

---

# Unsupervised Object Segmentation by Redrawing

---

**Mickaël Chen**

Sorbonne Université, CNRS, LIP6, F-75005, Paris, France  
mickael.chen@lip6.fr

**Thierry Artières**

Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France  
Ecole Centrale Marseille  
thierry.artieres@centrale-marseille.fr

**Ludovic Denoyer**

Facebook Artificial Intelligence Research  
denoyer@fb.com

## Abstract

Object segmentation is a crucial problem that is usually solved by using supervised learning approaches over very large datasets composed of both images and corresponding object masks. Since the masks have to be provided at pixel level, building such a dataset for any new domain can be very time-consuming. We present ReDO, a new model able to extract objects from images without any annotation in an unsupervised way. It relies on the idea that it should be possible to change the textures or colors of the objects without changing the overall distribution of the dataset. Following this assumption, our approach is based on an adversarial architecture where the generator is guided by an input sample: given an image, it extracts the object mask, then redraws a new object at the same location. The generator is controlled by a discriminator that ensures that the distribution of generated images is aligned to the original one. We experiment with this method on different datasets and demonstrate the good quality of extracted masks.

## 1 Introduction

Image segmentation aims at splitting a given image into a set of non-overlapping regions corresponding to the main components in the image. It has been studied for a long time in an unsupervised setting using prior knowledge on the nature of the region one wants to detect using e.g. normalized cuts and graph-based methods. Recently the rise of deep neural networks and their spectacular performances on many difficult computer vision tasks have led to revisit the image segmentation problem using deep networks in a fully supervised setting [5, 20, 58], a problem referred as semantic image segmentation.

Although such modern methods allowed learning successful semantic segmentation systems, their training requires large-scale labeled datasets with usually a need for pixel-level annotations. This feature limits the use of such techniques for many image segmentation tasks for which no such large scale supervision is available. To overcome this drawback, we follow here a very recent trend that aims at revisiting the unsupervised image segmentation problem with new tools and new ideas from the recent history and success of deep learning [55] and from the recent results of supervised semantic segmentation [5, 20, 58].

Building on the idea of scene composition [4, 14, 18, 56] and on the adversarial learning principle [17], we propose to address the unsupervised segmentation problem in a new way. We start by postulating an underlying generative process for images that relies on an assumption of independence between regions of an image we want to detect. This means that replacing one object in the image with another one, e.g. a generated one, should yield a realistic image. We use such a generative model as a backbone for designing an object segmentation model we call ReDO (ReDrawing of Objects), which outputs are then used to modify the input image by redrawing detected objects. Following ideas from adversarial learning, the supervision of the whole system is provided by a discriminator that is trained to distinguish between real images and fake images generated accordingly to the generative process. Despite being a simplified model for images, we find this generative process effective for learning a segmentation model.

The paper is organized as follows. We present related work in Section 2, then we describe our method in Section 3. We first define the underlying generative model that we consider in Section 3.2 and detail how we translate this hypothesis into a neural network architecture to learn a segmentation module in Section 3.3. Then we give implementation details in Section 4. Finally, we present experimental results on three datasets in Section 5 that explore the feasibility of unsupervised segmentation within our framework and compare its performance against a baseline supervised with few labeled examples.

## 2 Related Work

Image segmentation is a very active topic in deep learning that boasts impressive results when using large-scale labeled datasets. Those approaches can effectively parse high-resolution images depicting complex and diverse real-world scenes into informative semantics or instance maps. State-of-the-art methods use clever architectural choices or pipelines tailored to the challenges of the task [5, 20, 58].

However, most of those models use pixel-level supervision, which can be unavailable in some settings, or time-consuming to acquire in any case. Some works tackle this problem by using fewer labeled images or weaker overall supervision. One common strategy is to use image-level annotations to train a classifier from which class saliency maps can be obtained. Those saliency maps can then be exploited with other means to produce segmentation maps. For instance, WILDCAT [13] uses a Conditional Random Field (CRF) for spatial prediction in order to post-process class saliency maps for semantic segmentation. PRM [59], instead, finds pixels that provoke peaks in saliency maps and uses these as a reference to choose the best regions out of a large set of proposals previously obtained using MCG [2], an unsupervised region proposal algorithm. Both pipelines use a combination of a deep classifier and a method that take advantage of spatial and visual handcrafted image priors.

Co-segmentation, introduced by Rother et al. 2006 [46], addresses the related problem of segmenting objects that are shared by multiple images by looking for similar data patterns in all those images. Like the aforementioned models, in addition to prior image knowledge, deep saliency maps are often used to localize those objects [23]. Unsupervised co-segmentation [22], i.e. the task of covering objects of a specific category without additional data annotations, is a setup that resembles ours. However, unsupervised co-segmentation systems are built on the idea of exploiting features similarity and can't easily be extended to a class-agnostic system. As we aim to ultimately be able to segment very different objects, our approach instead relies on independence between the contents of different regions of an image which is a more general concept.

Fully unsupervised approaches have traditionally been more focused on designing handcrafted features or energy functions to define the desired property of *objectness*. Impressive results have been obtained when making full use of depth maps in addition to usual RGB images [44, 49] but it is much harder to specify good energy functions for purely RGB images. W-NET [55] extracts latent representations via a deep auto-encoder that can then be used by a more classic CRF algorithm. Kanezaki 2018 [28] further incorporate deep priors and train a neural network to directly minimize their chosen intra-region pixel distance. A different approach is proposed by Ji et al. 2019 [26] whose method finds clusters of pixels using a learned distance invariant to some known properties. Unlike ours, none of these approaches are learned entirely from data.

