DOI: 10.1111/cgf.14623
ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2022
D. L. Michels and S. Pirk
(Guest Editors)
Volume 41 (2022), Number 8

Differentiable Simulation for Outcome-Driven Orthognathic Surgery Planning

D. Dorda[1], D. Peter[2], D. Borer[1], N.B. Huber[2], I. Sailer[3], M. Gross[1], B. Solenthaler[1], B. Thomaszewski [1,2]

[1]ETH Zurich, Switzerland
[2]Align Technology
[3]University of Geneva, Switzerland

Abstract

Algorithms at the intersection of computer graphics and medicine have recently gained renewed attention. A particular interest are methods for virtual surgery planning (VSP), where treatment parameters must be carefully chosen to achieve a desired treatment outcome. FEM simulators can verify the treatment parameters by comparing a predicted outcome to the desired one. However, estimating the optimal parameters amounts to solving a challenging inverse problem. In current clinical practice it is solved manually by surgeons, who rely on their experience and intuition to iteratively refine the parameters, verifying them with simulated predictions.

We prototype a differentiable FEM simulator and explore how it can enhance and simplify treatment planning, which is ultimately necessary to integrate simulation-based VSP tools into a clinical workflow. Specifically, we define a parametric treatment model based on surgeon input, and with analytically derived simulation gradients we optimise it against an objective defined on the visible facial 3D surface. By using sensitivity analysis, we can easily explore the solution-space with first-order approximations, which allow the surgeon to interactively visualise the effect of parameter variations on a given treatment plan. The objective function allows landmarks to be freely chosen, accommodating the multiple methodologies in clinical planning. We show that even with a very sparse set of guiding landmarks, our simulator robustly converges to a feasible post-treatment shape.

1. Introduction

A long-standing goal at the intersection of computer graphics, mechanical engineering and medicine is to create accurate, robust and easy-to-use surgical simulators [Gro98]. Their physics-based outcome predictions allow surgeons to test different treatment plans *in silico*, and assure the patient of their post-treatment appearance, thus facilitating informed decision-making. In orthognathic and maxillofacial surgery, even minor changes in facial symmetry and proportion can have a large impact on perceived beauty and quality of life [Far94]. To guarantee the optimal outcome, the treatment plan must be especially precise [BH13]. This creates a need for advanced planning tools, which must nonetheless easily integrate into the clinical process [CMF*21].

Planning orthognathic surgery is labour intensive, and necessitates the collaboration of an interdisciplinary medical and technical team [CMF*21, LYW*21]. This team determines the goals of the treatment, i.e. the desired post-treatment state which maximises the patient's dental and facial aesthetics and function. Virtual surgery planning (VSP) provides a systematic and mathematically principled way of evaluating treatment plans, leading to better predictions of post-treatment outcome when compared to traditional methods [CMF*21]. Unfortunately, VSP workflows are already time-intensive in the measurement and planning stages. This has become the new barrier to adoption; for "simple" cases, traditional methods are still preferred [CMF*21].

Meanwhile, research on medical applications of FEM continues to enhance forward simulation, focusing on more accurate modelling of materials and physical effects [YDX*14, KKR*19, VWWSGB21]. A predetermined treatment plan serves as the input to these methods. However, when we take a broader look at the VSP workflow, we see that the initial step of VSP is determining a treatment target. It is then the treatment target which implicitly defines the treatment plan. In practice, several iterations of the treatment plan are created, each aiming to more accurately reach the target.

From this perspective, treatment planning becomes an inverse problem, where the treatment parameters which act as input to the simulator need to be optimised to yield an outcome as close as possible to the desired target. Developing an automatic method to optimise the treatment plan from a given target promises to simplify the way the tool is used in practice. Differentiable simulation offers a methodology for solving such problems, with many recent examples in graphics and robotics [IKKP17, KK19, SWR*21, HHD*21, ZKBT17, ZCT22].

We explore how to apply these techniques in the medical domain, and analyse the benefits they can yield. Specifically, we start with a parameterised model of double jaw surgery, and and define an objective function based on commonly measured clinical landmarks. We optimise the parameters of our treatment model, and show that the common landmarks are sufficient for driving an optimisation process. Further, we use sensitivity analysis to give first-order estimates of surgical outcome around a given equilibrium point. With this technique, the surgeon can rapidly explore the simulation space, which can benefit collaborative treatment planning and patient communication.

Our work's main contribution lies in formulating and solving an inverse treatment planning problem via differentiable simulation. This formulation is agnostic to the forward simulation model we have used—an improvement in simulation accuracy will yield corresponding improvements to the inverse problem solution. This has allowed us to test our method on a simple set of plausible surgical meshes from two models, which were crafted under the supervision of medical experts. Nonetheless, it means that validation against data from a diverse patient population is necessary before this method can be used clinically. Such diverse testing data would allow the forward simulation model to be tuned—taking into account material parameters and physical effects—and it would allow the treatment model to be validated against the full range of clinical conditions.

We are hopeful that differentiable simulation can be one of the keys to creating easy-to-use, widely available surgical simulation tools for future doctors.

2. Background and Related Work

Orthognathic surgery.

Orthognathic surgeries are standard treatments used to correct severe cases of jaw malocclusion - commonly categorised by Angle's classes - or facial asymmetry [RE14, Rey10, HNGK19]. These conditions, if left untreated, negatively affect oral and emotional health [Lar17, TLTdL22]. An example illustrating the occlusion classes is shown in Figure 1.

The treatment which we model, the double jaw osteotomy, can be used to address both Class II, Class III and asymmetric cases due to its powerful long-term effect on the facial profile, and low complication risk [BH13]. During this procedure, a section of the mandible (lower jaw-bone) and maxilla (upper jaw-bone) are cut, moved into a new position, and fixed in place with metal braces [Rey10].

Virtual surgery planning (VSP).

