Pathology Segmentation using Distributional Differences to Images of Healthy Origin

Simon Andermatt, Antal Horváth, Simon Pezold, and Philippe Cattin

Department of Biomedical Engineering, University of Basel, Allschwil, Switzerland

Abstract. Fully supervised segmentation methods require a large training cohort of already segmented images, providing information at the pixel level of each image. We present a method to automatically segment and model pathologies in medical images, trained solely on data labelled on the image level as either healthy or containing a visual defect. We base our method on CycleGAN, an image-to-image translation technique, to translate images between the domains of healthy and pathological images. We extend the core idea with two key contributions. Implementing the generators as residual generators allows us to explicitly model the segmentation of the pathology. Realizing the translation from the healthy to the pathological domain using a variational autoencoder allows us to specify one representation of the pathology, as this transformation is otherwise not unique. Our model hence not only allows us to create pixelwise semantic segmentations, it is also able to create inpaintings for the segmentations to render the pathological image healthy. Furthermore, we can draw new unseen pathology samples from this model based on the distribution in the data. We show quantitatively, that our method is able to segment pathologies with a surprising accuracy being only slightly inferior to a state-of-the-art fully supervised method, although the latter has per-pixel rather than per-image training information. Moreover, we show qualitative results of both the segmentations and inpaintings. Our findings motivate further research into weakly-supervised segmentation using image level annotations, allowing for faster and cheaper acquisition of training data without a large sacrifice in segmentation accuracy.

1 Introduction

Supervised segmentation in medical image analysis is an almost solved problem for many applications, where methodological improvements have a marginal effect on accuracy. Such methods depend on a large annotated training corpus, where pixelwise labels have to be provided by medical experts. In practice, such data are expensive to gather. In contrast, weakly labelled data, such as images showing a certain disease are easily obtainable, since they are created on a daily basis in medical practice. We want to take advantage of these data for pathology segmentation, by providing a means to finding the difference between healthy and pathological data distributions. We present a weakly supervised framework capable of pixelwise segmentation as well as generating samples from the pathology distribution, trained on images which are only annotated with a scalar binary label at the image level classifying them as healthy or pathological.

Our idea is inspired by CycleGAN [11], a recently proposed solution for unpaired image to image translation, where the combination of domain-specific generative adversarial networks (GANs) and so-called cycle consistency allow for robust translation. We call our adaptation PathoGAN and count the following contributions: We formulate a model capable of segmentation based on a single label per training sample. We simultaneously train two generative models, able to generate inpaintings at a localized region of interest to transform an image from one domain to the other. We are able to sample healthy anatomy as well as sample possible pathologies for a given anatomical structure. Furthermore, our method enforces domain-specific information to be encoded outside of the image, which omits adversarial "noise" common to CycleGAN [4] to some degree.

We show the performance of our implementation on 2d slices of the training data of the Brain Tumor Segmentation Challenge 2017 [9,2] and compare our segmentation performance to a supervised segmentation technique [1].

CycleGAN has been previously used to segment by transfering to another target modality, where segmentation maps are available (e.g. [10]), or applied to generate training from cheaply generated synthetic labelmaps [6]. Using a Wasserstein GAN, another method directly estimates an additive visual attribution map [3]. To our knowledge, there has not been a method that jointly learns to segment on one medical imaging modality using only image-level labels and generate new data using GANs for both healthy and pathological cases.

2 Methods

2.1 Problem Statement

We assume two image domains $\mathcal{A}$ and $\mathcal{B}$, where the former contains only images of healthy subjects and the latter consists of images showing a specific pathology. We seek to approximate the functions $G_{\mathcal{A}}$ and $G_{\mathcal{B}}$ that perform the mappings $(x_{\mathcal{A}}, \delta_{\mathcal{B}}) \mapsto \hat{y}_B$ and $(x_{\mathcal{B}}, \delta_{\mathcal{A}}) \mapsto \hat{y}_{\mathcal{A}}$, where $x_{\mathcal{A}}, \hat{y}_{\mathcal{A}} \in \mathcal{A}$ and $x_{\mathcal{B}}, \hat{y}_{\mathcal{B}} \in \mathcal{B}$. Vectors $\delta_{\mathcal{B}}$ and $\delta_{\mathcal{A}}$ encode the missing target image information (e.g. the pathology):

$$\hat{y}_{\mathcal{B}} = G_{\mathcal{A}}(x_{\mathcal{A}}, \delta_{\mathcal{B}}), \quad \hat{y}_{\mathcal{A}} = G_{\mathcal{B}}(x_{\mathcal{B}}, \delta_{\mathcal{A}}). \tag{1}$$

