

Deep Implicit Statistical Shape Models for 3D Medical Image Delineation

Ashwin Raju^{1,2}, Shun Miao¹, Dakai Jin¹, Le Lu¹, Junzhou Huang², Adam P. Harrison^{1*}

¹ PAII Inc, Bethesda, MD, USA

² University of Texas at Arlington, Arlington, TX, USA
{ashwin.raju93, jzhuang}@uta.edu

{shwinmiao, dakai.jin, tiger.lelu, adam.p.harrison}@gmail.com

Abstract

3D delineation of anatomical structures is a cardinal goal in medical imaging analysis. Prior to deep learning, statistical shape models (SSMs) that imposed anatomical constraints and produced high quality surfaces were a core technology. Today’s fully-convolutional networks (FCNs), while dominant, do not offer these capabilities. We present deep implicit statistical shape models (DISSMs), a new approach that marries the representation power of deep networks with the benefits of SSMs. DISSMs use an implicit representation to produce compact and descriptive deep surface embeddings that permit statistical models of anatomical variance. To reliably fit anatomically plausible shapes to an image, we introduce a novel rigid and non-rigid pose estimation pipeline that is modelled as a Markov decision process (MDP). Intra-dataset experiments on the task of pathological liver segmentation demonstrate that DISSMs can perform more robustly than four leading FCN models, including nnU-Net + an adversarial prior: reducing the mean Hausdorff distance (HD) by 7.5-14.3 mm and improving the *worst case* Dice-Sørensen coefficient (DSC) by 1.2-2.3%. More critically, cross-dataset experiments on an external and highly challenging clinical dataset demonstrate that DISSMs improve the *mean* DSC and HD by 2.1-5.9% and 9.9-24.5 mm, respectively, and the *worst-case* DSC by 5.4-7.3%. Supplemental validation on a highly challenging and low-contrast larynx dataset further demonstrate DISSM’s improvements. These improvements are over and above any benefits from representing delineations with high-quality surfaces.

Introduction

3D delineation is a fundamental task in medical imaging analysis. Currently, medical segmentation is dominated by fully-convolutional networks (FCNs) (Long, Shelhamer, and Darrell 2015), which segment each pixel or voxel in a bottom-up fashion. FCNs are well-suited to the underlying convolutional neural network (CNN) technology and are straightforward to implement using modern deep learning software. The current state has been made particularly plain by the dominance of nnU-Net (Isensee et al. 2021), *i.e.*, nothing-new U-Net, in the medical segmentation decathlon (MSD) challenge (Simpson et al. 2019). Yet, de-




Figure 1: Top Panel: DISSMs produce high quality surfaces without discretization. Its use of rich and explicit anatomical priors ensure robustness, even on highly challenging cross-dataset clinical samples. In this example, nnU-Net (Isensee et al. 2021) oversegments the cardiac region (green arrow) and mishandles a TACE-treated lesion (red arrow), causing a fragmented effect. Bottom Panel: a 2D t-SNE embedding (van der Maaten and Hinton 2008) of the DISSM shape latent space. Shapes closer together share similar features.

spite their undisputed abilities, FCNs do lack important features compared to prior technology. For one, a surface-based representation is usually the desired end product, but FCNs output masks, which suffer from discretization effects (see top panel of Fig. 1). These are particularly severe when large inter-slice distances come into play. Conversion to a smoothed mesh is possible, but it introduces its own artifacts. While this is an important drawback, an arguably more critical limitation is that current FCN pipelines typically operate with no shape constraints.

Shape priors are critical for ensuring anatomically plausible delineations, but the techniques and concepts of statistical shape models (SSMs), so integral prior to deep learning (Heimann and Meinzer 2009), have fallen out of favor. Despite the incontrovertible power of FCNs, they can produce egregious mistakes, especially when presented with

*Corresponding Author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

morbidities, scanners, or other scenarios not seen in training (again see top panel of Fig. 1). Because it is impossible to represent all clinical scenarios well enough in training datasets, priors can act as valuable regularizing forces. FCNs may also struggle with anatomical structures with low-contrast boundaries. Efforts have been made to incorporate anatomical priors with CNNs, but these either do not directly model shape and/or do not estimate rigid poses. Thus, they do not construct a true SSM. Even should their interoperation with CNNs not be an issue, classic SSMs also have their own disadvantages, as they mostly rely on the point distribution model (PDM) (Cootes et al. 1992), which requires determining a correspondence across shapes. Ideally, correspondences would not be required.

To fill these gaps, we introduce deep implicit statistical shape models (DISSMs), a new and deep approach to SSMs. Using the recently introduced deep implicit shape concept (Park et al. 2019; Chen and Zhang 2019; Mescheder et al. 2019), DISSMs learn a compact and rich latent space that can accurately and densely generate the signed distance functions (SDFs) of a representative sample of shapes. Importantly, correspondence between shapes is unnecessary, eliminating a major challenge with traditional SSMs. Statistics, *e.g.*, mean shape, principle components analysis (PCA), and interpolation, can be performed directly on the latent space (see bottom panel of Fig. 1). To fit an anatomically plausible shape to a given image, DISSMs use a CNN to determine rigid and non-rigid poses, thus marrying the representation power of deep networks with anatomically-constrained surface representations. Pose estimation is modelled as a Markov decision process (MDP) to determine a trajectory of rigid and non-rigid poses, where the latter are defined as PCA loadings of the DISSM latent space. To handle the intractably large search space of poses, DISSMs make use of marginal space learning (MSL) (Zheng et al. 2008; Zheng and Comaniciu 2014) and inverted episodic training, the latter a concept we introduce. A final constrained deep level set refinement (Michalkiewicz et al. 2019) captures any fine details not represented by DISSM latent shape space. At a high level, DISSMs share many philosophies with traditional SSMs, but they modernize these concepts in a powerful deep-learning framework.

