



HAL
open science

Vertebrae localization, segmentation and identification using a graph optimization and an anatomic consistency cycle

Di Meng, Edmond Boyer, Sergi Pujades

► **To cite this version:**

Di Meng, Edmond Boyer, Sergi Pujades. Vertebrae localization, segmentation and identification using a graph optimization and an anatomic consistency cycle. *Computerized Medical Imaging and Graphics*, 2023, 107, pp.102235. 10.1016/j.compmedimag.2023.102235 . hal-04397445

HAL Id: hal-04397445

<https://inria.hal.science/hal-04397445v1>

Submitted on 16 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike 4.0 International License

Vertebrae localization, segmentation and identification using a graph optimization and an anatomic consistency cycle

Di Meng*, Edmond Boyer, Sergi Pujades

Inria, Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, France

Abstract

Vertebrae localization, segmentation and identification in CT images is key to numerous clinical applications. While deep learning strategies have brought to this field significant improvements over recent years, transitional and pathological vertebrae are still plaguing most existing approaches as a consequence of their poor representation in training datasets. Alternatively, proposed non-learning based methods take benefit of prior knowledge to handle such particular cases. In this work we propose to combine both strategies. To this purpose we introduce an iterative cycle in which individual vertebrae are recurrently localized, segmented and identified using deep-networks, while anatomic consistency is enforced using statistical priors. In this strategy, the transitional vertebrae identification is handled by encoding their configurations in a graphical model that aggregates local deep-network predictions into an anatomically consistent final result. Our approach achieves the state-of-the-art results on the VerSe20 challenge benchmark, and outperforms all methods on transitional vertebrae as well as the generalization to the VerSe19 challenge benchmark. Furthermore, our method can detect and report inconsistent spine regions that do not satisfy the anatomic consistency priors. Our code and model are openly available for research purposes¹.

Keywords: Spine, Vertebrae, Segmentation, Labelling, Transitional vertebrae

1. Introduction

Segmenting, identifying and localizing vertebrae in CT images play an essential role in many clinical applications such as the spinal deformities assessment (Forsberg et al., 2013) or computer-assisted surgical interventions (Knez et al., 2016; Merloz et al., 1998). Each of these tasks is faced with challenges. For instance neighboring vertebrae share similar shapes which makes their identification uncertain. Moreover pathological spines can present abnormal shapes or the number of transitional vertebrae can be different among patients. These transitional vertebrae, i.e. the absence of T12 or occurrence of T13 or L6, are common and are reported to affect between 15% and 35% of the general population (Carrino et al., 2011; Uçar et al., 2013; Konin and Walz, 2010). However, they only impact one vertebra among the 24 in a spine and their occurrence percentages in standard vertebrae datasets appear to be therefore much lower, down to a few percentages in practice.

As a consequence, in the state-of-the-art methods, vertebrae segmentation and labelling are treated either independently, assuming one of the others is known, or with a chain of modules that address the tasks in a sequential manner (Fig. 1). The latter benefits from some partial




FIGURE 1: Existing methods vs. Ours. "I" stands for identification; "S" stands for segmentation; "L" stands for localization.

consistency between the tasks, but the errors are propagated through the modules. Vertebrae detection, usually as the first step of the pipeline, can be affected by the pathology, or metal implants or simply the prediction failures. Missing detections therefore cause empty segmentation and inconsistent labelling. Vertebrae localization, segmentation and identification are intrinsically interdependent and a common issue comes from inconsistencies that can propagate between them. In this work, we propose to cycle through the three tasks (Fig. 1) : iteratively localize, segment and identify individual vertebrae while enforcing the global anatomic consistency.

The current state-of-the-art methods, which are mostly based on deep networks trained over standard datasets, tend to experience important performance drops in the presence of transitional vertebrae and pathological cases.

*. Corresponding author :

Email address: di.meng@inria.fr (Di Meng)

1. https://gitlab.inria.fr/spine/vertebrae_segmentation

Alternatively, non-learning based methods take benefit of the prior knowledge on the full structure of the spine to handle such particular cases. In this work, we investigate the combination of such a strategy with a deep learning approach. Precisely we propose to enforce the global anatomic consistency with statistical priors and achieve the local individual prediction accuracy with the deep networks. On the one hand, in order to locate the pathological vertebrae which present non-standard shapes, we leverage the learned statistics of vertebrae volume and inter-vertebral distances. The vertebrae are estimated from the connected components of the spine mask and from the intervals of the detected vertebrae. These two shape representations quantify a part of the vertebra shape, which exhibit more robustness to pathological cases than shape or appearance models. On the other hand, in order to handle transitional vertebrae in the identification, we propose to encode the admissible configurations in a graphical model. Such a model leverages local deep-network predictions and aggregates them into an anatomically consistent result. Our experiments demonstrate that this strategy successfully handles local inaccurate predictions. Therefore it performs better than other methods with transitional vertebrae while providing state-of-the-art results on standard benchmarks such as VerSe20.

The contributions of this paper are summarized as follows :

- We propose an anatomic consistency cycle for the unified task of vertebrae localization, segmentation and identification. The three tasks cooperate and complement each other, achieving an anatomically coherent result. In case the anatomic consistency criteria were not met, the region of inconsistency is reported.
- We propose to leverage the deep neural networks with the statistical shape priors, especially the soft shape representations other than the image or appearance model. The rare cases, such as pathologies which are hard to be learned with the limited representation in the training set, are well detected.
- We extend the work on individual vertebra classification (Mohammed et al., 2020) and show significant improvement by including neighboring information. We experimentally quantify that solely using the shape of the vertebrae achieves higher or comparable performance than using as input the CT image.
- We introduce a graph to enforce the global consistency of the predicted individual vertebrae labels. It allows to detect anatomically rare cases, e.g. the transitional vertebrae T13, L6 and the absence of T12.

2. Related work

We start by reviewing the vertebrae labelling literature, addressing the localization and identification task. Then

we review the multi-label segmentation methods which in addition includes the vertebrae segmentation.

2.1. Vertebrae labelling

For vertebrae localization and identification, traditional methods build on shape models (Cai et al., 2015; Lindner et al., 2014) or graphical models (Glocker et al., 2012; Bromiley et al., 2016) to detect landmarks. In (Glocker et al., 2012) vertebrae locations are regressed by making assumptions on the shape appearance. Such models are powerful since they encode prior knowledge about the full structure of the spine. However, surface (Klinder et al., 2009; Aslan et al., 2010; Cai et al., 2015; Meng et al., 2020) or image based appearance models (Glocker et al., 2012; Bromiley et al., 2016), which make assumptions on exact shapes, are not robust to pathological cases (such as fractures or metal implants). Relative positions of vertebrae (Glocker et al., 2012, 2013; Wang et al., 2021) are shown to be less sensitive to these cases. In our work we propose to learn statistics on vertebrae relative positions and volumes. The volume quantifies a *part of the vertebra shape* that is robust to morphological deviations. Furthermore we learn statistics conditioned on the spine level (cervical, thoracic, lumbar), increasing hence their performance. Moreover we relax the need for a fixed threshold to detect vertebrae (Payer et al., 2020) and compute adaptive thresholds that automatically vary across spine levels and patients.

More recent methods build on machine learning and convolutional neural networks (CNN) in which hand-crafted features tend to be replaced by learned ones (Glocker et al., 2013; Chen et al., 2015). Sekuboyina et al. (2018) combine information across several 2D projections using a butterfly-like architecture and encode the local spine structure as an anatomic prior with an energy-based adversarial training. Liao et al. (2018) and Qin et al. (2020) develop a multi-label classification and localization network using FCN and residual blocks. They improve the classification branch with bidirectional recurrent neural network (Bi-RNN) to encode short and long range spatial and contextual information. In (Wang et al., 2021) a nnU-net keypoint detection model is trained to estimate 26 vertebrae activation maps (including the sacrum). Chen et al. (2019) propose a FCN variant that localizes vertebrae in original resolution and classifies them in down-sampled resolutions. These networks (Tao et al., 2022) directly output vertebrae locations and identifications and can be trained in an end-to-end manner. All these methods have difficulties with transitional vertebrae. The scarcity of data with rare cases makes the end-to-end learning approaches under-perform in their detection.

Other works in the literature perform vertebrae labelling in multiple stages. McCouat and Glocker (2019) employ separate CNNs to localize the vertebrae in 3D samples and identify them in 2D slices with a two-stage method. In order to find vertebrae, Jakubicek et al. (2020) detect first the spine and then track the spinal cord based on

the combination of a CNN and a growing sphere method. Payer et al. (2020) follow a similar strategy of coarse to fine vertebrae localization by detecting the spine in lower resolutions and each vertebra in higher resolution images. For identifying vertebrae, Chen et al. (2015) propose a joint learning model (J-CNN) that can classify the vertebrae labels and encode the pairwise conditional dependency at the same time. In our work, we take a two-step approach to localize and identify vertebrae. To identify individual vertebra we build on Mohammed et al. (2020) and use the segmented vertebrae shape instead of the raw CT image.

