ORIGINAL ARTICLE

Learning the shape of female breasts: an open-access 3D statistical shape model of the female breast built from 110 breast scans

Maximilian Weiherer[1,2] · Andreas Eigenberger[3,4] · Bernhard Egger[2] · Vanessa Brébant[4] · Lukas Prantl[4] · Christoph Palm[1,5]

Abstract

We present the *Regensburg Breast Shape Model* (RBSM)—a 3D statistical shape model of the female breast built from 110 breast scans acquired in a standing position, and the first publicly available. Together with the model, a fully automated, pairwise surface registration pipeline used to establish dense correspondence among 3D breast scans is introduced. Our method is computationally efficient and requires only four landmarks to guide the registration process. A major challenge when modeling female breasts from surface-only 3D breast scans is the non-separability of breast and thorax. In order to weaken the strong coupling between breast and surrounding areas, we propose to minimize the *variance* outside the breast region as much as possible. To achieve this goal, a novel concept called *breast probability masks* (BPMs) is introduced. A BPM assigns probabilities to each point of a 3D breast scan, telling how *likely* it is that a particular point belongs to the breast area. During registration, we use BPMs to align the template to the target as accurately as possible *inside* the breast region and only roughly outside. This simple yet effective strategy significantly reduces the unwanted variance outside the breast region, leading to better statistical shape models in which breast shapes are quite well decoupled from the thorax. The RBSM is thus able to produce a variety of different breast shapes as independently as possible from the shape of the thorax. Our systematic experimental evaluation reveals a generalization ability of 0.17 mm and a specificity of 2.8 mm. To underline the expressiveness of the proposed model, we finally demonstrate in two showcase applications how the RBSM can be used for surgical outcome simulation and the prediction of a missing breast from the remaining one. Our model is available at https://www.rbsm.re-mic.de/.

Keywords Statistical shape model · Non-rigid surface registration · Breast imaging · Surgical outcome simulation · Breast reconstruction surgery

✉ Christoph Palm
christoph.palm@oth-regensburg.de

1 Regensburg Medical Image Computing (ReMIC), Ostbayerische Technische Hochschule Regensburg (OTH Regensburg), Regensburg, Germany

2 Chair of Visual Computing, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Erlangen, Germany

3 Faculty of Mechanical Engineering, OTH Regensburg, Regensburg, Germany

4 University Center of Plastic, Aesthetic, Hand and Reconstructive Surgery, University Hospital Regensburg, Regensburg, Germany

5 Regensburg Center of Biomedical Engineering (RCBE), OTH Regensburg and Regensburg University, Regensburg, Germany

1 Introduction

Since the seminal work of Cootes et al. [14], *statistical shape models* became an emerging tool to capture natural shape variability within a given class of objects. As a result, a large number of shape models were developed during the last decades. The probably most well-known models were built for the human face using textured 3D face scans. Introduced by Blanz and Vetter [9], this class of statistical shape models is commonly known as *3D Morphable Models* (3DMMs) and includes models such as the *Basel Face Model* (BFM) and *Large-Scale Facial Model* (LSFM) presented by Paysan et al. [42] and Booth et al. [11], respectively. Well-studied applications of 3DMMs include face recognition, expression transfer between individuals, face animation, and 3D face reconstruction from a single 2D photograph [21]. Particu-

larly in the last decade, shape analysis also gained popularity in the field of computational anatomy, where statistical shape models are successfully used to model variations of anatomical objects such as bones and organs. Later, these models are utilized for a variety of medical applications including (but not limited to) image segmentation, surgical simulation, therapy planning, and motion analysis [3]. Despite the popularity of statistical shape models in the aforementioned areas, to date and to the best of our knowledge, no publicly available 3D statistical shape model of the female breast exists.

With breast cancer being the most common malignant neoplasm among women [52], successful *breast reconstruction surgery* (BRS) is crucial for patients undergoing mastectomy. In order to give patients a first impression about what their breast might look like after BRS, surgical outcomes are more and more often simulated using patient-specific 3D breast scans, acquired in a standing position (see Sect. 2 for an overview). Typically, simulations are performed using physically motivated deformable models of the breast. While these models take into account material properties and physical effects such as gravity, they may *not always* produce realistic-looking shapes as no prior knowledge in the form of example shapes is included [43,49]. Hence, simulated outcomes might be physically plausible, but definitely lack *statistical plausibility* in the sense that generated breast shapes are somehow *likely* or similar to those typically observed within a target population. As the ultimate goal of BRS is an outcome looking as *natural* as possible, we believe that simulation of surgical outcomes must not only rely on physically based deformable models but also should take into account statistical effects. Indeed, the phenomenon that humans tend to compare themselves with others clearly underlines the importance of the fact that simulated breast shapes should look similar to the breasts within the target population. As a first step toward combining physical *and* statistical plausibility of simulated breast surgery outcomes, we propose to use 3D statistical shape models built from natural-looking female breasts. In addition, by introducing such models into the breast shape domain, many of the aforementioned applications from other domains could be transferred to the breast as we will exemplary show later in this article.

To this end, this paper introduces the first publicly available 3D statistical shape model of the female breast built from 110 breast scans acquired in a standing position. Together with the model, we present a fully automated, pairwise registration pipeline especially tailored for 3D breast scans and its application in the context of statistical shape modeling. Our method is computationally efficient and requires only four landmarks to guide the registration process.
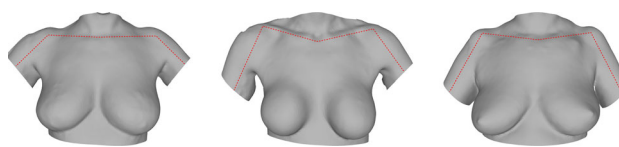
Fig. 1 Three typical 3D breast scans sampled from our database. Although a common pose was declared during data acquisition, a lot of pose variations are still present (indicated through a skeleton drawn in red). These mainly emerge from the arms and shoulders

1.1 Challenges

Compared to shape modeling of most parts of the human body, building a statistical shape model for the female breast imposes some new challenges as discussed in the following.

Data acquisition. Firstly, acquiring a sufficient amount of high-quality training data is challenging. Usually, 3D scanning and *manual* landmark detection is an uncomfortable situation for the participants, in which their upper body is required to be naked. Moreover, landmarks can be identified only through palpation and by using a regular tape measure, both requiring a physical examination in a clinical environment. In addition, during the *whole* examination, a specified posture needs to be held fixed ensuring a similar pose across all subjects. This can be very exhausting, especially if 3D breast scans are taken in a standing position in which both arms should ideally be held away from the body in order to capture the breast as isolated as possible. As a result, 3D scanning protocols used in clinical practice are often designed to be carried out relatively fast, thus lacking necessary precision for pose standardization, see Fig. 1. Note that the problem of only *quasi-similar* postures was also recently observed by Mazier et al. [38]. All in all, the aforementioned factors definitely hinder the implementation of large-scale, high-quality data surveys and might also explain why no publicly available data set of female 3D breast scans exits.

Correspondence estimation. Secondly, establishing dense correspondence among 3D breast scans by means of surface registration (rigid and non-rigid) is difficult due to the lack of reliable landmarks. In essence, only four valid landmarks can be used for non-rigid registration. These include both nipples as well as both lower breast poles (throughout this work defined as *the lowest (most caudal) point of the breast*, see Fig. 2). Anatomical landmark points such as the sternal notch or processus coracoideus cannot be used for non-rigid registration purposes as they would recover undesired, pose-dependent shape variations during statistical shape modeling. This way, the processus coracoideus disqualifies as its position depends strongly on the position of the arms and shoulders. On the other hand, the xiphoid, located at the center of the thorax, *could* technically be used for registra-

Fig. 2 A brief overview of common landmarks and key anatomical structures of the female breast and thorax (illustration adapted from [30])
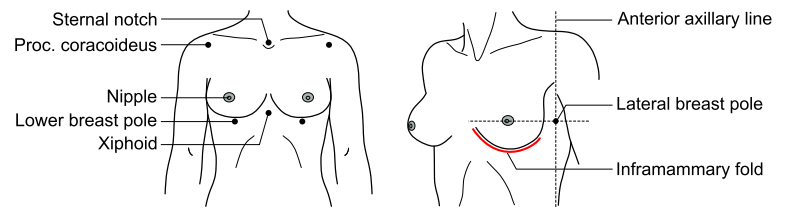
tion. However, it cannot be reliably and consistently located across a wide range of differently shaped breasts because of its dorsal tilt, effectively hindering the identification of a unique point. Another fundamental problem is the complete lack of robust landmarks on the lateral part of the breast. Although Hartmann et al. [28] defined a lateral breast pole as the orthogonal intersection between the anterior axillary line and a line passing through the nipple, this point cannot be used for non-rigid registration as its position is also affected by the pose of the arms and shoulders. Finally, already the initial, rigid alignment of 3D breast scans is challenging due to the lack of reliable landmarks that do not undergo large soft tissue deformations.

Non-separability of breast and thorax. Last but most importantly, the region of interest, i.e., the breast, obviously cannot be well separated from the rest of the thorax when considering 3D breast scans. This is primarily because the chest wall separating the breast from the thorax cannot be captured using 3D surface scanning devices, but also due to the reason that no commonly accepted and *exact* definition of the breast contour exists, see, e.g., [35]. A statistical shape model built from 3D breast scans will therefore *necessarily* capture also shape variations of the thorax, even after pose standardization. In particular, these include morphological shape variations of the underlying chest wall and upper abdomen, but also those emerging from the arms, armpits, and shoulders due to improper pose standardization (also seen from Fig. 1). If not reduced to a minimum, this will cause the following unwanted effect. Since breast and thorax shapes are tightly correlated in the subspace spanned by the model, the range of representable breast shapes is limited. The reason is that a considerably large part of the subspace accounts for the unwanted shape variations of the surrounding areas. Hence, the breast region should be decoupled from the thorax as much as possible in order to build an expressive and well-performing statistical shape model. Due the absence of an exact definition of the breast contour, we assume the breast region to *approximately* range from the lower breast pole to the second rib.

1.2 Contributions

As key contribution, this work presents the *Regensburg Breast Shape Model* (RBSM)—a 3D statistical shape model
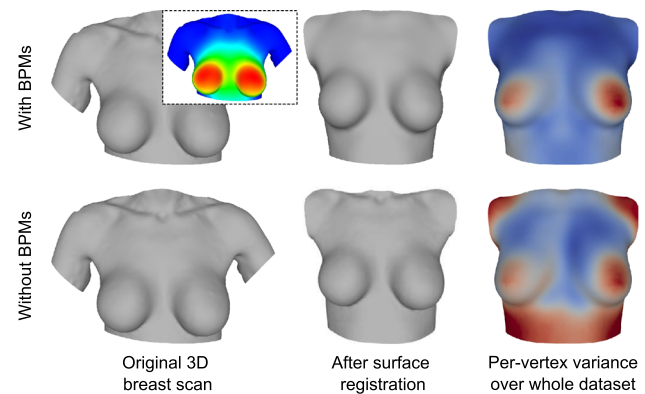

Fig. 3 The proposed concept of BPMs (top row) allows to minimize unwanted shape variations of the thorax by registering a template surface as precisely as possible *inside* the breast region and only roughly outside. The breast region is defined by BPMs in a probabilistic manner. (Top left, red areas correspond to a high probability of belonging to the breast region.) This simple yet effective strategy decouples the breast from the surrounding regions by reducing the variance outside the breast area. In the last column, the per-vertex variance over the whole dataset is visualized on the resulting mean shape. The regions showing the highest variance (red) are almost coincident with the breasts in our proposed BPM-based approach. Contrary, without BPMs (bottom row), a lot of unwanted variance is present in the surrounding regions, especially around the arms, armpits, shoulders, and upper abdomen. This implies a strong coupling between breast and thorax in the final statistical shape model

of the female breast. In order to weaken the strong coupling between the breast and surrounding regions, we propose to minimize the *variance* outside the breast region as much as possible. To achieve this goal, a novel concept called *breast probability masks* (BPMs) is introduced. A BPM assigns probabilities to each point of a 3D breast scan, telling how *likely* it is that a particular point belongs to the breast region. Later, during pairwise registration, we use the BPMs to align the template to the target as accurately as possible *inside* the breast region and only roughly outside. This way, only the most prominent and global shape variations outside the breast region will be recovered, effectively reducing the unwanted variance in these areas to a minimum. Figure 3 illustrates this idea.

To summarize, the contributions of this paper are threefold:

- We introduce the *Regensburg Breast Shape Model* (RBSM)—an open-access 3D statistical shape model of

the female breast built from 110 breast scans. It is available at https://www.rbsm.re-mic.de/.

- We propose a fully automated, pairwise registration pipeline used to establish dense correspondence among our 3D breast scans. It uses *breast probability masks* (BPMs) to decouple the breast from the surrounding regions as much as possible and requires only four landmarks to guide the registration process.
- We present two exemplary applications demonstrating how the RBSM can be used for surgical outcome simulation and the prediction of a missing breast from the remaining one.

