



Learning shape distributions from large databases of healthy organs: applications to zero-shot and few-shot abnormal pancreas detection

Rebeca Vétel, Clément Abi Nader, Alexandre Bône, Marie-Pierre Vullierme,
Marc-Michel Rohé, Pietro Gori, Isabelle Bloch

► To cite this version:

Rebeca Vétel, Clément Abi Nader, Alexandre Bône, Marie-Pierre Vullierme, Marc-Michel Rohé, et al.. Learning shape distributions from large databases of healthy organs: applications to zero-shot and few-shot abnormal pancreas detection. Medical Image Computing and Computer Assisted Interventions, Sep 2022, Singapore, Singapore. pp.464-473, <10.1007/978-3-031-16434-7_45>. <hal-03800705>

HAL Id: hal-03800705

<https://hal.science/hal-03800705v1>

Submitted on 10 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Learning shape distributions from large databases of healthy organs: applications to zero-shot and few-shot abnormal pancreas detection

Rebeca Vétel^{1,2,*}, Clément Abi-Nader², Alexandre Bône², Marie-Pierre Vullierme³, Marc-Michel Rohé², Pietro Gori¹, and Isabelle Bloch^{1,4}

¹ LTCI, Télécom Paris, Institut Polytechnique de Paris, France

² Guerbet Research, Villepinte, France

³ Department of Radiology, Hospital of Annecy-Genevois, Université de Paris, France

⁴ Sorbonne Université, CNRS, LIP6, Paris, France

rebeca.vetil@guerbet.com

Abstract. We propose a scalable and data-driven approach to learn shape distributions from large databases of healthy organs. To do so, volumetric segmentation masks are embedded into a common probabilistic shape space that is learned with a variational auto-encoding network. The resulting latent shape representations are leveraged to derive zero-shot and few-shot methods for abnormal shape detection. The proposed distribution learning approach is illustrated on a large database of 1200 healthy pancreas shapes. Downstream qualitative and quantitative experiments are conducted on a separate test set of 224 pancreas from patients with mixed conditions. The abnormal pancreas detection AUC reached up to 65.41% in the zero-shot configuration, and 78.97% in the few-shot configuration with as few as 15 abnormal examples, outperforming a baseline approach based on the sole volume.

Keywords: Shape Analysis · Anomaly Detection · Pancreas

1 Introduction

Anatomical alterations of organs such as the brain or the pancreas may be informative of functional impairments. For instance, hippocampal atrophy and duct dilatation are well-known markers of Alzheimer’s disease and pancreatic ductal adenocarcinoma, respectively [9, 15]. In these examples, quantifying anatomical differences bears therefore a great potential for determining the patient’s clinical status, anticipating its future progression or regression, and supporting the treatment planning.

Since the seminal work of Thompson [19], the computational anatomy literature proposed several Statistical Shape Modeling (SSM) approaches, which embed geometrical shapes into metric spaces where notions of distance and difference can be defined and quantified [2, 5, 12]. Taking advantage of these representations, statistical shape models were then proposed to perform group analyses

*Corresponding author: rebeca.vetil@guerbet.com

of shape collections. In particular, atlas models [16] learn geometrical distributions in terms of an “average” representative shape and associated variability, generalizing the Euclidean mean-variance analysis. In medical imaging, learning atlases from healthy examples allows for the definition of normative models for anatomical structures or organs, such as brain MRIs or subcortical regions segmented from neuroimaging data [10, 22], thus providing a natural framework for the detection of abnormal anatomies.

In practice, leveraging an atlas model to compute the likelihood of a given shape to belong to the underlying distribution either requires to identify landmarks [6], or to solve a registration problem [3]. To circumvent the computational cost of this shape embedding operation, the authors in [21] proposed to train an encoder network to predict registration parameters from image pairs. In [7, 14], the authors built on this idea and used the variational autoencoder (VAE) of [13] to learn the embedding space jointly with the atlas model, instead of relying on pre-determined parametrization strategies. However, the structure of the decoding network remained constrained by hyperparameter-rich topological assumptions, enforced via costly smoothing and numerical integration operators from a computational point of view.

