

# Implicit Neural Representation for Physics-driven Actuated Soft Bodies

LINGCHEN YANG, ETH Zurich, Switzerland  
BYUNGSOO KIM, ETH Zurich, Switzerland  
GASPARD ZOSS, DisneyResearch|Studios, Switzerland  
BARAN GÖZCÜ, ETH Zurich, Switzerland  
MARKUS GROSS, ETH Zurich, Switzerland  
BARBARA SOLENTHALER, ETH Zurich, Switzerland




Fig. 1. We present a method to control active soft bodies using an implicit neural representation. A continuous mapping from a spatial point  $\mathbf{x}$  in the material space to the control parameters  $\Omega_{\mathbf{x}}$  is established, rendering the method discretization agnostic and applicable to various soft body types (left). For faces, we consider both actuation and jaw kinematics for articulating high-fidelity expressions (right).

Active soft bodies can affect their shape through an internal actuation mechanism that induces a deformation. Similar to recent work, this paper utilizes a differentiable, quasi-static, and physics-based simulation layer to optimize for actuation signals parameterized by neural networks. Our key contribution is a general and implicit formulation to control active soft bodies by defining a function that enables a continuous mapping from a spatial point in the material space to the actuation value. This property allows us to capture the signal's dominant frequencies, making the method discretization agnostic and widely applicable. We extend our implicit model to mandible kinematics for the particular case of facial animation and show that we can reliably reproduce facial expressions captured with high-quality capture systems. We apply the method to volumetric soft bodies, human poses, and facial expressions, demonstrating artist-friendly properties, such as simple control over the latent space and resolution invariance at test time.

Authors' addresses: Lingchen Yang, ETH Zurich, Switzerland, lingchen.yang@inf.ethz.ch; Byungsoo Kim, ETH Zurich, Switzerland, kimby@inf.ethz.ch; Gaspard Zoss, DisneyResearch|Studios, Switzerland, gaspard.zoss@disneyresearch.com; Baran Gözcü, ETH Zurich, Switzerland, baran.goezcue@inf.ethz.ch; Markus Gross, ETH Zurich, Switzerland, grossm@inf.ethz.ch; Barbara Solenthaler, ETH Zurich, Switzerland, solenthaler@inf.ethz.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
0730-0301/2022/7-ART122 \$15.00  
<https://doi.org/10.1145/3528223.3530156>

CCS Concepts: • **Computing methodologies** → **Physical Simulation**; **Neural networks**.

Additional Key Words and Phrases: Differentiable Physics, Deep Learning, Digital Human

## ACM Reference Format:

Lingchen Yang, Byungsoo Kim, Gaspard Zoss, Baran Gözcü, Markus Gross, and Barbara Solenthaler. 2022. Implicit Neural Representation for Physics-driven Actuated Soft Bodies. *ACM Trans. Graph.* 41, 4, Article 122 (July 2022), 13 pages. <https://doi.org/10.1145/3528223.3530156>

## 1 INTRODUCTION

Active soft bodies can deform their shape through an internal actuation mechanism and play a key role in physics-based animation. The actuation mechanism induces an internal force that drives the motion of the discrete mesh elements of a body, and can be driven by artists for shape targeting problems [Klár et al. 2020], induced by the underlying musculature in the case of human bodies and faces [Sifakis et al. 2005], or controlled by embedded sensors in robotics applications [Bächer et al. 2021].

In order to guide a simulation towards a target shape, an inverse problem arises when the optimal actuation signals, which deform the simulation mesh to match a given target surface, must be found for each element. This problem can be solved with a differentiable simulator, either through a neural network approach [Fulton et al. 2019] or by directly differentiating the governing equations and thus inducing a strong physics prior [Du et al. 2022]. For the particular application of faces, Srinivasan et al. [2021] postulates that the

control parameters of the soft tissue actuation mechanism are low-dimensional and can be learned by a neural network that is coupled with a differentiable quasi-static physics solver. Their actuation-generative network is tied to a specific discretization, which requires the re-definition of the network each time the discretization of the input shape changes. It is thus not generally applicable to soft bodies of arbitrary shapes and resolution inputs.

In this work, we present a general implicit physics- and data-driven formulation to control active soft bodies. We build upon the critical assumption that we can learn a continuous function which maps the spatial point in the material space (undeformed space) to the actuation signal. We hypothesize that the trained model can effectively capture the dominant frequencies of the actuation signal and reproduce them at test time. Our approach offers the core advantage that the method is agnostic to the underlying shape representation. In other words, with our method, there is no need to manually re-define the network architecture or retrain the model if the underlying representation or resolution changes. These fundamental properties render the method generally applicable to arbitrary soft body inputs and reduce the required expert knowledge, allowing artists to generate new poses and animations efficiently. We underline our claims by showing results for volumetric soft bodies, human body motion, and facial expressions. For the particular case of faces, the motion of the soft tissue is highly dominated by the movement of an internal bone structure that defines Dirichlet boundary conditions in the simulator. Therefore, our implicit model optimizes not only actuation signals but also mandible kinematics.

The main contributions of our work are summarized as follows:

- An implicit neural method for computing actuation signals of active soft bodies, enabling applicability to arbitrary shapes.
- Extension of the implicit method to facial animations, where mandible kinematics are optimized and coupled with the physics solver via Dirichlet boundary conditions.
- Conditioning of the actuation network on a continuous resolution input to enable resolution invariance at test time.
- Improved performance due to a closed-form Hessian of the energy density function in the implicit model and a weight matrix modulation in the network design.
- Demonstration of the effectiveness and generality of our method on active objects of different types and resolutions.

## 2 RELATED WORK

We briefly review prior work on soft body and face simulation, focusing on inverse problems, data-driven methods, and implicit neural representations.

*Differentiable simulation.* Differentiable physics solvers enable the use of gradient-based methods for solving complex control problems. ChainQueen [Hu et al. 2019b] and DiffTaichi [Hu et al. 2019a] apply differentiability to the material point method. Hahn et al. [2019] and Geilinger et al. [2020] use differentiable implicit FEM simulators that systematically apply a combination of the adjoint method and sensitivity analysis to derive gradients. Qiao et al. [2021] introduces a differentiable framework to dynamically simulate a soft body with rigid skeleton actuation. Similar to them, we consider skeleton kinematics but restricted to the quasi-static state. Hence,

we adopt the quasi-static skeleton position for the target shape as the Dirichlet boundary condition in the simulation. Differentiable Projective Dynamics (DiffPD) with implicit time stepping is used in [Du et al. 2022]. A focus is put on optimizing the backpropagation to achieve fast gradient computations that outperform the standard Newton’s method. In our work, we adopt DiffPD as the backward solver for the optimization. Neural physics solvers are naturally differentiable. An autoencoder and latent-space dynamics are used in [Fulton et al. 2019] to simulate deformable bodies, and a graph neural network is applied to particle-based soft body dynamics in [Sanchez-Gonzalez et al. 2020].