The remainder of this paper is organized as follows: Sect. 2 briefly reviews some related work. Section 3 describes the entire model building pipeline used to construct the RBSM. In particular, it formally introduces the notion of a BPM, followed by a detailed description of the proposed registration pipeline. Section 4 presents an extensive evaluation of the RBSM in terms of the common metrics compactness, generalization, and specificity. Using the RBSM, two exemplary applications are showcased in Sect. 5. Finally, Sect. 6 discusses the results, whereas Sect. 7 concludes this article.

2 Related work

In this section, we briefly summarize some related work concerning statistical shape models of the female breast, breast surgery simulation, and popular techniques for pairwise surface registration used within the context of statistical shape modeling in general. Note that we have limited our review to the 3D case.

Statistical shape models of the female breast The literature about 3D statistical shape models of the female breast is sparse. In the early work of Seo et al. [45], a 3D statistical shape model was built from 28 breast scans with the goal of analyzing breast volume and surface measurements. However, they assume symmetric breasts obtained by simply mirroring the right breast and did not make their model publicly available. To date and to the best of our knowledge, this is the only work primarily addressing the construction of a statistical shape model from 3D breast scans, which is thus closest to our work. Besides, Ruiz et al. [44] utilized a 3D statistical shape model built from 310 breast scans for the validation of a novel weighted regularized projection method used for 3D reconstruction. As their focus did not lie on the construction of a well-performing statistical shape model of the breast, they did not provide detailed information about the registration method used to establish correspondence, the training data nor a comprehensive evaluation in terms of common metrics. Their model is also not publicly shared.

At last, few works exist attempting to construct a 3D statistical shape model from magnetic resonance imaging (MRI) data. Gallo et al. [23] applied *principal component analysis* to surface meshes extracted from 46 MRIs taken in prone position. Further, Gallego and Martel [22] developed a statistical shape model from 415 semi-automatically segmented breast MRIs for model-based breast segmentation.

Breast surgery simulation Most of the existing methods for preoperative breast surgery simulation are designed to simulate alloplastic, implant-based breast augmentation procedures, either for aesthetic reasons or after mastectomy as part of BRS. Typically, those methods first generate a patient-specific, geometric representation of the breast (using tetrahedral meshes, for instance). Afterward, the soft tissue deformation caused by implant insertion is simulated using a geometric and biomechanical model of the implant and breast, respectively.

As such, Roose et al. [43] used the tensor–mass model introduced by Cotin et al. [16], a combination of classical finite element and mass–spring models, for implant-based breast augmentation planning. De Heras Ciechomski et al. [19] proposed a web-based tool for breast augmentation planning which requires only 2D photographs and anthropometric measurements as input and allows the user to choose from a variety of different implants. Their method automatically reconstructs a 3D breast model and subsequently applies a tissue elastic model closely resembling the finite element model. Georgii et al. [24] utilize patient-specific finite element models generated from 3D breast scans. Combined with a novel mechanism called *displacement template*, geometric implant models are no longer required, thus breaking up the coupling between implant and enclosing breast.

Besides the simulation of breast augmentation procedures using implants, some methods exist especially addressing the simulation of BRS *without* implant insertion, for example, by means of autologous fat tissue. Based on Pascal's principle and volume conservation, Costa and Balaniuk [15] developed a novel approach for real-time physically based simulation of deformable objects, called *Long Elements Method* (LEM). An extension of LEM, known as the *Radial Elements Method* [7], was later used by Balaniuk et al. [6] for cosmetic and reconstructive breast surgery simulation. Williams et al. [55] employ a finite element approach incorporating the Mooney–Rivlin hyperelastic material model for the realistic simulation of soft tissue to simulate *transverse rectus abdominis myocutaneous* flap breast reconstruction.

Although the majority of the works on state-of-the-art breast surgery simulation utilize biomechanical models to predict the breast shape after surgery, an early attempt by Kim et al. [32] employs an example-based approach. Using a *sparse* set of 3D feature point pairs collected on 30 patients
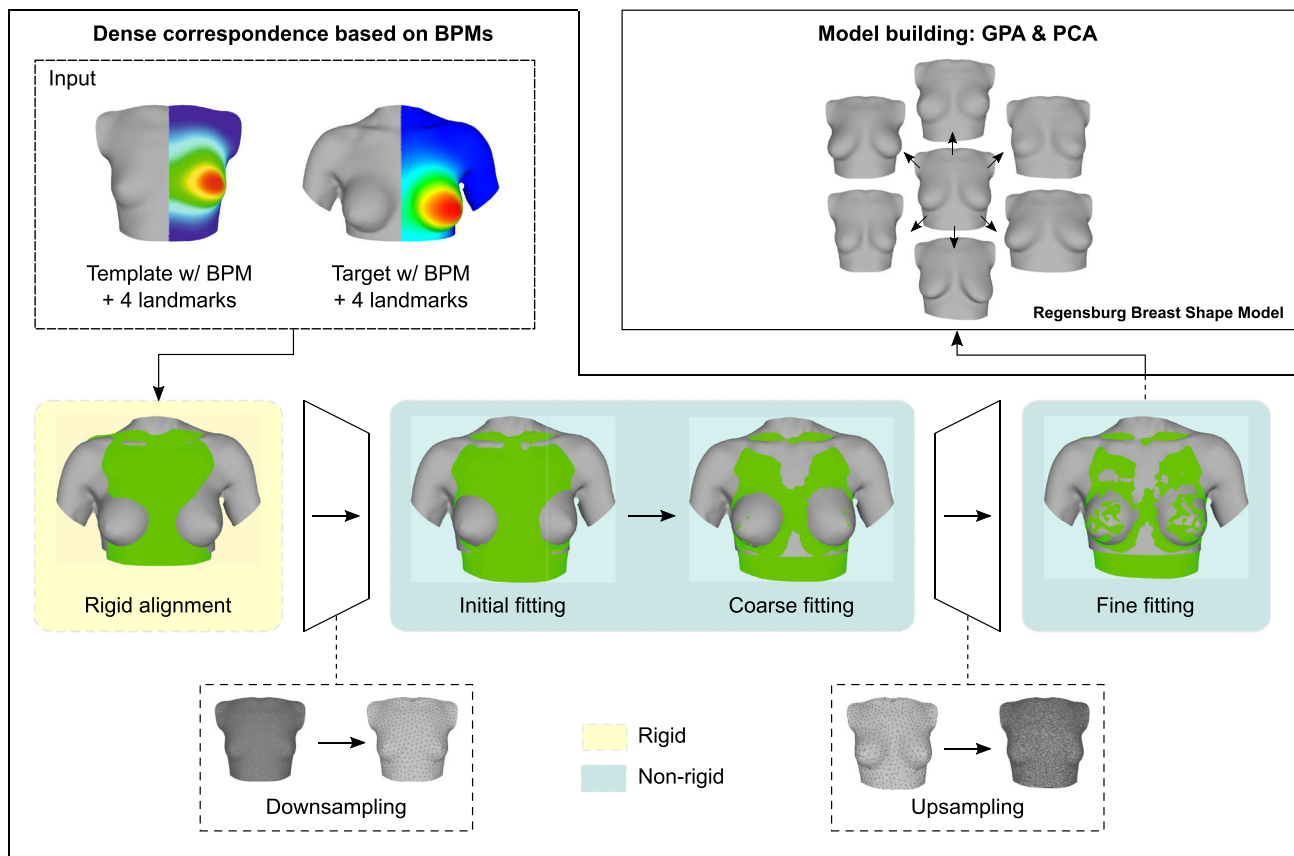
Fig. 4 An overview of the pipeline used to build the RBSM. We first establish dense correspondence by means of pairwise surface registration. Based on BPMs, our method starts by rigidly aligning the template to the target. Subsequently, a non-rigid alignment is applied in a hierar-chical, multi-resolution fashion. In this step, BPMs are used to precisely recover only shape variations of the breast. Finally, we perform classical *Generalized Procrustes Analysis* (GPA) and *principal component analysis* (PCA) to build the model

before and after surgery as training database, they developed a linear combination model to predict the surgical outcome.

For most of the existing methods, however, various authors highlighted the disagreement between simulated and actual outcomes [17,43] or only partially satisfying results for certain types of breast shapes [53].

Pairwise surface registration During the last few decades, countless algorithms were proposed tackling the problem of (pairwise) non-rigid surface registration. To date, however, none of them were used for 3D breast scan registration.

One of the most widely used methods is the *Optimal Step Non-rigid Iterative Closest Point* (NICP) framework proposed by Amberg et al. [4] and based on the work of Allen et al. [2]. NICP is the only method already used in the female breast shape domain to reconstruct a 3D breast model from a sequence of depth images [33]. Moreover, NICP was employed to construct the BFM and LSFM. A second class of non-rigid registration methods is based on splines, such as *Thin Plate Splines* (TPS) or B-splines. Pioneered by Bookstein [10], TPS were utilized by Paulsen et al. [41] to construct a statistical shape model of the human ear

canal. B-splines are extensively used in the *free-form deformations* framework and, among others, used for the creation of shape models of the human heart [39]. A third, recently introduced framework is based on *Gaussian Process Morphable Models* (GPMMs) introduced by Lüthi et al. [37]. GPMMs are statistical shape models themselves generalizing classical point-based models as proposed by Cootes et al. [14]. By means of GPMMs, expected deformations can be modeled using analytic covariance functions and later used as a *prior* for non-rigid surface registration, thus effectively reducing the search space. This approach was successfully used by Gerig et al. [25] to build an improved version of the BFM including facial expressions. Lastly, pairwise non-rigid surface registration algorithms are built upon the *as-rigid-as-possible* (ARAP) real-time mesh deformation framework proposed by Sorkine and Alexa [46]. This method was successfully transferred to the registration domain, yielding to similar methods constraining deformations to be *as conformal as possible* [59] or *as similar as possible* [31,57]. Recently, a variant of ARAP was utilized by Dai et al. [18] to build a shape model of the full human head.

3 Methodology

This section describes the entire pipeline used to build the RBSM. As outlined in Fig. 4, we start by establishing dense correspondence among our training data. Based on a novel concept called *breast probability masks* (Sect. 3.1), this is achieved by means of a fully automated, pairwise registration pipeline as proposed in Sect. 3.2. Finally, we follow the typical workflow used to build a point-based statistical shape model by applying *Generalized Procrustes Analysis* and *Principal Component Analysis* to the registered data set (briefly summarized in Sect. 3.3).

In what follows, 3D breast scans are represented using triangular surface meshes. A triangle mesh $\mathcal{M} = (V, E, \mathcal{P})$ is fully specified by a set of n vertices $V \subset \mathbb{N}$, edges $E \subset V \times V$, and an embedding $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2 \ldots, \mathbf{p}_n\} \subset \mathbb{R}^3$. Sometimes, however, instead of arranging points $\mathbf{p}_i$ into a set, it is more convenient to use a matrix notation $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n)^\top \in \mathbb{R}^{n \times 3}$. Hence, we will denote a triangle mesh either as $\mathcal{M} = (V, E, \mathcal{P})$ or equivalently as $\mathcal{M} = (V, E, \mathbf{P})$.

3.1 Breast probability masks

Given a 3D breast scan represented as triangle mesh $\mathcal{M} = (V, E, \mathcal{P})$, we call

$$p_{\mathcal{M}} : \mathcal{P} \longrightarrow (0, 1] \tag{1}$$

a *breast probability mask* (BPM). Technically, a BPM is a scalar field defined over $\mathcal{M}$ assigning each point $\mathbf{p}_i$ of a 3D breast scan a probability $p_{\mathcal{M}}(\mathbf{p}_i)$ telling how *likely* it is that $\mathbf{p}_i$ belongs to the breast region.

Concrete mapping As a concrete mapping for $p_{\mathcal{M}}$, we propose to use a normalized sum of *elliptical basis functions* (EBFs), centered at the nipples. We use EBFs instead of ordinary *radial basis functions* (RBFs) because we found that they better capture the natural teardrop shape of the breast (see Fig. 5 for a comparison between RBFs and EBFs). Technically, EBFs are a generalization of RBFs using the Mahalanobis distance instead of an ordinary vector norm. Formally, an EBF $\phi : [0, \infty) \longrightarrow \mathbb{R}$ centered at a point $\mathbf{c} \in \mathbb{R}^n$ is of the form $\phi(\mathbf{x}) = \phi(d_M(\mathbf{x}, \mathbf{c}))$. Here, d_M is the Mahalanobis distance, defined as

$$d_M(\mathbf{x}, \mathbf{c}) := \sqrt{(\mathbf{x} - \mathbf{c})^\top \mathbf{S}^{-1} (\mathbf{x} - \mathbf{c})}, \tag{2}$$

where $\mathbf{S} \in \mathbb{R}^{n \times n}$ is a symmetric positive definite matrix, also called *covariance matrix*. To stress that the Mahalanobis distance depends on $\mathbf{S}$, we write $d_M(\mathbf{x}, \mathbf{c}; \mathbf{S})$ in the following.