As proof of concept, we evaluate DISSM primarily on the problem of 3D pathological liver delineation from computed tomography (CT) scans, with supplemental validation on a challenging larynx segmentation task from low-contrast CT. For the former, we compare our approach to leading 2D (Harrison et al. 2017), hybrid (Li et al. 2018), 3D cascaded (Isensee et al. 2021), and adversarial learning (Yang et al. 2017) FCNs. When trained and tested on the MSD liver dataset (Simpson et al. 2019), DISSMs provide more robust delineations, improving the mean Hausdorff distance (HD) by 7.5-14.3mm. *This is over and above any benefits of directly outputting a high resolution surface.* More convincingly, we perform cross-dataset evaluation on an external dataset (97 volumes) that directly reflect clinical conditions (Raju et al. 2020a). DISSMs improve the mean Dice-Sørensen coefficient (DSC) and HD from 92.4% to 95.9% and from 34.1mm to 21.8mm, respectively, over the best

fully FCN alternative (nnU-Net). In terms of robustness, the worst-case DSC is boosted from 88.1% to 93.4%. Commensurate improvements are also observed on the larynx dataset. These results confirm the value, once taken for granted, of incorporating anatomical priors in 3D delineation. Our contributions are thus: 1) we are the first to introduce a true *deep* SSM model that outputs high resolution surfaces; 2) we build a new correspondence-free, compact, and descriptive anatomical prior; 3) we present a novel pose estimation scheme that incorporates inverted episodic training and MSL; and 4) we provide a more robust solution than leading FCNs in both intra- and cross-dataset evaluations. Our code is publicly shared¹.

Related Work

Anatomical Shape Priors: Because anatomical structures are highly constrained, delineations should match shape priors. SSMs were a workhorse for 3D medical segmentation and were most popularly realized as PDMs (Cootes et al. 1992), which requires determining a set of surface point correspondences across all shapes. Correspondence permitted statistical techniques, such as PCA, to model shape variability (Heimann and Meinzer 2009). However, dense 3D surface landmarks are not typically possible to define, and even when so, they are not always reliably present (Heimann and Meinzer 2009). The alternative of automatically generating correspondences is still a dense research topic (Heimann and Meinzer 2009), and all existing solutions are imperfect. While PDMs have been used with CNNs (Milletari et al. 2017; Bhalodia et al. 2018), these works make no attempt at determining rigid poses, thus they are only applicable in constrained and limited setups. In general the reliance on explicit surface representations, *e.g.* meshes, makes it difficult to integrate traditional SSMs with CNNs. SSMs based on implicit level set representations have been explored (Cremers, Rousson, and Deriche 2007), but the statistics must be collected across a dense 3D regular grid. The most popular approach today is to impose an implicit *prior* (different from implicit *representation*) using auto-encoders (Oktay et al. 2017; Ravishankar et al. 2017) or generative adversarial networks (Yang et al. 2017; Raju et al. 2020b; Cai et al. 2019), but these do not permit a controllable and interpretable prior and are not invariant to rigid similarity transforms, so a true shape prior is not constructed. Recent work has also used CNNs to deform an initial mesh (Yao et al. 2019; Wickramasinghe et al. 2020) or point-cloud (Cai et al. 2019) sphere, but the process is used to augment an existing voxel-based output, otherwise the deformation cannot capture all details (Wickramasinghe et al. 2020). Besides, these works offer no statistical description of anatomical priors. Like traditional SSMs, DISSMs explicitly enforce shape constraints via a model of anatomical variability, but using a deep implicit SDF *representation* that is 1) highly descriptive, 2) highly compact, and 3) requires no correspondences. DISSM is the first to offer a complete “deep” SSM solution that includes rigid and non-rigid pose estimation.

¹<https://github.com/AshStuff/dissm>

Deep Implicit Shapes: The concept of deep implicit shapes (Park et al. 2019; Chen and Zhang 2019; Mescheder et al. 2019) were introduced to represent shapes using a compact latent space and an SDF decoder. DISSMs adopt Park et al. (2019)’s auto-decoder formulation. Unlike these seminal works, DISSMs build a model of anatomical variability and propose a pose estimation strategy for 3D images.

Pose Estimation: Fitting an SSM to an image requires determining the rigid and non-rigid poses of the shape. Standard learning strategies, such as one shot regression (Štern, Ebner, and Urschler 2016; Gauriau et al. 2015) and exhaustive scanning (Zheng et al. 2015), yield sub-optimal results as the model relies on a prohibitively large number of samples (Ghesu et al. 2017) and only has one opportunity to produce pose estimations. A recent work directly regresses poses using deep implicit shapes, but these are from simulated 2D images on white backgrounds (Xu et al. 2019). DISSMs must operate in the challenging 3D medical setting, where datasets may be small and the contrast low. We formulate pose estimation as an agent-based MDP, where the agent navigates to a pose matching the present anatomical structure. This has connections to some recent registration approaches (Liao et al. 2017; Krebs 2017; Ma et al. 2017). However, estimating both non-rigid and rigid poses presents an extremely large search space. To deal with this, 1) we introduce inverted episodic training and 2) we employ a deep realization of MSL (Zheng and Comaniciu 2014; Zheng et al. 2008) to incrementally learn pose parameters. The use of MSL links DISSMs to a variety of traditional SSM works pre-dating modern deep networks.

Method

Fig. 2 illustrates the DISSM framework. As the bottom panel demonstrates, a CNN encoder predicts rigid and non-rigid poses, which, along with desired coordinates, are fed into a deep implicit multilayer perceptron (MLP) shape decoder to output corresponding SDF values. The encoder searches for the best pose using an MDP combined with MSL (top panel). We first outline the deep implicit shape model, discuss pose estimation, then describe the final local surface refinement.

Implicit Shape Representation

Deep implicit shapes (Park et al. 2019; Chen and Zhang 2019; Mescheder et al. 2019) are a recent and powerful implicit shape representation. We use Park et al. (2019)’s formulation to model organ shapes using an SDF, which, given a coordinate, outputs the distance to a shape’s surface:

$$SDF(\mathbf{x}) = s : \mathbf{x} \in \mathbb{R}^3, s \in \mathbb{R}, \quad (1)$$

where s is negative inside the shape and positive outside of it. The iso-surface of (1), *i.e.*, coordinates where it equals 0, corresponds to the shape surface. The insight of deep implicit shapes is that given a set of coordinate/SDF pairs *in some canonical or normalized space*, $\tilde{\mathcal{X}} = \{\tilde{\mathbf{x}}_i, s_i\}$, a deep MLP can be trained to approximate a shape’s SDF:

$$f_{\theta_S}(\tilde{\mathbf{x}}) \approx SDF(\tilde{\mathbf{x}}), \forall \tilde{\mathbf{x}} \in \Omega, \quad (2)$$

