



HAL
open science

High-level prior-based loss functions for medical image segmentation : a survey

Rosana El Jurdi, Caroline Petitjean, Paul Honeine, Veronika Cheplygina,
Fahed Abdallah

► **To cite this version:**

Rosana El Jurdi, Caroline Petitjean, Paul Honeine, Veronika Cheplygina, Fahed Abdallah. High-level prior-based loss functions for medical image segmentation : a survey. *Computer Vision and Image Understanding*, 2021, 210, pp.103248. 10.1016/j.cviu.2021.103248 . hal-03410506

HAL Id: hal-03410506

<https://normandie-univ.hal.science/hal-03410506v1>

Submitted on 22 Aug 2023

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 4.0 International License



ELSEVIER

Computer Vision and Image Understanding

journal homepage: www.elsevier.com

High-level Prior-based Loss Functions for Medical Image Segmentation: A Survey

Rosana El Jurdi^{a,d}, Caroline Petitjean^a, Paul Honeine^a, Veronika Cheplygina^{b,c}, Fahed Abdallah^{d,e}

^aNormandie Univ, INSA Rouen, UNIROUEN, UNIHAVRE, LITIS, Rouen, France

^bComputer Science Department, IT University of Copenhagen, Denmark

^cMedical Image Analysis group, Eindhoven University of Technology, Eindhoven, The Netherlands

^dUniversité Libanaise, Hadath, Beyrouth, Liban

^eICD, M2S, Université de technologie de Troyes, Troyes, France

ABSTRACT

Today, deep convolutional neural networks (CNNs) have demonstrated state of the art performance for supervised medical image segmentation, across various imaging modalities and tasks. Despite early success, segmentation networks may still generate anatomically aberrant segmentations, with holes or inaccuracies near the object boundaries. To mitigate this effect, recent research works have focused on incorporating spatial information or prior knowledge to enforce anatomically plausible segmentation. If the integration of prior knowledge in image segmentation is not a new topic in classical optimization approaches, it is today an increasing trend in CNN based image segmentation, as shown by the growing literature on the topic. In this survey, we focus on high level prior, embedded at the loss function level. We categorize the articles according to the nature of the prior: the object shape, size, topology, and the inter-regions constraints. We highlight strengths and limitations of current approaches, discuss the challenge related to the design and the integration of prior-based losses, and the optimization strategies, and draw future research directions.


© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

Medical image segmentation can be viewed as the process of making per-pixel predictions in an image, in order to identify organs or lesions from the background. Imaging modalities include mag-

netic resonance imaging (MRI), computed tomography (CT), nuclear medicine functional imaging, ultrasound imaging, microscopy, to name a few. The results of segmentation can be used to compute biomarkers or quantitative measurements, to compute three-dimensional anatomical models for image-guided surgery, and to design the radiation beam in radiotherapy planning, in order to spare healthy organs while intensifying the beam on the tumor. As such, the segmentation task is a key step in assisting early disease detection, diagnosis,

e-mail: rosana.el-jurdi@univ-rouen.fr (Rosana El Jurdi), caroline.petitjean@univ-rouen.fr (Caroline Petitjean), paul.honeine@univ-rouen.fr (Paul Honeine), v.cheplygina@tue.nl (Veronika Cheplygina), fahed.abdallah76@gmail.com (Fahed Abdallah)



(a) From left to right: neuronal image, ground truth segmentation, baseline segmentation, segmentation with prior information

(b) From left to right: ground truth segmentation of a brain lesion in MRI, baseline segmentation, segmentation with prior information

Fig. 1: Example illustrations on the relevance of incorporating topological priors in neuronal (a) and brain lesion segmentation (b). In (a), segments obtained by baseline method with no topological prior shows broken membranes and incorrect grouping of neuronal areas, contrary to prior model that preserves these topological particularities. In (b), baseline segmentation has prohibited holes which are avoided thanks to the segmentation with prior. Figures from (Hu et al., 2019) (a) and (Kervadec et al., 2019a) (b). *[permission pending]*

monitoring treatment and follow up.

Due to the recent advances in deep learning, computer vision tasks, including automated image segmentation, have experienced a major breakthrough. The main reason for the success of convolutional neural networks (CNNs) lies in their ability to learn hierarchical feature representations automatically through training directly from images, eliminating the need for handcrafted features. State-of-the-art architectures for image segmentation typically have an encoder-decoder structure (Long et al., 2015; Ronneberger et al., 2015), that allows for an end-to-end processing. One of the key component of CNN training is the loss function, as it drives the back-propagation of the error between the predicted and the reference label.

In segmentation networks, the cross-entropy (Ronneberger et al., 2015) and the Dice loss (Milletari et al., 2016) as well as their variants, are widely used. However, these loss functions do not take into account high-level features or topological properties of the objects to be segmented. As a result, CNN predictions with the above loss functions may suffer from anatomically aberrant errors, with holes, voids or high inaccuracies close to or-

gan boundaries (Fig. 1a and 1b). For example, the winner of the ACDC challenge, which focused on the automatic delineation of the cardiac structures, was an ensemble of U-Net inspired architectures proposed by Isensee et al. (2018). Even though this approach obtained the top accuracy, it was not able to prevent aberrant and anatomically impossible segmentation, in 41 patients out of 50 of the testing database (Bernard et al., 2018).

Embedding prior knowledge on the object, such as shape, appearance or location, into the deep learning networks can increase the network robustness and accuracy, while generating anatomically plausible segmentation, as an increasing number of research papers on the topic demonstrate. Prior knowledge is indeed available in the medical domain, due to conventions in patient position and similarity shared in body structure. It has been commonly exploited in variational approaches before deep learning (Nosrati and Hamarneh, 2016) and incorporating this knowledge into the loss function currently exhibits a growing research avenue. Prior-based loss functions are also helpful in weakly supervised settings, where only partial labels (e.g. scribbles or bounding boxes) are avail-

able (El Jurdi et al., 2019; El Jurdi et al., 2020; Kervadec et al., 2020). Interestingly, the design of high level prior loss functions is a research avenue that is quite specific to medical image segmentation. Medical images have their own specificity; there is much less variability in terms of object shape and appearance, compared to natural scenes. One single feature such as the size or shape characteristic cannot convey the variability and the complexity of the objects to be segmented in natural images; in addition the spatial context is extremely important in scene understanding. Research works on this topic do not integrate domain knowledge on the object shape, but have a broader strategy, more focused on the architecture design, to add contextual dependencies with attention mechanism (Yu et al., 2020) for example, or to perform shape estimation, from a set of collected shapes of the same category, with a special branch (Kuo et al., 2019).

In order to incorporate prior knowledge in a segmentation process, two major questions arise. First, one needs to define the type of prior information and the modeling of the prior. The term "prior knowledge" is quite vague and covers a breadth of notions, as can be observed from the literature. It may refer to spatial or context information, under the form of distance maps or the image gradient, to more complex knowledge, regarding the anatomy of the object of interest (such as shape and size for example) and the connectivity between regions.

Second, one needs to specify how to embed the prior into the segmentation network. While the network architecture can be modified to integrate priors, another way to incorporate prior into the segmentation frameworks is at the level of the loss function. The loss function indeed offers a versatile way to include anatomical constraints at differ-

ent scales, while maintaining interactions between regions as well as the computational efficiency of the backbone network. However, designing novel losses for CNN-based segmentation poses several specific challenges: e.g. how to translate the desired anatomical properties on the network output, which is a real-valued probability map, or how to insure differentiability or convexity of the designed prior-based loss.

The goal of this paper is to establish an overview of recent contributions that focus on incorporating prior knowledge at the loss function level, in deep networks for medical image segmentation. The main objective is to add to the understanding of the mechanisms and intuition behind the design and implementation of prior based losses. In our survey, we do not intend to cover all types of prior, but rather focus on "high-level" priors. We define high-level prior as anatomical features extracted from the object that can help in characterizing and interpreting it, by opposition to low level operator, such as gradient or distance maps. We have categorized the articles according to the nature of the prior, that may be the object shape, size, topology, and the inter-regions constraints. We seek to provide the reader with (i) what types of priors exist in the literature, how they are modeled and embedded into the loss function, (ii) the major challenges linked to the design of such prior-based losses, and (iii) their common training and optimization strategies. To do so, Google Scholar was queried for peer-reviewed publications that included tags such as constraint losses, medical image segmentation, prior-based losses, constraint neural networks and anatomical constraints. The papers were then filtered in terms of employing a deep CNN-backbone for segmentation and the novelty present in the design of a new prior loss

function.

The rest of the paper is organized as follows. We first review existing surveys in CNN-based medical image segmentation, and provide a short review of prior-based approaches in the pre-deep learning era in Section 2. Section 3 introduces the fundamentals of deep learning for medical image segmentation. Before diving into the heart of our survey, we briefly review in Section 4 approaches that make use of low-level prior information to enforce segmentation consistency, so as to position these approaches with respect to the scope of this survey. Section 5 contains the proposed categorization of prior-based loss functions, along with a review of the corresponding works. Section 6 exhibit the common challenges and weaknesses faced while designing such losses and discuss some future trends and perspectives. Finally, we conclude the survey in Section 7.

2. Related work

2.1. Surveys in medical image segmentation

Since the rise of convolutional neural networks in computer vision, various medical image segmentation surveys have been published (Hesamian et al., 2019; Jiang et al., 2017; Haque and Neubert, 2020; Razzak et al., 2018; Taghanaki et al., 2019; Havaei et al., 2016; Chen et al., 2020; Lei et al., 2020; Zhou et al., 2019; Renard et al., 2020).

In one of the most recent surveys presented by Lei et al. (2020), authors consider contributions according to the level of supervision. Within the fully supervised setting, the paper categorizes the methods according to backbone network, design of network blocks, and loss functions improvement. Some regularization losses via distance-based transforms are briefly reviewed; however,

there is no focus on high-level priors for loss functions. Other recent surveys in the field like (Renard et al., 2020) review methods with regards to their source of variability and degree of reproducibility. Zhou et al. (2019) focus on multimodal fusion strategies and approaches for medical data segmentation. Razzak et al. (2018) present a survey of medical image segmentation with deep learning and categorized methods in terms of convolutional neural network structure and training techniques (weakly-, semi- or fully-supervised). Taghanaki et al. (2019) report contributions in image segmentation for natural and medical images and categorized existing works according to six main categories: model architecture, image modality, loss functions, model types, supervision strategies and multi-task learning. Haque and Neubert (2020) present an overview of the different deep learning methods used within the biomedical image segmentation domain and categorized them according to the image modality, the segmentation architectures, and their evaluation metrics. Domain specific reviews also exist, focusing on a particular pathology or organ: (Magadza and Viriri, 2021; Chahal et al., 2020; Meyer et al., 2018; Havaei et al., 2016) focus on brain pathology, (Chen et al., 2020; Yue et al., 2019) on cardiac image segmentation and (Krithiga and Geetha, 2020; Debelee et al., 2020) on breast pathology.