Now, in order to define a concrete BPM using EBFs, let $\mathbf{p}_N^\tau \in \mathcal{P}$ denote the position of the left (L) and right (R)

nipple, respectively, and $\tau \in \{L, R\}$. We first construct two individual probability masks for the left and the right breast, given as

$$p_{\mathcal{M}}^\tau(\mathbf{p}_i) = \phi\left(d_M\left(\mathbf{p}_i, \mathbf{p}_N^\tau; \mathbf{S}_\tau\right)\right). \tag{3}$$

Hereby, we define $\phi : [0, \infty) \longrightarrow (0, 1]$ as

$$\phi(x) = \exp\left(-x^2\right). \tag{4}$$

Finally, the BPM for a whole 3D breast scan is given as the normalized sum

$$\begin{aligned} p_{\mathcal{M}}(\mathbf{p}_i) = \frac{1}{4} \Big(& p_{\mathcal{M}}^L(\mathbf{p}_i) + \hat{p}_{\mathcal{M}}^L(\mathbf{p}_i) \\ & + p_{\mathcal{M}}^R(\mathbf{p}_i) + \hat{p}_{\mathcal{M}}^R(\mathbf{p}_i) \Big), \end{aligned} \tag{5}$$

where

$$\hat{p}_{\mathcal{M}}^\tau(\mathbf{p}_i) = \phi\left(d_M\left(\mathbf{p}_i, \hat{\mathbf{p}}_N^\tau; \hat{\mathbf{S}}_\tau\right)\right) \tag{6}$$

are shifted BPMs of the left and right breast added to better mimic the teardrop shape, and $\hat{\mathbf{p}}_N^\tau = \mathbf{p}_N^\tau + \mathbf{t}_\tau$ with translation vectors $\mathbf{t}_\tau \in \mathbb{R}^3$.

Parameter selection In order to fully define a BPM, appropriate matrices $\mathbf{S}_\tau, \hat{\mathbf{S}}_\tau \in \mathbb{R}^{3 \times 3}$ and translation vectors $\mathbf{t}_\tau \in \mathbb{R}^3$ need to be chosen first. As such, a total of 30 values are required to be properly determined (six for each $\mathbf{S}_\tau$ and $\hat{\mathbf{S}}_\tau$, and three for each $\mathbf{t}_\tau$). To simplify that task, we assume diagonal covariance matrices and utilize previously expert-marked landmarks on the 3D breast scans. Specifically, denote the landmark points shown in Fig. 2 as $\mathbf{p}_{SN}, \mathbf{p}_{XI} \in \mathcal{P}$ for sternal notch and xiphoid, and $\mathbf{p}_{LaBP}^\tau, \mathbf{p}_{LBP}^\tau \in \mathcal{P}$ for left and right lateral and lower breast pole, respectively. We then define

$$\begin{aligned} \mathbf{S}_\tau = \frac{1}{2} \operatorname{diag}\Big(& d_G\left(\mathbf{p}_{LaBP}^\tau, \mathbf{p}_N^\tau\right) + d_G\left(\mathbf{p}_N^\tau, \mathbf{p}_{XI}\right), \\ & d_G\left(\mathbf{p}_N^\tau, \mathbf{p}_{LBP}^\tau\right), d_G\left(\mathbf{p}_{LaBP}^\tau, \mathbf{p}_N^\tau\right) \Big), \\ \hat{\mathbf{S}}_\tau = \frac{1}{2} \operatorname{diag}\Big(& d_G\left(\mathbf{p}_{LaBP}^\tau, \mathbf{p}_N^\tau\right) + d_G\left(\mathbf{p}_N^\tau, \mathbf{p}_{XI}\right), \\ & d_G\left(\mathbf{p}_N^\tau, \mathbf{p}_{SN}\right), d_G\left(\mathbf{p}_{LaBP}^\tau, \mathbf{p}_N^\tau\right) \Big), \\ \mathbf{t}_\tau = \mathbf{p}_N^\tau + & \frac{1}{5}\left(0, d_G\left(\mathbf{p}_N^\tau, \mathbf{p}_{SN}\right), 0\right), \end{aligned} \tag{7}$$

where d_G denotes the Geodesic distance between two points on the surface mesh. Note that $\mathbf{S}_\tau$ and $\hat{\mathbf{S}}_\tau$ differ only in the second diagonal element.

3.2 Registration of 3D breast scans

Following Fig. 4, the proposed pairwise registration pipeline is mainly composed of rigid alignment (Sect. 3.2.1) and non-

Fig. 5 From left to right: comparison between RBF, EBF, and a sum of two EBFs, illustrated as contour plots. While a simple RBF or EBF is not able to accurately mimic the typical teardrop shape of the breast, a sum of two EBFs comes close
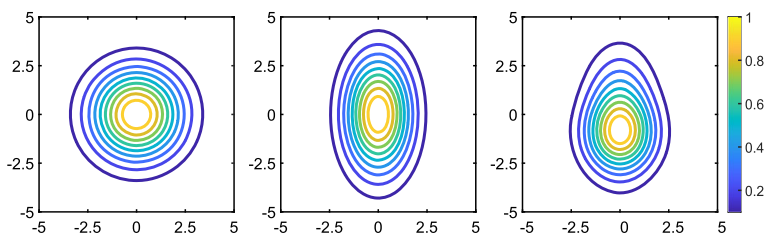
rigid alignment (Sect. 3.2.2). To speed up convergence, the latter is carried out in a hierarchical, multi-resolution fashion (Sect. 3.2.3).

Both phases make extensive use of BPMs in order to align a template surface $S = (V, E, \mathbf{P})$ to a target T as accurately as possible *inside* the breast region and only roughly outside, effectively decoupling the breast from the rest of the thorax by reducing the *variance* outside the breast region to a minimum. This is justified as the covariance $cov(x, y)$ becomes smaller if $var(x)$ or $var(y)$ is lowered, following from the well-known fact that $|cov(x, y)| \leq \sqrt{var(x)}\sqrt{var(y)}$ (which holds via the Cauchy–Schwarz inequality).

Finally, note that the target surface T can be given in any representation that allows for closest point search. We use a triangular surface mesh but write $T \subset \mathbb{R}^3$ for the sake of notational simplicity.

3.2.1 Rigid alignment

The overall goal of the rigid alignment is to move the template as close as possible to the *rigid part* of the target, which we define as *the thorax without the breast*. In particular, we expect that the thoraxes of two subjects without the breast region can be sufficiently well aligned if we assume the breast to be the only part of the thorax that deforms non-rigidly. Based on this assumption, the absence of suitable landmarks, and due to the fact that our initial 3D breast scans are already reasonably well aligned (see Sect. 4.1), we propose a modified version of the *Iterative Closest Point* (ICP) algorithm, originally introduced by Besl and McKay [8].

Essentially, compared to the standard version of the ICP algorithm, our modified version differs in the following three aspects: (i) A scaling factor is added to the rigid transformation effectively allowing for Euclidean similarity transformations [20,60]. Secondly, (ii) to ensure that only the rigid parts of the 3D breast scans are used for alignment, correspondences, where both points have a high probability belonging to the breast region, are discarded. This is implemented by thresholding the template and target BPMs. Finally, (iii) rotations are restricted to the x-axis corresponding to the transversal plane. Rotations around the y- and z-axis (sagittal and coronal plane) possibly introduced due to severe overweight in conjunction with an uneven distribution of abdominal fat could destroy the initial alignment and

lead to misalignment. In any case, asymmetries introduced due to the thorax should *not* affect the rigid alignment of the template.

3.2.2 Non-rigid alignment

Given the rigidly aligned template $S = (V, E, \mathbf{P})$, the goal of the non-rigid alignment is to gradually deform S into a new surface $S' = (V, E, \mathbf{P}')$ with identical topology such that S' is as close as possible to the target T *inside* the breast region. Following various authors including Jiang et al. [31] and Yamazaki et al. [57], we formulate our non-rigid registration problem using the following nonlinear energy functional

$$F(\mathbf{P}') = F_D(\mathbf{P}') + \alpha F_R(\mathbf{P}') + \beta F_L(\mathbf{P}'), \qquad (8)$$

where F_D is a distance term used to penalize the point-to-point distance between the template and target surface, F_R is a regularization term constraining deformations *as similar as possible*, and F_L constitutes a landmark term ensuring certain points to be matched. $\alpha, \beta \geq 0$ are weights controlling the individual contribution of each term to the cost function. Minimizing F finally leads to the new points $\mathbf{P}'$ of the deformed template surface S', i.e.,

$$\mathbf{P}' = \arg\min_{\mathbf{P}' \in \mathbb{R}^{n \times 3}} F(\mathbf{P}'). \qquad (9)$$

Adapting the strategy proposed by Allen et al. [2], instead of computing (9) only once, we minimize F several times but each time lowering the regularization weight α in (8). As later demonstrated by Amberg et al. [4], this strategy is able to recover the whole range of global and local non-rigid deformations efficiently. Following various authors [31,46], the optimization problem in (9) is solved using an alternating minimization (AM) approach as briefly summarized in Appendix A.

Distance term The distance term F_D is used to attract the template S to the target T. Assuming fixed correspondences between both surfaces, i.e., $\{(\mathbf{p}_1, \mathbf{q}_1), (\mathbf{p}_2, \mathbf{q}_2), \ldots, (\mathbf{p}_n, \mathbf{q}_n)\}$ with $\mathbf{q}_i \in T$ being the closest point to $\mathbf{p}_i$, the distance term can be written as

$$F_D(\mathbf{P}') = \frac{1}{2} \left\| \mathbf{C}(\mathbf{P}' - \mathbf{Q}) \right\|_F^2, \qquad (10)$$

where $\mathbf{C} := \operatorname{diag}(c_1, c_2, \ldots, c_n)$, $c_i \geq 0$ for all $i \in \{1, 2, \ldots, n\}$ are weights used to quantify the reliability of a match, and $\mathbf{Q} := (\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n)^\top \in \mathbb{R}^{n \times 3}$. Using the BPMs $p_{\mathcal{S}}$ and $p_{\mathcal{T}}$ of the template and target, we set

$$c_i = \frac{p_{\mathcal{S}}(\mathbf{p}_i) + p_{\mathcal{T}}(\mathbf{q}_i)}{2}. \tag{11}$$

This way, correspondences $(\mathbf{p}_i, \mathbf{q}_i)$ mapping from one breast region to the other have a greater impact on the overall distance term as $c_i \in (0, 1]$ becomes large in this case. Conversely, the influence tends to zero if $c_i \to 0$, i.e., if both points are less likely to belong to the breast region. As such, the deformation of points $\mathbf{p}_i$ on the template with a small value for c_i is mainly controlled by the regularization term, as previously described by Allen et al. [2].

Regularization term The regularization term F_R should prevent the template surface from shearing and distortion while simultaneously ensuring structure preservation and smooth deformations. To do so, we adapt the *consistent as-similar-as-possible* (CASAP) regularization technique in which deformations are constrained to be *locally* as similar as possible [31,57]. Specifically, given a local neighborhood $E_i \subset E$ around each point $\mathbf{p}_i$, the template surface is only allowed to move in terms of an Euclidean similarity transformation

$$\mathbf{p}'_j - \mathbf{p}'_k = s_i \mathbf{R}_i (\mathbf{p}_j - \mathbf{p}_k) \qquad \forall (j, k) \in E_i, \tag{12}$$

where $s_i > 0$ is a scaling factor and $\mathbf{R}_i \in SO(3)$ a rotation matrix. Following Chao et al. [13], we define E_i as *the set containing all (directed) edges of triangles incident to* $\mathbf{p}_i$, also known as *spokes-and-rims*. Finally, our CASAP regularization term reads

$$F_R(\mathbf{P}') = \frac{1}{2} \sum_{i=1}^{n} w_i$$
$$\left[\sum_{(j,k) \in E_i} w_{jk} \left\| \left(\mathbf{p}'_j - \mathbf{p}'_k \right) - s_i \mathbf{R}_i (\mathbf{p}_j - \mathbf{p}_k) \right\|_2^2 + \lambda \sum_{l \in N_i} w_{il} \| \mathbf{R}_i - \mathbf{R}_l \|_F^2 \right], \tag{13}$$

where weights $w_i > 0$ are added to individually control the amount of regularization for each particular point. As mentioned above, since the deformation of points $\mathbf{p}_i$ with a small value for c_i is mainly controlled by the regularization term, we define

$$w_i = \frac{1}{(h-1)c_i + 1} \quad \text{with} \quad \frac{1}{h} \leq w_i < 1 \tag{14}$$

for all $i \in \{1, 2, \ldots, n\}$ and $h \in \mathbb{N}^+$ (we used $h = 2$ throughout this paper). As seen, this strategy keeps points $\mathbf{p}_i$ of the template relatively stiff if (i) $\mathbf{p}_i$ has a low probability belonging to the breast region and (ii) if the corresponding point on the target is also not likely to be part of the breast region (because $w_i \to 1$ if $c_i \to 0$), thus effectively preventing the template from adapting too close to the target outside the breast region. Lastly, $N_i \subset V$ in (13) denotes the one-ring neighborhood of the i-th point and $w_{jk} \in \mathbb{R}$ are the popular cotangent weights, see, e.g., [12]. $\lambda \geq 0$ is usually set to $0.02A$, where $A \geq 0$ is the total surface area of $\mathcal{S}$ [34].