Orthognathic surgery is planned by an interdisciplinary team of surgeons, orthodontists and technicians. In its first stage, landmarks are identified on the soft-tissue and bony surface [RMV*20, LYW*21, Rey10]. The inter-landmark distances are compared to optimal reference values, which allows the treatment goal to be defined [PP70, Ste60, And15]. The treatment goal can be defined with respect to different sets of landmarks, which inadvertently prioritise different surgical outcomes [RMV*20]. To support the range

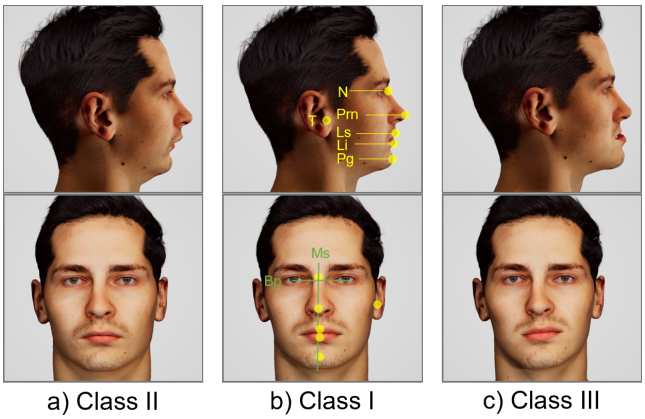

Figure 1: *Visualisation of occlusion classes and the landmarks which are used during treatment planning. The labelled landmarks are the nasion (N), pronasale (Prn), upper lip (Ls), lower lip (Li), Pogonion (Pg) and Tragion (T) [PP70]. The bipupilary and mid-sagittal lines are also labelled.*

of clinical methods, we do not rely on specific landmarks, but allow them to be freely chosen. For our experiments, we use a landmarking procedure based on [PP70] and [HNGK19] for malocclusion and asymmetry correction. The landmarks used are visualised in Figure 1, and are further discussed in Section 3.2.

Many software products have been built to facilitate VSP [LL15]. There are ongoing efforts to create a unified environment for performing all the necessary tasks of VSP, including medical image registration, 3D surface extraction, cephalometric analysis, virtual osteotomy, interactive bone movement, and post-operative soft-tissue prediction [LYW*21, YML*17]. Despite the value provided by outcome predictions, clinical software seldom supports the generation of simuation-based predictions. This can partly be attributed to the difficulty inherent to setting up and running the simulation [LL15]. Differentiable simulation offers a larger feature set, and the promise to simplify existing simulator functions, which incentivises its inclusion into future medical software.

Facial tissue modelling and outcome prediction.

Modelling and simulation of facial tissues in orthognathic surgery has been accomplished using data or biomechanics based methods. Early data-driven work relies on statistical analysis of the relationship between the displacement of soft and hard tissue landmarks in pre- and post- surgical states [AKFT15]. Modern approaches use machine learning to predict healthy bone shape [XDK*21] or soft-tissue [LKF*22]. Inference is usually fast and easy to perform. However, it comes at the cost of prohibitively large training datasets, the loss of an interactive environment for exploring the solution space, and no hard constraints on the physical viability of the predictions.

Physics-based mass-spring [KGC*96] and hyperelasitc [CLP03] models of soft tissue were amongst the first simulation approaches to be developed and tested. Hyperelastic modelling continues to

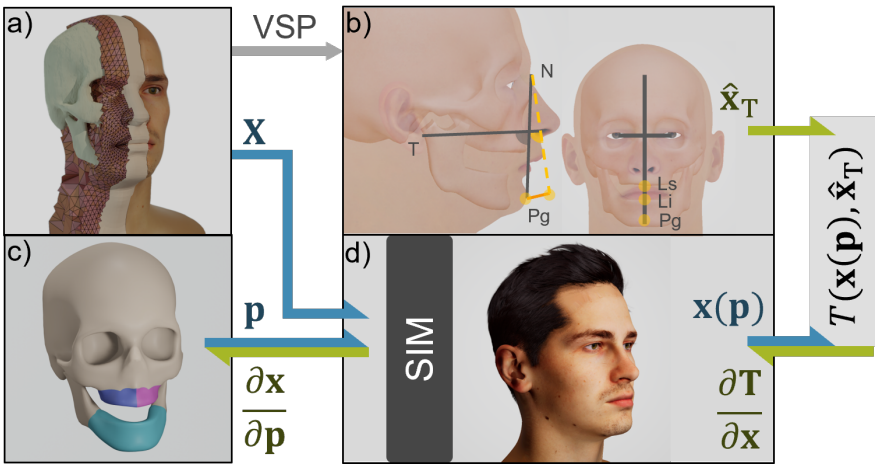

Figure 2: *An overview of our data, and pipeline. a) The anatomic model: tetrahedral mesh* $\boldsymbol{X}$ *and surface scan. b) Landmarks used during treatment planning (grey), and their goal positions (yellow) [PP70, Rey10] set the target vector* $\boldsymbol{x}_T$*. c) Treatment parameters* $\boldsymbol{p}$ *define blendshape-based deformation of surgically affected regions, shown in colour. The remaining bone surface is fixed in place. d) The forward simulator (blue arrows) predicts post treatment outcome* $\boldsymbol{x}(\boldsymbol{p})$*. The differentiable solver (green arrows) optimises* $\boldsymbol{p}$ *to minimise* $T(\boldsymbol{x}(\boldsymbol{p}), \hat{\boldsymbol{x}}_T)$*, the difference between prediction and goal.*

be used for recent works [APG*21, KKR*19, VWWSGB21]. Recent works often investigate previously unmodelled physical phenomena, such as time-dependent bone healing [VWWSGB21], lip and mucosa sliding [KKR*19], bone-cutting forces [YDX*14], or constrained motion models of the TMJ [OVT*08]. Several of these lines of enquiry parallel problems from computer graphics [ZBBB18]. Most commonly, material parameters are set from values reported in medical literature [XY15], and treatment parameters (i.e. bone movement) are assumed to be the input.

Differentiable simulation.

Differentiable physics solvers allow for complex inverse problems to be solved using gradient-based methods, and have found many applications in computer graphics and engineering. They have been used for shape optimisation [ZKBT17], or for estimating material parameters of soft robots [HBBC19] or cloth [LLK19]. There is increasing support for differentiable models of phenomena such as contact [GHZ*20], cutting [HMN*21], and plasticity [HHD*21].

In facial applications, Sifakis et al. develop an anatomic muscle model whose activations are optimised through a differentiable finite element simulation to match motion capture data [SNF05]. Ichim et al. generalise the muscle models and control a physics-based facial performance with blendshapes [IKKP17]. Kadleček and Kavan optimise facial material properties [KK19]. Impressive results for facial animation have been achieved by combining differentiable quasi-static solvers with neural networks [SWR*21]. Despite some speculation about the utility of differentiable simulation for medical applications [IKKP17, KK19], there have been no follow up works that elaborate on how to apply these techniques to the clinical process.