*Inverse Modeling.* Differentiable physics solvers are used to find minimizers for inverse problems in robotics control tasks [Degraeve et al. 2019; Macklin et al. 2020], cloth simulation [Liang et al. 2019; Qiao et al. 2020], and fluid dynamics [Holl et al. 2020; Tang et al. 2021]. In the context of digital humans, inverse problems are solved to identify mechanical properties of the face, in particular heterogeneous stiffness and prestrain, using 3D scans and imaging data [Kadleček and Kavan 2019]. A differentiable physics-based muscle model is applied to 3D facial poses and 2D images, alike [Bao et al. 2018]. Instead of explicitly modeling muscles, a simulation mesh is subdivided into tetrahedral elements where each can be activated [Ichim et al. 2017; Klár et al. 2020]. Ichim et al. [2017] additionally includes jaw kinematics, similar to our work. However, different from them, we explicitly build the differential connection between the network and the kinematics. Srinivasan et al. [2021] uses a neural network in combination with a differentiable soft body solver to learn an actuation mechanism such that a facial mesh is deformed into a desired expression after physics-based simulation. The network depends on the input representation, and hence needs to be manually changed if the input changes. We overcome this limitation in our work by using an implicit representation that enables a continuous mapping from a spatial point to a control value.

*Blendshapes and data-driven models.* Blendshapes are widely used for facial and human body animation [Lewis et al. 2014]. They approximate the space of valid poses, and new poses can be computed as a linear, weighted combination of the basis shapes. These early methods have been improved by data-driven models, which leverage high-quality surface scan datasets [Egger et al. 2020; Klehm et al. 2015]. More recently, neural networks have been used to disentangle the facial identity and expression, and jointly model the geometry and appearance [Chandran et al. 2020; Li et al. 2020]. Deep learning methods have also been applied to human bodies, either to model soft tissue deformation as a function of body motion [Casas and Otaduy 2018], in combination with body shapes [Loper et al. 2015; Pons-Moll et al. 2015], or by using nonlinear subspaces to encode tissue deformation [Santesteban et al. 2020].

*Coordinate-based MLPs.* Coordinate-based multilayer perceptrons (MLPs) [Tancik et al. 2020] have recently been very successful at generating high-fidelity representations of different visual problems. A series of works have adopted coordinate-based MLPs for 3D scene reconstruction [Mildenhall et al. 2020] or 3D objects [Mescheder et al. 2019; Park et al. 2019]. Mildenhall et al. [2020] demonstrates the effectiveness of architectures using a positional encoding scheme

for representing high-frequency signals and presents photo-realistic novel view synthesis. SIREN [Sitzmann et al. 2020] adopts periodic activation functions for implicit neural representations of intricate objects and their derivatives. Our network design is motivated by SIREN. However, SIREN’s limited performance on high-dimensional inputs makes their design not directly applicable to our problem setting. A couple of works [Chan et al. 2020; Or-El et al. 2021; Zhou et al. 2021] extend SIREN to condition on the input by using adaptive instance normalization (AdaIN) [Huang and Belongie 2017] in the feature space. In contrast to AdaIN, we condition the weight matrix on the latent code, which reduces memory and enables an efficient backpropagation.

### 3 INVERSE CONTROL

Active objects can affect their shape through an internal actuation mechanism. We take an example pose of an object as input, and solve the following inverse control problem: For each element of a soft body mesh we want to find an activation value, such that when used in a forward quasi-static physics solver, the rest pose of the body is deformed into the target shape. To solve this, we follow Srinivasan et al. [2021] and learn a low-dimensional control space, where a latent code  $\mathbf{z}$  is sampled to get the fine-grained actuation mechanism corresponding to a shape. We train the model via differentiable physics, which allows us to integrate the numerical solver into the training of neural networks where gradient information flows from the quasi-static simulator into the network. This general idea can be extended to complex bodies as well, such as in facial soft tissues, where deformations are heavily dominated by the underlying bone movement. An overview of our complete method is shown in Fig. 2.

We can generally define an inverse problem as follows. We want to learn simulation parameters  $\Omega$  from a set of observations  $\mathbf{s}$  such that after optimization, a simulator using  $\Omega$  can reconstruct  $\mathbf{s}$ . To achieve this, we first need to define the energy function  $E(\mathbf{u}, \Omega)$ , which describes the relation between the simulation mesh vertices  $\mathbf{u}$  and the parameters of interest  $\Omega$ . The quasi-static simulation requires us to find the minimizer of  $E$  over  $\mathbf{u}$  given  $\Omega$ , thereby finding the optimal positions  $\mathbf{u}^*$  that make the net force become zero everywhere:

$$\nabla_{\mathbf{u}} E(\mathbf{u}^*, \Omega) = \mathbf{0}. \quad (1)$$

With this implicit constraint between  $\mathbf{u}^*$  and  $\Omega$ , we can define the inverse problem as

$$\Omega^* = \arg \min_{\Omega} \mathcal{L}(\mathbf{u}^*(\Omega), \mathbf{s}), \quad (2)$$

where the loss function  $\mathcal{L}$  measures the difference between the simulation and the observation. The sensitivity matrix  $\partial \mathbf{u}^* / \partial \Omega$ , which plays a key role in the optimization, is derived by differentiating on both sides of Eqn. (1) with respect to  $\Omega$ :

$$\left. \frac{\partial \nabla_{\mathbf{u}} E}{\partial \mathbf{u}} \right|_{(\mathbf{u}^*, \Omega)} \frac{\partial \mathbf{u}^*}{\partial \Omega} + \left. \frac{\partial \nabla_{\mathbf{u}} E}{\partial \Omega} \right|_{(\mathbf{u}^*, \Omega)} = \mathbf{0}. \quad (3)$$

By omitting the explicit evaluation at  $(\mathbf{u}^*, \Omega)$ , we get a concise form:

$$\frac{\partial \mathbf{u}^*}{\partial \Omega} = - \underbrace{\left( \frac{\partial \nabla_{\mathbf{u}} E}{\partial \mathbf{u}} \right)^{-1}}_{\mathbf{H}_{\mathbf{u}}} \underbrace{\frac{\partial \nabla_{\mathbf{u}} E}{\partial \Omega}}_{\mathbf{H}_{\Omega}}. \quad (4)$$

This matrix is used to propagate the gradient accumulated on  $\mathbf{u}^*$  to  $\Omega$  by the chain rule when optimizing  $\mathcal{L}$ :

$$\frac{\partial \mathcal{L}}{\partial \Omega} = \frac{\partial \mathcal{L}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}^*}{\partial \Omega} = - \frac{\partial \mathcal{L}}{\partial \mathbf{u}} \mathbf{H}_{\mathbf{u}}^{-1} \mathbf{H}_{\Omega}. \quad (5)$$

### 4 IMPLICIT FORMULATION

Generally, we can implicitly represent any continuous mechanical property using a neural network  $\mathcal{N}$  that takes as input a spatial point  $\mathbf{x}$  in the material space (undeformed space) and outputs the corresponding property value, e.g.,  $\Omega_{\mathbf{x}} = \mathcal{N}(\mathbf{x})$ , as illustrated in Fig. 1. Here, we derive an implicit formulation of an internal actuation mechanism. To animate an active object, we control this mechanism via a network  $\mathcal{N}_{\mathbf{A}}$ , trained on target poses.