We encourage $G_{\mathcal{A}}, G_{\mathcal{B}}$ to produce results, such that the mappings are realistic (2), cycle-consistent (3), specific (4) and that only the affected part in the image is modified (5):

$$G_{\mathcal{A}}(x_{\mathcal{A}}, \delta_{\mathcal{B}}) \sim \mathcal{B}, \qquad G_{\mathcal{B}}(x_{\mathcal{B}}, \delta_{\mathcal{A}}) \sim \mathcal{A}, \tag{2}$$

$$G_{\mathcal{B}}(G_{\mathcal{A}}(x_{\mathcal{A}}, \delta_{\mathcal{B}}), \delta_{\mathcal{A}}) \approx x_{\mathcal{A}}, \qquad G_{\mathcal{A}}(G_{\mathcal{B}}(x_{\mathcal{B}}, \delta_{\mathcal{A}}), \delta_{\mathcal{B}}) \approx x_{\mathcal{B}}, \tag{3}$$

$$G_{\mathcal{B}}(x_{\mathcal{A}}, 0) \approx x_{\mathcal{A}}, \qquad G_{\mathcal{A}}(x_{\mathcal{B}}, 0) \approx x_{\mathcal{B}}, \tag{4}$$

$$\underset{\overline{G}_{\mathcal{A}}}{\arg\min} |x_{\mathcal{A}} - \overline{G}_{\mathcal{A}}(x_{\mathcal{A}}, \delta_{\mathcal{B}})| \approx G_{\mathcal{A}}, \qquad \underset{\overline{G}_{\mathcal{B}}}{\arg\min} |x_{\mathcal{B}} - \overline{G}_{\mathcal{B}}(x_{\mathcal{B}}, \delta_{\mathcal{A}})| \approx G_{\mathcal{B}}. \tag{5}$$

2.2 Model Topology

To fulfill Eqs. (2–5), we adopt the main setup and objective from CycleGAN: we employ two discriminators, $D_{\mathcal{A}}$ and $D_{\mathcal{B}}$ together with generators $G_{\mathcal{A}}$ and $G_{\mathcal{B}}$ to perform the translation from domain $\mathcal{A}$ to $\mathcal{B}$ and vice versa, formulating two generative adversarial networks (GANs) [7]. In both directions, the respective discriminator is trained to distinguish a real image from the output of the generator, whereas the generator is incentivized to generate samples that fool the discriminator.

In the remaining paper, we will use a short notation for equations which are applicable to both pathways (see Fig. 1) to overcome redundancy due to symmetrical components. We specify placeholder domains $\mathcal{X}$ and $\mathcal{Y}$, where for both pathways, $\mathcal{X}$ denotes the domain the initial image belongs to and $\mathcal{Y}$ stands for the target domain we want to transfer an image into. Domain $\mathcal{X}$ could therefore be either $\mathcal{A}$ or $\mathcal{B}$, and domain $\mathcal{Y}$ either $\mathcal{B}$ or $\mathcal{A}$, respectively. Furthermore, $x_{\mathcal{X}} \in \mathcal{X}$ denotes the image x in the respective initial domain $\mathcal{X}$, $G_{\mathcal{X}}(x_{\mathcal{X}}) = \hat{y}_{\mathcal{Y}}$ its transformation in the target domain $\mathcal{Y}$ and $G_{\mathcal{Y}}(\hat{y}_{\mathcal{Y}}) = \tilde{x}_{\mathcal{X}}$ its reconstruction in $\mathcal{X}$.

