



HAL
open science

SuperpixelGridMasks Data Augmentation: Application to Precision Health and Other Real-world Data

Karim Hammoudi, Adnane Cabani, Bouthaina Slika, Halim Benhabiles, Fadi Dornaika, Mahmoud Melkemi

► **To cite this version:**

Karim Hammoudi, Adnane Cabani, Bouthaina Slika, Halim Benhabiles, Fadi Dornaika, et al.. SuperpixelGridMasks Data Augmentation: Application to Precision Health and Other Real-world Data. Journal of Healthcare Informatics Research, 2023, 6, pp.442-460. 10.1007/s41666-022-00122-1 . hal-03951542

HAL Id: hal-03951542

<https://hal.science/hal-03951542v1>

Submitted on 24 Jan 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.

SuperpixelGridCut, SuperpixelGridMean and SuperpixelGridMix Data Augmentation

Karim Hammoudi^{1,2,*} Adnane Cabani^{3,*} Bouthaina Slika⁴ Halim Benhabiles⁵

Fadi Dornaika^{4,6} Mahmoud Melkemi^{1,2}

¹Université de Haute-Alsace, IRIMAS ²Université de Strasbourg

³Normandie Université, ESIGELEC-IRSEEM ⁴University of the Basque Country

⁵JUNIA, CNRS-IEMN, Université de Lille ⁶IKERBASQUE

karim.hammoudi@uha.fr

adnane.cabani@esigelec.fr

bslika001@ikasle.ehu.eus

halim.benhabiles@junia.com

fadi.dornaika@ehu.eus

mahmoud.melkemi@uha.fr




Figure 1. (b) and (c) show generated augmented images using the proposed *SuperpixelGridCut* and *SuperpixelGridMean* applied from the original image (a). (e) shows the generated augmented image using the *SuperpixelGridMix* over the original images (a) and (d).

Abstract

A novel approach of data augmentation based on irregular superpixel decomposition is proposed. This approach called *SuperpixelGridMasks* permits to extend original image datasets that are required by training stages of machine learning-related analysis architectures towards increasing their performances. Three variants named *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* are presented. These grid-based methods produce a new style of image transformations using the dropping and fusing of information. Extensive experiments using various image classification models and datasets show that baseline performances can be significantly outperformed using our methods. The comparative study also shows that our methods can overpass the performances of other data augmentations. Experimental results obtained over image recognition datasets of varied natures show the efficiency of these new methods. *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* codes are publicly available at <https://github.com/hammoudiproject/SuperpixelGridMasks>.

Keywords: classification, YOLO, CNN, *SuperpixelGridCut*, *SuperpixelGridMean*, *SuperpixelGridMix*, deep learning, XAI.

*Two authors contributed equally to this work (corresponding authors).

Dataset	Baseline (%)	Ours (%)
PASCAL VOC ¹	57.81	75.00 (+17.19)
Chest X-Ray Images ² (Kaggle)	78.68	88.14 (+9.46)

Table 1. Reported top-1 accuracies of our classification results using VGG19 over image recognition datasets of different natures.

1. Introduction

More than ever, machine learning is a field of great interest for the community of image recognition. To the past, a huge amount of approaches have been developed to perform semi-automatic and automatic image segmentation, image classification, object detection (e.g. YOLO) and so on. Nowadays, deep Convolutional Neural Networks (CNNs) appear as an ultimate resource towards operating recognition tasks over e.g. biomedical images, scene images, object images, document images.

In this context, the investigated topics are more and more refined and delicate to tackle. The CNN architectures often include millions of parameters, requiring an important number of images for performing their training, validation and test stages towards returning decisions with high precisions. In many cases, we face a lack of data due to the topics that are new and that require hand-made and fastidious data labeling tasks. An insufficient amount of

relevant data can lead to serious over-fitting problems. Data augmentation appears as a complementary technique for overcoming such problems.

Various types of data augmentation have been proposed in the literature. Frequent methods can be divided into three categories, i) the pixel-level transforms (e.g. colorimetry-based), ii) the spatial-level transforms (e.g. geometry-based), and iii) the information dropping and fusing (e.g., mask-based). For instance, pixel-level transforms can be obtained by applying conventional image filters such as blur noise, contrast or compression related ones. Spatial-level transforms can be to resize, rotate or flip images. The information dropping and fusing is the most recent category which includes transforms such as image masking or mixing. Each category contains transforms which can integrate parameters used with random values for exploration or optimization reasons.