As far as prior based loss functions are concerned, Taghanaki et al. (2019) briefly reviews a few works that integrate boundary and topological prior-based losses, in addition to presenting common loss functions, as well as their variants and combinations. (Chen et al., 2020) mention the idea of imposing anatomical constraints at the level of the loss function, as a solution that allows the network to learn features representing the under-

lying anatomical structures. However, this survey focuses only on cardiac image segmentation. Moreover, it includes works that combine structural and loss constraints. Thus, to the best of our knowledge, no survey deals specifically with prior-based loss functions in image segmentation.

2.2. Prior-based segmentation approaches in the pre-deep learning era

Among the segmentation methods that existed before the deep learning tsunami, optimization-based approaches have been hugely popular, due to their versatility and efficiency. They consist in obtaining the segmentation by optimizing an appropriate energy functional. In the case where the image domain is considered to be continuous, optimization-based approaches have embraced active contours, level sets, and deformable models in general (Xu et al., 2000). On the other hand, spatially discrete approaches consider the image as a graph, and include the well-known graph cut and normalized cuts approaches (Shi and Malik, 2000; Boykov and Funka-Lea, 2006), to name a few.

In order to counteract the effect of the noise, occlusion and low contrast in medical images, and to increase anatomical plausibility, researchers already used prior information to guide optimization-based segmentation algorithms, well before the advent of deep learning for computer vision. Indeed, optimization-based approaches allow to encode easily some constraints on the segmentation results. Independently from the spatial domain (whether it is spatial or continuous), the energy functional to be minimized comprises at least two terms: a data-fidelity term, related to the image to be segmented, and a regularization term that controls the appearance of the contour, e.g. smoothness. One convenient way to add prior

information is to embed it into an additional term in the objective functional or via a probabilistic formulation. The additional term contains a dissimilarity measure between the model feature and the segmented region feature. Depending on the targeted property, the additional term is designed similarly to the data term or to the smoothness term.

Prior information encompasses a breadth of various forms, as distinguished by Nosrati and Hamarneh (2016) in their review on the topic: they can be based on elementary image properties, such as intensity, color, and texture information, or more elaborate features on the object shape, such as topological and geometrical constraints (Vicente et al., 2008), moment priors (Ayed et al., 2008; Foulonneau et al., 2009), distance and adjacency constraints (Liu et al., 2008), as well as motion and model/atlas-based priors (Lorenzo-Valdes et al., 2002; Rohlfing et al., 2005). In our review, we will encounter types of prior which are similar to these. We believe that these past research works could be fruitfully explored to design new losses. For example, Mirikharaji and Hamarneh (2018) took inspiration from (Veksler, 2008) to design a star-shape prior. On the other hand, the shape template in (Slabaugh and Unal, 2005; Rousson and Paragios, 2002) or the popular statistical models in (Cootes et al., 1995; Heimann and Meinzer, 2009), are build based on ground truth segmentation maps and their corresponding images, even though they are not "learning" approaches per se. There is no such prior in the papers we review in our survey, since the CNN is trained with this type of data, and already learns some appearance and semantic information. For further information on the topic of prior information in optimization based segmentation approaches, in addition to

(Nosrati and Hamarneh, 2016), we refer the reader to Cremers et al. (2007) on shape prior segmentation in variational continuous frameworks, and (Grady, 2012) for spatially discrete frameworks.

In order to be easily optimized, newly-designed energy terms have to be convex (or submodular in spatially discrete frameworks). Interestingly, Nosrati and Hamarneh (2016) highlights the trade-off between the richness of the energy functional, and its optimizability in variational approaches: an accurate modelling of the underlying property will require highly complex or costly optimization. In deep segmentation networks, designing new loss functions includes specific challenges (developed in Section 6) that are not the same as the ones in optimization-based approaches, but our survey could also relate to the fidelity-optimizability trade-off in the reviewed papers.

3. Fundamentals of medical image segmentation with deep networks

The implementation of medical image segmentation networks require choosing an appropriate architecture, and training it so as to fix the network weights. Training is done by optimizing a loss function that models the problem at hand, here, pixel labeling. This implementation contains many degrees of freedom, among which the architecture and the loss function. We will briefly review below state-of-the-art fundamentals in medical image segmentation in this context.

3.1. Segmentation networks

One of the first CNN architectures to allow automatic end-to-end semantic segmentation is the fully convolutional network (FCN) (Long et al., 2015). The FCN has a structure that is derived from a deep classification model, such as VGG16,

AlexNet or GoogLeNet, by removing the corresponding classification layers, i.e. replacing their fully connected layers with convolutional ones, and plugging in an upsampling path that is dedicated to transforming coarse outputs into dense predictions. With its ability to extract multiscale features, FCN has set a milestone in segmentation approaches and paved the way for encoder-decoder segmentation networks. However, FCN often fails to consider global spatial information, and often result in fuzzy coarse-grained predictions (Ravishanker et al., 2017). Thus, it has given rise to many improved variants, among which U-Net.

The U-Net (Ronneberger et al., 2015) has an encoder/decoder structure, which has the particularity of being symmetric and having skip connections. The encoder part is a contracting path composed of stacked convolutional and max pooling layers, that are dedicated to capturing contextual information in order to detect objects present in the image. The decoder part is an expanding path composed of deconvolutions, also called transposed convolution, or bilinear upsampling layers, that help precise localization of patterns including contours and boundaries. As an image moves further into the contracting layers, it decreases in size but increases in depth of its learnt contextual features. In contrast, the decoder layers increase its input size but decreases its depth until it reaches its initial resolution, thus producing a segmentation map of similar dimensions to that of the input. To make use of both contextual and positional features, skip connections between the downsampling (encoder) and upsampling (decoder) paths are added. Skip connections concatenate symmetrically features from opposing convolution and deconvolution layers. Several variants of U-Net consist in changing the backbone model used for en-

coding, e.g. VGG and DenseNet, replacing deconvolution layers with super-resolution ones for more concise localization ability, or enhancing the architecture with modifications such as attention mechanisms, dense or residual connections. Since its publication in 2015, the U-Net quickly became a state-of-the-art architecture for medical image segmentation (Isensee et al., 2019).

In particular, extensions to 3D images have been proposed in the 3D-UNet model (Çiçek et al., 2016) and the V-Net model (Milletari et al., 2016). It is worth noting that several attempts have been made to revisit the U-Net by integrating some prior knowledge into the architecture, such as in (Oktay et al., 2018; El Jurdi et al., 2020); however, such architectural modifications are difficult to engineer with limited or specific use. They are not as flexible as injecting the prior at the loss function level, whose versatility comes from the plug-into-any-CNN-backbone property.

3.2. Loss functions for image segmentation

Given a training dataset of N images $\{\mathbf{x}_i\}_{i=1}^N$ and their corresponding ground truth masks $\{\mathbf{y}_i\}_{i=1}^N$, the goal is to train the segmentation network so that it can learn to approximate the “true” function f parameterized by θ , that maps the input image $\{\mathbf{x}\}$ to the predicted label map, i.e. such that $f(\mathbf{x}, \theta)$ represents a map with the label probability at each pixel. In the following, we will use $\hat{\mathbf{y}}_i = f(\mathbf{x}_i, \theta)$ and rely on Table 1 for mathematical notations.

Training the network boils down to finding the network parameters θ that minimize a loss function $\mathcal{L}(\theta)$. For sake of simplicity, we will drop the dependency in θ and denote by \mathcal{L} the loss function. The loss function \mathcal{L} reflects the problem at hand, i.e. is a data-fitting loss, that we note as \mathcal{L}_{fit} . It

is of the form:

$$\mathcal{L}_{fit} = \sum_{p \in \Omega} L(\hat{y}_p, y_p) \quad (1)$$

where L is a function that penalizes the discrepancy between the predicted pixel label (\hat{y}_p) and the ground truth label (y_p) for each pixel $p \in \Omega$. The shape of L defines how the error is computed. It is mainly derived from common norms, such as the Euclidean norm or the log-norm (cross entropy shape) for example. When based on a norm, the loss is continuous and differentiable, which allows it to be efficiently optimized during back-propagation. Properties of the loss shape are important to translate task specifications. Symmetry, for instance, ensures equal penalization between errors caused by over-segmentations and ones that are caused by under-segmentations (Charoenphakdee et al., 2019).

The standard segmentation losses are the cross-entropy (Ronneberger et al., 2015) and the soft approximation of the Dice score (Milletari et al., 2016). The cross-entropy L_{CE} is a widely used standard loss function that is formulated via the Kullback–Leibler divergence and computes the dissimilarity between the predicted probabilistic distribution and its corresponding target binary distribution. Its mathematical expression given in the case of K classes, is:

$$\mathcal{L}_{CE} = -\frac{1}{|\Omega|} \sum_{p \in \Omega} \sum_{r=1}^K y_p^r \log(\hat{y}_p^r). \quad (2)$$

Since each pixel is handled independently from its neighbors, problems may arise due to class imbalance, as training can be dominated by the most prevalent class. For this reason, multiple works proposed variants of cross-entropy, with weights according to class or pixel imbalance (Jang et al., 2018; Baumgartner et al., 2017). One important cross-entropy variant is the weighted cross-entropy

Table 1: Mathematical notations

Ω	spatial image domain, $\Omega \subset \mathbb{R}^2$ or \mathbb{R}^3
\mathcal{N}_p	set of neighboring pixels of pixel p
K	number of classes
\mathbf{y}	true label map of dimension $ \Omega $
$\widehat{\mathbf{y}}$	predicted label (probability) map of dimension $ \Omega $
\widehat{y}_p	predicted label of pixel p
y_p	true label of pixel p
y_p^r	binary value indicating whether pixel p belongs to class r or not
\widehat{y}_p^r	probability value indicating whether pixel p belongs to class r or not

(Ronneberger et al., 2015), which tackles the cross-entropy sensitivity towards class distributions. Denoted by L_{WCE} , it is defined as:

$$\mathcal{L}_{WCE} = -\frac{1}{|\Omega|} \sum_{p \in \Omega} \sum_{r=1}^K w_k y_p^r \log(\widehat{y}_p^r), \quad (3)$$

where the weighting factor w_i assigns more weight to recessive classes, thus enforcing a higher penalty on their corresponding errors. Another variant is the focal loss (Lin et al., 2017), which extends upon cross-entropy in order to deal with the extreme foreground-background class imbalance in images.

