BronchoPose: an analysis of data and model configuration for vision-based bronchoscopy pose estimation

Juan Borrego-Carazo[a,b,*], Carles Sanchez[a], David Castells-Rufas[b], Jordi Carrabina[b], Débora Gil[a]

[a]*Computer Vision Center, Universitat Autònoma de Barcelona, 08193 Cerdanyola del Vallès, Spain*
[b]*Microelectronics and Electronic Systems Department, Universitat Autònoma de Barcelona, 08193 Cerdanyola del Vallès, Spain*

Abstract

Vision-based bronchoscopy (VB) models require the registration of the virtual lung model with the frames from the videobronchoscopy to provide an effective guidance during the biopsy. The registration can be achieved by either tracking the position and orientation of the bronchoscopy camera, or by calibrating its deviation from the pose (position and orientation) simulated in the virtual lung model. Recent advances in neural networks and temporal image processing have provided new opportunities for guided bronchoscopy. However, such progress has been hindered by the lack of comparative experimental conditions.

In the present paper we share a novel synthetic dataset allowing for a fair comparison of methods. Moreover, this paper investigates several neural network architectures for the learning of temporal information at different levels of subject personalization. In order to improve orientation measurement, we also present a standartized comparison framework and a novel metric for camera orientation learning. Results on the dataset show that the proposed metric and architectures, as well as the standardized conditions, provide notable improvements to current state of the art camera pose estimation in videobronchoscopy.

Keywords: videobronchoscopy guiding, deep learning, architecture optimization, datasets, standardized evaluation framework, pose estimation

1. Introduction

Early detection is fundamental for lung cancer mortality reduction (Pastorino et al., 2019; de Koning et al., 2020). After a suspicious pulmonary lesion (PL) has been detected through a computed-tomography (CT) scan, decisive diagnosis can only be achieved through a biopsy. Recent advances in sensorics and imaging have improved the sensitivity yield of navigational bronchoscopy (NB) (Asano et al., 2014; Ishiwata et al., 2019), establishing it as a solid alternative to percutaneous approaches, which have a higher degree of medical complications (Han et al., 2018; Gould et al., 2013).

Among the different methods for NB, vision-based NB (VNB) stands out for its low cost, accessible configuration, and reliability. In such method, a virtual model of the patient pulmonary system is built from CT scans (Mori et al., 2008; Skalski et al., 2010; Gil et al., 2019) and the optimal path to the PL is defined. The physician should replicate this path during the intervectional bronchoscopy. Therefore, during navigation, VNB models require the registration of the virtual lung model with the frames from the videobronchoscopy to provide an effective guidance during the biopsy. The registration can be achieved by, either tracking the position and orientation of the bronchoscopy camera (Chien et al., 2020; Byrnes and Higgins, 2014; Luó et al., 2011, 2012a), or by callibrating its deviation from the pose (position and orientation) simulated in the virtual lung model.

---

*Corresponding author. Work performed during a research stage at Computer Vision Center.
Email address: `juan.borrego@uab.cat` (Juan Borrego-Carazo)

Method	Year	Image Size	Tracking Type	PE (mm)	AE (°)	CTF (%)
(Bricault et al., 1998)	1998	100x100	Local/Global	2	5	-
(Mori et al., 2001)	2001	-	Local	-	-	79
(Helferty and Higgins, 2002)	2002	-	Local	-	-	-
(Mori et al., 2002)	2002	410x410	Local	-	-	73.37
(Deligianni et al., 2004)	2004	454x487	Local	3 ± 2.26	2.18 ± 1.63	-
(Nagao et al., 2004)	2004	-	Local	-	-	77.79
(Shinohara et al., 2006)	2006	30x30	Local	-	-	76.4
(Khare et al., 2009)	2009	-	Local	-	-	-
(Khare and Higgins, 2010)	2010	-	Global	-	-	89
(Luó et al., 2011)	2011	362x370	Local	-	-	83.2
(Luo et al., 2011)	2011	30x30	Local	-	-	70.2
(Luó et al., 2012b)	2012	362x370	Local	3.72	10.2	-
(Luo and Mori, 2014)	2014	256x263	Local	4.5	12.3	-
(Shen et al., 2015)	2015	487x487	Local	8.48 ± 6.29	-	-
(Esteban-Lansaque et al., 2016)	2016	-	Global	-	-	-
(Visentini-Scarzanella et al., 2017)	2017	50x50	Local	1.5	-	-
(Shen et al., 2017)	2017	-	Local	5-15[1]	-	-
(Sganga et al., 2019a)	2019	-	Local	2.4	3.4	90.2
(Shen et al., 2019)	2019	307×313	Local	3.18 ± 2.34	-	-
(Zhao et al., 2020)	2020	256x256	Local	1.17	9.71	-
(Wang et al., 2020)	2020	440x440	Local	3.02	-	78.1
(Banach et al., 2021)	2021		Local	6.2 ± 2.9	-	-

Table 1: Comparison among bronchoscopic tracking studies with regards to data and evaluation characteristics. Notably, none of the methods share a dataset (currently there is no publicly available dataset for this task) or publish their code. Moreover, metrics, although aiming to measure the same quantities, are different or lacking in some cases. Metrics shown are umbrella terms for measuring the position error (PE), angle error (AE) and the number of correctly tracked frames. Tracking type (Khare and Higgins, 2010) refers to the type of information provided by the tracking: global type positions the bronchoscope in macro terms, e.g. 3rd bifurcation, while local, gives information with regards to position and angle of the bronchoscope.

Traditionally, the problem of image-based tracking in VNB has been solved through geometric and hand-crafted methods. Feature generation (Chien et al., 2020; Byrnes and Higgins, 2014; ?; Luó et al., 2012a) and similarity measures (Luo and Mori, 2014; Shen et al., 2015; Luo et al., 2011; Khare and Higgins, 2010) were commonly used for such purpose, accounting, however, with tracking errors and large execution times.

Recently, supervised data-intensive learning methods, such as neural networks (NNs) (Sganga et al., 2019b; Visentini-Scarzanella et al., 2017; Zhao et al., 2020; Shen et al., 2019), have been used for localization and tracking in bronchoscopies, providing better results than previous methods. Moreover, temporal learning techniques have recently been applied to other endoscopic modalites (Turan et al., 2017), but has not been appropriately tested in bronchoscopy. Additionally, depth information has lately been extensively used to improve tracking (Recasens et al., 2021; Banach et al., 2021; Shen et al., 2019; Liu et al., 2020), mixing it with generative neural networks (Zhao et al., 2020; Shen et al., 2019; Liu et al., 2020; Banach et al., 2021).

Despite these advances, there is a common obstacle among studies: results are affected by a lack of fair comparability due to the absence of public bronchoscopy datasets and the usage of appropriate metrics as a gold standard, Additionally, learning methods depend on high data availability, often hindering their application in data scarce environments. Table 1 summarizes such situation from bronchoscopic literature: none of the selected studies details public code or data as to allow for a fair comparison. With regards to metrics, position and angle metrics differ or lack in most studies, making further comparison difficult.