where the tilde denotes canonical coordinates. Because it outputs smoothly varying values (2) has no resolution limitations, and, unlike meshes, it does not rely on an explicit discretization scheme. In practice the resolution of any captured details are governed by the model capacity and the set of training samples within \mathcal{X}_i . While (2) may describe a single shape, it would not describe the anatomical variation across a set of shapes. To do this, we follow Park et al. (2019)’s elegant approach of auto-decoding latent variables. We set of K SDF samples from the same organ, $\mathcal{S} = \{\tilde{\mathcal{X}}_k\}_{k=1}^K$, and we create a corresponding set of latent vectors, $\mathcal{Z} = \{\mathbf{z}_k\}_{k=1}^K$. The deep MLP is modified to accept two inputs, $f_{\theta_S}(\tilde{\mathbf{x}}, \mathbf{z}_k)$, conditioning the output by the latent vector to specify which particular SDF is being modeled. We jointly optimize the network weights θ_S and latent vectors \mathcal{Z} to produce the best SDF approximations:

$$\arg \min_{\theta_S, \mathcal{Z}} \sum_k^K \left(\sum_{i=1}^{|\tilde{\mathcal{X}}_k|} \mathcal{L}(f_{\theta_S}(\tilde{\mathbf{x}}_i, \mathbf{z}_k), s_i) + \frac{1}{\sigma^2} \|\mathbf{z}_k\|_2^2 \right), \quad (3)$$

where \mathcal{L} is the L1 loss and the second term is a zero-mean Gaussian prior whose compactness is controlled by σ^2 .

The implicit shape model assumes shapes are specified in a canonical space. However, segmentation labels are usually given as 3D masks. To create canonical coordinate/SDF training pairs, we first perform within-slice interpolation (Albu, Beugeling, and Laurendeau 2008) on masks to remove the most egregious of discretization artifacts. We then convert the masks to meshes using marching cubes (Lewiner et al. 2003), followed by a simplification algorithm (Forstmann 2020). Each mesh is then rigidly aligned to an arbitrarily chosen anchor mesh using coherent point drift (Myronenko and Song 2010). Similar to Park et al. (2019), SDF and coordinate values are randomly sampled from the mesh, with regions near the surface much more densely sampled. SDF values are also scaled to fit within $[-1, 1]$. Based on the anchor mesh, an affine matrix that maps between canonical and pixel coordinates can be constructed, $\mathbf{x} = \mathbf{A}\tilde{\mathbf{x}}$. More details can be found in our supplementary (Raju 2021).

Once the shape decoder is trained, any latent vector can be inputted into $f_{\theta_S}(\cdot, \cdot)$ along with a set of coordinates to rasterize an SDF, which can then be rendered by extracting the iso-boundary. As the top row of Fig. 3 and bottom panel of Fig. 1 demonstrate, the latent space provides a rich description of shape variations. The mean latent vector, $\boldsymbol{\mu}$, produces an anatomically valid shape. A PCA can capture meaningful variation, *e.g.*, the first basis corresponds to stretching and flattening while the second controls the prominence of lobe protuberances. Interpolating between latent vectors produces reasonable shapes (bottom row of Fig. 3).

Pose Estimation

The next major step is use the compact and rich DISSM shape space to delineate an object boundary given an image, I . We assume a dataset of labelled images is available, allowing for the generation of coordinate/SDF pairs: $\mathcal{D} = \{I_k, \mathcal{X}_k\}_{k=1}^{K_D}$, where $\mathcal{X}_k = \{\mathbf{x}_i, s_i\}$ is specified using pixel coordinates. Note, we only assume a mask/SDF is

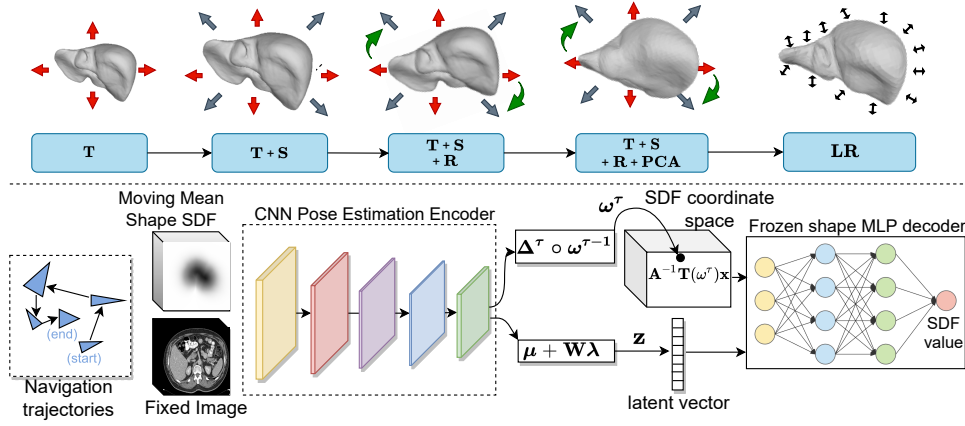


Figure 2: Overview of our deep implicit statistical shape modeling (DISSM) framework. Here **T**, **S**, **R**, **PCA**, **LR** are Translation, Scale, Rotation, PCA loadings and Local Refinement respectively.




Figure 3: DISSM shape embedding space on MSD liver dataset (Simpson et al. 2019). The top row shows the shapes generated from the mean latent vector and scaled version of the two first PCA bases (λ_1 and λ_2). The bottom row renders an interpolation between two selected latent vectors: $\mathbf{z} = (1 - \alpha)\mathbf{z}_0 + \alpha\mathbf{z}_1$.

present, and do not require *explicit* ground-truth rigid and non-rigid poses. We need to define: 1) the rigid-body transform from the canonical coordinates to the image space, and 2) the latent vector that specifies the shape variation. We denote the rigid-body transformation as $\mathbf{T}(\omega)$ with parameters $\omega = \{\mathbf{t}, \mathbf{s}, \mathbf{b}\}$, *i.e.*, translation, anisotropic scale, and rotation, respectively, where $\mathbf{t} \in \mathbb{R}^3$, $\mathbf{s} \in \mathbb{R}^3$, and $\mathbf{b} \in \mathbb{R}^6$. Here we use Zhou et al. (2019)’s recent six-dimensional parameterization of the rotation matrix, where we actually predict deviations from identity, *i.e.*, $\mathbf{I} + \mathbf{T}(\mathbf{b})$. We model shape variation using a truncated PCA basis that only captures salient variations: $\mathbf{z} = \boldsymbol{\mu} + \mathbf{W}\boldsymbol{\lambda}$. Unlike explicit SSMs, the PCA is performed on the latent space and does not require correspondences. We employ an encoder network, $g_{\theta_E}(I)$, to predict the rigid pose parameters, ω , and the non-rigid PCA loadings $\boldsymbol{\lambda}$. The parameters predicted by $g_{\theta_E}(\cdot)$ are fed into a frozen $f_{\theta_S}(\cdot, \cdot)$ to produce the object’s SDF:

$$SDF(\mathbf{x}) = f_{\theta_S}(\mathbf{A}^{-1}\mathbf{T}(\omega)\mathbf{x}, \boldsymbol{\mu} + \mathbf{W}\boldsymbol{\lambda}), \quad (4)$$

$$\omega, \boldsymbol{\lambda} = g_{\theta_E}(I). \quad (5)$$

While (4) and (5) could work in principle, directly predicting global pose parameters in one shot is highly sensitive

to any errors and we were unable to ever reach convergence. We instead interpret the encoder $g_{\theta_E}(\cdot)$ as an “agent” that, given an initial pose, ω^0 , generates samples along a trajectory by predicting corrections to the previous pose:

$$\Delta^\tau, \boldsymbol{\lambda}^\tau = g_{\theta_E}(I_k, \omega^{\tau-1}), \quad (6)$$

$$\omega^\tau = \Delta^\tau \circ \omega^{\tau-1}, \quad \text{if } \tau > 0, \quad (7)$$

where \circ denotes the composition of two rigid-body transforms and τ indicates the current step in the trajectory. An observation of $\omega^{\tau-1}$ is injected into the input of the encoder so that it is aware the previous pose to correct. To do this we rasterize the SDF corresponding to the mean shape, $SDF_{\boldsymbol{\mu}}$, *once*. After every step it is rigidly transformed using $\omega^{\tau-1}$ and fed as a second input channel into $g_{\theta_E}(\cdot, \cdot)$. The agent-based formulation turns the challenging one-step pose estimation task into a simpler multi-step correction task. Note in (6) we do not predict a trajectory for the PCA loadings. Unlike rigid pose estimation, which can use the transformed $SDF_{\boldsymbol{\mu}}$, it not clear how to best inject a concept of PCA state into the encoder without rasterizing a new SDF after every step. Since this is prohibitively expensive, the PCA loadings are directly estimated at each step. We break the search space down even further by first predicting rigid poses then predicting the PCA loadings, as detailed below.

Rigid Pose Estimation We first train the encoder $g_{\theta_E}(\cdot, \cdot)$ to predict rigid poses. In training we generate samples along a trajectory of T steps, which is referred to as an episode. The encoder is trained by minimizing a loss calculated over the episodes generated on the training data:

$$\arg \min_{\theta_E} \sum_{k=1}^{K_D} \sum_{i=1}^{|\mathcal{X}_k|} \sum_{\tau=1}^T \mathcal{L}(f_{\theta}(A^{-1}T(\omega^\tau)\mathbf{x}_i, \boldsymbol{\mu}), s_i), \quad (8)$$

where back-propagation is only executed on the encoder weights θ_E within each step τ , and the dependence on $g_{\theta_E}(\cdot)$ is implied through ω_τ . Note that (8) uses the mean latent vector, $\boldsymbol{\mu}$, to generate SDF predictions and the $\boldsymbol{\lambda}$ output in (6) is ignored for now. This training process shares similarities to deep reinforcement learning (DRL), particularly

Stage	Δ	η	ω^0
Trans.	\mathbf{t}	η_t	$\omega_{\mathcal{D}}$
Scale	$\{\mathbf{s}, \mathbf{t}\}$	$\eta_s \circ \eta'_t$	$\{\omega_{t,k}^T\}_{k=1}^{K_{\mathcal{D}}}$
Rot.	$\{\mathbf{b}, \mathbf{s}, \mathbf{t}\}$	$\eta_r \circ \eta'_s \circ \eta'_t$	$\{\omega_{s,k}^T\}_{k=1}^{K_{\mathcal{D}}}$
Non Rig.	$\{\mathbf{b}, \mathbf{s}, \mathbf{t}\}$	$\eta'_r \circ \eta'_s \circ \eta'_t$	$\{\omega_{r,k}^T\}_{k=1}^{K_{\mathcal{D}}}$

Table 1: MSL schedule used in DISSM.

in its formulation of the prediction step as an MDP. Unlike DRL, and similar to MDP registration tasks (Liao et al. 2017; Krebs 2017; Ma et al. 2017), there is no need for cumulative rewards because a meaningful loss can be directly calculated. At the start of training, the agent will not produce reliable trajectories but, as the model strengthens, the playing out of an *episode* of T steps for each training iteration will better sample meaningful states to learn from. To expose the agent to a greater set of states, we also inject random pose perturbations, η , after every step, modifying (6) and (7) to

$$\Delta^\tau, \lambda^\tau = g_{\theta_E}(I_k, \eta^\tau \circ \omega^{\tau-1}), \quad (9)$$

$$\omega^\tau = \Delta^\tau \circ \eta^\tau \circ \omega^{\tau-1}, \quad \text{if } \tau > 0. \quad (10)$$

A downside to episodic training is that each image is sampled consecutively T times, which can introduce instability and overfitting to the learning process. To avoid this we introduce *inverted episodic training*, altering the loop order to make each episodic step play out as an outer loop:

$$\arg \min_{\theta_E} \sum_{\tau=1}^T \sum_{k=1}^{K_{\mathcal{D}}} \sum_{i=1}^{|\mathcal{X}_k|} \mathcal{L}(f_{\theta}(\mathbf{A}^{-1}\mathbf{T}(\omega^\tau)\mathbf{x}_i, \boldsymbol{\mu}), s_i), \quad (11)$$

where ω^τ is saved for each sample after each iteration.

Marginal Space Learning: The MDP of (11) provides an effective sampling strategy, but it requires searching amongst all possible translation, scale, and rotation configurations, which is too large a search space. Indeed we were unable to ever reliably produce trajectories that converged. To solve this, DISSMs use a deep realization of marginal space learning (MSL) (Zheng and Comaniciu 2014). MSL decomposes the search process into a chain of dependant estimates, focusing on one set while marginalizing out the others. In practice (Tab. 1) this means that we first limit the search space by training the encoder to only predict a translation trajectory, \mathbf{t} , with the random perturbations also limited to only translation, *i.e.*, η_t . The initial pose is the mean location in the training set, denoted $\omega_{\mathcal{D}}$. Once trained, the translation encoder weights and final poses, $\{\omega_{t,k}^T\}_{k=1}^{K_{\mathcal{D}}}$, are used to initialize a scale encoder of identical architecture, but one that predicts scale corrections, \mathbf{s} , in addition to finetuning the translation. Importantly, to focus the search space on scale, the random *translation* perturbations are configured to be much smaller than before, which is represented by the prime modifier on η'_t . Finally, a rotation model is trained (while finetuning translation + scale with smaller perturbations). In inference, the rigid pose is estimated by successively applying the models of each stage, using the final pose of the previous step to initialize the next.