The nature of the considered topic, associated images and applicative conditions often help in the choice of the data augmentation methods that have to be experimented. For instance, if the recognition has to be operated through a camera system, then augmenting data by simulating a camera noise via a blur (category 1) could be efficient; if the images to classify represent faces, an horizontal image flip (category 2) could be a good option and a vertical flip not; if the localization of objects in the image of a scene plays an important role, augmenting data by cropping and masking image parts (category 3) can permit to increase the perception field of the considered network by occluding some image parts and forcing then the attention on potentially relevant object parts.

Motivation — Information dropping techniques (category 3) can present the advantage of forcing a considered network to explore image parts of interest through this trade-off in between image occluding and preserving. Although often efficient and flexible, it is worth mentioning that most of the masking data augmentations methods, i.e. CutOut methods usually exploit the dropping of squares or rectangles. However, in a real-world topic, most of the processed image sets appear like natural images. The presence of structures that are horizontally or vertically aligned with respect to image orientations rarely occurs. Also, information fusion techniques often aim to merge in one image the image parts of two or more images. It can result from whole image encapsulation or mapping of image parts. The advantage can lie in the generation of new images which embed discriminative characteristics of original image pairs belonging to the same class. Once again, although effective and flexible, most of these mixing data augmentation methods, i.e CutMix methods usually exploit quadrilateral structures.

In this paper, we propose SuperpixelGridMasks, namely *SuperpixelGridCut*, *SuperpixelGridMean* and

SuperpixelGridMix data augmentation methods. By comparison to previously presented methods, they have several advantages, namely

- to respectively perform dropping and fusing information while limiting the integration in the considered network of periodic linear and right-angles related noises due to the use of regular structures,
- to generate a 2D irregular structure which is purely built using natural image information (e.g. pixel intensities, textures or contours),
- to potentially produce structures which delimit real boundaries in between objects and then perform dropping and fusing of delineated object parts instead of dropping and fusing of inconsistent quadrilateral image parts.

2. Related Work

Most of the conventional data augmentation methods and variants from the categories of pixel-level transforms and spatial-level transforms are regrouped in [1]. Data augmentation methods acting on the information dropping and fusing are more and more present. Those previously described such as CutOut, CutMix as well as many variants are presented in [2], [3] and [4–15], respectively. These latter, which are also mainly discussed in this survey [16], are all potentially efficient although they all exploit regular and quadrilateral structures.

To the best of our knowledge, except for direct segmentation purposes, very few approaches exploit superpixel entities towards performing data augmentation tasks. In [1, 17], the authors proposed methods that augment data by transforming input images into superpixel representations which are densely posterized. In [18], a superpixel representation is produced. Then, each image is split into a set of subimages (i.e. individual superpixels) towards preparing patch-based processes. A method *Superpixel-Mix* is presented in [19]. A superpixel representation is produced and a swap is operated in between some superpixel regions of image pairs. This CutMix is locally operated for augmenting data towards segmentation tasks. In sum, current superpixel-based data augmentation either applies raw image posterization or applies localized operations such as image-to-patch decomposition or parsimonious superpixel mixing.

In our case, superpixel grid-based masks are proposed for the first time to perform global operations by sparsely applying dropping or fusion operations over a superpixel grid with a uniform distribution. This principle is the core of the proposed SuperpixelGridMasks approach which is described in detail in the next section.




Figure 2. Parameters of SuperpixelGridMasks approach: segmentation type t , random seed-based positioning s , nb. of superpixels q , ratio of processed superpixels r

3. Proposed Approach

The SuperpixelGridMasks approach is presented through three variants; namely *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix*. Their basis principles are presented in functions 1, 2 and 3, respectively. Common initial stage consists of using a superpixel segmentation technique over an original image by providing the number q of requested superpixels. A variety of superpixel segmentation techniques can be used to perform this task (see [20]). In our case, we selected a refined version which maximizes the matching in between the generated superpixel grid lines and intrinsic contour lines of the original images. By this way, subsequent dropping and fusing processes can operate with coherence on superpixels since they directly share boundaries with object boundaries. The principle consists then of parsing the obtained superpixel list.

