor to the Tekscan measuring electric circuit so as to cause minimal interference with the mechanical contact. The connector provided by Tekscan was too thick (~ 2 mm profile), and, as it is in immediate proximity of the sensor, found itself “jammed” in between the foot and the shoe, interfering with the contact pressures. We removed

Smoothing strength	Ergonomic score	Percentage reduced
0	5.04×10^7	90.4%
10^{11}	4.70×10^7	91.1%
10^{13}	3.40×10^7	93.5%
10^{15}	1.41×10^8	73.1%
10^{17}	7.94×10^8	-51.1%

TABLE 2

Ergonomic scores using different smoothing values. The initial scores for the five optimizations are all 5.25×10^8 . The entry in the last row is negative because the smoothing strength is too large, causing the optimizer to overly smoothen the shape and compromise ergonomics.

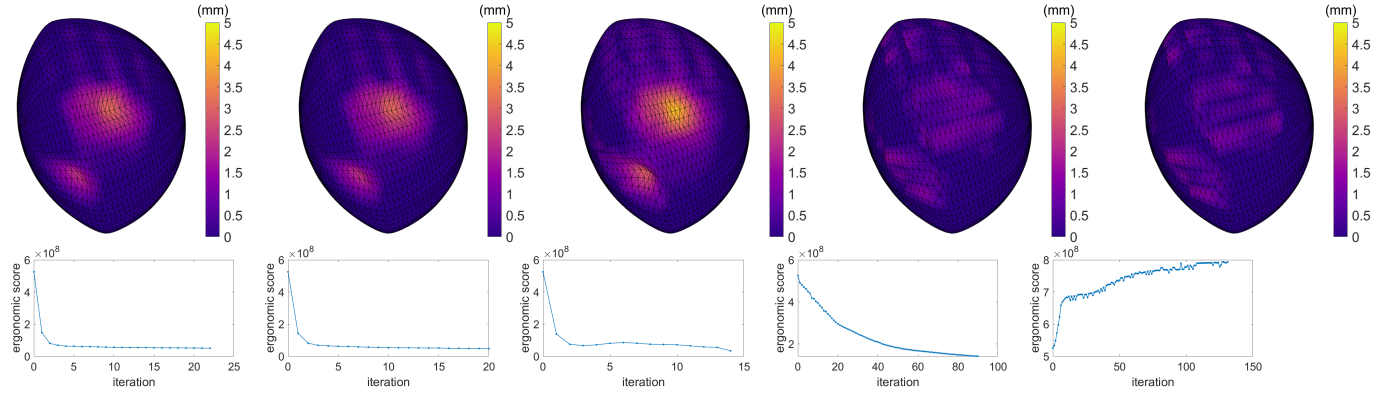


Fig. 11. Optimized shapes and ergonomic scores over the iterations of the optimization, using different smoothing values. Top row: optimized shape with color-coded displacements. Bottom row: ergonomic scores at each optimization iteration. The smoothing values used from the leftmost column and the rightmost column are 0, 10^{11} , 10^{13} , 10^{15} and 10^{17} , respectively.

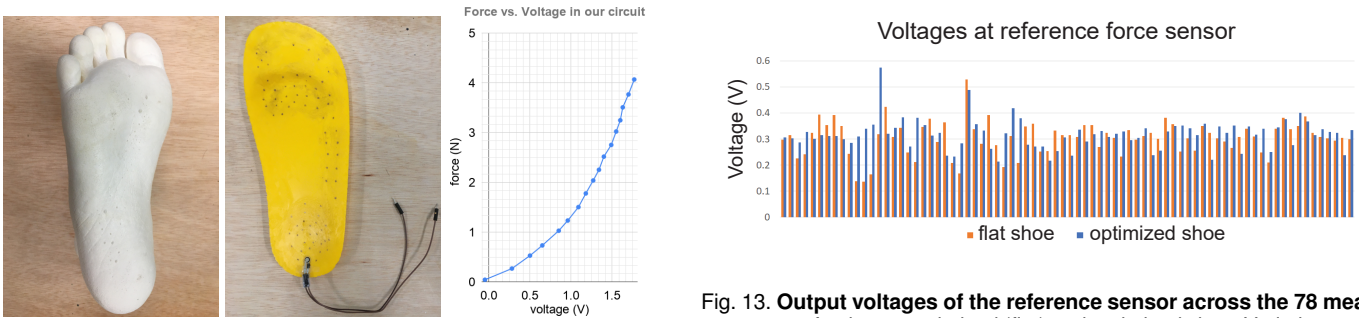


Fig. 12. **Manufactured models and the force sensor.** Left: the plastic foot (photo not to scale) manufactured using the Alja-Safe cloning process. This plastic foot is used only to stably scan the foot 3D shape; it is otherwise unused in our method. Middle: the 3D-printed plastic optimized shoe sole with sample locations marked. The reference force sensor is seen fixed at the bottom of the shoe sole. Right: the relationship between the contact force and output voltage.