Alternative approaches proposed to drop topological hypotheses by relying on variations of the AE or its variational counterparts [13] to learn normative models that are subsequently used to perform Anomaly Detection (AD). These methods compress and reconstruct images of healthy subjects to capture a normative model of organs [1, 23]. Yet, they are usually applied on the raw imaging data, thus they entail the risk of extracting features related to the intensity distribution of a dataset which are not necessarily specific to the organ anatomy. Therefore, regularization constraints [1, 4] are introduced to improve the detection performances compared to the vanilla AE. To further reduce the overfitting risk, these methods artificially increase the dataset size by working on 2D slices.

Given this context, we propose a VAE-based method to learn a normative model of organ shape that can subsequently be used to detect anomalies, thus bridging the gap between SSM and AD models. Although SSM methods with explicit modeling constraints proved effective to learn relevant shape spaces from relatively small collections of high-dimensional data, we propose to further reduce the set of underlying hypotheses and leave the decoding network unconstrained in its architecture. With the objective to learn normative shape models from collections of healthy organs, we argue that sufficiently large databases of relevant medical images can be constructed by pooling together different data sources, see [8] for instance. To reduce the risk of overfitting and focus on the anatomy of organs, the VAE is learned from 3D binary segmentation masks and is coupled with a shape-preserving data augmentation strategy consisting of translations, rotations and scalings. An approach to study and visualize group differences is also proposed.

Section 2 details the proposed method, which is then illustrated on a pancreas shape problem in Section 3. Section 4 discusses the results and concludes.

2 Methods

Modeling organ shape. We consider an image acquired via a standard imaging technique. For a given organ in the image, its anatomy can be represented by a binary segmentation mask $\mathbf{X} = \{x_i, i = 1 \dots d\}$ with $x_i \in \{0, 1\}$ and d the number of voxels in the image. We are interested in studying the shape of this organ, and assume that it is characterized by a set of underlying properties that can be extracted from the segmentation mask. Therefore, we hypothesize the following generative process for the segmentation mask:

$$p_\theta(\mathbf{X} \mid \mathbf{z}) = \prod_{i=1}^d f_\theta(\mathbf{z})_i^{x_i} (1 - f_\theta(\mathbf{z})_i)^{1-x_i} \quad (1)$$

where $0^0 = 1$ by convention, and \mathbf{z} is a latent variable generated from a prior distribution $p(\mathbf{z})$. This latent variable provides a low-dimensional representation of the segmentation mask embedding its main shape features. The function f_θ is a non-linear function mapping \mathbf{z} to a predicted probabilistic segmentation mask.

We are interested in inferring the parameters θ of the generative process, as well as approximating the posterior distribution of the latent variable \mathbf{z} given a segmentation mask \mathbf{X} . We rely on the VAE framework [13] to estimate the model parameters. Hence, we assume that $p(\mathbf{z})$ is a multivariate Gaussian with zero mean and identity covariance. We also introduce the approximate posterior distribution $q_\phi(\mathbf{z} \mid \mathbf{X})$ parameterized by ϕ , and optimize a lower bound \mathcal{L} of the marginal log-likelihood, which can be written for the segmentation mask \mathbf{X}^p of a subject p as:

$$\mathcal{L} = \mathbb{E}_{q_\phi(\mathbf{z} \mid \mathbf{X}^p)} [\log p_\theta(\mathbf{X}^p \mid \mathbf{z})] - KL[q_\phi(\mathbf{z} \mid \mathbf{X}^p) \mid p(\mathbf{z})], \quad (2)$$

where $q_\phi(\mathbf{z} \mid \mathbf{X}^p)$ follows a Gaussian distribution $\mathcal{N}(\mu_\phi(\mathbf{X}^p), \sigma_\phi^2(\mathbf{X}^p)\mathbf{I})$ with \mathbf{I} the identity matrix, and KL is the Kullback-Leibler divergence.

To capture shape features, we rely on a convolutional network and adopt the U-Net [17] encoder-decoder architecture without skip connections*. In practice, the number of convolutional layers and the convolutional blocks are automatically inferred thanks to the nnU-Net self-configuring procedure [11] (see Section A in the supplementary material for details). Due to this encoder-decoder architecture, the segmentation masks are progressively down-sampled to obtain low-resolution feature maps which are mapped through a linear transformation to the latent variable \mathbf{z} . The latent code is subsequently decoded by a symmetric path to reconstruct the original masks.