Landmark term The goal of the landmark term F_L is to keep certain positions (i.e., landmarks) fix during the registration process. Let $I \subset \mathbb{N}$ be an index set containing the indices of the m landmarks specified on the template surface $\mathcal{S}$. Define a matrix $\mathbf{D} \in \mathbb{R}^{m \times n}$ as

$$\mathbf{D} = (d_{ij}) := \begin{cases} 1, & \text{if } j \in I, \\ 0, & \text{otherwise} \end{cases} \tag{15}$$

for $i = 1, 2, \ldots, m$ and $j = 1, 2, \ldots, n$. Next, denote the corresponding landmarks on the target surface by $\{\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_m\} \subset \mathcal{T}$. Then, the landmark term is defined as

$$F_L(\mathbf{P}') = \frac{1}{2} \left\| \mathbf{D}\mathbf{P}' - \mathbf{Q}_L \right\|_F^2, \tag{16}$$

where $\mathbf{Q}_L := (\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_m)^\top \in \mathbb{R}^{m \times 3}$.

3.2.3 Multi-resolution fitting strategy

Following common practice and to speed up convergence, instead of applying the previously described non-rigid alignment only once, we employ a hierarchical, multi-resolution fitting strategy composed of initial fitting, coarse fitting, and fine fitting (see also Fig. 4).

Initial fitting Having a low-resolution instance of the rigidly aligned template at hand, the goal of the initial fitting is to roughly adapt the coarse template to the key features (i.e., landmarks) of the target. To do so, we strictly prioritize the landmark constraints and do not use BPMs in this phase.

Coarse fitting In this step, the initially fitted low-resolution template is gradually deformed toward the target.

Upsampling Next, the deformation obtained from the previous step is applied to the original, full-resolution template. This is achieved using a concept called *Embedded Deformation*, introduced by Sumner et al. [50]. In essence, the deformation of the coarse template obtained from the previ-

ous step is transferred to the template by linearly interpolating the transformation at each point.

Fine fitting Lastly, the upsampled template is fitted to the target again to produce the final result.

3.3 Model building

Once the data set is brought into correspondence, we follow the typical workflow used to build a classical point-based statistical shape model as proposed by Cootes et al. [14]. For notational simplicity, instead of stacking points $\mathcal{P}$ of a triangular mesh $\mathcal{M} = (V, E, \mathcal{P})$ into a matrix $\mathbf{P} \in \mathbb{R}^{n \times 3}$ as before, we use a vectorized representation, denoted as $\mathbf{x} = \text{vec}(\mathbf{P}) \in \mathbb{R}^{3n}$ in the following.

Briefly, given a set of k breast scans $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k\} \subset \mathbb{R}^{3n}$ in correspondence, we first perform *Generalized Procrustes Analysis* (GPA) as introduced by Gower [26]. GPA iteratively aligns the objects to the arithmetic mean $\bar{\mathbf{x}} \in \mathbb{R}^{3n}$ (successively estimated from the data) by using an Euclidean similarity transformation, effectively transforming the objects into the shape space. Secondly, *principal component analysis* (PCA) is carried out on the Procrustes-aligned shapes. Let $\{\lambda_1, \lambda_2, \ldots, \lambda_q\} \subset \mathbb{R}^+$ be the $q < k$ nonzero eigenvalues (also called *principal components* (PCs) in this context) of the empirical covariance matrix calculated from the data and sorted in a descending order. Denote the corresponding eigenvectors as $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_q\} \subset \mathbb{R}^{3n}$. Then, a statistical shape model can be interpreted as a linear function $M : \mathbb{R}^q \longrightarrow \mathbb{R}^{3n}$ defined as

$$M(\boldsymbol{\alpha}) = \bar{\mathbf{x}} + \mathbf{U}\boldsymbol{\alpha}, \tag{17}$$

where $\mathbf{U} := (\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_q) \in \mathbb{R}^{3n \times q}$. To ensure plausibility of the newly generated shapes, α_i is usually restricted to $|\alpha_i| \leq 3\sqrt{\lambda_i}$ for all $i \in \{1, 2, \ldots, q\}$. If a (possibly unseen) shape $\mathbf{x}' \in \mathbb{R}^{3n}$ is in correspondence with the model and properly aligned, it can be reconstructed from M in a least-squares sense by using

$$\boldsymbol{\alpha}^* = \mathbf{U}^{-1}\left(\mathbf{x}' - \bar{\mathbf{x}}\right) \tag{18}$$

as the model parameters, i.e., $\mathbf{x}' \approx M(\boldsymbol{\alpha}^*)$. The number $q < k$ of retained PCs is chosen so that the model typically represents a fixed proportion of the total variance, e.g., 98%.

4 Evaluation

Based on our 3D breast scan database which we present in Sect. 4.1, several experiments were conducted in order to evaluate the proposed statistical shape model (Sect. 4.2).

Table 1 An overview of the 3D breast scan database used to build the RBSM

	Mean ($\pm$ SD)	Range [min–max]
Age [years]	40.78 ($\pm$ 14.30)	18.00–83.00
BMI [kg/m^2]	23.66 ($\pm$ 3.57)	16.90–38.27
Breast volume [cc]		
Left	477.23 ($\pm$ 242.94)	70.60–1,258.90
Right	481.45 ($\pm$ 240.27)	80.00–1,609.30
Difference	80.12 ($\pm$ 81.34)	0.30–367.40

It includes 110 textured 3D breast scans

4.1 Data

Our database consists of 110 textured 3D breast scans collected at our institution (St. Josef Hospital Regensburg) using the portable Vectra H2$^{\text{TM}}$ scanning system (Canfield Scientific, New Jersey, USA). The H2 system employs photogrammetry to reconstruct a 3D surface mesh from a series of 2D images with a resolution of 3.5 mm and a maximal capture volume of $70 \times 41 \times 40$ cm^3 (width $\times$ height $\times$ depth). In our case, in total three photographs were taken of each subject: one from a frontal view and two from a lateral view ($\pm 45°$). Note that due to the standardized distance and angles from which the photographs are taken, the reconstructed 3D breast scans are already reasonably well aligned and consistently oriented.

A previously developed standardized scanning protocol [28] was used to ensure a common pose and same imaging conditions for all participants. In brief, all subjects were asked to stand in an upright posture and abduct both arms by an angle of $45°$. (A telescopic stick was used to support the subjects in holding their arms fixed at the specified angle.) This posture produces a natural-looking breast shape and allows to capture the whole breast as isolated as possible. Moreover, it is fast and easy to adopt for all participants regardless of age, body weight, or medical history.

A compact overview of some key parameters in our database is given in Table 1. In addition, 62 out of 110 participants have at least one ptotic breast. Of these, 15 women show different ptosis grades on the left and right breast. 72 out of 110 participants received some kind of breast surgery such as alloplastic or autologous breast reconstruction or augmentation.

4.2 Experiments and results

Following common practice, we evaluate our statistical shape model using the well-known metrics compactness, generalization, and specificity as proposed by Styner et al. [48]. Registration results are assessed by the classical distance-based
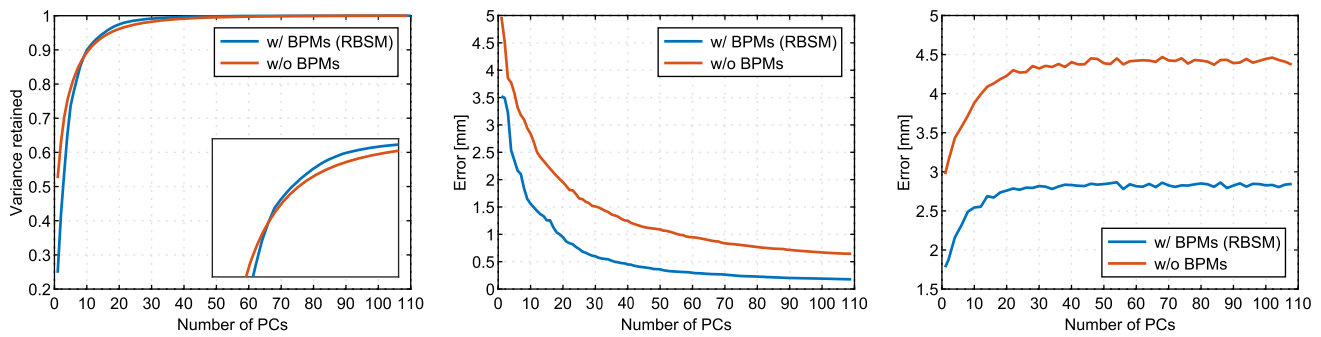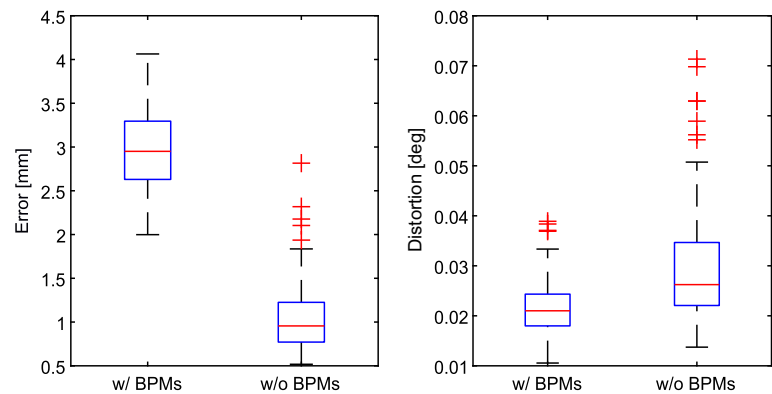
Fig. 6 From left to right: compactness, generalization, and specificity of statistical shape models built from 3D breast scans registered either with (i.e., the RBSM) or without using BPMs during registration

Fig. 7 MSE (left) and angle distortion (right), calculated from 3D breast scans registered with and without BPMs

mean squared error (MSE) between deformed template and target. In addition, the *angle distortion* between the deformed template $\mathcal{S}'$ and the original template $\mathcal{S}$ is determined in order to quantify the amount of shearing introduced due to non-rigid registration. This is achieved by simply averaging the absolute deviation between the inner-triangle angles of $\mathcal{S}$ and $\mathcal{S}'$ over all triangles [54].

For all experiments, the same template $\mathcal{S}$ was used during registration. In particular, the 3D breast scan of a healthy, non-operated subject was chosen and subsequently mirrored along the sagittal plane to produce a perfectly symmetrical template. After isotropic remeshing and Laplacian smoothing, a regular surface mesh with an average edge length of 2.05 mm was obtained. It consists of 30,924 vertices. The coarse, low-resolution version of our template was created using mesh simplification techniques and includes 646 vertices. All relevant parameters used for registration are listed in Appendix B.

The complete registration pipeline was implemented in C++. Statistical shape modeling was performed with an open-source framework called Scalismo (https://scalismo.org/), implemented in Scala and based on Statismo [36].

4.2.1 On the effect of BPMs

The first experiment should evaluate our novel BPM-based registration technique, its influence on the resulting statistical shape models, and, in particular, whether or not BPMs are able to weaken the strong coupling between the breast and surrounding regions. As such, we compare two different models: The first one is constructed from 3D breast scans brought into correspondence using the proposed approach based on BPMs (i.e., the RBSM). The second model is built from 3D breast scans registered without using BPMs.

Figure 6 shows the evaluation metrics of the resulting statistical shape models. As seen, although the model built without BPMs is more compact than with BPMs (i.e., the RBSM) when considering only the very first PCs, it is also the model that generalizes worse. Specifically, when using all 109 available PCs a generalization error of 0.65 mm is reported. For comparison, the RBSM shows a generalization error of only 0.17 mm when using the same number of PCs. Finally, the RBSM constantly achieves the lower specificity error of about 2.8 mm if more than 30 PCs are used.

