

Integrating Spatial Configuration into Heatmap Regression Based CNNs for Landmark Localization

Christian Payer^{1,2}

CHRISTIAN.PAYER@ICG.TUGRAZ.AT

¹ *Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria*

² *Ludwig Boltzmann Institute for Clinical Forensic Imaging, Graz, Austria*

Darko Štern²

DARKO.STERN@CFI.LBG.AC.AT

Horst Bischof¹

BISCHOF@ICG.TUGRAZ.AT

Martin Urschler^{2,1}

MARTIN.URSCHLER@CFI.LBG.AC.AT

Abstract

In many medical image analysis applications, often only a limited amount of training data is available, which makes training of convolutional neural networks (CNNs) challenging. In this work on anatomical landmark localization, we propose a CNN architecture that learns to split the localization task into two simpler sub-problems, reducing the need for large training datasets. Our fully convolutional SpatialConfiguration-Net (SCN) dedicates one component to locally accurate but ambiguous candidate predictions, while the other component improves robustness to ambiguities by incorporating the spatial configuration of landmarks. In our experimental evaluation, we show that the proposed SCN outperforms related methods in terms of landmark localization error on size-limited datasets.

Keywords: anatomical landmarks, localization, heatmap regression, spatial configuration

1. Introduction

Localization of anatomical landmarks is an important step in medical image analysis, e.g., in segmentation (Beichel et al., 2005), or registration (Johnson and Christensen, 2002). Unfortunately, locally similar structures often introduce difficulties due to ambiguity into landmark localization. To deal with these difficulties, machine learning based approaches often combine local landmark predictions with explicit handcrafted graphical models, aiming to restrict predictions to feasible spatial configurations. Thus, the landmark localization problem is simplified by separating the task into two successive steps. The first step is dedicated to locally accurate but potentially ambiguous predictions, while in the second step graphical models (Cootes et al., 1995; Felzenszwalb and Huttenlocher, 2005) eliminate ambiguities.

Recent advances in computer vision and medical imaging have mainly been driven by convolutional neural networks (CNNs) due to their superior capabilities to automatically learn important image features (LeCun et al., 2015). Unfortunately, CNNs typically need large amounts of training data. Especially in medical imaging, this requirement is hard to fulfill, due to ethical and financial concerns as well as time consuming expert annotations.

In this work, we show that the amount of required training data can be reduced with our proposed two-component SpatialConfiguration-Net (SCN), which follows the idea of handcrafted graphical models to split landmark localization into two successive steps. This

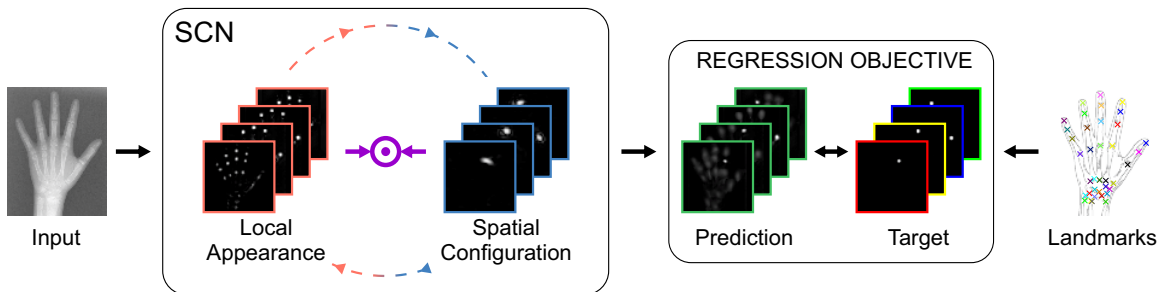


Figure 1: Landmark localization by regressing heatmaps for each landmark in our end-to-end trained fully convolutional SpatialConfiguration-Net (SCN).

extended abstract gives a short overview of the key concepts of our journal paper published in (Payer et al., 2019), while we refer the reader to the full paper for more detailed descriptions and more extensive evaluations on a variety of datasets.

2. Method

Our method for landmark localization is based on regressing heatmap images (Tompson et al., 2014), which encode the pseudo-probability of a landmark being located at a certain pixel position. With N being the total number of landmarks, we define the target heatmap image of a landmark L_i , $i = \{1, \dots, N\}$ as the d -dimensional Gaussian function $g_i(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ centered at the target landmark’s groundtruth coordinate $\mathbf{x}_i^* \in \mathbb{R}^d$.