Introduced by Milletari et al. (2016), the Dice loss L_{Dice} is a soft approximation of the well-known Dice metric, which penalizes the overlap mismatch between the predicted segmentation map and the corresponding target map. It can be computed in the general case with K classes (Sudre et al., 2017):

$$\mathcal{L}_{Dice} = 1 - 2 \frac{\sum_{r=1}^K \sum_{p \in \Omega} y_p^r \widehat{y}_p^r}{\sum_{r=1}^K \sum_{p \in \Omega} (y_p^r)^2 + \sum_{r=1}^K \sum_{p \in \Omega} (\widehat{y}_p^r)^2}. \quad (4)$$

The Dice loss is sensitive to segmentation errors when the segmented object is small. For this reason, some works have aimed at weighting the Dice loss (Yang et al., 2018; Sudre et al., 2017) in order to take into account the class imbalance, or extend-

ing upon it by accounting for background pixels, such as the Kappa coefficient inspired loss (Zhang et al., 2020).

3.3. Limitations of U-Net predictions

The cross-entropy and Dice losses, as well as their variants, and their combinations, are widely used in segmentation. However, these losses ignore high-level features or structures concerning the object of interest, such as their shape or topology. They also penalize all mistakes equally, regardless of their nature. In the same spirit, the U-Net original architecture (Ronneberger et al., 2015), does not leverage specific, anatomical or contextual constraints, nor does it exploit spatial relationships between organs. This is why many research works have explored the possibility of introducing prior information.

Incorporating prior information into segmentation encompasses a wealth of notions and covers various implementations; it can consist in introducing architecture modifications, adding constraints in the optimization problem, adding penalty terms in the loss function, or combining all these modifications. In the next section, Section 4, we review some approaches that take benefit from low-level prior, i.e. architectural constraints

or extracting features from the label maps, without making use of high-level knowledge in the loss function. Then, in Section 5, we review high-level based priors, that are integrated into the loss term, which is the core of our review.

4. Incorporating low-level prior in image segmentation

4.1. Feature extraction from label maps

One way to improve segmentation consistency is through conducting transformations on the ground-truth in order to obtain representations that are able to reveal geometric and contour attributes of the object of interest. Yang et al. (2019) exploit Laplacian filters in order to develop a boundary enhanced loss term that invokes the network to generate strong responses around boundary areas while producing a zero response in pixels that are at the periphery. Distance maps are also helpful, as a penalty-term (Caliva et al., 2019; Kervadec et al., 2019a) or added into the softmax, which is the normalization function based on exponential, at the end of the network (Petit et al., 2019). In the same context, Arif et al. (2018) introduce a shape aware loss function that constrains predictions to conform to permissible manifold in vertebrae segmentation. Different from the previous approaches, Mosinska et al. (2018) propose to leverage the topological information or shape descriptors present within the internal layers of VGG16 network, in order to close small gaps in thin structures (neuronal membranes) and alleviate topology mistakes. Kim and Ye (2020) introduce a loss term inspired by the Mumford-Shah functional in order to force each region to have similar pixel intensity, and to regularize the contour length. Lambert et al. (2021) propose a criterion of edge alignment, based on a weighted total

variation term. For further reading on the topic, we refer the reader to the survey (Lei et al., 2020), Section C-8.

4.2. Architectural constraints

Another approach to improving segmentation consistency with neural network predictions is through structural prior constraints (Trullo et al., 2017; Oda et al., 2018; Zotti et al., 2017; El Jurdi et al., 2020; Oktay et al., 2018). For example, Trullo et al. (2017) introduced a collaborative architecture in order to iteratively refine the posterior probability and provide information about neighboring organs. El Jurdi et al. (2020) integrated location and shape prior into the learning process, by introducing bounding filters at the level of the skip-connections in a U-Net model.

In Oktay et al. (2018), an autoencoder is integrated into a segmentation CNN, to act as regularisation model to constrain class label predictions. The integration is done via a so-called a loss, defined as the MSE between the encoded predicted map and the encoded true label map.

In many works, authors introduce both structural and loss constraints. For instance, Oda et al. (2018) extended upon cross-entropy to introduce a boundary enhanced loss function similar to that of Caliva et al. (2019) and Arif et al. (2018). However, instead of weighting by the errors through distance maps, Oda et al. (2018) added an extra decoder branch to the U-Net network in order to predict hard to segment boundaries. In the same manner, Zotti et al. (2017) integrated the center of mass and the contour prior into their loss function, which were obtained from an encoder-decoder structure trained end-to-end, along with the segmentation network.

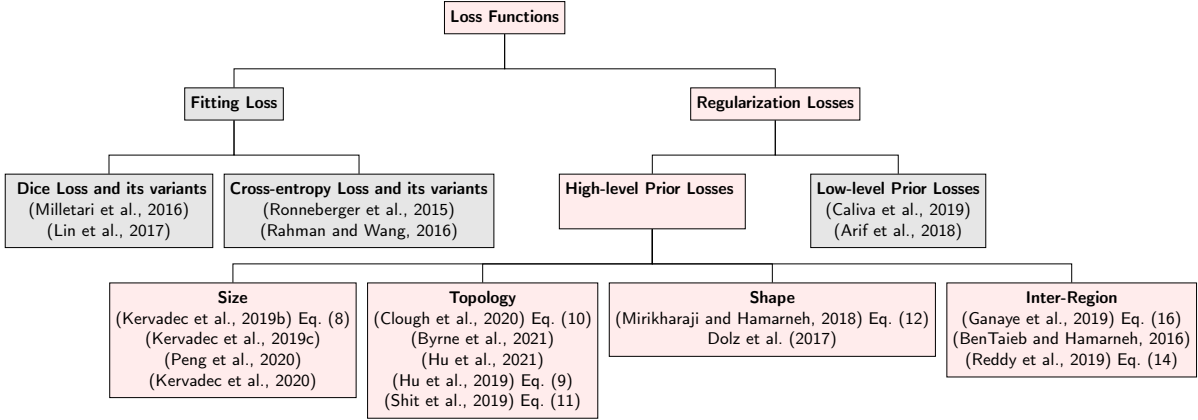


Fig. 2: Loss functions categorization tree. Loss functions can be data-fitting loss or regularization loss. Regularization or prior-based losses can be distinguished according to the nature of prior that they incorporate: low-level prior (Section 4.1) or high-level prior (Section 5). For each paper, the equation refers to the loss function to optimize.

5. High-level prior-based loss functions

In this section, we review the approaches that aim to integrate high-level prior for medical image segmentation at the level of the loss function, listed in Table 2. In order to contextualize the high-level prior-based loss functions, we present a categorization of loss functions, for medical image segmentation, in Fig. 2: we distinguish between data-fitting losses (as described in Section 3.2) that model the problem at hand and regularization losses. Prior-based losses are considered to act as regularization losses, and can be classified according to the nature of prior that they incorporate: low-level prior (as already reviewed in Section 4.1) or high-level prior (the current section), which is at the heart of this review. We have categorized the high-level priors (and subsequently the 13 reviewed papers in this section) according to the nature of the constraint: size constraint, topology, shape constraint and inter-regions constraints. We first start by the problem formulation.

5.1. Problem formulation

In the general problem formulation of finding the segmentation network parameters by ways of op-

timizing a loss function, one only has the data-fitting loss, as stated in Section 3.2. Training the network thus boils down to minimizing only the loss term of Eq. 1. The integration of a prior into the loss function can be seen as a constrained optimization problem (Marquez-Neila, 2017; Kervadec et al., 2019b; Peng et al., 2020; Pathak et al., 2015). Thus, in addition to the minimization of the data-fitting term formulated in Equation 1 some constraints C_j to be satisfied are added. The goal is to find the network parameters θ that minimize:

$$\mathcal{L} = \mathcal{L}_{Fit} \quad s.t. \quad C_j(\hat{y}) \leq 0 \quad j : 1, \dots, I \quad (5)$$

where I is the total number of constraints in the problem. The fitting loss function \mathcal{L}_{Fit} can be any of the common losses such as Dice or cross-entropy, as described in Section 3.2, whereas the constraints are mathematical representations of the prior, which relate to the number of connected components, the size of the organ, the topology etc. From here on, one can distinguish between a variety of optimization and training strategies. Moreover, optimization can be done either in a continuous domain where the formulated loss function is mainly derived from soft probabilities (Kervadec

et al., 2019b; Clough et al., 2020) or a discrete domain which directly targets hard-label assignments (Peng et al., 2020). One common method for solving constrained CNN training is through the method of Lagrange multipliers, also known as the penalty-based method (Kervadec et al., 2019b; Nocedal and Wright, 1999; Lillo et al., 1993). Such method models the constraint as a penalty term $\mathcal{L}_{penalty}$ in the loss function weighted by a parameter λ as follows:

$$\mathcal{L} = \mathcal{L}_{Fit} + \lambda \mathcal{L}_{penalty} \quad (6)$$

The additional loss term must be differentiable, convex and produces a value proportional to the degree of constraint violation. The weighting factor λ can be either predefined throughout training (static training) or fine-tuned along training (dynamic training). We now move on to our categorization by prior nature.

