

Marcel Bengs*, Nils Gessert and Alexander Schlaefer

4D spatio-temporal convolutional networks for object position estimation in OCT volumes

<https://doi.org/10.1515/cdbme-2020-0001>

Abstract: Tracking and localizing objects is a central problem in computer-assisted surgery. Optical coherence tomography (OCT) can be employed as an optical tracking system, due to its high spatial and temporal resolution. Recently, 3D convolutional neural networks (CNNs) have shown promising performance for pose estimation of a marker object using single volumetric OCT images. While this approach relied on spatial information only, OCT allows for a temporal stream of OCT image volumes capturing the motion of an object at high volumes rates. In this work, we systematically extend 3D CNNs to 4D spatio-temporal CNNs to evaluate the impact of additional temporal information for marker object tracking. Across various architectures, our results demonstrate that using a stream of OCT volumes and employing 4D spatio-temporal convolutions leads to a 30% lower mean absolute error compared to single volume processing with 3D CNNs.

Keywords: convolutional neural networks; optical coherence tomography; position estimation; spatio-temporal data.

Introduction

Minimally invasive surgery (MIS) enables fewer post-operative complications compared to open surgery, by significantly reducing the access incisions and surgery trauma [1]. However, performing MIS is a challenging task, due to a limited field of view and lacking perception of force feedback, which requires computer-assisted surgery, particularly precise surgery tool tracking. In this regard, several vision-based approaches using images and videos

have been proposed [2]. While 2D images and videos only provide 2D spatial information, typical tissue structures and object movements are inherently three dimensional. Therefore, for many medical applications using volumetric imaging is preferable or required, e.g. for prostate radiation therapy [3], or for precise pose estimation of a marker object [4]. Some modalities provide not only volumetric images, but also allow for imaging with a high temporal resolution, such as optical coherence tomography (OCT), and hence can be used as an imaging modality for an optical tracking system [5, 6].

To overcome the limitations of classical tracking approaches, which rely on handcrafted features limited to specific application scenarios such as skin [6] or eye motion tracking [7], deep learning has been proposed recently. In particular, 3D convolutional neural networks (CNNs) have shown promising results for precisely localizing small objects based on OCT-data [4]. This approach employed 3D CNNs on a single volumetric image, allowing to turn arbitrary small objects into a marker for pose estimation. However, as OCT allows for a temporal stream of OCT image volumes, it seems reasonable that the preceding image volumes at high volume rates may carry information on the object's motion. This leads to the challenging problem of 4D deep learning, which is largely unexplored so far and has only been addressed in a few applications such as functional magnetic resonance imaging [8], computed tomography [9] and OCT-based force estimation [10] as well as OCT-based tissue motion estimation [11].

In this work, we systematically extend 3D CNNs to 4D spatio-temporal data processing and evaluate whether a stream of OCT volumes improves object position estimation performance. Spatio-temporal processing with CNNs can be done by stacking multiple frames into the channel dimension [12], or by using full or factorized spatio-temporal convolutions [13, 14]. Even though these methods have shown promising performance for video analysis tasks [12, 13, 14], it is largely unclear how CNNs perform with 4D, as they have not been systematically studied. Therefore, we evaluate four widely used CNN architectures and consider several different types of convolutions for 4D data processing. We employ volume stacking, factorized, and full spatio-temporal

Marcel Bengs and Nils Gessert were contributed equally.

*Corresponding author: Marcel Bengs, Institute of Medical Technology and Intelligent Systems, Hamburg University of Technology, Hamburg, Germany, E-mail: marcel.bengs@tuhh.de

Nils Gessert and Alexander Schlaefer, Institute of Medical Technology and Intelligent Systems, Hamburg University of Technology, Hamburg, Germany

