

Learning a statistical full spine model from partial observations

Di Meng¹, Marilyn Keller², Edmond Boyer¹,
Michael Black², and Sergi Pujades¹

¹ Inria, Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, France

² Max Planck Institute for Intelligent Systems, Tübingen, Germany

Abstract. The study of the morphology of the human spine has attracted research attention for its many potential applications, such as image segmentation, bio-mechanics or pathology detection. However, as of today there is no publicly available statistical model of the 3D surface of the full spine. This is mainly due to the lack of openly available 3D data where the full spine is imaged and segmented. In this paper we propose to learn a statistical surface model of the full-spine (7 cervical, 12 thoracic and 5 lumbar vertebrae) from partial and incomplete views of the spine. In order to deal with the partial observations we use probabilistic principal component analysis (PPCA) to learn a surface shape model of the full spine. Quantitative evaluation demonstrates that the obtained model faithfully captures the shape of the population in a low dimensional space and generalizes to left out data. Furthermore, we show that the model faithfully captures the global correlations among the vertebrae shape. Given a partial observation of the spine, i.e. a few vertebrae, the model can predict the shape of unseen vertebrae with a mean error under 3 mm. The full-spine statistical model is trained on the VerSe 2019 public dataset and is publicly made available to the community for non-commercial purposes³.

Keywords: Spine statistical model · Vertebrae surface · Incomplete data

1 Introduction

A reliable spine model with accurate vertebrae structures is essential in numerous medical application such as image segmentation, orthopedics, anesthesiology and pathology quantification. A full spine model can help segment 3D volumetric images such as computed tomography (CT) or magnetic resonance imaging (MRI), as well as 2D projections like X-ray images, for instance by injecting prior knowledge on the global shape of the spine as a regularizer. Statistical models have also shown their interest in the generation of synthetic data for different modalities. A full spine model can also be used as a reference structure in clinical

³ https://gitlab.inria.fr/spine/spine_model

practice supporting the localization of other organs, as well as be used in the diagnosis of spine scoliosis.

The statistical shape model is one of the most employed statistical atlas [10]. Creating statistical models of shape using surface measurements has attracted a lot of attention since the early work of Blanz and Vetter [2]. While early methods mostly rely on Principal Component Analysis, non-linear methods have been also proposed to overcome the fact that unseen shapes cannot be expressed by linear combinations of the training samples [24].

In order to create a statistical surface model of the full spine using the classical methodology, the full spine of several subjects needs to be observed. However, observations of the full spine (e.g. CT images) are rare. Practitioners usually narrow down the field of view of the spine to one specific level for a detailed diagnosis as well as to reduce the dose given to the patient. Thus most work has been done to model the shape of the vertebrae which are consistently observed in the scans. For instance, lumbar vertebrae have received a special attention for low back pain as this area supports the greatest load of the spine [5]. Hollenbeck et al. [11] generated a statistical shape model of the lumbar region from 52 subjects using PCA with a focus on the L4-L5 and L5-S1 functional spinal units (FSU). Campbell et al. [6] used automated methods for landmark identification to create a statistical shape model of the lumbar spine, applied in biomechanics and population-based finite element modeling. A nonlinear SSM based on kernel PCA was investigated for 3D active shape model segmentation [15]. They used 7 CT scans covering vertebrae level from T10 to L3.

Another difficulty in the learning of statistical models of the spine is that the relative position of vertebrae highly depends on the patient's posture during data acquisition. Overlapping and excessive separation of neighboring structures may appear in the vertebrae segmentation. In order to tackle this problem, a model of the interspace of the vertebrae was developed [8], by considering the variations of the space between two surfaces. Similarly, our model only focuses on the relative translation of the vertebrae and disregards their relative rotation. Multi-vertebrae statistical models [4][19] capture the variations in shape and pose simultaneously, and reduce the number of registration parameters, allowing to help in the segmentation of a section of the vertebral column. These models were tested on 32 subjects and with a focus on the lumbar vertebrae section.

While most works have focused on the modeling of a specific part of the spine, a few works have modeled the full spine. For instance, Klinder et al. [13] designed an automatic framework segmenting vertebrae from arbitrary CT images with a full spine surface model. To create the model they first scanned a commercially available plastic phantom to create the template, and then they manually registered it to ten actual scans of the full spine. Mirzaalian et al. [17] learned a statistical shape model of the spine by independently learning three models, one for each level (cervical, thoracic and lumbar). Thus, their per-level models do not learn the shape correlations across the full spine.