Anomaly detection. We propose to learn a normative model of organ shapes by applying the VAE framework previously presented on the segmentation masks of a large cohort of N healthy patients, allowing the model to capture a low-dimensional embedding characteristic of a normal organ anatomy. In addition, we use a data augmentation procedure consisting of random translations, rotations and scalings, in order to be invariant to these transformations

*Code available at <https://github.com/rebeca-vetil/HealthyShapeVAE>.

and force the network to extract shape features. Based on this learned model, we propose two approaches to perform Anomaly Detection (AD) by leveraging the latent representation of normal organ shapes.

Zero-shot learning method. After training, the recognition model $q_\phi(\mathbf{z} | \mathbf{X})$ can be used to project the segmentation maps \mathbf{X}^p of the cohort of healthy subjects and obtain an empirical distribution of normal shapes in the latent space. We rely on this low-dimensional distribution of normality to detect abnormal shapes. To do so, we compute the mean of the healthy subjects projection, and define abnormality through the L2 distance to this mean latent representation.

Few-shot learning method. Another approach is to classify normal and abnormal shapes based on their low-dimensional representations. In practice, we project the segmentation maps from a set of healthy and pathological subjects in the latent space using the recognition model $q_\phi(\mathbf{z} | \mathbf{X})$. Therefore, we obtain for all these subjects a set of low-dimensional organ shape features that we can use to learn any type of classifier (*e.g.*, linear SVM).

Studying organ shapes differences. Our framework can also be used to study organ differences between groups. Let us consider a set of healthy and pathological subjects, as well as their segmentation masks. Based on the recognition model $q_\phi(\mathbf{z} | \mathbf{X})$, we can compute the average of the subject’s latent projection for each group, denoted by \mathbf{z}_{normal} and $\mathbf{z}_{abnormal}$, respectively. We consider the line of equation $(1 - t) \times \mathbf{z}_{normal} + t \times \mathbf{z}_{abnormal}$ with $t \in \mathbb{R}$. When moving along this line with increasing values of t , we progress from a healthy mean latent shape representation to a pathological one, and can reconstruct the corresponding segmentation mask using the probabilistic decoder $p_\theta(\mathbf{X} | \mathbf{z})$.

3 Experiments

In this section, we applied our method in the case of the pancreas. A normative model of pancreas shape was learned on a large cohort of healthy subjects, and was then leveraged for AD on an independent test cohort. Several configurations were proposed to assess the model performances, including the impact of the number of training subjects and of the latent space dimensionality on the AD performances. Detection with the few-shot learning method was performed using Support Vector Machine (SVM). Finally, we showed how the proposed framework can be used to visualize differences between the healthy and pathological pancreas.

Training. The training dataset \mathcal{D}^{Train} was created from our private cohort containing 2606 abdominal Portal CT scans of patients with potential liver cancer. To ensure the healthy condition and shape of the pancreas, several exclusion criteria were applied (see Section B in the supplementary material). Finally, 1200 portal CT scans were retained. To explore the influence of the number of samples seen during training, subsets \mathcal{D}_N with a growing number of subjects were created (see Table 1 in the supplementary material). For each \mathcal{D}_N , 80% and 20% of the samples were used for training and validation, respectively. Splitting was done such that the pancreas volume distribution was balanced across the splits.



Fig. 1. Examples of normal and abnormal pancreas shapes from \mathcal{D}^{Test} . Green and red figures are examples taken from $\mathcal{D}_{normal}^{Test}$ and $\mathcal{D}_{abnormal}^{Test}$, respectively.

Testing. The test database \mathcal{D}^{Test} was obtained by combining two datasets: i) a private dataset $\mathcal{D}_{abnormal}^{Test}$ containing 144 cases diagnosed with pancreatic cancer, and for whom the pancreas shape was evaluated as abnormal by an expert radiologist; ii) a public dataset $\mathcal{D}_{normal}^{Test}$ from The Cancer Imaging Archive (TCIA) containing 80 CT scans [18] of patients who neither had abdominal pathologies nor pancreatic lesions, and for whom the assumption of normal pancreas shape held. Centers, machines and protocols differed among the three datasets. Examples of normal and abnormal shapes can be seen in Figure 1.

