

Probabilistic Shape Models of Anatomy Directly from Images

Jadie Adams

Scientific Computing and Imaging Institute, University of Utah, UT, USA
School of Computing, University of Utah, UT, USA
jadie@sci.utah.edu

Abstract

Statistical shape modeling (SSM) is an enabling tool in medical image analysis as it allows for population-based quantitative analysis. The traditional pipeline for landmark-based SSM from images requires painstaking and cost-prohibitive steps. My thesis aims to leverage probabilistic deep learning frameworks to streamline the adoption of SSM in biomedical research and practice. The expected outcomes of this work will be new frameworks for SSM that (1) provide reliable and calibrated uncertainty quantification, (2) are effective given limited or sparsely annotated/incomplete data, and (3) can make predictions from incomplete 4D spatiotemporal data. These efforts will reduce required costs and manual labor for anatomical SSM, helping SSM become a more viable clinical tool and advancing medical practice.

Overview

Over the last decade, deep learning approaches to shape modeling from images have proven useful in diverse applications. Statistical shape analysis is particularly beneficial in medical studies, as it allows for comparing anatomy against population-based morphological characteristics, potentially aiding in pathology detection and disease diagnosis. Computational methods allow for optimizing SSM in the form of dense sets of landmarks or correspondence points, called point distribution models (PDMs) (Cates, Elhabian, and Whitaker 2017). However, this traditional pipeline of SSM is time-consuming and entails expert-driven manual steps such as anatomy segmentation, alignment, and optimization parameter tuning. These burdens have prevented SSM from becoming a staple of medical research.

Recently, deep networks have been used to alleviate this tedious workflow by predicting PDMs directly from unsegmented 3D images. If we could safely use deep learning to accurately predict SSMs from raw medical scans, costs and manual labor required would decrease, making SSM a more usable clinical tool. This would accelerate biomedical research and result in more accurate and accessible diagnosis and pathology detection for patients. Current deep learning solutions are lacking in three ways. First, they produce over-confident estimates and lack the necessary safeguard of well-calibrated uncertainty quantification. Second, they

require ample annotated medical imaging data for training. Third, there is no solution for predicting SSM from sparse 4D spatiotemporal data. I plan to address these gaps in the current literature via three research aims.

Aim 1 - Uncertainty Quantification

I aim to predict SSM from 3D images with estimates of uncertainty that are demonstratively reliable and well-calibrated, providing the necessary safeguard for using SSM in a clinical setting. A drawback of deep learning approaches is that they can produce overconfident estimates that cannot be blindly assumed to be accurate. Such estimates can have dangerous consequences in sensitive decision-making tasks such as clinical evaluations. Quantifying what the model does not know via uncertainty measures is necessary to determine the trustworthiness of model output to prevent unsafe predictions.

DeepSSM (Bhalodia et al. 2021) is a deep learning framework that entails performing Principal Component Analysis (PCA) on the correspondence points to obtain low-dimensional shape descriptors. This serves as a prior on the convolutional network and allows for population-driven data augmentation. In my previous work, Uncertain-DeepSSM (Adams, Bhalodia, and Elhabian 2020), I extended this framework to predict data-dependent aleatoric uncertainty (by adapting the network to predict intrinsic input variance) and model-dependent epistemic uncertainty (via a Monte Carlo dropout sampling for approximate variational inference). This approach relies on PCA to impose a shape prior, introducing a linearity assumption that could taint uncertainty quantification. I addressed this shortcoming via a variational information bottleneck formulation (VIB-DeepSSM) that relaxes the linearity assumption that could be marring uncertainty predictions with off-subspace error (Adams and Elhabian 2022). This formulation provided more accurate aleatoric uncertainty measures and generalizes better under a limited training budget.

In my future work, I will adapt the VIB approach to predict epistemic uncertainty and comprehensively compare uncertainty quantification methods, as many existing Bayesian formulations fail to report quantitatively on the calibration of predicted uncertainty. My efforts toward formalizing how best to quantify the quality of uncertainty measures will be relevant beyond the task of SSM.

Aim 2 - Data Scarcity

I will alleviate the burdensome requirement of a large annotated training set, increasing the potential for medical applications where data is limited. Another well-known drawback of deep learning frameworks is that training requires ample annotated data that matches the inference domain. Data scarcity is a crucial problem regarding volumetric medical imaging. It is often difficult to access a large cohort of scans of given anatomy, especially given an uncommon disease or pathology. Segmented scans are even more scarce as they require time and domain expertise.