In such a situation, with great progress but also lack of comparability, is of outstanding importance establishing ground points for enabling advances. In the case of a system for pose estimation in intervention guiding, it should address 3 key points: 1) definition of the most appropriate metric and comparison protocol for the evaluation of the estimated pose; 2) determination of the most accurate strategy for the processing of temporal information and 3) the highest generalization level (single or across subject) of models.
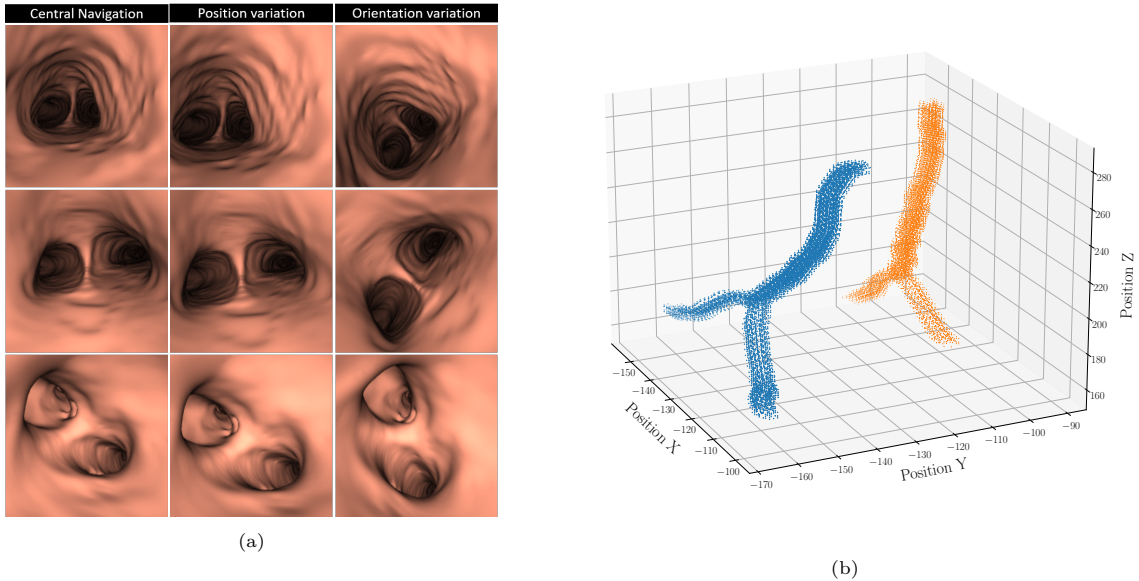
Figure 1: Details of the synthetic dataset for bronchoscopy tracking and calibration. 1a, example of synthetic frames from a trajectory from patient P18 lower left lobe. 1b, positions visited by the trajectories corresponding to two different patients: P18 and P20. Both cases only show points pertaining to lower right and left lobe trajectories.

Hence, the goal of this paper is to analyze the performance of different deep learning approaches for image based bronchoscopy tracking using a standardized and fair comparison framework. In particular, we contribute to the following aspects:

- **Synthetic Dataset.** We present a bronchoscopy navigation synthetic dataset based on real anatomies to enable fair comparison among methods with a cross-subject setting analysis, as well as, address the data requirements of learning methods.

- **Evaluation Protocols.** A study and comparison of rotation and position losses and metrics (including a novel one) for bronchoscopy navigation, which helps to establish better grounds for training and evaluation.

- **Processing of Temporal Information.** We investigate different solutions for neural network temporal learning. Models such as recurrent NNs (RNNs) (Hochreiter and Schmidhuber, 1997a; Schuster and Paliwal, 1997), pseudo-3D convolutions (Tran et al., 2018) or 3D convolutions (Carreira and Zisserman, 2017), could exploit temporal information in bronchoscopic videos, currently not explored, and provide new results.

- **Population Modelling.** We analyse the different options with regards to data usage for industrial applications: first, with a across patient setting, focused on the development of a general model, and secondly, with an intra-patient setting, by providing a specialized model.

Following sections are organized as follows. Section 2 describes the data generation and its characteristics. Section 3 presents the different metrics and losses for evaluation. Then, Section 4, defines and describes important concepts and composition of the proposed bronchoscopic system and its elements, altogether with a description of the temporal learning architectures and proposed population modelling. Next, Sections 5 and 6 present the experiments, their setup, and the main results obtained for the system while exploring their significance. Finally, Section 7 concludes the article with the main key points and important takeaways.

2. Dataset Generation

Virtual lung models are built from an own database of computed tomography scans [2] (Diez-Ferrer et al., 2016) using (Gil et al., 2019) to segment the airways. Virtual airways models are simulated using an own platform developed in C++ and VTK, BronchoX.

From the virtual models, bronchoscope trajectories are simulated from the trachea entrance up until the 4-6 level an covering upper-right, lower-right, upper-left and lower-left lobes. Trajectories are generated from the central navigation path through the luminal central line traversed using the arch-length parameter. Different increments in this parameter allow the simulation of varying velocities across the path.

For each central path, different variations, both, in position (between $[-2 : 1 : 2]$ voxels in each axis) and camera orientation (in the range $[-45 : 15 : 45]$ degrees of rotations around the navigation vector) are generated. The variation in camera position implicitly also modifies the camera point of view, since it is given by its position and a point in the central path at a distance Δd from the current point. The rotation around this navigation vector introduces a variation in the orientation of the image plane. This way, we simulate a full change in the camera central pose.

Finally, paths with neighbouring variations are randomly combined along the navigation arc-length parameter in order to simulate realistic trajectories. In total, our dataset has 876 trajectories per patient and lobe, amounting to a total of 842712 frames.

The dataset has as input values the synthetic frames from the camera view during the trajectory, and as ground truth source values the associated pose inside the VTK airway model coordinate system. The position is in voxel units, and camera view angles are presented in Euler angles.

Figure 1a shows some dataset examples. Each row are different carinas and each column represents variations in position and orientation from the central navigation. Figure 1b shows examples of lower left and right lobes for two different patients. Dataset will be made publicly available upon article acceptance.

Additionally, the data is processed before training to prepare the inputs and ground truths. For every two pair of images in a sequence, the difference in position and rotation between them is computed. Both components, the difference in position $\Delta p = (\Delta x, \Delta y, \Delta z) \in \mathbb{R}^3$, and the difference in orientation $\Delta o = (\Delta\alpha, \Delta\beta, \Delta\gamma) \in \mathbb{R}^3$, define the difference in pose, $\Delta\mathbf{P} = (\Delta p, \Delta o) \in \mathbb{R}^6$, which constitutes the ground truth of the system. That is, we predict the difference in position and orientation between two images, thus allowing both tracking and calibration. No standardization is applied to the ground truths, while images are standardized through mean subtraction and standard deviation division.

3. Metrics

In order to train and validate a system for pose estimation, we need metrics for assessing the error of the predicted rotation and position. As those are two separate components, we can have different metrics for each of them.