Fig. 13. **Output voltages of the reference sensor across the 78 measurements** for the un-optimized (flat) and optimized shoe. Variations are relatively small.

the connector and manually resoldered the wires, and then gently sanded the soldered electrical connections down to a $\sim 0.3\text{mm}$ profile using sand-paper. We measured each of the 78 locations one by one. This kept our design inexpensive, and also reduced contact interference. The Tekscan measuring circuit samples the contact force at the sensor at 10Hz. For each measuring location, the subject stood in the shoe for 1 minute, which gave us 600 pressure points for each sample location. The subject was standing on both legs and was told to keep a stable neutral standing posture. Out of the 600 pressures, we discarded the smallest and largest 50 values to reduce the noise. We then computed the average pressure and standard deviation at each sample location. We then interpolated the pressures onto the vertices of the shoe sole using linearly precise biharmonic weights [61], whereby each of the 78 sensor locations serves as a “handle”.

The contact pressures can be somewhat affected by variations in the human standing pose. To address this, we placed a second reference force sensor with an independent circuit at a fixed location on the shoe sole (Figure 12). We recorded a one-minute pressure signal for the reference sensor and computed its average pressure at the fixed location. For each of the 78 measuring points, this gives us a primary pressure p_i at the measuring location, and a pressure r_i at the

reference location. We then scale pressures p_i with a scalar α_i , where α_i is chosen so that $\alpha_i r_i$ is the same constant C for all i ; this compensates for variations in the leaning postures across the 78 measurements. We choose constant C so that, after the pressures $\alpha_i p_i$ are interpolated to the entire shoe sole (using linearly precise biharmonic weights [61]), the total net contact force on the foot equals half of the gravity force of the human subject. This convenient choice of C also avoids the need for a separate calibration step to calibrate the output voltages versus sensor resistance. Generally, we want the reference sensor pressures to be close to each other when the primary measuring sensor is placed at different locations, i.e., posture variations should be small; and this is indeed the case (Figure 13). We plot the resulting pressures in Figure 14. It can be seen that (1) the real and simulated pressures on the *un-optimized* shoe sole are qualitatively similar to each other, and (2) the real and simulated pressures on the *optimized* shoe sole are qualitatively similar to each other, and (3) the un-optimized shoe sole has substantial pressure hot-spots whereas the optimized pressures does not have hot-spots and the pressures are spread over a large surface area.

7 COMPARISON TO “SIT & RELAX”

We compared our work with the “Sit & Relax” method proposed by Leimer et al. [8]. This method conceptually immerses the supported body part into its negative image volume, in order to estimate a pressure distribution across the surface. During optimization, the supporting surface




Fig. 14. **Comparison between simulation and real measurements.** Top row: un-optimized (flat) shoe sole. Bottom row: optimized shoe sole. Left column: pressure computed using simulation. Middle column: pressure obtained using real measurements. Right column: standard deviation of the pressure obtained using real measurements during the 1-minute acquisition.

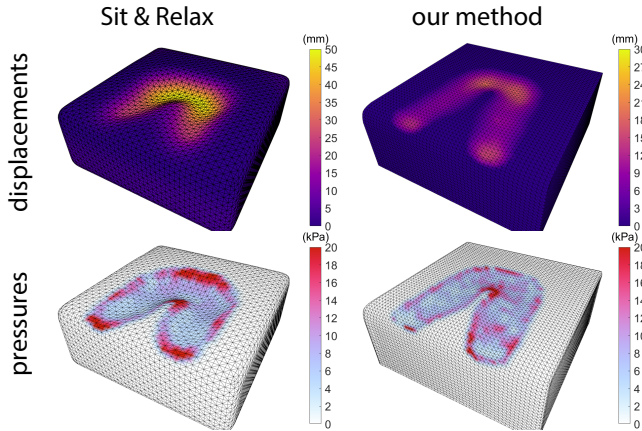


Fig. 15. **Comparison to Sit & Relax [8].** Left and right columns give the results optimized with Sit & Relax and our method, respectively. “Displacements” gives the magnitude of vertex displacements from the initial guess, and “pressures” are the contact pressures computed using physically based simulation against the output of either method. Because our method incorporates a physically based simulation, it gives a distribution with far fewer high-pressure points.