Since raw predictions from networks are not always accurate, various post-processing stages have been adopted to constrain vertebrae locations based on the anatomical ordering. Chen et al. (2015) use a shape regression model to correct the offsets from the deviation in the vertical axis. The model assumes that the coordinates distribution can be described by a quadratic form and is limited to coordinates on the vertical axis. Yang et al. (2017) introduce a chain-structure graphical model to depict the spatial relationship between vertebrae and regularize their locations with an L_1 norm to learn the best sparse representation. Chen et al. (2019) model the score maps interpolated along the 1-D spinal curve with a Hidden Markov Model to generate the optimized 1-D coordinates. Mader et al. (2018, 2019) learn an optimal conditional random field as a spatial regularizer to select a global optimal configuration over all the landmarks. In our approach, we contribute a new graphical model to select a global optimal configuration among all the individual predictions, within which the transitional vertebrae are explicitly modeled.

2.2. Multi-label vertebrae segmentation

For the segmentation task, earlier methods mainly used statistical models (Klinder et al., 2009; Rasoulian et al., 2013), active shape models (Graham et al., 1995; Benjeloun et al., 2011; Al Arif et al., 2016), graph cuts (Aslan et al., 2010), and level sets (Tsai et al., 2003; Lim et al., 2014). Recent learning based approaches have nevertheless demonstrated better performances. Al Arif et al. (2017) use a deep convolutional neural network for cervical vertebrae segmentation and introduce a shape-aware term in the loss function. The approach is semi-automatic and the vertebrae locations are given manually. Sekuboyina et al. (2017) propose a two-stage approach for lumbar vertebrae segmentation in which they first regress the bounding box of the lumbar region and then segment and label each vertebra locally. In a similar fashion, we compute a binary segmentation of the full spine that is used as strong support in locating the region of interest.

Several methods achieve very good performances by assuming that a specific part of the spine is observed (Janssens et al., 2018; Al Arif et al., 2018a,b). However, these methods are not robust to arbitrary field of views. Lessmann et al. (2019) handle this issue by iteratively applying convolutional networks in the CT images. The vertebrae are segmented and identified as they are progressi-

vely found with a sliding window. This method relies on the first detection and hardly recovers from failures that occur during this first stage. To address this limitation, Altini et al. (2021) propose a semi-automatic method that requires the user to provide the identity of the first top vertebra and the total number of vertebrae.

Masuzawa et al. (2020) propose a multi-stage framework where, first, the bounding boxes of cervical, thoracic and lumbar vertebrae are found, then the vertebrae in each bounding box are segmented and identified in an iterative manner. This approach reduces the intra-class variance from identifying one out of 24 vertebrae into identifying one out of 7 cervical, 12 thoracic and 5 lumbar vertebrae. Similarly, Mohammed et al. (2020) quantify that a two-step approach has a higher identification rate. In our work we leverage this two-step approach and extend the work of Mohammed et al. (2020) by including neighbouring vertebrae as input.

Payer et al. (2020) use a three-step approach to first localize the spine, then simultaneously locate and identify each vertebra and finally segment each individual vertebra. A fixed threshold is used for missing vertebrae. Kim et al. (2021) propose a very similar approach to detect and segment vertebrae in X-ray images. In our method we propose an anatomic consistency cycle enforcing a coherent final result. The segmentation module builds on the segmentation network of Payer et al. (2020) and the refinement strategy of Lessmann et al. (2019).

3. Method

We propose to cycle through the three tasks of localization, segmentation and identification in order to enforce anatomic consistencies among them (Sec. 3.1), namely consistencies among locations and segmentations (Sec. 3.2) and their identifications (Sec. 3.4). The overview of the method is shown in Fig 2.

The input to our method is a CT volume with an arbitrary field of view, resolution and anatomic position, potentially imaging a human spine. The CT volume is resampled into an isotropic resolution of $1 \times 1 \times 1\text{mm}^3$ and re-oriented into the same anatomic orientation where the networks and statistical priors are learned. The obtained result will be transformed back to the original orientation and resolution.

3.1. Anatomic Consistency Cycle

Our method starts by segmenting a global *spine mask* (Sec. 3.3.1) which defines the region where the vertebrae are. The anatomical constraints (Sec. 3.2) comparing the individual and the global vertebrae masks allow to find locations where, potentially, individual vertebra are missing. Then, given these locations, additional vertebra can be segmented by the individual vertebra segmentor (Sec. 3.3.2). Their location and segmentation are then refined with an iterative location-segmentation refinement

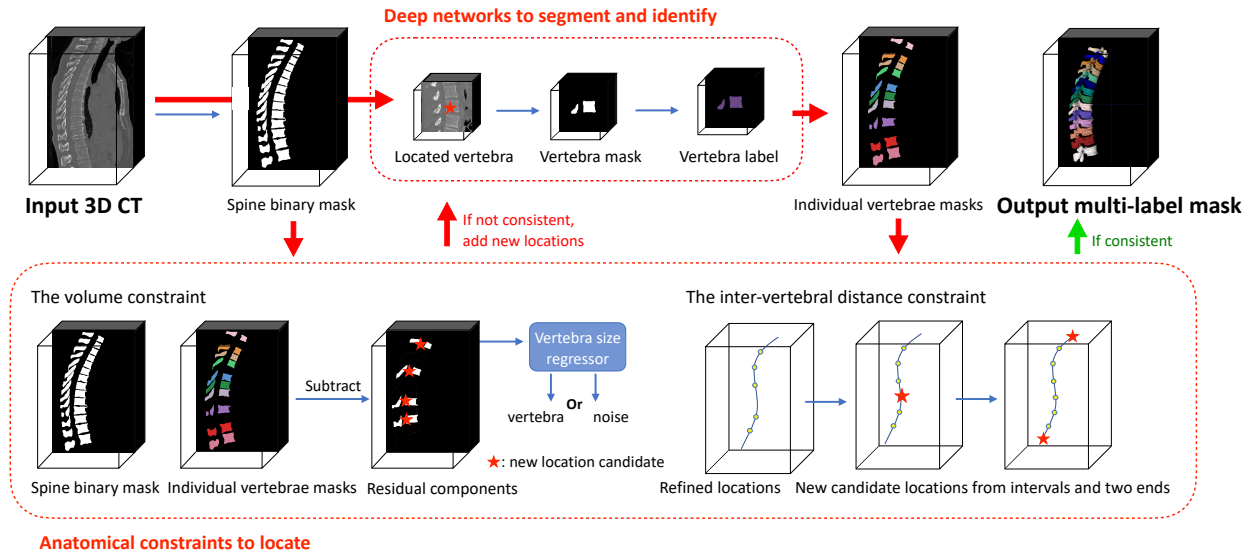


FIGURE 2: The method overview. Given a 3D CT as input, the spine is segmented supporting as a reference of where all the vertebrae are located. The anatomical constraints are applied leveraging the spine binary mask to locate individual vertebrae. Each vertebra is then segmented and identified using the deep neural networks. The set of individual vertebra location, segmentation and identification is cycled through until the consistency criteria are met. (Arrows in bold indicate the flow of the cycle)

scheme (Sec. 3.3.3). The vertebra is further identified using the location and segmentation mask (Sec. 3.4). The obtained identifications allow, in turn, to enforce finer anatomic consistency constraints (Sec. 3.2) and to detect possible new candidate locations. In the latter case, the new vertebrae go through the segmentation and identification steps described before. The process ends when the proposed consistency criteria are satisfied. It also stops when the set of detected locations, segmentation masks and identifications does not change over the cycle. The remaining inconsistencies, such as a failure in the location-segmentation refinement or an inconsistency in the anatomic constraints, are reported in the result.