Non-Rigid Pose Estimation Once a rigid pose is determined, anatomically plausible deformations can then be estimated. We initialize the weights and poses of the non-rigid encoder using the translation + scale + rotation rigid model, modifying (11) to now incorporate the PCA basis:

$$\arg \min_{\theta_E} \sum_{\tau=1}^T \sum_{k=1}^{K_{\mathcal{D}}} \sum_{i=1}^{|\mathcal{X}_k|} \mathcal{L}(f_{\theta}(\mathbf{A}^{-1}\mathbf{T}(\omega^\tau)\mathbf{x}_i, \boldsymbol{\mu} + \mathbf{W}\lambda^\tau), s_i) + \frac{1}{\sigma^2} \|\boldsymbol{\mu} + \mathbf{W}\lambda^\tau\|_2^2. \quad (12)$$

As Table 1 indicates, the random rigid perturbations are configured to be small in magnitude to confine the search space to primarily the PCA loadings.

Surface Refinement

Like classic SSMs (Heimann and Meinzer 2009), non-rigid pose estimation provides a robust and anatomically plausible prediction, but it may fail to capture very fine details. We execute local refinements using an FCN model, $r = h_{\theta_R}(I, SDF_{\lambda})$, that accepts a two-channel input comprising the 3D image and the rasterized SDF after the non-rigid shape estimation. Its goal is to refine the SDF_{λ} to better match the ground truth SDF. To retain an implicit surface representation, we adapt portions of a recent deep level set loss (Michalkiewicz et al. 2019):

$$\begin{aligned} \mathcal{L}_r = & \sum_{\mathbf{x} \in \Omega_b} (SDF(\mathbf{x})^2 \cdot \delta_{\epsilon}(SDF_{\lambda}(\mathbf{x}) + r(\mathbf{x})))^{1/2} \\ & + \lambda_1 \sum_{\mathbf{x} \in \Omega_b} (|\|\nabla(SDF_{\lambda}(\mathbf{x}) + r(\mathbf{x}))\| - 1|) \\ & + \lambda_2 \sum_{\mathbf{x} \in \Omega_b} |\max(0, r(\mathbf{x}) - \rho)|, \end{aligned} \quad (13)$$

where SDF is the ground truth SDF and δ_{ϵ} is a differentiable approximation of the Dirac-delta function. The first term penalizes mismatches between the iso-boundaries of the refined SDF ground-truth. The second term ensures a unit gradient everywhere, guaranteeing that it remains a proper SDF. See Michalkiewicz et al. (2019) for more details. The third term ensures the refinement does not deviate too much from SDF_{λ} beyond a margin, ρ , otherwise it is free to deviate without penalty. Following standard level set practices, we only produce refinements within a narrow band, Ω_b , around the SDF_{λ} iso-boundary, which is also represented in the loss of (13). Finally, in addition to standard data augmentations to I , we also independently augment SDF_{λ} with random rigid and non-rigid pose variations, enriching the model’s training set.

Experiments

Liver Dataset: We focus on delineating pathological livers from venous-phase CTs. We use the size-131 training set of the MSD liver dataset (Simpson et al. 2019), splitting it randomly into training, testing, and validation using proportions of 70%, 20%, and 10%, respectively. Tumor masks were merged into the liver to create a complete pathological liver mask. Intra-dataset results are important, but we

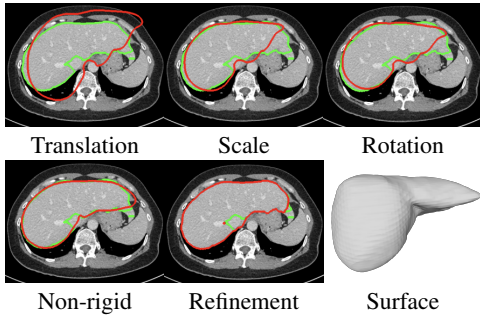


Figure 4: The DISSM pose estimation process. Red and green contours represent the prediction and ground truth, respectively.

Model	DSC (%)	ASSD (mm)	HD (mm)
2D-PHNN	95.9 ± 2.7	2.6 ± 1.2	35.7 ± 16.1
nnU-Net	96.4 ± 1.9	1.7 ± 0.9	29.1 ± 12.4
H-DenseUNet	96.3 ± 2.1	1.9 ± 1.1	30.4 ± 13.9
Adv. Shape Prior	96.0 ± 2.4	2.1 ± 1.3	28.9 ± 12.8
DISSM w/o refine	96.1 ± 0.9	1.5 ± 0.7	23.4 ± 12.2
DISSM w refine	96.5 ± 0.7	1.1 ± 0.7	21.4 ± 11.8

Table 2: Quantitative results on the MSD liver dataset.

take a further step to evaluate robustness and generalizability on unseen data from true and challenging clinical scenarios. To do this, we also evaluate on the external test set of Raju et al. (2020a), which comprises 97 venous-phase CTs. The dataset was sampled directly from a hospital archive with minimal curation and includes challenging scenarios not seen in the MSD dataset, *i.e.*, new lesion types and new co-morbidities. Demographics are also different (Asian population vs. mostly Western population in MSD). Thus, the external dataset helps reveal whether the DISSM anatomical priors can provide robustness against unavoidable new and unseen clinical scenarios.

We compare DISSM against very strong FCN alternatives. 1) 2D P-HNN (Harrison et al. 2017), used as an FCN backbone for a recent semi-supervised liver segmentation method (Raju et al. 2020a). 2) HDenseUNet (Li et al. 2018), a leader of the LiTS liver segmentation challenge (Bilic et al. 2019) that uses a cascade of 2D and 3D FCNs. 3) nnU-Net (Isensee et al. 2021), the winner of the MSD (Simpson et al. 2019) and KiTS (Heller et al. 2020) challenges. We use its dual model 3D cascade setting, which performed best on the liver task of the MSD challenge. 4) nnU-Net augmented with an adversarial anatomical prior that follows Yang et al. (2017). For its discriminator, we use Raju et al. (2020a)’s more modern version, which has already proven effective on the clinical liver dataset. For all, we use their published implementations, including recommended resolutions, pre- and post-processing, and data augmentations. For quantitative comparisons we threshold the DISSM SDF surface to produce a mask and measure the DSC, average symmetric surface distance (ASSD), and HD scores against the original masks. *Note, this is a disadvantageous setup for DISSM, as the original masks suffer from stair-like discretization ef-*




Figure 5: Box-whisker plot of DSCs on the MSD dataset.

fects, which DISSM aims to rectify.