Other probabilistic models, different from the shape surface models, such as probabilistic atlas [20], graph models [23], hidden markov models [9] and hier-

archical models[22][26][7] have also been proposed. For example Ruiz et al. [20] proposed a probabilistic atlas of the spine. By co-registering 21 CT scans, a probability map is created which can be used to segment and detect the vertebrae with a special focus on ribs suppression. Schmidt et al. [23] proposed a probabilistic graphical model for the location and identification of the vertebrae in MR images. In both cases full spines were observed at train time and the proposed methods cannot be used to infer the shape of the full spine from a partial observation.

As full spines are not publicly available, we learn our novel statistical surface model from partial spine observations using Probabilistic Principal Component Analysis (PPCA). In the medical domain, PPCA has been used for craniofacial statistical modeling [16]. However, in their work full observations of the skull were used for training. In contrast, our proposed model never observed a complete spine. A robust shape model can capture the characteristics of individual vertebrae and the anatomy of the spine. It also provides the appearance information and relative positions of the observed objects. However, existing statistical models have focused on a specific region of the spine or were learned from full spine observations. Unfortunately, none of them are publicly available.

Our full spine statistical model is composed of 24 vertebrae including cervical, thoracic and lumbar. The model is learned from the openly available VerSe 2019 Dataset [21] containing partial observations and annotations of the spine. Our model captures the strong shape correlations across the individual vertebrae of the spine. Given an arbitrary field-of-view and incomplete scan, our model is able to predict the full spine adapting to the shape variations of the subject. The model is openly available to the community for non-commercial purposes.

2 Method

The inputs to our method are i) 80 manually segmented CT volumes \mathbf{V}^i , where $i \in [1, 80]$ denotes the volume index, together with the annotations of the individual vertebra from the VerSe Dataset [21] and ii) 24 artist created mesh templates \mathbf{T}_k with $k \in [1, \dots, 24]$ being the vertebra index, one for each individual vertebra. Each template \mathbf{T}_k has N_k 3D vertices and the full spine mesh including all 24 vertebrae has 127.294 vertices and 254.636 faces. All template meshes \mathbf{T}_k are symmetric with respect to the medial plane. From the segmented volumes \mathbf{V}^i we extract meshes using the Marching Cubes approach. We name them scans and note them as \mathcal{S}_k^i . Let us note that: i) in the VerSe Dataset none of the volumes contains the 24 vertebra; ii) some vertebra are imaged at the boundary of the CT volume and thus only provide an incomplete observation. These scans have holes and are not watertight; we note them $\tilde{\mathcal{S}}_k^i$. An overview of the VerSe Dataset vertebrae is shown in Fig. 1.

The creation of the statistical spine shape model consists of several steps. First, in Sec. 2.1 we detail how we perform a non-rigid registration of each template \mathbf{T}_k to all watertight scans $\{\mathcal{S}_k^i - \tilde{\mathcal{S}}_k^i\}$. To bring all registrations to a coherent frame, we compute a rigid transformation [1] between the registration and the

template. We name these coherent registrations \mathcal{I}_k^i . Then, for each vertebra k , we compute Principal Component Analysis (PCA) on all the available registrations. As we have few observations, we exploit the medial plane symmetry: for each registration we create the symmetric registration. With the registrations and their symmetric we obtain 24 (symmetric) mean shapes \mathbf{T}_k^μ and their 24 individual principal vectors \mathbf{B}_k . With the obtained PCA, we perform a second non-rigid registration of each template \mathbf{T}_k as explained in Sec. 2.2. This time all scans \mathcal{S}_k^i , including the non-watertight, are registered. We initialize the template k vertices with the mean shape \mathbf{T}_k^μ and use the computed shape space \mathbf{B}_k to constrain the registration process. We compute the rigid transformation to obtain the coherent registrations and keep the corresponding rotation R_k^i and translation t_k^i . We exploit again the medial plane symmetry and obtain a new set of registrations \mathcal{A}_k^i . From these registrations, we use Probabilistic Principal Component Analysis (PPCA) (Sec. 2.3) and we learn the full-spine shape model (Sec 2.4) defined by \mathbf{T}^μ and \mathbf{B} .

2.1 Initial registration

We first perform a non-rigid registration of each template \mathbf{T}_k to all watertight scans $\{\mathcal{S}_k^i - \bar{\mathcal{S}}_k^i\}$. Because the scans are watertight, we want all template vertices to be close to the scan surface, and vice versa, all scan vertices to be close to the template surface. We effectively enforce this constraint by computing the vertex to surface distance from a point set to a surface and define the energy $E_{p2m}(\mathcal{S}, \mathbf{T})$, which accounts for the distance of the vertices of \mathcal{S} to the mesh surface \mathbf{T} . The registration has three steps. Translation optimization, translation and rotation optimization and free form optimization. To simplify notation we drop the volume indices i and keep k to indicate the per-vertebrae process.