Residual Generator In order to segment pathologies, we seek to only modify a certain part of the image. In contrast to CycleGAN, we model the transformation G from one domain to the other as a residual or inpainting p which is exchanged with part l of the original image. We achieve this by directly estimating $n + 1$ feature maps $r_{\mathcal{X}}$ using a residual generator $Z_{\mathcal{X}}$ within $G_{\mathcal{X}}$, where n is the number of image channels used. We obtain labelmap $l_{\mathcal{X}}$ and inpaintings $p_{\mathcal{X}}$, activating $r_{\mathcal{X}}^{(0)}$ with a sigmoid and each $r_{\mathcal{X}}^{(i)}$ with a tanh activation:

$$l_{\mathcal{X}} = \mathrm{S}\left(r_{\mathcal{X}}^{(0)} + \epsilon\right), \quad p_{\mathcal{X}}^{(i-1)} = \tanh\left(r_{\mathcal{X}}^{(i)}\right), \quad r_{\mathcal{X}} = Z_{\mathcal{X}}(x_{\mathcal{X}}, \delta_{\mathcal{Y}}), \qquad (6)$$

where $\mathrm{S}\left(y\right) = \frac{1}{1+e^{-y}}$ and $i > 0$. With $\epsilon \sim \mathcal{N}(0, I)$, we turn $r_{\mathcal{X}}^{(0)} + \epsilon$ into samples from $\mathcal{N}(r_{\mathcal{X}}^{(0)}, I)$ using the reparameterization trick [8]. This allows reliable calculations of $l_{\mathcal{X}}$ only for large absolute values of $r_{\mathcal{X}}^{(0)}$, forcing $l_{\mathcal{X}}$ to be binary and intensity information to be encoded in the inpaintings. We set ϵ to zero during testing. From $l_{\mathcal{X}}$ and $p_{\mathcal{X}}$ we compute the translated result $\hat{y}_{\mathcal{Y}}$, supposedly in domain $\mathcal{Y}$ now:

$$\hat{y}_{\mathcal{Y}} = l_{\mathcal{X}} \odot p_{\mathcal{X}} + (1 - l_{\mathcal{X}}) \odot x_{\mathcal{X}} = G_{\mathcal{X}}(x_{\mathcal{X}}, \delta_{\mathcal{Y}}).$$

In the following, we detail the computation of $r_{\mathcal{A}}$ and $r_{\mathcal{B}}$ using the two networks $Z_{\mathcal{A}}$ and $Z_{\mathcal{B}}$ for the two possible translation directions. Both translation pathways are visualized in Fig. 1.

$\boldsymbol{Z_{\mathcal{A}} : \mathcal{A} \rightarrow \mathcal{B}}$ To map from healthy to pathological data, we estimate $r_{\mathcal{A}}$ (and thus $l_{\mathcal{A}}, p_{\mathcal{A}}$) using a variational autoencoder (VAE) [8]. First, we employ

(a) $\mathcal{A} \to \mathcal{B} \to \mathcal{A}$

(b) $\mathcal{B} \to \mathcal{A} \to \mathcal{B}$

Fig. 1. Proposed architecture: *top to bottom:* directions $\mathcal{A} \to \mathcal{B} \to \mathcal{A}$ and $\mathcal{B} \to \mathcal{A} \to \mathcal{B}$; $x_{\mathcal{A}}, x_{\mathcal{B}}$ are samples from the two data distributions $\mathcal{A}$ (healthy), $\mathcal{B}$ (pathological) and $\delta_{\mathcal{B}} \sim \mathcal{N}(0, I)$. Red and blue triangles depict decoder and encoder networks. A red square illustrates a simple generator. $\Delta_{\mathcal{B}}$ and $\Gamma_{\mathcal{A}}$ encode features inside and outside the pathological region. $\Delta_{\mathcal{B}}, \Gamma_{\mathcal{A}}$ and $E_{\mathcal{A}}$ form a variational autoencoder, which acts as the residual generator $Z_{\mathcal{A}}$ used in $G_{\mathcal{A}}$. On the other hand, the residual generator $Z_{\mathcal{B}}$ in $G_{\mathcal{B}}$ is implemented as one component, as information about the missing healthy structure is completely inferred from the surroundings without explicitly encoding it.