Larynx Dataset: We also perform validation on larynx segmentation from CT, a critical organ-at-risk (OAR) for head and neck cancer radiotherapy (Brouwer et al. 2015). This task is highly challenging due to very low contrast boundaries in CT. We compare against the SOARS method and dataset of Guo et al. (2020) (142 CTs), who reported a recent and computationally intensive approach that relies on stratified learning and neural architecture search. However, SOARS is also designed to segment other OARs at the same time, so results are not apples-to-apples. For this reason, we additionally compare against nnU-Net (Isensee et al. 2021), trained only to segment the larynx. More details on this dataset can be found in the supplementary (Raju 2021).

Implementation Details: We kept settings as similar as possible for the two datasets. The shape decoder structure and hyperparameters follow that of Park et al. (2019) and we use a size 256 latent variable. For the pose encoders, $g_{\theta_E}(\cdot, \cdot)$, use the 3D encoder of a 3D U-Net (Çiçek et al. 2016), with 4 downsampling layers and global averaging pooling to produce Δ^τ and λ . For the liver we estimate the first 28 PCA components (72% of the variance). Larynx shape variations are more constrained, so we estimate the first 12 components (95% variance). The number of inverted episodic steps, T, for training the translation, scale, rotation, and non-rigid encoders was 7, 15, 15, and 15, respectively. The translation encoder was trained on a coarsely sampled volume. After it converged, we cropped volumes to encompass the maximum organ size and trained the remaining pose encoders on higher resolution volumes. Finally, we use a patch-based 3D U-Net (Çiçek et al. 2016) as the local surface refinement model, $h_{\theta_R}(\cdot, \cdot)$. Full training details, including a complete listing of all hyperparameters, can be found in the supplementary (Raju 2021).

Liver Results: DISSM consumes 12 – 13s to fully delineate a 3D volume. Fig. 4 overlays the DISSM results on top of a CT scan after each stage. As can be seen, each stage progressively improves the result. After the non-rigid PCA loading, the delineation quality is already quite high, capturing a lobe curve not represented by the mean shape. The local refinement improves results even further by capturing




Figure 6: Qualitative comparison of DISSM (second row) versus nnU-Net (first row). Red and green contours represent the prediction and ground truth delineations, respectively. The first and second-to-fourth columns are drawn from the MSD and clinical dataset, respectively. Fourth row is worst-case performance for DISSM. Blue arrows in the top row shows nnU-Net mispredictions. Results from other methods can be found in the Supplemental.

Model	DSC (%)	ASSD (mm)	HD (mm)
2D-PHNN	90.1 ± 5.1	3.9 ± 1.4	46.3 ± 21.1
nnU-Net	92.4 ± 3.3	3.6 ± 1.1	34.1 ± 17.3
HDenseUNet	92.1 ± 3.7	3.3 ± 1.3	36.2 ± 16.7
Adv. Shape Prior	93.8 ± 1.7	3.1 ± 1.1	31.7 ± 14.3
DISSM w/o refine	95.7 ± 1.8	2.6 ± 1.1	24.7 ± 12.6
DISSM w refine	95.9 ± 1.6	2.3 ± 0.9	21.8 ± 12.1

Table 3: Cross-dataset results on the clinical liver dataset.

fine-grained boundary curvatures. An ablation study can be found in the supplementary (Raju 2021).

Table 2 outlines the performance on the MSD data. As can be seen, all models perform quite well, but DISSM exhibits less variability in DSC and ASSD, indicating better robustness. This is cogently illustrated by the HD numbers, *i.e.*, the worst-case distances for each volume. DISSM dramatically improves the HD numbers by roughly 26% to 40%, resulting in much more reliable delineations. This robustness can be best seen by the box and whisker DSC plot of Fig. 5, which shows DISSM posting better worst-case and third-quartile performance. The visual impact of these improvements can be seen in the first column of Fig. 6.

While the above demonstrate that DISSM can provide superior intra-dataset robustness, the clinical results are even more telling. As Table 3 highlights, the competitor model performances drop drastically on the clinical dataset, underscoring the difficulty of operating when morbidities, scanners, patient populations, and practices can vary in unanticipated ways. In contrast, DISSM’s performance is much more stable, still posting very good numbers. Compared to the competitors, DISSM boosts the mean DSC score by 2.1 to 5.8% and reduces the HD by 31 to 53%. As the box and whisker plot of Fig. 7 shows, DISSM also provides much better worst-case performance and smaller spread, even when compared against the adversarial prior. The best FCN-only competitor, *i.e.*, nnU-Net posts a worst-case DSC performance of 88.1%, whereas DISSM’s is a much better 93.2%. The second to fourth columns of Fig. 6 illus-

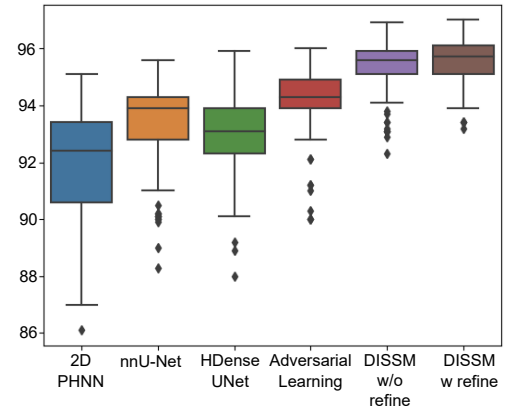


Figure 7: Box-whisker plot of DSCs on the clinical dataset.

Model	DSC (%)	HD (mm)
SOARS	56.7 ± 17.1	9.0 ± 7.1
nnU-Net	58.6 ± 14.7	8.6 ± 4.4
DISSM w/o refine	59.4 ± 9.3	7.8 ± 5.3
DISSM w refine	60.9 ± 5.9	7.1 ± 5.7

Table 4: Larynx dataset results.

trate challenging clinical examples, where nnU-Net leaked into the cardiac region, failed to segment a treated lesion, and failed to handle a patient with splenomegaly. Note, the fourth row represents the worst-case result for DISSM.

Larynx Results: Table 4 outlines mean DSC and HD scores on the larynx dataset, which is all that SOARS reported (Guo et al. 2020). As can be seen, DISSM outperforms the competitors, posting a mean DSC of 60.9%, which is roughly 2% better than the next best result (nnU-Net). More notably, the standard deviation is significantly reduced (from 14.7 to 5.9), indicating that DISSM is much more reliable. These robustness benefits on a delineation task with challenges distinct from pathological livers, *i.e.*, low contrast boundaries, provide further evidence of the value of the DISSM SSM. Qualitative examples can be found in the supplemental (Raju 2021).