We adopt the energy density function  $\Psi$  from shape targeting [Klár et al. 2020] for the internal actuation model, given as

$$\Psi(\mathbf{F}, \mathbf{A}) = \min_{\mathbf{R} \in \text{SO}(3)} \|\mathbf{F} - \mathbf{R}\mathbf{A}\|_F^2, \quad (6)$$

where  $\mathbf{F}$  is the local deformation gradient and  $\mathbf{A}$  is a symmetric  $3 \times 3$  actuation matrix, whose eigenvectors indicate the muscle contraction directions and eigenvalues the magnitudes of the local deformation. The rotation matrix  $\mathbf{R}$  factors away the rotational difference between  $\mathbf{F}$  and  $\mathbf{A}$  and makes  $\Psi$  rotationally-invariant. We assign an actuation matrix to every spatial point  $\mathbf{x}$  in the object’s material space, i.e.,  $\mathbf{A}(\mathbf{x}) = \mathcal{N}_{\mathbf{A}}(\mathbf{x})$ .

For the following equations we first define  $\text{vec}(\cdot)$  as row-wise flattening of a matrix into a vector, and the expanded matrix  $\hat{\mathbf{A}}$  as:

$$\hat{\mathbf{A}} = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A} \end{bmatrix}. \quad (7)$$

Using  $\check{\mathbf{f}} = \text{vec}(\mathbf{F})$ ,  $\check{\mathbf{r}} = \text{vec}(\mathbf{R})$ , and  $\hat{\mathbf{A}}\check{\mathbf{r}} = \text{vec}(\mathbf{R}\mathbf{A})$ , the continuous energy function  $\tilde{E}$  in the simulation is defined as

$$\tilde{E} = \int_{\mathcal{D}^0} \frac{1}{2} \|\check{\mathbf{f}}(\mathbf{x}) - \hat{\mathbf{A}}(\mathbf{x})\check{\mathbf{r}}^*(\mathbf{x})\|_2^2 dV, \quad (8)$$

where  $\mathcal{D}^0$  denotes the material space of the object, and  $\check{\mathbf{r}}^*$  is the vectorized rotation matrix for point  $\mathbf{x}$ , precomputed from the polar decomposition of  $\mathbf{F}\mathbf{A}$  [Klár et al. 2020]. Following the standard practices of the Finite Element method, we discretize  $\mathcal{D}^0$  using tiny elements connected by nodal vertices  $\mathbf{u}$ :

$$\tilde{E} \approx \sum_e \int_{\mathcal{D}_e^0} \frac{1}{2} \|\check{\mathbf{f}}(\mathbf{x}) - \hat{\mathbf{A}}(\mathbf{x})\check{\mathbf{r}}^*(\mathbf{x})\|_2^2 dV \quad (9)$$

$$\approx \sum_e \frac{V_e}{N} \sum_i \frac{1}{2} \|\check{\mathbf{f}}(\mathbf{x}_{e,i}) - \hat{\mathbf{A}}(\mathbf{x}_{e,i})\check{\mathbf{r}}^*(\mathbf{x}_{e,i})\|_2^2 \quad (10)$$

$$E(\mathbf{u}, \mathcal{A}) = \sum_e \frac{V_e}{N} \sum_i \frac{1}{2} \underbrace{\|\mathbf{G}(\mathbf{x}_{e,i})\mathbf{u}_e - \hat{\mathbf{A}}(\mathbf{x}_{e,i})\check{\mathbf{r}}^*(\mathbf{x}_{e,i})\|_2^2}_{\Psi_{e,i}}, \quad (11)$$

where  $\mathcal{A}$  denotes all the sampled actuation matrices  $\mathbf{A}$  from the network  $\mathcal{N}_{\mathbf{A}}$ , i.e.,  $\mathbf{A}(\mathbf{x}) = \mathcal{N}_{\mathbf{A}}(\mathbf{x})$ ,  $e$  denotes an element, and  $\mathcal{D}_e^0$  indicates the continuous region inside  $e$  while  $V_e$  is its volume. We sample  $N$  points inside each  $\mathcal{D}_e^0$  to approximate the integral (Eqn. (10)). The deformation gradient  $\check{\mathbf{f}}$  at each point can be approximated by the concatenated nodal vertices  $\mathbf{u}_e$ , associated with the

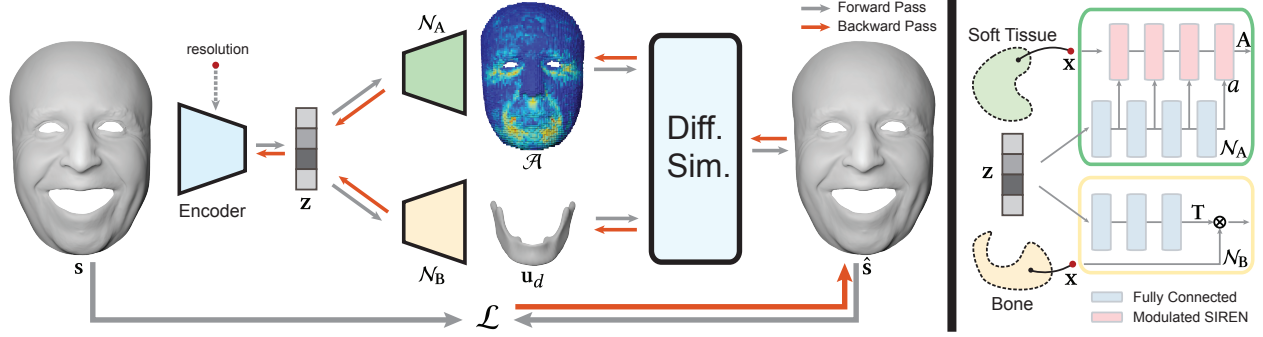


Fig. 2. Overview of our method. Using a set of observations  $s$  (target poses) we learn actuation signals  $\mathcal{A}$  and mandible kinematics  $\mathbf{u}_d$  (for faces only), such that when using these parameters in a forward pass the simulation output  $\hat{s}$  matches ground truth. We implicitly represent the two mechanical properties using the networks  $\mathcal{N}_A$  and  $\mathcal{N}_B$ , and couple them with a differentiable quasi-static soft body simulator to allow gradient information to flow from the solver to the networks. The encoder is a global shape descriptor and outputs a latent code  $z$ .

element  $e$  where  $\mathbf{x}_{e,i}$  resides, via the corresponding mapping matrix  $\mathbf{G}$  (Eqn. (11)). We adopt hexahedral elements and use projective dynamics [Bouaziz et al. 2014] for minimizing  $E$ .