Parameters related to our SuperpixelGridMasks are displayed in Figure 2. A sample of superpixels is then processed (dropped or mixed) depending on a selected occupancy threshold r , $r \in [0, 1]$. This rate of processed superpixel occupancy is approximately equal to the rate of pixels to process over the augmented image when superpixel sizes are close. If r is equal to 0.5, then it means that we target the processing of half of superpixels. If r is equal to 0.3, then it means that we target the processing of 30% of the list of superpixels. To perform a uniform superpixel processing, we proceed by randomly drawing a variable p_i from a uniform distribution over $[0, 1]$ for each

Function 1: SuperpixelGridCut(I, q, r)

Input : I : original image; q : nb of Superpixels;
 r : rate of processed superpixel occupancy
Output: J : generated SuperpixelGridCut image

```

J ← I;
/* Superpixel segmentation */
List Superpixel ← Superpixel_segmentation(J, q);
/* Superpixel parsing & cutting */
for (i = 1; i ≤ Superpixel.size(); i++) do
    p = rand(0, 1);
    if (p < r) then
        Superpixel(i) ← Black
return J

```

Function 2: SuperpixelGridMean(I, q, r)

Input : I : original image; q : nb of Superpixels;
 r : rate of processed superpixel occupancy
Output: J : generated SuperpixelGridMean image

```

J ← I;
Integer AVG;
/* Superpixel segmentation */
List Superpixel ← Superpixel_segmentation(J, q);
/* Superpixel parsing & averaging */
for (i = 1; i ≤ Superpixel.size(); i++) do
    p = rand(0, 1);
    if (p < r) then
        AVG ← mean(Superpixel(i));
        Superpixel(i) ← AVG;
return J

```

superpixel i . The processing decision d is finally given by Eq.(1).

$$d = \begin{cases} \text{processed superpixel} & \text{if } p_i \in [0, r] \\ \text{unprocessed otherwise} & \end{cases} \quad (1)$$

if p is inferior to the selected r , then the considered superpixel will be processed. Otherwise, the superpixel will be untouched. This random pulling of p permits to process the superpixel in a sparse manner and with a uniform distribution. By this way, the threshold r also permits to control the contiguity level of processed superpixels.

From this step, function 1 and function 2 directly colorize the retained sample of superpixels to process in black or with average color values, respectively while the function 3 replaces the texture of the retained sample of superpixels to process by the one of another original image of the dataset. In this latter process, the basis version

Function 3: SuperpixelGridMix(I,J,q,r)

```
Input : I,J: original images; q: nb of Superpixels;
r: rate of processed superpixel occupancy
Output: K: generated SuperpixelGridMix image

K ← I;
/* Superpixel segmentation of K */
List Superpixel ← Superpixel_segmentation(K,q);
/* Superpixel parsing */
for (i = 1; i ≤ Superpixel.size(); i++) do
    p = rand(0, 1);
    if (p < r) then
        /* Pixel parsing & copying */
        for
            (j = 1; j ≤ Superpixel(i).size(); j++)
                do
                    PixelK(j) ← PixelJ(j)
return K
```

employs a second image that belongs to the same image class than the first image and a correspondence in between pixel positions is preserved during the texture mapping.

The methods *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* can be unified using the expression:

$$K = (1 - M) \odot I + M \odot J \quad (2)$$

where \odot is the element-wise operation, $A = (Superpixel_A(1), \dots, Superpixel_A(q))^t$, $A \in \{I, J, K\}$. $M = (M(1), \dots, M(q))^t$ with

$$M(i) = \begin{cases} 0 & \text{if } Superpixel(i) \text{ untouched} \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

In Eq.(2), function 1 uses $J = 0$ and function 2 employs the image J defined by:

$$Superpixel_J(i) = mean(Superpixel_I(i)) \quad (4)$$

Finally, in case of function 3, the image J is an image different than I , selected from the dataset.

Besides, by applying the *SuperpixelGridCut* (or *SuperpixelGridMean*), a one-to-one augmentation is obtained in the sense that each original image can produce one augmented image. By this way, the size of a considered training image set can then be multiplied by 2. By applying the *SuperpixelGridMix* to a training set of N images, each image can potentially be mixed with $N - 1$ images. The number of images in the training set after augmentation can reach a number:

$$a = N + N \times (N - 1) = x^2 \quad (5)$$

where $N^2 - N$ images are augmented ones. We remind that in this latter case, each original image of the training set has its own generated superpixel grid. These calculations assume the use of our augmentation methods in their conventional forms.