3.2. Anatomic consistency constraints for vertebrae localization

Vertebrae are naturally ordered in the spine and their relative locations and consecutive sizes are heavily correlated. With the objective to exploit such prior information we compute statistics over individual vertebrae volumes and inter-vertebral distances. Given these priors we enforce then two consistency criteria. First, the *spine mask* should be similar to the union of individual vertebrae masks. The segmentation of a *spine mask* is, in general, a straightforward task which allows identifying the area where individual vertebrae should lie. In practice, the vertebra location is explored by computing the mass center of the *spine mask* at the first cycle. At each following cycle, the difference between the *spine mask* and all individual vertebrae masks is computed to obtain residual connected components. Vertebrae volume statistics are then considered to decide whether a residual is a vertebra or just noise. The locations are estimated from the mass centers of the residual connected

components. We purposely chose to use vertebrae volumes as they are more robust to morphological deviations than surface or image priors. Second, distances between all detected locations should follow the inter-vertebral distances statistics. New locations are estimated from the intervals. **The vertebrae volume constraint.** Given a residual connected component, the idea is to check its volume by considering the volumes of the neighboring vertebrae. To this purpose a regressor is trained over consecutive vertebrae to predict the volume of the previous or next vertebrae along the spine. The learned coefficients are in Sup Mat.

If the residual volume size is larger than a fraction of the predicted size (50% in our experiments), it is considered at least a vertebra and its location at its mass center is added to the vertebrae list. The residual is otherwise discarded as noise. When the vertebra identification is available, one can condition the regressors by spine levels (cervical, thoracic or lumbar), which improves the accuracy of the prediction. Although the described volume constraint allows locating vertebrae, it should be noticed that if the *spine mask* presents failures, i.e. on abnormal vertebrae, locations can still be missed. It then motivates the inter-vertebral distance constraint.

The inter-vertebral distance constraint. This constraint builds on the fact that distances between locations are well structured. To this aim, we use two statistical models of the distance between vertebrae. The first is a Gaussian distribution for each anatomic group capturing the mean and variance of the distances between consecutive vertebrae. It is used as a prior to detect abnormally large distances. In addition, linear regressors are used to predict the inter-vertebral distance from the neighbouring

ones. In practice they predict the inter-vertebral distance from : i) the previous vertebrae distance, ii) the next one and iii) using both sides. These regressors are also conditioned on each anatomic group. Details can be found in Sup Mat. When a larger inter-vertebral distance is detected, new candidate locations are added in between. The number of new candidates is adapted depending on the current distance and the predicted one.

We also leverage identification to check if the spine extremes are complete. If C1 (or L5) is not yet found, a location is predicted up (or down) using the two most top (down) locations. If it is inside the image field of view, it is added.

3.3. Spine and vertebra segmentation

To segment bones (i.e. the full spine or individual vertebrae) we use a 3D Attention U-net. The backbone architecture is a 3D U-net (Ronneberger et al., 2015) with attention blocks (Oktay et al., 2018) embedded in the skip connections. The network input is a 3D CT image patch and the output is a probability mask for each voxel indicating if it belongs to the bone or not.

To train the network we use a regression loss \mathbf{L} that combines the Dice coefficient and the Mean Squared Error, the objective being to minimize the difference between our prediction and the ground truth. We note the input x , the output $G(x)$, the ground truth patch as y and define

$$\mathbf{L}(G(x), y) = 1 - \frac{2|G(x) \cdot y|}{|G(x)|^2 + |y|^2} + \lambda \|y - G(x)\|_2, \quad (1)$$

3.3.1. Spine segmentation.

As the input CT has an arbitrary size and the network accepts only a fixed size input, we use patches of size $96 \times 96 \times 96$. In Eq. 1 λ is empirically set to 10 and kept constant during the training. At inference time, the trained model is applied in a sliding mode over the whole CT volume with a stride length of 24 (one quarter of the patch size) and for each voxel we obtain 64 predictions. We aggregate the results in each voxel by averaging the 64 probabilities and binarize the result with a 0.5 threshold.

3.3.2. Individual vertebra segmentation.

The goal of this step is to obtain an individual vertebra segmentation from a detected vertebra location. We use as input a crop of the CT image centered at the given 3D vertebra location with an auxiliary channel in which the location is encoded (Payer et al., 2020). This channel contains the 3D location converted into a 3D probability map using a Gaussian kernel with $\sigma = 20$. The output is a one-channel probability map which is then binarized with a 0.5 threshold to obtain the final binary mask. The input patch size is $128 \times 128 \times 128$ which allows to cover any full shaped vertebra in an arbitrary orientation. The λ in Eq. 1 is set to 20 during training. Although the input is in PIR orientation, some vertebrae, especially cervical, have

important rotations around the Z-Axis. To make the network robust to these cases at inference time, we augment the training data with vertebrae locations translated with an offset $t \in [-10, +10]^3$ mm as well as vertebrae rotated by $[\min(-50, \theta - 50), \max(50, \theta + 50)]$ degrees around the Z-Axis, where θ is the vertebra rotation w.r.t. the axial plane. The angle θ is computed using the current and next vertebrae locations. Each sample is augmented 162 times.

3.3.3. location-segmentation refinement scheme.

The locations from the anatomic constraints (Sec. 3.2) are approximate locations. To refine the locations, we follow a similar idea as in (Lessmann et al., 2019). They propose an iterative approach in which a segmentation patch is navigated towards the center of the segmented vertebra until a complete vertebra is seen. In our method, the approximate location already anchors the segmentor to the target vertebra. A refined location can in turn be computed from the mass center of the segmented mask. Through the iterative process, both the location and segmentation get stable and consistent. If the the segmentation mask is empty, the initial location is still added, as statistical priors are given a higher priority to deal with the network failure. At the same time, we report the inconsistency of potential surgical implants or pathology.

3.4. Vertebrae identification

To identify vertebrae we propose to combine both a local and global reasoning. The local reasoning is based on an individual classification network that predicts one of the 24 possible labels for each vertebra. On top of it, the global reasoning assumes that consecutive vertebrae are *sorted* in the spine and aggregates the local predictions into a global, consistent identification.

3.4.1. Individual vertebra classification.

Our individual vertebra classification module builds on (Mohammed et al., 2020) where a 3D VGG network is used to classify a single vertebra into one of the 24 categories. The input of the network is a binary *patch* of size $128 \times 128 \times 128$ that contains a vertebra segmentation mask. The output is the N dimensional vector of all class probabilities. The prediction proceeds in two stages. First, a spine level, i.e., cervical, thoracic or lumbar vertebra, is predicted ($N = 3$). Second, the individual vertebra goes into a per spine level classifier to predict the identity within the level ($N = 7/12/5$). As shown in (Mohammed et al., 2020), the hierarchical method obtains a higher accuracy than directly classifying into one of the 24 categories.

In our work, we extend this method (Mohammed et al., 2020) by adding the neighboring vertebrae to the input, while keeping the target vertebra in the center of the *patch*, which effectively adds more context for the prediction. At inference time, two predictions are performed : one with the *spine mask* and a second with the union of all individual vertebrae segmentations. Both predictions are averaged. The final prediction has the probability $Pv_i^j \in [0, 1]$,




FIGURE 3: Vertebrae identification global graph. An example with 4 vertebrae. Edges in orange are regular connections between nodes. Edges in blue connecting T12 to T12, and L5 to L5, allow for the presence of T13 and L6. The edge between T11 and L1 allows for T12 to be absent. The result with red edges corresponds to a spine where T10, T11, T12 and T13 are observed. (Best viewed in color)

with $j \in [1, 24]$. When an individual segmentation mask is empty, e.g. metal implants can disturb the individual segmentor, a uniform distribution over all vertebrae is used ($Pv_i^j = 1/24$). In order to deal with transitional vertebrae, such as T13 and L6, for which only few samples appear in the training set - 1 T13 and 13 L6 in practice - we consider T13 as T12 and L6 as L5 and leave their fine identification to the global reasoning.

3.4.2. Global graph optimization.

Some vertebrae can still be misidentified by the local approach. In order to improve the classification, we propose a global strategy that enforces the natural ordering of vertebrae, i.e. consecutive vertebrae should have consecutive labels. We cast the problem as a shortest path search in a graph which identifies an optimal set of labels, as illustrated in Fig. 3. Given a list with n consecutive locations, a graph with $24 \times n$ nodes is created. Each of the n columns represents a located vertebra, and each of the 24 rows represents a vertebra class. Each node carries a unary cost and each edge carries a binary cost. An optimal path is chosen by minimizing the sum of the unary costs and binary costs, as shown in Eq. 2. We populate the unary cost of each node with the cost obtained from the individual vertebra classification probability (Sec. 3.4.1) ($1 - Pv_i^j$). Additionally, as the 3-group predictions are robust, we also add a group cost Ug_i^j to each unary to prevent group swapping. It is set to 5 for the predicted group and 0 for else. The edges encode the fact that two consecutive vertebrae need to have consistent classes. In practice, node N_i^j with $i \in [1, n]$ and $j \in [1, 24]$ is only connected with N_{i+1}^{j+1} . These edges are shown in orange in Fig. 3. The binary cost B_i is set to 1 for consecutive nodes and infinity for else, which effectively enforce that consecutive vertebrae get consecutive labels. Three configurations require special attention in this process, which correspond

to transitional vertebrae : i) The presence of T13, ii) The lack of T12 and iii) The presence of L6. To handle the presence of T13, an edge is added between two consecutive T12 nodes. For the lack of T12, an edge is added between two consecutive nodes belonging to T11 and L1. The presence of L6 is dealt with an edge between two consecutive L5 nodes. These three special edges are shown in blue in Fig. 3 and are given a higher cost. To complete the graph, two extra nodes with no cost are added to be used as the source (src) and the destination (dst) in the optimization. The src node is connected to all nodes of the first vertebrae and the dst node is connected to all nodes of the last vertebrae.

$$path_{src \rightarrow dst} : \text{minimize} \left(\sum_{i=1}^n [(1 - Pv_i^j) + Ug_i^j] + B_i \right) \quad (2)$$

Once the graph is created and the costs are populated, we compute the shortest path in the graph using the classical Dijkstra algorithm (Dijkstra et al., 1959). As a post process, we check if repeated instances of T12 (or L5) are obtained and adjust the class of the second node to T13 (or L6) accordingly.

