

Light In The Black: An Evaluation of Data Augmentation Techniques for COVID-19 CT's Semantic Segmentation

Bruno A. Krinski, Daniel V. Ruiz, and Eduardo Todt

Department of Informatics
Federal University of Paraná (UFPR) – Curitiba, PR – Brazil

{bakrinski, dvruiz, todt}@inf.ufpr.br

Abstract. *With the COVID-19 global pandemic, computer-assisted diagnoses of medical images have gained much attention, and robust methods of Semantic Segmentation of Computed Tomography (CT) became highly desirable. Semantic Segmentation of CT is one of many research fields of automatic detection of COVID-19 and has been widely explored since the COVID-19 outbreak. In this work, we propose an extensive analysis of how different data augmentation techniques improve the training of encoder-decoder neural networks on this problem. Twenty different data augmentation techniques were evaluated on five different datasets. Each dataset was validated through a five-fold cross-validation strategy, thus resulting in over 3,000 experiments. Our findings show that spatial level transformations are the most promising to improve the learning of neural networks on this problem.*

Resumo. *Com a COVID-19, diagnósticos de imagens médicas assistidos por computador ganharam muita atenção, e métodos robustos de Segmentação Semântica de Tomografia Computadorizada (TC) tornaram-se altamente desejáveis. A Segmentação Semântica de TC é um dos muitos campos de pesquisa de detecção automática da COVID-19 e foi amplamente explorado desde o surto da COVID-19. Neste trabalho, propomos uma análise extensiva sobre o quanto diferentes técnicas de aumento de dados contribuem para melhorar o treinamento de redes neurais codificador-decodificador sobre este problema. Vinte técnicas diferentes de aumento de dados foram avaliadas em cinco conjuntos de dados diferentes. Cada conjunto de dados foi validado através de uma estratégia de validação cruzada de cinco subconjuntos, resultando assim em mais de 3.000 experimentos. Nossas descobertas mostram que as transformações de nível espacial são as mais promissoras para melhorar o aprendizado das redes neurais sobre este problema.*

1. Introduction

Since 2019 the world has struggled with the new coronavirus (COVID-19) pandemic, with millions of infections and deaths worldwide [Wang et al. 2020]. Until now, there are a total of 427,169,421 global cases and a total of 5,902,878 global deaths [of Medicine 2022] (updated February 22th, 2022). Due to the virus's quick dissemination, early diagnosis is highly desirable for faster treatment and tracking infected people [Chen et al. 2020a]. Automatic detection of COVID-19 infections in Computed Tomographs (CTs) shows

to be a great help for early diagnoses [Shi et al. 2021], with the Semantic Segmentation [Cao and Bao 2020] of CTs being widely explored since the COVID-19 outbreak [Shi et al. 2021]. Deep Learning based techniques and Deep Neural Networks achieved impressive results in the segmentation of COVID-19 CTs [Shi et al. 2021, Krinski et al. 2021]. However, it has two limiting factors. The first one is that labeling Semantic Segmentation is a labor-intensive and timing-consuming process, and each pixel of the image must receive the correct label. Otherwise, the network could converge to wrong results [Shi et al. 2021, Cao and Bao 2020]. In addition to that, labeling CT segmentation datasets must be made by highly specialized doctors to properly label the lesion regions of the image [Shi et al. 2021].

With new approaches being proposed quickly, an urgency aggravated by the global pandemic, the need for a proper evaluation becomes apparent. A broad benchmark of architectures was presented by [Krinski et al. 2021], and one of their conclusions was that the models' generalization was impaired by the small number of samples on the field's datasets which also suffer from class imbalance introducing some bias to the models. Data augmentation can mitigate this issue; however, the influence of data augmentation during training was left out.

In this work, we propose an extensive analysis of how different data augmentation techniques improve the training of encoder-decoder neural networks on this problem. Twenty different data augmentation techniques were evaluated, see section 4, in three distinct experiments using five CTs datasets: MedSeg [MedSeg 2021], Zenodo [Jun et al. 2020], CC-CCII [Zhang et al. 2020], MosMed [Morozov et al. 2020], Ricord1a [Tsai et al. 2020]. Each dataset was validated through a five-fold cross-validation strategy, thus resulting in over 3,000 experiments. The code for running these same experiments is publicly available¹.

2. Related Work

Data augmentation aims to generate a synthetic image by applying different operations to a preexisting labeled image [Ruiz et al. 2020b]. The most common operation are variations of an affine transformation such as flip, translate, rotate, scale. The Random Erasing [Zhong et al. 2020] is one example of data augmentation that adds information to the original image. In this technique, a rectangle with random values is positioned in the image. This rectangle's height, width, and center points are random values. This data augmentation helps the network learning process be more robust to object occlusions. Also, it reduces the overfitting in the training step. The CutMix [Yun et al. 2019] follows the same idea of the Random Erasing. However, instead of using constant values or even random ones, the technique mixes two images by adding an image A to some region of image B. This reduces information loss and encourages generalization.

Following the same line of CutMix [Yun et al. 2019], the study presented by [Summers and Dinneen 2019] evaluated several different non-linear mixing algorithms. The authors showed that non-linear mixing algorithms are also effective as linear mixing data augmentations. [Hendrycks et al. 2020] proposed a mixing data augmentation called AugMix. In AugMix, sequences of data augmentations are applied in parallel,

¹<https://github.com/VRI-UFPR/SparkInTheDarkLars2021>

generating a different image for each data augmentation sequence. In the end, an element-wise convex combination is applied to mix all generated images. In [Kisantal et al. 2019], the authors proposed a data augmentation technique for small objects to improve such objects' detection and segmentation. The authors applied "copy and pasting" strategy to create several copies of the objects of interest. They showed that this strategy increases the number of anchor boxes generated by the Mask-RCNN [He et al. 2017], which helps the network to learn and detect small objects. The ANDA [Ruiz et al. 2019] and IDA [Ruiz et al. 2020a] techniques follow the idea of introducing new objects, however since those are techniques focused on the generic problem of Saliency Object Detection (SOD), some additional operations are necessary such as Image Inpainting to erase the original object and some additional computation to choose which combination of background and object produce a significant saliency and the affine transformations to be applied to the new object that will replace the original one.