Further augmented images can be obtained by varying parameters shown in Figure 2 such as instances of the random positioning for processed superpixels (e.g. using other seeds s), the superpixel quantity q , the threshold of processed superpixel occupancy r or even the superpixel grid type t by applying other superpixel segmentation methods.


4. Experimental Results and Evaluation

4.1. Datasets, Implementation Details and Protocol

For our experimental needs, we exploited two image datasets of different natures that are publicly available, namely a *PASCAL VOC* dataset and the *Chest X-Ray Images* dataset. *PASCAL VOC*¹ corresponds to a standardized image set for object class recognition e.g. towards urban scene analysis. For simulating conditions for which a classification problem is new and has a limited number of labeled images, we used the *PASCAL VOC*¹ dataset which is composed of four image classes; namely, motorbikes, bicycles, people, and cars in arbitrary pose. A subset of 679 color images having varied resolutions has been experimented. Besides, *Chest X-Ray Images*² is a dataset which is exploited in biomedical studies for pneumonia analysis. This latter includes two specific image classes; namely, normal and viral or bacterial pneumonia. *Chest X-Ray Image* contains 5856 gray-scale images having relatively high resolutions. It is worth mentioning that the visual classification of such a dataset is extremely difficult for a non-expert since one single image type is represented, namely chest X-Ray images. Hence, discussed augmentation methods are experimented over two challenging classification tasks.

A VGG19 classification architecture has been exploited since this latter is often used as a comparison baseline in image classification studies. In our case, the objective is not to outperform state-of-the-art results by comparing various CNNs performances over a specific classification problem, but observing the basis performance over a reference classification architecture and the impact of data augmentation methods over varied datasets.

By using the selected *PASCAL VOC* and *Chest X-Ray Images* datasets; *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* have then been experimented over binary and multi-class image classification problems. The dataset baselines have been calculated using VGG19 without data augmentation. For measuring the impact of our methods, the training stage has been done from scratch



S_0 : Starting state, baseline known
 S_1 : Baseline overpassed
 S_2 : Inferior or equal to the baseline result
 a_0 : Data augmentation & training
 a_1 : Result double-checking
 a_2 : Tuning augmentation & training

Figure 3. State diagram which models an instance of our protocol for classification performance searching. Considered experimental states and actions are represented.

when using the selected PASCAL VOC since some of its object classes (e.g. vehicles) are already well represented in ImageNet [21]. However, our VGG19 architecture uses an ImageNet pretraining for classifying the Chest X-Ray Images dataset because of its high inter-class similarity.

SuperpixelGridCut, *SuperpixelGridMean* and *SuperpixelGridMix* have employed the SLIC superpixel segmentation method [22] for generating the superpixel grids with close superpixel sizes. The data augmentations *SuperpixelGridCut* and *SuperpixelGridMean* have been applied to the original images of our training sets permitting them to multiply their sizes by 2. Regarding the *SuperpixelGridMix* data augmentation, it is worth mentioning that it can directly be applied in between original images if they share the same size but a strategy has to be adopted if the image size differs from one image to the other. In our implementation, a segmented original image (i.e. having a generated superpixel grid) uses a second original image for a *SuperpixelGridMix* if this latter has a size that is superior or equal to the segmented one. By this way, border effects are avoided during the processing. Also, we apply this method on the set of consecutive image pairs from the training sets. In this case, the number of augmented images is approximately equal to the size of the original training sets, respectively.

Figure 3 displays a one-pass scenario of our protocol for classification performance searching. By this way, experimented states and actions are concisely presented. Let us consider a starting state S_0 for which the baseline classification result obtained from a query dataset and a

selected model are known. First action a_0 consists of applying a query data augmentation method and training a model. This action can conduct to overpass the baseline result (state S_1); otherwise, the result is inferior or equal to the baseline (state S_2). Next action a_1 aims to double-check this latter result. This can conduct to overpass the baseline result (state S_1); otherwise, if the result is still inferior to the baseline (state S_2), then action a_2 consists of tuning the augmentation method and training. This can lead to obtain a result better than the baseline (state S_1); otherwise the process is restarted (state S_0) e.g. acting by using another data augmentation method.

4.2. Analysis of Generated SuperpixelGridMasks

Training image sets of previously described datasets have been augmented using *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix*. The SLIC Superpixel segmentation method has been used. Low and high quantities of generated superpixels as well as low and high contiguity level associated with the processed superpixels have been experimented. A sample of generated *SuperpixelGridMask* images are shown in Figures 1, 4 and 5 with varied respective datasets.