3. Methods

We create a bespoke differentiable simulator based on the finite element method, and show two applications that simplify surgical planning. In Section 3.1, we outline the forward simulation, including data processing and the parametric treatment model. In Section 3.2, we describe how treatment targets are set. We then calculate simulation gradients and use them for parameter optimisation and rapid first-order outcome prediction.

3.1. Simulation

Data processing. Our anatomical model is made from a cone-beam CT (CBCT) and a 3D facial surface scan [BBB*10, BHB*11]. We create a tetrahedral simulation mesh of the facial soft tissue, which we treat as the volume bounded by the skin surface on the outside and the bone surface on the inside. The rest-state vertex positions are referred to as $\mathbf{X} \in \mathbb{R}^{N\times 3}$, where $N \simeq 4 \times 10^4$ is the number of vertices. Vertex positions of the deformed mesh are referred to as $\mathbf{x}$. The facial scan is used to generate high-quality visualisations of the low-resolution deformed simulation mesh, using an embedding technique similar to Mezger et al. [MTPS09]. Since established processes are used for extracting surface data from CBCTs, registering them to the face scan, creating a tetrahedral simulation mesh and rendering our results, we leave the details of these steps to the Appendix.

Vertex types. Our simulation mesh models the facial soft tissue, and we distinguish several types of vertices on its boundary. Vertices corresponding to the points on the bone whose position is unaffected by treatment (e.g. on the skull) are fixed in-place using a Dirichlet boundary condition. Using the bone surface in this manner lets us keep the skull hollow, drastically reducing the number of vertices in the simulation. Vertices which lie on the osteotomised

bone surface, and move during surgery, are referred to as active vertices $\mathbf{X}_\mathrm{A}$. Vertices on which the treatment goal is defined are referred to as target vertices $\mathbf{X}_\mathrm{T}$. In our examples, these vertices are the landmarks on which treatments are defined, and we place them on the surface of the skin. These collections of vertices are visualised in the pipeline Figure 2b and 2c, and their role in the simulation is described in the sections below.

Treatment Model.

We model a simultaneous bilateral sagittal split osteotomy (BSSO) and a two-segmental LeFort I osteotomy [Rey10]. The bony regions affected by these procedures were identified with the help of a surgeon (see Figure 2c), who also guided us in sculpting several surgically plausible modifications to the jaw shape. By taking the difference between the starting position and the plausibly modified meshes, we obtain a blendshape for the boney surface of our simulation mesh. We can compute the goal positions of the active vertices as

$$\hat{\mathbf{x}}_\mathrm{A}(\mathbf{p}) = \mathbf{X}_\mathrm{A} + \sum_{i=1}^{N_p} p_i \mathbf{B}_i \tag{1}$$

Where each $\mathbf{B}_i$ is a blendshape vector, and N_p is the total number of blendshapes. The overall parameter vector defining all transformations is denoted by $\mathbf{p}$. Our process optimises for $\mathbf{p}$, thus finding combination of blendshape weights which result in a soft-tissue deformation that best satisfies the treatment goal. The only constraint on the treatment model is that it must be differentiable with respect to $\mathbf{p}$. As such, blendshapes were chosen to represent the treatment because they allow us to easily set-up and demonstrate the inverse problem optimisation.

Forward simulation.

Our forward simulation uses a linear FEM, and takes as input $\mathbf{X}$, $\mathbf{X}_\mathrm{A}$, and $\mathbf{p}$. We use a quasistatic simulation scheme, where each set of inputs maps to an equilibrium state $\mathbf{x}(\mathbf{p})$ defined implicitly through an energy minimisation problem (note $\mathbf{x}$'s dependence on $\mathbf{p}$ in the equilibrium state). The total energy of the system is defined as

$$E_\mathrm{tot}(\mathbf{x}, \mathbf{X}, \mathbf{p}) = E_\mathrm{NH}(\mathbf{x}, \mathbf{X}) + E_\mathrm{A}(\mathbf{x}_\mathrm{A}, \hat{\mathbf{x}}_\mathrm{A}(\mathbf{p})), \tag{2}$$

where E_NH is a Neo-Hookean elastic energy induced by the deformation of the soft-tissue elements, and E_A is a virtual potential energy which drives the active vertices to their goal position.

For the elastic energy, we use an hyper-elastic material [SB12], whose parameters are identified from literature as 0.9 g/ml, 5 kPa, and 0.47 for density, Young's modulus and Poisson's ratio respectively [XY15], and are assumed to be homogeneous and isotropic.

The virtual potential energy is an L_2 penalty between the target and current positions of the active vertices, with $k = 10^3$ as a stiffness weight which ensures that active vertices end up imperceptibly close to their goals:

$$E_\mathrm{A} = \frac{1}{2} k \|\mathbf{x}_\mathrm{A} - \hat{\mathbf{x}}_\mathrm{A}(\mathbf{p})\|_2^2 \quad . \tag{3}$$

At an equilibrium point (for a specific choice of parameters) $\mathbf{x}(\mathbf{p})$, the elastic and virtual forces balance, which corresponds to the minimum energy state, and satisfies the equilibrium condition

$$\frac{\partial E_\mathrm{tot}}{\partial \mathbf{x}} = \frac{\partial E_\mathrm{NH}}{\partial \mathbf{x}} + \frac{\partial E_\mathrm{A}}{\partial \mathbf{x}} = 0 \quad . \tag{4}$$

We find the local minimum of Equation 2 by applying Newton's method. We first calculate a search direction by solving the linear system $\frac{\partial^2 E}{\partial \mathbf{x}^2} \delta\mathbf{x} = \frac{\partial E_\mathrm{tot}}{\partial \mathbf{x}}$. If the stiffness matrix $\frac{\partial^2 E}{\partial \mathbf{x}^2}$ is not positive-definite, we iteratively apply diagonal regularisation to make it so. After solving for $\delta\mathbf{x}$, we use a backtracking line search to find the step size which decreases E_tot.

3.2. Inverse Simulation

Selecting treatment targets.

Through consulting maxillofacial surgeons and their literature [Rey10, PP70], we choose between three and seven commonly used, plausible landmarks to guide the optimisation. The number of landmarks varies depending on the complexity of the case. A symmetric Class II malocclusion requires only three landmarks to define a treatment goal, with additional landmarks being used for an asymmetric Class III case. The most important landmarks used are the pogonion (Pg), the lower lip (Li) and upper lip points (Ls). The tragion (T) and Nasion (N), as well as the bipupilarly and midsaggital lines are also used to set optimal target locations, with respect to healthy reference values [PP70, Rey10]. The constructions, which were done in 3D, are visualised in Figure 3. Landmarks defined like this becomes our vector of targets $\hat{\mathbf{x}}_\mathrm{T}$ and corresponding simulation vertices $\mathbf{X}_\mathrm{T}$.