We start by computing a translation \mathbf{t} , so that $\mathbf{T}_k + \mathbf{t}$ is close to \mathcal{S} by minimizing

$$E(\mathcal{S}, \mathbf{T}_k, \mathbf{t}) = E_{p2m}(\mathcal{S}, \mathbf{T}_k + \mathbf{t}) + \lambda_{m2s} E_{p2m}(\mathbf{T}_k + \mathbf{t}, \mathcal{S}) \quad (1)$$

w.r.t. to \mathbf{t} , where $\lambda_{m2s} = 1$. Next we optimize for a 3D rotation, parametrized by a 3D Rodriguez vector \mathbf{r} . Given a 3D vector \mathbf{r} and a mesh \mathbf{T} , we use $R(\mathbf{T}, \mathbf{r})$ to describe the 3D rotated mesh. The CT scans from the VerSe Dataset (see Sec. 3) have a consistent patient orientation encoded in the dicom metadata. After reorienting all the CTs along the same orientation, we can thus initialize the rotation vector r ($r = [-2, 0, 2]$) with the same value for all $i \in [1, 80]$. This value corresponds to the rotation between the template frame and the acquisition frame defined by the radiologist. Then we optimize

$$E(\mathcal{S}, \mathbf{T}_k, \mathbf{t}, \mathbf{r}) = E_{p2m}(\mathcal{S}, R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}) + \lambda_{m2s} E_{p2m}(R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}, \mathcal{S}) \quad (2)$$

w.r.t. to \mathbf{t} and \mathbf{r} , where \mathbf{t} is initialized with the result of (1). Next we allow all vertices of $R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}$ to freely deform to best match \mathcal{S}_k^i . This free-form deformation is represented as an additive 3D vector added to each vertex that

we note \mathbf{f} . To regularize the position of these displacements we use a coupling term on edges $E_{cpl}(\mathbf{T}, \mathbf{T} + \mathbf{f})$, enforcing the edges of the registration to be close to the edges of the initial template shape. To compute E_{cpl} we use the same energy term defined in Eq. 8 from [3] and optimize

$$\begin{aligned} E(\mathcal{S}, \mathbf{T}_k, \mathbf{t}, \mathbf{r}, \mathbf{f}) = & E_{p2m}(\mathcal{S}, (R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}) + \mathbf{f}) + \\ & + \lambda_{m2s} E_{p2m}((R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}) + \mathbf{f}, \mathcal{S}) + \\ & + \lambda_{cpl} E_{cpl}((R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}), (R(\mathbf{T}_k, \mathbf{r}) + \mathbf{t}) + \mathbf{f}) \quad (3) \end{aligned}$$

w.r.t. to \mathbf{f} and by keeping \mathbf{t} and \mathbf{r} fixed. The weight λ_{cpl} was empirically chosen to leverage the contribution of each term into the final energy, namely $\lambda_{cpl} = 0.01$ for all experiments. Then we compute a rigid transformation [1] between the obtained registration and the template and obtain \mathcal{I}_k^i . From these registrations we compute a per vertebrae mean template \mathbf{T}_k^μ and perform PCA. We note by \mathbf{B}_k the 10 first PCA principal components of vertebra k .

2.2 PCA-guided registration

In the PCA-guided registration we first compute \mathbf{t} and \mathbf{r} for non-watertight scans $\bar{\mathcal{S}}_k^i$, by optimizing (1) and (2). We set $\lambda_{m2s} = 0$ as not every vertex in the template must explain a scan point. Then, we use the individual shape space to parametrize a mesh as $\mathbf{T}_k^\mu + \beta \mathbf{B}_k$ where β is a 10-dimensional vector. We solve for \mathbf{t} , \mathbf{r} and β to minimize

$$\begin{aligned} E(\mathcal{S}, \mathbf{T}_k^\mu, \mathbf{t}, \mathbf{r}, \beta) = & E_{p2m}(\mathcal{S}, R(\mathbf{T}_k^\mu, \mathbf{r}) + \mathbf{t} + \beta \mathbf{B}_k) + \\ & + \lambda_{m2s} E_{p2m}(R(\mathbf{T}_k^\mu, \mathbf{r}) + \mathbf{t} + \beta \mathbf{B}_k, \mathcal{S}) \quad (4) \end{aligned}$$

with $\lambda_{m2s} = 1$ if the scan is watertight and $\lambda_{m2s} = 0$ otherwise. Then we perform a free vertices optimization following (3) by coupling the vertices to the solution obtained from (4). We compute the rigid transformation [1] between the obtained registration and the mean template \mathbf{T}_k^μ and obtain the final unposed registrations \mathcal{A}_k^i . These registrations will be used to learn the full spine model.