Our work instead follows a more recent trend by inferring scene decomposition directly from data. Stemming from DRAW [19], many of those approaches [4, 14] use an attention network to read a region of an image and a Variational Auto-encoder (VAE) to partially reconstruct the image in an iterative process in order to flesh out a meaningful decomposition. LR-GAN [56] is able to generate

simple scenes recursively, building object after object, and Sbai et al. 2018 [48] decompose an image into single-colored strokes for vector graphics. While iterative processes have the advantage of being able to handle an arbitrary number of objects, they are also more unstable and difficult to train. Most of those can either only be used in generation [56], or only handle very simple objects [4, 14, 18]. As a proof of concept, we decided to first ignore this additional difficulty by only handling a set number of objects but our model can naturally be extended with an iterative composition process. This choice is common among works that, like ours, focus on other aspects of image compositionality. Van Steenkiste et al. 2018 [52] advocates for a generative framework that accounts for relationship between objects. While they do produce masks as part of their generative process, they cannot segment a given image. Closer to our setup, the very recent IODINE [18] propose a VAE adapted for multi-objects representations. Their learned representations include a scene decomposition, but they need a costly iterative refinement process whose performance have only been demonstrated on simulated datasets and not real images. Like ours, some prior work have tried to find segmentation mask by recomposing new images. SEIGAN [42] and Cut & Paste [45] learns to separate object and background by moving the region corresponding to the object to another background and making sure the image is still realistic. These methods however, need to have access to background images without objects, which might not be easy to obtain.

Our work also ties to recent research in disentangled representation learning. Multiple techniques have been used to separate information in factored latent representations. One line of work focuses on understanding and exploiting the innate disentangling properties of Variational Auto-Encoders. It was first observed by  $\beta$ -VAE [21] that VAEs can be constrained to produce disentangled representations by imposing a stronger penalty on the Kullback-Leibler divergence term on the VAE loss. FactorVAE [30] and  $\beta$ -TCVAE [7] extract a total correlation term from the KL term of the VAE objective and specifically re-weight it instead of the whole KL term. In a similar fashion, HFVAE [15] introduces a hierarchical decomposition of the KL term to impose a structure on the latent space. A similar property can be observed with GAN-based models, as shown by InfoGAN [9] which forces a generator to map a code to interpretable features by maximizing the mutual information between the code and the output. Using adversarial training is also a good way to split and control information in latent embeddings. Fader Networks [32] uses adversarial training to remove specific class information from a vector. This technique is also used in adversarial domain adaptation [16, 36, 51] to align embeddings from different domains. Similar methods can be used to build factorial representations instead of simply removing information [6, 10, 11, 38]. Like our work, they use adversarial learning to match an implicitly predefined generative model but for purposes unrelated to segmentation.

### 3 Method

#### 3.1 Overview

A segmentation process  $F$  splits a given image  $\mathbf{I} \in \mathbb{R}^{W \times H \times C}$  into a set of non-overlapping regions.  $F$  can be described as a function that assigns to each pixel coordinate of  $\mathbf{I}$  one of  $n$  regions. The problem is then to find a correct partition  $F$  for any given image  $\mathbf{I}$ . Lacking supervision, a common strategy is to define properties one wants the regions to have, and then to find a partition that produces regions with such properties. This can be done by defining an energy function and then finding an optimal split. The challenge is then to accurately describe and model the statistical properties of meaningful regions as a function one can optimize.

We address this problem differently. Instead of trying to define the right properties of regions at the level of each image, we make assumptions about the underlying generative process of images in which the different regions are explicitly modeled. Then, by using an adversarial approach, we learn the parameters of the different components of our model so that the overall distribution of the generated images matches the distribution of the dataset. We detail the generative process in the section 3.2, while the way we learn  $F$  is detailed in Section 3.3.

#### 3.2 Generative Process

We consider that images are produced by a generative process that operates in three steps: first, it defines the different regions in the image i.e the organization of the scene (*composition step*). Then, given this segmentation, the process generates the pixels for each of the regions **independently** (*drawing step*). At last, the resulting regions are assembled into the final image (*assembling step*).

Let us consider a scene composed of  $n - 1$  objects and one background we refer to as object  $n$ . Let us denote  $\mathbf{M}^k \in \{0, 1\}^{W \times H}$  the mask corresponding to object  $k$  which associates one binary value to each pixels in the final image so that  $\mathbf{M}_{x,y}^k = 1$  iff the pixel of coordinate  $(x, y)$  belongs to object  $k$ . Note that, since one pixel can only belong to one object, the masks have to satisfy  $\sum_{k=1}^n \mathbf{M}_{x,y}^k = 1$  and the background mask  $\mathbf{M}^n$  can therefore easily be retrieved computed from the object masks as  $\mathbf{M}^n = 1 - \sum_{k=1}^{n-1} \mathbf{M}^k$ .

The pixel values of each object  $k$  are denoted  $\mathbf{V}^k \in \mathbb{R}^{W \times H \times C}$ . Given that the image we generate is of size  $W \times H \times C$ , each object is associated with an image of the same size but only the pixels selected by the mask will be used to compose the output image. The final composition of the objects into an image is computed as follows:

$$\mathbf{I} \leftarrow \sum_{k=1}^n \mathbf{M}^k \odot \mathbf{V}^k.$$

To recap, the underlying generative process described previously can be summarized as follow: i) first, the masks  $\mathbf{M}^k$  are chosen together based on a mask prior  $p(\mathbf{M})$ . ii) Then, for each object independently, the pixel values are chosen based on a distribution  $p(\mathbf{V}^k | \mathbf{M}^k, k)$ . iii) Finally, the objects are assembled into a complete image.