Optimisation.

The target vectors let us define our objective function $T(\mathbf{x}(\mathbf{p}), \hat{\mathbf{x}}_\mathrm{T})$ and formulate a constrained optimisation problem

$$\begin{aligned} \min_{\mathbf{p}} \quad & T(\mathbf{x}(\mathbf{p}), \hat{\mathbf{x}}_\mathrm{T}) = \frac{1}{2} \|\mathbf{x}_\mathrm{T}(\mathbf{p}) - \hat{\mathbf{x}}_\mathrm{T}\|_2^2 \\ \text{s.t.} \quad & \mathbf{f}(\mathbf{x}(\mathbf{p})) = \mathbf{0}, \end{aligned} \tag{5}$$

where T is the L_2 distance between target landmarks at the equilibrium point and the desired target landmark location, and $\mathbf{f}$ is the static equilibrium constraint from Equation 4 expressed in terms of nodal forces.

We solve this problem using a Sparse Gauss-Newton solver [ZCT22]. Using this technique, the problem is reduced to solving the sparse system

$$\begin{bmatrix} \frac{\partial^2 T}{\partial \mathbf{x}^2} & 0 & \frac{\partial \mathbf{f}}{\partial \mathbf{x}}^\mathrm{T} \\ 0 & 0 & \frac{\partial \mathbf{f}}{\partial \mathbf{p}}^\mathrm{T} \\ \frac{\partial \mathbf{f}}{\partial \mathbf{x}} & \frac{\partial \mathbf{f}}{\partial \mathbf{p}} & 0 \end{bmatrix} \begin{bmatrix} \delta\mathbf{x} \\ \delta\mathbf{p} \\ \delta\lambda \end{bmatrix} = \begin{bmatrix} 0 \\ -\frac{\mathrm{d}T}{\mathrm{d}\mathbf{p}} \\ 0 \end{bmatrix} \tag{6}$$

for $\delta\mathbf{p}$. An appropriate update step is determined through line search, thus allowing us to efficiently converge to a solution.

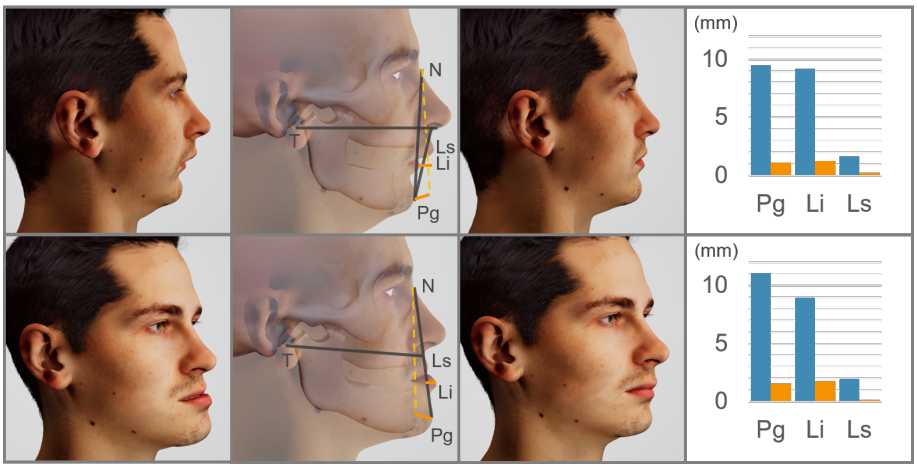

Figure 3: *Two examples presenting Angle's classes II (top) and III (bottom). The parameters of our surgical model can be effectively optimised, such that the selected landmarks after simulation are as close to their targets as possible. a) Input mesh. b) Sparse landmarks and targets for guiding the bone optimisation. c) Optimised face shape. d) Distance to the target landmarks is reduced on average from 6.3 to 1.6 mm.*

We obtain analytic expressions for the gradients $\frac{\partial \mathbf{f}}{\partial \mathbf{p}}$ and $\frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ through manual derivation and automatic differentiation [MGH*12]. We note that by the linearity of the derivative, the expressions for $\frac{\partial \mathbf{f}}{\partial \mathbf{p}}$ and $\frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ become sums of the Hessians of the Neo-Hookean and Potential energies. This allows us to split the computation into easily implementable chunks, and flexibly add more gradient information should additional energies be included into the simulation model.

First-order perturbation.

Additionally, starting from an equilibrium state $\mathbf{x}(\mathbf{p}_0)$, we can approximate the solution at $\mathbf{p}_1 = \mathbf{p}_0 + \mathrm{d}\mathbf{p}$, where $\mathrm{d}\mathbf{p}$ is a small step in the parameter space. We do this by explicitly calculating the sensitivity matrix, and linearly updating $\mathbf{x}(\mathbf{p}_0)$:

$$\begin{aligned} \mathbf{S} &= -\left(\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right)^{-1} \frac{\partial \mathbf{f}}{\partial \mathbf{p}} \\ \tilde{\mathbf{x}}(\mathbf{p}_1) &= \mathbf{x}(\mathbf{p}) + \mathbf{S} \cdot \mathrm{d}\mathbf{p}, \end{aligned} \tag{7}$$

where we denote $\tilde{\mathbf{x}}(\mathbf{p}_1)$ as the first-order approximate solution. This technique could be used in clinical practice, for example, to quickly view many estimates $\tilde{\mathbf{x}}(\mathbf{p}_1)$, and only perform the computationally intensive calculation of the true equilibrium $\mathbf{x}(\mathbf{p}_1)$ for the most promising predictions.

4. Results and Discussion

We test our method on two types of simulation meshes, presenting Angle's classes II and III [Rey10]. Our class III mesh also contains strong asymmetry. We design these meshes under the supervision of maxillofacial surgeons. With their aid we also identify the facial landmarks, and record their target positions using a process reflective of clinical practice.

Treatment parameter optimisation.