is raised to make contact with the surface of the body part, with each point’s affinity weighted by the corresponding negative-image pressure. The smoothness term and an underlying Catmull-Clark spline model act as counterbalancing regularizers. Since the pressure map is computed for a support surface that is never actually realized (the negative image), and is frozen at the beginning of the optimization instead of varying with the evolving support surface, the Sit & Relax approach leads to un-ergonomic high-pressure regions, visualized in red in Figure 15. In contrast, our method recomputes the pressure map in each iteration, using a physically realistic, differentiable model of deformable-rigid contact, and achieves a much more evenly distributed pressure map without sharp peaks. The ergonomic score (Equation 11) for the Sit & Relax chair is 8.40×10^{11} . In contrast, the score for our optimized chair (starting from a flat surface) is 1.39×10^{11} , or 83.4% less. We also experimented with initializing our method with

the optimized Sit & Relax surface. This produced a score of 2.36×10^{11} , still not quite as good as with the flat initialization. We conjecture that starting with an unbiased flat surface allows the gradient descent to reach a better local minimum in this case. Note that these visualizations and scores assume that our simulation model closely approximates real-world pressure distributions and hence can be used as a basis for comparison of the two methods. We validated this in Section 6. Since the Sit & Relax optimization involves no physical simulation in the inner loop, it is certainly orders of magnitude faster than ours. Speeding up our optimization is an interesting avenue for future work; for example, by parameterizing our supporting surface with Catmull-Clark splines.

8 COMPARISON TO “POSE TO SEAT”

A more recent method by Leimer et al. – “Pose to Seat” [9] – also optimizes supporting surfaces via inferred pressure maps. The main contribution of this paper, compared to “Sit & Relax”, is handling an articulated human model. For their contact model, they assume *all* downward-facing vertices are in contact, effectively modeling the pressure distribution on a negative image mold and keeping it unchanged throughout the geometric optimization of the support surface, just like Sit & Relax. In determining the pressure distribution, they do not consider the other contact object (the supporting surface), deformable elasticity, or contact simulation. In contrast, our method does not handle articulation, but *does* model a (differentiable) physically-based contact pressure map specific to a supporting surface, and the distributed pressure map is updated repeatedly in the inner loop of our shape optimization.

Since these aspects (articulation vs distributed contact) are orthogonal aspects of the problem, it is difficult to do an apples-to-apples comparison with this paper. Further, the support synthesis step of Leimer et al. [9] (which uses the fixed pressure distribution as a precomputed importance map but is otherwise completely non-physical) is designed to support the *entire* human body, using a handcrafted template with pre-assigned patches per body part. This algorithm is highly specific and it is not clear how it could be adapted to our scenarios.

Hence, we chose to compare to Leimer et al. [9] by studying the pressure maps produced by following their assumptions, without considering articulations. To aid their method, we explicitly computed more realistic sets of contact vertices for their method, by bringing the (non-elastic) foot in contact with candidate supporting surfaces and marking the foot vertices within a small threshold distance of the supports. In this fashion, we ran their method with two different sets of body vertices in contact (Figure 16). We transferred their contact pressures to the shoe sole for visualization and comparison to our method (Figure 16). Our experiment demonstrates that our pressure maps are much closer to the real measured pressures than those obtained using [9], despite our contact modeling improvements to their method (Figure 16). Considering all downward-facing vertices to be in contact, as the original Leimer et al. “Pose to Seat” approach does, yields a nearly uniform, even less




Fig. 16. **Comparison to Pose to Seat [9].** The left, middle and right columns give the pressure maps on the shoe sole computed by Leimer et al. [9], real-world experiments and our optimization method, respectively. While Leimer et al. assume *all* downwards-facing vertices are in contact, we aided their method by supplying more realistic contact regions (without considering elasticity). We investigated two contact sets of vertices: “flat sole” and “optimized sole” give the contact pressures between the foot and the flat shoe sole, and between the foot and our optimized shoe sole. Because our method incorporates physically based simulation and our human foot has a deformable layer, our method produces contact pressures that are qualitatively and quantitatively closer to the measured pressures.

realistic pressure distribution that is independent of the actual supporting surface.