The network is set up to regress N heatmaps *simultaneously* by minimizing the differences between predicted heatmaps $h_i(\mathbf{x})$ and the corresponding target heatmaps $g_i(\mathbf{x})$ in an end-to-end manner (Ronneberger et al., 2015; Shelhamer et al., 2017). In network inference, we obtain the predicted coordinate $\hat{\mathbf{x}}_i \in \mathbb{R}^d$ of each landmark L_i by taking the coordinate, where the heatmap has its highest value.

2.1. SpatialConfiguration-Net

The fundamental concept of the SpatialConfiguration-Net (SCN) is the interaction between its two components (see Fig. 1). The first component takes the image as input to generate locally accurate but potentially ambiguous *local appearance* heatmaps $h_i^{\text{LA}}(\mathbf{x})$. Motivated by handcrafted graphical models for eliminating these potential ambiguities, the second component takes the predicted candidate heatmaps $h_i^{\text{LA}}(\mathbf{x})$ as input to generate inaccurate but unambiguous *spatial configuration* heatmaps $h_i^{\text{SC}}(\mathbf{x})$.

For N landmarks, the set of predicted heatmaps $\mathbb{H} = \{h_i(\mathbf{x}) \mid i = 1 \dots N\}$ is obtained by element-wise multiplication \odot of the corresponding heatmap outputs $h_i^{\text{LA}}(\mathbf{x})$ and $h_i^{\text{SC}}(\mathbf{x})$ of the two components:

$$h_i(\mathbf{x}) = h_i^{\text{LA}}(\mathbf{x}) \odot h_i^{\text{SC}}(\mathbf{x}). \quad (1)$$

This multiplication is crucial for the SCN, as it forces both of its components to generate a response on the location of the target landmark \mathbf{x}_i^* , i.e., both $h_i^{\text{LA}}(\mathbf{x})$ and $h_i^{\text{SC}}(\mathbf{x})$ deliver

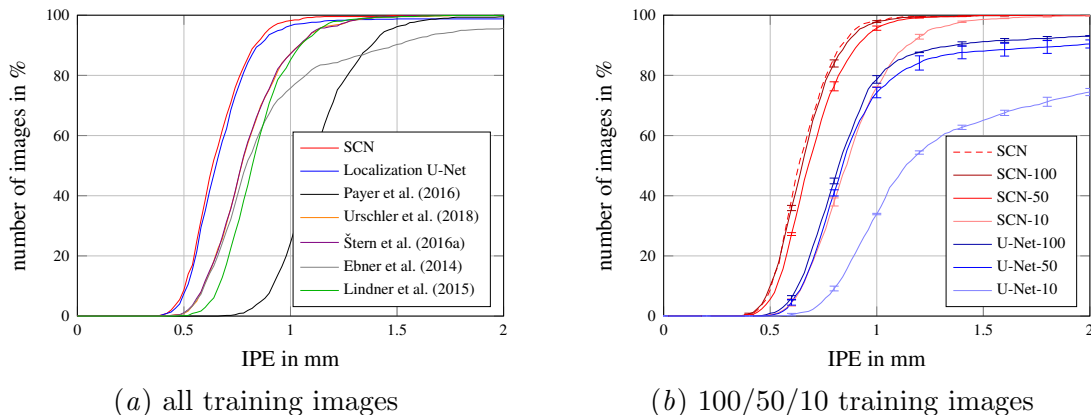


Figure 2: Cumulative distributions of the point-to-point error for 895 radiographs. (a) shows results compared with other state-of-the-art methods. (b) shows results of SCN and localization U-Net for reduced numbers of training images.

responses for \mathbf{x} close to \mathbf{x}_i^* , while on all other locations one component may have a response as long as the other one does not have one.

3. Experiments and Results

We evaluate our proposed SCN on a dataset of 895 radiographs of left hands with 37 annotated characteristic landmarks on finger tips and bone joints. We compare our SCN to state-of-the-art random regression forests (Ebner et al., 2014; Lindner et al., 2015; Štern et al., 2016; Urschler et al., 2018), our previous CNN-based method of (Payer et al., 2016), and our implementation of a localization U-Net for heatmap regression. Results of the image-specific point-to-point errors for three-fold cross validation of the 895 radiographs are shown in Fig. 2. When using all training images, our SCN outperforms all other compared methods. Additionally, when drastically reducing the number of training images to 100, 50, and 10, respectively, our SCN greatly outperforms the localization U-Net. This confirms that splitting the localization task into predicting accurate but potentially ambiguous *local appearance* heatmaps and inaccurate but unambiguous *spatial configuration* heatmaps is especially useful when dealing with only limited amounts of training data.