From these landmark targets, our optimiser identifies the best treatment parameters under the chosen simulation model. These parameters, which define the post-treatment bone shape and thus the soft-tissue surface, yield the minimum distance between landmarks and targets. The results of this process are shown in Figure 3. Additionally, Figure 4 displays an additional subject with differently presenting Class II and III malocclusion, where by setting the same treatment goal, we arrive at the same optimal bone position from vastly different initial geometries. Our experiments converge on average in just one or two optimisation steps. The average computation time is around 20 minutes on a consumer laptop with a 2.6GHz Intel i7-9750H CPU.

It is worth discussing the final non-zero distance-to-landmark error, and how it pertains to the desired behaviour of a parameter-recommendation system. We point out that it is trivial to relax the physical constraints of the simulation such that zero-error is reached, if one allows for wildly unrealistic deformations and treatment plans. The challenge lies in finding an appropriately "stiff" model which effectively converges to the set of all and only feasible treatments. In traditional forward orthognathic simulators, the

Figure 4: *The optimal treatment target (in yellow) is reliably identified starting from vastly different types of malocclusion using just seven landmarks. Malocclusions from left to right: severe class II, moderate class III, severe class III with asymmetry.*

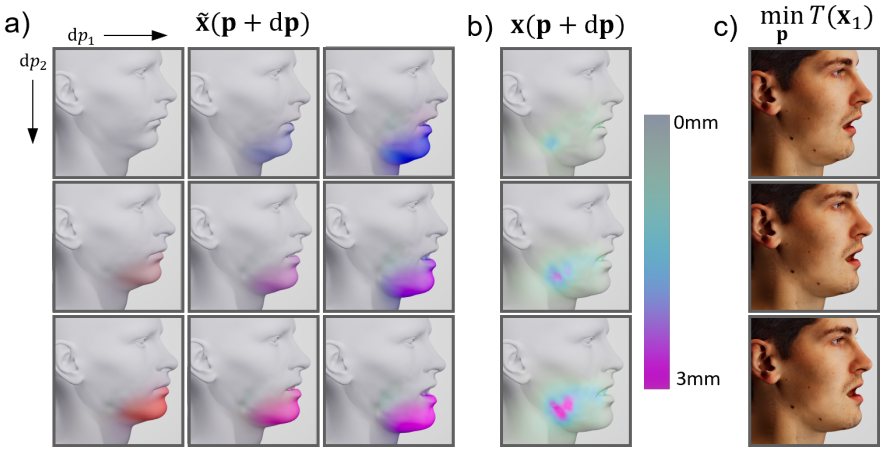

Figure 5: *a) Estimated equilibrium state for perturbed $d\boldsymbol{p}$, showing how increasing one element of $\boldsymbol{p}$ changes the estimate. We vary blendshapes responsible for the downwards and sideways translation of the mandible, and colour the surface displacement in each direction as blue or red respectively. b) True equilibrium state for rightmost column of a), recomputed using forward simulation. Colour shows the estimate's error. c) Optimising $\boldsymbol{p}$ to find a true equilibrium state as close as possible to the predicted $\tilde{\boldsymbol{x}}$. No error plot was generated, due to the optimal parameters being imperceptibly similar to those from b).*

surgeons manually constrain the simulated treatments into a "reasonable" range through their domain knowledge.

Conversely, a treatment model with insufficient degrees of freedom yields larger "optimal" errors than what could be achieved in practice. Analysis of larger sets of patient data may reveal deformation not effectively captured in our current blendshape model. Identifying which treatment models are most appropriate for inverse simulation is an unsolved problem.

Interactive first-order prediction.

The predictions of our first order model are visualised in Figure 5. We see that increasing the parameters responsible for opening and protruding the mandible, we obtain reasonable estimates in a fraction of the time it takes to find the equilibria via simulation. Computing the sensitivity matrix takes 5 seconds for our meshes. This matrix needs to be precomputed only a single time (for a given equilibrium), after which the first-order predictions can be calculated in real time. In contrast, calculating the true equilibrium takes an average of 32s for $|dp_i| = 0.05$ - which produces a rather modest deformation of 4mm in the simulation mesh, as measured at the affected surface vertices.

By comparing $\tilde{\mathbf{x}}(\mathbf{p}_1)$ to $\mathbf{x}(\mathbf{p}_1)$, we test the accuracy of our prediction. As seen in Figure 5b, an overall maximum error of 3mm is seen for a total facial displacement of 20 mm from $\mathbf{x}_0$, measured at the pogonion.

It is then possible to confidently explore predictions in this range, and at the accuracy boundary, recompute the true equilibrium point, or even run the optimisation to find the point in simulation space which is closest to the estimate (as shown in Figure 5c). This synergy lends itself to a rapid, interactive workflow, that could be used by clinicians and patients to easily visualise the dependence between treatment plan and outcome in a physically-based manner.

Future Work.

We have uncovered several unique challenges which arise when solving the inverse problem inherent to treatment planning. They allow us to highlight the knowledge gaps in this cross-domain application.

First, the priority for the VSP application is for the inverse problem to have a clinically viable optimal solution. This would involve further investigation into treatment models which robustly express all and only feasible treatments. Current forward simulation models rely on rigid body transforms, whose parameters are constrained by surgeon intuition to the manifold of "reasonable treatments". Meanwhile, our blendshape model has surgical prior knowledge built in. However, it may not necessarily represent all the surgically feasible transformations of a patients jaw. Additionally, extending it to new patients is non-trivial. The possibilities for treatment models are numerous, and data-driven approaches could be effective at identifying a valid treatment parametrisation across a wide population.

Other meaningful treatment variables which merit further investigation are related to tissue cutting, removal, and damage. Choosing the surface of the bone cut and how much bone to remove has important implications on the achievable post-treatment outcomes and the risks associated with damaging fragile structures like nerves and arteries. Though there is previous work in treatment planning with risk mitigation [HMP*12, CAR*09], extending this to a setting with inverse treatment-parameter optimisation is non-trivial. It would also be interesting to further automate the treatment planning process, for example by combining our approach with methods which suggest post-treatment targets [XDK*21] thereby alleviating the need for manually specifying treatment landmarks. Finally, although our work addresses the unique challenges that arise in a typical setting of maxillofacial surgery, our algorithmic framework straightforwardly generalizes to other osteotomies of the body, and extensions to further domains such as orthopaedic

or plastic surgery are feasible. The differentiable framework could also be applied not just to treatment planning, but conceivably to post-operative outcome validation and evaluation as well.