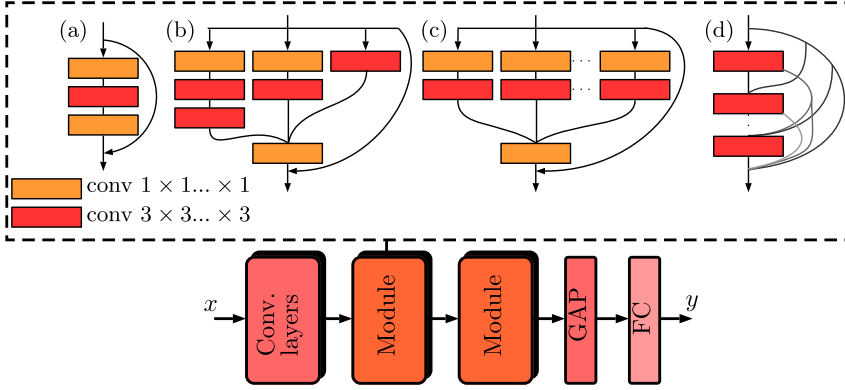


Figure 1: Each of our custom architecture consists of an initial part with five convolutional layers, followed by architecture modules that represent subsequent architecture blocks. Note, the first block in each module downsamples the input dimensions by a factor of two. The different architecture blocks are (a) ResNet, (b) Inception (c), ResNeXt, and (d) Densenet. We use a global average pooling (GAP) layer after the last module, and the output is directly fed into a fully connected output layer (FC).

convolutions, and compare the position estimation performance to single volume processing. For systematic evaluation of our methods, we consider the problem of position estimation of a marker object, with a specialized OCT setup which enables fast acquisition of sufficient 4D data with a well-defined ground-truth.

Materials and methods

Network architectures

We evaluate four different methods with four different architectures to predict the current position of a marker object using a stream of OCT volumes. Similar to a previous approach [4], we define our own architectures following the architecture principles of four widely used state-of-the-art architectures, ResNet, Inception, ResNeXt, and Densenet. Each of our custom architecture consists of an initial part with five convolutional layers, followed by architecture modules, shown in Figure 1. Note, the number of building blocks inside the modules are tuned based on validation performance. For each architecture, we evaluate four different types of convolutions, see Figure 2.

First, we consider a previous approach on marker object tracking [4], and use 3D convolutions applied to single volumetric images, which is our baseline. (3D)

Second, we stack multiple consecutive volumes into the channel dimension of the network’s input and use a 3D convolution. (3D-C)

Third, we examine factorized spatio-temporal convolutions [14], which split a full spatio-temporal convolution into a temporal and a spatial convolution. Every single spatio-temporal convolution is replaced by two successive factorized 4D convolutions. Note, there are no native implementation of 4D operations available for standard libraries such as Tensorflow or PyTorch. Hence, we implement our custom 4D convolution and pooling operation in Tensorflow, using multiple native 3D convolution and pooling operations across multiple time-shifted volumes. (F-4D)

Fourth, we consider 4D spatio-temporal convolutions and replace each 3D convolution and 3D pooling with the corresponding 4D counterparts. (4D)

The networks are trained for 350 epochs with a batch size of 18 and Adam for optimization of the mean squared error (MSE) loss function.

Data set

For data acquisition, we use a commercially available swept-source OCT device (OMES, OptoRes), a second scanning stage with two mirror galvanometers, an achromatic lens, a marker object, and a holder for the marker object. The marker object is made of a polyoxymethylene block with a size of 1 mm^3 . The whole setup is shown in Figure 3. We consider volumes with a size of $32 \times 32 \times 32$ with a corresponding field of view (FOV) of $3 \times 3 \times 3.5 \text{ mm}$, and an acquisition speed of 833 volumes per second. Our OCT setup is enhanced with a second scanning stage with two mirror galvanometers controlled by stepper motors, which enable to shift the FOV in the lateral dimensions. Also, a third motor shifts the FOV in the axial dimension, by setting the pathlength of the OCT’s reference arm. In this way, our OCT-setup allows for shifting the FOV in all spatial directions without moving the scan head. This can be utilized for automatic OCT volume acquisition and ground-truth annotation. In particular, instead of moving the

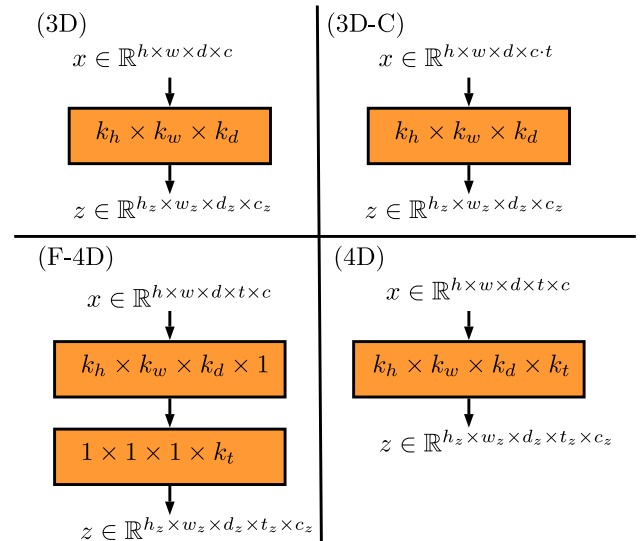


Figure 2: The different convolutions we employ: (3D) 3D spatial convolution; (3D-C) 3D convolution, with temporal information stacked into the channel dimension; (F-4D) Factorized 4D spatio-temporal convolution; (4D) 4D spatio-temporal convolution.