$\Psi_{e,i}$  in Eqn. (11) is the key for deriving  $\mathbf{H}_u$  in Eqn. (4), since  $\mathbf{H}_u$  is the accumulation of all Hessians  $\nabla^2 \Psi_{e,i}$ . We omit  $(x, e, i, *)$  except  $\mathbf{u}_e$  for simplicity. As a result of projective dynamics,  $\nabla \Psi = \mathbf{G}^\top (\mathbf{G} \mathbf{u}_e - \hat{\mathbf{A}} \mathbf{f})$ . Taking the derivative of  $\nabla \Psi$ , we get

$$\frac{\partial \nabla \Psi}{\partial \mathbf{u}_e} = \mathbf{G}^\top \mathbf{G} - \mathbf{G}^\top \hat{\mathbf{A}} \frac{\partial \mathbf{f}}{\partial \mathbf{u}_e} \quad (12)$$

$$= \mathbf{G}^\top \mathbf{G} - \mathbf{G}^\top \hat{\mathbf{A}} \frac{\partial \mathbf{f}}{\partial \mathbf{A}} \hat{\mathbf{A}} \mathbf{G} \quad (13)$$

$$= \mathbf{G}^\top \mathbf{G} - \mathbf{G}^\top \hat{\mathbf{A}} \mathbf{H}_R \hat{\mathbf{A}} \mathbf{G}. \quad (14)$$

Note that  $\mathbf{r}$  comes from the polar decomposition of  $\hat{\mathbf{A}} \mathbf{f}$ ,  $\mathbf{H}_R$  is the rotation gradient, namely the derivative with respect to  $\hat{\mathbf{A}} \mathbf{f}$ . Luckily, the closed form for  $\mathbf{H}_R$  has already been derived in [Kim and Eberle 2020], which can be constructed from its off-the-shelf three eigenvectors  $\mathbf{q}$  and eigenvalues  $\lambda$ , *i.e.*,  $\mathbf{H}_R = \sum_i \lambda_i \mathbf{q}_i \mathbf{q}_i^\top$ . Similarly, we have  $\partial \nabla \Psi / \partial \hat{\mathbf{a}}$  as

$$\frac{\partial \nabla \Psi}{\partial \hat{\mathbf{a}}} = -\mathbf{G}^\top \hat{\mathbf{A}} \mathbf{H}_R \hat{\mathbf{F}} - \mathbf{G}^\top \hat{\mathbf{R}}, \quad (15)$$

which can be used for constructing  $\mathbf{H}_\Omega$  in Eqn. (4) ( $\Omega$  denotes actuation  $\mathcal{A}$  here). The definition for the expanded matrices  $\hat{\mathbf{F}}$  and  $\hat{\mathbf{R}}$ , and more details on the derivation can be found in Supplemental Material. The sensitivity matrix will be used to calculate the gradient accumulated on the sampled actuation matrices, building the connection between the network  $\mathcal{N}_A$  and the simulator. Compared to [Srinivasan et al. 2021] where  $\mathbf{H}_u$  and  $\mathbf{H}_\Omega$  are evaluated with 9 iterations of sensitivity analysis, our closed form requires a single iteration only. In general, the key advantage of our implicit formulation is that it is agnostic to the shape representation, which renders the technique widely applicable.

## 5 NETWORK ARCHITECTURE

For the actuation network  $\mathcal{N}_A$  we adopt SIREN [Sitzmann et al. 2020] due to its excellent representational power. The most effective part of SIREN is its sine activation function  $\mathbf{h} \mapsto \sin(\mathbf{W}\mathbf{h} + \mathbf{b})$ , where

$\mathbf{h}$  is the output hidden vector from a previous layer, and  $\mathbf{W}$  and  $\mathbf{b}$  are trainable parameters in the current layer. Since we found that SIREN does not handle high-dimensional inputs well, *i.e.*, directly concatenating the latent code  $z$  and spatial point  $\mathbf{x}$  as input, we instead condition  $\mathbf{W}$  on  $z$  as

$$\mathbf{W}_{i,j} = a_i \cdot \hat{\mathbf{W}}_{i,j}, \quad (16)$$

where  $a_i$  is the modulating coefficient decoded from  $z$  with a tiny MLP,  $\hat{\mathbf{W}}$  is the shared weight matrix in this layer, and  $i$  and  $j$  enumerate the dimensions of input and output feature vectors, respectively. The motivation for this conditioning mechanism is that a signal (actuation) can be approximated as the linear combination ( $\mathbf{W}$ ) of a set of basis functions (the output  $\mathbf{h}$  of the previous SIREN layer). One advantage of modulating the weight matrix  $\mathbf{W}$  instead of  $\mathbf{h}$  is to save memory. Since we aim at sampling a large number of points ( $\sim 268k$  in our case) all at once for the simulation, modulating  $\mathbf{h}$  will cache a hidden vector for each sampled points used in backpropagation; modulating  $\mathbf{W}$  on the other hand entails only one such operation, which will be shared for each sampled point.

### 5.1 Training Strategy

Similar to [Srinivasan et al. 2021], we use a two-stage training strategy. In the first stage, we calculate a plausible approximation of the actuation from the target poses, which will be used to pretrain  $\mathcal{N}_A$  without the differentiable simulator. Specifically, we approximate the actuation by dragging our volumetric simulation mesh into the given target shape with virtual zero-rest-length springs. The resulting local deformations are used to initialize the actuation tensors  $\mathcal{A}$ , as in [Srinivasan et al. 2021]. This pretraining serves as a warm-up and can be used to identify the appropriate dimension of the latent code  $z$ .

In the second stage, we fine-tune the parameters of  $\mathcal{N}_A$  using our simulator-integrated pipeline. The loss function  $\mathcal{L}$  for the inverse problem is defined as

$$\mathcal{L} = \sum_i |\hat{\mathbf{s}}_i - \mathbf{s}_i| + \alpha |1 - \hat{\mathbf{n}}_i^\top \mathbf{n}_i|, \quad (17)$$



where  $i$  refers to the vertices in the observation space,  $\mathbf{s}$  and  $\mathbf{n}$  indicate the observed vertex positions and normals,  $\hat{\mathbf{s}}$  and  $\hat{\mathbf{n}}$  are the vertex positions and normals mapped from the simulation space  $\mathbf{u}$ , and  $\alpha \in \mathbb{R}$  is a control weight.

## 5.2 Continuous Resolution

The main motivation for our presented model is its generalization capability across shapes and resolutions. Resolution invariance is a prerequisite for practical use, as it enables a decision on the resolution at test time depending on the given requirements on speed or geometric fidelity. Our network captures the dominant frequencies of the signal, and thus transfer learning on a different resolution (1 epoch) would be sufficient to make the network consistent with the new discretization. Alternatively, we propose to condition the actuation-generative network  $\mathcal{N}_A$  on a continuous resolution input, such that we only need one forward pass to get the optimal actuation. To do so, we add another small MLP on top of  $\mathcal{N}_A$ , that takes as input the resolution value, and adds up the output to the shape latent code  $\mathbf{z}$ . A detailed overview of the final architecture can be found in the Supplemental Material.