9 CONCLUSION

We presented a method to design the shape of supporting surfaces to maximize the ergonomics of their contact against the human body. While previous ergonomic work mostly focused on geometric design, we contribute a physically grounded approach that models elastic contact using large-deformation FEM nonlinear elasticity and uses measured real human tissue properties. We validated our approach using 3D printing and confirmed that the optimized pressure distributions qualitatively match those observed in the real world.

We modeled our human as a passive deformable object in a representative pose. While this proved to be a physically plausible and computationally tractable model, method accuracy could be improved by modeling articulated human bodies and adding active human motor control. To make the optimization tractable and differentiable, we model contact using a penalty-based scheme. Although our penalty forces are not differentiable at the point where a vertex leaves contact, the non-differentiable set is small (has measure zero) and we were in practice able to mitigate the lack of differentiability by taking sufficiently small steps in our line search. Constraint-based contact models are a good avenue for future work. Our optimizer only requires the ergonomic score and its gradient; we did not pursue Hessians. In our work, we are interested in contact whereby the supporting surface touches the human in a chosen region of interest of the human body. Tangential sliding is not desirable and we prevented it by locking the two tangential translational

degrees of freedom of the human. This also had an important pragmatic benefit of not needing a frictional model. Friction forces are notoriously difficult to differentiate. Our optimization spends most of the time in computing the static equilibrium against the supporting surface. Exploring contact models, including friction, that permit rapid optimization and are readily differentiable, yet they accelerate the optimization, is an important avenue for future work. While we launch our optimization from multiple initial guesses, much like most optimization methods, our optimization scheme is local and does not necessarily find the globally optimal solution. Our optimizer treats all surface vertices of the supporting surface as the optimized degrees of freedom; lower-dimensional deformation models (such as, for example, linearly precise biharmonic weights [61]) could be used instead. In the future, more accurate internal human models could be pursued, such as those involving anatomical shapes of muscles and subcutaneous fat. Because the human body is not equally sensitive to pressure in all parts of the body, our ergonomic function could be equipped with a “pain map” [34]. As a common choice in ergonomics literature, our ergonomic function incorporates contact pressures; but ergonomics is also related to posture comfort, such as muscle and joint torques. Furthermore, while most surveyed ergonomic paper advocate for contact pressures as the ergonomic metric, an alternative discomfort metric to explore is three-dimensional elastic strain in the human body.

APPENDIX A CONTACT GRADIENTS AND HESSIAN

Equation 10 defines the energy of collision and the collision force can be computed through its gradient. To use an implicit integrator, the gradient of the force (the hessian of the energy) is also needed. Suppose the three vertices of the triangle that contains the closest point to \mathbf{p}_a on the supporting surface are \mathbf{p}_1 , \mathbf{p}_2 and \mathbf{p}_3 . Then, the position of $\mathbf{p}_{\text{anchor}}$ is a function of \mathbf{p}_a , \mathbf{p}_1 , \mathbf{p}_2 and \mathbf{p}_3 . Observe that \mathbb{R}^3 can be decomposed into 7 disjoint Voronoi regions, one for each vertex, edge and the face interior. The query point \mathbf{p}_a is in one of the regions. As we perturb \mathbf{p}_a , we will remain in the same Voronoi region (unless on the boundary). This means that the closest point will remain the same vertex, slide on the edge, or slide on the face, depending on the specific case; and hence the gradient and hessian formulas must be derived separately for each case. Although the gradient is not continuous across Voronoi region boundaries, such a construction is vastly superior to not making the gradient Voronoi-region aware; it made the difference of the optimizer converging slowly to not at all, to consistently converging.