Figure 7 summarizes the registration results, quantified in terms of MSE and angle distortion. As expected, registration without using BPMs clearly achieves the lower MSE of
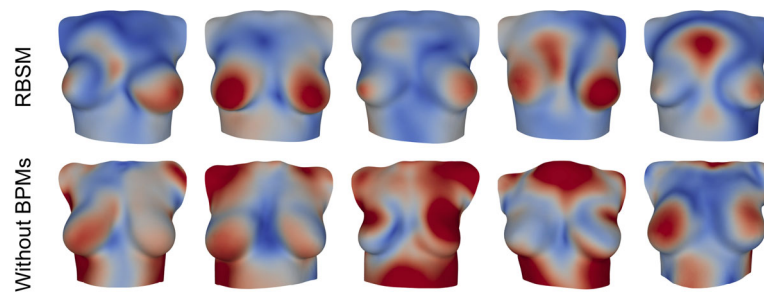
Fig. 8 Comparison between five random samples from the RBSM (top row, generated from 3D breast scans registered using BPMs) and five random samples from a model constructed without using BPMs (bottom row). The samples are colored according to the distance from the mean $\bar{\mathbf{x}}$. As such, the red areas underwent a high deformation, whereas blue regions are rather stiff. For samples drawn from the RBSM, the regions showing the highest variation are almost always the breast region, indicating a solid decoupling between breast shape and thorax
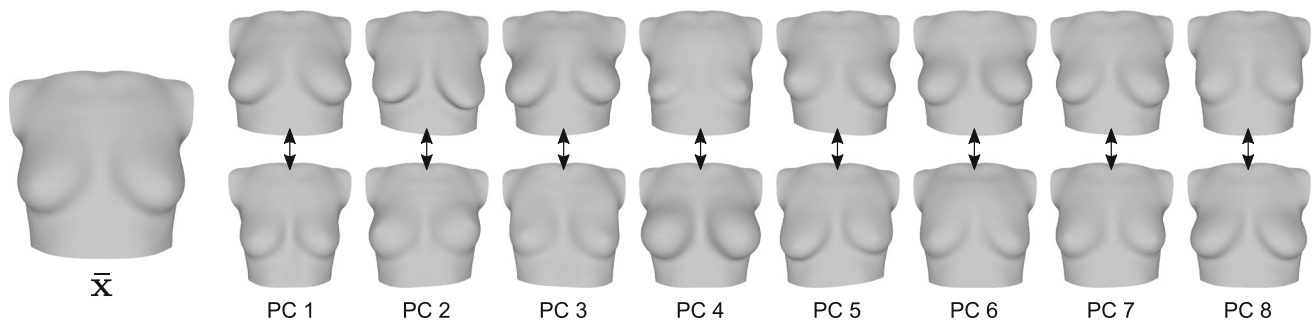
Fig. 9 The first eight principal modes of variation from the RBSM, visualized by either adding (top row) or subtracting (bottom row) $3\sqrt{\lambda_i}\mathbf{u}_i$ from the mean $\bar{\mathbf{x}}$, displayed on the left. Together, they represent about 85% of the total variance present in the dataset

1.05 mm as the goal here is to register the template to the target as close as possible. For comparison, the MSE achieved when using BPMs is 2.98 mm. However, it is important to remember that the goal of the BPM-based registration is to align the template as precisely as possible to the target only *inside* the breast region. Therefore, the overall MSE between template and target might not reflect the actual registration quality well (see, e.g., Fig. 3). In terms of angle distortion, only a small difference is noticeable between registration with and without BPMs, respectively. In particular, registration without BPMs caused the angle distortion to increase by only 0.008 degrees on average.

In order to further investigate to which extend BPM-based registration is able to weaken the strong coupling between breast region and thorax, Fig. 8 shows some random samples from the RBSM and the model built from 3D breast scans registered without BPMs. The samples are colored according to their distance to the mean shape, effectively providing a *measure of variation*. Interestingly, the areas of the highest variation in the samples drawn from the RBSM are almost always located at the breasts. The surrounding regions are rather stiff, indicating low variance and a quite well decoupling between breast and thorax. Hence, the RBSM is able to produce a variety of breast shapes without altering the whole thorax too much, also reflected in the principal modes

of variation shown in Fig. 9. Contrary, the area of the highest variation in the samples generated from the model built without BPMs is oftentimes *not* the breast region.

4.2.2 How much data are needed?

An often arising question in the context of statistical shape modeling concerns the amount of training data needed to build a reasonably well-performing model. Generally, this question very much depends on the amount of variability samples from a particular class of objects are expected to show. As a rule of thumb, a good training set should always reflect the whole bandwidth of possible variations likely to occur within a target population. The goal of this experiment is to test how much data is needed to build a well-performing statistical shape model of the female breast. To do so, we randomly sample 30, 60, and 90 breast scans from our database and subsequently compare the resulting models with the model containing all 110 breast scans (i.e., the RBSM). To avoid sampling bias, the whole procedure was repeated three times, and results were averaged.

Figure 10 shows the results in terms of compactness, generalization, and specificity. Regarding compactness, it can be clearly seen that the model built with 30 breast scans is the most compact one, followed by the models constructed
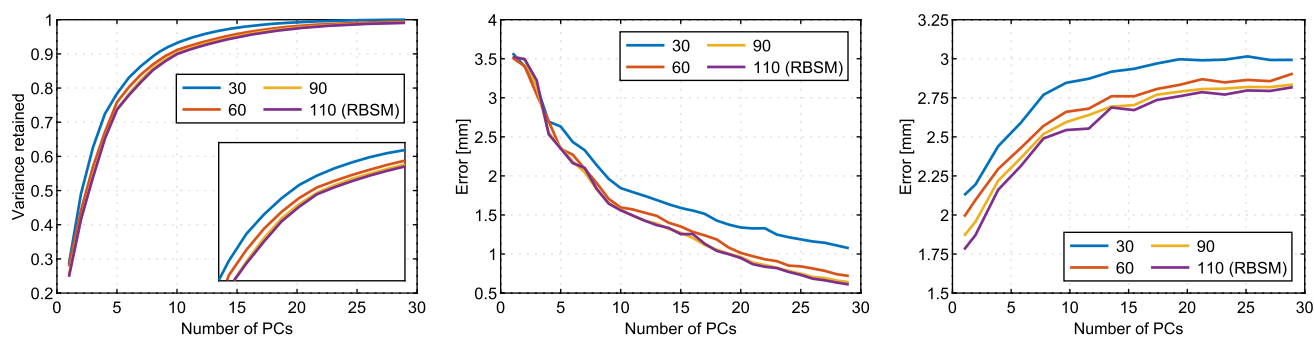
Fig. 10 From left to right: compactness, generalization, and specificity of statistical shape models built from randomly sampled subsets of the training data
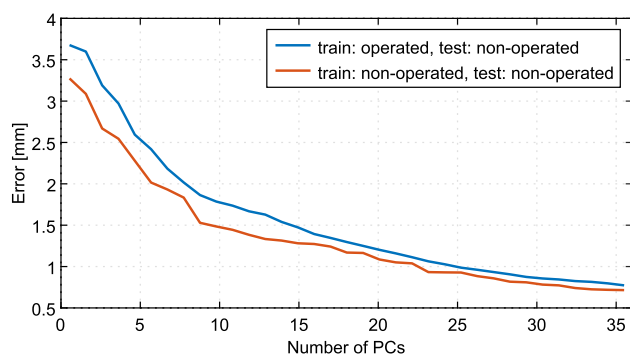

Fig. 11 Generalization ability of a statistical shape model built from 37 operated breasts, calculated using the 38 non-operated breasts as test set. For comparison purposes, a second model was trained with the same number of non-operated breasts. Its generalization ability was evaluated on the 38 non-operated breasts using leave-one-out cross-validation

from 60 and 90 breast scans and the RBSM. On the other hand, the model built from only 30 breast scans shows the worst generalization ability (about 1.07 mm when using all 29 PCs available). The model containing 90 breast scans and the RBSM achieved very similar results, both showing a generalization error of about 0.6 mm when considering the first 29 PCs. As expected, whereas the RBSM is the most specific one, the model learned from only 30 breast scans clearly performs the worst.

4.2.3 Building a statistical shape model from operated breasts

As already mentioned earlier, a quite big proportion of our 3D breast scans (72 out of 110) result from subjects who already received some kind of breast surgery. It is thus natural to assume that the RBSM might be biased toward operated breasts and is therefore not able to explain non-operated breasts well. The goal of this experiment is thus to verify how well a model built from operated breasts is able to reconstruct 3D breast scans from participants that did not undergo breast surgery. In particular, a statistical shape model is built from a

random subset of 37 (out of 72) operated breasts. Afterward, generalization ability is evaluated using the 38 non-operated breasts as test set. Again, to avoid sampling bias, the sampling procedure was repeated three times, and results were averaged. 37 out of the 38 non-operated breasts are the basis for a second model, calculating the generalization ability using leave-one-out cross-validation.

The results are shown in Fig. 11. As it can be seen, although our database is obviously biased toward operated breasts, a model built only from operated breasts *is* able to explain non-operated breasts quite well. Specifically, starting with a generalization error of over 3.5 mm, the generalization ability increases constantly as the number of PCs increase. Finally, a generalization error of 0.77 mm is achieved when using 36 PCs. For comparison, the model constructed from 37 non-operated breasts is able to reconstruct non-operated breasts only slightly better with an error of 0.71 mm when using the same number of PCs.

4.2.4 Generalization ability using clinical parameters

Our final experiment evaluates the generalization ability of the RBSM using three common, *clinical* parameters typically obtained on the breast. This way, we aim to show whether or not reconstructions obtained from the RBSM not only well explain unseen objects in general, but also preserve meaningful clinical parameters. Technically, this is quite similar to the generalization ability as defined by Styner et al. [48]. However, instead of calculating point-to-point distances between an unseen 3D breast scan and its corresponding reconstruction from the RBSM, we compare three anthropometric measurements obtained on each mesh. The following three distances between expert-annotated landmarks are measured:

- sternal notch to left nipple (SN-NL),
- sternal notch to right nipple (SN-NR), and
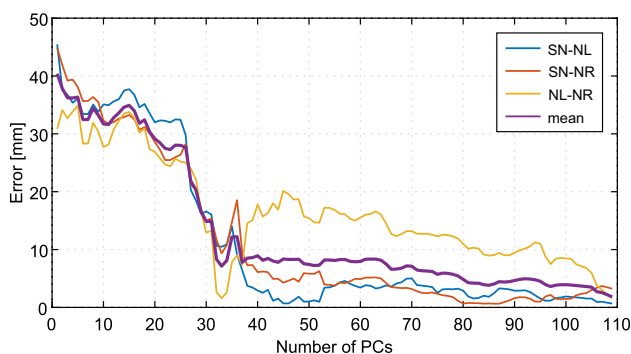- left nipple to right nipple (NL-NR).

Fig. 12 Generalization ability of the proposed RBSM, evaluated by comparing anthropometric measurements obtained on an original 3D breast scan and its corresponding reconstruction from the model. The average generalization error of the three measurements considered is also shown
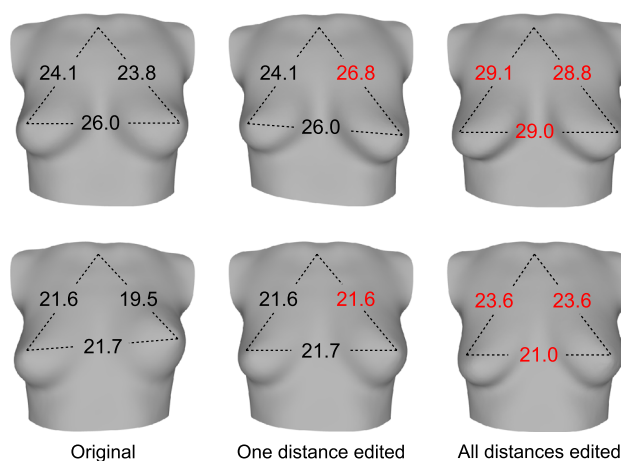

Fig. 13 Systematic breast shape editing for two examples (top row and bottom row) using the RBSM and three common anthropometric distances as features. Note that all measurements are given in centimeters

Afterward, the absolute difference between the ground-truth distance taken on the original 3D breast scan and the distance obtained from the reconstruction is calculated for each of the three anthropometric measurements. Similar to ordinary generalization based on point-to-point distances, we employ leave-one-out cross-validation for all 110 breast scans to calculate our newly defined generalization ability.

The results are shown in Fig. 12. As it can be seen, using less than 30 PCs leads to a high generalization error for all three anthropometric measurements, ranging between 15 mm and 45 mm. However, as the number of PCs increases, the generalization error drops significantly and remains low for the SN-NL and SN-NR distances. Interestingly, only the NL-NR distance increases again up to 20 mm. In summary, all three anthropometric measurements are best preserved when using 109 PCs, showing an average generalization error of 1.82 mm. The individual generalization errors when using 109 PCs are 0.66 mm (SN-NL), 3.24 mm (SN-NR), and 1.57 mm (NL-NR).

5 Applications

To further underline the expressiveness of the proposed model, this section demonstrates two exemplary applications for the RBSM that may be used for breast surgery simulation. The first application (Sect. 5.1) is based on the feature editing framework [2,9] and allows to specifically manipulate clinical parameters on the breast. In the second application (Sect. 5.2), we demonstrate how the RBSM can be used to predict a missing breast from the remaining one by utilizing *posterior* shape models [1].