In particular, Figure 4 shows augmentations generated by applying the proposed *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* over urban and semi-urban scene images. Top rows and bottom rows show our results with two selected superpixel quantities, $q = 200$ and $q = 500$, respectively. The selected ratio of processed superpixels is $r = 0.4$. We can see that the images contain foreground, object and background parts which mainly are road, car and forest or building facade textures.

We can observe that the superpixel segmentation has permitted to segment object parts (on the contrary of quadrilateral segmentation operated by cutting techniques). For instance, in the first row, car parts (a wheel, part of doors, driver’s window) are distinguishable when using the *SuperpixelGridCut*. The *SuperpixelGridMean* which sparsely averages superpixels acts on the images like a slight blurring of object parts. It looks like low variations of textures and shapes of object parts at local positions (see e.g. the deformed tree trunks or road lines). Besides, we can observe that the *SuperpixelGridMix* which exploited two images having different backgrounds (forest and building facade, respectively) has created an image which looks like having a new background appearing as a hedge of shrubs at the bottom of the building facade. We emphasize that the superpixel grid can be more or less regular when the superpixels are small and located in homogeneous areas (e.g. see road superpixels at the bottom row).

In particular, Figure 5 shows augmentations generated by applying the proposed *SuperpixelGridCut*,




Figure 4. Augmentations generated by applying the proposed *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* over urban and semi-urban scene images. Top rows and bottom rows show our results with two selected superpixel quantities, $q = 200$ and $q = 500$, respectively. The selected ratio of processed superpixels is $r = 0.4$.

SuperpixelGridMean and *SuperpixelGridMix* over X-Ray images. Top rows and bottom rows show our results with two selected superpixel quantities, $q = 200$ and $q = 1000$, respectively. The selected ratio of processed superpixels is $r = 0.4$. We can see that the X-Ray superpixels contain well delimited object parts (e.g. lungs, ribs). Proposed *SuperpixelGridMasks* have permitted to generate object-related variants of Chest X-Ray images towards further experiments.

Besides, we emphasize that *SuperpixelGridCut* is the most penalizing data augmentation method by comparison to *SuperpixelGridMean* or *SuperpixelGridMix*. Indeed, *SuperpixelGridCut* occludes a set of superpixels for considered original images with black masks whereas *SuperpixelGridMean* or *SuperpixelGridMix*, either averaged superpixels or mixed superpixels, respectively.

In Figure 6, we analyzed the behavior of our *SuperpixelGridCut* augmentation method over original images of the *Chest X-Ray Images* dataset. To this end, we compared this latter method (see a sample in Figure 6b) with the CutOut method which also proceeds

by masking an image region by a black mask (see a sample in Figure 6a). For a fair comparison, the ratio of masking of both is approximately equal to 0.2. For comparing, samples of original training images and corresponding augmented images have been passed through a data dimensionality reduction method. In particular, PCA (Principal Component Analysis) and t-SNE (t-distributed stochastic neighbor embedding) [23] have been applied to generate a distribution of points towards visualization and analysis of image samples (400 images used for each generated view).

In Figure 6c, we can observe that two dissociated clusters of compact points are formed; one corresponds to the not augmented images (normal and pneumonia), the other one corresponds to the images augmented by using CutOut. The Figure 6d shows a single cluster of points with not augmented images regrouped at its center and images augmented by *SuperpixelGridCut* homogeneously distributed in its periphery. Although CutOut and *SuperpixelGridCut* representations mask in black an approximately similar proportion for one considered image,