In the Grid Mask [Chen et al. 2020b], a mask composed of black squares uniformly distributed is generated and applied on top of the image. This augmentation follows the same idea of Random Erasing, which forces the network to learn from occluded regions in the image and reduces overfitting. The advantage over others that randomly remove regions from the image is that random algorithms can remove relevant regions from the image. The InstaBoost [Fang et al. 2019] uses an inpainting technique [Bertalmio et al. 2001] to remove the interest object from the image and place it in another region in the image. An appearance consistency heatmap [Field et al. 1993] is used to estimate the new region of the image where the object will be placed.

The authors of [Liu et al. 2019] proposed a data augmentation based on image-to-image divided into two steps: training and deployment. The proposed method uses several images of different classes in the training step, called source images, and learns to translate the images between these classes. Then, in the deployment step, a small set of images from the target class is used, with the proposed method being able to translate from the source classes to the target class. For each image to be mixed in the Super-Mix [Dabouei et al. 2020], a mixing binary mask with the saliency information of the respective image is generated. Then, a teacher model already trained in the problem is applied to mix the images, optimize the position of the salient region in the mixed image, and ensure that both salient regions are presented in the final mixed image.

These Data Augmentations (DAs) are, in general, for generic segmentation problems. However, there is no proper comparison of DA methods for the COVID-19 segmentation problem. In this work, we focus on performing an extensive comparison of generic DA methods in the approached problem.

3. Metrics and Datasets

The models were trained and evaluated across five different CTs datasets: MedSeg [MedSeg 2021], Zenodo [Jun et al. 2020], CC-CCII [Zhang et al. 2020], MosMed [Morozov et al. 2020], and Ricord1a [Tsai et al. 2020]. The MedSeg has 929 images and labels for four classes, with the following pixel proportion: Background (0.98563), Ground Glass Opacity (GGO) (0.01072), Consolidation (0.00351), and Pleural Effusion (0.0001). The Zenodo dataset has 3,520 images and labels for four classes, with the following pixel proportion: Background (0.89893), Left Lung (0.04331), Right Lung

(0.04923), and Infections (0.00852). The MosMed dataset is composed of 2,049 images, with labels for two classes, with the following pixel proportion: Background (0.99810) and GGO-Consolidation (0.00189). Ricord dataset is divided into three sets: 1a, 1b, and 1c. The set 1a is the only one with segmentation masks and has 9,166 images with labels for two classes, with the following pixel proportion: Background (0.95295) and Infections (0.04704). We also used a sub-set of CC-CCII with segmentation masks composed of 750 images and has labels for four classes, with the following pixel proportion: Background (0.87152), Lung Field (0.11691), GGO (0.00802), and Consolidation (0.00353).

One of the problems pointed out in [Krinski et al. 2021] was the class imbalance due to several images with just the background class; in fact, recent work has shown that several problems suffer from class imbalance [Laroca et al. 2021, Laroca et al. 2022]. Therefore, to mitigate this problem, in the first step of this evaluation, we removed from the datasets images with no lesion in the ground-truth mask. In the CC-CCII dataset, 4 images were removed; in the MedSeg dataset, 457 images were removed; in the MosMed, 1,264 images were removed; in the Zenodo dataset, 547 were removed; and in the Ricord1a, no image was removed. The datasets were split into 80% for training and 20% for testing. Then, a five-fold cross-validation strategy further divided the training set between training and validating sets. The metrics used for evaluation were the F-score described by equation 1 and Intersection over Union (IoU) described by equation 2.

$$F - score = \frac{TruePositive}{TruePositive + \frac{FalsePositive + FalseNegative}{2}} \quad (1)$$

$$IoU = \frac{intersection}{union} \quad (2)$$

4. Experiments

In general, when dealing with the COVID-19 CT segmentation problem it is usual to either completely neglect a dedicated evaluation of the impact of DA techniques or merely report using a limited set of generic DA, not optimized or specially designed for medical images, like Flip and Rotation operations, such as in [Zhao et al. 2021, Qiblawey et al. 2021, Raj et al. 2021, Müller et al. 2020, Chen et al. 2020c, Xu et al. 2020]. To properly measure the impact of data augmentation on the COVID-19 CT segmentation problem, we evaluated twenty data augmentation techniques.

In this work we evaluated the following twenty data augmentation techniques: CLAHE, Coarse Dropout, Elastic Transform, Emboss, Flip, Gaussian Blur, Grid Distortion, Grid Dropout, Image Compression, Median Blur, Optical Distortion, Piecewise Affine Transformation, Posterize, RBC, Random Crop, Random Gamma, Random Snow, Rotate, Sharpen, Shift Scale Rotate. Figure 1 illustrates the twenty data augmentation techniques applied to a CT image 1(a). All data augmentation techniques evaluated here are available in the Albumentation library [Buslaev et al. 2020].