2.3 PPCA on the registrations

For each volume index i , we do not have the full set of registration \mathcal{A}_k^i , as some k are not observed in the volume i (see Fig. 1). Thus we use Probabilistic Principal Component Analysis (PPCA) [25], a variant of PCA dealing with missing data. We construct a matrix with size $N \times S$ values, where $N = \sum_{k=1}^{24} 3N_k$, and $S = 160$ is the number of volumes used (80 volumes plus their symmetric version). Each column i is the concatenation of the 24 vectors, each given by the registration \mathcal{A}_k^i vertices, if it is available, or an empty vector if the data is missing. The created matrix has 57% missing values. We compute PPCA using the publicly available implementation in [18] and obtain the mean spine \mathbf{T}^μ and the associated shape space \mathbf{B} .

2.4 The spine model

The architecture of the spine model is a graphical model mesh architecture parametrized by the shape parameters and the pose parameters, where each vertebrae has its own translation and rotation. The shape parameters β apply additive offsets to the initial vertebrae mesh \mathbf{T}^μ . The offsets are the PCA principal components learned with the PPCA and the poses and translations are rigidly applied to each vertebra. To constrain the locations of the vertebrae, we learn a distribution over the relative positions of neighboring vertebrae. For each pair of neighboring vertebrae k and $k + 1$, we use R_k^i , t_k^i and $t_{(k+1)}^i$ to compute the relative translation of vertebra $k + 1$ w.r.t. vertebra k . Then for each 23 pairs $k, k + 1$ we fit a Gaussian model to the observed relative translations. These relative translations distribution accounts for the thickness of the inter-vertebral disk. However, as the relative rotation between vertebrae highly depends on the patient posture, we do not learn a pose prior on the relative rotations.

3 VerSe challenge dataset

To train our model we use the publicly available MICCAI Challenge 2019 VerSe dataset [21]. It provides 80 CT scans with manually annotated voxel-wise labels, and 40 CT scans without labels. We use the 80 annotated subjects to create the spine model. Most volumes have a voxel size $1 \times 1 \times 1 \text{ mm}^3$ and a few $1 \times 1 \times 2 \text{ mm}^3$. We resample all volumes to have a fixed $1 \times 1 \times 1 \text{ mm}^3$ sampling. An overview of the available observations from the full spine in the VerSe dataset is shown in Fig. 1. Each row in the figure represents a vertebrae, C1 to T5 from top to bottom. Each column represents a CT volume. As we can see, each CT scan only covers a few vertebrae and no full spine is observed in any CT volume. In addition, the number of volumes including a given vertebra is not balanced across different vertebrae. The dataset is dominated by lumbar observations and few cervical vertebrae are observed. The section including the junction between cervical and thoracic vertebrae is critically sparse.

In our work we aim at learning the shape correlations across the full spine and to evaluate our model we use cross validation. In order to make sure that the junction between the cervical and thoracic vertebrae is always seen at train time, we label a volume *junction volume* if C5 to T2 are observed in the volume, and *not junction volume* otherwise. We use this junction label to create a balanced cross validation 8-fold setting for all our experiments.

As our method is trained on a few samples, we carefully inspected the input data. We noticed some misannotations in the ground truth labels. For example, some misannotations appear as floating points far away from the corresponding vertebra, as well as lumbar volumes having cervical labels. We cleaned the ground truth meshes by keeping only the biggest connected component in the mesh extracted from the CT volume. This procedure is standard as the evaluation criteria of the VerSe dataset only considers the biggest component in the volume [21]. We also observed a few vertebrae containing metal implants. As their shape is severely biased, we manually removed them from the dataset.