5.2. Size constraint

The size of an organ is a feature that has a known range of variability. In (Kervadec et al., 2019b; Pathak et al., 2015), the idea is to integrate this information into the segmentation process, and to constraint the predicted organ area to be in this known size range. The problem is to estimate the organ size from a soft probability map. Pathak et al. (2015) are the first to propose a size constraint optimized through biconvex Lagrangian dual methods. They formulate the ground truth as a latent distribution. Then, they alternate between bringing the probability distribution to be as close as possible to the ground truth distribution, given fixed model parameters on one hand, and optimize model parameters via gradient descent given known latent distribution on the other hand. They experiment their method on natural

images. In (Kervadec et al., 2019b), the authors address the problem of medical image segmentation and argue that the dual optimization problem is computationally intractable when applied to neural networks. As a result, it is more convenient to integrate the size prior directly at the level of the loss function, under the form of a differentiable penalty term and optimize model parameters accordingly. Thus, Kervadec et al. (2019b) do not threshold the predicted label map, but rather estimate the area with the summation of the probabilities over the whole image domain:

$$A(\hat{\mathbf{y}}) = \sum_{p \in \Omega} \hat{y}_p \quad (7)$$

Then prior knowledge is used to impose a lower bound a and a higher bound b on the organ size. A penalty loss function that integrates these bounds is proposed as follows:

$$\mathcal{L}_{size} = \begin{cases} (A(\hat{\mathbf{y}}) - a)^2 & \text{if } A(\hat{\mathbf{y}}) \leq a, \\ (A(\hat{\mathbf{y}}) - b)^2 & \text{if } A(\hat{\mathbf{y}}) \geq b, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

The proposed loss \mathcal{L}_{size} is implemented in a weakly supervised setting for cardiac segmentation. In this case prior knowledge is used in order to overcome the problem of partial label absence.

Methods of Kervadec et al. (2019b) and Pathak et al. (2015) explore the optimization of neural networks under constraints, given a continuous domain. However, conducting optimization from a continuous perspective may not guarantee that the discrete constraints are satisfied. The issue of the discrete nature of anatomical constraints have led to discrete optimization of neural networks in (Peng et al., 2020), where the authors investigated the alternating direction method of multipliers algorithm (ADMM), in order to incorporate boundary smoothness and size constraints. The ADMM

method is a variant of the augmented Lagrangian scheme, which allows the decoupling of the continuous optimization of neural network parameters by gradient descent, from the discrete optimization of size constraints.

Both (Kervadec et al., 2020) and (Kervadec et al., 2019c) aim to exploit bounding box prior as means of extracting size and tightness constraints. With bounding box annotations, they are able to constrain the organ size and location, inside the bounding box. However, Kervadec et al. (2019c) integrate bounding box prior to conduct fully supervised medical image segmentation, their work in Kervadec et al. (2020) demonstrate the ability of exploiting this prior in order to permit segmentation under weak supervision with performance close to full supervision. To integrate bounding box prior, authors argue that the segmented region should be sufficiently close to the sides of the bounding boxes. Thus, each horizontal or vertical line that is parallel to the sides of the bounding box is to intersect the target segment at least once. As a result, the sum of pixels along the line should be greater than the sum of pixels belonging to the label (Figure 3). To integrate the multiple constraints, authors adopted a Lagrangian optimization method with log-barrier extensions (Kervadec et al., 2019c). The method involves introducing a standard log-barrier function (Boyd and Vandenberghe, 2004) that evades the need for dual optimizations and their issues. Optimization under the log-barrier extensions have been introduced previously (Chouzenoux et al., 2019); however, it is still a novel research direction in medical image segmentation.




Fig. 3: Tightness prior. (i) Camel image. Given that any segment (pink stripe) is made up of w lines and that each line crosses the camel at least once, any segment (horizontal or vertical) of width w crosses the camel in at least w pixels, as illustrated in (ii), (iii), (iv) and (v). Figures from (Kervadec et al., 2019c). *[permission pending]*

5.3. Topology constraints

Topology is concerned with the properties of spatial objects by abstracting their connectivity, while ignoring their detailed form (Ségonne and Fischl, 2015). In this section, we present works which are based on explicit topology modeling, through the use of Betti numbers, a measure of topological structures (Hu et al., 2019; Clough et al., 2020; Hu et al., 2021), and skeletonization (Shit et al., 2019).

Betti numbers are topological invariants. They are determined for a dimension k : Betti number β_k is the number of k -dimensional features on a topological surface. For example, β_0 represents the number of connected components and β_1 the number of holes (in 2D images, only these two Betti numbers are useful). Betti numbers are discrete, obtained on thresholded binary predictions, and as such cannot be used directly for CNN training. Instead, Hu et al. (2019) and Clough et al. (2020) have exploited the notion of persistent homology to integrate this prior onto the loss term. Persistent homology is a transformation that encodes the evolution of topological structures of nested spaces. In our case, the persistent homology consists in thresholding the prediction map with a linear sequence of threshold values and summarize these information in concise format. In particular,

depending on the threshold, some structures, that may be 0-dimensional or 1-dimensional, are born or die in the image.

Hu et al. (2019) exploit homology via persistent diagrams (Dgm). Each persistent diagram contains a finite number of 2D dots ($p = (b, d)$) corresponding to a topological structure, that is born at threshold b and killed at a threshold d . Then, the goal is to find the best one-to-one correspondence noted as γ , between the sets of dots corresponding to the ground truth persistent diagram ($Dgm(g)$) and to the predicted persistent diagram ($Dgm(f)$), by minimizing the squared distance between them as follows:

$$\mathcal{L}_{Topo1} = \sum_k \sum_{p_i^k \in Dgm(f)} \left(b(p_i^k) - b(\gamma(p_i^k)) \right)^2 + \left(d(p_i^k) - d(\gamma(p_i^k)) \right)^2 \quad (9)$$

Following an idea similar to persistence diagrams, topological structures can be represented through persistence barcodes, as in (Clough et al., 2020) and (Byrne et al., 2021) (Figure 4). Here, the authors make the hypothesis that Betti numbers for the segmented object, denoted as β_k^* , are known. The idea is to consider that longest bars have the most meaningful topological features in the data, the length being computed as the difference between birth and death. Ideally, longest bars should have a length reaching 1 indicating the persistence of the topological feature throughout the entire threshold space. Hence, their loss aims to maximize the longest bars (first term in Eq. 10), and to get rid of transient components, corresponding to shortest bars (second term in Eq. 10) as




Fig. 4: Persistent barcode diagram. The probability map on the left contains three visible regions of high intensity, which correspond to the three persistent 0-dimensional features shown as red bars in the right diagram, with threshold values corresponding to birth and death of a topological feature on the x-axis. The map also contains a loop of high intensity, corresponding to the one persistent 1-dimensional feature, shown here as a green bar on the barcode diagram. Figures from (Clough et al., 2020). [permission pending]

follows:

$$\mathcal{L}_{Topo2} = \sum_k \left(\sum_{l=1}^{\beta_k^*} \left(1 - |b(p_l^k) - d(p_l^k)| \right)^2 \right) + \sum_{l=\beta_k^*+1}^{\infty} |b(p_l^k) - d(p_l^k)|^2 \quad (10)$$

Whereas (Clough et al., 2020) investigates this loss within the binary segmentation setting, Byrne et al. (2021) extend this work for multi-class segmentation of cardiac MR images.

As an alternative to persistent homology, authors of (Hu et al., 2021) exploit notions of discrete Morse theories (Milnor, 2016) in order to propose a novel approach that identifies critical topological structures and preserves desired Betti numbers. Morse theories base on the assumptions that network outputs or probability maps are rather terrain functions characterized with ridges and valleys, representing critical topological structures. Their proposed loss identifies these structures and enforces higher penalties along them.

Another important concept that reveals topological properties of objects is the skeleton. Skeletonization is the process of obtaining compact rep-

representations of images and objects while still preserving topological properties. The idea of (Shit et al., 2019) is to constrain the skeleton of the predicted map to match the skeleton of the ground truth map. This prior is used in the segmentation of vessels and neurons in both 2D and 3D. While the skeleton of a shape can be obtained with various approaches, the authors underline that using a discrete operation such as the Euclidean distance transform does not allow to obtain a differentiable approximation. Thus, they propose to use morphological thinning, a sequence of dilations and erosions. To handle the soft probability map values, erosions and dilations are replaced by their grey-scale equivalent (min and max filters), giving rise to ‘soft-skeletonization’. In the CNN, iterative min- and max-pooling is applied as a proxy for morphological erosion and dilation. Once the skeleton is computed, an appropriate prior loss term can be designed.

Let s and \widehat{s} be the ground truth and the predicted skeleton respectively, of size $|\Omega|$. The sensitivity (or recall) between the predicted segmentation and ground truth skeleton is introduced as $T_{sens}(s, \widehat{y}) = |s \cap \widehat{y}|/|s|$. Likewise, the precision between the ground truth mask y and the predicted skeleton \widehat{s} is defined as: $T_{prec}(\widehat{s}, y) = |\widehat{s} \cap y|/|\widehat{s}|$. The cDice is defined as the F1-score between precision T_{prec} and sensitivity T_{sens} as follows:

$$\mathcal{L}_{cDice} = 2 \frac{T_{prec}(\widehat{s}, y) T_{sens}(s, \widehat{y})}{T_{prec}(\widehat{s}, y) + T_{sens}(s, \widehat{y})}, \quad (11)$$

Interestingly, for all three approaches (Hu et al., 2019; Clough et al., 2020; Shit et al., 2019) results show the topological prior do not fully outperform the no-prior approaches, as measured with regional metrics such as Dice, but they increase specific topological metrics, such as the cDice accuracy in Shit et al. (2019).

5.4. Shape constraint

There are numerous shape descriptors, such as geometric features, moments, shape transforms, or based on scale-space theory or polygonal approximation. These descriptors are usually computed on shapes which are represented by a set of point coordinates or with a binary map. The difficulty of integrating such descriptors into the loss terms stems from the fact that they must be computed on real-valued probability maps, i.e. the network output, instead of binary ones. Inspired by (Veksler, 2008), Mirikharaji and Hamarneh (2018) propose a loss that forces the segmented region to have a star shape, for the task of segmenting dermoscopic skin lesion. An organ is said to have a star shape if, for any point p inside the object, all the pixels q lying on the straight-line segment connecting p to the object center c , are inside the object (Figure 5). Let l_{pc} be the line segment connecting pixel p to the object center c , and q be any pixel incident on line l_{pc} . The proposed loss is expressed as:

$$\mathcal{L}_{star} = \sum_{p \in \Omega} \sum_{q \in l_{pc}} \delta_{p,q} \cdot |y_p - \widehat{y}_p| \cdot |\widehat{y}_p - \widehat{y}_q| \quad (12)$$

where $\delta_{p,q}$ is the Kronecker symbol defined as:

$$\delta_{p,q} = \begin{cases} 1 & \text{if } y_p = y_q; \\ 0 & \text{otherwise.} \end{cases}$$

Star shape prior is a way to promote convexity for organ shapes. The star-shape loss registers significance improvement on segmentation performance, given a variety of networks such as U-Net and ResNet-DUC (Wang et al., 2018).