4. Experiments

We evaluate our method using the VerSe20 MICCAI challenge dataset (Löffler et al., 2020; Sekuboyina et al., 2020, 2021), which is currently the largest such spine dataset. It is split into a public training set, a public testset and a hidden testset. Each set consists of 100 CT scans. We follow the challenge setting using the training set data and the associated annotations to set up our method : 80 random CT scans from the training set are used to train the networks and estimate the anatomical prior statistics, the remaining 20 for validation. The setting is applied throughout all experiments.

4.1. VerSe20 challenge results.

We evaluate our method on the VerSe20 Challenge public and hidden testsets with 200 CT subjects in total and adopted the metrics of the challenge evaluation protocol (Sekuboyina et al., 2021) : The *Dice coefficient (DSC)* and the *Hausdorff Distance (HD)* evaluate the correctly identified segmentation performance in terms of voxel and surface similarity respectively ; The *Identification rate (ID rate)* and the *Mean Localization Distance (MLD)* evaluate the accuracy of the labelling task, where MLD measures the Euclidean distance of the predicted location to the GT location and 20 mm is the criterion for a valid identification. In the challenge metrics, when a vertebra is missed, MLD and HD are not computed. Thus there is a trade-off between ID rate and MLD (same for DSC and HD). While more positive detections lead to a better ID rate, a worse MLD can be obtained if the detections are less accurate.

	public testset				hidden testset			
	labelling results		segmentation results		labelling results		segmentation results	
	ID rate(%)	MLD(mm)	DSC(%)	HD(mm)	ID rate(%)	MLD(mm)	DSC(%)	HD(mm)
Zhang A.	94.93	2.99	88.82	7.62	96.22	2.59	89.36	7.92
Payer et al. (2020)	95.06	2.90	91.65	5.80	92.82	2.91	89.71	6.06
Chen et al. (2020)	95.61	1.98	91.72	6.14	96.58	1.38	91.23	7.15
ours	96.59	2.21	92.53	7.03	95.38	2.43	91.11	6.69
ours(a)	96.02	2.22	92.06	6.82	94.93	2.34	90.85	6.53
ours(b)	93.61	3.10	89.76	7.80	89.41	3.81	85.45	8.39

TABLE 1: Quantitative comparison of our method with the top ranked methods on the VerSe20 challenge (Sekuboyina et al., 2021) as well as with alternative (a) in which the anatomic consistency constraints are not conditioned on the identification and alternative (b) where the anatomic consistency priors are not used for detection.

	public testset	hidden testset
ours	91.04	89.70
Payer et al. (2020)	85.96	89.59
Zhang A.	87.15	87.35
Chen et al. (2020)	84.21	87.01

TABLE 2: Methods evaluation (Dice score %) on transitional vertebrae. Sekuboyina et al. (2021)

Table 1 shows the quantitative results of the proposed method and the top-scoring methods on the benchmark. Our method is on par with the best performing method in the leader-board (Chen et al., 2020) while achieving the highest accuracy on transitional cases, as shown in Table 2, which is our main objective. In that respect, we observe that Chen et al. (2020) did not win the challenge, as one important criterion in the challenge was the performance in handling the transitional vertebrae - 6 cases with T13 (2/2/2 in Train/Public/Hidden), 47 cases with L6 (15/15/17) and 8 cases with absent T12(3/4/1). Following the challenge guidelines we computed the Dice score only on scans with transitional vertebrae and the obtained results are presented in Table 2. Our method consistently outperforms all existing methods.

Figure 4 displays qualitative results of the method. Severely fractured vertebrae (T8, T9 and L1) are well segmented and identified in Fig. 4-(a). Metal-inserted vertebrae are successfully handled in Fig. 4-(e,f). In (b), the location-segmentation refinement (Lessmann et al., 2019) fails due to the occurrence of metal and T4 is not segmented. Thanks to the anatomic priors, T4, T5 and T6 can still be properly located and identified. The cross next to the label indicates the detected inconsistency. The transitional vertebrae such as L6 and T13 are detected in Fig 4-(c,d).

To gain more insights into the limitations of our method, we inspected the results that obtain Dice scores lower than 90% and observed that most failures are caused by wrong identification. Two typical failure scenarios are observed. First, all the vertebrae are identified with one label shifted. The confusions are in thoracolumbar area. To investigate if the identification error is from the individual

vertebra prediction or the graph optimization, we inspect the probabilities of the individual prediction. We observe that the individual vertebra classification network holds high confidence in its outputs. This implies that either the network is not capable enough to extract the implicit features of each vertebra, or the network is not capable to classify and differentiate them. The limitation points to the network architecture design, where we use a baseline backbone architecture *3D vgg16* (Sec. 3.4). Further improvements can be made towards the optimization of the individual vertebra classification network.

Second, the identification fails on challenging transitional vertebra and the consequent labels are shifted. The graph explicitly models the transitional vertebrae. Nevertheless, it is still a challenging task. One possible reason for the failure is the imbalance of the presence of the transitional vertebrae in the training and test sets. Transitional vertebrae are rare in the dataset. Moreover they are unproportionally distributed in the training and test sets to challenge the algorithm. In our method, the cost of bridging the abnormal transitions is higher and weighted, where the weights for the three transitions (presence of T13, L6 and the absence of T12) are turned with the training set. It does not generalize the best performance to the test sets. The improvement can be explored in learning the weights for the transitional vertebrae and the other individual vertebrae.

4.2. Ablation study.

We perform ablation studies on the proposed vertebrae identification and the anatomic consistency cycle.

4.2.1. Vertebrae identification.

To evaluate the benefit of using the context and the graph (Sec. 3.4), we compare four identification settings : (a) Using a single vertebra binary mask (Mohammed et al., 2020), (b) adding the neighbouring vertebrae, (c) adding the neighbouring vertebrae and using the graph. To compare to image based identification approaches we trained a classification method using the CT image as input (d). For a) the identification accuracy is 70.50%. When including the neighboring vertebrae masks, the accuracy is increased

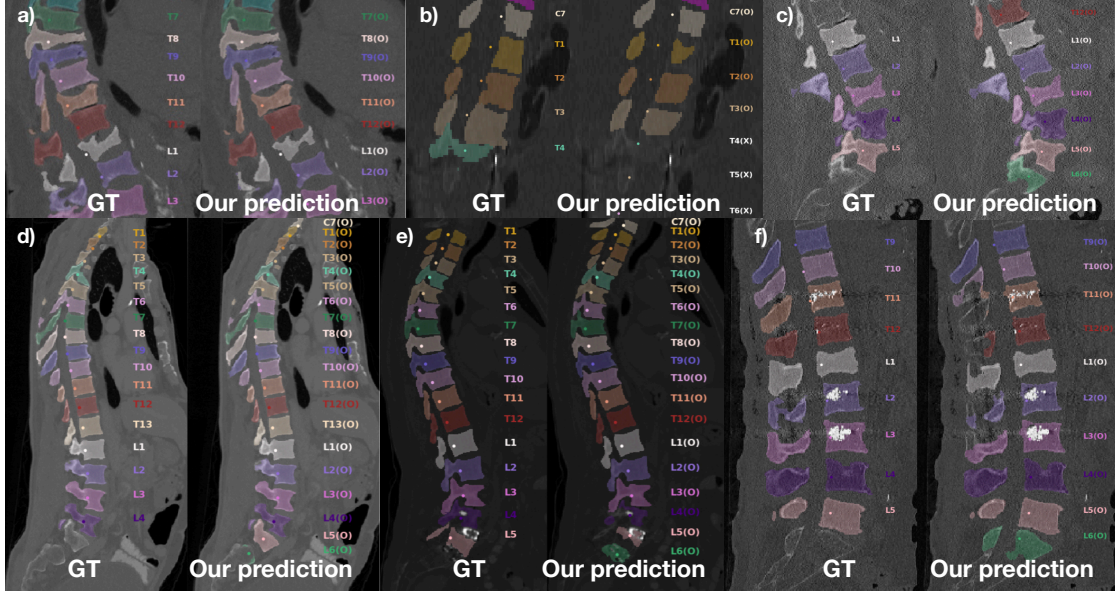


FIGURE 4: Qualitative results on fractured (a), metal-inserted (b,e,f) and transitional (c,d) vertebrae.

to 84.55%. After the graph optimization, 97.36% accuracy is obtained. When using the CT image patch as input instead of its segmentation, a lower identification accuracy (82.36%) is observed compared to the binary mask of the vertebrae (84.55%). This can be attributed to the fact that the masks provide shape discriminative cues on the vertebrae.