A.1 Vertex anchor

When $\mathbf{p}_{\text{anchor}}$ is a vertex \mathbf{p}_i , we have

$$\mathbf{p}_{\text{anchor}} = \mathbf{p}_i, \tag{24}$$

$$\frac{\partial E_{\text{coll}}}{\partial \mathbf{p}_a} = 2k_{\text{coll}}(\mathbf{p}_a - \mathbf{p}_i)^T, \tag{25}$$

$$\frac{\partial^2 E_{\text{coll}}}{\partial \mathbf{p}_a^2} = 2k_{\text{coll}} \cdot \mathbf{I}. \tag{26}$$

A.2 Edge anchor

When $\mathbf{p}_{\text{anchor}}$ is an interior point of an edge with endpoints \mathbf{p}_i and \mathbf{p}_j , we have

$$\mathbf{p}_{\text{anchor}} = \frac{(\mathbf{p}_a - \mathbf{p}_i) \cdot (\mathbf{p}_i - \mathbf{p}_j)\mathbf{p}_i + (\mathbf{p}_a - \mathbf{p}_j) \cdot (\mathbf{p}_j - \mathbf{p}_i)\mathbf{p}_j}{\|\mathbf{p}_i - \mathbf{p}_j\|_2}, \quad (27)$$

$$\frac{\partial E_{\text{coll}}}{\partial \mathbf{p}_a} = 2k_{\text{coll}}(\mathbf{p}_a - \mathbf{p}_{\text{anchor}})^T \left(I - \frac{(\mathbf{p}_i - \mathbf{p}_j)(\mathbf{p}_i - \mathbf{p}_j)^T}{\|\mathbf{p}_i - \mathbf{p}_j\|_2^2} \right), \quad (28)$$

$$\frac{\partial^2 E_{\text{coll}}}{\partial \mathbf{p}_a^2} = 2k_{\text{coll}} \cdot \left(I - \frac{(\mathbf{p}_i - \mathbf{p}_j)(\mathbf{p}_i - \mathbf{p}_j)^T}{\|\mathbf{p}_i - \mathbf{p}_j\|_2^2} \right)^2. \quad (29)$$

A.3 Face anchor

When $\mathbf{p}_{\text{anchor}}$ is an interior point of the face, we have

$$\mathbf{p}_{\text{anchor}} = \mathbf{p}_a - ((\mathbf{p}_a - \mathbf{p}_1) \cdot \mathbf{n})\mathbf{n}, \quad (30)$$

$$\frac{\partial E_{\text{coll}}}{\partial \mathbf{p}_a} = 2k_{\text{coll}}(\mathbf{p}_a - \mathbf{p}_{\text{anchor}})^T \cdot \mathbf{n}\mathbf{n}^T, \quad (31)$$

$$\frac{\partial^2 E_{\text{coll}}}{\partial \mathbf{p}_a^2} = 2k_{\text{coll}} \cdot \mathbf{n}\mathbf{n}^T, \quad (32)$$

where \mathbf{n} is the unit normal of the triangle.

ACKNOWLEDGMENTS

This research was sponsored in part by Adobe Research, NSF (IIS-1911224), a USC Annenberg Fellowship to Danyong Zhao and Yijing Li, and Bosch Research. Part of the work was carried out while Li was an intern at Adobe. We also gratefully acknowledge Sourav Kumar Bose and Uday Kusupati, who were student leads on early research in this direction at IIT Bombay.