In (Dolz et al., 2017), the idea is to enforce compactness. This property is represented by the ratio of the square of the perimeter to the shape area, that is required to be as small as possible. Here the boundary length requires to estimate a discrete form of this ratio functional, not with the usual




Fig. 5: Star shape prior. (a) Star shape object O w.r.t. the supplied object center c (red dot). (b) Example of star shape constraint violation. (c) Example when the star-shape prior loss is required. Figures from (Mirikharaji and Hamarneh, 2018) [permission pending].

continuous variable \mathbf{y} , but with a discrete binary vector $\widehat{\mathbf{z}} \in \{0, 1\}^{|\Omega|}$:

$$\mathcal{L}_{compact} = \frac{P(\widehat{\mathbf{z}})^2}{A(\widehat{\mathbf{z}})} \quad (13)$$

where $A(\widehat{\mathbf{z}})$ and $P(\widehat{\mathbf{z}})$ represent the predicted organ area and boundary respectively. The area is computed according to Equation 7 and boundary is proportional to the number of neighboring pixels with different labels, and thus computed as: $P(\widehat{\mathbf{z}}) \propto \sum_{p \in \Omega} \sum_{q \in \mathcal{N}_p} |\widehat{z}_p - \widehat{z}_q|$.

The proposed loss is dimensionless, unbiased and position independent. However, due to the discrete nature of the prior involved, optimization of this compactness-based loss comes with great challenges. For this reason, Dolz et al. (2017) alternate between optimizing the network parameters with SGD and optimizing the discretely-constrained segmentation labels, via ADMM.

5.5. Inter-regions constraint losses

In the case of multi-label segmentation, specific interactions between regions, known a priori, can be authorized or forbidden: adjacency relations between organs are handled in (Ganaye et al., 2019), while (BenTaieb and Hamarneh, 2016) propose solutions to enforce regions exclusion and inclusion.

Focusing on gland histology images, BenTaieb and Hamarneh (2016) have identified that the cell and the object to be segmented, is made of two




Fig. 6: Inter region constraint prior. Three anatomical objects, (WT, violet), (TC, green) and (ET, pink), have $2^3 = 8$ possible combinations of existence. Given the correct anatomical topology specified in (i), the validity map V can then be derived for each of the 8 cases as shown in the last row of the table. Figure from (Reddy et al., 2019) [permission pending]

nested regions: one region (the lumen) is contained in another one (the epithelium). The authors integrated this spatial property by introducing a binary validity indicator map. A validity indicator map V returns 1 if a given label y_p corresponds to a topologically-valid assignment, and zero otherwise. In this way, the network not only penalizes incorrect label assignment per pixel, but also penalize incorrect label hierarchy. Their loss term is based on a V -weighted cross-entropy, defined as follows:

$$\mathcal{L}_{inter} = \sum_{p \in \Omega} \sum_{r=1}^K -y_p^r \log(\widehat{y}_p^r) \cdot V \quad (14)$$

where

$$V = \begin{cases} 1 & \text{if topology is in accordance;} \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, their method exploits a boundary smoothness term that takes into consideration the difference between probabilities of pixels corresponding to same labels, to the proposed loss. The two constraints are combined via penalty-based optimization. This approach has been ap-

plied to brain tumor segmentation in 3D MR images Figure 6 (Reddy et al., 2019). We note that the inclusion-exclusion loss (and the boundary smoothness loss) are not optimized alone, but are added to a fitting loss, which may be cross-entropy (BenTaieb and Hamarneh, 2016) or Dice (Reddy et al., 2019) and the total loss function, which is the sum of all three losses, is optimized through regular stochastic gradient descent.

Both publications mentioned above validate the importance of the proposed prior loss in different tasks and modalities. However, the use of the prior loss does not compensate for the need of fitting losses such as Dice or cross-entropy. The presence of such losses is generally necessary for the convergence of the segmentation framework. Moreover, the method still depends on penalty-based optimization to balance out the two constraints, which does not accommodate the interplay and relations between the different constraints (Kervadec et al., 2020).

In (Ganaye et al., 2019), the authors propose a loss that takes into consideration the relationships between neighboring anatomical objects. From the ground truth label maps, one can define an adjacency matrix \mathbf{A} of general binary term A_{ij} between regions, that represents whether two regions can be adjacent or not. Then the set of forbidden label connections can be defined as follows: $F = \{(i, j) | A_{ij} = 0\}$.

However, an adjacency measure computed from the outputs of the CNN, which are probability maps and not label maps, is more difficult to define. When two regions i and j should not be connected, i.e. $(i, j) \in F$, then the probability for a pixel and its neighbors to belong to i and j must be close to zero. Let $\hat{y}_p(i)$ (resp. $\hat{y}_q(j)$) be the probability of pixel p (resp. q) to belong to class

i (resp. j), Ganaye et al. (2019) propose to model the constraint by the product $\hat{y}_p(i)\hat{y}_q(j)$. The adjacency measure is then:

$$a_{ij} = \sum_{p \in \Omega} \sum_{q \in \mathcal{N}_p} \hat{y}_p(i)\hat{y}_q(j) \quad (15)$$

Following this, the proposed loss consists of forcing all the forbidden adjacency relationships, with respect to the relations defined in the set of impossible transitions, F :

$$\mathcal{L}_{adj} = \sum_{(i,j) \in F} a_{ij} \quad (16)$$

The authors then solve the constrained optimization problem via the penalty-based method mentioned in Section 5.1. In the experiments, datasets with number of regions ranging from 20 to 135 are investigated. Interestingly, as the shape and size constraints, the proposed adjacency loss does not require the ground truth segmentation – it just requires the definition F , thus the segmentation method can be evaluated in a semi-supervised framework. Model performance in terms of the Dice metric shows no significant improvement, when compared to the baseline method (i.e. the same segmentation method without the proposed loss term). However, incorporating adjacency constraints registers considerable improvement with regards to the distance metrics (Hausdorff and Mean Distance Metric). These results are consistent over all datasets (in 2D and 3D) and settings, i.e. with full or semi supervision. Depending on the applications, one limitation of the approach may be the assumption that all patients have the same inter-organ connectivity.

6. Discussion

In addition to the common challenges in deep networks training, such as overfitting, scarcity

Table 2: List of reviewed papers with respect to the category of the prior types: topology, size, shape, inter-regions priors. Evaluation metrics: std refers to the standard evaluation metrics such as pixel-wise accuracy (pA), dice similarity coefficients (DSC), Hausdorff distance (HD); specif. means that the paper uses prior-specific metrics. SGD: stochastic gradient descent. ADMM: Alternating direction method of multipliers.

Paper	Architecture	Optimized features	Dim.	Class nb	Training strategy	Evaluation metrics		Supervision	Constraint nb	Optimization domain	Optimization method
						std	specif				
Size											
(Peng et al., 2020)	U-Net, E-Net	Binary segmentation proposals and Lagrangian multipliers	2D 3D	1	N/A	DSC	-	weak	2	Discrete	ADMM
(Kervadee et al., 2020)	U-net E-Net	Combined tightness and size prior constraints	2D	1	static	DSC	-	weak	3	Continuous	Lagrangian multiplier with Log-barrier extensions
(Kervadee et al., 2019b)	U-net E-Net	Lower and upper bounds on organ size	2D 3D	1	dynamic	DSC		weak	1-2	Continuous	Penalty based approach via SGD
Topology											
(Byrne et al., 2021)	U-Net	Betti number	2D 3D	3	static	DSC	yes	full	1	Continuous	Penalty based approach via SGD
(Hu et al., 2021)	U-Net, DIVE, Mosin	Morse skeleton	2D 3D	1	static with refinement	pA	yes	full	1	Continuous	Penalty based approach via SGD
(Clough et al., 2020)	U-Net	Betti number	2D 3D	1	static	DSC	yes	semi	1	Continuous	Penalty based approach via SGD
(Hu et al., 2019)	U-Net, DIVE, Mosin	Birth and death of topological components (persistent homology)	2D 3D	1	static with refinement	pA	yes	full	1	Continuous	Penalty based approach via SGD
(Shir et al., 2019)	U-Net FCN	Skeletons	2D 3D	1	static	DSC, pA	yes	full	1	Continuous	Penalty based approach via SGD
Shape											
(Mirikharaji and Hamarneh, 2018)	U-Net ResNet-DUC	Star-shapeness	2D	1	static	DSC, pA, Jacard, Spec, Sens	-	full	1	Continuous	Penalty based approach via SGD
(Dolz et al., 2017)	FCNN	Ratio of boundary length to area	3D	1	N/A	DSC	-	full	1	Discrete	ADMM
Inter-regions											
(Ganaye et al., 2019)	SD-Net	Prohibited and permissible connectivity between organs	2D 2.5D 3D	20-135	dynamic	DSC, HD	yes	full, semi	1	Continuous	Penalty based approach via SGD
(Reddy et al., 2019)	FCN	Boundary smoothness and containment/exclusion properties	2D	3	static	DSC, HD	-	full	2	Continuous	Penalty based approach via SGD
(BenTaieb and Hamarneh, 2016)	Alexnet-FCN FCN-8 U-Net	Boundary smoothness and containment/exclusion properties	2D	2	static	DSC, pA	-	full	2	Continuous	Penalty based approach via SGD

of annotated data, class imbalance, and gradient divergence which are extensively discussed in (Hesamian et al., 2019; Litjens et al., 2017; Havaei

et al., 2016) for example, there are particular challenges when dealing with a prior-based term and its incorporation into the loss function. In this sec-

tion, we summarize and highlight a number of key aspects of embedding a prior based loss function into a segmentation network.