5. Conclusion

As simulation for surgical planning becomes more powerful, we must also focus on making it easier to use. Differentiable simulation can be effectively applied to inverse problems which arise in a variety of medical tasks. Here, we demonstrate its utility for orthognathic surgery planning, by showing how treatment parameters can be optimised from clinical measurements, how the equilibrium state can be quickly explored through a first-order gradient-informed approximation, and how the two systems can synergise with each other.

These prototypes pave the way towards high fidelity inverse simulation. There are many promising avenues for follow-up work. With a well constructed, diverse set of patient data, it will be possible to validate the forward simulation, enhance the treatment model, and aim towards clinical application. The benefits of fully-fledged differentiable simulation systems applied to clinical problems cannot be overstated. Therefore, we implore the research community to explore the topic of optimisation-based surgical planning, and look forward to seeing the innovation they deliver.

Acknowledgements

We thank our resident medical experts, Andrey Latyshev, Luís Pereira Azevedo and Magdalena Blankenburg for useful discussion and clinical knowledge transfer.

References

[AKFT15] ABE N., KURODA S., FURUTANI M., TANAKA E.: Data-based prediction of soft tissue changes after orthognathic surgery: clinical assessment of new simulation software. *Int J Oral Maxillofac Surg 44*, 1 (Jan 2015), 90–96. 2

[And15] ANDREWS L. F.: The 6-elements orthodontic philosophy: Treatment goals, classification, and rules for treating. *Am J Orthod Dentofacial Orthop 148*, 6 (Dec 2015), 883–887. 2

[APG*21] ALCAÑIZ P., PÉREZ J., GUTIÉRREZ A., BARREIRO H., VILLALOBOS A., MIRAUT D., ILLANA C., GUIÑALES J., OTADUY M. A.: Soft-Tissue Simulation for Computational Planning of Orthognathic Surgery. *J Pers Med 11*, 10 (Sep 2021). 3

[ARV07] AMBERG B., ROMDHANI S., VETTER T.: Optimal step nonrigid ICP algorithms for surface registration. In *CVPR* (2007), IEEE Computer Society. 9

[BBB*10] BEELER T., BICKEL B., BEARDSLEY P., SUMNER B., GROSS M.: High-quality single-shot capture of facial geometry. *ACM Trans. on Graphics (Proc. SIGGRAPH) 29*, 3 (2010), 40:1–40:9. 3, 9

[BH13] BUCHANAN E. P., HYMAN C. H.: LeFort I Osteotomy. *Semin Plast Surg 27*, 3 (Aug 2013), 149–154. 1, 2

[BHB*11] BEELER T., HAHN F., BRADLEY D., BICKEL B., BEARDSLEY P., GOTSMAN C., SUMNER R. W., GROSS M.: High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph. 30* (August 2011), 75:1–75:10. URL: http://doi.acm.org/10.1145/2010324.1964970. 3, 9

[CAR*09] CHENTANEZ N., ALTEROVITZ R., RITCHIE D., CHO L., HAUSER K. K., GOLDBERG K., SHEWCHUK J. R., O'BRIEN J. F.: Interactive simulation of surgical needle insertion and steering. *ACM Trans. Graph. 28*, 3 (jul 2009). URL: https://doi.org/10.1145/1531326.1531394, doi:10.1145/1531326.1531394. 6

[CLP03] CHABANAS M., LUBOZ V., PAYAN Y.: Patient specific finite element model of the face soft tissues for computer-assisted maxillofacial surgery. *Medical Image Analysis 7*, 2 (2003), 131–151. Mathematical Methods in Biomedical Image Analysis - MMBIA 2001. URL: https://www.sciencedirect.com/science/article/pii/S1361841502001081, doi:https://doi.org/10.1016/S1361-8415(02)00108-1. 2

[CMF*21] CHEN Z., MO S., FAN X., YOU Y., YE G., ZHOU N.: A meta-analysis and systematic review comparing the effectiveness of traditional and virtual surgical planning for orthognathic surgery: Based on randomized clinical trials. *J Oral Maxillofac Surg 79*, 2 (2021), 471.e1–471.e19. URL: https://www.sciencedirect.com/science/article/pii/S0278239120311113, doi:https://doi.org/10.1016/j.joms.2020.09.005. 1

[Epi] EPIC GAMES: Unreal engine. URL: https://www.unrealengine.com. 9

[Far94] FARKAS L.: *Anthropometry of the head and face*, second ed. Raven Press, 1994. 1

[GHZ*20] GEILINGER M., HAHN D., ZEHNDER J., BÄCHER M., THOMASZEWSKI B., COROS S.: Add: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph. 39*, 6 (nov 2020). URL: https://doi.org/10.1145/3414685.3417766, doi:10.1145/3414685.3417766. 3

[Gro98] GROSS M. H.: Computer graphics in medicine: From visualization to surgery simulation. *SIGGRAPH Comput. Graph. 32*, 1 (feb 1998), 53–56. URL: https://doi.org/10.1145/279389.279462, doi:10.1145/279389.279462. 1

[HBBC19] HAHN D., BANZET P., BERN J. M., COROS S.: Real2sim: Visco-elastic parameter estimation from dynamic motion. *ACM Trans. Graph. 38*, 6 (nov 2019). URL: https://doi.org/10.1145/3355089.3356548, doi:10.1145/3355089.3356548. 3

[HHD*21] HUANG Z., HU Y., DU T., ZHOU S., SU H., TENENBAUM J. B., GAN C.: Plasticinelab: A soft-body manipulation benchmark with differentiable physics. *CoRR abs/2104.03311* (2021). URL: https://arxiv.org/abs/2104.03311, arXiv:2104.03311. 1, 3

[HMN*21] HEIDEN E., MACKLIN M., NARANG Y. S., FOX D., GARG A., RAMOS F.: DiSECt: A Differentiable Simulation Engine for Autonomous Robotic Cutting. In *Proceedings of Robotics: Science and Systems* (Virtual, July 2021). doi:10.15607/RSS.2021.XVII.067. 3

[HMP*12] HERGHELEGIU P.-C., MANTA V.-I., PERIN R., BRUCKNER S., GRÖLLER E.: Biopsy Planner - Visual Analysis for Needle Pathway Planning in Deep Seated Brain Tumor Biopsy. *Computer Graphics Forum* (2012). doi:10.1111/j.1467-8659.2012.03101.x. 6

[HNGK19] HOSSEINZADEH NIK T., GHOLAMREZAEI E., KESHVAD M. A.: Facial asymmetry correction: From conventional orthognathic treatment to surgery-first approach. *J Dent Res Dent Clin Dent Prospects 13*, 4 (2019), 311–320. 2