REFERENCES

- [1] V. Putz-Anderson, B. P. Bernard, S. E. Burt, L. L. Cole, C. Fairfield-Estill, L. J. Fine, K. A. Grant, C. Gjessing, L. Jenkins, J. J. Hurrell Jr et al., "Musculoskeletal disorders and workplace factors," *National Institute for Occupational Safety and Health (NIOSH)*, vol. 104, 1997.
- [2] P. Baudisch and S. Mueller, "Personal fabrication," *Foundations and Trends in Human-Computer Interaction*, vol. 10, no. 3-4, pp. 165–293, 2017.
- [3] S. Chaudhuri, D. Ritchie, J. Wu, K. Xu, and H. Zhang, "Learning generative models of 3D structures," *Computer Graphics Forum*, vol. 39, no. 2, pp. 643–666, 2020.
- [4] H. Grabner, J. Gall, and L. Van Gool, "What makes a chair a chair?" in *Proc. CVPR*, 06 2011, pp. 1529–1536.
- [5] V. G. Kim, S. Chaudhuri, L. Guibas, and T. Funkhouser, "Shape2pose: Human-centric shape analysis," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2014)*, vol. 33, no. 4, 2014.
- [6] Y. Jiang, H. Koppula, and A. Saxena, "Hallucinated humans as the hidden context for labeling 3D scenes," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2993–3000.
- [7] R. Hu, Z. Yan, J. Zhang, O. Van Kaick, A. Shamir, H. Zhang, and H. Huang, "Predictive and generative neural networks for object functionality," *ACM Trans. on Graphics (Proc. of SIGGRAPH 2018)*, vol. 37, no. 4, 2018.
- [8] K. Leimer, M. Birsak, F. Rist, and P. Musialski, "Sit & relax: Interactive design of body-supporting surfaces," *Computer Graphics Forum*, vol. 37, no. 7, pp. 349–359, 2018.
- [9] K. Leimer, A. Winkler, S. Ohrhallinger, and P. Musialski, "Pose to seat: Automated design of body-supporting surfaces," *Computer-Aided Geometric Design*, vol. 79, pp. 79:1–1, 2020.
- [10] D. James and D. Pai, "Modeling deformation of linear elastostatic objects," *Haptic Rendering: Foundations, Algorithms, and Applications*, pp. 395–420, 2008.
- [11] M. Bendsoe and O. Sigmund, *Topology Optimization*. Springer-Verlag Berlin Heidelberg, 2004.
- [12] H. Liu, Y. Hu, B. Zhu, W. Matusik, and E. Sifakis, "Narrow-band topology optimization on a sparsely populated grid," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH Asia 2018)*, vol. 37, no. 6, pp. 251:1–251:14, 2018.
- [13] T. Langlois, A. Shamir, D. Dror, W. Matusik, and D. I. W. Levin, "Stochastic structural analysis for context-aware design and fabrication," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH Asia 2016)*, vol. 35, no. 6, pp. 226:1–226:13, 2016.
- [14] E. Ulu, J. Mccann, and L. B. Kara, "Lightweight structure design under force location uncertainty," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2017)*, vol. 36, no. 4, 2017.
- [15] N. P. van Dijk, K. Maute, M. Langelaar, and F. Van Keulen, "Level-set methods for structural topology optimization: a review," *Structural and Multidisciplinary Optimization*, vol. 48, no. 3, pp. 437–472, 2013.
- [16] M. Shugrina, A. Shamir, and W. Matusik, "Fab forms: Customizable objects for fabrication with validity and geometry caching," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2015)*, vol. 34, no. 4, 2015.
- [17] M. E. Yumer, S. Chaudhuri, J. K. Hodgins, and L. B. Kara, "Semantic shape editing using deformation handles," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2015)*, vol. 34, no. 4, pp. 86:1–86:12, 2015.
- [18] N. Umetani, T. Igarashi, and N. J. Mitra, "Guided exploration of physically valid shapes for furniture design," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2012)*, vol. 31, no. 4, 2012.
- [19] Z. Wang, P. Song, F. Isvoranu, and M. Pauly, "Design and structural optimization of topological interlocking assemblies," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH Asia 2019)*, vol. 38, no. 6, 2019.
- [20] M. Bäcker, E. Whiting, B. Bickel, and O. Sorkine-Hornung, "Spin-It: Optimizing moment of inertia for spinnable objects," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2014)*, vol. 33, no. 4, pp. 96:1–96:10, 2014.
- [21] R. Prévost, E. Whiting, S. Lefebvre, and O. Sorkine-Hornung, "Make It Stand: Balancing shapes for 3D fabrication," *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH 2013)*, vol. 32, no. 4, pp. 81:1–81:10, 2013.
- [22] N. Umetani, Y. Koyama, R. Schmidt, and T. Igarashi, "Pteromys: Interactive design and optimization of free-formed free-flight model airplanes," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2014)*, vol. 33, no. 4, 2014.
- [23] L. Wang and E. Whiting, "Buoyancy optimization for computational fabrication," *Computer Graphics Forum (Proc. of Eurographics)*, vol. 35, no. 2, 2016.
- [24] D. Chen, D. I. Levin, W. Matusik, and D. M. Kaufman, "Dynamics-aware numerical coarsening for fabrication design," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2017)*, vol. 34, no. 4, 2017.
- [25] E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner, "Data-driven estimation of cloth simulation models," *Computer Graphics Forum (Proc. of Eurographics)*, vol. 31, no. 2, 2012.
- [26] J. Pérez, B. Thomaszewski, S. Coros, B. Bickel, J. A. Canabal, R. Sumner, and M. A. Otaduy, "Design and fabrication of flexible rod meshes," *ACM Trans. Graph.*, vol. 34, no. 4, 2015.
- [27] J. Montes, B. Thomaszewski, S. Mudur, and T. Popa, "Computational design of skintight clothing," *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2020)*, vol. 39, no. 4, 2020.
- [28] M. Bucki, V. Luboz, A. Perrier, E. Champion, B. Diot, N. Vuillerme, and Y. Payan, "Clinical workflow for personalized foot pressure ulcer prevention," *Medical engineering & physics*, vol. 38, no. 9, pp. 845–853, 2016.
- [29] V. Luboz, M. Bailet, C. B. Grivot, M. Rochette, B. Diot, M. Bucki, and Y. Payan, "Personalized modeling for real-time pressure ulcer prevention in sitting posture," *Journal of tissue viability*, vol. 27, no. 1, pp. 54–58, 2018.
- [30] U. G. S. Administration, "Ergonomic seating adjustment guide," https://www.gsa.gov/cdnstatic/GSA_Ergonomic_Seat_Adjustment_Guide.pdf, 2020.
- [31] G. Kyung and M. A. Nussbaum, "Specifying comfortable driving postures for ergonomic design and evaluation of the driver workspace using digital human models," *Ergonomics*, vol. 52, no. 8, pp. 939–953, 2009.