This process makes an assumption of independence between the colors and textures of the different objects composing a scene. While this is a naive assumption, as colorimetric values such as exposition, brightness, or even the real colors of two objects, are often related, this simple model still serves as a good prior for our purposes.

### 3.3 From Generative Process to Object Segmentation

Now, instead of considering a purely generative process where the masks are generated following a prior  $p(\mathbf{M})$ , we consider the inductive process where the masks are extracted directly from any input image  $\mathbf{I}$  through the function  $F$  which is the object segmentation function described previously. The role of  $F$  is thus to output a set of masks given any input  $\mathbf{I}$ . The new generative process acts as follows: i) it takes a random image in the dataset and computes the masks using  $F(\mathbf{I}) \rightarrow \mathbf{M}_1, \dots, \mathbf{M}_n$ , and ii) it generates new pixel values for the regions in the image according to a distribution  $p(\mathbf{V}^k | \mathbf{M}^k, k)$ . iii) It aggregates the objects as before.

In order for output images to match the distribution of the training dataset, all the components (i.e  $F$  and  $p(\mathbf{V}^k | \mathbf{M}^k, k)$ ) are learned adversarially following the GAN approach. Let us define  $D : \mathbb{R}^{W \times H \times C} \rightarrow \mathbb{R}$  a discriminator function able to classify images as *fake* or *real*. Let us denote  $G_F(\mathbf{I}, \mathbf{z}_1, \dots, \mathbf{z}_n)$  our generator function able to compose a new image given an input image  $\mathbf{I}$ , an object segmentation function  $F$ , and a set of vectors  $\mathbf{z}_1, \dots, \mathbf{z}_n$  each sampled independently following a prior  $p(\mathbf{z})$  for each object  $k$ , background included. Since the pixel values of the different regions are considered as independent given the segmentation, our generator can be decomposed in  $n$  generators denoted  $G_k(\mathbf{M}^k, \mathbf{z}_k)$ , each one being in charge of deciding the pixel values for one specific region. The complete image generation process thus operates in three steps:

- 1)  $\mathbf{M}^1, \dots, \mathbf{M}^n \leftarrow F(\mathbf{I})$  (composition step)
- 2)  $\mathbf{V}^k \leftarrow G_k(\mathbf{M}^k, \mathbf{z}_k)$  for  $k \in \{1, \dots, n\}$  (drawing step)
- 3)  $G_F(\mathbf{I}, \mathbf{z}_1, \dots, \mathbf{z}_n) = \sum_{k=1}^n \mathbf{M}^k \odot \mathbf{V}^k$  (assembling step).

Provided the functions  $F$  and  $G_k$  are differentiable, they can thus be learned by solving the following adversarial problem:

$$\min_{G_F} \max_D \mathcal{L} = \mathbb{E}_{\mathbf{I} \sim p_{data}} [\log D(\mathbf{I})] + \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{z}_1, \dots, \mathbf{z}_n \sim p(\mathbf{z})} [\log(1 - D(G_F(\mathbf{I}, \mathbf{z}_1, \dots, \mathbf{z}_n)))] .$$

Therefore, in practice we have  $F$  output soft masks in  $[0, 1]$  instead of binary masks. Also, in line with recent GAN literature [3, 39, 50, 57], we choose to use the hinge version of the adversarial loss

[35, 50] instead, and obtain the following formulation:

$$\begin{aligned}\max_{\mathbf{G}_F} \mathcal{L}_G &= \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{z}_1, \dots, \mathbf{z}_n \sim p(\mathbf{z})} [\mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_1, \dots, \mathbf{z}_n))] \\ \max_{\mathbf{D}} \mathcal{L}_D &= \mathbb{E}_{\mathbf{I} \sim p_{data}} [\min(0, -1 + \mathcal{D}(\mathbf{I}))] \\ &\quad + \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{z}_1, \dots, \mathbf{z}_n \sim p(\mathbf{z})} [\min(0, -1 - \mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_1, \dots, \mathbf{z}_n)))].\end{aligned}$$

Still, as it stands, the learning process of this model may fail for two reasons. First, it does not have to extract a meaningful segmentation in regards to the input  $\mathbf{I}$ . Indeed, since the values of all the output pixels will be generated,  $\mathbf{I}$  can be ignored entirely to generate plausible pictures. For instance, the segmentation could be the same for all the inputs regardless of input  $\mathbf{I}$ . Second, it naturally converges to a trivial extractor  $F$  that puts the whole image into a single region, the other regions being empty. We thus have to add additional constraints to our model.

**Constraining mask extraction by redrawing a single region.** The first constraint aims at forcing the model to extract meaningful region masks instead of ignoring the image. To this end, we take advantage of the assumption that the different objects are independently generated. We can, therefore, replace only one region at each iteration instead of regenerating all the regions. Since the generator now has to use original pixel values from the image in the reassembled image, it cannot make arbitrary splits. The generation process becomes as follows:

$$\begin{aligned}1) \mathbf{M}^1, \dots, \mathbf{M}^n &\leftarrow F(\mathbf{I}) && \text{(composition step)} \\ 2) \mathbf{V}^k &\leftarrow \mathbf{I} \text{ for } k \in \{1, \dots, n\} \setminus \{\mathbf{i}\} \\ \mathbf{V}^{\mathbf{i}} &\leftarrow \mathbf{G}_{\mathbf{i}}(\mathbf{M}^{\mathbf{i}}, \mathbf{z}_{\mathbf{i}}) && \text{(drawing step)} \\ 3) \mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i}) &= \sum_{k=1}^n \mathbf{M}^k \odot \mathbf{V}^k && \text{(assembling step),}\end{aligned}$$