Conclusion

Deep implicit statistical shape models (DISSMs) use a deep implicit model to construct statistical and correspondence-free anatomical priors and directly outputs high-quality surfaces, rather than voxelized masks. For pose estimation, DISSM proposes a robust MDP that incorporates MSL and inverted episodic training. DISSM is the first to integrate a true SSM with deep learning technology. Cross dataset evaluations on pathological liver segmentation demonstrate that DISSM outperforms leading FCNs, *e.g.*, nnU-Net (Isensee et al. 2021), improving the mean and worst-case DSC by 3.5% and 5.1%, respectively. Supplemental validation on a challenging larynx dataset further confirmed the value of DISSM. While continued maturation is necessary, DISSM represents a new and promising approach to 3D medical imaging delineation.

References

- Albu, A. B.; Beugeling, T.; and Laurendeau, D. 2008. A Morphology-Based Approach for Interslice Interpolation of Anatomical Slices From Volumetric Images. *IEEE Transactions on Biomedical Engineering*, 55(8): 2022–2038.
- Bhalodia, R.; Elhabian, S. Y.; Kavan, L.; and Whitaker, R. T. 2018. DeepSSM: A Deep Learning Framework for Statistical Shape Modeling from Raw Images. In Reuter, M.; Wachinger, C.; Lombaert, H.; Paniagua, B.; Lüthi, M.; and Egger, B., eds., *Shape in Medical Imaging*, Lecture Notes in Computer Science, 244–257. Cham: Springer International Publishing. ISBN 978-3-030-04747-4.
- Bilic, P.; et al. 2019. The Liver Tumor Segmentation Benchmark (LiTS). *CoRR*, abs/1901.04056.
- Brouwer, C. L.; Steenbakkers, R. J.; Bourhis, J.; Budach, W.; Grau, C.; Grégoire, V.; Van Herk, M.; Lee, A.; Maingon, P.; Nutting, C.; et al. 2015. CT-based delineation of organs at risk in the head and neck region: DAHANCA, EORTC, GORTEC, HKNPCSG, NCIC CTG, NCRI, NRG Oncology and TROG consensus guidelines. *Radiotherapy and Oncology*, 117(1): 83–90.
- Cai, J.; Xia, Y.; Yang, D.; Xu, D.; Yang, L.; and Roth, H. 2019. End-to-End Adversarial Shape Learning for Abdomen Organ Deep Segmentation. In Suk, H.-I.; Liu, M.; Yan, P.; and Lian, C., eds., *Machine Learning in Medical Imaging*, Lecture Notes in Computer Science, 124–132. Cham: Springer International Publishing. ISBN 978-3-030-32692-0.
- Chen, Z.; and Zhang, H. 2019. Learning Implicit Fields for Generative Shape Modeling. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5932–5941. Long Beach, CA, USA: IEEE. ISBN 978-1-72813-293-8.
- Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S. S.; Brox, T.; and Ronneberger, O. 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In Ourselin, S.; Joskowicz, L.; Sabuncu, M. R.; Unal, G.; and Wells, W., eds., *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, 424–432. Cham: Springer International Publishing. ISBN 978-3-319-46723-8.
- Cootes, T. F.; Taylor, C. J.; Cooper, D. H.; and Graham, J. 1992. Training Models of Shape from Sets of Examples. In Hogg, D.; and Boyle, R., eds., *BMVC92*, 9–18. London: Springer London. ISBN 978-1-4471-3201-1.
- Cremers, D.; Rousson, M.; and Deriche, R. 2007. A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape. *International Journal of Computer Vision*, 72(2): 195–215.
- Forstmann, S. 2020. Fast-Quadric-Mesh-Simplification. <https://github.com/sp4cerat/Fast-Quadric-Mesh-Simplification>.
- Gauriau, R.; Cuingnet, R.; Lesage, D.; and Bloch, I. 2015. Multi-organ localization with cascaded global-to-local regression and shape prior. *Medical image analysis*, 23(1): 70–83.
- Ghesu, F.-C.; Georgescu, B.; Zheng, Y.; Grbic, S.; Maier, A.; Hornegger, J.; and Comaniciu, D. 2017. Multi-scale deep reinforcement learning for real-time 3D-landmark detection in CT scans. *IEEE transactions on pattern analysis and machine intelligence*, 41(1): 176–189.
- Guo, D.; Jin, D.; Zhu, Z.; Ho, T.-Y.; Harrison, A. P.; Chao, C.-H.; Xiao, J.; and Lu, L. 2020. Organ at Risk Segmentation for Head and Neck Cancer Using Stratified Learning and Neural Architecture Search. 4223–4232.
- Harrison, A. P.; Xu, Z.; George, K.; Lu, L.; Summers, R. M.; and Mollura, D. J. 2017. Progressive and Multi-path Holistically Nested Neural Networks for Pathological Lung Segmentation from CT Images. In Descoteaux, M.; Maier-Hein, L.; Franz, A.; Jannin, P.; Collins, D. L.; and Duchesne, S., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, 621–629. Cham: Springer International Publishing. ISBN 978-3-319-66179-7.
- Heimann, T.; and Meinzer, H.-P. 2009. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis*, 13(4): 543–563.
- Heller, N.; Isensee, F.; Maier-Hein, K. H.; Hou, X.; Xie, C.; Li, F.; Nan, Y.; Mu, G.; Lin, Z.; Han, M.; et al. 2020. The state of the art in kidney and kidney tumor segmentation in contrast-enhanced CT imaging: Results of the KiTS19 Challenge. *Medical Image Analysis*, 101821.
- Isensee, F.; Jaeger, P. F.; Kohl, S. A. A.; Petersen, J.; and Maier-Hein, K. H. 2021. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2): 203–211. Number: 2 Publisher: Nature Publishing Group.
- Krebs, J. e. a. 2017. Robust non-rigid registration through agent-based action learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 344–352. Springer.
- Lewiner, T.; Lopes, H.; Vieira, A. W.; and Tavares, G. 2003. Efficient Implementation of Marching Cubes’ Cases with Topological Guarantees. *Journal of Graphics Tools*, 8(2): 1–15.
- Li, X.; Chen, H.; Qi, X.; Dou, Q.; Fu, C. W.; and Heng, P. A. 2018. H-DenseUNet: Hybrid Densely Connected UNet for Liver and Tumor Segmentation From CT Volumes. *IEEE Transactions on Medical Imaging*, 37(12): 2663–2674.
- Liao, R.; Miao, S.; de Tournemire, P.; Grbic, S.; Kamen, A.; Mansi, T.; and Comaniciu, D. 2017. An artificial agent for robust image registration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ma, K.; Wang, J.; Singh, V.; Tamersoy, B.; Chang, Y.-J.; Wimmer, A.; and Chen, T. 2017. Multimodal image registration with deep context reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 240–248. Springer.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy Networks: Learning 3D