## 6 FACIAL ANIMATION

For the specific application of our method to faces, we have extended the underlying geometry with bone structures. We can articulate diverse expressions by only considering the relative motion between the skull and the mandible. Thus, the skull is fixed all the time, and we include the mandible kinematics in the optimization.

### 6.1 Bone Kinematics

Different from the actuation mechanism acting on the inside of the soft tissue, the bone structure is located at the boundary and defines the Dirichlet condition in the simulation. We introduce a second network  $\mathcal{N}_B$ , which takes a spatial point on the bone region of the object as input and outputs the transformed position.

For learning the parameters of  $\mathcal{N}_B$ , we divide the simulation vertices  $\mathbf{u}$  into two parts  $\mathbf{u} = [\mathbf{u}_c, \mathbf{u}_d]$ , where  $\mathbf{u}_c$  denotes the nodal vertices located inside the soft tissue, and  $\mathbf{u}_d$  those on the bone region which defines the Dirichlet condition. Given that  $\mathbf{u}_d$  is calculated with  $\mathcal{N}_B$ , a quasi-static solution for  $\mathbf{u}_c^*$  satisfies

$$\nabla_{\mathbf{u}_c} E([\mathbf{u}_c^*, \mathbf{u}_d], \mathcal{A}) = \mathbf{0}. \quad (18)$$

With such an implicit constraint between  $\mathbf{u}_c^*$  and  $\mathbf{u}_d$ , we can derive the sensitivity matrix  $\partial \mathbf{u}_c^* / \partial \mathbf{u}_d$  by taking the derivative of Eqn. (18) on both sides with respect to  $\mathbf{u}_d$ :

$$\frac{\partial \mathbf{u}_c^*}{\partial \mathbf{u}_d} = - \underbrace{\left( \frac{\partial \nabla_{\mathbf{u}_c} E}{\partial \mathbf{u}_c} \right)^{-1}}_{\mathbf{H}_{cc}} \underbrace{\frac{\partial \nabla_{\mathbf{u}_c} E}{\partial \mathbf{u}_d}}_{\mathbf{H}_{cd}}, \quad (19)$$

where  $\mathbf{H}_{cc}$  and  $\mathbf{H}_{cd}$  are the submatrices of the Hessian  $\mathbf{H}_u$  (which we derive in the Sec. 4). They can be constructed by only keeping the relevant rows and columns of  $\mathbf{H}_u$ . We use the sensitivity matrix to calculate the gradient accumulated on the evaluated Dirichlet condition, building the connection between the network  $\mathcal{N}_B$  and the simulator.

## 6.2 Extended Architecture

The architecture for facial animation consists of three parts as illustrated in Fig. 2; an encoder, an actuation-generative network  $\mathcal{N}_A$ , and a jaw-generative network  $\mathcal{N}_B$ . The encoder outputs a latent code  $\mathbf{z}$  representing an input face shape  $\mathbf{s}$ .  $\mathcal{N}_A$  is conditioned on  $\mathbf{z}$  to generate the actuation matrix  $\mathbf{A}$  for each input point in the soft tissue domain.  $\mathcal{N}_B$  is also conditioned on  $\mathbf{z}$  to produce the transformed bone position for each input point in the mandible domain.

For  $\mathcal{N}_B$ , we assume rigid motion, and use an MLP to regress the transformation matrix  $\mathbf{T}$ , from  $\mathbf{z}$  such that  $\mathcal{N}_B(\mathbf{x}, \mathbf{z}) = \mathbf{T}(\mathbf{z})\mathbf{x}$ . Note that  $\mathbf{T}$  is shared for each input point but differs from expression to expression. During the first training stage, we approximate the mandible position using the tracking method presented in Zoss et al. [2019]. The second stage then includes  $\mathcal{N}_B$  in the optimization. More details can be found in the Supplemental Material.

## 7 IMPLEMENTATION

The two main designs of our pipeline are the differentiable quasi-static physics solver and the conditional coordinate-based network. For the simulator, we use a hexahedral mesh, leveraging the advantage that the discretization is straightforward and can be robustly applied. We explicitly handle non-manifold topological features by duplicating vertices [Mitchell et al. 2015] around a cut, e.g., around the lips to enable mouth opening. We use Projective Dynamics [Bouaziz et al. 2014] as our forward solver, and differentiable Projective Dynamics [Du et al. 2022] as backward solver. Specifically, we uniformly sample 8 points inside each hexahedron to estimate the energy function  $E$ . We employ a highly optimized singular value decomposition method [McAdams et al. 2011] for polar decomposition, and use cuSPARSE from CUDA Toolkit to solve large sparse linear systems on the GPU. As in [Du et al. 2022] we utilize the adjoint method to avoid the explicit calculation of  $\mathbf{H}_u^{-1}\mathbf{H}_\Omega$ , and adopt the scalability optimization from [Srinivasan et al. 2021] to implicitly evaluate  $\mathbf{H}_\Omega$ . The optimization is terminated when either the relative progress is less than 1e-6 or if the maximum of 300 iterations is reached.

For the network architecture, we use our proposed modulated SIREN layer as the backbone for the actuation-generative network  $\mathcal{N}_A$ . For the other layers of  $\mathcal{N}_A$  and  $\mathcal{N}_B$  we use a fully connected layer with a LeakyReLU nonlinearity. We have chosen  $\alpha = 0$  for the normal weight in the loss function of Eqn. (17) for both the starfish and human body examples, and  $\alpha = 1$  for the face examples, as we found that the inclusion of the normal constraint positively affects the fidelity of the resulting wrinkles and facial details. We use ADAM optimizer [Kingma and Ba 2015] to jointly train our networks. For resolution invariance, we use the number of sampled points on the simulation mesh as the resolution indicator. We adopt the network pretrained on one resolution (268K), and use our simulator integrated pipeline to enable continuous resolution conditioning via transfer learning. We uniformly sample 20 resolutions from 42K to 268K for training and 25 different resolutions for testing. We run all the experiments using PyTorch [Paszke et al. 2019] on an NVIDIA RTX 2080Ti with the simulator integrated as a layer. For additional information on the architecture and training we refer to the Supplemental Material.

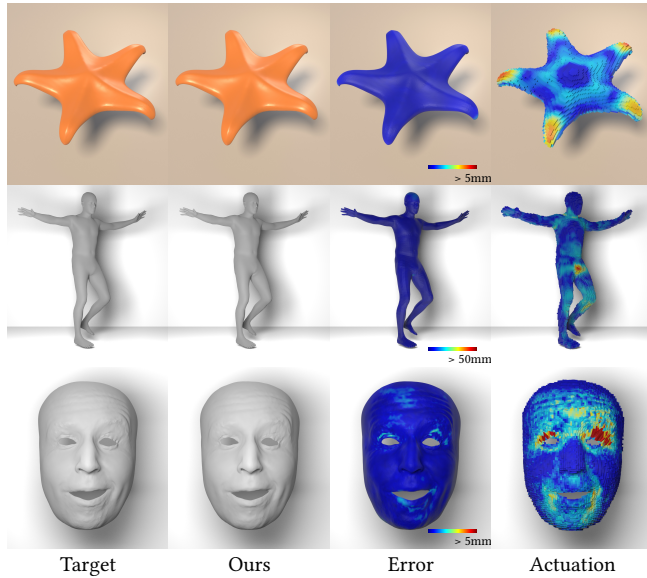


Fig. 3. Training results using the starfish, human body and face datasets. From left to right: target pose, simulation result, reconstruction error, optimized actuation magnitudes.