The most common choice are either the mean squared error (MSE), when the metric is a loss function, or its equivalent, euclidean norm (L2), when it is an evaluation metric. Although they naturally fit the euclidean space of positions, these functions do not necessarily suit the rotation space. An alternative used in the literature is the *direction error* (DE) (Merritt et al., 2013),

$$DE = \cos^{-1}(\mathbf{v}_E \cdot \mathbf{v}_{GT}) \tag{1}$$

where $\mathbf{v}_E$ and $\mathbf{v}_{GT}$ are, respectively, the estimated and ground truth direction vectors. Such direction vectors are computed from the rotation matrix, R, and a unitary direction vector, $\mathbf{u}$, as:

---

[2] CVC CPAP Study Database

(a) Base network

(b) Base network plus recurrence

(c) Base network plus convolutional recurrence
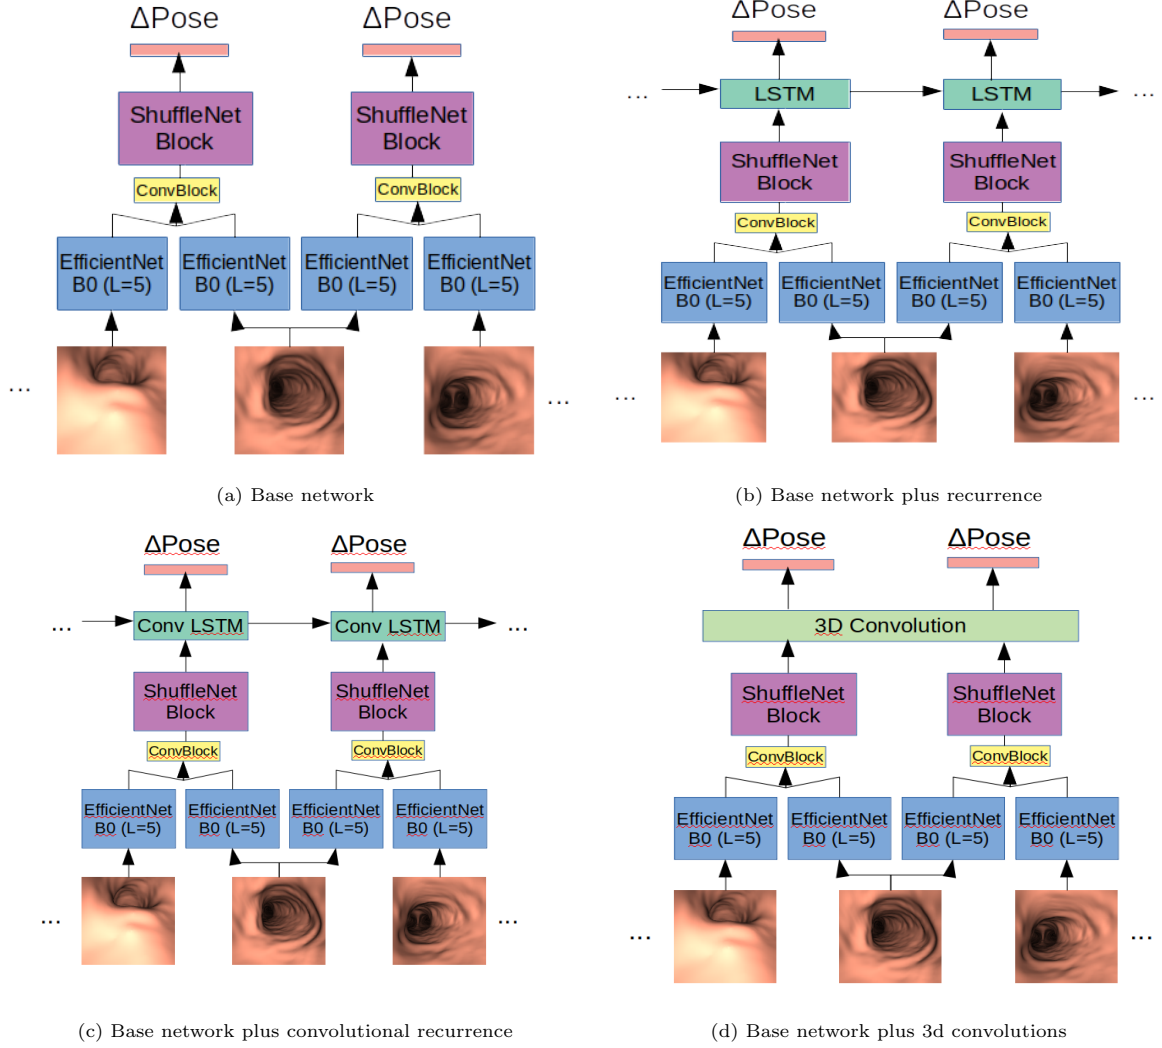
(d) Base network plus 3d convolutions

Figure 2: Architectures for bronchoscopy callibration and tracking. (a) is the baseline network without temporal information management, as in (Sganga et al., 2019a). (b, c, d) include different mechanisms to manage temporal information across predictions.

$$\mathbf{v}_r = R \cdot \mathbf{u} =$$

$$\begin{bmatrix} c_\beta c_\gamma & -c_\beta s_\gamma & s_\beta \\ s_\alpha s_\beta c_\gamma + c_\alpha s_\gamma & -s_\alpha s_\beta s_\gamma + c_\alpha c_\gamma & -s_\alpha c_\beta \\ -c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma & c_\alpha s_\beta s_\gamma + s_\alpha s_\gamma & c_\alpha c_\beta \end{bmatrix} \begin{bmatrix} u_x \\ u_y \\ u_z \end{bmatrix}$$

where, c_i and s_i are, respectively, the cosinus and sinus of angle i. A common choice for $\mathbf{u}$ is the look-at vector of the virtual camera, usually given by the x-axis, so that:

$$\mathbf{v}_r = R \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} c_\beta c_\gamma \\ s_\alpha s_\beta c_\gamma + c_\alpha s_\gamma \\ -c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma \end{bmatrix}$$

The main issue with such function is that the choice of $\mathbf{u}$ affects the perceived rotations. That is, by selecting an specific $\mathbf{u}$, the system is oblivious to the rotations through that direction.

To remedy such problem, we present an alternative metric, the *cosinus error* (CE):

$$
\begin{aligned}
CE = \frac{1}{3}((1 - \cos(\Delta\alpha_E - \Delta\alpha_{GT})) + \\
(1 - \cos(\Delta\beta_E - \Delta\beta_{GT})) + \\
(1 - \cos(\Delta\gamma_E - \Delta\gamma_{GT})))
\end{aligned}
$$

where $(\alpha_E, \beta_E, \gamma_E)$, $(\alpha_{GT}, \beta_{GT}, \gamma_{GT})$ are, respectively the estimated and ground truth Euler angles.

Importantly, all metrics can be used both as loss for training, as well as metric for evaluation of performance.

4. Relative Pose Estimation

Given a sequence of L image pairs at time t, $(I_1^t, I_2^t) \in \mathcal{D}_1 \times \mathcal{D}_2$, for $\mathcal{D}_1$, $\mathcal{D}_2$ a source and target domains, a pose estimation system can be formulated as a function, $f_{\mathcal{P}}$, predicting the difference in pose between the image pairs:

$$
f_{\mathcal{P}} : (I_1^t, I_2^t)_{t=0}^{t=L} \longrightarrow (\Delta\mathbf{P}^t)_{t=0}^{t=L}, \quad I_i^t \in \mathcal{D}_i, \Delta P^t \in \mathbb{R}^6 \tag{2}
$$

where each domain $\mathcal{D}_i = \mathbb{R}^{C_i \times H_i \times W_i}$ represents RGB images ($C_i = 3$) of size $H_i \times W_i$, and the difference in pose for each time t is a vector $\Delta P^t = (\Delta p^t, \Delta o^t) = (\Delta x^t, \Delta y^t, \Delta z^t, \Delta\alpha^t, \Delta\beta^t, \Delta\gamma^t)$ representing the difference in (x, y, z) position coordinates and (α, β, γ) Euler angles.