The encoder-decoder network chosen to evaluate the dataset augmentations was the RegNetx-002 [Xu et al. 2021] encoder and U-net++ [Zhou et al. 2018] decoder. Since the encoders achieved close results in the comparison performed in [Krinski et al. 2021],




Figure 1. Illustration of twenty data augmentation techniques applied in a CT image. In 1(a), the original image is presented. The data augmentations are illustrated in the following sequence: Contrast Limited Adaptive Histogram Equalization (CLAHE) 1(b), Coarse Dropout 1(c), Elastic Transform 1(d), Emboss 1(e), Flip 1(f), Gaussian Blur 1(g), Grid Distortion 1(h), Grid Dropout 1(i), Image Compression 1(j), Median Blur 1(k), Optical Distortion 1(l), Piecewise Affine Transformation 1(m), Posterize 1(n), Random Brightness Contrast (RBC) 1(o), Random Crop 1(p), Random Gamma 1(q), Random Snow 1(r), Rotate 1(s), Sharpen 1(t), Shift Scale Rotate 1(u).

the RegNetx-002 was chosen due to being the network with a smaller number of parameters, making the RegNetx-002 faster for training. The U-net++ was chosen because it achieved the highest F-score compared with other decoders [Krinski et al. 2021]. The evaluation of how data augmentation affects the results of different encoders and decoders was left for future evaluation. Also, all experiments were evaluated through a five-fold cross-validation strategy. In total, we performed three sets of data augmentation evaluation, with each set varying network training parameters. The varied parameters were the number of epochs trained, the learning rate, and learning rate decay.

In the first evaluation of data augmentations, the architecture was trained for 50

Table 1. Results of the first data augmentation evaluation in the test set. The blue-colored values indicate the best F-scores, and the red-colored values indicate the best IoUs values. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. The green highlighted values show the data augmentations in which the P-value achieved values lower than 0.05, and the null hypothesis was rejected.

Probability	Augmentation	CC-CCII		MedSeg		MosMed		Ricord1a		Zenodo	
		F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU
	No Augmentation	0.6419	0.6011	0.4763	0.4289	0.7859	0.7211	0.8956	0.8372	0.8497	0.8152
0.1	CLAHE	0.6363	0.5962	0.4754	0.4278	0.7927	0.7268	0.8933	0.8340	0.8490	0.8142
	Coarse Dropout	0.6410	0.6004	0.4756	0.4281	0.7904	0.7244	0.8941	0.8350	0.8488	0.8140
	Elastic Transform	0.6454	0.6048	0.4781	0.4305	0.7977	0.7311	0.8915	0.8322	0.8496	0.8149
	Emboss	0.6409	0.6007	0.4749	0.4275	0.7922	0.7265	0.8960	0.8374	0.8490	0.8143
	Flip	0.6458	0.6044	0.4787	0.4309	0.7957	0.7296	0.8858	0.8256	0.8468	0.8110
	Gaussian Blur	0.6419	0.6011	0.4758	0.4283	0.7892	0.7238	0.8942	0.8353	0.8494	0.8148
	Grid Distortion	0.6443	0.6036	0.4764	0.4291	0.7967	0.7308	0.8911	0.8314	0.8502	0.8156
	Grid Dropout	0.6409	0.5999	0.4748	0.4272	0.7916	0.7265	0.8954	0.8367	0.8493	0.8146
	Image Compression	0.6427	0.6019	0.4769	0.4295	0.7936	0.7276	0.8950	0.8361	0.8499	0.8155
	Median Blur	0.6423	0.6015	0.4764	0.4290	0.7851	0.7201	0.8955	0.8371	0.8496	0.8151
	Optical Distortion	0.6439	0.6032	0.4768	0.4292	0.7932	0.7269	0.8922	0.8330	0.8500	0.8155
	Piecewise Affine	0.6430	0.6025	0.4756	0.4282	0.8000	0.7335	0.8909	0.8320	0.8499	0.8154
	Posterize	0.6446	0.6039	0.4763	0.4289	0.7954	0.7285	0.8931	0.8338	0.8492	0.8147
	RBC	0.6452	0.6039	0.4780	0.4303	0.7859	0.7202	0.8930	0.8336	0.8493	0.8146
	Random Crop	0.6425	0.6020	0.4757	0.4285	0.7863	0.7218	0.8952	0.8367	0.8497	0.8151
	Random Gamma	0.6421	0.6017	0.4759	0.4280	0.7907	0.7252	0.8933	0.8345	0.8496	0.8149
	Random Snow	0.6415	0.6010	0.4737	0.4263	0.7887	0.7231	0.8936	0.8346	0.8492	0.8143
	Rotate	0.6473	0.6064	0.4775	0.4299	0.7998	0.7335	0.8893	0.8297	0.8475	0.8118
Sharpen	0.6426	0.6024	0.4760	0.4285	0.7885	0.7225	0.8941	0.8350	0.8494	0.8149	
Shift Scale Rotate	0.6452	0.6043	0.4785	0.4307	0.8001	0.7336	0.8924	0.8330	0.8494	0.8146	
0.2	CLAHE	0.6386	0.5978	0.4723	0.4252	0.7844	0.7192	0.8955	0.8367	0.8491	0.8142
	Coarse Dropout	0.6379	0.5976	0.4744	0.4270	0.7931	0.7273	0.8920	0.8329	0.8489	0.8141
	Elastic Transform	0.6454	0.6050	0.4785	0.4303	0.8015	0.7345	0.8839	0.8229	0.8491	0.8141
	Emboss	0.6420	0.6015	0.4755	0.4280	0.7927	0.7270	0.8942	0.8355	0.8487	0.8137
	Flip	0.6465	0.6054	0.4771	0.4292	0.7991	0.7328	0.8889	0.8289	0.8444	0.8080
	Gaussian Blur	0.6421	0.6015	0.4755	0.4288	0.7876	0.7230	0.8928	0.8338	0.8494	0.8146
	Grid Distortion	0.6476	0.6070	0.4802	0.4325	0.8038	0.7368	0.8886	0.8288	0.8496	0.8148
	Grid Dropout	0.6398	0.5990	0.4740	0.4264	0.7868	0.7209	0.8855	0.8250	0.8487	0.8140
	Image Compression	0.6436	0.6022	0.4776	0.4299	0.7965	0.7310	0.8963	0.8378	0.8499	0.8155
	Median Blur	0.6431	0.6025	0.4762	0.4289	0.7888	0.7231	0.8951	0.8364	0.8497	0.8150
	Optical Distortion	0.6455	0.6045	0.4781	0.4300	0.7941	0.7279	0.8938	0.8352	0.8503	0.8159
	Piecewise Affine	0.6480	0.6072	0.4771	0.4297	0.8049	0.7385	0.8899	0.8304	0.8495	0.8147
	Posterize	0.6417	0.6014	0.4772	0.4299	0.7967	0.7317	0.8923	0.8330	0.8497	0.8151
	RBC	0.6419	0.6009	0.4761	0.4280	0.7875	0.7212	0.8867	0.8266	0.8493	0.8145
	Random Crop	0.6456	0.6044	0.4752	0.4279	0.7943	0.7287	0.8953	0.8368	0.8497	0.8150
	Random Gamma	0.6415	0.6006	0.4761	0.4279	0.7935	0.7274	0.8916	0.8323	0.8493	0.8146
	Random Snow	0.6422	0.6015	0.4711	0.4237	0.7875	0.7215	0.8883	0.8284	0.8484	0.8133
	Rotate	0.6458	0.6050	0.4781	0.4302	0.8010	0.7340	0.8813	0.8203	0.8473	0.8114
Sharpen	0.6418	0.6009	0.4770	0.4289	0.7852	0.7189	0.8954	0.8370	0.8495	0.8148	
Shift Scale Rotate	0.6482	0.6070	0.4806	0.4326	0.8004	0.7340	0.8856	0.8253	0.8491	0.8141	