- [32] —, “Driver sitting comfort and discomfort (part ii): Relationships with and prediction from interface pressure,” *International Journal of Industrial Ergonomics*, vol. 38, no. 5, pp. 526–538, 2008.
- [33] M. P. De Looze, L. F. Kuijt-Evers, and J. Van Dieen, “Sitting comfort and discomfort and the relationships with objective measures,” *Ergonomics*, vol. 46, no. 10, pp. 985–997, 2003.
- [34] L. Savonnet, X. Wang, and S. Duprey, “Finite element models of the thigh-buttock complex for assessing static sitting discomfort and pressure sore risk: a literature review,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 21, no. 4, pp. 379–388, 2018.
- [35] D. K. Pai, A. Rothwell, P. Wyder-Hodge, A. Wick, Y. Fan, E. Larionov, D. Harrison, D. R. Neog, and C. Shing, “The human touch: measuring contact with real human soft tissues,” *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2018)*, vol. 37, no. 4, p. 58, 2018.
- [36] J. E. Chadwick, D. R. Haumann, and R. E. Parent, “Layered construction for deformable animated characters,” in *Proc. of ACM SIGGRAPH 1989*, 1989, pp. 243–252.
- [37] M. Pauly, D. K. Pai, and L. J. Guibas, “Quasi-rigid objects in contact,” in *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 2004, pp. 109–119.
- [38] D. Terzopoulos and A. Witkin, “Physically based models with rigid and deformable components,” *IEEE Computer Graphics and Applications*, vol. 8, no. 6, pp. 41–51, 1988.
- [39] D. Metaxas and D. Terzopoulos, “Dynamic deformation of solid primitives with constraints,” in *Proceedings of ACM SIGGRAPH 1992*, 1992, pp. 309–312.
- [40] M.-P. Cani-Gascuel, “Layered deformable models with implicit surfaces,” in *Proceedings of Graphics Interface Conf.*, 1998, pp. 201–208.
- [41] N. Galoppo, M. A. Otaduy, P. Mecklenburg, M. Gross, and M. C. Lin, “Fast simulation of deformable models in contact using dynamic deformation textures,” in *Symp. on Computer Animation (SCA)*, 2006, pp. 73–82.
- [42] N. Galoppo, M. A. Otaduy, S. Tekin, M. Gross, and M. C. Lin, “Soft articulated characters with fast contact handling,” *Computer Graphics Forum*, vol. 26, no. 3, pp. 243–253, 2007.
- [43] M. Kim, G. Pons-Moll, S. Pujades, S. Bang, J. Kim, M. J. Black, and S.-H. Lee, “Data-driven physics for human soft tissue animation,” *ACM Trans. Graph.*, vol. 36, no. 4, 2017.
- [44] C. Romero, M. A. Otaduy, D. Casas, and J. Perez, “Modeling and estimation of nonlinear skin mechanics for animated avatars,” *Computer Graphics Forum (Proc. of Eurographics)*, vol. 39, no. 2, 2020.
- [45] G. Saul, M. Lau, J. Mitani, and T. Igarashi, “Sketchchair: An all-in-one chair design system for end users,” in *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction*, ser. TEI ’11, 2011, pp. 73–80.
- [46] Y. Zheng, H. Liu, J. Dorsey, and N. J. Mitra, “Ergonomics-inspired reshaping and exploration of collections of models,” *IEEE Trans. on Visualization and Computer Graphics*, vol. 22, no. 6, pp. 1732–1744, 2016.
- [47] Artec3D, “Eva Scanner, <http://www.artec3d.com>,” 2018.
- [48] TurboSquid, “TurboSquid,” 2020, www.turbosquid.com.
- [49] T. Brochu and R. Bridson, “Robust topological operations for dynamic explicit surfaces,” *SIAM Journal on Scientific Computing*, vol. 31, no. 4, pp. 2472–2493, 2009.
- [50] H. Si, “TetGen: A Quality Tetrahedral Mesh Generator and a 3D Delaunay Triangulator,” 2011.
- [51] G. Taubin, “A signal processing approach to fair surface design,” in *Proc. of ACM SIGGRAPH 1995*, pp. 351–358.
- [52] G. Irving, J. Teran, and R. Fedkiw, “Invertible Finite Elements for Robust Simulation of Large Deformation,” in *Symp. on Computer Animation (SCA)*, 2004, pp. 131–140.
- [53] R. Kikuuwe, H. Tabuchi, and M. Yamamoto, “An edge-based computationally efficient formulation of saint venant-kirchhoff tetrahedral finite elements,” *ACM Trans. on Graphics*, vol. 28, no. 1, pp. 1–13, 2009.
- [54] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical recipes: The art of scientific computing*, 3rd ed. Cambridge University Press, Cambridge, UK, 2007.
- [55] J. Barbič and D. L. James, “Six-dof haptic rendering of contact between geometrically complex reduced deformable models,” *IEEE Transactions on Haptics*, vol. 1, no. 1, pp. 39–52, 2008.
- [56] J. Liang, M. Lin, and V. Koltun, “Differentiable cloth simulation for inverse problems,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32, 2019, pp. 772–781.
- [57] FujiFilm, “Fujifilm Prescale Pressure Measurement Film, <https://fujiprescalefilm.com>,” 2020.
- [58] Tekscan, “<http://www.tekscan.com>,” 2020.
- [59] AljaSafe, “SmoothOn Inc.” 2020, www.smooth-on.com.
- [60] Artec3D, “Spider Scanner, <http://www.artec3d.com>,” 2018.
- [61] Y. Wang, A. Jacobson, J. Barbič, and L. Kavan, “Linear subspace design for real-time shape deformation,” *ACM Trans. on Graphics (TOG)*, vol. 34, no. 4, p. 57, 2015.