5.1 Breast shape editing

So far, our statistical shape model provides a convenient way to generate new breast shapes by simply varying its parameter $\boldsymbol{\alpha} \in \mathbb{R}^q$ within a suitable range. However, a major drawback of this parameterization is the non-interpretability of $\boldsymbol{\alpha}$ in the sense that it does not correlate with any meaningful features of the breast. This clearly hinders the generation of breast shapes with certain properties or the ability to alter an existing shape according to some clinical parameters, for instance. To this end, the feature editing framework proposed by Blanz and Vetter [9] and extended by Allen et al. [2] overcomes this drawback by linearly relating features with shape parameters. Briefly, being l feature values $\{f_{i1}, f_{i2}, \ldots, f_{il}\} \subset \mathbb{R}^l$ for each individual given and stacked into a feature vector $\mathbf{f}_i := (f_{i1}, f_{i2}, \ldots, f_{il}, 1) \in \mathbb{R}^{l+1}$, $i = 1, 2, \ldots, k$, a matrix $\mathbf{F} := (\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_k) \in \mathbb{R}^{(l+1) \times k}$ is defined. Arranging the shape parameters of each individual in a matrix $\mathbf{A} := (\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \ldots, \boldsymbol{\alpha}_k) \in \mathbb{R}^{q \times k}$ and following Allen et al. [2], $\mathbf{F}$ and $\mathbf{A}$ can be related by means of an unknown transformation matrix $\mathbf{M} \in \mathbb{R}^{q \times (l+1)}$ satisfying

$$\mathbf{MF} \overset{!}{=} \mathbf{A} \implies \mathbf{M} = \mathbf{F}^+ \mathbf{A}, \tag{19}$$

where $\mathbf{F}^+$ denotes the Moore–Penrose inverse of $\mathbf{F}$. Using the matrix $\mathbf{M}$, the shape of an individual can be edited by simply providing new feature values. If $\Delta \mathbf{f}$ denotes the component-wise difference between a target feature vector and the actual feature vector of an individual, the new shape parameter $\boldsymbol{\alpha}'$ is obtained as $\boldsymbol{\alpha}' = \boldsymbol{\alpha} + \mathbf{M} \Delta \mathbf{f}$. The edited shape is given by $M(\boldsymbol{\alpha}')$.

For our exemplary application, we use the clinical parameters SN-NL, SN-NR, and NL-NR as features and set $q = k-1$. Two exemplary results are shown in Fig. 13. The RBSM
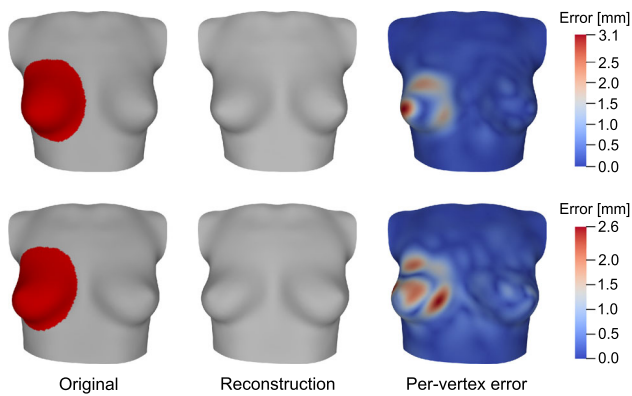
Fig. 14 Predicting the right breast (marked as *missing* in red) from the left one for two examples (top row and bottom row) using the RBSM

not only allows specific manipulation of important anthropometric distances but also produces at the same time plausible and natural-looking breast shapes, confirmed by the clinical experts involved in this project.

5.2 Breast shape prediction

While our first application is designed to alter key clinical parameters of the breast, this section demonstrates how the RBSM can be used to predict a missing breast from the existing one by means of *posterior* shape models [1].

Given an unseen 3D breast scan in correspondence with the template and represented as $\mathbf{x} \in \mathbb{R}^{3n}$, the area enclosing the missing breast is marked first. Denote the indices of marked points as $J \subset \mathbb{N}$. Following Albrecht et al. [1], the remaining $r := n - |J|$ points provide known observations which we use to compute a posterior mean $\bar{\mathbf{x}}_p \in \mathbb{R}^{3n}$. It is the likeliest reconstruction of the missing breast. Formally, denote as $\bar{\mathbf{x}}_* \in \mathbb{R}^{3r}$ and $\mathbf{U}_* \in \mathbb{R}^{3r \times q}$ the sub-vector and sub-matrix obtained by removing those entries from $\bar{\mathbf{x}}$ and $\mathbf{U}$ corresponding to the selected points in J. Similarly, $\mathbf{x}_* \in \mathbb{R}^{3r}$ represents the target 3D breast scan after removing the marked points. Then,

$$\bar{\mathbf{x}}_p = \bar{\mathbf{x}} + \mathbf{U} \left(\mathbf{U}_*^\top \mathbf{U}_* + \sigma^2 \mathbf{I}_q \right)^{-1} \mathbf{U}_*^\top \left(\mathbf{x}_* - \bar{\mathbf{x}}_* \right), \qquad (20)$$

where $\sigma^2 \geq 0$ is a small variance accounting for the deviation of $\mathbf{x}_*$ from the model [1] and $\mathbf{I}_q \in \mathbb{R}^{q \times q}$ is the identity matrix.

Two exemplary results are shown in Fig. 14. Note that, compared to simply mirroring the remaining breast, the prediction obtained from the RBSM is equipped with a statistical probability. Hence, it may not be the most symmetrical result, but, more importantly, the likeliest and probably most natural one. This is crucial especially in BRS, where the goal is to reconstruct a breast looking as natural as possible.

6 Discussion

The following section discusses the most important insights gained from experimental evaluation and also summarizes some limitations of our method.

6.1 Decoupling between breast and thorax

The experimental evaluation showed that BPM-based registration is able to decouple the breast region from the thorax well. The areas with the highest variation in random samples drawn from the RBSM are almost always coincident with the breast region whereas the thorax is kept relatively stiff. Hence, a variety of different breast shapes can be generated independently from the surrounding areas or without altering the thorax too much, implying a quite well decoupling between breast and thorax. This can be also verified from Fig. 3, showing that the unwanted variance of the surrounding regions is properly reduced when using BPM-based registration. Furthermore, the fact that our feature editing application is able to change only the requested features without altering the shape of the thorax supports our observations.

Finally, we also want to note that there might be an alternative strategy for breaking up the strong correlation between breast and surrounding regions. In particular, the framework proposed by Wilms et al. [56] allows constructing statistical shape models with a locality assumption by manipulating the sample covariance matrix. Distant areas are decoupled by explicitly setting nonzero covariances between those points to zero. This way, covariances between points in the breast region and points outside could be set to zero, effectively decoupling those regions. A similar but probably easier way to localize statistical shape models is by using GPMMs as described by Lüthi et al. [37]. We believe, however, that minimizing variances instead of covariances (between pairs of points) is a more intuitive way to weaken the coupling between breast and thorax. Note that there is a strong relation between both approaches since $|\text{cov}(x, y)| \leq \sqrt{\text{var}(x)}\sqrt{\text{var}(y)}$.

6.2 Landmarks guiding the registration process

Using only four landmarks (both nipples as well as both lower breast poles) for non-rigid registration is quite uncommon. Most of the existing methods need a lot more points to guide the registration process reasonably well. For instance, the registration of facial 3D scans for the construction of statistical shape models is usually guided by about 60 to 80 landmarks, see, e.g., [11,25]. Due to human anatomy, collecting such a huge amount of landmarks on the female breast is challenging and a fundamental problem of 3D breast scan registration. Except for the nipples (and partially the lower breast poles), there are no landmarks that can be *reli-*

ably detected only through visual inspection and consistently across a wide range of differently-shaped breasts.

Besides anatomical points, additional landmarks based on easy-detectable, non-bony structures may be used. This way, the lower breast pole was simply defined as *the lowest (most caudal) point of the breast*. From a medical point of view, this is indeed a valid definition of a landmark that is also used in clinical practice. Based on our collected data, however, it turns out that this definition is not sufficiently precise to be used for registration purposes. The actually determined lower breast poles were oftentimes *not* accurately located at the lowest point of the breast. As a result, the initially detected points had to be manually refined in about 40% of the subjects.

Another fundamental issue is the non-existence of easy-detectable landmarks on the lateral part of the breast. The absence of those points poses a serious problem in our 3D breast scan registration. In particular, when not already quite well aligned initially, the template has no chance of being pulled laterally due to the missing guidance.

To conclude, detecting landmarks on 3D breast scans reliably is challenging. Ultimately, since a physical examination is required to detect anatomical landmarks, automatic landmark detection algorithms based on 3D breast scans (only capturing the surface of the body) are almost impossible to develop. Additionally, it is important to note that even in some participants clearly detectable landmark points would be present, one has to keep in mind that those landmarks always need to be consistently located across *all* subjects. By means of the proposed registration pipeline, however, we demonstrated that 3D breast scan registration *is* generally possible by providing only four landmarks.

6.3 BPM-based registration

Our quantitative evaluation clearly indicates that BPM-based registration produces superior results than registration without BPMs, affecting both, registration *and* model quality. When using BPMs, only a little distortion was introduced. In addition, although only four landmarks are provided, pretty well correspondences are established. We observed only minor correspondence errors when randomly sampling from the RBSM.

Regarding the EBF-based representation of the BPMs, we can conclude that EBFs are quite well suited to capture the natural teardrop shape of the breast. However, in some cases, we observed that the current BPMs are not always able to properly capture the *entire* breast region, especially if the breast does not follow a typical teardrop shape. The reason is the restriction of the covariance matrices (occurring in the Mahalanobis distance) to be diagonal, pulling off a lot of flexibility for the BPM to capture unusual breast shapes. We expect that if the full covariance matrix would

be specified, better BPMs can be constructed. However, this will clearly hinder the automatic parameter selection as 30 values need to be carefully determined for each 3D breast scan.

Lastly, it is important to note that our rigid alignment will completely fail to estimate the correct transformation if the initial alignment is bad. This is a well-known behavior of the ICP algorithm as it does not optimize for a globally optimal transformation. Since our 3D breast scans are initially quite well aligned and consistently oriented due to the standardized data acquisition using photogrammetry, we did not run into this problem.

6.4 3D breast scan database

Although our 3D breast scan database contains nearly twice as many operated than non-operated breasts (72 operated, 38 non-operated), the RBSM is likely to show only a minor bias toward operated breasts. This is supported by the fact that a model built from a randomly chosen subset of 37 operated breasts is able to reconstruct non-operated breasts with an almost similar error than a model trained *and* tested on the same number of non-operated breasts. Specifically, an absolute difference of 0.06 mm between both reconstructions was reported. Intuitively, this suggests that the operated breasts contained in our database look quite natural (or, at least, similar to non-operated ones), making it challenging to decide whether a particular patient received breast surgery or not. It is, however, important to note that this experiment was conducted by considering only a subset of 37 operated breasts. This is because otherwise, a model built from an equal amount of non-operated breasts could not be constructed for comparison purposes. On the other hand, the RBSM contains 72 operated breasts. However, even if the absolute difference between a reconstruction obtained from a model trained on operated breasts and one that includes only non-operates breasts scales *exponentially* with the number of operated breasts considered, an absolute difference of 0.23 mm would be reported for the RBSM. It is thus valid to assume that the bias introduced due to operated breasts will not hinder the RBSM from explaining non-operated breasts sufficiently well.

Regarding the amount of data, our experiments confirmed what was expected: the more data, the better. The RBSM (containing all 110 breast scans) clearly outperforms the other models trained only on a subset of the data. Our experiments do not allow a conclusion about whether or not 110 breast scans are already enough in order to represent the vast amount of possible breast shapes well. However, the difference between 90 and 110 breast scans in terms of generalization and specificity is quite small (see Fig. 10) strengthening the impression of being near convergence.

Besides, we believe that in order to provide better pose-standardized training data, a more advanced technical setup is needed in which the movement of arms and shoulders is further constrained. This, however, would lengthen the whole data acquisition which might then become unfeasible to be applied in a clinical environment.

7 Conclusion and outlook

This paper proposed the RBSM—the first publicly available 3D statistical shape model for the female breast, learned from 110 breast scans. Our extensive evaluation reveals a generalization ability of 0.17 mm and a specificity error of about 2.8 mm when using all 109 PCs available. Together with the model, a fully automated, pairwise surface registration pipeline was presented which requires only four landmarks to guide the registration process. In order to break up the strong coupling between breast and thorax and thus being able to capture the actual breast shape as isolated as possible, we proposed to minimize the unwanted variance of the surrounding regions by means of a novel concept called BPMs. Defined over a surface mesh and based on EBFs, BPMs allow for a probabilistic definition of the breast region and hence overcomes the difficulty in determining an exact delineation of the breast. Later, BPMs are incorporated into the registration pipeline in order to align the template as accurately as possible inside the breast region and only roughly outside. With our experiments, we could show that this strategy effectively reduces the unwanted variance outside the breast region and hence, decouples the breast from the surrounding areas very well. As a result, the RBSM is able to generate a variety of different breast shapes quite independently from the thorax.