4. Conclusion

In conclusion, we have shown how to combine information of *local appearance* and *spatial configuration* into a single end-to-end trained network for landmark localization. Our generic architecture achieves state-of-the-art results in terms of localization error, even when only limited amounts of training images are available. We are currently looking into extending our SCN regarding occluded structures and multi-object localization, and into adapting our SCN for semantic segmentation problems (see (Payer et al., 2018) for preliminary results), where structural constraints may be used in a similar manner.

References

- Reinhard Beichel, Horst Bischof, Franz Leberl, and Milan Sonka. Robust Active Appearance Models and Their Application to Medical Image Analysis. *IEEE Trans. Med. Imaging*, 24(9):1151–1169, sep 2005. doi: 10.1109/TMI.2005.853237.
- Tim F. Cootes, Christopher J. Taylor, David H. Cooper, and Jim Graham. Active Shape Models-Their Training and Application. *Comput. Vis. Image Underst.*, 61(1):38–59, jan 1995. doi: 10.1006/cviu.1995.1004.
- Thomas Ebner, Darko Štern, René Donner, Horst Bischof, and Martin Urschler. Towards Automatic Bone Age Estimation from MRI: Localization of 3D Anatomical Landmarks. In *Proc. Med. Image Comput. Comput. Interv.*, pages 421–428. Springer, 2014. doi: 10.1007/978-3-319-10470-6_53.
- Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Pictorial Structures for Object Recognition. *Int. J. Comput. Vis.*, 61(1):55–79, 2005. doi: 10.1023/B:VISI.0000042934.15159.49.
- Hans J. Johnson and Gary E. Christensen. Consistent Landmark and Intensity-Based Image Registration. *IEEE Trans. Med. Imaging*, 21(5):450–461, 2002. doi: 10.1109/TMI.2002.1009381.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep Learning. *Nature*, 521(7553):436–444, 2015. doi: 10.1038/nature14539.
- Claudia Lindner, Paul A. Bromiley, Mircea C. Ionita, and Tim F. Cootes. Robust and Accurate Shape Model Matching Using Random Forest Regression-Voting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1862–1874, sep 2015. doi: 10.1109/TPAMI.2014.2382106.
- Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Regressing Heatmaps for Multiple Landmark Localization Using CNNs. In *Proc. Med. Image Comput. Comput. Interv.*, pages 230–238. Springer, 2016. doi: 10.1007/978-3-319-46723-8_27.
- Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Multi-label Whole Heart Segmentation Using CNNs and Anatomical Label Configurations. In *Stat. Atlases Comput. Model. Hear. ACDC MMWHS Challenges. STACOM 2017.*, pages 190–198. Springer, 2018. doi: 10.1007/978-3-319-75541-0_20.
- Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Integrating Spatial Configuration into Heatmap Regression Based CNNs for Landmark Localization. *Med. Image Anal.*, 54:207–219, may 2019. doi: 10.1016/j.media.2019.03.007.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proc. Med. Image Comput. Comput. Interv.*, pages 234–241. Springer, 2015. doi: 10.1007/978-3-319-24574-4_28.
- Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):640–651, apr 2017. doi: 10.1109/TPAMI.2016.2572683.

Darko Štern, Thomas Ebner, and Martin Urschler. From Local to Global Random Regression Forests: Exploring Anatomical Landmark Localization. In *Proc. Med. Image Comput. Comput. Interv.*, pages 221–229. Springer, 2016. doi: 10.1007/978-3-319-46723-8_26.

Jonathan Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint Training of a Convolutional Network and a Graphical Model for Human Pose Estimation. In *Adv. Neural Inf. Process. Syst.*, pages 1799–1807, 2014.

Martin Urschler, Thomas Ebner, and Darko Štern. Integrating Geometric Configuration and Appearance Information into a Unified Framework for Anatomical Landmark Localization. *Med. Image Anal.*, 43:23–36, jan 2018. doi: 10.1016/j.media.2017.09.003.