encoders $\Gamma_\mathcal{A}$ and $\Delta_\mathcal{B}$ to encode anatomical information around and inside the pathological region:

$$\gamma_\mathcal{A} = \Gamma_\mathcal{A}(x_\mathcal{A}), \qquad \delta_\mathcal{B} = \Delta_\mathcal{B}(l'_\mathcal{B} \odot x'_\mathcal{B}),$$

where $x_\mathcal{A}$ is our healthy image, $l'_\mathcal{B}$ and $x'_\mathcal{B}$ are the labelmap and pathological image of the previous transformation and $\delta_\mathcal{B}, \gamma_\mathcal{A} \sim \mathcal{N}(0, I)$. If $l'_\mathcal{B}$ and $x'_\mathcal{B}$ are not available because $x_\mathcal{A}$ is a real healthy image, we simply sample $\delta_\mathcal{B}$. Finally, a decoder $E_\mathcal{A}$ is applied to $\gamma_\mathcal{A}$ and $\delta_\mathcal{B}$:

$$r_\mathcal{A} = E_\mathcal{A}(\gamma_\mathcal{A}, \delta_\mathcal{B}).$$

The residual generator $Z_\mathcal{A}$ is hence composed of encoder $\Gamma_\mathcal{A}$ and decoder $E_\mathcal{A}$, where $\delta_\mathcal{B}$ is either sampled or calculated with the additional encoder $\Delta_\mathcal{B}$:

$$Z_\mathcal{A}(x_\mathcal{A}, \delta_\mathcal{B}) = E_\mathcal{A}(\Gamma_\mathcal{B}(x_\mathcal{A}), \delta_\mathcal{B}),$$

$Z_\mathcal{B} : \mathcal{B} \to \mathcal{A}$ To generate healthy samples from pathological images, we use the residual generator $Z_\mathcal{B}$ directly on the input as introduced in [11] and estimate r directly:

$$r_\mathcal{B} = Z_\mathcal{B}(x_\mathcal{B}).$$

Here, we omit $\delta_\mathcal{A}$, since the location and appearance of missing healthy tissue can be inferred from $x_\mathcal{B}$. We also omit using an encoding bottleneck due to possible information loss and less accurate segmentation.

2.3 Objective

To train this model, a number of different loss terms are necessary. In the following, we explain the individual components using $\hat{y}_.$ and $\tilde{x}_.$ to denote results from the translated and reconstructed images respectively (e.g. mapping $x_\mathcal{X}$ into $\mathcal{Y}$ results in $\hat{y}_\mathcal{Y}$, translating it back results in $\tilde{x}_\mathcal{X}$). We use $\lambda_.$ to weight the contribution of different loss terms.

CycleGAN As in CycleGAN [11], we formulate a least squares proxy GAN loss, which we minimize with respect to $G_\mathcal{X}$ and maximize with respect to $D_\mathcal{Y}$:

$$\mathcal{L}_{\text{GAN}}(D_\mathcal{Y}, G_\mathcal{X}, x_\mathcal{X}, x_\mathcal{Y}) = \mathbb{E}[(D_\mathcal{Y}(x_\mathcal{Y}))^2] + \mathbb{E}[(1 - D_\mathcal{Y}(G_\mathcal{X}(x_\mathcal{X}, \delta_\mathcal{Y})))^2]. \quad (7)$$

Likewise, to make mappings reversible, we add the cycle-consistency loss:

$$\mathcal{L}_{\text{CC}}(G_\mathcal{X}, G_\mathcal{Y}, x_\mathcal{X}) = \lambda_{\text{CC}} \|G_\mathcal{Y}(G_\mathcal{X}(x_\mathcal{X}, \delta_\mathcal{Y}), \delta_\mathcal{X}) - x_\mathcal{X}\|_1. \quad (8)$$

$\mathcal{L}_{\text{GAN}}$ and $\mathcal{L}_{\text{CC}}$ encourage the properties defined in Eqs. (2) and (3).