In case I_1^t, I_2^t are consecutive frames of the same path, $f_{\mathcal{P}}$ is modelling a tracker, and the change in pose through the sequence would be given by accumulating the differences estimated across the video:

$$
\mathbf{P}_L = \mathbf{P}^0 + \sum_{t=0}^{L} \Delta\mathbf{P}^t \tag{3}
$$

where $\mathbf{P}^0$ corresponds to the initial pose vector, and $\mathbf{P}^L$ to the pose vector after the accumulated changes of the L pose differences.

If I_1^t, I_2^t correspond to frames of different paths, $f_{\mathcal{P}}$ would be modelling a pose callibration.

The network loss is given by the addition of the position and rotation metrics:

$$
\mathfrak{L} = \mathfrak{L}_{\mathfrak{p}} + \mathfrak{L}_{\mathfrak{o}} \tag{4}
$$

where p refers to position and o to orientation. The position loss $\mathfrak{L}_{\mathfrak{p}}$ is given by the MSE error, while for the orientation loss $\mathfrak{L}_{\mathfrak{o}}$ we used the three metrics (MSE, DE and the proposed CE) described in the previous section.

In any case (tracking or callibration), if the input sequence has more than two image pairs ($L > 0$), there are several ways of processing such temporal information in order to improve the difference in pose estimation. All architectures follow Equation 3 and their scheme can be found at Figure 2. The different architectures used are: a base network working only between two images, and 3 different ways of incorporating temporal information across frames. Next, configuration details for each architecture are presented, altogether with the population modeling approach.

4.1. Base Network (Baseline)

The baseline network is a static estimation of pose differences from a single image pair ($L = 0$ in Equation 2). Each image is passed through a convolution backbone, specifically an EfficientNet-B0 (Tan and Le, 2020). Then both feature maps are concatenated and passed through a convolutional block and a

ShuffleNet block (Ma et al., 2018), to obtain a suitable performance/latency trade-off. Finally, the resulting feature map is flattened and passed through a fully connected layer to obtain the difference pose prediction between the two images. An illustration of the overall components of the base network can be found at Figure 2a.

4.2. Recurrence (Baseline + LSTM)

The previous network does not include temporal management. To be able to include such information, the first modification to the base network consists in the addition of a recurrent LSTM (Hochreiter and Schmidhuber, 1997b) (Long Short Term Memory) module after the ShuffleNet block. Such convolutional plues recurrent network type can well exploit temporal information, and thus it has been successfully applied to video tasks, such as object tracking (Ning et al., 2017), action recognition (Ullah et al., 2018) or video captioning (Jin et al., 2019).

For every pair of images in the sequence we obtain a group of feature maps. Each one is flattened and a vector, $\mathbf{v}$, of dimension (L, F), is built, where L is the number of image pairs and F the size of the flattened feature maps. Such vector, $\mathbf{v}$, is the input for the LSTM block. At each step, l from the sequence of L image pairs, the output vector from the LSTM cell is forwarded to a fully connected layer, which finally delivers the difference pose vector. An illustration of the mentioned module is found at Figure 2b.

4.3. Convolutional recurrence (Baseline + ConvRNN)

In the previous architecture recurrence required flattening the feature maps so they could be fed to the LSTM. Such flattening destroys visual relations present in the feature maps, thus losing information. To avoid such loss, a possibility is to use a convolutional LSTM (Shi et al., 2015), where vectorial operations are substituted by convolution ones. In such way, we are able to maintain the visual structure during recurrence. Flattening, however, is still needed to produce the pose prediction through a fully connected layer after Conv-LSTM. In Figure 2c, an illustration of the overall structure is presented.

4.4. 3D Convolution (Baseline + 3D)

An alternative to recurrence for managing temporal information is 3D convolutions (Ji et al., 2018; Guo et al., 2019). Once all the feature representations from all image pairs in the sequence are generated, a 4D tensor of size (L, C, H, W) can be built.

Such tensor is fed to a block of two (Conv3D, BatchNorm, ReLU) layers. Working at once with the feature maps coming from all image pairs allows to learn relations among them, producing improvements in angle and position estimation. After the 3D convolution block, the result is flattened and fed to a fully connected layer to predict the final difference pose prediction. In Figure 2d an illustration of the overall network can be seen.

4.5. Approaches for Population

Once the data has been prepared as specified in Section 2, we define two different approaches to prepare our training and validation scheme. The purpose is to establish the attainable degree of generalization of the models developed in two different industrial settings.

First, the population or cross-subject setting, in which a patient is selected as validation and the rest of patients are used for training. This case, in an industrial environment, would correspond to building a general model with different patients and expect enough generalization of the model as to apply it directly to an unseen patient.

And second, the personalized setting, where validation is performed over the same patients as training, but with different sequences. Such procedure, in an industrial environment, implies that every time we include a patient, models should be retrained. Hence, although the procedure would be more cumbersome, less generalization effort is expected from models.

Baseline	Position Error	Rotation Error		
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	9.657 ± 5.777	**19.037 ± 51.06**	1.608 ± 0.658	**0.803 ± 0.357**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	8.426 ± 4.569	45.196 ± 46.648	1.566 ± 0.63	1.018 ± 0.411
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	**7.138 ± 4.547**	35.029 ± 47.231	**1.477 ± 0.65**	0.925 ± 0.937
Baseline + LSTM				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	4.962 ± 2.739	**19.779 ± 50.898**	1.733 ± 0.567	**0.775 ± 0.362**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	6.714 ± 3.495	45.229 ± 47.895	**1.528 ± 0.565**	1.020 ± 0.424
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	**4.515 ± 3.636**	36.255 ± 45.839	1.670 ± 0.718	0.791 ± 0.411
Baseline + 3D				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	7.798 ± 4.176	**19.499 ± 52.023**	**1.502 ± 0.726**	**0.740 ± 0.359**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	8.487 ± 4.757	40.985 ± 45.632	1.572 ± 0.683	1.020 ± 0.411
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	**5.068 ± 2.869**	36.445 ± 48.947	1.678 ± 0.534	0.7595 ± 0.356
Baseline + ConvRNN				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	4.550 ± 3.352	**21.163 ± 50.795**	1.609 ± 0.665	**0.669 ± 0.348**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	6.088 ± 4.686	42.226 ± 46.591	1.643 ± 0.642	0.963 ± 0.441
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	**4.487 ± 3.945**	34.881 ± 48.553	**1.487 ± 0.663**	0.809 ± 0.394

Table 2: Personalized data scheme. Results for the different loss and architecture combinations. Values show mean and standard deviation computed among all the paths in the validation set. Results are computed after accumulating tracking through the whole validation paths and only by evaluating last position and orientation, as indicated in Equation 3. Bold font indicates best value inside group of row and column. Red indicates best of column, and blue second best of column.