Danyong Zhao is a PhD student in computer science at the University of Southern California. He obtained his BS degree from Tsinghua University. His research interests are in computer graphics, physically based animation and contact.



Yijing Li holds a PhD in Computer Science, Viterbi School of Engineering, University of Southern California. His research interests include physically based modeling, geometric shape modeling and computer animation.



Siddhartha Chaudhuri is a Senior Research Scientist in the Creative Intelligence Lab at Adobe Research, and an Assistant Professor (on leave) of computer science at IIT Bombay. His work combines geometric analysis, machine learning, and UI innovation to make sophisticated 3D modeling accessible even to non-expert users. He also studies related foundational problems in geometry processing. In addition to many publications at leading venues, he is the original author of the commercial 3D

character modeling tool Adobe Fuse, and the author of the open-source Thea toolkit for geometric computing.



Timothy Langlois is a Senior Research Scientist in the Creative Intelligence Lab at Adobe Research. His main research interest is leveraging physical simulation to assist creativity. Tim has published in multiple fields including physical acoustics, stochastic structural analysis, immersive audio for AR/VR, and robust elastodynamic simulation.



Jernej Barbič is an associate professor of computer science at USC. In 2011, MIT Technology Review named him one of the Top 35 Innovators under the age of 35 in the world (TR35). Jernej published over 40 papers on nonlinear solid deformation modeling, collision detection and contact, and interactive design of deformations and animations. He is the author of Vega FEM, an efficient free C/C++ software physics library for deformable object simulation. Jernej is a Sloan Fellow (2014).