Variational Autoencoder A variational autoencoder (VAE) is trained by minimizing the KL-divergence of the distribution $q(z|x)$ of encoding z to some assumed distribution and the expected reconstruction error $\log p(x|z)$, where x is the data. In contrast to a classical VAE, we use two distinct encoding vectors $\gamma_{\mathcal{A}}$ and $\delta_{\mathcal{B}}$, encoding the healthy and the pathological part, and produce two separate results, the labelmap $l_{\mathcal{A}}$ and the inpainting $p_{\mathcal{A}}$. We directly calculate the KL-divergence for our two encodings:

$$\mathcal{L}_{\mathrm{KL}}(G_{\mathcal{A}}, G_{\mathcal{B}}, x_{\mathcal{A}}, x_{\mathcal{B}}) = \mathrm{KL}[q(\gamma_{\mathcal{A}}|x_{\mathcal{A}})||\mathcal{N}(0, I)] + \mathrm{KL}[q(\delta_{\mathcal{B}}|x_{\mathcal{B}}, \hat{l}_{\mathcal{B}})||\mathcal{N}(0, I)]. \tag{9}$$

For the expected reconstruction error, we assume that l and p follow approximately a Bernoulli and a Gaussian distribution ($\mathcal{N}(\mu, I)$). We selectively penalize the responsible encoding, by using separate loss functions for the residual region and the rest. Unfortunately, we only ever have access to the ground truth of one of these regions, since we do not use paired data. We solve this, by using the relevant application in the network, where individual ground truths are available, to calculate the approximation of the marginal likelihood lower bound:

$$\mathcal{L}_{\mathrm{VAE}}(G_{\mathcal{X}}, G_{\mathcal{Y}}, x_{\mathcal{X}}, x_{\mathcal{Y}}) =$$
$$- \frac{\lambda_{\mathrm{VAE}}}{N} \sum_{m=1}^{N} (\log p(\tilde{l}_{\mathcal{Y}}|\gamma_{\mathcal{X}}, \delta_{\mathcal{Y}}) + \log p(\hat{p}_{\mathcal{X}}|\gamma_{\mathcal{X}}) + \log p(\tilde{p}_{\mathcal{Y}}|\delta_{\mathcal{Y}})), \tag{10}$$

where $\hat{p}_{\mathcal{X}}$ denotes the inpainting used to translate the original image $x_{\mathcal{X}}$ to domain $\mathcal{Y}$ and N is the total number of pixels. $\tilde{p}_{\mathcal{X}}$ is the inpainting produced when translating an already translated image $\hat{y}_{\mathcal{X}}$ that originated from $\mathcal{Y}$ back to that domain. Similarly, $\hat{l}_{\mathcal{X}}$ and $\tilde{l}_{\mathcal{X}}$ denote the respective labelmaps:

$$\log p(\hat{p}_{\mathcal{X}}|\gamma_{\mathcal{X}}) = \frac{||(1 - \hat{l}_{\mathcal{X}})(\hat{p}_{\mathcal{X}} - x_{\mathcal{X}})||_2}{\omega_1}, \tag{11}$$

$$\log p(\tilde{p}_{\mathcal{X}}|\delta_{\mathcal{Y}}) = \frac{||\tilde{l}_{\mathcal{Y}}(\tilde{p}_{\mathcal{X}} - x_{\mathcal{Y}})||_2}{\omega_2}, \tag{12}$$

where $\omega_1 = \frac{\sum(1 - \hat{l}_{\mathcal{X}}) + \varepsilon}{N}$ and $\omega_2 = \frac{\sum(\tilde{l}_{\mathcal{Y}}) + \varepsilon}{N}$ are considered constant during optimization, with $\varepsilon > 0$. Finally, we use the labelmap produced by the other generator responsible for the opposite transformation $\hat{l}_{\mathcal{Y}}$ as ground truth for $\tilde{l}_{\mathcal{X}}$, where we consider $\hat{l}_{\mathcal{Y}}$ constant in this term:

$$\log p(\tilde{l}_{\mathcal{X}}|\gamma_{\mathcal{X}}, \delta_{\mathcal{Y}}) = \hat{l}_{\mathcal{Y}} \log \tilde{l}_{\mathcal{X}} + (1 - \hat{l}_{\mathcal{Y}}) \log(1 - \tilde{l}_{\mathcal{X}}). \tag{13}$$

To restrict the solution space of our model, we use $\mathcal{L}_{\mathrm{VAE}}$ for both directions.