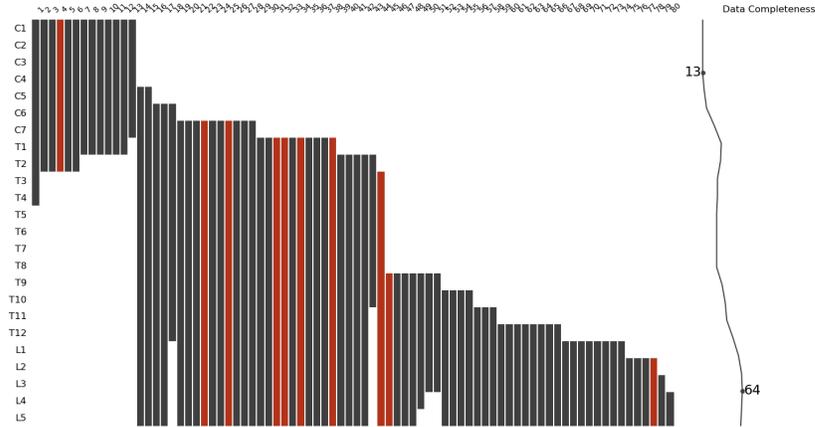


Fig. 1. Overview of the VerSe dataset. Rows represent one of the 24 individual vertebrae (C1 to L5). Column are the observed patients. No full spine is observed for any patient. The columns in red are the test set of one of the 8-fold used for cross validation.

4 Experiments

To demonstrate the accuracy and benefits from the spine model we perform two evaluations. In a first set of experiments we assess the metric accuracy of the registrations, the created model and the compacity of the learned shape space. In a second evaluation we assess the accuracy of our model to predict the shape of missing vertebrae.

4.1 Metric Accuracy

Before we learn the spine model, we first show that the registrations closely capture the input data shape with sub-voxel accuracy. Then we measure the compacity and the generalization power of the full spine model.

Individual vertebrae registrations accuracy. To assess the accuracy of the individual registrations obtained in Sec. 2.2 we compute the point to mesh distance between the scans \mathcal{S}_k^i and their registrations \mathcal{A}_k^i . For each individual vertebra we aggregate the mean distance per scan by taking the mean over all available volumes. We report the mean error and standard deviation for each vertebra in Tab. 1. We aggregate the data by weighting each vertebrae errors with the number of samples and obtained 0.24 mm errors in mean and 0.20 mm std. Let us recall that the input voxel size is $1 \times 1 \times 1 \text{ mm}^3$, thus the registrations faithfully capture the input data shape with a sub-voxel accuracy.

	me	std		me	std		me	std		me	std		me	std
C1	0.20	0.16	C2	0.20	0.16	C3	0.19	0.15	C4	0.19	0.15	C5	0.19	0.15
C6	0.18	0.15	T1	0.21	0.17	T2	0.21	0.17	T3	0.21	0.16	T4	0.21	0.16
T5	0.21	0.17	T6	0.21	0.17	T7	0.22	0.19	T8	0.23	0.19	T9	0.23	0.20
T10	0.23	0.20	T11	0.25	0.22	T12	0.25	0.21	L1	0.25	0.21	L2	0.26	0.22
L3	0.27	0.23	L4	0.28	0.25	L5	0.29	0.26						

Table 1. Vertebrae registration error. me: mean error. std: standard deviation (mm).

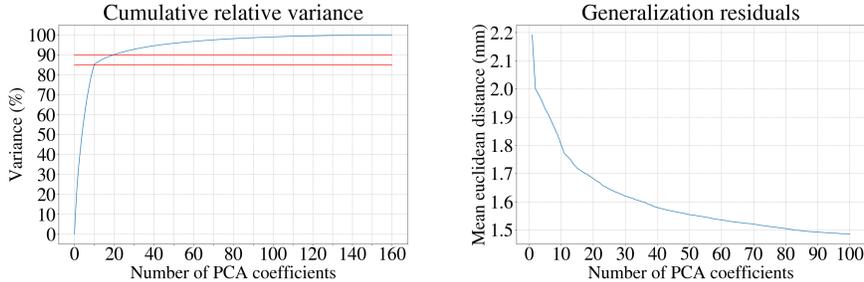


Fig. 2. Left: the cumulative variance of the full-spine model wrt the number of used shape components. Right: the generalization error of the model (cross validation) wrt to the number of used shape components.

Full-Spine model accuracy. To assess the compacity of the learned spine shape space we show on the left of Fig. 2 the cumulative variance of the shape space. The first 10 shape components capture 85% of the variance, and 90% variance is captured by the first 20. Let us recall that each registration has 127.294 vertices and 254.636 faces. To assess how well the model can capture the shape of unseen spines, we perform an 8-fold cross validation (see Sec. 3). We learn the shape space with a train set and evaluate it on the unseen test set. For an unseen, potentially partial, registration, we complete the missing vertebrae by using the learned PPCA, we then project it into the PCA space and reconstruct it with a selected number of components. We then measure the vertex to vertex error between the registration and the reconstruction. Only original data, i.e. not completed by the PPCA, is used in the metrics. On the right of Fig. 2 the generalization errors w.r.t. the number of used components are presented. Using 10 components we obtain a mean reconstruction error of 1.80 mm, and using 20 components we obtain a mean error of 1.68 mm.