I will investigate two research directions to address data scarcity: model-based data augmentation and probabilistic transfer learning. DeepSSM presented a statistic-preserving data augmentation method that entails sampling from a PCA subspace. In Uncertain-DeepSSM, I enhanced this approach by using non-parametric kernel density estimation to fit the subspace, providing a more accurate representation of complex, nonlinear data. In either case, this augmentation technique imposes a linear assumption on the data by pre-projecting data on a linear subspace using PCA. To relax this restrictive assumption, I aim to explore techniques such as regularized adversarial training and few-shot learning. Data augmentation could be integrated into the training process to reduce workflow complexity and introduce challenging examples to the primary learning task. I will also explore the issue of data scarcity via Bayesian learning approaches that entail transfer learning, such as the Deep Weight Prior (Atanov et al. 2019). These approaches are fruitful to explore as they make deep models more robust to limited data and naturally allow for epistemic uncertainty quantification.

Aim 3 - Sparse 4D SSM

I aim to predict SSM from 4D spatiotemporal data with missing time points, providing the opportunity for dynamic and longitudinal research studies. Many clinical studies involve spatial and temporal information, for example, evaluating dynamic shape over an organ cycle or analyzing longitudinal changes to anatomy pre and post-surgery. While deep models that predict shape from spatiotemporal data have been explored, such as 4D segmentation from video, to the best of my knowledge there is no such model that predicts SSM from unsegmented incomplete 4D data.

As a preliminary step toward this goal, I formalized a novel SSM optimization scheme that produces landmarks that are in correspondence both across the population (inter-subject) and across time-series (intra-subject) (Adams et al. 2022). This workflow still requires the burdensome steps such as segmentation, but it provides a data-driven solution to spatiotemporal SSM. While the optimization objective disentangles subject and time correspondence, it does not fully capture their interaction as a time-sequence generative modeling approach would. In my future work, I will integrate such a generative model into the optimization scheme to inform landmark position updates. This improvement will additionally allow for inferring or imputing missing time points. By allowing for the use of incomplete sequences, fewer scans will be necessary, and partially corrupted data

will be usable, further reducing the required time and cost.

This work will provide a technique for creating training data and the theoretical foundation to explore generative deep learning approaches. Progress has been made in imputation via preprocessing with convolution-recurrent autoencoders and generative adversarial networks, but such techniques have not been applied to shape modeling. My long-term goal is to create an end-to-end 4D model that performs imputation using a latent representation when required.

Ongoing and Future Work

I have made progress toward all three aims, and I have a plan to complete the remaining work. In Fall 2022, I will improve the optimization approach to spatiotemporal SSM that I introduced in (?), allowing for sparse input data. This work will serve as a baseline for deep generative methods. In Spring 2023, I will complete the comparison of my proposed uncertainty quantification techniques and other existing state-of-the-art Bayesian methods in predicting SSM from images. In tandem, I will continue grappling with data scarcity, as it is linked to probabilistic approaches. In my final year, 2024, I will tie all this work together by exploring probabilistic deep learning solutions for predicting spatiotemporal SSM directly from sparse sequences of images.

Throughout my work, I will continue to externally validate the efficacy of my proposed methods in downstream medical tasks, demonstrating the value of accessible SSM. The collaborators I have in orthopedics, cardiac, and neuroscience research at the University of Utah provide me the potential to test on diverse datasets and tasks. I will further promote progress in AI research at large by providing open-source access to my methods via a user-friendly software interface. Upon successfully completing my research, the viability of using SSM in clinical research and onsite tools would improve, potentially impacting patient quality of life. Innovation in these directions will additionally benefit numerous other critical computer vision applications.

References

- Adams, J.; Bhalodia, R.; and Elhabian, S. Y. 2020. Uncertain-DeepSSM: From Images to Probabilistic Shape Models. *CoRR*, abs/2007.06516.
- Adams, J.; and Elhabian, S. 2022. From Images to Probabilistic Anatomical Shapes: A Deep Variational Bottleneck Approach. *arXiv preprint arXiv:2205.06862*.
- Adams, J.; Khan, N.; Morris, A.; and Elhabian, S. 2022. Spatiotemporal Cardiac Statistical Shape Modeling: A Data-Driven Approach. *arXiv preprint*.
- Atanov, A.; Ashukha, A.; Struminsky, K.; Vetrov, D.; and Welling, M. 2019. Deep Weight Prior. *arXiv:1810.06943*.
- Bhalodia, R.; Elhabian, S.; Adams, J.; Tao, W.; Kavan, L.; and Whitaker, R. 2021. DeepSSM: A Blueprint for Image-to-Shape Deep Learning Models.
- Cates, J.; Elhabian, S.; and Whitaker, R. 2017. Shapeworks: particle-based shape correspondence and visualization software. In *Statistical Shape and Deformation Analysis*, 257–298. Elsevier.