epochs with a learning rate of 0.001. The learning rate was divided by 10 every 10 epochs. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. As presented in Table 1, most of the data augmentations did not improve the F-score and the IoU. The MosMed dataset was the only dataset that applying data augmentations improved the results, with improvements in the F-score of 1% and 2% in most of the data augmentations applied. In the Zenodo dataset, most of the augmentations achieved similar results with the baseline. However, the Grid Distortion with the probability of 0.1 and the Optical Distortion with the probability of 0.2 improved the F-score by 1%. In the MedSeg, only the Shift Scale Rotate augmentation with the probability of 0.2 achieved better results with an F-score 1% higher than the baseline. The CC-CCII and Ricord1a datasets did not achieve improvements with data augmentations. To perform a statistical analysis of

the data augmentation evaluation, the one-sided Wilcoxon signed-rank test was applied with the null hypothesis as the F-scores of the distribution without data augmentation are greater than the distributions with data augmentation.

The datasets CC-CCII and MosMed were the datasets most sensitive to data augmentation and achieved better F-scores in seven data augmentations when the data augmentation was applied with probability 0.1 and eight data augmentations when applied with probability 0.2. In the MosMed dataset, the null hypothesis was rejected in most of the data augmentations applied in both probabilities. In the MedSeg, the null hypotheses were rejected in only one data augmentation when applied with probability 0.1 and two data augmentations when applied with probability 0.2. The Ricord1a and Zenodo did not achieve statistical significance to reject the null hypotheses in any data augmentation.

Table 2. Results of the second data augmentation evaluation in the test set. The blue-colored values indicate the best F-scores, and the red-colored values indicate the best IoUs values. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. The green highlighted values show the data augmentations in which the P-value achieved values lower than 0.05, and the null hypothesis was rejected.