4.2 Missing vertebrae shape prediction accuracy

Next, we aim at quantifying how well the model captures the global correlations across the spine. With the learned model, given a partial observation of the spine we can reconstruct the missing spine shape. We setup the following experiment to measure the accuracy of the shape of the predicted vertebrae.

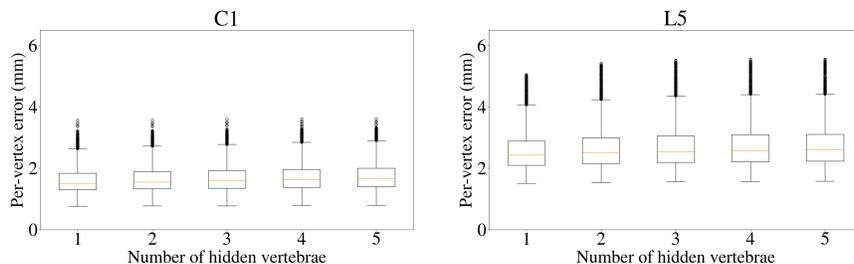


Fig. 3. Left: reconstructing C1 from partial spine observations. When one vertebra is hidden C1 is not observed, when 5 vertebrae are hidden, C1, C2, C3, C4 and C5 are not observed. Right: reconstructing L5 from partial spine observations. When one vertebra is hidden L5 is not observed, when 5 vertebrae are hidden, L5, L4, L3, L2 and L1 are not observed. The reconstruction errors slightly increase as less data is used for the shape prediction.

For a cross validation fold, we learn the spine model on the train test. Then we take all registrations obtained in Sec 2.2 in the test fold containing the cervical vertebrae C1 to C6. We first mask one vertebra starting from C1 and reconstruct it using the rest of the volume vertebrae. We then remove two vertebrae (C1 and C2) and reconstruct them using the rest. Progressively we remove cervical vertebrae until C5, and use the remaining ones for reconstruction. We measure the vertex to vertex distance between the reconstructed vertebrae and the original registration. In an analogous way, we perform the same experiment for lumbar vertebrae, by progressively removing L5 to L1.

We use the 8-fold cross validation scheme (see Sec. 3) and evaluate on all volumes for which the data is available. In Fig. 3 we report the per vertex mean error over all volumes for C1 and L5. In the left plot, the reconstruction errors on C1, when masking 1, 2, 3, 4 and 5 vertebrae are reported. In the right plot, the results for the L5 reconstruction experiment are presented. In Fig. 4 we present the per vertex mean errors on the mean template.

We aggregate the per vertex errors for all C1 and obtain reconstruction errors of 1.61 mm when one vertebra is masked, 1.65 mm when two vertebrae are masked, 1.68 mm when three vertebrae are masked, 1.72 mm when four vertebrae are masked and 1.75 mm when five vertebrae are masked. For the L5 vertebrae we obtain 2.57 mm when one vertebra is masked, 2.65 mm when two vertebrae are masked, 2.70 mm when three vertebrae are masked, 2.73 mm when four vertebrae are masked and 2.74 mm when five vertebrae are masked.

A trend in the increase of the error is observed as more vertebrae are masked, but the errors do not drastically increase. This trend can be seen in Fig. 3 as well as in the aggregated errors. These experiments show that the shape of all vertebrae is strongly correlated throughout the spine, and by observing a subset of the spine, the learned model can infer a plausible shape for the missing vertebrae.