In our future work, we ultimately plan to combine physically motivated deformable models and statistical shape models of the breast in order to enable more realistic and statistically plausible surgical outcome simulation for BRS. The proposed RBSM is seen as a first step toward this goal. Moreover, although the results of the proposed sample applications look visually pleasing and promising from a medical point of view, further evaluation is needed in order to assess the practical impact.

Besides the two exemplary applications shown in this paper, there is room for plenty of different applications, extensions, and improvements of our model. Inspired by recent work of Göpper et al. [27], an interesting future application could include a thoracic wall into the RBSM to enable breast volume estimation from 3D surface scans. Furthermore, since our model is generative, it could be conveniently used for data augmentation in the machine learning domain. The RBSM may also be used as a prior for surface registration algorithms, effectively reducing the search space of possible deformations. Following the ideas of Booth et al. [11], so-called *bespoke models* (statistical shape models trained from a dedicated subset of the training set) may be built, for example, for different classes of BMIs, ptosis grades, or breast volumes. Additionally, instead of one global model accounting for the whole breast, two individual models for the left and right breast could be built and subsequently combined. Lastly, our rigid alignment needs to be improved to be able to deal with an arbitrary initial alignment. For instance, more robust ICP variants such as Go-ICP [58] may be used or advanced feature detectors (based on heat or wave kernel signatures [5,51] or multi-scale mean curvatures [40], for instance) combined with a robust outlier detection to find reliable correspondences for transformation calculation.

Declarations

Appendix A. Nonlinear optimization of *F* using AM

In this section, we briefly summarize the alternating minimization (AM) approach used to solve the nonlinear optimization problem in (9). Before we proceed by adapting the AM method to our specific problem, we note that the regularization term F_R is actually supposed to have three parameters. Besides the new points $\mathbf{P}'$, the scales $\{s_1, s_2, \ldots, s_n\}$ as well as rotations $\{\mathbf{R}_1, \mathbf{R}_2, \ldots, \mathbf{R}_n\}$ are also unknown. Consequently, the optimization procedure needs to take into account all three variables. AM does this by independently

optimizing for each variable, while the remaining variables are fixed.

Appendix A.1. Optimization for scale and rotation

To begin with, we first perform some algebraic manipulations to simplify further calculations. Additionally, according to Horn et al. [29], however, using the formulation of (12) leads to an asymmetry in the determination of the optimal scale factor during the optimization process. That is, if we transform the surface $\mathcal{S}$ into $\mathcal{S}'$, then the inverse scaling factor from $\mathcal{S}'$ to $\mathcal{S}$ is in general *not* $1/s$ as one would expect. We thus use the following symmetric version as suggested by Horn et al. [29], which is given by

$$\frac{1}{\sqrt{s_i}} \left(\mathbf{p}'_j - \mathbf{p}'_k \right) = \sqrt{s_i} \mathbf{R}_i \left(\mathbf{p}_j - \mathbf{p}_k \right) \qquad \forall (j,k) \in E_i. \tag{A.1}$$

Using the fact that $\|\mathbf{Q} - \mathbf{R}\|_F^2 = 6 - 2\operatorname{tr}\left(\mathbf{R}^\top \mathbf{Q}\right)$ for all $\mathbf{Q}, \mathbf{R} \in SO(3)$, we can finally rewrite (13) into

$$F_R(\mathbf{P}') = \frac{1}{2} \sum_{i=1}^n w_i$$
$$\left[\sum_{(j,k)\in E_i} w_{jk} \left(\frac{1}{s_i} \left\| \mathbf{e}'_{jk} \right\|_2^2 - 2 \left(\mathbf{e}'_{jk} \right)^\top \mathbf{R}_i \mathbf{e}_{jk} + s_i \left\| \mathbf{e}_{jk} \right\|_2^2 \right) \right.$$
$$\left. + 3\lambda \sum_{l\in N_i} w_{il} - \lambda \sum_{l\in N_i} w_{il} \operatorname{tr}\left(\mathbf{R}_l^\top \mathbf{R}_i \right) \right], \tag{A.2}$$

where we defined $\mathbf{e}'_{jk} := \mathbf{p}'_j - \mathbf{p}'_k$ and $\mathbf{e}_{jk} := \mathbf{p}_j - \mathbf{p}_k$. We will use this symmetric representation of F_R for the calculation of the optimal scale and rotation, both presented next.

Scale Taking partial derivative

$$\frac{\partial F_R}{\partial s_i} = \frac{w_i}{2} \sum_{(j,k)\in E_i} w_{jk} \left(-\frac{1}{s_i^2} \left\| \mathbf{e}'_{jk} \right\|_2^2 + \left\| \mathbf{e}_{jk} \right\|_2^2 \right) \tag{A.3}$$

and setting it to zero yield

$$s_i = \sqrt{\frac{\sum_{(j,k)\in E_i} w_{jk} \left\| \mathbf{e}'_{jk} \right\|_2^2}{\sum_{(j,k)\in E_i} w_{jk} \left\| \mathbf{e}_{jk} \right\|_2^2}} \tag{A.4}$$

since $w_i, s_i > 0$.

Rotation Next, we solve for the optimal rotation $\mathbf{R}_i$. Dropping terms of F_R that do not depend on $\mathbf{R}_i$, using the fact that $\operatorname{tr}\left(\mathbf{v}\mathbf{w}^\top\right) = \mathbf{w}^\top \mathbf{v}$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, and exploiting some well-known facts of the trace, we are remained with

$$\mathbf{R}_i = \operatorname*{arg\,min}_{\mathbf{R}_i \in SO(3)}$$
$$-\left(\sum_{(j,k)\in E_i} w_{jk} \operatorname{tr}\left(\mathbf{R}_i e_{jk} \left(\mathbf{e}'_{jk} \right)^\top \right) \right.$$
$$\left. + \frac{\lambda}{2} \sum_{l\in N_i} w_{il} \operatorname{tr}\left(\mathbf{R}_i \mathbf{R}_l^\top \right) \right) \tag{A.5}$$
$$= \operatorname*{arg\,max}_{\mathbf{R}_i \in SO(3)} \operatorname{tr}$$
$$\left(\mathbf{R}_i \underbrace{\left(\sum_{(j,k)\in E_i} w_{jk} \mathbf{e}_{jk} \left(\mathbf{e}'_{jk} \right)^\top + \frac{\lambda}{2} \sum_{l\in N_i} w_{il} \mathbf{R}_l^\top \right)}_{=:\mathbf{S}_i \in \mathbb{R}^{3\times 3}} \right)$$
$$= \operatorname*{arg\,max}_{\mathbf{R}_i \in SO(3)} \operatorname{tr}(\mathbf{R}_i \mathbf{S}_i).$$

This problem can be efficiently solved using a singular value decomposition of $\mathbf{S}_i$, see, e.g., [47].

Appendix A.2. Optimization for new points

To solve for the new points $\mathbf{P}'$, we first differentiate F w.r.t. $\mathbf{P}'$ and then set the derivative to zero. Taking derivative of F leads to

$$\nabla F(\mathbf{P}') = \nabla F_D(\mathbf{P}') + \alpha \nabla F_R(\mathbf{P}') + \beta \nabla F_L(\mathbf{P}'). \tag{A.6}$$

The first and the last term is obvious since they are just ordinary least squares objectives. The derivative of the regularization term in the middle of (A.6) is a bit more involved. We have

$$\frac{\partial F_R}{\partial \mathbf{p}'_i} = 2w_i \sum_{j\in N_i} w_{ij} \left(\left(2 + \frac{|G|}{2} \right) \mathbf{e}'_{ij} \right.$$
$$\left. - \left(s_i \mathbf{R}_i + s_j \mathbf{R}_j + \frac{1}{2} \sum_{k\in G} s_k \mathbf{R}_k \right) \mathbf{e}_{ji} \right), \tag{A.7}$$

where we defined $G := \{v \in N_i : (j,i) \in E_v\} \subset V$. Intuitively, the set G contains all vertices $v \in N_i$ *opposite* to the edge (j,i). Note that for triangle meshes and non-boundary edges $|G| = 2$, for boundary edges $|G| = 1$. Setting derivative to zero yields

$$w_i \sum_{j\in N_i} w_{ij} \mathbf{e}'_{ij}$$
$$= w_i \underbrace{\frac{1}{2 + \frac{|G|}{2}} \sum_{j\in N_i} w_{ij} \left(s_i \mathbf{R}_i + s_j \mathbf{R}_j + \frac{1}{2} \sum_{k\in G} s_k \mathbf{R}_k \right) \mathbf{e}_{ji}}_{=:\mathbf{h}_i \in \mathbb{R}^3} \tag{A.8}$$

and hence

$$\nabla F_R(\mathbf{P}') = \mathbf{0} \quad \Longleftrightarrow \quad \mathbf{WLP}' = \mathbf{WH}, \qquad (A.9)$$

where $\mathbf{W} := \mathrm{diag}(w_1, w_2, \ldots, w_n) \in \mathbb{R}^{n \times n}$, $\mathbf{H} := (\mathbf{h}_1, \mathbf{h}_2, \ldots, \mathbf{h}_n) \in \mathbb{R}^{n \times 3}$, and $\mathbf{L} \in \mathbb{R}^{n \times n}$ is the Laplacian matrix, see, e.g., [12]. Putting everything together finally yields the following sparse linear system

$$\nabla F(\mathbf{P}') = \mathbf{0} \quad \Longleftrightarrow \quad \underbrace{\begin{pmatrix} \mathbf{C} \\ \alpha\mathbf{WL} \\ \beta\mathbf{D} \end{pmatrix}}_{=:\mathbf{A}} \mathbf{P}' = \underbrace{\begin{pmatrix} \mathbf{CQ} \\ \alpha\mathbf{WH} \\ \beta\mathbf{Q}_L \end{pmatrix}}_{=:\mathbf{B}} \qquad (A.10)$$

with $\mathbf{A} \in \mathbb{R}^{3n \times n}$ and $\mathbf{B} \in \mathbb{R}^{n \times 3}$. This can be quite efficiently solved using iterative algorithms or direct solvers based on Cholesky decomposition, for example.

Appendix B. Parameters for registration

This section lists relevant parameters we used during rigid and non-rigid registration.

Rigid alignment The template and target BPMs are thresholded using a value of 0.2. Our modified ICP algorithm terminates if the relative change of MSE between template and target is less than 0.001 or a maximum of 150 iterations is reached.

Non-rigid alignment For all three phases, initial, coarse, and fine fitting $\alpha = 250$ and $\beta = 10^3$ was used. However, during initial fitting, no BPMs are used and only one iteration is performed. For the coarse and fine fitting steps, α is divided by 1.1 and 2 in each iteration, respectively. As a termination criterion, we utilized the distance term F_D and terminate if its value becomes less than 1.9 or if $\alpha < 1$.