## 8 EXPERIMENTS

We demonstrate our results using different soft body types. We first describe the datasets, followed by a discussion of the results.

### 8.1 Datasets

*Volumetric soft body.* We generate a synthetic dataset using a starfish surface mesh consisting of 25K vertices and 50K faces. The voxelized representation contains 27K hexahedra and 32K vertices, which we use as our simulation mesh. The dynamic sequence used for learning is generated following the method of DiffPD [Du et al. 2022], where the constitutive model consists of the co-rotational elasticity and the volume-preserving term. We generate 200 frames in total, from which 180 are used for training and 20 for testing. We sample 216K points at once to evaluate the actuation values for the simulation. We use PCA dimensionality reduction to express each shape of the sequence as a 16 dimensional vector, which is then

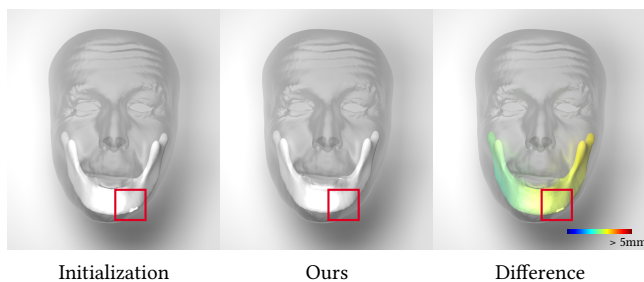


Fig. 4. Optimized jaw position (middle) and color-coded magnitude of the movement (right) compared to the initialization (left).




Fig. 5. Simulation of unseen target poses. From left to right: target pose, simulated result, reconstruction error, optimized actuation magnitudes.

used as input to the encoder. Training takes about 1 day, and testing takes about 10s for 1 forward pass.

*Human body motion.* We use a dancing sequence of the AMASS dataset [Mahmood et al. 2019]. We sample 1600 frames for training and 100 frames for testing. We create a voxelized version of the template body mesh, containing 42K hexahedra and 51K vertices. We sample 336K points at once to evaluate the actuation values for the simulation. We use the provided skeleton positions to get the Dirichlet boundary conditions in the simulation, but do not optimize the kinematic chain for simplicity. The skeleton vertices on the simulation mesh can be obtained by forward kinematics. We directly use the 63 dimensional SMPL body pose parameters [Loper et al. 2015] as the global pose descriptor, which is followed by 3 fully connected layers. Training takes about 5 days. Testing takes about 15s for 1 forward pass.

*Facial expressions.* We use a subset of the facial performance dataset introduced in Zoss et al. [2020]. This subset contains 23 expression blendshapes and 641 frames of diverse dialog and facial workout performances for one subject. We extract 100 frames as testing set and use the remainder for training. As this dataset contains only surface meshes, we create a voxelized version of the reference (neutral) mesh, containing 33K hexahedra and 46K vertices. In our experiments with this dataset we sample 268K points at once to evaluate the actuation values for the simulation. We use the 23 expression blendshapes to fit a per-frame blendweight vector which we use as input to our encoder. We fit a template mandible mesh to the neutral face following the method proposed by [Zoss et al. 2018] and approximate the mandible position for each expression using the tracking method presented in [Zoss et al. 2019]. Training takes about 3 days. Testing takes about 11s for 1 forward pass.




Fig. 6. Latent space interpolation between two selected expressions.

## 8.2 Results

*Target shape matching.* Fig. 3 shows the training results for the three datasets. In all examples, our model can reliably reproduce the target shapes with the optimized parameters, as indicated by the low reconstruction error. The actuation signals are color-coded on the hexahedral mesh on the right, where blue and red indicate a low and high value, respectively. Note that the network contains only 0.3M parameters but has the representational power to reliably reproduce fine details such as wrinkles. In contrast, the number of parameters of explicit models depends on the simulation mesh (discretization) and is typically two orders of magnitude larger. For example, Srinivasan et al. [2021] uses about 43M parameters in the first three layers of the decoder (not including the final and sparse layer with custom connectivity).

For the face we optimize both actuation signals and mandible kinematics. Thanks to the jaw-generative network we are able to directly compute a jaw position that is compatible with the actuation (Fig. 4). This was not possible in previous work [Srinivasan et al. 2021], where the boundary conditions associated with skeletal attachments are presumed known at all time instances. As expected, the differentiable connection between the jaw and simulator reduces the error (see the Ablation Study below) and even successfully resolves collision artifacts (highlighted in the red box) of the given initialization [Zoss et al. 2019].

*New poses.* With our trained model we can fit new expressions by keeping the weights of  $\mathcal{N}_A$  and  $\mathcal{N}_B$  fixed and optimizing for the latent code that yields the best fit, as in [Srinivasan et al. 2021]. To do so, we first use our encoder to get an initialized latent code, then optimize it for 10 iterations. Fig. 5 illustrates the result for unseen target poses and the associated low errors.

In Fig. 6 we interpolate between two selected expressions, *i.e.*, from smiling to angrily surprised, indicating smooth transitions despite

the non-linearity of the deformation. The actuation and jaw position are also consistently interpolated.




Fig. 7. Results with four unseen input resolutions used at test time (from left to right: 51K, 117K, 192K, 258K).

*Continuous resolution.* As demonstrated in Fig. 7, our method successfully generalizes to continuous resolution input. We show results for four different, unseen input resolutions ranging from 51K to 258K. More results are shown in the supplemental video. The resolution invariance property paves the way for efficient practical use where the resolution can be chosen at test time depending on the requirements at hand.



Table 1. Average vertex error (mm) for different methods

Method	Baseline-Stage1	Baseline-NDiffJaw	Ours
Error	0.88	0.42	0.37

*Comparison with Blendshapes.* Fig. 8 compares our method with blendshapes. The error maps show clearly the superiority of our method, which can be accounted to the non-linearity of our model. Visually, this property is important for reproducing geometric details and wrinkles, which in this example are conspicuous on the cheek.




Fig. 8. Comparison of blendshapes (left) and our method (right).

*Ablation Study.* We first show the necessity of our normal constraint in the loss function for the face example. For this, we train a network without the normal constraint (Baseline-NNorm). Fig. 9 shows the difference between Baseline-NNorm and ours. The inclusion of the normal constraint positively affects the fidelity of the resulting wrinkles and facial details.