Figure 3: The experimental setup: Marker object (left); OCT setup (right). The marker object is attached to a holder.

marker object, we move the FOV of the OCT and the current motor positions represent the relative marker position in the FOV.

Next, we repeat the following steps and define a set of target motor positions that shift the FOV, representing smooth marker movements. First, a set of 60 to 90 target positions n_j are randomly generated for the three stepper motors. Then, piecewise cubic spline interpolation $f: \mathbb{R}_+ \rightarrow \mathbb{R}^3, \tau \mapsto f(\tau)$ is used to obtain a smooth function connecting the target positions, $f(\tau_j) = n_j$. Afterward, 500 motor points are sampled from the piecewise cubic spline function $f(\tau)$ with equidistant parameter values τ . Note, this does not lead to equidistant data points, due to the curvature of the spline function. We repeat this procedure, to obtain the full data set with 7,000 examples. Afterward, we acquire one volumetric image for each target motor positions, that serve directly as ground truth annotation. In summary, we use 5000 volumes for training and 1000 each for validating and testing our models. For our experiments, we evaluate a sequences length of five consecutive volumes. The corresponding target $t \in \mathbb{R}^3$ refers to the last position in one sequence.

Results

We report the mean absolute error (MAE) and relative mean absolute error (rMAE) for our experiments in Table 1. The MAE is given in μm based on a calibration between galvo motor steps and image coordinates. The rMAE is relative to the target’s standard deviation. Overall, using temporal data improves performance for all architectures, while 4D spatio-temporal convolutions perform best. On average the inference times are 6 ms and 20 ms for 3D and 4D architectures, respectively.

Discussion and conclusion

Our results in Table 1 show that using a sequence of volumes consistently outperforms single volume usage. This agrees with our expectation that a temporal stream of volumetric images should improve position estimates. Analyzing the different types of temporal processing shows that increasing complexity of the 4D image processing results in better predictions. Stacking the volume sequence in the channel dimension already improves performance by 15% on average compared to using a single volumetric input, while the number of parameters remain similar. This indicates that

Table 1: Comparison of the different architectures with the different types of convolutions.

Type	MAE (μm)	rMAE	Parameters
ResNet			
3D	15.87 \pm 14.40	0.013 \pm 0.011	409755
3D-C	13.39 \pm 10.96	0.011 \pm 0.009	411483
F-4D	12.36 \pm 10.15	0.010 \pm 0.008	454575
4D	11.79 \pm 9.79	0.009 \pm 0.008	1137459
Inception			
3D	17.65 \pm 15.48	0.014 \pm 0.012	428521
3D-C	14.83 \pm 11.84	0.012 \pm 0.009	430249
F-4D	13.23 \pm 11.36	0.010 \pm 0.009	475006
4D	11.87 \pm 9.66	0.009 \pm 0.008	1161568
ResNeXt			
3D	16.96 \pm 15.53	0.013 \pm 0.012	392367
3D-C	13.00 \pm 12.16	0.010 \pm 0.010	394095
F-4D	12.32 \pm 10.99	0.010 \pm 0.009	432903
4D	11.87 \pm 10.93	0.009 \pm 0.009	1012215
Densenet			
3D	16.03 \pm 13.69	0.013 \pm 0.011	406723
3D-C	14.39 \pm 11.57	0.011 \pm 0.009	445139
F-4D	12.51 \pm 10.07	0.010 \pm 0.008	420205
4D	11.54 \pm 9.51	0.009 \pm 0.008	1080683

even with processing only at the network’s input, valuable temporal information can be extracted. Note that temporal information is lost after the first convolution operation, because no temporal convolutional operation is performed [13]. Using 4D factorized convolutions instead improves performance by 25% on average compared to using a single volumetric input, and the number of parameters is only increased by less than 11%. This shows that the 4D data structure can be leveraged by factorized convolutions similar to previous findings on 3D spatio-temporal data [14]. Finally, full 4D spatio-temporal convolutions lead to the best performance, demonstrating that 4D CNNs are able to extract valuable spatio-temporal features from 4D data. Moreover, our methods perform consistently across different network architectures. Notably, the type of network architecture only has a minor impact on the errors, while Densenet results in the lowest overall error. The more costly 4D convolutions also affect inference times, which would be important for real-time tracking. While the 3D CNNs can provide position estimates