Preprocessing. The first step consisted in obtaining the pancreas segmentation masks. For the public dataset $\mathcal{D}_{normal}^{Test}$, we used the reference pancreas segmentation masks provided by TCIA. For $\mathcal{D}_{normal}^{Train}$ and $\mathcal{D}_{abnormal}^{Test}$, the masks were obtained semi-automatically using an in-house segmentation algorithm derived from the nnU-Net, and validated by a radiologist with 25 years of expertise in abdominal imaging. Finally, all the masks were resampled to $1 \times 1 \times 2 \text{ mm}^3$ in (x, y, z) directions, and centered in a volume of size $192 \times 128 \times 64$ voxels.

Zero-shot AD. We trained our model on the different datasets \mathcal{D}_N^{Train} , with a growing number of latent dimensions L ranging from 16 to 1024 (denoted by $L_{16} \dots L_{1024}$). For each experiment, we applied the zero-shot AD procedure, as previously explained, on \mathcal{D}^{Test} . We report the Area Under the Curve (AUC), in %, in Table 1. Increasing the dimension of the latent space L improved the classification performances on each dataset \mathcal{D}_N . Moreover, for each dataset size the best result was consistently obtained when L was set at the maximum value L_{1024} . We also observed that the effect of the latent space dimension on the performances seemed to attenuate as the dataset size increased. Indeed, we observed that when going from L_{16} to L_{1024} , the mean AUC for \mathcal{D}_{300} , \mathcal{D}_{600} , \mathcal{D}_{900} , \mathcal{D}_{1200} improved by 10.7, 5.9, 4.0 and 3.1 points, respectively. Regarding the effect of the database size, we observed that increasing the training set size seemed to globally improve the AUC scores. For instance, going from \mathcal{D}_{300} to \mathcal{D}_{600} increased the classification performances for all the experiments, particularly for L_{16} which gained 9.3 points. This beneficial effect of both larger training sets and latent dimension was also observed on the Dice score between the original and reconstructed segmentation masks (see Table 2 in the supplementary material). Thus, for the following experiments, we chose the model trained on \mathcal{D}_{1200} with a latent dimension L_{1024} as it gave the best results in terms of AUC and Dice scores.

To visualize the separation between normal and abnormal shapes, we projected each subject from \mathcal{D}^{Test} using the recognition model $q_\phi(\mathbf{z}|\mathbf{X})$. Based on the

Table 1. Results for zero-shot AD. For each experiment, corresponding to a specific training size \mathcal{D} and latent space dimension L , we report the mean and standard deviation of AUC scores (in %) obtained by bootstrapping with 10000 repetitions. Best results by line are underlined and by column are in **bold**.

	L_{16}	L_{64}	L_{256}	L_{1024}
\mathcal{D}_{300}	51.51 \pm 0.37	59.08 \pm 0.37	62.16 \pm 0.37	<u>62.17</u> \pm 0.37
\mathcal{D}_{600}	59.24 \pm 0.38	60.97 \pm 0.36	64.32 \pm 0.36	<u>65.11</u> \pm 0.36
\mathcal{D}_{900}	60.77 \pm 0.37	62.64 \pm 0.37	64.04 \pm 0.36	<u>64.81</u> \pm 0.36
\mathcal{D}_{1200}	62.28 \pm 0.36	61.74 \pm 0.37	62.58 \pm 0.37	<u>65.41</u> \pm 0.36

subjects’ latent representation, we applied three dimensionality reduction techniques, namely Principal Components Analysis (PCA), t-distributed Stochastic Neighbor Embedding (t-SNE) and Isomap. Results are displayed in Figure 2, on which each point represents the latent projection of a test subject reduced on a 2D plane. We observed that, independently of the projection technique, normal and abnormal shapes tended to be separated in two different clusters.




Fig. 2. 2D reduction of the latent representation of the test subjects. The 80 samples from $\mathcal{D}_{normal}^{Test}$ are in green, and the 144 samples from $\mathcal{D}_{abnormal}^{Test}$ are in red.