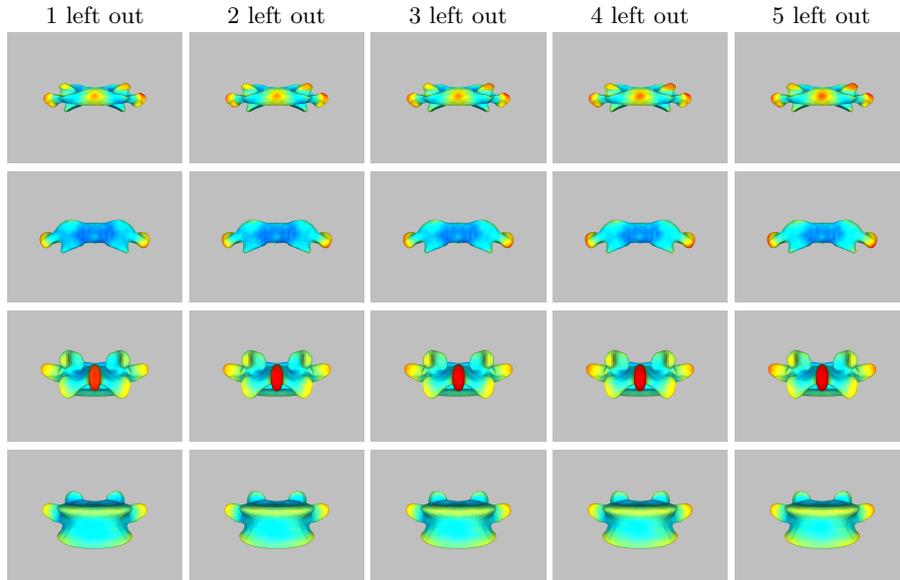


Fig. 4. Per vertex mean reconstruction error when $N \in [1,2,3,4,5]$ vertebrae are left out, visualized on the mean shape. First two rows: C1 reconstructions front and rear views. Red is 4 mm. Blue is 0mm. Third and fourth rows: L5 reconstructions front and rear views. Red is 6 mm. Blue is 0mm. Errors slightly increase as less vertebrae are used for the reconstruction.

5 Conclusion

In this paper we present a statistical surface model of the full-spine learned solely from partial and incomplete views of the spine. In order to deal with the partial observations we use probabilistic principal component analysis (PPCA) to learn a surface shape model of the full spine. The nature of the spine, specifically the correlations across the vertebrae shape, make the use of PPCA particularly well-suited. The quantitative evaluation demonstrates that the obtained model faithfully captures the shape of the population in a low dimensional space and generalizes to left out data. The model is openly available to the community for non-commercial purposes at https://gitlab.inria.fr/spine/spine_model.

While in this paper we presented an application of the model to predict the shape of missing vertebrae, future work will address the application of the model in other tasks, such as the generation of synthetic data, its registration to 2D X-ray images and its use in the segmentation and identification of vertebrae in volumetric images.

6 Acknowledgments and Disclosures

We thank Florence Forbes and Jakob Verbeek for insightful discussions. Di Meng’s work was funded by the SPINE PDCA project. Sergi Pujades work was funded by the ANR SEMBA project. Michael J. Black has received research gift funds from Intel, Nvidia, Adobe, Facebook, and Amazon. While he is a part time employee of Amazon and has financial interests in Amazon and Meshcapade GmbH, his research was performed solely at, and funded solely by, MPI.

References

1. Arun, K. Somani, Thomas S. Huang, and Steven D. Blostein. "Least-squares fitting of two 3-D point sets." *IEEE Transactions on pattern analysis and machine intelligence* 5 (1987): 698-700.
2. Blanz, Volker, and Thomas Vetter. "A morphable model for the synthesis of 3D faces." In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pp. 187-194. 1999.
3. Bogo, Federica, Michael J. Black, Matthew Loper, and Javier Romero. "Detailed full-body reconstructions of moving people from monocular RGB-D sequences." In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2300-2308. 2015.
4. Bossa, Matas N., and Salvador Olmos. "Multi-object statistical pose+ shape models." In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1204-1207. IEEE, 2007.
5. Bruners, Philipp, Tobias Penzkofer, Markus Nagel, Robert Elfring, Nina Gronloh, Thomas Schmitz-Rode, Rolf W. Gnther, and Andreas H. Mahnken. "Electromagnetic tracking for CT-guided spine interventions: phantom, ex-vivo and in-vivo results." *European radiology* 19, no. 4 (2009): 990-994.
6. Campbell, J. Q., and A. J. Petrella. "Automated finite element modeling of the lumbar spine: Using a statistical shape model to generate a virtual population of models." *Journal of biomechanics* 49, no. 13 (2016): 2593-2599.
7. Cai, Yunliang, Said Osman, Manas Sharma, Mark Landis, and Shuo Li. "Multi-modality vertebra recognition in arbitrary views using 3D deformable hierarchical model." *IEEE transactions on medical imaging* 34, no. 8 (2015): 1676-1693.
8. Castro-Mateos, Isaac, Jose M. Pozo, Marco Perea, Karim Lekadir, Aron Lazary, and Alejandro F. Frangi. "Statistical interspace models (SIMs): application to robust 3D spine segmentation." *IEEE transactions on medical imaging* 34, no. 8 (2015): 1663-1675.
9. Glocker, Ben, Johannes Feulner, Antonio Criminisi, David R. Haynor, and Ender Konukoglu. "Automatic localization and identification of vertebrae in arbitrary field-of-view CT scans." In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 590-598. Springer, Berlin, Heidelberg, 2012.
10. Heimann, Tobias, and Hans-Peter Meinzer. "Statistical shape models for 3D medical image segmentation: a review." *Medical image analysis* 13, no. 4 (2009): 543-563.
11. Hollenbeck, Justin FM, Christopher M. Cain, Jill A. Fattor, Paul J. Rullkoetter, and Peter J. Laz. "Statistical shape modeling characterizes three-dimensional shape and alignment variability in the lumbar spine." *Journal of biomechanics* 69 (2018): 146-155.