where  $\mathbf{i}$  designates the index of the only region to redraw and is sampled from  $\mathcal{U}(n)$ , the discrete uniform distribution on  $\{1, \dots, n\}$ . The new learning objectives are as follows:

$$\begin{aligned}\max_{\mathbf{G}_F} \mathcal{L}_G &= \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{i} \sim \mathcal{U}(n), \mathbf{z}_{\mathbf{i}} \sim p(\mathbf{z})} [\mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i}))] \\ \max_{\mathbf{D}} \mathcal{L}_D &= \mathbb{E}_{\mathbf{I} \sim p_{data}} [\min(0, -1 + \mathcal{D}(\mathbf{I}))] + \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{i} \sim \mathcal{U}(n), \mathbf{z}_{\mathbf{i}} \sim p(\mathbf{z})} [\min(0, -1 - \mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i})))]\end{aligned}$$

**Conservation of Region Information.** The second constraint is that given a region  $\mathbf{i}$  generated from a latent vector  $\mathbf{z}_{\mathbf{i}}$ , the final image  $\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i})$  must contain information about  $\mathbf{z}_{\mathbf{i}}$ . This constraint is designed to prevent the mask extractor  $F$  to produce empty regions. Indeed, if region  $\mathbf{i}$  is empty, i.e.  $\mathbf{M}_{x,y}^{\mathbf{i}} = 0$  for all  $x, y$ , then  $\mathbf{z}_{\mathbf{i}}$  cannot be retrieved from the final image. Equivalently, if  $\mathbf{z}_{\mathbf{i}}$  can be retrieved, then region  $\mathbf{i}$  is not empty. This information conservation constraint is implemented through an additional term in the loss function. Let us denote  $\delta_k$  a function which objective is to infer the value of  $\mathbf{z}_k$  given any image  $\mathbf{I}$ . One can learn such a function simultaneously to promote conservation of information by the generator. This strategy is similar to the mutual information maximization used in InfoGAN. [9].

The final complete process is illustrated in Figure 1 and correspond to the following learning objectives:

$$\begin{aligned}\max_{\mathbf{G}_F, \delta} \mathcal{L}_G &= \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{i} \sim \mathcal{U}(n), \mathbf{z}_{\mathbf{i}} \sim p(\mathbf{z})} [\mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i})) - \lambda_z ||\delta_{\mathbf{i}}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i})) - \mathbf{z}_{\mathbf{i}}||_2^2] \\ \max_{\mathbf{D}} \mathcal{L}_D &= \mathbb{E}_{\mathbf{I} \sim p_{data}} [\min(0, -1 + \mathcal{D}(\mathbf{I}))] + \mathbb{E}_{\mathbf{I} \sim p_{data}, \mathbf{i} \sim \mathcal{U}(n), \mathbf{z}_{\mathbf{i}} \sim p(\mathbf{z})} [\min(0, -1 - \mathcal{D}(\mathbf{G}_F(\mathbf{I}, \mathbf{z}_{\mathbf{i}}, \mathbf{i}))],\end{aligned}$$

where  $\lambda_z$  is a fixed hyper-parameter that controls the strength of the information conservation constraint. Note that the constraint is necessary for our model to find non trivial solutions, as otherwise, putting the whole image into a single region is both optimal and easy to discover for the neural networks. The final learning algorithm follows classical GAN schema [3, 17, 39, 57] by updating the generator and the discriminator alternatively with the update functions presented in Algorithm 1.

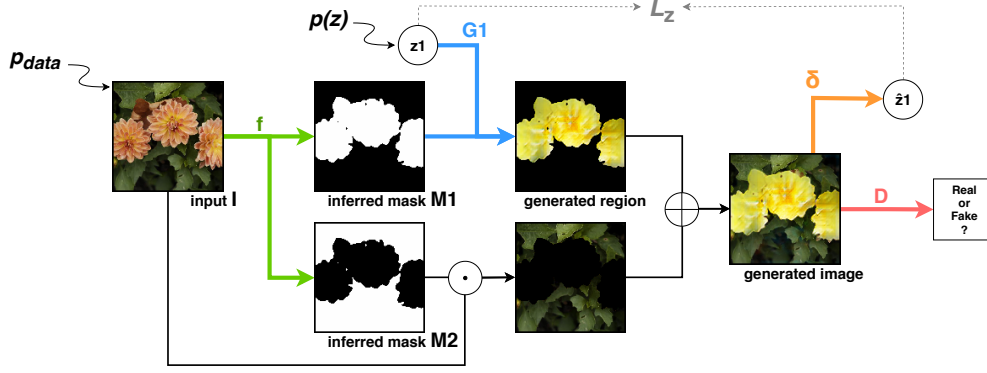


Figure 1: Example generation with  $G_f(\mathbf{I}, \mathbf{z}_i, \mathbf{i})$  with  $\mathbf{i} = 1$  and  $n = 2$ . Learned functions are in color.