4.2.2. Anatomic consistency cycle.

To highlight the benefits of the different parts of the proposed anatomic priors we perform two experiments. In the first one (*ours(a)* in Tab. 1) we do not use the identification to condition the vertebrae statistics. Instead of learning specialized regressors for each spine level, we use a fixed threshold to determine if a residual mask is a potential true vertebra (we use half of the smallest vertebra volume size in the training set : $7820/2$). To decide if there is a missing vertebra given a distance between 2 we use a fixed threshold of 50 mm (Payer et al., 2020). The overall scores show that the improvement when using conditional statistics is minor but not negligible : without conditioning 8 vertebrae are missed in the public testset and 6 in the hidden test set while with the identification information no vertebra is missed in public testset and only 4 in hidden testset. In the second experiment (*ours(b)* in Tab. 1) the anatomic consistency priors (Sec. 3.2) are not used. Without the constraints on vertebrae volume and the relative locations the performance significantly drops.

5. Generalization and limitation

We further evaluate our method on other spine datasets to investigate the generalization capability and the limitation of the method. These datasets include the publicly available datasets with associated benchmark where

we can compare with the state-of-the-art approaches, and in-house datasets. We provide both quantitative and qualitative results.

VerSe19 dataset. The VerSe19 dataset (Sekuboyina et al., 2021), released as part of the VerSe challenge in MICCAI 2019, consists of 80 training CT scans with the associated annotation and 40 testing CT scans for public and hidden testsets respectively. The dataset is rich with varying field of views. To evaluate the generalization capability of our method, we perform the experiment of the VerSe20 challenge (Sekuboyina et al., 2021) where the approach developed and trained on the VerSe20 training set is applied to the VerSe19 hidden testset. Table 3 shows that our method outperforms all approaches and generalizes well to VerSe19 testset.

Lumbar vertebra segmentation CT image database. The LumbarSeg dataset consists of 10 scans of healthy subjects and associated annotation for lumbar vertebrae. The slice resolution is between 0.28 mm to 0.79 mm and the slice thickness is between 0.72 mm to 1.53 mm. We follow the evaluation protocol of Lessmann et al. (2019) and use it for an external evaluation of our supervised method. Scans from this dataset were only used for evaluation and were not part of the training set. For the segmentation we obtain Dice score (%) of 95.3 ± 0.3 which is consistent with the performance on VerSe20 dataset. For identification accuracy, we obtained 45 correctness out of 50 vertebrae. One subject failed on the *one-label-shift* issue. The L5 was recognized as L6 and the rest vertebrae above it are with one label offset. For other methods, Korez et al. (2015a) achieved on-par performance when using a different evaluation protocol, where they split the 10 CT subjects into train and test sets. Chu et al. (2015) evaluated their me-

Dataset	Method	Dice(%)	ID rate(%)	ASSD(mm)	HD(mm)
VerSe19	ours	90.84	-	-	-
	Chen D.#	86.44	-	-	-
	Payer C.#	84.11	-	-	-
	Zhang A.#	85.42	-	-	-
LumbarSeg	ours	95.3±0.3	90	0.3±0.1	7.4±1.5
	Lessmann et al. (2019)#	96.5±0.8	100	0.2±0.0	-
	Korez et al. (2015a)	95.3±1.4	-	0.3±0.1	-
	Chu et al. (2015)	91.0±7.0	-	0.9±0.3	7.3±2.2
	Ibragimov et al. (2014)	93.6±1.1	-	0.8±0.1	-
xVertSeg	ours	94.6±0.9	100	0.6±0.1	13.06±6.71
	Lessmann et al. (2019)#	94.6±2.2	100	0.3±0.2	-
	Janssens et al. (2018)	95.7±0.8	100	0.4±0.1	4.32±2.60
	Sekuboyina et al. (2017)	94.3±2.8	-	-	-
	Chuang et al. (2019)#	88.5	100	-	-
	Cheng et al. (2021)	87.7±3.5	100	-	-
CSI2014Seg	ours	93.4±0.8	-	0.7±0.1	6.6±1.2
	Kolařík et al. (2019)	97.1±0.0	-	0.3±0.1	-
	Lessmann et al. (2019)#	96.3±1.3	81	0.1±0.1	-
	Lessmann et al. (2018)#	94.8±1.6	-	0.3±0.1	-
	Cheng et al. (2021)#	95.3±1.4	-	-	4.0±2.1
	Korez et al. (2015a)	94.6±2.0	-	0.3±0.1	-
	Korez et al. (2015b)	93.1±2.0	-	-	-
Hammernik et al. (2015)	93.0±4.0	-	-	-	

same training/evaluation settings as ours

TABLE 3: Quantitative results of the state-the-of-art methods on publicly available spine CT datasets.

thod with a leave-one-out cross validation. Only vertebral bodies were segmented and evaluated. Quantitative results can be found in Table 3.

xVertSeg dataset. The xVertSeg dataset (Ibragimov et al., 2017), released as part of the xVertSeg challenge in MICCAI 2016, consists of 15 train CT volumes with ground truth segmentation of the lumbar vertebrae (into five classes, L1-L5) and 10 test CT volumes. The ground truth segmentation of the test set is not publicly available. The dataset includes non-fractured vertebrae and vertebrae with fractures of different morphological grades and cases. The scans were reconstructed to in-plane resolutions of 0.29 mm to 0.80 mm and slice thickness of 1.0 mm to 1.9 mm. The scans are with varying field of views, while mostly observing the lumbar area. There are methods that are evaluated using the same dataset, but with different evaluation/training protocols. Janssens et al. (2018) conducted a leave-three-out cross-validation study. Sekuboyina et al. (2017) used the 15 training set to train their model and evaluate the performance on test data, where the ground truth is in-house generated. Cheng et al. (2021) divided the 15 training data into two parts : 10 images for training and 5 for testing. The split information was not explicitly provided. To conduct a comparative evaluation, we use the scans 1 to 5 for evaluation and the remaining 10 scans for training, which is the same setting introduced in (Lessmann et al., 2019; Chuang et al., 2019). Quantitative results are shown in Table 3.

CSI2014 segmentation dataset. The CSI2014Seg dataset, released as part of the spine segmentation challenge in MICCAI 2014, consists of 15 CT scans of healthy subjects. Full thoracic and lumbar vertebrae (17 vertebrae in total) are observed and their corresponding segmentations are provided. The scans were reconstructed to in-plane resolutions of 0.31 mm to 0.36 mm and slice thickness of 0.7 mm to 1.0 mm. The dataset is split into training set of 10 scans and test set of 5 scans. We evaluate our method following the challenge train/test split, same as Lessmann et al. (2018, 2019) and Cheng et al. (2021). Only the 10 training data were used for other approaches. Hammernik et al. (2015) and Korez et al. (2015a,b) performed a leave-one-out 10-fold cross validation and reported the average accuracy over 10 experiments. Kolařík et al. (2019) evaluated their approach using a leave-one-out 3-fold cross validation.

EOS imaging test data. We received 3 test data from EOS imaging, France. All the three CTs are scanning the full spine from C1 to L5. Manual references are provided by their experts. The first two subjects are elder patients, both with contrast agents injected in the spinal cords. We applied the method trained on VerSe20 training set. For Dice scores (%) we obtain 85.88 ± 3.11 and 91.05 ± 1.68 respectively. For HD (mm) we obtain 7.33 ± 2.01 and 6.30 ± 2.69 respectively. The third subject is a young adult with scoliosis. We obtain Dice score (%) of 93.38 ± 2.00 and HD (mm) of 4.45 ± 4.86 . The obtained accuracy indicates that our method generalizes well to unseen datasets and




FIGURE 5: Results visualization of special cases ; (a) Vertebrae segmentation and identification of a full spine with scoliosis ; (b) An axial slice of the vertebra. Green line is the GT segmentation, red line is our prediction. The contrast agent is segmented as bone by our method. (c) Promising predictions on extremely poor-quality CT scan.

provides adequate segmentation and identification.