Identity Loss We apply an identity loss [11] on labelmap $l_{\mathcal{X}, x_{\mathcal{Y}}}$ which results from feeding $G_{\mathcal{X}}$ with the wrong input $x_{\mathcal{Y}}$. In this case $G_{\mathcal{X}}$ should not change anything, since the input is already in the desired domain $\mathcal{Y}$:

$$\mathcal{L}_{\mathrm{Idt}}(G_{\mathcal{X}}, x_{\mathcal{Y}}) = \lambda_{\mathrm{Idt}} ||l_{\mathcal{X}, x_{\mathcal{Y}}}||_1. \tag{14}$$

Relevancy Loss By now, we have defined all necessary constraints for a successful translation between image domains. The remaining constraints restrict the location and amount of change, $l_\mathcal{X}$. Fulfilling Eq. (5), we want to entice label map $l_\mathcal{X}$ to be only set at locations of a large difference between inpainting $p_\mathcal{X}$ and image $x_\mathcal{X}$ and penalize large label maps in general:

$$\mathcal{L}_\mathrm{R}(G_\mathcal{X}, x_\mathcal{X}) = \lambda_\mathrm{R} \left[|| - \log(1 - l_\mathcal{X}^2)||_1 - \frac{||l_\mathcal{X}(x_\mathcal{X} - p_\mathcal{X})||_1}{||l_\mathcal{X}||_1} \right]. \tag{15}$$

In order to not reward exaggerated pathology inpaintings, we consider $(x_\mathcal{X} - p_\mathcal{X})$ constant in this expression.

Full PathoGAN Objective combining all loss terms for direction $\mathcal{X}$ to $\mathcal{Y}$ as $\mathcal{L}_{\mathcal{X} \to \mathcal{Y}}$, we can finally define:

$$\mathcal{L}_{\mathcal{X} \to \mathcal{Y}} = \mathcal{L}_\mathrm{GAN} + \mathcal{L}_\mathrm{CC} + \mathcal{L}_\mathrm{VAE} + \mathcal{L}_\mathrm{Idt} + \mathcal{L}_\mathrm{R}, \tag{16}$$

$$\mathcal{L}_\mathrm{PathoGAN} = \mathcal{L}_{\mathcal{A} \to \mathcal{B}} + \mathcal{L}_{\mathcal{B} \to \mathcal{A}} + \lambda_\mathrm{KL} \mathcal{L}_\mathrm{KL}(x_\mathcal{A}, x_\mathcal{B}). \tag{17}$$

2.4 Network structure

We describe the networks used in our generators $G_\mathcal{A}$ and $G_\mathcal{B}$ using a shorthand notation which we will define in the following. We denote the number of output feature maps as `f`, the number of inpaintings including labelmap as `r`, the smallest image width and height before reshaping to a vector as `i` and the encoding size as `z`. The smallest image width and height in our case is 15 and we chose an encoding length of 256. We denote the reshaping operation from a square image to a vector as `Q2F` and the inverse operation as `F2Q`. We further denote a convolution operation with stride 1 and a kernel size of `k` as `Ck-f`. We use a lower-case `c` in `ck-f` if the convolution is followed by an instance norm and an activation with an exponential linear unit (ELU) [5]. A downsampling operation of a stride 2 convolution with kernel size 3, followed with an instance norm and following ELU is coded as `df`, whereas an upsampling operation using a transposed convolution followed by an instance norm and an ELU activation is denoted as `uf`. We describe a fully connected layer with `l(f)` and a residual block as `Rf`. A residual block results in a sum of its input with a residual computation of a convolution layer of kernel size 3, an instance norm, an ELU and a convolution layer of kernel size 3. Finally, a trailing `t` or `e` denotes an additional tanh or ELU activation, respectively.

The generator $G_\mathcal{A}$ consists of the residual generator $Z_\mathcal{A}$, which is composed of two encoder networks $(\varGamma_\mathcal{A}, \varDelta_\mathcal{B})$ and a decoder network $(E_\mathcal{A})$. Both $\varGamma_\mathcal{A}$ and $\varDelta_\mathcal{B}$ are defined as:

```
c7-64,d128,d256,d512,d1024,C1-15,Q2F,l(z*i)t,l(2*z).
```

The output of these encoders is split in half to produce the mean and log variance vectors of length 256 each.