Baseline	Position Error	Rotation Error		
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	0.447 ± 0.276	**1.142 ± 8.339**	**0.649 ± 0.513**	**0.147 ± 0.201**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	0.459 ± 0.282	2.529 ± 8.465	0.830 ± 0.590	0.472 ± 0.392
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	**0.395 ± 0.268**	1.749 ± 8.508	0.735 ± 0.571	0.188 ± 0.237
Baseline + LSTM				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,MSE}$	**0.353 ± 0.221**	**1.092 ± 8.349**	**0.611 ± 0.501**	**0.136 ± 0.195**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	0.421 ± 0.241	2.516 ± 8.488	0.816 ± 0.602	0.446 ± 0.371
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$	0.371 ± 0.240	1.888 ± 8.157	0.790 ± 0.599	0.213 ± 0.250
Baseline + 3D				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{MSE} + \mathcal{L}_{MSE}$	0.443 ± 0.295	**0.916 ± 8.341**	**0.503 ± 0.416**	**0.093 ± 0.149**
$\mathcal{L}_{MSE} + \mathcal{L}_{DE}$	0.467 ± 0.289	2.293 ± 8.841	0.880 ± 0.590	0.386 ± 0.327
$\mathcal{L}_{MSE} + \mathcal{L}_{CE}$	**0.397 ± 0.252**	1.592 ± 8.508	0.724 ± 0.559	0.167 ± 0.221
Baseline + ConvRNN				
Loss	L_2	L_2	DE	CE
$\mathcal{L}_{MSE} + \mathcal{L}_{MSE}$	0.368 ± 0.245	**1.165 ± 8.348**	**0.655 ± 0.502**	**0.153 ± 0.197**
$\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,DE}$	0.432 ± 0.260	2.480 ± 8.449	0.888 ± 0.635	0.452 ± 0.361
$\mathcal{L}_{MSE} + \mathcal{L}_{CE}$	**0.355 ± 0.265**	1.80 ± 8.522	0.732 ± 0.537	0.195 ± 0.242

Table 3: Personalized data scheme. Results for the different loss and architecture combinations evaluated at every image pair. Values show mean and standard deviation computed among all the pairs of images in the paths of the validation set. Bold font indicates best value inside group of row and column. Red indicates best of column, and blue second best of column.

		Position Error	Rotation Error		
Network	Loss	L_2	L_2	DE	CE
Baseline + LSTM	$\mathcal{L}_{MSE} + \mathcal{L}_{MSE}$	16.398 ± 7.215	29.683 ± 50.605	1.659 ± 0.691	0.863 ± 0.391
Baseline + 3D	$\mathcal{L}_{MSE} + \mathcal{L}_{MSE}$	18.245 ± 8.158	24.104 ± 52.498	1.583 ± 0.734	0.717 ± 0.355
Baseline + ConvRNN	$\mathcal{L}_{MSE} + \mathcal{L}_{CE}$	14.941 ± 7.886	37.561 ± 48.519	1.580 ± 0.666	0.848 ± 0.388

Table 4: Cross-subject data scheme. Results for loss and architecture combinations that obtained best results in the personalized data scheme. Results are computed after accumulating tracking through the whole validation path and only by evaluating last position and orientation, as indicated in Equation 3. Values show mean and standard deviation computed among all the paths in the validation set. Red indicates best value in column, and blue column's second best.

		Position Error	Rotation Error		
Network	Loss	L_2	L_2	DE	CE
Baseline + LSTM	$\mathcal{L}_{\mathfrak{p}MSE} + \mathcal{L}_{\mathfrak{o}MSE}$	0.680 ± 0.319	1.901 ± 9.913	1.017 ± 0.659	0.324 ± 0.279
Baseline + 3D	$\mathcal{L}_{\mathfrak{p}MSE} + \mathcal{L}_{\mathfrak{o}MSE}$	0.711 ± 0.339	1.314 ± 9.905	0.739 ± 0.547	0.172 ± 0.203
Baseline + ConvRNN	$\mathcal{L}_{\mathfrak{p}MSE} + \mathcal{L}_{\mathfrak{o}CE}$	0.678 ± 0.334	2.348 ± 9.965	1.133 ± 0.712	0.392 ± 0.305

Table 5: Cross-subject data scheme. Results for loss and architecture combinations that obtained best results in the personalized data scheme. Results are computed for every pair of images. Values show mean and standard deviation computed among all the pairs of images in the paths of the validation set. Red indicates best of column, and blue column's second best.

5. Experimental setup

In this section we present the experimental details and the proposed experiments using the different configurations stated in Section 4.

Three different experiments are defined:

1. **Model Optimization.** Loss and architecture combinations are trained and validated on a personalized data setting.
2. **Model generalization**. Best architecture and loss combinations from Experiment 1 are selected. They are trained and tested on a population data scheme with a leave-one-out validation structure.
3. **Model comparison with SoA**. To assess the advantages of the proposed strategy, the best model selected in the first and second experiments is compared to a SoA method.

For each experiment we have selected different groups of data. Specifically, in the case of personalized data scheme, we have selected 15 trajectories per patient and lobe for training, reserving a total of 3 for validation. Each sequence has been divided into a set of pairs of images. In the case of temporal information experiments, we have selected the same number of sequences, but sequences have been split into chunks of 10 image pairs. Additionally, in the population data scheme, we have selected 15 trajectories per patient and lobe for 4 patients for training, and 15 trajectories of 1 patient for validation.

Networks have been trained using two NVIDIA RTX 2080ti, with Pytorch Lightning as a framework (Falcon and The PyTorch Lightning team, 2019). Adam optimizer with a learning rate of $1e^{-4}$ has been used and early stopping based on validation loss evolution has been applied. Dropout has been added to the ShuffleNet blocks to avoid overfitting. In the case of temporal learning, networks have been trained using truncated back propagation through time (T-BPTT). A batch size of 512 has been used in all the experiments. With regards to evaluation metrics, the same functions used as losses are used for evaluation. Namely, L_2 for position, and L_2, DE, and CE, for orientation.

6. Results and discussion

In the next section we present the results corresponding to the three experiments stated in Section 5. Code used to build the system and reproduce experiment results will be published upon acceptance.

Network	Param. (M)	L_2 (voxel)	CE
OffsetNet (Sganga et al., 2019c)	43.6	5.694 ± 4.634	0.994 ± 0.342
BronchoTrack	**14.1**	$\mathbf{4.994 \pm 4.442}$	$\mathbf{0.808 \pm 0.354}$

Table 6: Comparison with state of the art (SoA) method (Sganga et al., 2019c) for bronchoscopy tracking in terms of position error (L_2), angle error, and the number of parameters. Results are computed after accumulating tracking through the whole test paths and only by evaluating last position and orientation, as indicated in Equation 3. Values show mean and standard deviation computed among all the paths in the test set.
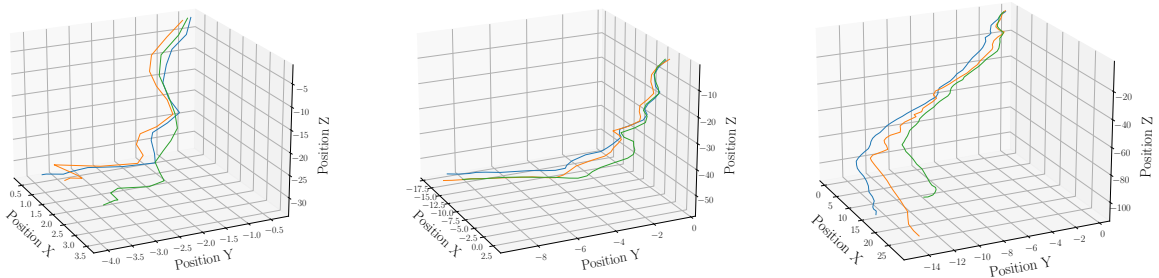
Figure 3: Sample position tracking comparison with SoA method, OffsetNet (green), our method (orange), and the ground truth (blue). Axis coordinates are defined in voxel units, with the (0, 0, 0) being in the trachea.