By inspecting the results slice by slice, we found that the contrast agents injected in the patients’ spinal cord are segmented as bones, as shown in Figure 5-(b). The brightness caused by the contrast agent presents different and new features that the training set did not observe, suggesting that our method is not robust to this contrast agent. Nevertheless, our method is well generalized to the scoliotic spine. Figure 5-(a) shows the segmented spine which presents a moderate curve (25 – 40 degrees). By virtue of the data augmentation strategy, our method is robust to the tilted vertebrae.

CHU test data. This test set is from CHU Grenoble, France. We are blind to this test set in the way that the hospital took our packed implementation and the model trained on VerSe20 training set and tested it on their own. We present the statistics of the results that we received. The test data consists of 675 CT scans associated with the manual annotations from the doctor. All the scans include the abdomino-pelvic cavity, where part of the spine is present, but are not restricted to this field of view.

Our results obtain 96.5% Dice score. A significant fact was reported that the doctor corrected their initial annotation by adding 25 L6. They are missed in the manual segmentation but detected by our method. It indicates the normality of sacralization/lumbarization in the population. Around 10% L6 are found in the test set which is coherent with the statistics (Uçar et al., 2013).

Most errors appear in the L5/Sacrum area. Our method tends to look for L6 in the field of view, some S1 which are not separated from the sacrum are being segmented and identified as L6. A few failure cases are observed in abnormal subjects or the ones with severe deformities, such as severe scoliosis and kyphosis as shown in Fig. 6. The spine with severe kyphosis has up to 90 degrees curvature, which means the axial plane of the vertebrae is orthogonal to the axial plane of the CT volume. In our al-




FIGURE 6: Our method challenges on severe scoliosis and kyphosis.

gorithm, we pre-process all the input scans into the same anatomical orientation. The spine is allowed to have some rotations (50 degrees tolerance) as we augmented the training data. The trained network is not robust to the severe scoliosis where the vertebrae are largely rotated. To lift this limitation, future work can augment the training data with larger variance.

Overall, our method achieved considerable results on this large test set. For some extreme cases such as metal-inserted vertebrae, fractured vertebrae and very poor-quality scans (Fig 5-(c)), it also provided promising outputs.

6. Discussion

We presented an anatomic consistency cycle for vertebrae localization, segmentation and identification in arbitrary field-of-view CT images. Compared with most existing methods, which use a chain of modules to perform

the three tasks, the cycling process avoids the accumulated errors from the sequential blocks. Moreover, we regard the proposed anatomic consistency cycle a general framework for multi-label segmentation in medical images, in the sense that the individual modules can be replaced. For instance, the segmentation or the identification network can be customized (Payer et al., 2020; Tao et al., 2022). We use the baseline networks such as U-net (Ronneberger et al., 2015) and vgg (Simonyan and Zisserman, 2014), as designing a new architecture is beyond the scope of this work. The framework can be generalized to other parts of the body, e.g. ribs, as the statistics of the anatomy can be encoded and learned. Additionally, the cycle with the defined consistency criteria can be used as a post-processing step to refine the result.

In the consistency cycle, we proposed to leverage the anatomical priors with deep neural networks. The deep networks are proven to achieve high accuracy on learning the large amount of data. However, it is not robust to the rare and minority cases, e.g. pathological or transitional vertebrae. While the conventional statistical methods encode the spine structure and shape priors, it shows the benefit in handling particular cases. However, it is less generalized to the large dataset. We proposed to take the advantages of both strategies. Especially, statistical priors are given a higher priority than the network prediction. For instance, the vertebra location is still added to satisfy the anatomic consistency even the network fails to segment some abnormal vertebra, i.e. metal-inserted vertebrae.

Compared with the sequential pipelines, the cycling process is less efficient but more accurate on the abnormally rare cases. It improves the low limit of the overall performance. Regarding the trade-off between the accuracy and the efficiency, the accuracy can be encouraged to be pursued ahead of the efficiency in medical imaging applications, especially for the computer-aided diagnosis. As less mistakes should be tolerated. Moreover, higher accuracy is in fact saving the time for post correction. The inference time of the method can be improved through engineering.

For the identification, we employed the vertebrae morphology as input to the classification network. Most existing works use the CT image as input to classify vertebrae with the intuition that the original scan provided more information. We experimentally showed that the vertebrae shape is representable compared with the image patch, higher/comparable accuracy was achieved in identifying the individual vertebra. The advantage of classifying vertebra from its shape is modality independent. The trained classifier is generalized and not restricted to one modality (e.g. CT or MRI) particularly. However, the prerequisite of using the shape is to segment the spine from the original scan. Vertebrae with incomplete or corrupted shape would thus affect the precision of the prediction. Even with the completely segmented vertebrae, the individual predictions may have errors. It motivated us to combine the local predictions with a global reasoning. The proposed

graph enforces the consistency over individual predictions. It tolerates a certain number of inaccurate individual predictions and sorts out an optimal configuration. Another key feature of the proposed graph is that it explicitly models the transitional vertebrae, which affect $\sim 25\%$ of the general population but are limited present in the publicly available training set. Especially for the high prevalence of the L6 case, some practitioners are unaware of them as we show for the in-house test set. When annotating data and identifying vertebrae, a reverse counting approach is used, assuming the last vertebrae is L5. With the graph, our method well handles these corner cases while maintaining a good performance on general data.

Considering the training time, our method consists of 6 networks : 4 identification networks (group, cervical, thoracic and lumbar) and 2 segmentation networks (the spine and vertebra). Training a single network takes around 22 hours on two Quadro RTX 5000 (16G each). Considering the total time and GPU cost of our method, inference times on VerSe20 public testset (103 CT scans) vary between 1.5 mins to 248 mins with a median runtime of 26 mins. Among them, the spine binary segmentation costs a median runtime of 9 mins while the cycle costs a median time of 24 mins. For the in-house dataset of 675 CT scans, the inference time varies between 7 – 90 mins ,with a median runtime of 16.2 mins. The huge difference comes from the different scan sizes. As the cycle processes the vertebra individually in a patch, more vertebrae in the view devote more inference time. Also the time is heavily devoted to the spine binary segmentation, as the spine is segmented with slidingly overlapped patches. It can be drastically improved by considering parallelization over patches. Nevertheless, the method is GPU efficient with no more than 3 GB occupation.

7. Conclusion

In this work we presented a new approach to localize, segment and identify vertebrae in CT images. The approach combines deep-learning networks with statistical consistency priors on vertebrae, improving robustness to transitional vertebrae and pathological cases. An anatomic consistency cycle is proposed that aggregates task specific deep-networks while enforcing statistical priors. The evaluation on the standard VerSe20 challenge benchmark demonstrates the interest of the proposed strategy, in particular with transitional vertebrae which are present in a meaningful part of the population. The extension of the proposed anatomic consistency cycle to other anatomic structures with similar specific cases is currently under investigation.

Acknowledgments

This work was jointly supported by the Fonds Unique Interministériel [grant number 1701595901-N00013834],

the Auvergne Rhône Alpes region and the Grenoble Alpes metropole. Sergi Pujades' work was funded by the ANR SEMBA project. The networks training in this work was performed using HPC resources from GENCI-IDRIS (Grant 2021-[AD011012208]). We thank Dr. Kazu Hasegawa for the test dataset and the EOS imaging team for the manual expert segmentations used in the assessment. We thank Charlene Dumont for testing the code and the CHU Grenoble for the test dataset with manual expert segmentations.