[IKKP17] ICHIM A. E., KADLECEK P., KAVAN L., PAULY M.: Phace: physics-based face modeling and animation. *ACM Trans. Graph. 36*, 4 (2017), 153:1–153:14. URL: https://doi.org/10.1145/3072959.3073664, doi:10.1145/3072959.3073664. 1, 3

[KGC*96] KOCH R. M., GROSS M. H., CARLS F. R., VON BÜREN D. F., FANKHAUSER G., PARISH Y. I. H.: Simulating facial surgery using finite element models. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), SIGGRAPH '96, Association for Computing Machinery, p. 421–428. URL: https://doi.org/10.1145/237170.237281, doi:10.1145/237170.237281. 2

[KK19] KADLEČEK P., KAVAN L.: Building accurate physics-based face

models from data. *Proc. ACM Comput. Graph. Interact. Tech. 2*, 2 (jul 2019). URL: https://doi.org/10.1145/3340256, doi:10.1145/3340256. 1, 3

[KKR*19] KIM D., KUANG T., RODRIGUES Y. L., GATENO J., SHEN S. G. F., WANG X., DENG H., YUAN P., ALFI D. M., LIEBSCHNER M. A. K., XIA J. J.: A New Approach of Predicting Facial Changes following Orthognathic Surgery using Realistic Lip Sliding Effect. *MICCAI 11768* (Oct 2019), 336–344. 1, 3

[Lar17] LARSEN M. K.: Indications for orthognathic surgery - a review. *oral health and dental management 2017* (2017), 1–13. 2

[LKF*22] LAMPEN N., KIM D., FANG X., XU X., KUANG T., DENG H. H., BARBER J. C., GATENO J., XIA J., YAN P.: Deep learning for biomechanical modeling of facial tissue deformation in orthognathic surgical planning. *Int J Comput Assist Radiol Surg 17*, 5 (May 2022), 945–952. 2

[LL15] LIN H.-H., LO L.-J.: Three-dimensional computer-assisted surgical simulation and intraoperative navigation in orthognathic surgery: A literature review. *Journal of the Formosan Medical Association 114*, 4 (2015), 300–307. URL: https://www.sciencedirect.com/science/article/pii/S0929664615000686, doi:https://doi.org/10.1016/j.jfma.2015.01.017. 2

[LLK19] LIANG J., LIN M., KOLTUN V.: Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems* (2019), Wallach H., Larochelle H., Beygelzimer A., d'Alché-Buc F., Fox E., Garnett R., (Eds.), vol. 32, Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper/2019/file/28f0b864598a1291557bed248a998d4e-Paper.pdf. 3

[LYW*21] LEE S. J., YOO J. Y., WOO S. Y., YANG H. J., KIM J. E., HUH K. H., LEE S. S., HEO M. S., HWANG S. J., YI W. J.: A Complete Digital Workflow for Planning, Simulation, and Evaluation in Orthognathic Surgery. *J Clin Med 10*, 17 (Sep 2021). 1, 2

[MGH*12] MONAGAN M. B., GEDDES K. O., HEAL K. M., LABAHN G., VORKOETTER S. M., DEVITT J., HANSEN M., REDFERN D., RICKARD K., ET AL.: *Maple V Programming Guide: For Release 5*. Springer Science & Business Media, 2012. 5

[MTPS09] MEZGER J., THOMASZEWSKI B., PABST S., STRASSER W.: Interactive physically-based shape editing. *Computer Aided Geometric Design 26*, 6 (2009), 680–694. doi:https://doi.org/10.1016/j.cagd.2008.09.009. 3, 9

[OVT*08] OLSZEWSKI R., VILLAMIL M. B., TREVISAN D. G., NEDEL L. P., FREITAS C. M., REYCHLER H., MACQ B.: Towards an integrated system for planning and assisting maxillofacial orthognathic surgery. *Computer Methods and Programs in Biomedicine 91*, 1 (2008), 13–21. URL: https://www.sciencedirect.com/science/article/pii/S0169260708000485, doi:https://doi.org/10.1016/j.cmpb.2008.02.007. 3

[PP70] PECK H., PECK S.: A concept of facial esthetics. *Angle Orthod 40*, 4 (Oct 1970), 284–318. 2, 3, 4

[RE14] RAMIREZ E., ELENIUS J.: *Indications and Frequency of Orthognathic Surgery in Sweden – a Questionnaire Survey*. PhD thesis, 2014. URL: http://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-97848. 2

[Rey10] REYNEKE J.: *Essentials of Orthognathic Surgery*. Quintessence books. Quintessence Pub., 2010. URL: https://books.google.pl/books?id=TV7kTgEACAAJ. 2, 3, 4, 5

[RMV*20] RASMUSSEN C. M., MEYER P. J., VOLZ J. E., VAN ESS J. M., SALINAS T. J.: Facial Versus Skeletal Landmarks for Anterior-Posterior Diagnosis in Orthognathic Surgery and Orthodontics: Are They the Same? *J Oral Maxillofac Surg 78*, 2 (Feb 2020), 1–287. 2

[SB12] SIFAKIS E., BARBIC J.: Fem simulation of 3d deformable solids: A practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses* (New York, NY, USA, 2012), SIGGRAPH '12, Association for Computing Machinery. URL: https://doi.org/10.1145/2343483.2343501, doi:10.1145/2343483.2343501. 4

[SNF05] SIFAKIS E., NEVEROV I., FEDKIW R.: Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Trans. Graph. 24*, 3 (jul 2005), 417–425. URL: https://doi.org/10.1145/1073204.1073208, doi:10.1145/1073204.1073208. 3

[Ste60] STEINER C. C.: The use of cephalometrics as an aid to planning and assessing orthodontic treatment: Report of a case. *American Journal of Orthodontics 46*, 10 (1960), 721–735. URL: https://www.sciencedirect.com/science/article/pii/0002941660901457, doi:https://doi.org/10.1016/0002-9416(60)90145-7. 2

[SWR*21] SRINIVASAN S. G., WANG Q., ROJAS J., KLÁR G., KAVAN L., SIFAKIS E.: Learning active quasistatic physics-based models from data. *ACM Trans. Graph. 40*, 4 (jul 2021). URL: https://doi.org/10.1145/3450626.3459883, doi:10.1145/3450626.3459883. 1, 3