We examine the importance of our simulator-integrated pipeline using the face dataset. To do this, we train two more networks with different simplifications: one trained with only stage 1 (Baseline-Stage1) and another trained without the differentiable connection between bone and simulator (Baseline-NDiffJaw). Average vertex error is reported in Table (1). It can be seen that our full model produces the best accuracy. Note that the error difference between our full model and Baseline-NDiffJaw is not trivial, since we average the error over 44K vertices. The differentiable connection between the bone and simulator mainly reduces the error around the mandible region, as shown in Fig. 10.

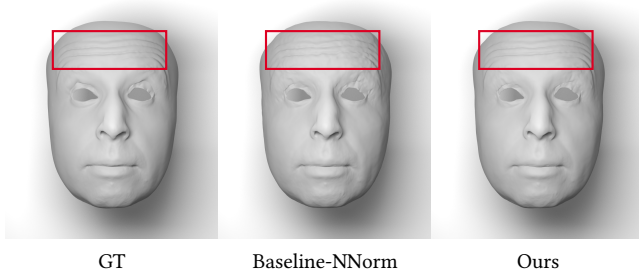


Fig. 9. Comparison with Baseline-NNorm (without the normal constraint).

## 9 CONCLUSION

In this work, we have presented a physics- and data-driven method for computing control parameters of actuated soft bodies. We have demonstrated that the implicit representation offers the unique




Fig. 10. Comparison with Baseline-NDiffJaw. We show two different examples.

advantage of being shape invariant. This property enables the application of the method to different soft body input meshes, as demonstrated for three different datasets. Moreover, it renders the method agnostic to the discretization resolution used during training and testing, which is often a prerequisite in practical settings. For shapes that require an internal bone structure, such as the human face, we can additionally use an implicit formulation to parameterize the bone kinematics. Putting rigorously derived physics priors into the core of the learning facilitates the model to reproduce non-linear deformations and fine geometric details like wrinkles.

While our implicit formulation is general across soft body models, it relies on some specific design choices. First, the derived closed-form Hessian matrices are specific to the shape targeting constitutive model and thus not generally applicable to other models. Similarly, the actuation initialization would need to be changed with a different constitutive model. Second, our implicit model of bone kinematics is targeted for the mandible and thus only supports a single rigid transformation. In the human body examples, we use the provided skeleton position to define the boundary conditions in the simulation but do not optimize for the kinematic chain. Our method could be extended to multiple bones and non-rigid transformations in future work. Third, our method only considers actuation and bone kinematics. Additional mechanical parameters could be taken into account, such as in Kadleček and Kavan [2019] where heterogeneous stiffness and prestrain are optimized for the face.

As discussed by Srinivasan et al. [2021], another noticeable limitation of our model is that there are no explicit guarantees that the optimized mechanical parameters have an anatomical meaning. However, the chosen design alleviates the need for an anatomically correct geometric model and thus substantially reduces manual modeling labor, while still generating physically plausible deformations. In accordance with related work, our method does not generalize well across different subjects as the implicit function encodes the actuation for a specific shape. While we can map the optimized actuation signals to another input shape, this will yield larger errors than if we were to retrain the model with the respective facial input scans or body poses. Therefore, future work is needed to support important downstream applications such as expression retargeting.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive comments. We thank Daniel Dorda for his valuable feedback. The work is supported by the Swiss National Science Foundation under Grant No.: 200021\_197136.