Similarly, the decoder network $E_\mathcal{A}$ is defined as:

```
l(i*i)e,l(i*i),F2Q,c3-1024,u512,u256,C7-256e,R256,R256,R256,
R256,R256,R256,R256,R256,R256,u128,u64,C7-r.
```

The generator $G_\mathcal{B}$ consists of the residual generator network $Z_\mathcal{B}$, which is slightly adjusted compared to the generators used in [11]. Here, we use exponential linear units instead of rectified linear units and double each layers' number of feature maps:

```
c7-64,d128,d256,R256,R256,R256,R256,R256,R256,R256,R256,
    u128,u64,C7-r.
```

We use the same network architecture as in CycleGAN [11] for the discriminator networks.

2.5 Data

We include all training patients of Brats2017 and normalize each brain scan for its non-zero voxels to follow $\mathcal{N}(0, 1/3)$, and clip the resulting intensities to $[-1, 1]$. We select all transverse slices from 60 to 100 in caudocranial direction. In order to create two distinct datasets and relying on the manual segmentations, we label slices without pathology as *healthy*, with more than 20 pixels segmented as *pathological* slices, and discard the rest. For training, we select 1 500 unaffected and 6 000 pathological slices from a total of 1 755 and 9 413 respectively[1].

We augment our training data by randomly mirroring the samples, applying a random rotation sampled from $U[-0.1, 0.1]$ and a random scaling $s = 1.1^r$ where r is sampled from $U[-1, 1]$. Furthermore we apply random deformations as described in [1], with a grid spacing of 128 and sample both components of the grid deformation vectors from $\mathcal{N}(0, 5)$.

Table 1. Segmentation Results. *Columns:* Dice, 95th percentile Hausdorff distance (HD95), average Hausdorff distance (AVD) and volumetric Dice per-patient (Dice PP) by stacking all evaluated slices. *Rows:* Scores are shown as mean±std(median) for the weakly-supervised PathoGAN (proposed) and the fully-supervised MDGRU, applied to training (Tr) and testing (Te) data.

|	Dice (in %)	HD95 (in pixel)	AVD (in pixel)	Dice PP (in %)
PathoGAN, (Tr)	$72.4 \pm 24.4(81.0)$	$40.6 \pm 30.7(38.0)$	$10.3 \pm 15.4(4.7)$	$77.4 \pm 14.4(81.2)$
PathoGAN, (Te)	$72.9 \pm 23.8(81.4)$	$39.4 \pm 29.9(37.6)$	$9.4 \pm 13.7(4.6)$	$77.4 \pm 14.4(81.7)$
MDGRU, (Tr)	$87.8 \pm 20.0(94.4)$	$3.7 \pm 9.7(1.0)$	$1.0 \pm 4.7(0.2)$	$90.8 \pm 8.8(93.3)$
MDGRU, (Te)	$86.3 \pm 21.3(93.6)$	$3.9 \pm 9.5(1.0)$	$1.1 \pm 4.9(0.2)$	$90.6 \pm 9.5(93.1)$

---

[1] Thus we would like to stress that the manual segmentations were only used to create the two image domains, but not for the actual training.

3 Results

Since the BratS evaluation is volumetric and comparing performance is difficult, we also train a supervised segmentation technique on our data. We chose MD-GRU [1] for this task, a multi-dimensional recurrent neural network, due to code availability and consistent state-of-the-art performance on different datasets.

We trained PathoGAN[2] for 119 epochs using batches of 4 and $\lambda_{\mathrm{KL}} = 0.1, \lambda_{\mathrm{R}} = 0.5, \lambda_{\mathrm{Idt}} = 1, \lambda_{\mathrm{CC}} = 5$ and $\lambda_{\mathrm{VAE}} = 1$. We trained MDGRU[3] as defined in [1], using batches of 4 and $27\,500$ iterations. Table 1 shows the results on the pathological training and test data. Figure 2 shows an exemplary sample from the test data. On the left, the input data, the generated inpaintings and the translation result are displayed. On the right, the manual segmentation for "whole tumor" and generated labelmaps of the weakly-supervised PathoGAN and fully-supervised MDGRU are presented.

Fig. 2. Qualitative Results of one example from the testing data. *Left columns, top to bottom:* the four available image channels $x_{\mathcal{B}}^{(i)}$, the generated inpaintings $\hat{p}_{\mathcal{B}}^{(i)}$ and the translated images $\hat{y}_{\mathcal{A}}^{(i)}$. *Right column, top to bottom:* The manual segmentation l_{GS}, the probability maps $\hat{l}_{\mathcal{B}}$ from PathoGAN (weakly-supervised, proposed) and l_{M} from MDGRU (fully-supervised) for whole tumor.