Références

- Al Arif, S., Gundry, M., Knapp, K., Slabaugh, G., 2016. Improving an active shape model with random classification forest for segmentation of cervical vertebrae. In : International Workshop on Computational Methods and Clinical Applications for Spine Imaging. Springer, pp. 3–15.
- Al Arif, S. M. R., Knapp, K., Slabaugh, G., 2017. Shape-aware deep convolutional neural network for vertebrae segmentation. In : International Workshop on Computational Methods and Clinical Applications in Musculoskeletal Imaging. Springer, pp. 12–24.
- Al Arif, S. M. R., Knapp, K., Slabaugh, G., 2018a. Fully automatic cervical vertebrae segmentation framework for x-ray images. *Computer methods and programs in biomedicine* 157, 95–111.
- Al Arif, S. M. R., Knapp, K., Slabaugh, G., 2018b. Spnet : Shape prediction using a fully convolutional neural network. In : International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 430–439.
- Altini, N., De Giosa, G., Fragasso, N., Coscia, C., Sibilano, E., Prencipe, B., Hussain, S. M., Brunetti, A., Buongiorno, D., Guerriero, A., et al., 2021. Segmentation and identification of vertebrae in ct scans using cnn, k-means clustering and k-nn. In : Informatics. Vol. 8. Multidisciplinary Digital Publishing Institute, p. 40.
- Aslan, M. S., Ali, A., Chen, D., Arnold, B., Farag, A. A., Xiang, P., 2010. 3d vertebrae segmentation using graph cuts with shape prior constraints. In : 2010 IEEE International Conference on Image Processing. IEEE, pp. 2193–2196.
- Benjelloun, M., Mahmoudi, S., Lecron, F., 2011. A framework of vertebra segmentation using the active shape model-based approach. *International journal of biomedical imaging* 2011.
- Bromiley, P. A., Kariki, E. P., Adams, J. E., Cootes, T. F., 2016. Fully automatic localisation of vertebrae in ct images using random forest regression voting. In : international workshop on computational methods and clinical applications for spine imaging. Springer, pp. 51–63.
- Cai, Y., Osman, S., Sharma, M., Landis, M., Li, S., 2015. Multimodality vertebra recognition in arbitrary views using 3d deformable hierarchical model. *IEEE transactions on medical imaging* 34 (8), 1676–1693.
- Carrino, J. A., Campbell Jr, P. D., Lin, D. C., Morrison, W. B., Schweitzer, M. E., Flanders, A. E., Eng, J., Vaccaro, A. R., 2011. Effect of spinal segment variants on numbering vertebral levels at lumbar mr imaging. *Radiology* 259 (1), 196–202.
- Chen, D., Bai, Y., Zhao, W., Ament, S., Gregoire, J., Gomes, C., 2020. Deep reasoning networks for unsupervised pattern de-mixing with constraint reasoning. In : International Conference on Machine Learning. PMLR, pp. 1500–1509.
- Chen, H., Shen, C., Qin, J., Ni, D., Shi, L., Cheng, J. C., Heng, P.-A., 2015. Automatic localization and identification of vertebrae in spine ct via a joint learning model with deep neural networks. In : International conference on medical image computing and computer-assisted intervention. Springer, pp. 515–522.
- Chen, Y., Gao, Y., Li, K., Zhao, L., Zhao, J., 2019. vertebrae identification and localization utilizing fully convolutional networks and a hidden markov model. *IEEE Transactions on Medical Imaging* 39 (2), 387–399.
- Cheng, P., Yang, Y., Yu, H., He, Y., 2021. Automatic vertebrae localization and segmentation in ct with a two-stage dense-u-net. *Scientific Reports* 11 (1), 1–13.
- Chu, C., Belavý, D. L., Armbrecht, G., Bansmann, M., Felsenberg, D., Zheng, G., 2015. Fully automatic localization and segmentation of 3d vertebral bodies from ct/mr images via a learning-based method. *PloS one* 10 (11), e0143327.
- Chuang, C.-H., Lin, C.-Y., Tsai, Y.-Y., Lian, Z.-Y., Xie, H.-X., Hsu, C.-C., Huang, C.-L., 2019. Efficient triple output network for vertebral segmentation and identification. *IEEE Access* 7, 117978–117985.
- Dijkstra, E. W., et al., 1959. A note on two problems in connexion with graphs. *Numerische mathematik* 1 (1), 269–271.
- Forsberg, D., Lundström, C., Andersson, M., Vavruch, L., Tropp, H., Knutsson, H., 2013. Fully automatic measurements of axial vertebral rotation for assessment of spinal deformity in idiopathic scoliosis. *Physics in Medicine & Biology* 58 (6), 1775.
- Glocker, B., Feulner, J., Criminisi, A., Haynor, D. R., Konukoglu, E., 2012. Automatic localization and identification of vertebrae in arbitrary field-of-view ct scans. In : International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 590–598.
- Glocker, B., Zikic, D., Konukoglu, E., Haynor, D. R., Criminisi, A., 2013. Vertebrae localization in pathological spine ct via dense classification from sparse annotations. In : International conference on medical image computing and computer-assisted intervention. Springer, pp. 262–270.
- Graham, J., Cooper, D., Taylor, C., Cootes, T., 1995. Active shape models their training and applications. *Computer Vision and Image Understanding* 61, 38–59.
- Hammernik, K., Ebner, T., Stern, D., Urschler, M., Pock, T., 2015. Vertebrae segmentation in 3d ct images based on a variational framework. In : Recent advances in computational methods and clinical applications for spine imaging. Springer, pp. 227–233.
- Ibragimov, B., Korez, R., Likar, B., Pernuš, F., Xing, L., Vrtovec, T., 2017. Segmentation of pathological structures by landmark-assisted deformable models. *IEEE transactions on medical imaging* 36 (7), 1457–1469.
- Ibragimov, B., Likar, B., Pernuš, F., Vrtovec, T., 2014. Shape representation for efficient landmark-based segmentation in 3-d. *IEEE transactions on medical imaging* 33 (4), 861–874.
- Jakubicek, R., Chmelik, J., Jan, J., Ourednicek, P., Lambert, L., Gavelli, G., 2020. Learning-based vertebra localization and labeling in 3d ct data of possibly incomplete and pathological spines. *Computer methods and programs in biomedicine* 183, 105081.
- Janssens, R., Zeng, G., Zheng, G., 2018. Fully automatic segmentation of lumbar vertebrae from ct images using cascaded 3d fully convolutional networks. In : 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). IEEE, pp. 893–897.
- Kim, K. C., Cho, H. C., Jang, T. J., Choi, J. M., Seo, J. K., 2021. Automatic detection and segmentation of lumbar vertebrae from x-ray images for compression fracture evaluation. *Computer Methods and Programs in Biomedicine* 200, 105833.
- Klinder, T., Ostermann, J., Ehm, M., Franz, A., Kneser, R., Lorenz, C., 2009. Automated model-based vertebra detection, identification, and segmentation in ct images. *Medical image analysis* 13 (3), 471–482.
- Knez, D., Likar, B., Pernuš, F., Vrtovec, T., 2016. Computer-assisted screw size and insertion trajectory planning for pedicle screw placement surgery. *IEEE transactions on medical imaging* 35 (6), 1420–1430.
- Kolařík, M., Burget, R., Uher, V., Říha, K., Dutta, M. K., 2019. Optimized high resolution 3d dense-u-net network for brain and spine segmentation. *Applied Sciences* 9 (3), 404.
- Konin, G., Walz, D., 2010. Lumbosacral transitional vertebrae : classification, imaging findings, and clinical relevance. *American Journal of Neuroradiology* 31 (10), 1778–1786.
- Korez, R., Ibragimov, B., Likar, B., Pernuš, F., Vrtovec, T., 2015a. A framework for automated spine and vertebrae interpolation-based detection and model-based segmentation. *IEEE transactions on medical imaging* 34 (8), 1649–1662.