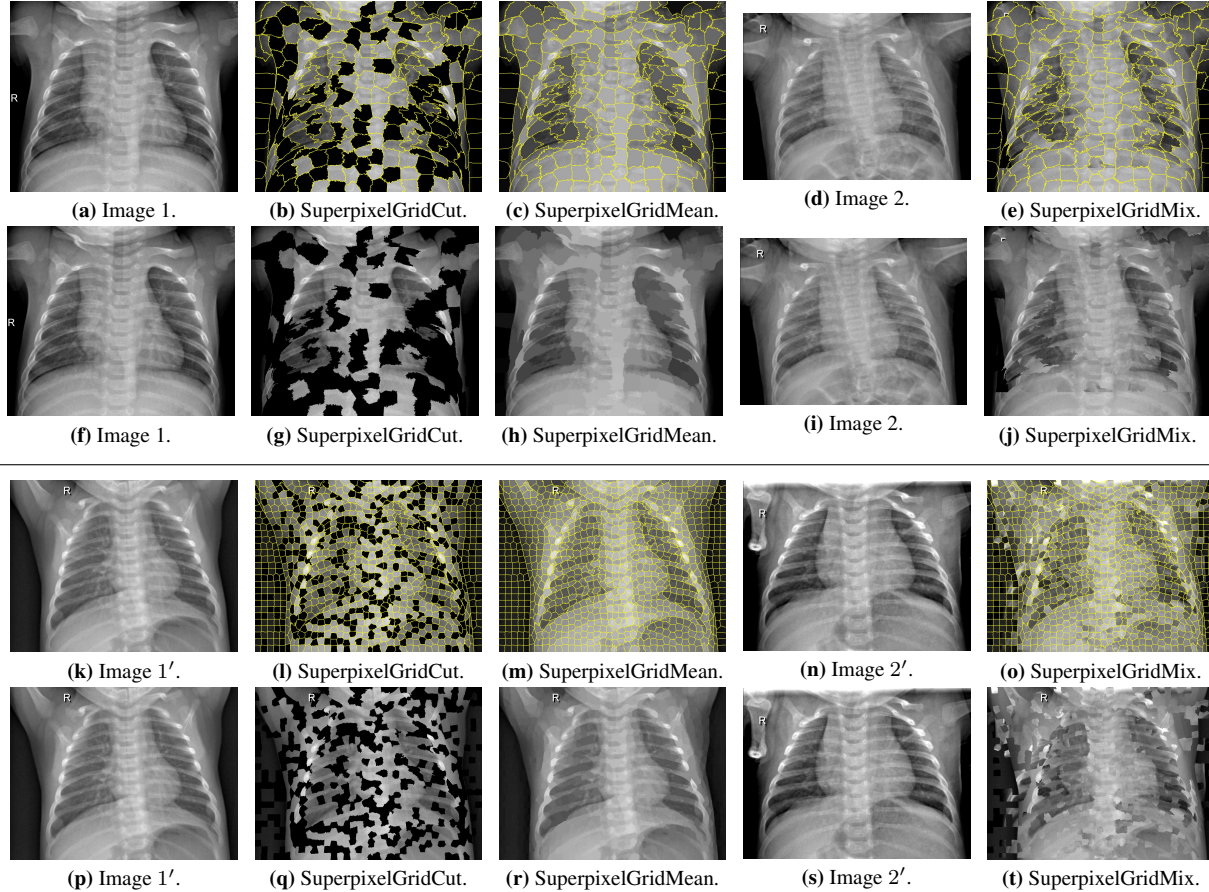


Figure 5. Augmentations generated by applying the proposed *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* over X-Ray images. Top rows and bottom rows show our results with two selected superpixel quantities, $q = 200$ and $q = 1000$, respectively. The selected ratio of processed superpixels is $r = 0.4$.

their applications to a same set of original images conduct to the generation of augmented datasets with different distribution characteristics.

4.3. Analysis of Classification Results

Table 1 reports top-1 accuracies of our classification results using VGG19 over a PASCAL VOC dataset and the Chest X-Ray Images dataset.

More precisely, Table 2 presents classification results of our proposed data augmentation methods and diverse ones obtained using VGG19 over a PASCAL VOC dataset using varied parametrization. The first row shows a baseline accuracy which is equal to 57.81%. The augmentation methods are divided into two parts, the methods proceeding using a single image and the methods mixing two images. In the first part, we remark that *SuperpixelGridMean* provides the best accuracy of 75.00%. *CutOut* and *SuperpixelGridCut* obtain close accuracies of 65.63% and 64.06%, respectively. *Horizontal flip* provides an accuracy of 62.50% while the *Adjust brightness* equals

Augmentation	Parameters	Accuracy (%)
Baseline	-	57.81
Horizontal flip [1]	-	62.50
Adjust brightness [1]	delta=0.1	57.81
CutOut [2]	r=0.2	65.63
CutOut [2]	r=0.4	65.63
SuperpixelGridCut (Ours)	(q=100, r = 0.2)	64.06
SuperpixelGridMean (Ours)	(q=1000, r = 0.4)	75.00 (+17.19)
CutMix [3]	r=0.2	67.19
CutMix [3]	r=0.4	60.94
SuperpixelGridMix (Ours)	(q=100, r = 0.2)	75.00 (+17.19)