Probability	Augmentation	CC-CCII		MedSeg		MosMed		Ricord1a		Zenodo	
		F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU
	No Augmentation	0.6447	0.6048	0.4780	0.4308	0.7938	0.7292	0.9031	0.8472	0.8511	0.8174
0.1	CLAHE	0.6431	0.6025	0.4781	0.4305	0.7939	0.7288	0.9033	0.8476	0.8512	0.8174
	Coarse Dropout	0.6429	0.6016	0.4769	0.4304	0.7985	0.7335	0.9027	0.8464	0.8510	0.8174
	Elastic Transform	0.6499	0.6095	0.4817	0.4342	0.8065	0.7404	0.8993	0.8421	0.8514	0.8176
	Emboss	0.6458	0.6054	0.4777	0.4308	0.7953	0.7296	0.9039	0.8482	0.8509	0.8171
	Flip	0.6493	0.6087	0.4812	0.4344	0.8115	0.7451	0.8990	0.8417	0.8497	0.8149
	Gaussian Blur	0.6448	0.6045	0.4799	0.4329	0.7917	0.7265	0.9035	0.8473	0.8511	0.8173
	Grid Distortion	0.6506	0.6102	0.4814	0.4341	0.8121	0.7459	0.9022	0.8452	0.8523	0.8188
	Grid Dropout	0.6412	0.6011	0.4802	0.4324	0.7927	0.7279	0.9019	0.8453	0.8511	0.8174
	Image Compression	0.6462	0.6052	0.4787	0.4315	0.7961	0.7318	0.9033	0.8472	0.8508	0.8170
	Median Blur	0.6462	0.6060	0.4765	0.4299	0.7959	0.7314	0.9023	0.8458	0.8516	0.8179
	Optical Distortion	0.6487	0.6079	0.4819	0.4341	0.8050	0.7391	0.9035	0.8473	0.8516	0.8181
	Piecewise Affine	0.6495	0.6092	0.4819	0.4349	0.8018	0.7370	0.9022	0.8454	0.8522	0.8186
	Posterize	0.6469	0.6070	0.4796	0.4326	0.7985	0.7329	0.9021	0.8457	0.8515	0.8178
	RBC	0.6439	0.6037	0.4805	0.4326	0.7940	0.7276	0.9025	0.8460	0.8512	0.8174
	Random Crop	0.6470	0.6057	0.4772	0.4303	0.7915	0.7264	0.9009	0.8441	0.8511	0.8173
	Random Gamma	0.6449	0.6043	0.4783	0.4305	0.7949	0.7300	0.9020	0.8455	0.8510	0.8171
	Random Snow	0.6435	0.6037	0.4770	0.4301	0.7958	0.7296	0.9027	0.8463	0.8511	0.8173
	Rotate	0.6522	0.6115	0.4810	0.4341	0.8073	0.7416	0.9007	0.8439	0.8501	0.8158
	Sharpen	0.6441	0.6036	0.4778	0.4306	0.7915	0.7269	0.9026	0.8467	0.8512	0.8177
	Shift Scale Rotate	0.6504	0.6096	0.4823	0.4350	0.8060	0.7397	0.9026	0.8460	0.8518	0.8181
0.2	CLAHE	0.6408	0.6005	0.4764	0.4293	0.7919	0.7273	0.9025	0.8461	0.8512	0.8173
	Coarse Dropout	0.6420	0.6014	0.4776	0.4307	0.7952	0.7303	0.9033	0.8471	0.8512	0.8176
	Elastic Transform	0.6526	0.6121	0.4840	0.4366	0.8137	0.7469	0.9012	0.8439	0.8517	0.8178
	Emboss	0.6430	0.6026	0.4790	0.4316	0.7990	0.7326	0.9043	0.8484	0.8510	0.8170
	Flip	0.6501	0.6091	0.4813	0.4341	0.8066	0.7408	0.8977	0.8397	0.8486	0.8135
	Gaussian Blur	0.6431	0.6026	0.4775	0.4300	0.7920	0.7276	0.9021	0.8457	0.8509	0.8171
	Grid Distortion	0.6543	0.6139	0.4835	0.4365	0.8102	0.7441	0.9007	0.8432	0.8524	0.8189
	Grid Dropout	0.6408	0.6006	0.4778	0.4308	0.7950	0.7298	0.9019	0.8452	0.8507	0.8167
	Image Compression	0.6460	0.6057	0.4782	0.4308	0.7962	0.7310	0.9048	0.8494	0.8512	0.8175
	Median Blur	0.6467	0.6063	0.4769	0.4302	0.7964	0.7305	0.9022	0.8455	0.8510	0.8173
	Optical Distortion	0.6492	0.6085	0.4821	0.4346	0.8057	0.7396	0.9030	0.8466	0.8518	0.8183
	Piecewise Affine	0.6482	0.6079	0.4815	0.4343	0.8106	0.7444	0.9024	0.8454	0.8516	0.8179
	Posterize	0.6472	0.6063	0.4793	0.4318	0.7930	0.7280	0.9027	0.8463	0.8512	0.8174
	RBC	0.6415	0.6008	0.4791	0.4312	0.7995	0.7333	0.9019	0.8453	0.8512	0.8174
	Random Crop	0.6449	0.6049	0.4771	0.4299	0.7963	0.7309	0.9034	0.8474	0.8513	0.8175
	Random Gamma	0.6459	0.6051	0.4772	0.4301	0.7929	0.7274	0.9025	0.8462	0.8512	0.8173
	Random Snow	0.6416	0.6018	0.4752	0.4280	0.7961	0.7291	0.9028	0.8467	0.8507	0.8166
	Rotate	0.6526	0.6121	0.4831	0.4357	0.8119	0.7456	0.8945	0.8362	0.8498	0.8153
	Sharpen	0.6466	0.6063	0.4778	0.4310	0.7867	0.7209	0.9035	0.8474	0.8515	0.8180
	Shift Scale Rotate	0.6534	0.6127	0.4843	0.4369	0.8135	0.7462	0.8964	0.8385	0.8517	0.8180

In the second evaluation of data augmentations, the architecture was trained for 100 epochs with a learning rate of 0.001. The learning rate was divided by 10 every 20 epochs. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. As presented in Table 2, the MosMed achieved the most significant improvements, with the Grid Distortion with probability 0.1 and the Elastic Transform with probability 0.2 increasing the F-score by 2% compared with the baseline. However, unlike the first evaluation, the MosMeg achieved better F-scores in only seven augmentations instead of fourteen augmentations.

Table 3. Results of the third data augmentation evaluation in the test set. The blue-colored values indicate the best F-scores, and the red-colored values indicate the best IoUs values. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. The green highlighted values show the data augmentations in which the P-value achieved values lower than 0.05, and the null hypothesis was rejected.