6.1. The nature of the prior

The high-level prior, as we defined it, expresses high-level features regarding the object of interest, with interpretable insight with respect to the organ geometry or anatomy. This prior can stem from medical knowledge (e.g. organ size range, organ connectivity) and as such, can be used in a weakly and semi-supervised learning context to improve performance. Sometimes, the prior has to rely on features extracted from the ground truth label maps, see for example the Betti numbers or the skeletonization process. In this case, its usage is restricted to full supervision.

6.2. The challenge of soft probability maps

One major challenge is to compute features from soft probability maps. A binary map expressing the object shape is much easier to characterize with usual shape features (e.g. circularity, compactness, isoperimetric ratio, skeleton). However, thresholding the probability map to make it binary can render the loss function non-differentiable. Some parameters can be estimated from probability maps, e.g. the predicted organ size in (Kervadec et al., 2019b). Other features require to resort to a discrete optimization scheme, such as the predicted organ boundary length (Dolz et al., 2017) (see Section 6.3). Another way of dealing with the soft probability maps is to impose a series of thresholds, to monitor topological changes (Clough et al., 2020; Hu2). However, their method is not generic and cannot be applied to all prior properties. These issues become more complicated as the type of prior handled becomes more complex, and loss

functions often end up being non-convex or hard to optimize.

6.3. Continuous vs discrete optimization strategies

A common and simple way to integrate constraints within a continuous domain is through the penalty based method. Penalty based method involves formalizing the constraint as an addition penalty loss term in the main loss function weighted by a parameter λ which may be statically or dynamically defined through training. The novel loss term which includes the main per-pixel fitting loss and the novel penalty loss are then optimized using regular stochastic gradient descent. Many works adopt the penalty-based method while dealing with anatomical prior for its ease of formulation and use (Shit et al., 2019; Kervadec et al., 2019b; Hu et al., 2019; Clough et al., 2020; BenTaieb and Hamarneh, 2016; Mirikharaji and Hamarneh, 2018). Despite this simplicity, penalty based approaches may not guarantee constraint satisfaction. Moreover, they require careful fine-tuning of their weighting terms, which may not be convenient in the case where multiple constraints occur and where one constraint may over-shadow the others.

One way to deal with multiple constraint optimization, demonstrated in (Kervadec et al., 2020, 2019c), is through Lagrangian optimization with log-barrier extensions. The method involves introducing a standard log-barrier function that avoids the need for dual optimizations and their issues (Boyd and Vandenberghe, 2004). The method then integrates these constraints into the log-barrier function and solves the optimization process in an unconstrained manner via stochastic gradient descent. Unlike the penalty based approach, the log-barrier approach does not yield null gradi-

ents or cause oscillations between competing constraints. It is rather characterized by stable gradients that insures training stability. Optimization in a discrete domain can also be insightful, given the discrete nature of anatomical properties of organs. Where continuous optimization makes assumptions on the soft probabilities in order to estimate differentiable functions, discrete optimization involves extracting the features in their discrete form from model predictions, and optimizing them relative to the ground truth. One way to perform discrete optimization is through the ADMM method (Peng et al., 2020). The ADMM algorithm generally aims at separating the optimization of the network parameters under SGD from the optimization of discrete constraint segmentation labels. Discretely optimizing networks generally benefits from the ability to solve sub-problems, either continuous or discrete, separately, and insures global optimum, which can improve solutions within a single gradient step and at higher convergence speed.

6.4. Relationship between organs and loss design


The studied papers address various segmentation problems, as listed in Table 3, that shows the targeted organs or objects to be segmented, and the datasets used in each paper. Size and shape constraint mostly concern single instance organs, that have convex or a blob shape, such as the prostate, the cardiac ventricles, the aorta, the esophagus, or skin lesions (see for example the Promise, ACDC and ISIC datasets). Topology priors are mostly used for thin, curvilinear objects such as neuron membranes, vessels or, at a higher scale, the myocardium of the left ventricle which has a ring shape (Fig. 7a). Inter-regions priors can help in problems of multiclass segmentation.

The adjacency constraints have shown efficiency in full body segmentation (120 regions in Anatomy3 dataset), and for multiple regions segmentation in brain segmentation (20 and 35 regions in IBSRv2 and MICCAI12 dataset). Exclusion and inclusion priors are helpful whenever there is a hierarchy in region membership. Their use is illustrated in applications, both at the microscopic and macroscopic levels. In microscopy images, cells can be composed of several layers: for example, gland cells (found in the GlaS dataset) are made of the inside region, called the glandular lumen, and the outer region, identified as the epithelial boundary. In the BraTs dataset, brain tumors in MRI are made of enhancing tumor (the deepest level), that is surrounded by the tumor core, itself surrounded by a region identified as whole tumor (Fig. 7b).

6.5. Future trends

The design of the prior loss is facing requirements, concerning the differentiability of the loss terms, or at least the computational complexity that must remain reasonable. Lately, as described in this survey, some advances have been made, that explored optimization techniques such as ADMM or the barrier functions (see column 'Opt method' in Table 3) and that have allowed to incorporate loss terms which are not directly optimizable by SGD. We believe future progress will originate from using advanced constraint optimization techniques, stemming from both equality or inequality constraints optimization.

Variational and optimization based approaches for image segmentation from the pre-deep learning era can also provide key pointers on how to model prior information regarding object shape and appearance. Researchers can rely on decades of works on the topic to find inspiration to design losses for



(a) Topological priors in retinal vessels (left), neuronal membrane (middle), myocardium of the left ventricle (right). Top: images, bottom: ground truth.

(b) Inter-region prior: adjacency constraint in brain and full body regions (top), inclusion-exclusion relationships (bottom)

Fig. 7: Example of targeted segmentation objects, that can benefit from topological priors (a), inter-region priors (b). Figures from (Hu2; Clough et al., 2020; Reddy et al., 2019; Ganaye, 2019). *[permission pending]*

their segmentation problem. In this regard, the majority of the priors are only presented in the case of binary segmentation, except for the inter-regions priors obviously. There is still work to be done to address multiple organ segmentation, with priors that are dedicated to organ shape, size or topology.

Prior based losses present promising behavior with regards to their ability to compensate for the need for full annotations, and are thus very useful in weakly and semi-supervised segmentation frameworks. This is already the case with the papers reviewed in this survey, see Table 2. More generally, embedding prior information plays a role in other applications in medical imaging. For example, when using data augmentation or self-supervised approaches, we (explicitly or implicitly) make assumptions about what samples in our training set should look like, what structure they should have, and so forth. To truly make progress, we believe it is important that such assumptions are made explicit, and to discuss these

methods from the perspective of more traditional methods which more heavily relied on such priors.

At last, we have noticed that sometimes, prior losses might not improve the segmentation results, when measured by generic metrics such as the Dice coefficient, see for example (Shit et al., 2019; Hu et al., 2019; Clough et al., 2020) (column 'Evaluation metric' in Table 2). In this case, the authors propose specific metrics, which are in close relation with the proposed loss, and thus are not generic. However, the prior losses in general generate more plausible regions, that are hardly measurable by the generic metrics. There is a wide variety of metrics for medical image segmentation as highlighted in a survey (Taha and Hanbury, 2015), that reviews overlap based, volume based, pair-counting based, information theoretic based, probabilistic based, and spatial distance based metrics - but no metric related to anatomical plausibility is reported. To our knowledge there is no metric that conveys the notion of anatomical plausibility; however we believe this is an important, open topic, to

Table 3: Targeted segmentation objects and datasets used in the presented papers. UKb: UK Biobank. MIC12: MICCAI 2012 workshop on multi-atlas labeling. Anat3: Anatomy3. p: proprietary dataset.

		Size			Topology			Shape	Inter-regions					
		(Peng et al., 2020)	(Kervadec et al., 2019c)	(Kervadec et al., 2019b)	(Hu et al., 2021)	(Byrne et al., 2021)	(Clough et al., 2020)	(Hu et al., 2019)	(Shit et al., 2019)	(Mirikharaji and Hamameh, 2018)	(Dolz et al., 2017)	(BenTateb and Hamameh, 2016)	(Ganaye et al., 2019)	(Reddy et al., 2019)
CT	Full body Aorta/eso										p	Anat3		
MRI	Brain Prostate Cardiac	Promise	Promise	Promise	ACDC	UKb					p	MIC12 IBSRv2	BraTs	
Photo	Skin lesion									ISIC				
Microscopy					ISBI12		ISBI12						GlaS	
					ISBI13		ISBI13							
					DRIVE		DRIVE	DRIVE						

address in the future, to show the benefits of these prior-based losses.

7. Conclusions

In this paper, we presented a survey of the current state-of-the-art methods regarding high level prior based losses, for medical image segmentation. We have proposed a categorization where we grouped these methods according to the type of prior that they incorporate: size, topology, shape, and inter-region relations. We have further characterized these methods based on the type of features they optimized, the architecture they use, optimization strategy and the anatomical object that they target. We have discussed the challenges involved with the design: the fact that the prior has to be extracted from soft probability map, the optimization constraints, and the design of the loss with respect to the object of interest. Finally, we have presented some future trends that could over-

come the limitations of current research works and hope they can be useful to foster research in this promising field.

Acknowledgments

The authors would like to acknowledge the National Council for Scientific Research of Lebanon (CNRS-L) and the *Agence Française de la Francophonie* (AUF) for granting a doctoral fellowship to Rosana El Jurdi, as well as the ANR (project API, grant ANR-18-CE23-0014). This work is part of the DAISI project, co-financed by the European Union with the European Regional Development Fund (ERDF) and by the Normandy Region. This research was conducted as part of a collaboration with the Eindhoven university of Technology under the support of the PHC Van Gogh project WeSmile.