Table 2. Classification results of our proposed data augmentation methods and diverse ones obtained using VGG19 over a PASCAL VOC¹ dataset using varied parametrization.

the baseline accuracy. In the second part, *CutMix* reaches an accuracy of 67.19% and *SuperpixelGridMix* provides an accuracy of 75.00%. Hence, *SuperpixelGridMix* as well as *SuperpixelGridMean* obtain the best results amongst the




Figure 6. Subfigures 6a and 6b show images from the *Chest X-Ray Images* when augmented by CutOut and SuperpixelGridCut with a similar masking ratio of approximately 0.2. Subfigures 6c and 6d show the distribution of points (images) generated via t-SNE by using 100 *Normal* and 100 *Bacterial pneumonia* images without augmentation (Originals) compared to their generated CutOut and SuperpixelGridCut augmentations, respectively.

considered augmentation methods and +17.19% compared to the baseline accuracy of 57.81%.

Table 3 presents classification results of our proposed data augmentation methods and diverse ones obtained using VGG19 over the *Chest X-Ray Images* dataset using varied parametrization. More precisely, the first row shows a baseline accuracy which is equal to 78.68%. The augmentation methods are divided into two parts, the methods proceeding using a single image and the methods mixing two images. In the first part, we remark that *SuperpixelGridMean* provides the best accuracy of 83.65%. The random brightness reaches an accuracy of 82.37%. CutOut (2 parametrizations) and *SuperpixelGridCut* obtain close accuracies of 80.60%, 80.28% and 80.44%, respectively. In the second part,

Augmentation	Parameters	Accuracy (%)
Baseline	-	78.68
Horizontal flip [1]	-	79.32
Random brightness [1]	range=[0.5,2.0]	82.37
CutOut [2]	r=0.2	80.60
CutOut [2]	r=0.4	80.28
SuperpixelGridCut (Ours)	(q=1000, r = 0.2)	80.44
SuperpixelGridMean (Ours)	(q=1000, r = 0.2)	83.65 (+4.97)
CutMix [3]	r=0.2	85.09
CutMix [3]	r=0.4	79.80
SuperpixelGridMix (Ours)	(q=1000, r = 0.4)	88.14 (+9.46)

Table 3. Classification results of our proposed data augmentation methods and diverse ones obtained using VGG19 over the *Chest X-Ray Images*² (Kaggle) dataset using varied parametrization.

CutMix provides to the best an accuracy of 85.09% and *SuperpixelGridMix* reaches an accuracy of 88.14%. This latter obtains the best results amongst the considered augmentation methods and +9.46% compared to the baseline accuracy of 78.68%.

We observe that our methods can overpass baseline results. Our experiments show that *SuperpixelGridCut* and CutOut can provide close performances. In this study, *SuperpixelGridMean* outperforms CutOut and *SuperpixelGridMix* is the most efficient method amongst all the considered data augmentation methods.

5. Conclusion

A new efficient and flexible object-related data augmentation approach has been presented. Proposed *SuperpixelGridCut*, *SuperpixelGridMean* and *SuperpixelGridMix* variants act limiting the inclusion of linear, right-angled and regular noises in the considered image analysis networks. A bunch of CutOut and CutMix-related data augmentation methods that currently exploit linear structures could easily adopt our approach in order to adapt their existing models towards targeting better performances. The source codes of each proposed method has publicly been made available online for permitting to anyone the easy augmentation of his/her image datasets and further experiments towards better performances. A SuperpixelGridMask branch is introduced in the category of dropping and fusing data augmentation methods with a high potential of generalization.

Future works may investigate our SuperpixelGridMasks and potential variants towards improving YOLO-based detectors. Besides, since who can do more can do less,

¹A PASCAL VOC dataset: http://host.robots.ox.ac.uk/pascal/VOC/databases.html#VOC2005_1

²Dataset *Chest X-Ray Images*: <https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia>

CutMean and variants could be experimented since they do not seem to be addressed in the current literature.