References

1. Albrecht, T., Lüthi, M., Gerig, T., Vetter, T.: Posterior shape models. Med. Image Anal. **17**, 959–973 (2013)
2. Allen, B., Curless, B., Popovic, Z.: The space of human body shapes: reconstruction and parameterization from range scans. ACM Trans. Graph **22** (2003)
3. Ambellan, F., Lamecker, H., von Tycowicz, C., Zachow, S.: Statistical shape models: understanding and mastering variation in anatomy. Adv. Exp. Med. Biol. **1156**, 67–84 (2019)
4. Amberg, B., Romdhani, S., Vetter, T.: Optimal step nonrigid icp algorithms for surface registration. In: CVPR, pp. 1–8 (2007)
5. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: a quantum mechanical approach to shape analysis. In: ICCV Workshops, pp. 1626–1633 (2011)
6. Balaniuk, R., Costa, I., Melo, J.: Cosmetic breast surgery simulation. In: VIII Symposium on Virtual Reality, pp. 387–396 (2006)
7. Balaniuk, R., Salisbury, K.: Soft-tissue simulation using the radial elements method. In: Ayache, N., Delingette, H. (eds.) Surgery Simulation and Soft Tissue Modeling, pp. 48–58. Springer, Berlin, Heidelberg (2003)
8. Besl, P., McKay, N.: A method for registration of 3-d shapes. IEEE Trans. Pattern Anal. Mach. Intell. **14**, 239–256 (1992)
9. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH 99: Proceedings of the 26th annual conference on computer graphics and interactive techniques, pp. 187–194, ACM Press/Addison-Wesley Publishing Co. (1999)
10. Bookstein, F.: Principal warps: thin-plate splines and the decomposition of deformations. IEEE Trans. Pattern Anal. Mach. Intell. **11**, 567–585 (1989)
11. Booth, J., Roussos, A., Zafeiriou, S., Ponniah, A., Dunaway, D.: A 3d morphable model learnt from 10,000 faces. In: CVPR, pp. 5543–5552 (2016)
12. Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Levy, B.: Polygon Mesh Processing. A K Peters, USA (2010)
13. Chao, I., Pinkall, U., Sanan, P., Schröder, P.: A simple geometric model for elastic deformations. ACM Trans. Graph **29** (2010)
14. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models–their training and application. Comput. Vis. Image Underst. **61**, 38–59 (1995)
15. Costa, I., Balaniuk, R.: Lem - an approach for real time physically based soft tissue simulation. In: ICRA, pp. 2337–2343 (2001)
16. Cotin, S., Delingette, H., Ayache, N.: A hybrid elastic model for real-time cutting, deformations, and force feedback for surgery training and simulation. Vis. Comput. pp. 437–452 (2000)
17. Cruz, N.: Breast augmentation: patient satisfaction with 3d simulation of surgical outcomes. Plast. Reconstr. Surg. **136**, 161 (2015)
18. Dai, H., Pears, N., Smith, W., Duncan, C.: Statistical modeling of craniofacial shape and texture. Int. J. Comput. Vis. **128**, 547–571 (2020)
19. De Heras Ciechomski, P., Constantinescu, M., Garcia, J., Olariu, R., Dindoyal, I., Le Huu, S., Reyes, M.: Development and implementation of a web-enabled 3d consultation tool for breast augmentation surgery based on 3d-image reconstruction of 2d pictures. J. Med. Internet Res. **14**, e21 (2012)
20. Du, S., Zheng, N., Ying, S., You, S., Wu, Y.: An extension of the icp algorithm considering scale factor. In: ICIP, pp. 193–196 (2007)
21. Egger, B., Smith, W., Tewari, A., Wuhrer, S., Zollhoefer, M., Beeler, T., Bernard, F., Bolkart, T., Kortylewski, A., Romdhani, S., Theobalt, C., Blanz, V., Vetter, T.: 3d morphable face models - past, present, and future. ACM Trans. Graph **39**, 1–38 (2020)
22. Gallego, C., Martel, A.: Automatic model-based 3d segmentation of the breast in mri. In: Proc. SPIE 7962, Medical Imaging 2011: Image Processing (2011)
23. Gallo, G., Guarnera, G., Catanuto, G.: Human breast shape analysis using pca. In: Third International Conference on Bio-inspired Systems and Signal Processing, pp. 163–167 (2010)
24. Georgii, J., Eder, M., Bürger, K., Klotz, S., Ferstl, F., Kovacs, L., Westermann, R.: Computational tool for preoperative breast augmentation planning in aesthetic plastic surgery. J. Biomed. Health Inform. **18**, 907–919 (2014)
25. Gerig, T., Morel-Forster, A., Blumer, C., Egger, B., Lüthi, M., Schönborn, S., Vetter, T.: Morphable face models - an open framework. In: International Conference Automatic Face Gesture Recognition, pp. 75–82 (2017)
26. Gower, J.: Generalized procrustes analysis. Psychometrika **40**, 33–51 (1975)
27. Göpper, M., Neubauer, J., Kalash, Z., Stark, G., Simunovic, F.: Improved accuracy of breast volume calculation from 3d surface imaging data using statistical shape models. PLoS ONE **15**, e0233586 (2020)
28. Hartmann, R., Weiherer, M., Schiltz, D., Seitz, S., Lotter, L., Anker, A., Palm, C., Prantl, L., Brebant, V.: A novel method of outcome assessment in breast reconstruction surgery: Comparison of autol-

ogous and alloplastic techniques using three-dimensional surface imaging. Aesth. Plast. Surg. **44**, 1980–1987 (2020)

29. Horn, B., Hilden, H., Negahdaripour, S.: Closed-form solution of absolute orientation using orthonormal matrices. J. Opt. Soc. Am. **5**, 1127–1135 (1988)

30. Hwang, K., Yong Park, J., Won Hwang, S.: A consideration of breast imagery in art as depicted through western painting. Arch. Plast. Surg. **42**, 226–231 (2015)

31. Jiang, T., Qian, K., Liu, S., Wang, J., Yang, X., Zhang, J.: Consistent as-similar-as-possible non-isometric surface registration. Vis. Comput. **33**, 891–901 (2017)

32. Kim, Y., Lee, K., Kim, W.: 3d virtual simulator for breast plastic surgery. Comp. Anim. Virtual Worlds **19**, 512–526 (2008)

33. Lacher, R., Vasconcelos, F., Williams, N., Rindermann, G., Hipwell, J., Hawkes, D., Stoyanov, D.: Nonrigid reconstruction of 3d breast surfaces with a low-cost rgbd camera for surgical planning and aesthetic evaluation. Med. Image Anal. **53**, 11–25 (2019)

34. Levi, Z., Gotsman, C.: Smooth rotation enhanced as-rigid-as-possible mesh animation. IEEE Trans. Vis. Comput. Graph **21**, 264–277 (2014)

35. Lotter, L., Brebant, V., Eigenberger, A., Hartmann, R., Mueller, K., Baringer, M., Prantl, L., Schiltz, D.: "topographic shift": a new digital approach to evaluating topographic changes of the female breast. Arch. Gynecol. Obstet. **303**, 515–520 (2021)

36. Lüthi, M., Blanc, R., Albrecht, T., Gass, T., Goksel, O., Büchler, P., Kistler, M., Bousleiman, H., Reyes, M., Cattin, P., Vetter, T.: Statismo–a framework for pca based statistical models. Insight J. (2012)

37. Lüthi, M., Gerig, T., Jud, C., Vetter, T.: Gaussian process morphable models. IEEE Trans. Pattern Anal. Mach. Intell. **40**, 1860–1873 (2018)

38. Mazier, A., Ribes, S., Gilles, B., Bordas, S.: A rigged model of the breast for preoperative surgical planning. J. Biomech. **128**, 110645 (2021)

39. Ordas, S., Oubel, E., Leta, R., Carreras, F., Frangi, A.: A statistical shape model of the heart and its application to model-based segmentation. In: Manduca, A., Hu, X.P. (eds.) Proc. SPIE 6511, Medical Imaging 2007: Physiology, Function, and Structure from Medical Images, 65111K. SPIE - The International Society for Optical Engineering, Washington (2007)

40. Panozzo, D., Puppo, E., Rocca, L.: Efficient multiscale curvature and crease estimation. In: Skala, V., Hitzer, E.M. (eds.) 2nd International Workshop on Computer Graphics, Computer Vision and Mathematics, GraVisMa 2010 - Workshop Proc., pp. 9–16. Vaclav Skala - Union Agency, Plzen (2010)

41. Paulsen, R., Larsen, R., Nielsen, C., Laugesen, S., Ersboll, B.: Building and testing a statistical shape model of the human ear canal. In: MICCAI, pp. 373–380 (2002)

42. Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3d face model for pose and illumination invariant face recognition. In: 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, pp. 296–301. IEEE (2009)

43. Roose, L., De Maerteleire, W., Mollemans, W., Maes, F., Suetens, P.: Simulation of soft-tissue deformations for breast augmentation planning. In: Harders, M., Székely, G. (eds.) Biomedical Simulation, pp. 197–205 . Springer, Berlin, Heidelberg (2006)

44. Ruiz, G., Ramon, E., Garcia, J., Sukno, F., Gonzalez Ballester, M.: Weighted regularized statistical shape space projection for breast 3d model reconstruction. Med. Image Anal. **47**, 164–179 (2018)

45. Seo, H., Cordier, F., Hong, K.: A breast modeler based on analysis of breast scans. Comp. Anim. Virtual Worlds **18**, 141–151 (2007)

46. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Eurographics Symposium on Geometry Processing, pp. 109–116 (2007)

47. Sorkine-Hornung, O., Rabinovich, M.: Least-squares rigid motion using svd. Computing **1**, 1–5 (2017)

48. Styner, M., Rajamani, K., Nolte, L.-P., Zsemlye, G., Szekely, G., Taylor, C., Davies, R.: Evaluation of 3d correspondence methods for model building. In: Inf. Process Med. Imag. pp. 63–75 (2003)

49. Su, Z.: Statistical Shape Modelling: Automatic Shape Model Building. PhD thesis, University College London (2011)

50. Sumner, R., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. ACM Trans. Graph **26**, 80–86 (2007)

51. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. Comput. Graph. Forum **28**, 1383–1392 (2009)

52. Sung, H., Ferlay, J., Siegel, R., Laversanne, M., Soerjomataram, I., Jemal, A., Bray, F.: Global cancer statistics 2020: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. CA Cancer J. Clin. **71**(3), 209–249 (2021)

53. Vorstenbosch, J., Islur, A.: Correlation of prediction and actual outcome of three-dimensional simulation in breast augmentation using a cloud-based program. Aesth. Plast. Surg. **41**, 481–490 (2017)

54. Wang, Z., Luo, Z.-X., Zhang, J.-L., Saucan, E.: Arap++: an extension of the local/global approach to mesh parameterization. Frontiers Inf. Technol. Electronic Eng. **17**, 501–515 (2016)

55. Williams, C., Kakadaris, I., Ravi-Chandar, K., Miller, M., Patrick, C.: Simulation studies for predicting surgical outcomes in breast reconstructive surgery. In: MICCAI, pp. 9–16 (2003)

56. Wilms, M., Handels, H., Ehrhardt, J.: Multi-resolution multi-object statistical shape models based on the locality assumption. Med. Image Anal. **38**, 17–29 (2017)

57. Yamazaki, S., Kagami, S., Mochimaru, M.: Non-rigid shape registration using similarity-invariant differential coordinates. In: 3DV, pp. 191–198 (2013)

58. Yang, J., Li, H., Campbell, D., Jia, Y.: Go-icp: A globally optimal solution to 3d icp point-set registration. IEEE Trans. Pattern Anal. Mach. Intell. **38**, 2241–2254 (2016)

59. Yoshiyasu, Y., Ma, W.-C., Yoshida, E., Kanehiro, F.: As-conformal-as-possible surface registration. Comput. Graph. Forum. **33**(5), 257–267 (2014)

60. Zinßer, T., Schmidt, J., Niemann, H.: Point set registration with integrated scale estimation. In: PRIP, pp. 116–119 (2005)

Maximilian Weiherer is a PhD student at the Chair of Visual Computing at Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). He received his MSc in Computer Science from the Ostbayerische Technische Hochschule Regensburg (OTH Regensburg) in 2021. He is also affiliated with the Regensburg Medical Image Computing (ReMIC) lab at OTH Regensburg and a member of the Digital Plastic Surgery (DPS) research group. His research interests include statistical shape modeling and non-rigid surface registration.

Andreas Eigenberger received his MSc in Medical Engineering from the OTH Regensburg in 2019. He is currently working as a PhD student at the University Center for Plastic-, Aesthetic-, Hand- and Reconstructive Surgery. Additionally, he is a fellow of the Friedrich Naumann foundation. His research interests include autologous fat transfer, mesenchymal stem cells, 3D printing, and 3D reconstruction of surgical outcome.

Bernhard Egger is a junior professor heading the cognitive computer vision lab at the Chair of Visual Computing at Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). Before joining FAU, he was a postdoc in Josh Tenenbaum's Computational Cognitive Science Lab at the Department of Brain and Cognitive Sciences at MIT and Polina Golland's Medical Vision group at MIT Computer Science & Artificial Intelligence Lab. He did his PhD on facial image annotation and interpretation in unconstrained images in the Graphics and Vision Research Group at the University of Basel. His research interests span 3D computer vision, inverse rendering, statistical shape modeling, and computational visual perception with a focus on human faces.

Vanessa Brébant is a Chief Senior Physician at the University Center for Plastic-, Aesthetic-, Hand- and Reconstructive Surgery in Regensburg. As medical specialist for plastic and aesthetic surgery as well as hand surgery, her research interests are breast reconstruction, microsurgery, and quality assurance of surgical outcome using 3D scanning methods. Since 2020, Dr. Brebant is the head of the interdisciplinary research group Digital Plastic Surgery (DPS).

Lukas Prantl is head of the University Center for Plastic-, Aesthetic-, Hand- and Reconstructive Surgery and the Applied Stem Cell Research Center at the University Regensburg. He is medical specialist for surgery, plastic surgery, and aesthetic surgery. As part of his research, he develops innovative methods for cell enriched lipotransfer, stem cell isolation from adipose tissue, mesenchymal stem cell therapy, and matrix-based systems for cell integration. He is visiting assistant at the Stanford University, Rudn University, and Shanghai University. Currently, Professor Prantl is the president of the German society for plastic, reconstructive, and aesthetic surgery.

Christoph Palm was graduated in Computer Science and received his PhD (Dr. rer. nat.) from RWTH Aachen University in 1997 and 2003, respectively. He was postdoc at Research Center Jülich and Honorary Research Fellow at University College London. Currently, Professor Palm is full professor for Medical Image Computing at Ostbayerische Technische Hochschule Regensburg (OTH Regensburg), where he heads the lab Regensburg Medical Image Computing (ReMIC) and is director of the Regensburg Center of Health Sciences and Technology. His research interests include all facets of medical image computing, especially machine learning, computer-aided diagnosis, human–computer interaction, shape modeling, and 3D visualization.