Few-shot AD. We trained a linear SVM classifier on the latent representation of \mathcal{D}^{Test} with stratified k-fold cross-validation. We varied the number k of folds to test the performances of the classifier depending on the train/test samples ratio. Experiments ranged from a 0.05 train/test ratio to a leave-one-out cross-validation and are presented in Table 2. We noticed that using only 8 healthy and 15 abnormal training samples increased the performance to 78.9%. We also observed that the AUC scores and the balanced accuracy increased with the number of training samples, reaching a maximum of 91.1% and 83.2% respectively in the leave-one-out configuration.

Comparison with a baseline method. We compared our approach with a baseline method classifying shapes based on their volume. We applied this method on \mathcal{D}^{Test} with bootstrap sampling and obtained an average AUC of 51% with a 95% confidence interval of [49.9; 51.7], below the maximum AUC scores

Table 2. Results for few-shot AD. For each experiment, we indicate the number of training samples, as well as the number of abnormal samples (in brackets). We report the means and standard deviations for AUC (in %) and balanced accuracy (in %), obtained by bootstrapping with 10000 repetitions.

Train/Test ratio	0.05	0.11	0.25	1	223
Number of training samples	12 (8)	23 (15)	45 (29)	112 (72)	223 (144)
AUC	66.02 \pm 0.02	78.97 \pm 0.03	81.87 \pm 0.07	86.95 \pm 0.23	91.18 \pm 0.19
Balanced Accuracy	67.78 \pm 0.03	70.88 \pm 0.05	73.96 \pm 0.10	75.49 \pm 0.37	83.26 \pm 0.34

of 65.4% and 91.1% previously reported in the zero-shot and few-shot cases, respectively. We also compared the proposed method to two SSM methods: active shape models (ASM) [6] and Large Deformation Diffeomorphic Metric Mapping (LDDMM) using the Deformetrica software [3]. Details and results, reported in Section D in the supplementary material, show that our approach outperforms these two state-of-the-art methods in both zero and few-shot configurations.

Studying pancreas shapes differences between groups. To model differences in the pancreas shape between healthy and pathological groups, we applied the procedure presented in Section 2 on the subjects from $\mathcal{D}_{normal}^{Test}$ and $\mathcal{D}_{abnormal}^{Test}$. Figure 3 shows the pancreas shapes obtained for different values of t . When going from a healthy towards a pathological latent representation, we observed a shrinkage of the shape in the body for the generated pancreas.

4 Discussion and conclusion

We presented a method based on a VAE to learn a normative model of organ shape. We hypothesized that such a model could be learned from large databases of healthy subjects. The method was applied in the case of the pancreas, for which morphological changes can be a marker of disease.

We empirically observed that large training sets and latent dimensions were beneficial to the model in terms of AD performances. Our results also demonstrated that the model captured features that distinguished between normal and




Fig. 3. Generated pancreas shapes. Pancreas shapes generated by decoding latent representations lying on the line of equation $(1 - t) \times \mathbf{z}_{normal} + t \times \mathbf{z}_{abnormal}$.

abnormal shapes in the latent space, as illustrated in Figure 2. From a quantitative point of view, we observed in the zero-shot case - *i.e.* without supervision - that the best model obtained an AUC score of $65.41 \pm 0.36\%$, which significantly outperformed a naive model classifying shapes based on their volume. In the few-shot experiments, we obtained a mean AUC score of 77.4% by training a SVM with only 8 healthy subjects and 15 pathological subjects. These findings highlight the discriminating properties of the latent normative model of pancreas shape estimated by our model. Moreover, classification performances reached up to 91.1% AUC and 83.2% balanced accuracy when training the classifier on 223 samples in a leave-one-out fashion. These results are in line with [15], where the authors reported a balanced accuracy of 85.2% on their private dataset. Yet, our approach differs from theirs by its paradigm. Instead of training a supervised model for joint shape representation and classification, we propose to learn a normative model of shape. The advantage of this approach is that it does not require different types of patients to be trained but solely a database of healthy subjects. Moreover, it can be used in an unsupervised manner (*cf.* zero-shot) or with few labeled data (*cf.* few-shot), with good performances in both cases.