References

- [1] A. Buslaev, V. I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, and A. A. Kalinin, "Albumentations: Fast and flexible image augmentations," *Information*, vol. 11, no. 2, 2020. [2](#), [7](#), [8](#)
- [2] T. Devries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," *CoRR*, vol. abs/1708.04552, 2017. [2](#), [7](#), [8](#)
- [3] S. Yun, D. Han, S. Chun, S. Oh, Y. Yoo, and J. Choe, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, (Los Alamitos, CA, USA), pp. 6022–6031, IEEE Computer Society, nov 2019. [2](#), [7](#), [8](#)
- [4] S. Huang, X. Wang, and D. Tao, "Snapmix: Semantically proportional mixing for augmenting fine-grained data," in *AAAI*, 2021. [2](#)
- [5] C. Zhao and Y. Lei, "Intra-class cutmix for unbalanced data augmentation," in *2021 13th International Conference on Machine Learning and Computing, ICMLC 2021*, (New York, NY, USA), p. 246–251, Association for Computing Machinery, 2021. [2](#)
- [6] A. Bochkovskiy, C. Wang, and H. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," *CoRR*, vol. abs/2004.10934, 2020. [2](#)
- [7] P. Chen, S. Liu, H. Zhao, and J. Jia, "Gridmask data augmentation," *CoRR*, vol. abs/2001.04086, 2020. [2](#)
- [8] S. Feng, S. Yang, Z. Niu, J. Xie, M. Wei, and P. Li, "Grid cut and mix: flexible and efficient data augmentation," in *Twelfth International Conference on Graphics and Image Processing (ICGIP 2020)* (Z. Pan and X. Hei, eds.), vol. 11720, pp. 656 – 662, International Society for Optics and Photonics, SPIE, 2021. [2](#)
- [9] M. B. Pereira and J. A. d. Santos, "Chessmix: Spatial context data augmentation for remote sensing semantic segmentation," 2021. [2](#)
- [10] J.-H. Kim, W. Choo, and H. O. Song, "Puzzle mix: Exploiting saliency and local statistics for optimal mixup," in *Proceedings of the 37th International Conference on Machine Learning* (H. D. III and A. Singh, eds.), vol. 119 of *Proceedings of Machine Learning Research*, pp. 5275–5285, PMLR, 13–18 Jul 2020. [2](#)
- [11] D. Walawalkar, Z. Shen, Z. Liu, and M. Savvides, "Attentive cutmix: An enhanced data augmentation approach for deep learning based image classification," 2020. [2](#)
- [12] A. F. M. S. Uddin, M. S. Monira, W. Shin, T. Chung, and S.-H. Bae, "Saliencymix: A saliency guided data augmentation strategy for better regularization," *ArXiv*, vol. abs/2006.01791, 2021. [2](#)
- [13] L. Yang, X. Li, B. Zhao, R. Song, and J. Yang, "Recursivemix: Mixed learning with history," 2022. [2](#)
- [14] C.-L. Li, K. Sohn, J. Yoon, and T. Pfister, "Cutpaste: Self-supervised learning for anomaly detection and localization," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9659–9669, 2021. [2](#)
- [15] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, "AugMix: A simple data processing method to improve robustness and uncertainty," *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [16] H. Naveed, "Survey: Image mixing and deleting for data augmentation," *CoRR*, vol. abs/2106.07085, 2021. [2](#)
- [17] Y. Zhang, L. Yang, H. Zheng, P. Liang, C. Mangold, R. G. Loreto, D. P. Hughes, and D. Z. Chen, "Spda: Superpixel-based data augmentation for biomedical image segmentation," in *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning* (M. J. Cardoso, A. Feragen, B. Glocker, E. Konukoglu, I. Oguz, G. Unal, and T. Vercauteren, eds.), vol. 102 of *Proceedings of Machine Learning Research*, pp. 572–587, PMLR, 08–10 Jul 2019. [2](#)
- [18] I. Acción, F. Argüello, and D. B. Heras, "Dual-window superpixel data augmentation for hyperspectral image classification," *Applied Sciences*, vol. 10, no. 24, 2020. [2](#)
- [19] G. Franchi, N. Belkhir, M. L. Ha, y. hu, A. Bursuc, V. Blanz, and A. Yao, "Robust Semantic Segmentation with Superpixel-Mix," in *The British Machine Vision Conference (BMVC)*, (Online, United Kingdom), Nov. 2021. [2](#)
- [20] M. Wang, X. Liu, Y. Gao, X. Ma, and N. Q. Soomro, "Superpixel segmentation: A benchmark," *Signal Processing: Image Communication*, vol. 56, pp. 28–39, 2017. [3](#)
- [21] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. [5](#)
- [22] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2274–2282, 2012. [5](#)
- [23] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 11 2008. [6](#)