6.1. Model Optimization

Table 2 shows results (mean $\pm$ standard deviation) in the validation set for the different loss and architecture combinations evaluated with the different metrics using the personalized data scheme. Predictions for each trajectory have been accumulated and the metrics are evaluated on the resulting accumulation.

Regardless of the loss, the best values with regards to position error are achieved, with notable difference, by those architectures with temporal information management, indicating the importance of it. Moreover, the combination of ME + CE losses provides the first and second best results, proving its usefulness for better positioning. With regards to rotation error, we can point several interesting points. First, temporal management information structures notably improve CE metric result, with the a great improvement with Convolutional RNN. Second, MSE seems to provide better results in both L_2 and CE metrics, and with the exception of DE. And, finally, in the case of DE and L_2 temporal information information management architectures seem not to provide any benefits.

Differently to Table 2, Table 3 shows the results for the same configurations but evaluating over image pairs and not whole trajectories. In this case, the loss combination MSE + CE still provides the best position estimation and best orientation differences are provided by the MSE + MSE loss, in concordance with results in Table 2 . In this case, however, the 3D convolution architecture shows superior performance in orientation estimation.

In general two conclusions can be obtained from Table 2 and Table 3 results. First, temporal management information structures help improve both rotation and position predictions, specially in the case of ConvRNN architecture for position and 3D for orientation. Secondly, the loss combination MSE + CE delivers better results for position prediction, while MSE + MSE improves rotation precision.

6.2. Model generalization

In Tables 4 and 5 we show results for the cross-subject or population data scheme defined in Section 4.5. Architecture and loss combinations are selected from the results in the personalized data scheme evaluation. Table 4 shows results evaluated at the last position and orientation after accumulating trajectory changes, and Table 5 shows results for each image pair.

As seen, with regards to position error, the combination of ConvRNN and MSE + CE loss delivers notably improved results compared to the other two alternatives, showcasing its better generalization capabilities with unseen patients. With regards to orientation errors, the 3D convolution mechanism provides, with difference, best results in orientation, both with metrics L_2 and CE. Followed by ConvRNN in the accumulated case and by LSTM in the image-pair case. Hence, it can be stated that convRNN plus MSE + CE as loss is the best option for position prediction, while 3D convolutions plus MSE + MSE is the best one for rotation prediction.

Importantly, another fact that can be observed is the difference in results between the personalized (Tables 2 and 3) and the population data schemes (Tables 4 and 5). In the former, results are much better for both for orientation and position with metric L_2, although the effect in CE and DE is smaller. Such point showcases the benefits of the personalized model, and the usefulness of retraining the model with each patient information.

6.3. Comparison to State-of-the-art

Finally, we compare our model with current SoA for bronchoscopy tracking, OffsetNet (Sganga et al., 2019b), implementing it as described in the original publication. For such comparison we build on our best results in previous sections and choose as loss the combination $\mathcal{L}_{\mathfrak{p}\,MSE} + \mathcal{L}_{\mathfrak{o}\,CE}$, and as architecture, the convolutional recurrent network. We call this network BronchoTrack. Both networks are trained and tested in an intra-patient setting, using sequences of 10 steps long for training, and full sequences for testing.

In Table 6 we show results for the comparison in the test set. As shown our network is better for both position and orientation tracking. Moreover, it has a notably reduced number of parameters: more than 3x times less parameters.

In Figure 3 we further compare both models by showing the accumulated position tracking for sample trajectories in the test set. As seen our model is able to closely match the ground truth sequence while OffsetNet deviates more and struggles to follow the ground truth path.

7. Conclusions

In the present study we have presented several contributions to bronchoscopy tracking. We have built a synthetic dataset to allow for fair comparison between methods and be able to train data-hungry models. Two different population approaches for training learning methods have been analyzed, concluding with the benefits of a personalized setting. We also have experimented with the configuration of a neural network model with regards to the loss function and the temporal information management. Finally, when comparing to a SoA method, better results have been obtained while using less parameters. All in all, there are still important next steps to take to improve current results. Such steps could involve the study of transfer learning to real bronchoscopy videos, for example with the use of generative adversarial networks, or the adaptation to the specific conditions of a medical setting in terms of hardware and resources.

CRediT Authorship Contribution Statement

Juan Borrego-Carazo: Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Roles/Writing - original draft, Writing - review & editing. **Carles Sánchez**: Conceptualization, Data curation, Formal analysis, Supervision, Visualization, Roles/Writing - original draft, Writing - review & editing. **David Castells-Rufas**: Conceptualization, Methodology, Supervision, Validation, Writing - review & editing. **Jordi Carrabina**: Funding acquisition, Project administration, Resources, Supervision, Writing - review & editing. **Débora Gil**: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Supervision, Validation, Roles/Writing - original draft, Writing - review & editing.

Declaration of Competing Interest

Acknowledgements

References

Asano, F., Eberhardt, R., Herth, F.J.F., 2014. Virtual bronchoscopic navigation for peripheral pulmonary lesions. Respiration; International Review of Thoracic Diseases 88, 430–440. doi:10.1159/000367900.

Banach, A., King, F., Masaki, F., Tsukada, H., Hata, N., 2021. Visually Navigated Bronchoscopy using Three Cycle-Consistent Generative Adversarial Network for Depth Estimation. Medical Image Analysis , 102164.

Bricault, I., Ferretti, G., Cinquin, P., 1998. Registration of real and CT-derived virtual bronchoscopic images to assist transbronchial biopsy. IEEE transactions on medical imaging 17, 703–714.

Byrnes, P.D., Higgins, W.E., 2014. Construction of a multimodal CT-video chest model, in: Medical Imaging 2014: Image-Guided Procedures, Robotic Interventions, and Modeling, International Society for Optics and Photonics. p. 903607. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9036/903607/Construction-of-a-multimodal-CT-video-chest-model/10.1117/12.2041609.short, doi:10.1117/12.2041609.

Carreira, J., Zisserman, A., 2017. Quo vadis, action recognition? a new model and the kinetics dataset, in: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6299–6308.

Chien, J.C., Lee, J.D., Su, E., Li, S.H., 2020. A Bronchoscope Localization Method Using an Augmented Reality Co-Display of Real Bronchoscopy Images with a Virtual 3D Bronchial Tree Model. Sensors 20, 6997. URL: https://www.mdpi.com/1424-8220/20/23/6997, doi:10.3390/s20236997.

Deligianni, F., Chung, A., Yang, G.z., 2004. Patient-specific bronchoscope simulation with pq-space-based 2D/3D registration. Computer Aided Surgery 9, 215–226. URL: https://doi.org/10.3109/10929080500144927, doi:10.3109/10929080500144927.

Diez-Ferrer, M., Gil, D., Carreño, E., Padrones, S., Aso, S., Vicens, V., Noelia, C., Lisbona, R.L., Sanchez, C., Borras, A., 2016. Positive airway pressure-enhanced CT to improve virtual bronchoscopic navigation. Chest 150, 1003A. Publisher: Elsevier.