- Korez, R., Ibragimov, B., Likar, B., Pernuš, F., Vrtovec, T., 2015b. Interpolation-based shape-constrained deformable model approach for segmentation of vertebrae from ct spine images. In : Recent Advances in Computational Methods and Clinical Applications for Spine Imaging. Springer, pp. 235–240.
- Lessmann, N., Van Ginneken, B., De Jong, P. A., Išgum, I., 2019. Iterative fully convolutional neural networks for automatic vertebra segmentation and identification. *Medical image analysis* 53, 142–155.
- Lessmann, N., van Ginneken, B., Išgum, I., 2018. Iterative convolutional neural networks for automatic vertebra identification and segmentation in ct images. In : *Medical Imaging 2018 : Image Processing*. Vol. 10574. International Society for Optics and Photonics, p. 1057408.
- Liao, H., Mesfin, A., Luo, J., 2018. Joint vertebrae identification and localization in spinal ct images by combining short-and long-range contextual information. *IEEE transactions on medical imaging* 37 (5), 1266–1275.
- Lim, P. H., Bagci, U., Bai, L., 2014. A robust segmentation framework for spine trauma diagnosis. In : *Computational Methods and Clinical Applications for Spine Imaging*. Springer, pp. 25–33.
- Lindner, C., Bromiley, P. A., Ionita, M. C., Cootes, T. F., 2014. Robust and accurate shape model matching using random forest regression-voting. *IEEE transactions on pattern analysis and machine intelligence* 37 (9), 1862–1874.
- Löffler, M. T., Sekuboyina, A., Jacob, A., Grau, A.-L., Scharr, A., El Hussein, M., Kallweit, M., Zimmer, C., Baum, T., Kirschke, J. S., 2020. A vertebral segmentation dataset with fracture grading. *Radiology : Artificial Intelligence* 2 (4), e190138.
- Mader, A. O., Lorenz, C., Berg, J. v., Meyer, C., 2019. Automatically localizing a large set of spatially correlated key points : a case study in spine imaging. In : *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 384–392.
- Mader, A. O., Lorenz, C., Bergtholdt, M., von Berg, J., Schramm, H., Modersitzki, J., Meyer, C., 2018. Detection and localization of spatially correlated point landmarks in medical images using an automatically learned conditional random field. *Computer Vision and Image Understanding* 176, 45–53.
- Masuzawa, N., Kitamura, Y., Nakamura, K., Iizuka, S., Simo-Serra, E., 2020. Automatic segmentation, localization, and identification of vertebrae in 3d ct images using cascaded convolutional neural networks. In : *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 681–690.
- McCouat, J., Glocker, B., 2019. Vertebrae detection and localization in ct with two-stage cnns and dense annotations. *arXiv preprint arXiv :1910.05911*.
- Meng, D., Keller, M., Boyer, E., Black, M., Pujades, S., 2020. Learning a statistical full spine model from partial observations. In : *International Workshop on Shape in Medical Imaging*. Springer, pp. 122–133.
- Merloz, P., Tonetti, J., Pittet, L., Coulomb, M., Lavallee, S., Troccaz, J., Cinquin, P., Sautot, P., 1998. Computer-assisted spine surgery. *Computer Aided Surgery : Official Journal of the International Society for Computer Aided Surgery (ISCAS)* 3 (6), 297–305.
- Mohammed, E., Meng, D., Pujades, S., 2020. Morphology-based individual vertebrae classification. In : *International Workshop on Shape in Medical Imaging*. Springer, pp. 134–144.
- Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., et al., 2018. Attention u-net : Learning where to look for the pancreas. *arXiv preprint arXiv :1804.03999*.
- Payer, C., Stern, D., Bischof, H., Urschler, M., 2020. Coarse to fine vertebrae localization and segmentation with spatialconfiguration-net and u-net. In : *VISIGRAPP (5 : VISAPP)*. pp. 124–133.
- Qin, C., Yao, D., Zhuang, H., Wang, H., Shi, Y., Song, Z., 2020. Residual block-based multi-label classification and localization network with integral regression for vertebrae labeling. *arXiv preprint arXiv :2001.00170*.
- Rasoulalian, A., Rohling, R., Abolmaesumi, P., 2013. Lumbar spine segmentation using a statistical multi-vertebrae anatomical shape+ pose model. *IEEE transactions on medical imaging* 32 (10), 1890–1900.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net : Convolutional networks for biomedical image segmentation. In : *International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.
- Sekuboyina, A., Hussein, M. E., Bayat, A., Löffler, M., Liebl, H., Li, H., Tetteh, G., Kukačka, J., Payer, C., Štern, D., et al., 2021. Verse : a vertebrae labelling and segmentation benchmark for multi-detector ct images. *Medical image analysis* 73, 102166.
- Sekuboyina, A., Rempfler, M., Kukačka, J., Tetteh, G., Valentinitzsch, A., Kirschke, J. S., Menze, B. H., 2018. Btrfly net : Vertebrae labelling with energy-based adversarial learning of local spine prior. In : *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 649–657.
- Sekuboyina, A., Rempfler, M., Valentinitzsch, A., Menze, B. H., Kirschke, J. S., 2020. Labeling vertebrae with two-dimensional reformations of multidetector ct images : an adversarial approach for incorporating prior knowledge of spine anatomy. *Radiology : Artificial Intelligence* 2 (2), e190074.
- Sekuboyina, A., Valentinitzsch, A., Kirschke, J. S., Menze, B. H., 2017. A localisation-segmentation approach for multi-label annotation of lumbar vertebrae using deep nets. *arXiv preprint arXiv :1703.04347*.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv :1409.1556*.
- Tao, R., Liu, W., Zheng, G., 2022. Spine-transformers : Vertebra labeling and segmentation in arbitrary field-of-view spine cts via 3d transformers. *Medical Image Analysis* 75, 102258.
- Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W. E., Willsky, A., 2003. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE transactions on medical imaging* 22 (2), 137–154.
- Uçar, D., Uçar, B. Y., Coşar, Y., Emrem, K., Gümüşsuyu, G., Mutlu, S., Mutlu, B., Çağan, M. A., Mertsoy, Y., Gümüş, H., 2013. Retrospective cohort study of the prevalence of lumbosacral transitional vertebra in a wide and well-represented population. *Arthritis* 2013.
- Wang, F., Zheng, K., Lu, L., Xiao, J., Wu, M., Miao, S., 2021. Automatic vertebra localization and identification in ct by spine rectification and anatomically-constrained optimization. In : *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5280–5288.
- Yang, D., Xiong, T., Xu, D., Huang, Q., Liu, D., Zhou, S. K., Xu, Z., Park, J., Chen, M., Tran, T. D., et al., 2017. Automatic vertebra labeling in large-scale 3d ct using deep image-to-image network with message passing and sparsity regularization. In : *International conference on information processing in medical imaging*. Springer, pp. 633–644.

Supplementary Material

A. Vertebrae statistics

Vertebrae volume regressor. Vertebrae volumes vary across the cervical, thoracic and lumbar while the consecutive vertebrae volumes are heavily related. We learn the statistics of the vertebrae volumes to predict one from the neighboring (previous or next) one. To be specific, for each anatomic group (cervical, thoracic and lumbar), two linear regressors are trained to predict the vertebra volume given the previous and the next vertebra respectively : $S_i = a * S_{i-1} + c_1$ and $S_i = b * S_{i+1} + c_2$. We use the identification of the neighboring vertebra to select the specialized regressor. The learned coefficients are ($a = 1.03, c_1 = 1471, b = 0.92, c_2 = 497$) for cervical, ($a = 1.03, c_1 = 1354, b = 0.94, c_2 = -140$) for thoracic and ($a = 1.05, c_1 = 981, b = 0.94, c_2 = -269$) for lumbar.

Inter-vertebral distance regressor. The vertebrae locations are consistent and well structured, we learn two forms of the inter-vertebral distance statistics. One is the Gaussian distribution for each anatomic group, the other is linear regressors predicting the inter-vertebral distance given either the previous, next or both-side distances.

For the learned coefficients of the Gaussian distributed inter-vertebral distance, we have ($\mu = 16.77, \sigma = 2.18$) for cervical, ($\mu = 23.32, \sigma = 3.55$) for thoracic, and ($\mu = 32.68, \sigma = 2.84$) for lumbar. The distance G_i is considered as the distance between two consecutive vertebrae when $\mu - 3\sigma < G_i < \mu + 3\sigma$.

To predict the inter-vertebral distance given the neighboring distances, we train three types of the regressors : using the both-side neighboring distances $G_i = m_1 * G_{i-1} + n_1 * G_{i+1} + K_1$, using the previous distance $G_i = m_2 * G_{i-1} + K_2$ or using the next distance $G_i = n_2 * G_{i+1} + k_3$. The latter two are intended for the top and bottom vertebrae who do not have the distances from both sides. The learned coefficient can be found in Table.A1.

	m_1	n_1	k_1	m_2	k_2	n_2	k_3
Cervical	0.55	0.45	-0.08	0.92	2.40	0.98	-0.13
Thoracic	0.57	0.44	-0.24	0.93	2.29	0.97	-0.07
Lumbar	0.56	0.46	-0.73	0.95	1.96	0.96	0.23

TABLE A1: Learned coefficients of the inter-vertebral distance regressors.

We then compute the Mean Relative Error (MRE) between the relative distance from the locations and the distance from the regressor. The distance is considered as normal when $\mu_{mre} - 3 * \sigma_{mre} < MRE < \mu_{mre} + 3 * \sigma_{mre}$. The learned μ_{mre} and σ_{mre} for each anatomic group can be found in Table. A2.

	both sides		previous		next	
	μ_{mre}	σ_{mre}	μ_{mre}	σ_{mre}	μ_{mre}	σ_{mre}
Cervical	9.13	2.86	10.20	2.05	12.13	3.36
Thoracic	2.42	1.43	3.96	1.56	4.17	0.88
Lumbar	2.04	0.93	4.45	1.55	5.16	1.71

TABLE A2: MRE mean and standard deviation of the inter-vertebral distance regressors.