## REFERENCES

- Moritz Bächer, Espen Knoop, and Christian Schumacher. 2021. Design and Control of Soft Robots Using Differentiable Simulation. *Current Robotics Reports* 2, 2 (6 2021), 211–221. <https://doi.org/10.1007/s43154-021-00052-7>
- Michael Bao, Matthew Cong, Stéphane Grabli, and Ronald Fedkiw. 2018. High-Quality Face Capture Using Anatomical Muscles. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (12 2018). <http://arxiv.org/abs/1812.02836>
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: fusing constraint projections for fast simulation. *ACM Transactions on Graphics* 33, 4 (7 2014), 1–11. <https://doi.org/10.1145/2601097.2601116>
- Dan Casas and Miguel A. Otaduy. 2018. Learning Nonlinear Soft-Tissue Dynamics for Interactive Avatars. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 1 (7 2018), 1–15. <https://doi.org/10.1145/3203187>
- Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2020. pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. *CVPR 2021 (Oral)* (12 2020). <http://arxiv.org/abs/2012.00926>
- Prashanth Chandran, Derek Bradley, Markus Gross, and Thabo Beeler. 2020. Semantic Deep Face Models. In *2020 International Conference on 3D Vision (3DV)*. IEEE, 345–354. <https://doi.org/10.1109/3DV50981.2020.00044>
- Jonas Degraeve, Michiel Hermans, Joni Dambre, and Francis wyffels. 2019. A Differentiable Physics Engine for Deep Learning in Robotics. *Frontiers in Neurorobotics* 13 (3 2019). <https://doi.org/10.3389/fnbot.2019.00006>
- Tao Du, Kui Wu, Pingchuan Ma, Sebastian Wah, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. 2022. DiffPD: Differentiable Projective Dynamics. *ACM Transactions on Graphics* 41, 2 (4 2022), 1–21. <https://doi.org/10.1145/3490168>
- Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 2020. 3D Morphable Face Models—Past, Present, and Future. *ACM Transactions on Graphics* 39, 5 (10 2020), 1–38. <https://doi.org/10.1145/3395208>
- Lawson Fulton, Vismay Modi, David Duvenaud, David I. W. Levin, and Alec Jacobson. 2019. Latent-space Dynamics for Reduced Deformable Simulation. *Computer Graphics Forum* 38, 2 (5 2019), 379–391. <https://doi.org/10.1111/cgf.13645>
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Transactions on Graphics* 39, 6 (11 2020), 1–15. <https://doi.org/10.1145/3414685.3417766>
- David Hahn, Pol Banzet, James M. Bern, and Stelian Coros. 2019. Real2Sim: visco-elastic parameter estimation from dynamic motion. *ACM Transactions on Graphics* 38, 6 (11 2019), 1–13. <https://doi.org/10.1145/3355089.3356548>
- Philipp Holl, Vladlen Koltun, and Nils Thuerey. 2020. Learning to Control PDEs with Differentiable Physics. *International Conference on Learning Representations* (1 2020). <http://arxiv.org/abs/2001.07457>
- Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. 2019a. DiffTaichi: Differentiable Programming for Physical Simulation. *International Conference on Learning Representations* (10 2019). <http://arxiv.org/abs/1910.00935>
- Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B. Tenenbaum, William T. Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. 2019b. ChainQueen: A Real-Time Differentiable Physical Simulator for Soft Robotics. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 6265–6271. <https://doi.org/10.1109/ICRA.2019.8794333>
- Xun Huang and Serge Belongie. 2017. Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization. *ICCV 2017 (Oral)* (3 2017). <http://arxiv.org/abs/1703.06868>
- Alexandru-Eugen Ichim, Petr Kadleček, Ladislav Kavan, and Mark Pauly. 2017. Phace: Physics-based Face Modeling and Animation. *ACM Transactions on Graphics* 36, 4 (7 2017), 1–14. <https://doi.org/10.1145/3072959.3073664>
- Petr Kadleček and Ladislav Kavan. 2019. Building Accurate Physics-based Face Models from Data. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 2, 2 (7 2019), 1–16. <https://doi.org/10.1145/3340256>
- Theodore Kim and David Eberle. 2020. Dynamic deformables: implementation and production practicalities. In *ACM SIGGRAPH 2020 Courses*. ACM, New York, NY, USA, 1–182. <https://doi.org/10.1145/3388769.3407490>
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (12 2015). <http://arxiv.org/abs/1412.6980>
- Gergely Klár, Andrew Moffat, Ken Museth, and Eftychios Sifakis. 2020. Shape Targeting: A Versatile Active Elasticity Constitutive Model. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks*. ACM, New York, NY, USA, 1–2. <https://doi.org/10.1145/3388767.3407379>
- Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. 2015. Recent Advances in Facial Appearance Capture. *Computer Graphics Forum* 34, 2 (5 2015), 709–733. <https://doi.org/10.1111/cgf.12594>
- J P Lewis, Ken Anjyo, Taehyun Rhee, Mengjie Zhang, Fred Pighin, and Zhigang Deng. 2014. Practice and Theory of Blendshape Facial Models. In *Eurographics 2014 - State of the Art Reports*, Sylvain Lefebvre and Michela Spagnuolo (Eds.). The Eurographics Association. <https://doi.org/10.2312/egst.20141042>
- Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, and Hao Li. 2020. Learning Formation of Physically-Based Face Attributes. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 3407–3416. <https://doi.org/10.1109/CVPR42600.2020.00347>
- Junbang Liang, Ming C. Lin, and Vladlen Koltun. 2019. Differentiable Cloth Simulation for Inverse Problems. In *Advances in Neural Information Processing Systems*. <https://doi.org/10.5555/3454287.3454357>
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics* 34, 6 (11 2015), 1–16. <https://doi.org/10.1145/2816795.2818013>
- M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and T.Y. Kim. 2020. Primal/Dual Descent Methods for Dynamics. *Computer Graphics Forum* 39, 8 (12 2020), 89–100. <https://doi.org/10.1111/cgf.14104>
- Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. 2019. AMASS: Archive of Motion Capture as Surface Shapes. *ICCV* (4 2019). <http://arxiv.org/abs/1904.03278>
- Aleka McAdams, Andrew Selle, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. 2011. *Computing the singular value decomposition of 3x3 matrices with minimal branching and elementary floating point operations*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (12 2019). <http://arxiv.org/abs/1812.03828>
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*. 405–421. [https://doi.org/10.1007/978-3-030-58452-8\\_124](https://doi.org/10.1007/978-3-030-58452-8_124)
- Nathan Mitchell, Court Cutting, and Eftychios Sifakis. 2015. GRIDiron: An interactive authoring and cognitive training foundation for reconstructive plastic surgery procedures. *ACM Transactions on Graphics* 34, 4 (7 2015), 1–12. <https://doi.org/10.1145/2766918>
- Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. 2021. StyleSDF: High-Resolution 3D-Consistent Image and Geometry Generation. (12 2021). <http://arxiv.org/abs/2112.11427>
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (1 2019). <http://arxiv.org/abs/1901.05103>
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32. 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- Gerard Pons-Moll, Javier Romero, Naureen Mahmood, and Michael J. Black. 2015. Dyna: a model of dynamic human shape in motion. *ACM Transactions on Graphics* 34, 4 (7 2015), 1–14. <https://doi.org/10.1145/2766993>
- Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming C. Lin. 2021. Differentiable Simulation of Soft Multi-body Systems. *Advances in Neural Information Processing Systems* 34 (2021).
- Yi-Ling Qiao, Junbang Liang, Vladlen Koltun, and Ming C. Lin. 2020. Scalable Differentiable Physics for Learning and Control. *International Conference on Machine Learning* (7 2020). <http://arxiv.org/abs/2007.02168>
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W. Battaglia. 2020. Learning to Simulate Complex Physics with Graph Networks. *International Conference on Machine Learning* (2 2020). <http://arxiv.org/abs/2002.09405>
- Igor Santesteban, Elena Garces, Miguel A. Otaduy, and Dan Casas. 2020. SoftSMPL: Data-driven Modeling of Nonlinear Soft-tissue Dynamics for Parametric Humans. *Computer Graphics Forum* 39, 2 (5 2020), 65–75. <https://doi.org/10.1111/cgf.13912>
- Eftychios Sifakis, Igor Neverov, and Ronald Fedkiw. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Transactions on Graphics* 24, 3 (7 2005), 417–425. <https://doi.org/10.1145/1073204.1073208>
- Vincent Sitzmann, Julien N P Martel, Alexander W Bergman, David B Lindell, and Gordon Wetzstein. 2020. Implicit Neural Representations with Periodic Activation Functions. In *NeurIPS 2020*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/>



- paper/2020/hash/53c04118df112c13a8c34b38343b9c10-Abstract.html
- Sangeetha Grama Srinivasan, Qisi Wang, Junior Rojas, Gergely Klár, Ladislav Kavan, and Eftychios Sifakis. 2021. Learning active quasistatic physics-based models from data. *ACM Transactions on Graphics* 40, 4 (8 2021), 1–14. <https://doi.org/10.1145/3450626.3459883>
- Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. 2020. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. *NeurIPS 2020 (Spotlight)* (6 2020). <http://arxiv.org/abs/2006.10739>
- Jingwei Tang, Vinicius C. Azevedo, Guillaume Cordonnier, and Barbara Solenthaler. 2021. Honey, I Shrunk the Domain: Frequency-aware Force Field Reduction for Efficient Fluids Optimization. *Computer Graphics Forum* 40, 2 (5 2021), 339–353. <https://doi.org/10.1111/cgf.142637>
- Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. 2021. CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. (10 2021). <http://arxiv.org/abs/2110.09788>
- Gaspard Zoss, Thabo Beeler, Markus Gross, and Derek Bradley. 2019. Accurate markerless jaw tracking for facial performance capture. *ACM Transactions on Graphics* 38, 4 (7 2019), 1–8. <https://doi.org/10.1145/3306346.3323044>
- Gaspard Zoss, Derek Bradley, Pascal Bérard, and Thabo Beeler. 2018. An empirical rig for jaw animation. *ACM Transactions on Graphics* 37, 4 (8 2018), 1–12. <https://doi.org/10.1145/3197517.3201382>
- Gaspard Zoss, Eftychios Sifakis, Markus Gross, Thabo Beeler, and Derek Bradley. 2020. Data-driven extraction and composition of secondary dynamics in facial performance capture. *ACM Transactions on Graphics* 39, 4 (7 2020). <https://doi.org/10.1145/3386569.3392463>