Probability	Augmentation	CC-CCII		MedSeg		MosMed		Ricord1a		Zenodo	
		F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU	F-score	IoU
	No Augmentation	0.6317	0.5893	0.4658	0.4183	0.7736	0.7089	0.9133	0.8610	0.8496	0.8153
0.1	CLAHE	0.6288	0.5873	0.4618	0.4149	0.7698	0.7068	0.9140	0.8617	0.8498	0.8155
	Coarse Dropout	0.6275	0.5863	0.4647	0.4179	0.7831	0.7178	0.9142	0.8620	0.8496	0.8152
	Elastic Transform	0.6330	0.5912	0.4679	0.4208	0.7907	0.7249	0.9136	0.8610	0.8505	0.8163
	Emboss	0.6282	0.5863	0.4654	0.4182	0.7758	0.7108	0.9139	0.8618	0.8499	0.8157
	Flip	0.6348	0.5922	0.4686	0.4214	0.7933	0.7278	0.9130	0.8603	0.8493	0.8147
	Gaussian Blur	0.6310	0.5890	0.4675	0.4202	0.7751	0.7110	0.9145	0.8625	0.8498	0.8156
	Grid Distortion	0.6328	0.5911	0.4704	0.4229	0.7907	0.7241	0.9144	0.8621	0.8509	0.8169
	Grid Dropout	0.6321	0.5900	0.4637	0.4171	0.7809	0.7158	0.9139	0.8616	0.8495	0.8152
	Image Compression	0.6296	0.5876	0.4653	0.4182	0.7725	0.7086	0.9142	0.8620	0.8477	0.8144
	Median Blur	0.6311	0.5886	0.4659	0.4193	0.7720	0.7088	0.9141	0.8620	0.8504	0.8165
	Optical Distortion	0.6328	0.5911	0.4696	0.4219	0.7862	0.7205	0.9141	0.8620	0.8496	0.8160
	Piecewise Affine	0.6339	0.5926	0.4675	0.4206	0.7923	0.7266	0.9147	0.8625	0.8504	0.8163
	Posterize	0.6308	0.5889	0.4645	0.4179	0.7762	0.7113	0.9141	0.8621	0.8502	0.8159
	RBC	0.6275	0.5858	0.4628	0.4159	0.7674	0.7029	0.9131	0.8607	0.8497	0.8154
	Random Crop	0.6295	0.5874	0.4626	0.4156	0.7764	0.7112	0.9128	0.8605	0.8498	0.8156
	Random Gamma	0.6301	0.5883	0.4627	0.4159	0.7744	0.7099	0.9139	0.8616	0.8494	0.8150
	Random Snow	0.6278	0.5863	0.4610	0.4146	0.7675	0.7032	0.9141	0.8620	0.8496	0.8151
	Rotate	0.6350	0.5929	0.4719	0.4243	0.7922	0.7268	0.9135	0.8609	0.8490	0.8143
	Sharpen	0.6322	0.5895	0.4650	0.4177	0.7740	0.7099	0.9146	0.8626	0.8499	0.8157
	Shift Scale Rotate	0.6381	0.5961	0.4710	0.4236	0.7962	0.7298	0.9139	0.8614	0.8506	0.8165
0.2	CLAHE	0.6272	0.5854	0.4649	0.4177	0.7663	0.7027	0.9140	0.8618	0.8497	0.8154
	Coarse Dropout	0.6286	0.5870	0.4625	0.4160	0.7822	0.7174	0.9137	0.8614	0.8495	0.8152
	Elastic Transform	0.6373	0.5956	0.4714	0.4235	0.7925	0.7268	0.9122	0.8589	0.8508	0.8166
	Emboss	0.6288	0.5870	0.4639	0.4171	0.7745	0.7101	0.9142	0.8622	0.8495	0.8151
	Flip	0.6365	0.5944	0.4696	0.4226	0.7935	0.7283	0.9127	0.8597	0.8487	0.8138
	Gaussian Blur	0.6297	0.5880	0.4659	0.4188	0.7755	0.7120	0.9140	0.8619	0.8501	0.8160
	Grid Distortion	0.6370	0.5951	0.4725	0.4249	0.7992	0.7327	0.9140	0.8612	0.8512	0.8172
	Grid Dropout	0.6304	0.5887	0.4649	0.4177	0.7795	0.7150	0.9133	0.8608	0.8492	0.8148
	Image Compression	0.6312	0.5888	0.4661	0.4186	0.7754	0.7118	0.9144	0.8625	0.8498	0.8155
	Median Blur	0.6297	0.5882	0.4681	0.4212	0.7779	0.7145	0.9144	0.8624	0.8501	0.8159
	Optical Distortion	0.6335	0.5914	0.4691	0.4217	0.7855	0.7205	0.9147	0.8628	0.8508	0.8168
	Piecewise Affine	0.6344	0.5927	0.4726	0.4253	0.7940	0.7280	0.9145	0.8621	0.8509	0.8168
	Posterize	0.6303	0.5886	0.4647	0.4177	0.7733	0.7098	0.9144	0.8625	0.8497	0.8156
	RBC	0.6273	0.5860	0.4668	0.4193	0.7748	0.7096	0.9131	0.8603	0.8498	0.8155
	Random Crop	0.6315	0.5894	0.4625	0.4157	0.7694	0.7049	0.9137	0.8615	0.8497	0.8154
	Random Gamma	0.6313	0.5893	0.4629	0.4155	0.7722	0.7074	0.9132	0.8610	0.8496	0.8152
	Random Snow	0.6258	0.5841	0.4599	0.4131	0.7709	0.7067	0.9145	0.8622	0.8484	0.8142
	Rotate	0.6368	0.5948	0.4739	0.4259	0.8004	0.7338	0.9127	0.8596	0.8495	0.8148
	Sharpen	0.6323	0.5901	0.4668	0.4196	0.7750	0.7119	0.9147	0.8628	0.8495	0.8154
	Shift Scale Rotate	0.6380	0.5960	0.4712	0.4243	0.7967	0.7308	0.9134	0.8602	0.8511	0.8171

The CC-CCII achieved better F-scores with the Grid Distortion, Rotate, and Shift Scale Rotate. Also, MedSeg had better F-scores in nine augmentations with the probability of 0.1 and seven augmentations with the probability of 0.2. Data augmentations

did not improve the F-score in the Zenodo and Ricord1a datasets. The statistical analysis pointed out that the CC-CCII rejected the null hypotheses in seven data augmentations in both probabilities. Training for 100 epochs made the MedSeg achieve better F-scores in eleven data augmentations with probability 0.1 and seven data augmentations with probability 0.2. The results achieved in the MosMed dataset got worse when compared with the results presented in Table 1, with only seven data augmentations rejecting the null hypotheses. In the Ricord1a and Zenodo datasets, besides the average F-score of the data augmentations being very close to the average F-score without data augmentation, in the Ricord1a dataset, the null hypothesis was rejected in two data augmentations with probability 0.2, and in the Zenodo dataset, the null hypotheses were rejected in three data augmentations with probability 0.1 and two data augmentations with probability 0.2.