[TLTdL22] TUK J. G., LINDEBOOM J. A., TAN M. L., DE LANGE J.: Impact of orthognathic surgery on quality of life in patients with different dentofacial deformities: longitudinal study of the Oral Health Impact Profile (OHIP-14) with at least 1 year of follow-up. *Oral Maxillofac Surg 26*, 2 (Jun 2022), 281–289. 2

[VWWSGB21] VAUTRIN A., WESSELING M., WIRIX-SPEETJENS R., GOMEZ-BENITO M. J.: Time-dependent in silico modelling of orthognathic surgery to support the design of biodegradable bone plates. *Journal of the Mechanical Behavior of Biomedical Materials 121* (2021), 104641. URL: https://www.sciencedirect.com/science/article/pii/S1751616121003180, doi:https://doi.org/10.1016/j.jmbbm.2021.104641. 1, 3

[XDK*21] XIAO D., DENG H. H., KUANG T., MA L., LIU Q., CHEN X., LIAN C., LANG Y., KIM D., GATENO J., SHEN S. G., SHEN D., YAP P., XIA J. J.: A self-supervised deep framework for reference bony shape estimation in orthognathic surgical planning. *CoRR abs/2109.05191* (2021). URL: https://arxiv.org/abs/2109.05191, arXiv:2109.05191. 2, 6

[XY15] XU M., YANG J.: Human facial soft tissue thickness and mechanical properties: A literature review. vol. Volume 1A: 35th Computers and Information in Engineering Conference of *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. V01AT02A045. URL: https://doi.org/10.1115/DETC2015-46363. 3, 4

[YDX*14] YANPING L., DEDONG Y., XIAOJUN C., XUDONG W., GUOFANG S., CHENGTAO W.: Simulation and Evaluation of a Bone Sawing Procedure for Orthognathic Surgery Based on an Experimental Force Model. *Journal of Biomechanical Engineering 136*, 3 (02 2014). 034501. URL: https://doi.org/10.1115/1.4026104, arXiv:https://asmedigitalcollection.asme.org/biomechanical/article-pdf/136/3/034501/6089539/bio_136_03_034501.pdf, doi:10.1115/1.4026104. 1, 3

[YML*17] YUAN P., MAI H., LI J., HO D. C., LAI Y., LIU S., KIM D., XIONG Z., ALFI D. M., TEICHGRAEBER J. F., GATENO J., XIA J. J.: Design, development and clinical validation of computer-aided surgical simulation system for streamlined orthognathic surgical planning. *Int J Comput Assist Radiol Surg 12*, 12 (Dec 2017), 2129–2143. 2

[ZBBB18] ZOSS G., BRADLEY D., BÉRARD P., BEELER T.: An empirical rig for jaw animation. *ACM Trans. Graph. 37*, 4 (jul 2018). URL: https://doi.org/10.1145/3197517.3201382, doi:10.1145/3197517.3201382. 3

[ZCT22] ZEHNDER J., COROS S., THOMASZEWSKI B.: Sgn: Sparse gauss-newton for accelerated sensitivity analysis. *ACM Transactions on Graphics 41*, 1 (Feb 2022), 1–10. URL: http://dx.doi.org/10.1145/3470005, doi:10.1145/3470005. 1, 4

[ZKBT17] ZEHNDER J., KNOOP E., BÄCHER M., THOMASZEWSKI B.: Metasilicone: Design and fabrication of composite silicone with desired mechanical properties. *ACM Trans. Graph. 36*, 6 (nov 2017). URL: https://doi.org/10.1145/3130800.3130881, doi:10.1145/3130800.3130881. 1, 3

6. Appendix: Data Pipeline

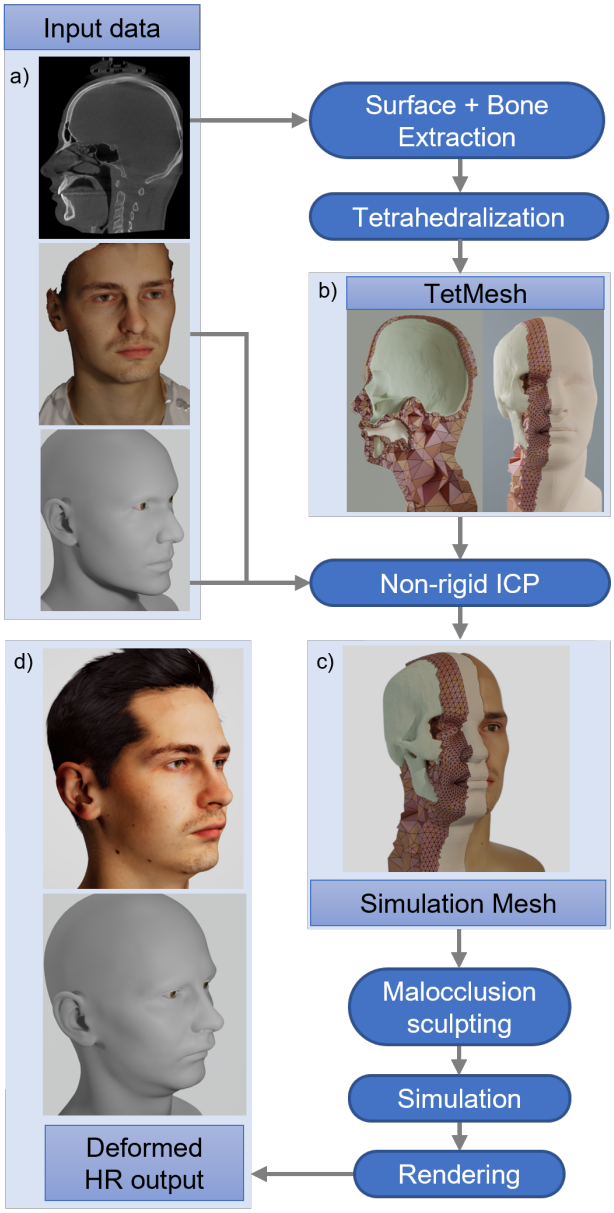

Figure 6: *We first collect a CBCT, 3D face scan [BBB*10, BHB*11], as well as a high-resolution template of a full head [Epi] (a). We extract the face and bone surfaces and create a tetrahedral mesh (b), which is registered with the face scan and template head [ARV07] (c). Supervised by a doctor we sculpt malocclusions to create an initial states from which we run the optimisation and simulation, before we render the deformed high-resolution head model [MTPS09, Epi] (d).*