---

**Algorithm 1** Networks update functions

---

- 1: **procedure** GENERATORUPDATE
  - 2:   sample data  $\mathbf{I} \sim p_{data}$ ,
  - 3:   sample region  $\mathbf{i} \sim \text{Uniform}(\{1, \dots, n\})$
  - 4:   sample noise vector  $\mathbf{z}_i \sim p(\mathbf{z})$
  - 5:    $\mathbf{I}_{gen} \leftarrow G_f(\mathbf{I}, \mathbf{z}_i, \mathbf{i})$  ▷ generate image
  - 6:    $\mathcal{L}_z \leftarrow -\|\delta_i(\mathbf{I}_{gen}) - \mathbf{z}_i\|_2$  ▷ compute information conservation loss
  - 7:    $\mathcal{L}_G \leftarrow D(\mathbf{I}_{gen})$  ▷ compute adversarial loss
  - 8:   update  $G_f$  with  $\nabla_{G_f}[\mathcal{L}_G + \mathcal{L}_z]$
  - 9:   update  $\delta_i$  with  $\nabla_{\delta_i} \mathcal{L}_z$
  - 10: **procedure** DISCRIMINATORUPDATE
  - 11:   sample datapoints  $\mathbf{I}_{real}, \mathbf{I}_{input} \sim p_{data}$
  - 12:   sample region  $\mathbf{i} \sim \text{Uniform}(\{1, \dots, n\})$
  - 13:   sample noise vector  $\mathbf{z}_i \sim p(\mathbf{z})$
  - 14:    $\mathbf{I}_{gen} \leftarrow G_f(\mathbf{I}_{input}, \mathbf{z}_i, \mathbf{i})$  ▷ generate image
  - 15:    $\mathcal{L}_D \leftarrow \min(0, -1 + D(\mathbf{I}_{real})) + \min(0, -1 - D(\mathbf{I}_{gen}))$  ▷ compute adversarial loss
  - 16:   update  $D$  with  $\nabla_D \mathcal{L}_D$
- 

## 4 Implementation

We now provide some information about the architecture of the different components (additional details are given in Supplementary materials). As usual with GAN-based methods, the choice of a good architecture is crucial. We have chosen to build on the GAN and the image segmentation literature and to take inspiration from the neural network architectures they propose.

For the mask generator  $F$ , we use an architecture inspired by PSPNet [58]. The proposed architecture is a fully convolutional neural network similar to one used in image-to-image translation [60], to which we add a Pyramid Pooling Module [58] whose goal is to gather information on different scales via pooling layers. The final representation of a given pixel is thus encouraged to contain local, regional, and global information at the same time.

The region generators  $G_k$ , the discriminator  $D$  and the network  $\delta$  that reconstructs  $\mathbf{z}$  are based on SAGAN [57] that is frequently used in recent GAN literature [3, 37]. Notably, we use spectral normalization [39] for weight regularization for all networks except for the mask provider  $F$ , and we use self-attention [57] in  $G_k$  and  $D$  to handle non-local relations. To both promote stochasticity in our generators and encourage our latent code  $\mathbf{z}$  to encode for texture and colors, we also use conditional batch-normalization in  $G_k$ . The technique has emerged from style modeling for style transfer tasks [12, 43] and has since been used for GANs as a mean to encode for style and to improve stochasticity [1, 8, 54]. All parameters of the different  $\delta_k$  functions are shared except for their last layers.

As it is standard practice for GANs [3], we use orthogonal initialization [47] for our networks and ADAM [31] with  $\beta = (0, .9)$  as optimizer. Learning rates are set to  $10^{-4}$  except for the mask network F which uses a smaller value of  $10^{-5}$ . We sample noise vectors  $\mathbf{z}_i$  of size 32 (except for MNIST where we used vectors of size 16) from  $\mathcal{N}(0, I_d)$  distribution. We used mini-batches of size 25 and ran each experiment on a single NVidia Tesla P100 GPU. Despite our conservation of information loss, the model can still collapse into generating empty masks at the early steps of the training. While the regularization does alleviate the problem, we suppose that the mask generator F can collapse even before the network  $\delta$  learns anything relevant and can act as a stabilizer. As the failures happen early and are easy to detect, we automatically restart the training should the case arise.

We identified  $\lambda_z$  and the initialization scheme as critical hyper-parameters and focus our hyper-parameters search on those. More details, along with specifics of the implementation used in our experiments are provided as Supplementary materials. The code, dataset splits and pre-trained models are also available open-source <sup>1</sup>.

## 5 Experiments

### 5.1 Datasets

We present results on three natural image datasets and one toy dataset. All images have been resized and then cropped to  $128 \times 128$ .

**Flowers** dataset [40, 41] is composed of 8189 images of flowers. The dataset is provided with a set of masks obtained via an automated method built specifically for flowers [40]. We split into sets of 6149 training images, 1020 validation and 1020 test images and use the provided masks as ground truth for evaluation purpose only.

**Labeled Faces in the Wild** dataset [25, 33] is a dataset of 13233 faces. A subpart of the funneled version [24] has been segmented and manually annotated [27], providing 2927 groundtruth masks. We use the non-annotated images for our training set. We split the annotated images between validation and testing sets so that there is no overlap in the identity of the persons between both sets. The test set is composed of 1600 images, and the validation set of 1327 images.

The **Caltech-UCSD Birds 200 2011 (CUB-200-2011)** dataset [53] is a dataset containing 11788 photographs of birds. We use 10000 images for our training split, 1000 for the test split, and the rest for validation.

As a sanity check, we also build a toy dataset **colored-2-MNIST** in which each sample is composed of an uniform background on which we draw two colored MNIST [34] numbers: one odd number and one even number. Odd and even numbers have colors sampled from different distributions so that our model can learn to differentiate them. For this dataset, we set  $n = 3$  as there are three components.

As an additional experiment, we also build a new dataset by fusing Flowers and LFW datasets. This new **Flowers+LFW** dataset has more variability, and contains different type of objects. We used this dataset to demonstrate that ReDO can work without label information on problems with multiple categories of objects.