References

- Arif, A., Rahman, S.M.M., Knapp, K., Slabaugh, G., 2018. Shape-aware deep convolutional neural network for vertebrae segmentation, in: *Computational Methods and Clinical Applications in Musculoskeletal Imaging*, pp. 12–24.
- Ayed, I.B., Li, S., Islam, A., Garvin, G., Chhem, R., 2008. Area prior constrained level set evolution for medical image segmentation, in: *Medical Imaging 2008: Image Processing*, International Society for Optics and Photonics. p. 691402.
- Baumgartner, C.F., Koch, L.M., Pollefeys, M., Konukoglu, E., 2017. An exploration of 2d and 3d deep learning techniques for cardiac mr image segmentation, in: *STACOM@MICCAI*.
- BenTaieb, A., Hamarneh, G., 2016. Topology Aware Fully Convolutional Networks for Histology Gland Segmentation, in: *MICCAI*. volume 9901, pp. 460–468. URL: http://link.springer.com/10.1007/978-3-319-46723-8_53, doi:10.1007/978-3-319-46723-8_53.
- Bernard, O., Lalonde, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P., Cetin, I., Lekadir, K., Camara, O., Gonzalez Ballester, M.A., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V.A., Krishnamurthi, G., Rohé, M., Pennec, X., Sermesant, M., Isensee, F., Jäger, P., Maier-Hein, K.H., Full, P.M., Wolf, I., Engelhardt, S., Baumgartner, C.F., Koch, L.M., Wolterink, J.M., Išgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P., 2018. Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging* 37, 2514–2525.
- Boyd, S., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press. doi:10.1017/CB09780511804441.
- Boykov, Y., Funka-Lea, G., 2006. Graph cuts and efficient nd image segmentation. *International journal of computer vision* 70, 109–131.
- Byrne, N., Clough, J.R., Montana, G., King, A.P., 2021. A persistent homology-based topological loss function for multi-class cnn segmentation of cardiac mri, in: Puyol Anton, E., Pop, M., Sermesant, M., Campello, V., Lalonde, A., Lekadir, K., Suinesiaputra, A., Camara, O., Young, A. (Eds.), *Statistical Atlases and Computational Models of the Heart. M&Ms and EMIDEC Challenges*, Springer International Publishing, Cham. pp. 3–13.
- Caliva, F., Iriondo, C., Martinez, A.M., Majumdar, S., Pedoia, V., 2019. Distance map loss penalty term for semantic segmentation, in: *International Conference on Medical Imaging with Deep Learning – Extended Abstract Track*, London, UK. URL: <https://openreview.net/forum?id=B1eIcvS45V>.
- Chahal, P.K., Pandey, S., Goel, S., 2020. A survey on brain tumor detection techniques for mr images. *Multimedia Tools and Applications* 79, 21771–21814. URL: <https://doi.org/10.1007/s11042-020-08898-3>, doi:10.1007/s11042-020-08898-3.
- Charoenphakdee, N., Lee, J., Sugiyama, M., 2019. On Symmetric Losses for Learning from Corrupted Labels, in: *ICML*, PMLR, PMLR. pp. 961–970.
- Chen, C., Qin, C., Qiu, H., Tarroni, G., Duan, J., Bai, W., Rueckert, D., 2020. Deep learning for cardiac image segmentation: A review. *Frontiers in Cardiovascular Medicine* 7, 25. URL: <https://www.frontiersin.org/article/10.3389/fcvm.2020.00025>, doi:10.3389/fcvm.2020.00025.
- Chouzenoux, E., Corbineau, M.C., Pesquet, J.C., 2019. A Proximal Interior Point Algorithm with Applications to Image Processing. *Journal of Mathematical Imaging and Vision* URL: <https://hal.archives-ouvertes.fr/hal-02120005>.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S., Brox, T., Ronneberger, O., 2016. 3d u-net: Learning dense volumetric segmentation from sparse annotation, in: *MICCAI*, pp. 424–432. URL: <http://lmb.informatik.uni-freiburg.de/Publications/2016/CABR16>.
- Clough, J., Byrne, N., Oksuz, I., Zimmer, V.A., Schnabel, J.A., King, A., 2020. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1doi:10.1109/TPAMI.2020.3013679.
- Cootes, T., Cooper, D., Taylor, C., Graham, J., 1995. Active shape models - their training and application. *Computer Vision and Image Understanding* 61, 38–59.
- Cremers, D., Rousson, M., Deriche, R., 2007. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *International Journal of Computer Vision* 72, 195–215.
- Debelee, T.G., Schwenker, F., Ibenthal, A., Yohannes, D., 2020. Survey of deep learning in breast cancer image analysis. *Evolving Systems* 11, 143–163. URL: <https://doi.org/10.1007/s12530-019-09297-2>, doi:10.1007/s12530-019-09297-2.

- s12530-019-09297-2.
- Dolz, J., Ben Ayed, I., Desrosiers, C., 2017. Unbiased shape compactness for segmentation, in: MICCAI, pp. 755–763.
- El Jurdi, R., Petitjean, C., Honeine, P., Abdallah, F., 2019. Organ Segmentation in CT Images With Weak Annotations: A Preliminary Study, in: 27th GRETSI Symposium on Signal and Image Processing, Lille, France.
- El Jurdi, R., Petitjean, C., Honeine, P., Abdallah, F., 2020. BB-UNet: U-Net with Bounding Box Prior. *IEEE Journal of Selected Topics in Signal Processing* 14, 1189 – 1198.
- Foulonneau, A., Charbonnier, P., Heitz, F., 2009. Multi-reference shape priors for active contours. *International journal of computer vision* 81, 68.
- Ganaye, P.A., 2019. A priori et apprentissage profond pour la segmentation en imagerie cérébrale. Ph.D. thesis. Université de Lyon.
- Ganaye, P.A., Sdika, M., Triggs, B., Benoit-Cattin, H., 2019. Removing segmentation inconsistencies with semi-supervised non-adjacency constraint. *Med Image Anal* 58, 101551.
- Grady, L., 2012. Targeted image segmentation using graph methods, in: Lézoray, O., Grady, L. (Eds.), *Image Processing and Analysis with Graphs*. CRC Press, pp. 111–135.
- Haque, I., Neubert, J., 2020. Deep learning approaches to biomedical image segmentation. *Informatics in Medicine Unlocked* 18, 100297. URL: <http://www.sciencedirect.com/science/article/pii/S235291481930214X>, doi:<https://doi.org/10.1016/j.imu.2020.100297>.
- Havaei, M., Guizard, N., Larochelle, H., Jodoin, P.M., 2016. Deep Learning Trends for Focal Brain Pathology Segmentation in MRI. pp. 125–148. URL: https://doi.org/10.1007/978-3-319-50478-0_6, doi:10.1007/978-3-319-50478-0_6.
- Heimann, T., Meinzer, H.P., 2009. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis* 13, 543–563.
- Hesamian, M.H., Jia, W., He, X., Kennedy, P., 2019. Deep learning techniques for medical image segmentation: Achievements and challenges. *Journal of Digital Imaging* 32. doi:10.1007/s10278-019-00227-x.
- Hu, X., Li, F., Samaras, D., Chen, C., 2019. Topology-preserving deep image segmentation, in: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc.. pp. 5657–5668. URL: <https://proceedings.neurips.cc/paper/2019/file/2d95666e2649fcfc6e3af75e09f5adb9-Paper.pdf>.
- Hu, X., Wang, Y., Fuxin, L., Samaras, D., Chen, C., 2021. Topology-Aware Segmentation Using Discrete Morse Theory. arXiv:2103.09992 [cs] URL: <http://arxiv.org/abs/2103.09992>. arXiv: 2103.09992 version: 1.
- Isensee, F., Jaeger, P.F., Full, P.M., Wolf, I., Engelhardt, S., Maier-Hein, K.H., 2018. Automatic Cardiac Disease Assessment on cine-MRI via Time-Series Segmentation and Domain Specific Features, in: *Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges*, pp. 120–129. doi:10.1007/978-3-319-75541-0_13.
- Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., et al., 2019. nnu-net: Self-adapting framework for u-net-based medical image segmentation, in: *Bildverarbeitung für die Medizin 2019*. Springer, pp. 22–22.
- Jang, Y., Hong, Y., Ha, S., Kim, S., Chang, H., 2018. Automatic segmentation of lv and rv in cardiac mri, in: *MICCAI workshop on Statistical Atlases and Computational Models of the Heart*, Springer Verlag, Germany. pp. 161–169. doi:10.1007/978-3-319-75541-0_17. 8th International Workshop on Statistical Atlases and Computational Models of the Heart, STACOM 2017, Held in Conjunction with MICCAI 2017 ; Conference date: 10-09-2017 Through 14-09-2017.
- Jiang, F., Grigorev, A., Rho, S., Tian, Z., Fu, Y., Sori, W.J., Khan, A., Liu, S., 2017. Medical image semantic segmentation based on deep learning. *Neural Comput Appl* 29, 1257–1265.
- Kervadec, H., Bouchtiba, J., Desrosiers, C., Granger, E., Dolz, J., Ben Ayed, I., 2019a. Boundary loss for highly unbalanced segmentation, in: *Medical Imaging with Deep Learning*, PMLR, London, UK. pp. 285–296.
- Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., Ayed, I.B., 2019b. Constrained-cnn losses for weakly supervised segmentation. *Medical Image Analysis* 54, 88 – 99. doi:<https://doi.org/10.1016/j.media.2019.02.009>.
- Kervadec, H., Dolz, J., Wang, S., Granger, E., ben Ayed, I., 2020. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision, in: *Medical Imaging with Deep Learning*. URL: <https://openreview.net/forum?id=V0QMC3rZtL>.
- Kervadec, H., Dolz, J., Yuan, J., Desrosiers, C., Granger,