In the third evaluation, the architecture was trained for 100 epochs with a learning rate of 0.0001. The learning rate was divided by 10 every 25 epochs. Two probabilities of applying the data augmentation were evaluated: 0.1 and 0.2. As presented in Table 3, in the MosMed, the Shift Scale Rotate with the probability of 0.1 and the Rotate with the probability of 0.2 increased the F-score by 2% compared with the baseline. Also, the MosMed achieved the best F-scores in nine augmentations with probability 0.1 and eight augmentations with the probability of 0.2, pointing this training configuration with the highest effectiveness for this dataset. Also, this training configuration showed significant effectiveness in the Zenodo dataset, which achieved the best F-scores in six augmentations with probability 0.1 and seven augmentations with probability 0.2.

However, this training configuration did not significantly affect other datasets. The MedSeg achieved higher F-scores in only three augmentations with a probability of 0.1 and five augmentations with a probability of 0.2. Also, the CC-CCII and Ricord1a datasets did not achieve improvements with any data augmentation. Statistical analysis also was performed and, besides the average F-score of the data augmentations being very close to the baseline without data augmentation, the Zenodo and Ricord1a datasets achieved the highest number of data augmentations with the null hypotheses rejected. The MosMed was the dataset with the most promising results, with the null hypotheses rejected in twelve data augmentations. The CC-CCII and MedSeg had fewer data augmentations with the null hypotheses rejected compared with previous experiments, showing that this training configuration is not proper for data augmentations in these datasets.

5. Conclusion

These three experiments demonstrated that, although necessary, the generic data augmentation techniques evaluated did not majorly improve the results in the COVID-19 segmentation problem. The MosMed achieved the most significant improvements with data augmentation, with the F-score being up to 2% higher in comparison with baseline. This dataset was the most sensitive to data augmentation techniques due to its imbalance problem. Data augmentations also improved the CC-CCII and MedSeg dataset results, but it was necessary to train the network for more epochs, and the data augmentations only achieved 1% of improvements in the F-score. The Ricord1a and Zenodo datasets were the most challenging and did not show improvements with data augmentations. Although they achieved statistical significance to reject the null hypotheses in many data augmentations, the average F-score slightly improved.

Another result of these data augmentation experiments is that spatial level transformations such as Elastic Transform, Flip, Grid Distortion, Piecewise Affine, Rotate, and Shift Scale Rotate were the operations that showed to be the most favorable data augmentations to this problem. These data augmentations improved the results in most of the experiments performed and thus are the most promising techniques for future experiments with data augmentation. The evaluation of more domain-specific data augmentation techniques was left for future works.

Acknowledgment

The authors would like to thank the Coordination for the Improvement of Higher Education Personnel (CAPES) for the PhD scholarship. We gratefully acknowledge the founders of the publicly available datasets, the support of NVIDIA Corporation with the donation of the GPUs used for this research and the C3SL-UFPR group for the computational cluster infrastructure.

References

- Bertalmio, M., Bertozzi, A. L., and Sapiro, G. (2001). Navier-stokes, fluid dynamics, and image and video inpainting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I.
- Buslaev, A., Iglovikov, V. I., Khvedchenya, E., Parinov, A., Druzhinin, M., and Kalinin, A. A. (2020). Albuementations: Fast and flexible image augmentations. *Information*, 11(2).
- Cao, F. and Bao, Q. (2020). A survey on image semantic segmentation methods with convolutional neural network. In *2020 International Conference on Communications, Information System and Computer Engineering (CISCE)*. IEEE.
- Chen, M., Tu, C., Tan, C., Zheng, X., Wang, X., Wu, J., Huang, Y., Wang, Z., Yan, Y., Li, Z., Shan, H., Liu, J., and Huang, J. (2020a). Key to successful treatment of covid-19: accurate identification of severe risks and early intervention of disease progression. *medRxiv*, DOI:10.1101/2020.04.06.20054890.
- Chen, P., Liu, S., Zhao, H., and Jia, J. (2020b). Gridmask data augmentation. arXiv preprint, arXiv:2001.04086.
- Chen, X., Yao, L., and Zhang, Y. (2020c). Residual attention u-net for automated multi-class segmentation of covid-19 chest ct images. arXiv preprint, arXiv:2004.05645.
- Dabouei, A., Soleymani, S., Taherkhani, F., and Nasrabadi, N. M. (2020). Supermix: Supervising the mixing data augmentation. arXiv preprint, arXiv:2003.05034.
- Fang, H.-S., Sun, J., Wang, R., Gou, M., Li, Y.-L., and Lu, C. (2019). Instaboost: Boosting instance segmentation via probability map guided copy-pasting. arXiv preprint, arXiv:1908.07801.
- Field, D. J., Hayes, A., and Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Research*, 33(2):173–193.
- He, K., Gkioxari, G., Dollar, P., and Girshick, R. (2017). Mask r-CNN. In *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE.

- Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer, J., and Lakshminarayanan, B. (2020). AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Jun, M. et al. (2020). Covid-19 ct lung and infection segmentation dataset. Zenodo. Available at: <https://zenodo.org/record/3757475>. Accessed: 2022-02-22.
- Kisantal, M., Wojna, Z., Murawski, J., Naruniec, J., and Cho, K. (2019). Augmentation for small object detection. In *9th International Conference on Advances in Computing and Information Technology (ACITY 2019)*. Airc Publishing Corporation.
- Krinski, B. A., Ruiz, D. V., and Todt, E. (2021). Spark in the dark: Evaluating encoder-decoder pairs for COVID-19 CT's semantic segmentation. In *2021 Latin American Robotics Symposium (LARS)*. IEEE.
- Laroca, R., Araujo, A. B., Zanlorensi, L. A., De Almeida, E. C., and Menotti, D. (2021). Towards image-based automatic meter reading in unconstrained scenarios: A robust and efficient approach. *IEEE Access*, 9:67569–67584.
- Laroca, R., Cardoso, E. V., Lucio, D. R., Estevam, V., and Menotti, D. (2022). On the cross-dataset generalization in license plate recognition. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, pages 166–178.
- Liu, M.-Y., Huang, X., Mallya, A., Karras, T., Aila, T., Lehtinen, J., and Kautz, J. (2019). Few-shot unsupervised image-to-image translation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE.
- MedSeg (2021). Covid-19 ct segmentation dataset. Available at: <http://medicalsegmentation.com/covid19/>. Accessed: 2021-05-03.
- Morozov, S., Andreychenko, A., Pavlov, N., Vladzimirskyy, A., Ledikhova, N., Gombolevskiy, V., Blokhin, I., Gelezhe, P., Gonchar, A., and Chernina, V. (2020). Mosmeddata: Chest ct scans with covid-19 related findings dataset. *medRxiv*, DOI:10.1101/2020.05.20.20100362.
- Müller, D., Rey, I. S., and Kramer, F. (2020). Automated chest ct image segmentation of covid-19 lung infection based on 3d u-net. arXiv preprint, arXiv:2007.04774.
- of Medicine, J. H. U. (2022). Coronavirus resource center. Available at: <https://coronavirus.jhu.edu/>. Accessed: 2022-02-22.
- Qiblawey, Y., Tahir, A., Chowdhury, M. E. H., Khandakar, A., Kiranyaz, S., Rahman, T., Ibtihaz, N., Mahmud, S., Al-Madeed, S., and Musharavati, F. (2021). Detection and severity classification of covid-19 in ct images using deep learning. arXiv preprint, arXiv:2102.07726.
- Raj, A. N. J., Zhu, H., Khan, A., Zhuang, Z., Yang, Z., Mahesh, V. G. V., and Karthik, G. (2021). ADID-UNET—a segmentation model for COVID-19 infection from lung CT scans. *PeerJ Computer Science*, 7:e349.
- Ruiz, D. V., Krinski, B. A., and Todt, E. (2019). ANDA: A novel data augmentation technique applied to salient object detection. In *2019 19th International Conference on Advanced Robotics (ICAR)*, pages 487–492.

- Ruiz, D. V., Krinski, B. A., and Todt, E. (2020a). IDA: Improved data augmentation applied to salient object detection. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 210–217.
- Ruiz, D. V., Salomon, G., and Todt, E. (2020b). Can giraffes become birds? an evaluation of image-to-image translation for data generation. *Anais do Computer on the Beach*, 11(1):176–182. DOI: 10.14210/cotb.v11n1.p176-182, also available as arXiv preprint, arXiv:2001.03637.
- Shi, F., Wang, J., Shi, J., Wu, Z., Wang, Q., Tang, Z., He, K., Shi, Y., and Shen, D. (2021). Review of artificial intelligence techniques in imaging data acquisition, segmentation, and diagnosis for COVID-19. *IEEE Reviews in Biomedical Engineering*, 14:4–15.
- Summers, C. and Dinneen, M. J. (2019). Improved mixed-example data augmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE.
- Tsai, E. et al. (2020). Medical imaging data resource center - rsna international covid radiology database release 1a - chest ct covid+ (midrc-ricord-1a). DOI: 10.7937/VTW4-X588.
- Wang, C., Horby, P. W., Hayden, F. G., and Gao, G. F. (2020). A novel coronavirus outbreak of global health concern. *The Lancet*, 395(10223):470–473.
- Xu, J., Pan, Y., Pan, X., Hoi, S., Yi, Z., and Xu, Z. (2021). Regnet: Self-regulated network for image classification. arXiv preprint, arXiv:2101.00590.
- Xu, Z., Cao, Y., Jin, C., Shao, G., Liu, X., Zhou, J., Shi, H., and Feng, J. (2020). Gasnet: Weakly-supervised framework for covid-19 lesion segmentation. arXiv preprint, arXiv:2010.09456.
- Yun, S., Han, D., Chun, S., Oh, S. J., Yoo, Y., and Choe, J. (2019). CutMix: Regularization strategy to train strong classifiers with localizable features. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE.
- Zhang, K. et al. (2020). Clinically applicable AI system for accurate diagnosis, quantitative measurements, and prognosis of COVID-19 pneumonia using computed tomography. *Cell*, 181(6):1423–1433.e11.
- Zhao, X., Zhang, P., Song, F., Fan, G., Sun, Y., Wang, Y., Tian, Z., Zhang, L., and Zhang, G. (2021). D2a u-net: Automatic segmentation of covid-19 lesions from ct slices with dilated convolution and dual attention mechanism. arXiv preprint, arXiv:2102.05210.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. (2020). Random erasing data augmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):13001–13008.
- Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N., and Liang, J. (2018). UNet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11. Springer International Publishing.