Esteban-Lansaque, A., Sánchez, C., Borràs, A., Diez-Ferrer, M., Rosell, A., Gil, D., 2016. Stable Anatomical Structure Tracking for Video-Bronchoscopy Navigation, in: Shekhar, R., Wesarg, S., González Ballester, M.A., Drechsler, K., Sato, Y., Erdt, M., Linguraru, M.G., Oyarzun Laura, C. (Eds.), Clinical Image-Based Procedures. Translational Research in Medical Imaging, Springer International Publishing, Cham. pp. 18–26. doi:10.1007/978-3-319-46472-5_3.

Falcon, W., The PyTorch Lightning team, 2019. PyTorch Lightning. URL: https://github.com/PyTorchLightning/pytorch-lightning, doi:10.5281/zenodo.3828935.

Gil, D., Sanchez, C., Borras, A., Diez-Ferrer, M., Rosell, A., 2019. Segmentation of distal airways using structural analysis. PLOS ONE 14, e0226006. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0226006, doi:10.1371/journal.pone.0226006.

Gould, M.K., Donington, J., Lynch, W.R., Mazzone, P.J., Midthun, D.E., Naidich, D.P., Wiener, R.S., 2013. Evaluation of individuals with pulmonary nodules: When is it lung cancer?: Diagnosis and management of lung cancer: American College of Chest Physicians evidence-based clinical practice guidelines. Chest 143, e93S–e120S.

Guo, S., Lin, Y., Li, S., Chen, Z., Wan, H., 2019. Deep Spatial–Temporal 3D Convolutional Neural Networks for Traffic Data Forecasting. IEEE Transactions on Intelligent Transportation Systems 20, 3913–3926. doi:10.1109/TITS.2019.2906365.

Han, Y., Kim, H.J., Kong, K.A., Kim, S.J., Lee, S.H., Ryu, Y.J., Lee, J.H., Kim, Y., Shim, S.S., Chang, J.H., 2018. Diagnosis of small pulmonary lesions by transbronchial lung biopsy with radial endobronchial ultrasound and virtual bronchoscopic navigation versus CT-guided transthoracic needle biopsy: A systematic review and meta-analysis. PLOS ONE 13, e0191590. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0191590, doi:10.1371/journal.pone.0191590.

Helferty, J.P., Higgins, W.E., 2002. Combined endoscopic video tracking and virtual 3D CT registration for surgical guidance, in: Proceedings. International Conference on Image Processing, IEEE. pp. II–II.

Hochreiter, S., Schmidhuber, J., 1997a. Long Short-Term Memory. Neural Computation 9, 1735–1780. doi:10.1162/neco.1997.9.8.1735.

Hochreiter, S., Schmidhuber, J., 1997b. LSTM can solve hard long time lag problems. Advances in neural information processing systems , 473–479.

Ishiwata, T., Gregor, A., Inage, T., Yasufuku, K., 2019. Advances in interventional diagnostic bronchoscopy for peripheral pulmonary lesions. Expert review of respiratory medicine 13, 885–897.

Ji, S., Zhang, C., Xu, A., Shi, Y., Duan, Y., 2018. 3D Convolutional Neural Networks for Crop Classification with Multi-Temporal Remote Sensing Images. Remote Sensing 10, 75. URL: https://www.mdpi.com/2072-4292/10/1/75, doi:10.3390/rs10010075.

Jin, T., Li, Y., Zhang, Z., 2019. Recurrent convolutional video captioning with global and local attention. Neurocomputing 370, 118–127. URL: https://www.sciencedirect.com/science/article/pii/S0925231219311804, doi:10.1016/j.neucom.2019.08.042.

Khare, R., Higgins, W.E., 2010. Toward image-based global registration for bronchoscopy guidance, in: Medical Imaging 2010: Visualization, Image-Guided Procedures, and Modeling, International Society for Optics and Photonics. p. 762510.

Khare, R., Yu, K.C., Higgins, W.E., 2009. Improved navigation for image-guided bronchoscopy, in: Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling, International Society for Optics and Photonics. p. 72612J. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7261/72612J/Improved-navigation-for-image-guided-bronchoscopy/10.1117/12.811142.short, doi:10.1117/12.811142.

de Koning, H.J., van der Aalst, C.M., de Jong, P.A., Scholten, E.T., Nackaerts, K., Heuvelmans, M.A., Lammers, J.W.J., Weenink, C., Yousaf-Khan, U., Horeweg, N., 2020. Reduced lung-cancer mortality with volume CT screening in a randomized trial. New England Journal of Medicine 382, 503–513.

Liu, X., Berg, J., King, F., Hata, N., 2020. Computer vision-guided bronchoscopic navigation using dual CNN-generated depth images and ICP registration, in: Medical Imaging 2020: Image-Guided Procedures, Robotic Interventions, and Modeling, International Society for Optics and Photonics. p. 113152C. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11315/113152C/Computer-vision-guided-bronchoscopic-navigation-using-dual-CNN-generated-depth/10.1117/12.2549719.short, doi:10.1117/12.2549719.

Luo, X., Kitasaka, T., Mori, K., 2011. ManiSMC: A New Method Using Manifold Modeling and Sequential Monte Carlo Sampler for Boosting Navigated Bronchoscopy, in: Fichtinger, G., Martel, A., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011, Springer, Berlin, Heidelberg. pp. 248–255. doi:10.1007/978-3-642-23626-6_31.

Luo, X., Mori, K., 2014. A discriminative structural similarity measure and its application to video-volume registration for endoscope three-dimensional motion tracking. IEEE transactions on medical imaging 33, 1248–1261.

Luó, X., Feuerstein, M., Deguchi, D., Kitasaka, T., Takabatake, H., Mori, K., 2012a. Development and comparison of new hybrid motion tracking for bronchoscopic navigation. Medical image analysis 16, 577–596.

Luó, X., Feuerstein, M., Kitasaka, T., Mori, K., 2012b. Robust bronchoscope motion tracking using sequential Monte Carlo methods in navigated bronchoscopy: dynamic phantom and patient validation. International Journal of Computer Assisted Radiology and Surgery 7, 371–387. URL: https://doi.org/10.1007/s11548-011-0645-6, doi:10.1007/s11548-011-0645-6.

Luó, X., Feuerstein, M., Kitasaka, T., Natori, H., Takabatake, H., Hasegawa, Y., Mori, K., 2011. On scale invariant features and sequential Monte Carlo sampling for bronchoscope tracking, in: Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling, International Society for Optics and Photonics. p. 79640Q. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7964/79640Q/On-scale-invariant-features-and-sequential-Monte-Carlo-sampling-for/10.1117/12.877515.short, doi:10.1117/12.877515.

Ma, N., Zhang, X., Zheng, H.T., Sun, J., 2018. ShuffleNet V2: Practical Guidelines for Efficient CNN Architecture Design. arXiv:1807.11164 [cs] URL: http://arxiv.org/abs/1807.11164. arXiv: 1807.11164.

Merritt, S.A., Khare, R., Bascom, R., Higgins, W.E., 2013. Interactive CT-Video Registration for the Continuous Guidance of Bronchoscopy. IEEE Transactions on Medical Imaging 32, 1376–1396. doi:10.1109/TMI.2013.2252361.