### 5.2 Results

To evaluate our method ReDO, we use two metrics commonly used for segmentation tasks. The pixel classification accuracy (Acc) measures the proportion of pixels that have been assigned to the correct region. The intersection over union (IoU) is the ratio between the area of the intersection between the inferred mask and the ground truth over the area of their union. In both cases, higher is better. Because ReDO is unsupervised and we can't control which output region corresponds to which object or background in the image, we compute our evaluation based on the regions permutation that matches the ground truth the best. For model selection, we used IoU computed on a held out labeled validation set. When available, we present our evaluation on both the training set and a test set as, in an unsupervised setting, both can be relevant depending on the specific use case. Results are presented in Table 1 and show that ReDO achieves reasonable performance on the three real-world datasets.

<sup>1</sup><https://github.com/mickaelChen/ReDO>



Figure 2: Generated samples (not cherry-picked, zoom in for better visibility). For each dataset, the columns are from left to right: 1) input images, 2) ground truth masks, 3) masks inferred by the model for object one, 4-7) generation by redrawing object one, 8-11) generation by redrawing object two. As we keep the same  $z_i$  on any given column, the color and texture of the redrawn object is kept constant across rows. More samples are provided in Supplementary materials. Faces from the LFW dataset have been anonymized, in visualisations only, to protect personality rights.



Figure 3: Results on LFW + Flowers dataset, arranged as in Figure 2. As  $z$  is kept constant on a column across all rows, we can observe that  $z$  codes for different textures depending on the class of the image even though the generator is never given this information explicitly. Faces from the LFW dataset have been anonymized, in visualisations only, to protect personality rights.

We also compared the performance of ReDO, which is unsupervised, with a supervised method, keeping the same architecture for  $F$  in both cases. We analyze how many training samples are needed to reach the performance of the unsupervised model (see Figure 4). One can see that the unsupervised results are in the range of the ones obtained with a supervised method, and usually outperform supervised models trained with less than around 50 or 100 examples depending on the dataset. For instance, on the LFW Dataset, the unsupervised model obtains about 92% of accuracy and 79% IoU and the supervised model needs 50-60 labeled examples to reach similar performance.

We provide random samples of extracted masks (Figure 2) and the corresponding generated images with a redrawn object or background. Note that our objective is not to generate appealing images but to learn an object segmentation function. Therefore, ReDO generates images that are less realistic than the ones generated by state-of-the-art GANs. Focus is, instead, put on the extracted masks, and we can see the good quality of the obtained segmentation in many cases. Best and worst masks, as well as more random samples, are displayed in Supplementary materials.

We also trained ReDO on the fused Flowers+LFW dataset without labels. We re-used directly the hyper-parameters we have used to fit the Flowers dataset without further tuning and obtained, as preliminary results, a reasonable accuracy of 0.856 and an IoU of 0.691. This shows that ReDO is able to infer class information from masks even in a fully unsupervised setup. Samples are displayed in Figure 3.

Dataset	Train Acc	Train IoU	Test Acc	Test IoU
LFW	-	-	$0.917 \pm 0.002$	$0.781 \pm 0.005$
CUB	$0.840 \pm 0.012$	$0.423 \pm 0.023$	$0.845 \pm 0.012$	$0.426 \pm 0.025$
Flowers*	$0.886 \pm 0.008$	$0.780 \pm 0.012$	$0.879 \pm 0.008$	$0.764 \pm 0.012$
Flowers+LFW	-	-	0.856	0.691

Table 1: Performance of ReDO in accuracy (Acc) and intersection over union (IoU) on retrieved masks. Means and standard deviations are based on five runs with fixed hyper-parameters. LFW train set scores are not available since we trained on unlabeled images. \*Please note that segmentations provided along the original Flowers dataset [41] have been obtained using an automated method. We display samples with top disagreement masks between ReDO and ground truth in Supplementary materials. In those cases, we find ours to provide better masks.

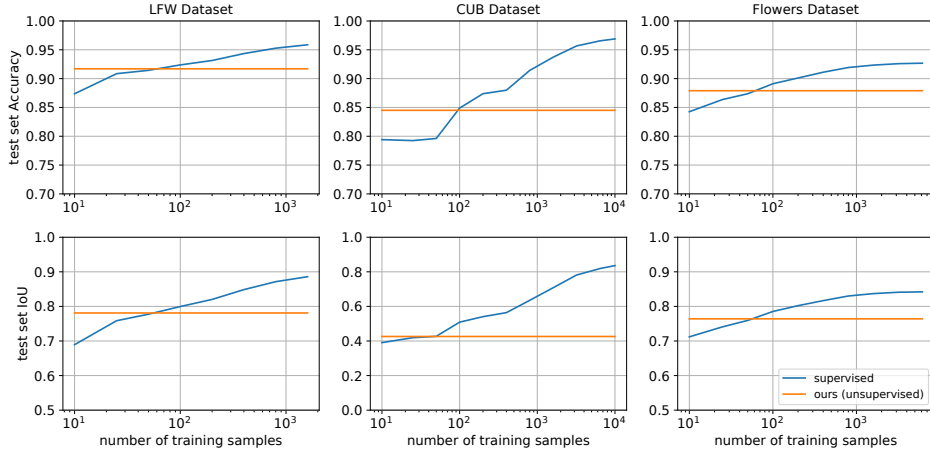


Figure 4: Comparison with supervised baseline as a function of the number of available training samples.