---

[2] Our implementation is based on `https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix`.

[3] We use the implementation of MDGRU at `https://github.com/zubata88/mdgru`.

4 Discussion

The results in Fig. 2 indicate that our relative weighting of the two inpainting reconstruction losses results in better reconstruction inside the tumor region than outside. The labelmaps of the supervised method compared to ours in Fig. 2 show great agreement, and both are relatively close to the gold standard. As the 95th-percentile and average Hausdorff measures in Table 1 show, there are some outliers in our proposed method, due to its weakly-supervised nature. Compared to the fully-supervised method, Dice scores for PathoGAN are about 10% smaller for both the per-slice and the per-patient case. Slightly inferior results for the proposed method are not surprising, though, given the drastically reduced information density during its training (one label per image for PathoGAN, one label per pixel for fully-supervised MDGRU). Also, it is important to remember that we segment with the only criterion of being not part of the healthy distribution, which could vary from the subjective measures used to manually segment data. The increase in accuracy and decrease in standard deviation in the per-patient case for both methods is most likely caused by the inferior segmentation performance in slices showing little pathology. The per-patient Dice of the supervised method is in the range of the top methods of BraTS 2017. Although not directly comparable, this suggests that we can use our computed supervised scores as good state-of-the-art reference to compare our results to.

We did only scratch the surface on the possible applications of our proposed formulation. Future work will include unaffected samples that are actually healthy. Furthermore, the model architecture could be drastically simplified using one discriminator for both directions, allowing for larger generator networks as well as using multiple discriminators at different scales to find inpaintings that are not just locally but also globally consistent with the image. A restriction to slices is unfortunate but necessary due to memory requirements. A generalisation of our approach to volumetric data would make it feasible for more real clinical applications.

Conclusion We presented a new generative pathology segmentation model capable of handling a plethora of tasks: First and foremost, we presented a weakly supervised segmentation method for pathologies in 2D medical images, where it is only known if the image is affected by the pathology and thus no pixel-wise label or classification is provided. Furthermore, we were able to sample from both our healthy as well as our pathology model. We showed qualitatively and quantitatively, that we are able to produce compelling results, motivating further research towards actual clinical applications of PathoGAN.

5 Acknowledgements

References

1. Andermatt, S., Pezold, S., Cattin, P.: Multi-dimensional Gated Recurrent Units for the Segmentation of Biomedical 3D-Data. In: International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis. pp. 142–151. Springer (2016)
2. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J., Freymann, J., Farahani, K., Davatzikos, C.: Advancing The Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features. Nature Scientific Data (2017), [in press]
3. Baumgartner, C.F., Koch, L.M., Tezcan, K.C., Ang, J.X., Konukoglu, E.: Visual feature attribution using wasserstein gans. arXiv preprint arXiv:1711.08998 (2017)
4. Chu, C., Zhmoginov, A., Sandler, M.: CycleGAN, a Master of Steganography. arXiv:1712.02950 [cs, stat] (Dec 2017)
5. Clevert, D.A., Unterthiner, T., Hochreiter, S.: Fast and accurate deep network learning by exponential linear units (elus). arXiv preprint arXiv:1511.07289 (2015)
6. Fu, C., Lee, S., Ho, D.J., Han, S., Salama, P., Dunn, K.W., Delp, E.J.: Fluorescence microscopy image segmentation using convolutional neural network with generative adversarial networks. arXiv preprint arXiv:1801.07198 (2018)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)
8. Kingma, D.P., Welling, M.: Stochastic gradient vb and the variational auto-encoder. In: Second International Conference on Learning Representations (2014)
9. Menze, B.H., Jakab, A., Bauer, S., et al: The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). IEEE Transactions on Medical Imaging 34(10), 1993–2024 (Oct 2015)
10. Tsaftaris, S.A.: Adversarial image synthesis for unpaired multi-modal cardiac data. In: Simulation and Synthesis in Medical Imaging: Second International Workshop, SASHIMI 2017, Held in Conjunction with MICCAI 2017. p. 3. Springer (2017)
11. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. pp. 2223–2232 (2017)