- E., Ayed, I.B., 2019c. Constrained deep networks: Lagrangian optimization via log-barrier extensions. *arXiv:1904.04205*.
- Kim, B., Ye, J.C., 2020. Mumford-Shah Loss Functional for Image Segmentation With Deep Learning. *IEEE Trans. Image Process.* 29, 1856–1866.
- Krithiga, R., Geetha, P., 2020. Breast cancer detection, segmentation and classification on histopathology images analysis: A systematic review. *Archives of Computational Methods in Engineering* URL: <https://doi.org/10.1007/s11831-020-09470-w>, doi:10.1007/s11831-020-09470-w.
- Kuo, W., Angelova, A., Malik, J., Lin, T.Y., 2019. Shape-mask: Learning to segment novel objects by refining shape priors, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9207–9216.
- Lambert, Z., Petitjean, C., Guyader, C.L., 2021. A geometrically-constrained deep network for CT image segmentation, in: *IEEE International Symposium on Biomedical Imaging (ISBI)*.
- Lei, T., Wang, R., Wan, Y., Zhang, B., Meng, H., Nandi, A.K., 2020. Medical image segmentation using deep learning: A survey. *arXiv:2009.13120*.
- Lillo, W.E., Loh, M.H., Hui, S., Zak, S.H., 1993. On solving constrained optimization problems with neural networks: a penalty method approach. *IEEE Transactions on Neural Networks* 4, 931–940. doi:10.1109/72.286888.
- Lin, T.Y., Goyal, P., Girshick, R.B., He, K., Dollár, P., 2017. Focal loss for dense object detection, in: *ICCV*, pp. 2999–3007. URL: <https://doi.org/10.1109/ICCV.2017.324>.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Medical Image Analysis* 42, 60–88. URL: <http://www.sciencedirect.com/science/article/pii/S1361841517301135>, doi:10.1016/j.media.2017.07.005.
- Liu, X., Veksler, O., Samarabandu, J., 2008. Graph cut with ordering constraints on labels and its applications, in: *IEEE CVPR*, pp. 1–8.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation, in: *CVPR*, pp. 3431–3440. doi:10.1109/CVPR.2015.7298965.
- Lorenzo-Valdes, M., Sanchez-Ortiz, G., Mohiaddin, R., Rueckert, D., 2002. Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration, in: *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer Berlin / Heidelberg, Tokyo, Japan. pp. 642–650.
- Magadza, T., Viriri, S., 2021. Deep learning for brain tumor segmentation: A survey of state-of-the-art. *Journal of Imaging* 7. URL: <https://www.mdpi.com/2313-433X/7/2/19>, doi:10.3390/jimaging7020019.
- Marquez-Neila, P., S.M.F.P., 2017. Imposing Hard Constraints on Deep Networks: Promises and Limitations., in: *CVPR Workshop on Negative Results in Computer Vision*.
- Meyer, P., Noblet, V., Mazzara, C., Lallement, A., 2018. Survey on deep learning for radiotherapy. *Computers in Biology and Medicine* 98. doi:10.1016/j.combiomed.2018.05.018.
- Milletari, F., Navab, N., Ahmadi, S., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: *2016 Fourth International Conference on 3D Vision (3DV)*, pp. 565–571. doi:10.1109/3DV.2016.79.
- Milnor, J., 2016. *Morse Theory.* (AM-51). Princeton University Press. URL: <https://doi.org/10.1515/9781400881802>, doi:doi:10.1515/9781400881802.
- Mirikharaji, Z., Hamarneh, G., 2018. Star shape prior in fully convolutional networks for skin lesion segmentation, in: *MICCAI*, pp. 737–745. doi:10.1007/978-3-030-00937-3_84.
- Mosinska, A., Márquez-Neila, P., Kozinski, M., Fua, P., 2018. Beyond the pixel-wise loss for topology-aware delineation, in: *CVPR*, pp. 3136–3145. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Mosinska_Beyond_the_Pixel-Wise_CVPR_2018_paper.html.
- Nocedal, J., Wright, S.J. (Eds.), 1999. *Penalty, Barrier, and Augmented Lagrangian Methods.* Springer New York, New York, NY. pp. 488–525. URL: https://doi.org/10.1007/0-387-22742-3_17, doi:10.1007/0-387-22742-3_17.
- Nosrati, M.S., Hamarneh, G., 2016. Incorporating prior knowledge in medical image segmentation: a survey. *CoRR* abs/1607.01092. URL: <http://arxiv.org/abs/1607.01092>, arXiv:1607.01092.
- Oda, H., Roth, H.R., Chiba, K., Sokolić, J., Kitasaka, T., Oda, M., Hinoki, A., Uchida, H., Schnabel, J.A., Mori, K., 2018. Besnet: Boundary-enhanced segmentation of cells in histopathological images, in: *MICCAI*, pp. 228–236.
- Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M., Bai, W., Caballero, J., Cook, S.A., de Marvao, A., Dawes,

- T., O'Regan, D.P., Kainz, B., Glocker, B., Rueckert, D., 2018. Anatomically constrained neural networks (acnns): Application to cardiac image enhancement and segmentation. *IEEE Transactions on Medical Imaging* 37, 384–395. doi:10.1109/TMI.2017.2743464.
- Pathak, D., Krähenbühl, P., Darrell, T., 2015. Constrained convolutional neural networks for weakly supervised segmentation, in: *ICCV*, pp. 1796–1804. doi:10.1109/ICCV.2015.209.
- Peng, J., Kervadec, H., Dolz, J., Ayed, I.B., Pedersoli, M., Desrosiers, C., 2020. Discretely-constrained deep network for weakly supervised segmentation. *Neural networks : the official journal of the International Neural Network Society* 130, 297–308.
- Petit, O., Thome, N., Soler, L., 2019. Biasing Deep ConvNets for Semantic Segmentation of Medical Images with a Prior-driven Prediction Function, in: *Medical Imaging with Deep Learning*.
- Rahman, M.A., Wang, Y., 2016. Optimizing intersection-over-union in deep neural networks for image segmentation, in: *ISVC*.
- Ravishankar, H., Venkataramani, R., Thiruvankadam, S., Sudhakar, P., Vaidya, V., 2017. Learning and incorporating shape models for semantic segmentation, in: *MICCAI*, pp. 203–211.
- Razzak, M., Naz, S., Zaib, A., 2018. Deep learning for medical image processing: Overview, challenges and the future, in: *Lecture Notes in Computational Vision and Biomechanics*, pp. 323–350. doi:10.1007/978-3-319-65981-7_12.
- Reddy, C., Gopinath, K., Lombaert, H., 2019. Brain tumor segmentation using topological loss in convolutional networks, in: *MIDL*, London, UK. URL: <https://openreview.net/forum?id=Syxck140tE>.
- Renard, F., Guedria, S., De Palma, N., Vuillerme, N., 2020. Variability and reproducibility in deep learning for medical image segmentation. *Scientific Reports* 10, 1–16.
- Rohlfing, T., Brandt, R., Menzel, R., Russakoff, D., Maurer, C., 2005. Quo vadis, atlas-based segmentation ?, in: *The Handbook of Medical Image Analysis: Segmentation and Registration Models*. Kluwer. chapter 11.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: *MICCAI*, pp. 234–241.
- Rousson, M., Paragios, N., 2002. Shape priors for level set representations. *European Conf. on Computer Vision* 2, 78–92.
- Ségonne, F., Fischl, B., 2015. Integration of Topological Constraints in Medical Image Segmentation. Springer US. pp. 245–262. URL: https://doi.org/10.1007/978-0-387-09749-7_13, doi:10.1007/978-0-387-09749-7_13.
- Shi, J., Malik, J., 2000. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence* 22, 888–905.
- Shit, S., Paetzold, J.C., Sekuboyina, A., Zhylka, A., Ezhov, I., Unger, A., Pluim, J.P.W., Tetteh, G., Menze, B.H., 2019. cIDice - a Topology-Preserving Loss Function for Tubular Structure Segmentation, in: *Medical Imaging Meets NeurIPS 2019 Workshop*.
- Slabaugh, G., Unal, G., 2005. Graph cuts segmentation using an elliptical shape prior, in: *International Conference on Image Processing (ICIP)*, pp. 1222–1225.
- Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Cardoso, M.J., 2017. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations, in: *DLMIA/ML-CDS@MICCAI*.
- Taghanaki, S.A., Abhishek, K., Cohen, J.P., Cohen-Adad, J., Hamarneh, G., 2019. Deep Semantic Segmentation of Natural and Medical Images: A Review. arXiv:1910.07655 [cs, eess] URL: <http://arxiv.org/abs/1910.07655>. arXiv: 1910.07655.
- Taha, A.A., Hanbury, A., 2015. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging* 15, 29. URL: <http://www.biomedcentral.com/1471-2342/15/29>.
- Trullo, R., Petitjean, C., Ruan, S., Dubray, B., Nie, D., Shen, D., 2017. Joint segmentation of multiple thoracic organs in CT images with two collaborative deep architectures. *MICCAI'17 workshop Deep Learning in Medical Image Analysis*.
- Veksler, O., 2008. Star shape prior for graph-cut image segmentation, in: *ECCV*, pp. 454–467.
- Vicente, S., Kolmogorov, V., Rother, C., 2008. Graph cut based image segmentation with connectivity priors, in: *IEEE Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8.
- Wang, P., Chen, P., Yuan, Y., Liu, D., Huang, Z., Hou, X., Cottrell, G., 2018. Understanding convolution for semantic segmentation, in: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1451–1460.
- Xu, C., Pham, D.L., Prince, J.L., 2000. Image segmentation using deformable models, in: *Handbook of medical imaging*. SPIE. volume 2, pp. 129–174.
- Yang, S., Kweon, J., Kim, Y.H., 2019. Major vessel segmen-

- tation on x-ray coronary angiography using deep networks with a novel penalty loss function, in: International Conference on Medical Imaging with Deep Learning – Extended Abstract Track, London, UK. URL: <https://openreview.net/forum?id=H11Th8unKN>.
- Yang, X., Bian, C., Yu, L., Ni, D., Heng, P.A., 2018. Class-balanced deep neural network for automatic ventricular structure segmentation, in: STACOM@MICCAI, pp. 152–160.
- Yu, C., Wang, J., Gao, C., Yu, G., Shen, C., Sang, N., 2020. Context prior for scene segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Yue, Q., Luo, X., Ye, Q., Xu, L., Zhuang, X., 2019. Cardiac segmentation from lge mri using deep neural network incorporating shape and spatial priors, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), Medical Image Computing and Computer Assisted Intervention – MICCAI 2019, Springer International Publishing, Cham. pp. 559–567.
- Zhang, J., Petitjean, C., Ainouz, S., 2020. Kappa loss for skin lesion segmentation in fully convolutional network, in: IEEE ISBI, Iowa City USA, pp. 2001–2004.
- Zhou, T., Ruan, S., Canu, S., 2019. A review: Deep learning for medical image segmentation using multi-modality fusion. Array 3-4, 100004. URL: <https://www.sciencedirect.com/science/article/pii/S2590005619300049>, doi:<https://doi.org/10.1016/j.array.2019.100004>.
- Zotti, C., Luo, Z., Humbert, O., Lalande, A., Jodoin, P., 2017. Gridnet with automatic shape prior registration for automatic MRI cardiac segmentation, in: Statistical Atlases and Computational Models of the Heart STACOM, Held in Conjunction with MICCAI, Quebec City, Canada, pp. 73–81. doi:10.1007/978-3-319-75541-0_8.