## 6 Conclusion

We presented a novel method called ReDO for unsupervised learning to segment images. Our proposal is based on the assumption that if a segmentation model is accurate, then one could edit any real image by replacing any segmented object in a scene by another one, randomly generated, and the result would still be a realistic image. This principle allows casting the unsupervised learning of image segmentation as an adversarial learning problem. Our experimental results obtained on three datasets show that this principle works. In particular, our segmentation model is competitive with supervised approaches trained on a few hundred labeled examples.

Our future work will focus on handling more complex and diverse scenes. As mentioned in Section 2, our model could generalize to an arbitrary number of objects and objects of unknown classes via iterative design and/or class agnostic generators. Currently, we are mostly limited by our ability to effectively train GANs on those more complicated settings but rapid advances in image generation [3, 29, 37] make it a reasonable goal to pursue in a near future. Meanwhile, we will be investigating the use of the model in a semi-supervised or weakly-supervised setup. Indeed, additional information would allow us to guide our model for harder datasets while requiring fewer labels than fully supervised approaches. Conversely, our model could act as a regularizer by providing a prior for any segmentation tasks.

## Acknowledgments

This work was supported by the French National Research Agency projects LIVES (grant number ANR-15-CE23-0026-03) and "Deep in France" (grant number ANR-16-CE23-0006).

## References

- [1] Amjad Almahairi, Sai Rajeshwar, Alessandro Sordoni, Philip Bachman, and Aaron Courville. Augmented cycleGAN: Learning many-to-many mappings from unpaired data. In *International Conference on Machine Learning*, pages 195–204, 2018.
- [2] P. Arbeláez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *Computer Vision and Pattern Recognition*, 2014.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- [4] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 801–818, 2018.
- [6] Mickael Chen, Ludovic Denoyer, and Thierry Artières. Multi-view data generation without view supervision. In *International Conference on Learning Representations*, 2018.
- [7] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pages 2610–2620, 2018.
- [8] Ting Chen, Mario Lucic, Neil Houlsby, and Sylvain Gelly. On self modulation for generative adversarial networks. In *International Conference on Learning Representations*, 2019.
- [9] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016.
- [10] Emily L Denton and vighnesh Birodkar. Unsupervised learning of disentangled representations from video. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4414–4423. Curran Associates, Inc., 2017.
- [11] Chris Donahue, Akshay Balsubramani, Julian McAuley, and Zachary C. Lipton. Semantically decomposing the latent spaces of generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- [12] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *Proc. of ICLR*, 2, 2017.
- [13] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. Wildcat: Weakly supervised learning of deep convnets for image classification, pointwise localization and segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [14] SM Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Geoffrey E Hinton, et al. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, pages 3225–3233, 2016.
- [15] Babak Esmaeili, Hao Wu, Sarthak Jain, Alican Bozkurt, N Siddharth, Brooks Paige, Dana H Brooks, Jennifer Dy, and Jan-Willem Meent. Structured disentangled representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2525–2534, 2019.
- [16] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 1180–1189. JMLR. org, 2015.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [18] Klaus Greff, Raphaël Lopez Kaufmann, Rishabh Kabra, Nick Watters, Chris Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. *arXiv preprint arXiv:1903.00450*, 2019.

- [19] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: a recurrent neural network for image generation. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 1462–1471. JMLR. org, 2015.
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [21] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [22] Kuang-Jui Hsu, Yen-Yu Lin, and Yung-Yu Chuang. Co-attention cnns for unsupervised object co-segmentation. In *IJCAI*, pages 748–756, 2018.
- [23] Kuang-Jui Hsu, Yen-Yu Lin, and Yung-Yu Chuang. Deepco 3: Deep instance co-segmentation by co-peak search and co-saliency detection. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [24] Gary B. Huang, Vidit Jain, and Erik Learned-Miller. Unsupervised joint alignment of complex images. In *ICCV*, 2007.
- [25] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [26] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information distillation for unsupervised image segmentation and clustering. In *International Conference on Computer Vision (ICCV)*, 2019.
- [27] Andrew Kae, Kihyuk Sohn, Honglak Lee, and Erik Learned-Miller. Augmenting CRFs with Boltzmann machine shape priors for image labeling. In *CVPR*, 2013.
- [28] Asako Kanezaki. Unsupervised image segmentation by backpropagation. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018.
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018.
- [30] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International Conference on Machine Learning*, pages 2654–2663, 2018.
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [32] Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic Denoyer, et al. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems*, pages 5967–5976, 2017.
- [33] Gary B. Huang Erik Learned-Miller. Labeled faces in the wild: Updates and new reporting procedures. Technical Report UM-CS-2014-003, University of Massachusetts, Amherst, May 2014.
- [34] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [35] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- [36] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018.
- [37] Mario Lucic, Michael Tschannen, Marvin Ritter, Xiaohua Zhai, Olivier Bachem, and Sylvain Gelly. High-fidelity image generation with fewer labels. *arXiv preprint arXiv:1903.02271*, 2019.
- [38] Michael F Mathieu, Junbo Jake Zhao, Junbo Zhao, Aditya Ramesh, Pablo Sprechmann, and Yann LeCun. Disentangling factors of variation in deep representation using adversarial training. In *Advances in Neural Information Processing Systems*, pages 5040–5048, 2016.
- [39] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- [40] Maria-Elena Nilsback and Andrew Zisserman. Delving into the whorl of flower segmentation. In *BMVC*, volume 2007, pages 1–10, 2007.