with up to 166 Hz, our 4D CNNs still achieve 50 Hz. Considering that there are no optimized 4D operations available yet, these results are promising for real-time applications such as motion tracking. Overall, we provide a comprehensive study of 4D spatio-temporal CNNs in comparison to their 3D counterparts and demonstrate that position estimations of an object can be improved significantly when a stream of volumes is used. As our methods are generic, they can be easily transferred to other tasks or imaging modalities where sequences of volumetric images are of interest, e.g., motion tracking based on volumetric ultrasound or magnetic particle imaging.

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission.

Research funding: None declared.

Competing interest: Authors state no conflict of interest.

Informed consent: Informed consent is not applicable.

Ethical approval: The conducted research is not related to either human or animals use.

References

1. Dogangil G, Davies B, Rodriguez y Baena F. A review of medical robotics for minimally invasive soft tissue surgery. *Proc Inst Mech Eng Pt H J Eng Med* 2010;224:653–79.
2. Bouget D, Allan M, Stoyanov D, Jannin P. Vision-based and markerless surgical tool detection and tracking: a review of the literature. *Med Image Anal* 2017;35:633–54.
3. Chinnaiyan P, Tomé W, Patel R, Chappell R, Ritter M. 3D-ultrasound guided radiation therapy in the post-prostatectomy setting. *Technol Canc Res Treat* 2003;2:455–8.
4. Gessert N, Schlüter M, Schlaefer A. A deep learning approach for pose estimation from volumetric OCT data. *Med Image Anal* 2018; 46:162–79.
5. Schlüter M, Otte C, Saathoff T, Gessert N, Schlaefer A. Feasibility of a markerless tracking system based on optical coherence tomography. In: *Medical imaging 2019: image-guided procedures, robotic interventions, and modeling: SPIE*; 2019, vol 10951:1095107.
6. Laves MH, Schoob A, Kahrs LA, Pfeiffer T, Huber R, Ortmaier T. Feature tracking for automated volume of interest stabilization on 4d-oct images. In: *Medical imaging 2017: image-guided procedures, robotic interventions, and modeling: SPIE*; 2017, vol 10135:101350W.
7. Camino A, Zhang M, Gao SS, Hwang TS, Sharma U, Wilson DJ, et al. Evaluation of artifact reduction in optical coherence tomography angiography with real-time tracking and motion correction technology. *Biomed Optic Express* 2016;7: 3905–15.
8. Bengs M, Gessert N, Schlaefer A. 4D spatio-temporal deep learning with 4D fMRI data for autism spectrum disorder classification. In: *Medical Imaging with Deep Learning, MIDL 2019 Conference*; 2019:1–4. <http://hdl.handle.net/11420/4299>, <https://doi.org/10.15480/882.2732>.
9. Clark D, Badea C. Convolutional regularization methods for 4d, x-ray ct reconstruction. In: *Medical imaging 2019: physics of medical imaging: International Society for Optics and Photonics*; 2019, vol 10948:109482A.
10. Gessert N, Bengs M, Schlüter M, Schlaefer A. Deep learning with 4d spatio-temporal data representations for oct-based force estimation. *Med Image Anal* 2020:101730. <https://doi.org/10.1016/j.media.2020.101730>.
11. Bengs M, Gessert N, Schlüter M, Schlaefer A. Spatio-temporal deep learning methods for motion estimation using 4D OCT image data. *Int J CARS* 2020;15:943–52.
12. Pfister T, Simonyan K, Charles J, Zisserman A. Deep convolutional neural networks for efficient pose estimation in gesture videos. In: *Asian Conference on Computer Vision: Springer*; 2014:538–52 pp.
13. Tran D, Bourdev L, Fergus R, Torresani L, Paluri M. Learning spatiotemporal features with 3D convolutional networks. In: *ICCV*; 2015:4489–97 pp.
14. Qiu Z, Yao T, Mei T. Learning spatio-temporal representation with pseudo-3D residual networks. In: *ICCV. IEEE*; 2017: 5534–42 pp.