Finally, we also showed that our framework could be used to study and visualize the morphological differences between the organ shape of different clinical groups, based on an exploration of the latent space. The anatomical changes observed in Figure 3 seemed to concur with clinical evidence as the shrinkage suggests partial parenchymal atrophy [20]. This hypothesis would require further medical evaluation, and could be the subject of a proper clinical validation.

Acknowledgments This work was partly funded by a CIFRE grant from ANRT # 2020/1448.

References

1. Baur, C., Denner, S., Wiestler, B., Navab, N., Albarqouni, S.: Autoencoders for unsupervised anomaly segmentation in brain MR images: a comparative study. *Medical Image Analysis* **69**, 101952 (2021)
2. Beg, M., Miller, M., Trouné, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* **61**(2), 139–157 (2005)
3. Bône, A., Louis, M., Martin, B., Stanley, D.: Deformetrica 4: An open-source software for statistical shape analysis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 3–13. Springer (2018)
4. Chen, X., Konukoglu, E.: Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders. In: *MIDL* (2018)
5. Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing* **5**(10), 1435–1447 (1996)
6. Cootes, T.F., Taylor, C.J., Graham, J.: Active shape models-their training and application. *Computer Vision and Image Understanding* **61**(1), 38–50 (1995)
7. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 729–738. Springer (2018)
8. Dufumier, B., Gori, P., Victor, J., Grigis, A., Wessa, M., Brambilla, P., Favre, P., Polosan, M., McDonald, C., Piguet, C.M., et al.: Contrastive learning with continuous proxy meta-data for 3D MRI classification. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 58–68. Springer (2021)
9. Fox, N., Warrington, E., Freeborough, P., Hartikainen, P., Kennedy, A., Stevens, J., Rossor, M.N.: Presymptomatic hippocampal atrophy in Alzheimer’s disease: a longitudinal MRI study. *Brain* **119**(6), 2001–2007 (1996)
10. Gori, P., Colliot, O., Marrakchi-Kacem, L., Worbe, Y., Poupon, C., Hartmann, A., Ayache, N., Durrleman, S.: A Bayesian Framework for Joint Morphometry of Surface and Curve meshes in Multi-Object Complexes. *Medical Image Analysis* **35**, 458–474 (2017)
11. Isensee, F., et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
12. Kendall, D.G.: Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society* **16**(2), 81–121 (1984)
13. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. In: *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada* (2014)
14. Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T.: Learning a probabilistic model for diffeomorphic registration. *IEEE Transactions on Medical Imaging* **38**(9), 2165–2176 (2019)
15. Liu, F., Xie, L., Xia, Y., Fishman, E., Yuille, A.: Joint shape representation and classification for detecting PDAC. In: *International Workshop on Machine Learning in Medical Imaging*. pp. 212–220. Springer (2019)
16. Pennec, X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision* **25**(1), 127–154 (2006)
17. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)

18. Roth, H., et al.: Data from pancreas-CT. The Cancer Imaging Archive (2016). <https://doi.org/10.7937/K9/TCIA.2016.tNB1kqBU>
19. Thompson, D.W.: On Growth and Form. Cambridge Univ. Press (1917)
20. Yamao, K., Takenaka, M., Ishikawa, R., Okamoto, A., Yamazaki, T., Nakai, A., Omoto, S., Kamata, K., Minaga, K., Matsumoto, I., et al.: Partial pancreatic parenchymal atrophy is a new specific finding to diagnose small pancreatic cancer (≤ 10 mm) including carcinoma in situ: comparison with localized benign main pancreatic duct stenosis patients. *Diagnostics* **10**(7), 445 (2020)
21. Yang, X., Kwitt, R., Styner, M., Niethammer, M.: Quicksilver: Fast predictive image registration – A deep learning approach. *NeuroImage* **158**, 378–396 (2017)
22. Zhang, M., Singh, N., Fletcher, P.T.: Bayesian estimation of regularization and atlas building in diffeomorphic image registration. In: IPMI. vol. 23, pp. 37–48 (2013)
23. Zimmerer, D., Isensee, F., Petersen, J., Kohl, S.A.A., Maier-Hein, K.: Unsupervised anomaly localization using variational auto-encoders. In: MICCAI. pp. 289–297 (2019)