- Reconstruction in Function Space. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4455–4465. Long Beach, CA, USA: IEEE. ISBN 978-1-72813-293-8.
- Michalkiewicz, M.; Pontes, J. K.; Jack, D.; Baktashmotlagh, M.; and Eriksson, A. 2019. Implicit Surface Representations As Layers in Neural Networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4742–4751. Seoul, Korea (South): IEEE. ISBN 978-1-72814-803-8.
- Milletari, F.; Rothberg, A.; Jia, J.; and Sofka, M. 2017. Integrating Statistical Prior Knowledge into Convolutional Neural Networks. In Descoteaux, M.; Maier-Hein, L.; Franz, A.; Jannin, P.; Collins, D. L.; and Duchesne, S., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, Lecture Notes in Computer Science, 161–168. Cham: Springer International Publishing. ISBN 978-3-319-66182-7.
- Myronenko, A.; and Song, X. 2010. Point Set Registration: Coherent Point Drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12): 2262–2275.
- Oktay, O.; Ferrante, E.; Kamnitsas, K.; Heinrich, M.; Bai, W.; Caballero, J.; Guerrero, R.; Cook, S.; de Marvao, A.; Dawes, T.; O’Regan, D.; Kainz, B.; Glocker, B.; and Rueckert, D. 2017. Anatomically Constrained Neural Networks (ACNN): Application to Cardiac Image Enhancement and Segmentation. *IEEE Transactions on Medical Imaging*, PP.
- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 165–174.
- Raju. 2021. Deep Implicit Statistical Shape Models for 3D Medical Image Delineation. *CoRR*, abs/2104.02847.
- Raju, A.; Cheng, C.-T.; Huo, Y.; Cai, J.; Huang, J.; Xiao, J.; Lu, L.; Liao, C.; and Harrison, A. P. 2020a. Co-heterogeneous and Adaptive Segmentation from Multi-source and Multi-phase CT Imaging Data: A Study on Pathological Liver and Lesion Segmentation. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J.-M., eds., *Computer Vision – ECCV 2020*, 448–465. Cham: Springer International Publishing. ISBN 978-3-030-58592-1.
- Raju, A.; Ji, Z.; Cheng, C. T.; Cai, J.; Huang, J.; Xiao, J.; Lu, L.; Liao, C.; and Harrison, A. P. 2020b. User-Guided Domain Adaptation for Rapid Annotation from User Interactions: A Study on Pathological Liver Segmentation. In Martel, A. L.; Abolmaesumi, P.; Stoyanov, D.; Mateus, D.; Zuluaga, M. A.; Zhou, S. K.; Racoceanu, D.; and Joskowicz, L., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Lecture Notes in Computer Science, 457–467. Cham: Springer International Publishing. ISBN 978-3-030-59710-8.
- Ravishankar, H.; Venkataramani, R.; Thiruvankadam, S.; Sudhakar, P.; and Vaidya, V. 2017. Learning and Incorporating Shape Models for Semantic Segmentation. In Descoteaux, M.; Maier-Hein, L.; Franz, A.; Jannin, P.; Collins, D. L.; and Duchesne, S., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, Lecture Notes in Computer Science, 203–211. Cham: Springer International Publishing. ISBN 978-3-319-66182-7.
- Simpson, A. L.; et al. 2019. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *CoRR*, abs/1902.09063.
- Štern, D.; Ebner, T.; and Urschler, M. 2016. From local to global random regression forests: Exploring anatomical landmark localization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 221–229. Springer.
- van der Maaten, L.; and Hinton, G. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86): 2579–2605.
- Wickramasinghe, U.; Remelli, E.; Knott, G.; and Fua, P. 2020. Voxel2Mesh: 3D Mesh Model Generation from Volumetric Data. In Martel, A. L.; Abolmaesumi, P.; Stoyanov, D.; Mateus, D.; Zuluaga, M. A.; Zhou, S. K.; Racoceanu, D.; and Joskowicz, L., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Lecture Notes in Computer Science, 299–308. Cham: Springer International Publishing. ISBN 978-3-030-59719-1.
- Xu, Q.; Wang, W.; Ceylan, D.; Mech, R.; and Neumann, U. 2019. DISN: Deep Implicit Surface Network for High-quality Single-view 3D Reconstruction. In *NeurIPS*.
- Yang, D.; Xu, D.; Zhou, S. K.; Georgescu, B.; Chen, M.; Grbic, S.; Metaxas, D.; and Comaniciu, D. 2017. Automatic Liver Segmentation Using an Adversarial Image-to-Image Network. In Descoteaux, M.; Maier-Hein, L.; Franz, A.; Jannin, P.; Collins, D. L.; and Duchesne, S., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, Lecture Notes in Computer Science, 507–515. Cham: Springer International Publishing. ISBN 978-3-319-66179-7.
- Yao, J.; Cai, J.; Yang, D.; Xu, D.; and Huang, J. 2019. Integrating 3D Geometry of Organ for Improving Medical Image Segmentation. In Shen, D.; Liu, T.; Peters, T. M.; Staib, L. H.; Essert, C.; Zhou, S.; Yap, P.-T.; and Khan, A., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Lecture Notes in Computer Science, 318–326. Cham: Springer International Publishing. ISBN 978-3-030-32254-0.
- Zheng, Y.; Barbu, A.; Georgescu, B.; Scheuering, M.; and Comaniciu, D. 2008. Four-Chamber Heart Modeling and Automatic Segmentation for 3-D Cardiac CT Volumes Using Marginal Space Learning and Steerable Features. *IEEE Transactions on Medical Imaging*, 27(11): 1668–1681.
- Zheng, Y.; and Comaniciu, D. 2014. *Marginal space learning for medical image analysis*. Springer.
- Zheng, Y.; Liu, D.; Georgescu, B.; Nguyen, H.; and Comaniciu, D. 2015. 3D deep learning for efficient and robust landmark detection in volumetric data. In *International conference on medical image computing and computer-assisted intervention*, 565–572. Springer.
- Zhou, Y.; Barnes, C.; Lu, J.; Yang, J.; and Li, H. 2019. On the Continuity of Rotation Representations in Neural Networks. In *Proceedings of CVPR*, 9. Long Beach, CA: IEEE.