Mori, K., Deguchi, D., Hasegawa, J.i., Suenaga, Y., Toriwaki, J.i., Takabatake, H., Natori, H., 2001. A method for tracking the camera motion of real endoscope by epipolar geometry analysis and virtual endoscopy system, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 1–8.

Mori, K., Deguchi, D., Kitasaka, T., Suenaga, Y., Hasegawa, Y., Imaizumi, K., Takabatake, H., 2008. Improvement of accuracy of marker-free bronchoscope tracking using electromagnetic tracker based on bronchial branch information, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 535–542.

Mori, K., Deguchi, D., Sugiyama, J., Suenaga, Y., Toriwaki, J.i., Maurer Jr, C.R., Takabatake, H., Natori, H., 2002. Tracking of a bronchoscope using epipolar geometry analysis and intensity-based image registration of real and virtual endoscopic images. Medical Image Analysis 6, 321–336.

Nagao, J., Mori, K., Enjouji, T., Deguchi, D., Kitasaka, T., Suenaga, Y., Hasegawa, J.i., Toriwaki, J.i., Takabatake, H., Natori, H., 2004. Fast and Accurate Bronchoscope Tracking Using Image Registration and Motion Prediction, in: Barillot, C., Haynor, D.R., Hellier, P. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004, Springer, Berlin, Heidelberg. pp. 551–558. doi:10.1007/978-3-540-30136-3_68.

Ning, G., Zhang, Z., Huang, C., Ren, X., Wang, H., Cai, C., He, Z., 2017. Spatially supervised recurrent convolutional neural networks for visual object tracking, in: 2017 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 1–4. doi:10.1109/ISCAS.2017.8050867. iSSN: 2379-447X.

Pastorino, U., Silva, M., Sestini, S., Sabia, F., Boeri, M., Cantarutti, A., Sverzellati, N., Sozzi, G., Corrao, G., Marchianò, A., 2019. Prolonged lung cancer screening reduced 10-year mortality in the MILD trial: new confirmation of lung cancer screening efficacy. Annals of Oncology 30, 1162–1169.

Recasens, D., Lamarca, J., Fácil, J.M., Montiel, J.M.M., Civera, J., 2021. Endo-Depth-and-Motion: Localization and Reconstruction in Endoscopic Videos using Depth Networks and Photometric Constraints. arXiv preprint arXiv:2103.16525 .

Schuster, M., Paliwal, K., 1997. Bidirectional recurrent neural networks. IEEE Transactions on Signal Processing 45, 2673–2681.

doi:10.1109/78.650093.

Sganga, J., Eng, D., Graetzel, C., Camarillo, D., 2019a. Offsetnet: Deep learning for localization in the lung using rendered images, in: 2019 International Conference on Robotics and Automation (ICRA), IEEE. pp. 5046–5052.

Sganga, J., Eng, D., Graetzel, C., Camarillo, D.B., 2019b. Autonomous driving in the lung using deep learning for localization. arXiv preprint arXiv:1907.08136 .

Sganga, J., Eng, D., Graetzel, C., Camarillo, D.B., 2019c. Deep learning for localization in the lung. arXiv preprint arXiv:1903.10554 .

Shen, M., Giannarou, S., Shah, P.L., Yang, G.Z., 2017. BRANCH:Bifurcation Recognition for Airway Navigation based on struCtural cHaracteristics, in: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (Eds.), Medical Image Computing and Computer-Assisted Intervention - MICCAI 2017, Springer International Publishing, Cham. pp. 182–189. doi:10.1007/978-3-319-66185-8_21.

Shen, M., Giannarou, S., Yang, G.Z., 2015. Robust camera localisation with depth reconstruction for bronchoscopic navigation. International Journal of Computer Assisted Radiology and Surgery 10, 801–813. URL: https://doi.org/10.1007/s11548-015-1197-y, doi:10.1007/s11548-015-1197-y.

Shen, M., Gu, Y., Liu, N., Yang, G.Z., 2019. Context-Aware Depth and Pose Estimation for Bronchoscopic Navigation. IEEE Robotics and Automation Letters 4, 732–739. doi:10.1109/LRA.2019.2893419.

Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.k., Woo, W.c., 2015. Convolutional LSTM Network: a machine learning approach for precipitation nowcasting, in: Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, MIT Press, Cambridge, MA, USA. pp. 802–810.

Shinohara, R., Mori, K., Deguchi, D., Kitasaka, T., Suenaga, Y., Takabatake, H., Mori, M., Natori, H., 2006. Branch identification method for CT-guided bronchoscopy based on eigenspace image matching between real and virtual bronchoscopic images, in: Medical Imaging 2006: Physiology, Function, and Structure from Medical Images, International Society for Optics and Photonics. p. 614314. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/6143/614314/Branch-identification-method-for-CT-guided-bronchoscopy-based-on-eigenspace/10.1117/12.654451.short, doi:10.1117/12.654451.

Skalski, A., Socha, M., Duplaga, M., Duda, K., Zieliński, T., 2010. 3D segmentation and visualisation of mediastinal structures adjacent to tracheobronchial tree from CT data, in: Information Technologies in Biomedicine. Springer, pp. 523–534.

Tan, M., Le, Q.V., 2020. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. arXiv:1905.11946 [cs, stat] URL: http://arxiv.org/abs/1905.11946. arXiv: 1905.11946.

Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M., 2018. A Closer Look at Spatiotemporal Convolutions for Action Recognition. arXiv:1711.11248 [cs] URL: http://arxiv.org/abs/1711.11248. arXiv: 1711.11248 version: 3.

Turan, M., Almalioglu, Y., Gilbert, H., Sari, A.E., Soylu, U., Sitti, M., 2017. Endo-VMFuseNet: Deep Visual-Magnetic Sensor Fusion Approach for Uncalibrated, Unsynchronized and Asymmetric Endoscopic Capsule Robot Localization Data. arXiv:1709.06041 [cs] URL: http://arxiv.org/abs/1709.06041. arXiv: 1709.06041.

Ullah, A., Ahmad, J., Muhammad, K., Sajjad, M., Baik, S.W., 2018. Action Recognition in Video Sequences using Deep Bi-Directional LSTM With CNN Features. IEEE Access 6, 1155–1166. doi:10.1109/ACCESS.2017.2778011.

Visentini-Scarzanella, M., Sugiura, T., Kaneko, T., Koto, S., 2017. Deep monocular 3D reconstruction for assisted navigation in bronchoscopy. International Journal of Computer Assisted Radiology and Surgery 12, 1089–1099. URL: https://doi.org/10.1007/s11548-017-1609-2, doi:10.1007/s11548-017-1609-2.

Wang, C., Oda, M., Hayashi, Y., Villard, B., Kitasaka, T., Takabatake, H., Mori, M., Honma, H., Natori, H., Mori, K., 2020. A visual SLAM-based bronchoscope tracking scheme for bronchoscopic navigation. International Journal of Computer Assisted Radiology and Surgery 15, 1619–1630. URL: https://doi.org/10.1007/s11548-020-02241-9, doi:10.1007/s11548-020-02241-9.

Zhao, C., Shen, M., Sun, L., Yang, G.Z., 2020. Generative Localization With Uncertainty Estimation Through Video-CT Data for Bronchoscopic Biopsy. IEEE Robotics and Automation Letters 5, 258–265. doi:10.1109/LRA.2019.2955941.