12. Kadoury, Samuel, Hubert Labelle, and Nikos Paragios. "Automatic inference of articulated spine models in CT images using high-order Markov Random Fields." *Medical image analysis* 15, no. 4 (2011): 426-437.
13. Klinder, Tobias, Jrn Ostermann, Matthias Ehm, Astrid Franz, Reinhard Kneser, and Cristian Lorenz. "Automated model-based vertebra detection, identification, and segmentation in CT images." *Medical image analysis* 13, no. 3 (2009): 471-482.
14. Korez, Robert, Bulat Ibragimov, Botjan Likar, Franjo Pernu, and Toma Vrtovec. "A framework for automated spine and vertebrae interpolation-based detection and model-based segmentation." *IEEE transactions on medical imaging* 34, no. 8 (2015): 1649-1662.
15. Kirschner, Matthias, Meike Becker, and Stefan Wesarg. "3D active shape model segmentation with nonlinear shape priors." In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 492-499. Springer, Berlin, Heidelberg, 2011.
16. Lüthi, Marcel, Thomas Albrecht, and Thomas Vetter. "Building shape models from lousy data." In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 1-8. Springer, Berlin, Heidelberg, 2009.
17. Mirzaalian Hengameh, Michael Wels, Tobias Heimann, B. Michael Kelm, and Michael Suehling. "Fast and robust 3D vertebra segmentation using statistical shape models." In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pp. 3379-3382. IEEE, 2013.
18. Probabilistic PCA python implementation <https://github.com/shergreen/pyppca>
19. Rasoulia, Abtin, Robert Rohling, and Purang Abolmaesumi. "Lumbar spine segmentation using a statistical multi-vertebrae anatomical shape+ pose model." *IEEE transactions on medical imaging* 32, no. 10 (2013): 1890-1900.
20. RuizEspaa, Silvia, Juan Domingo, Antonio DazParra, Esther Dura, Vctor D'OcnAlcaiz, Estanislao Arana, and David Moratal. "Automatic segmentation of the spine by means of a probabilistic atlas with a special focus on ribs suppression." *Medical physics* 44, no. 9 (2017): 4695-4707.
21. Sekuboyina et al. (2020), "VerSe: A Vertebrae Labelling and Segmentation Benchmark". arXiv eprint: 2001.09193. URL: arXiv:2001.09193
22. Seifert, Sascha, Adrian Barbu, S. Kevin Zhou, David Liu, Johannes Feulner, Martin Huber, Michael Suehling, Alexander Cavallaro, and Dorin Comaniciu. "Hierarchical parsing and semantic navigation of full body CT data." In *Medical Imaging 2009: Image Processing*. International Society for Optics and Photonics, 2009.
23. Schmidt, Stefan, Jrg Kappes, Martin Bergtholdt, Vladimir Pekar, Sebastian Dries, Daniel Bystrov, and Christoph Schnrr. "Spine detection and labeling using a parts-based graphical model." In *Biennial International Conference on Information Processing in Medical Imaging*, pp. 122-133. Springer, Berlin, Heidelberg, 2007.
24. Stacklies, Wolfram, Henning Redestig, Matthias Scholz, Dirk Walther, and Joachim Selbig. "pcaMethods: a bioconductor package providing PCA methods for incomplete data." *Bioinformatics* 23, no. 9 (2007): 1164-1167.
25. Tipping, Michael E., and Christopher M. Bishop. "Probabilistic principal component analysis." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61, no. 3 (1999): 611-622.
26. Zhan, Yiqiang, Dewan Maneesh, Martin Harder, and Xiang Sean Zhou. "Robust MR spine detection using hierarchical learning and local articulated model." In *International conference on medical image computing and computer-assisted intervention*, pp. 141-148. Springer, Berlin, Heidelberg, 2012.