- [41] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008.
- [42] Pavel Ostyakov, Roman Suvorov, Elizaveta Logacheva, Oleg Khomenko, and Sergey I Nikolenko. Seigan: Towards compositional image generation by simultaneously learning to segment, enhance, and inpaint. *arXiv preprint arXiv:1811.07630*, 2018.
- [43] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [44] Trung T Pham, Thanh-Toan Do, Niko Sünderhauf, and Ian Reid. Scenecut: joint geometric and object segmentation for indoor scenes. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018.
- [45] Tal Remez, Jonathan Huang, and Matthew Brown. Learning to segment via cut-and-paste. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 37–52, 2018.
- [46] Carsten Rother, Tom Minka, Andrew Blake, and Vladimir Kolmogorov. Cosegmentation of image pairs by histogram matching-incorporating a global constraint into mrfs. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 1, pages 993–1000. IEEE, 2006.
- [47] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *2nd International Conference on Learning Representations, ICLR 2014*, 2014.
- [48] Othman Sbati, Camille Couprie, and Mathieu Aubry. Vector image generation by learning parametric layer decomposition. *arXiv preprint arXiv:1812.05484*, 2018.
- [49] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012.
- [50] Dustin Tran, Rajesh Ranganath, and David M Blei. Deep and hierarchical implicit models. *arXiv preprint arXiv:1702.08896*, 7, 2017.
- [51] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [52] Sjoerd van Steenkiste, Karol Kurach, and Sylvain Gelly. A case for object compositionality in deep generative models of images. *arXiv preprint arXiv:1810.10340*, 2018.
- [53] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [54] Xiaolong Wang and Abhinav Gupta. Generative image modeling using style and structure adversarial networks. In *European Conference on Computer Vision*, pages 318–335. Springer, 2016.
- [55] Xide Xia and Brian Kulis. W-net: A deep model for fully unsupervised image segmentation. *arXiv preprint arXiv:1711.08506*, 2017.
- [56] Jianwei Yang, Anitha Kannan, Dhruv Batra, and Devi Parikh. LR-GAN: layered recursive generative adversarial networks for image generation. In *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [57] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018.
- [58] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [59] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [60] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.

## Supplementary Material for Unsupervised Object Segmentation by Redrawing

We provide architectural details, hyper-parameters discussion and additional samples output of our model.

### Architectural details

mask network $f(\mathbf{I})$	nonlinearities	output size
Image $\mathbf{I}$		3x128x128
Conv 7x7 (reflect. pad 3)	Instance Norm, ReLU	16x128x128
Conv 3x3 (stride 2, pad 1)	Instance Norm, ReLU	32x64x64
Conv 3x3 (stride 2, pad 1)	Instance Norm, ReLU	64x32x32
Residual Bloc (Instance Norm, ReLU)		64x32x32
Residual Bloc (Instance Norm, ReLU)		64x32x32
Residual Bloc (Instance Norm, ReLU)		64x32x32
Pyramid Pooling Module		68x32x32
Upsample		68x64x64
Conv 3x3 (pad 1)	Instance Norm, ReLU	34x64x64
Upsample		34x128x128
Conv 3x3 (pad 1)	Instance Norm, ReLU	17x128x128
Conv 3x3 (reflect. pad 3)	sigmoid (if $n = 2$ ) or softmax	$n$ x128x128

Table 2: Architecture of mask network  $f$ . Overall architecture is similar to Cycle-GAN for image translation, except with less residual blocs but a Pyramid Pooling Module introduced by PSPNet.

discriminator network $D$ and encoder $\delta$	output size
image input $\mathbf{I}$	3x128x128
Down Res Bloc (ReLU)	64x64x64
Self-Attention Bloc	
Down Res Bloc (ReLU)	128x32x32
Down Res Bloc (ReLU)	256x16x16
Down Res Bloc (ReLU)	512x8x8
Down Res Bloc (ReLU)	1024x4x4
Res Bloc (ReLU)	1024x4x4
Spatial sum pooling	1024x1x1
Linear	1 for $D$ , 32 for $\delta$

Table 3: Architecture of discriminator network  $D$  and encoder  $\delta$ .

## Hyperparameters

We discuss some notable hyper-parameters and architectural choices. We identified  $\lambda_z$  and the initialization scheme as critical hyper-parameters and focus our search on those, in addition to the standard learning rates search. The number of channels in our generators is also important.

- Learning rates are set to  $10^{-4}$  except for mask network for which we use a smaller learning rate ( $10^{-5}$ ). We searched for learning rates independently for each neural network.
- Batch size were set between 20 and 25, as much as we could fit on a single Nvidia Tesla P100 GPU.
- We chose smaller  $size_z$  for each  $\mathbf{z}_k$  (16 for c2-MNIST and 32 for the other datasets) than what is usually found in GAN literature so that the vectors could be reasonably retrieved by  $\delta$ .
- As expected, we found that  $\lambda_z$  were important to tune for the stability of our training procedure. We use  $\lambda_z = \frac{5}{n * size_z}$  for all datasets except LFW where  $\lambda_z = \frac{15}{n * size_z}$ .
- Adequate orthogonal initialization is critical. We found that our model worked best when initialized with a gain at 0.8, and not at all when set at .2 and 1.4.
- We tested smaller numbers of channels for our generators  $G_k$  since each generator only have to model a specific type of object. We still use the same number ( $ch = 64$ ) as the other networks for LFW and colored-2-MNIST but reduced to  $ch = 32$  for CUB and Flowers.
- We used spectral normalization for all our networks except the mask network on which a weight decay of  $10^{-4}$  is applied instead.
- The use of pyramid pooling produce masks of significantly better quality than standard residual network.
- The use of self-attention in both D and G might not be mandatory but more experiments would be needed to conclude.
- We found that having different ratio of discriminator updates and generator updates didn't help.

## Additional output masks

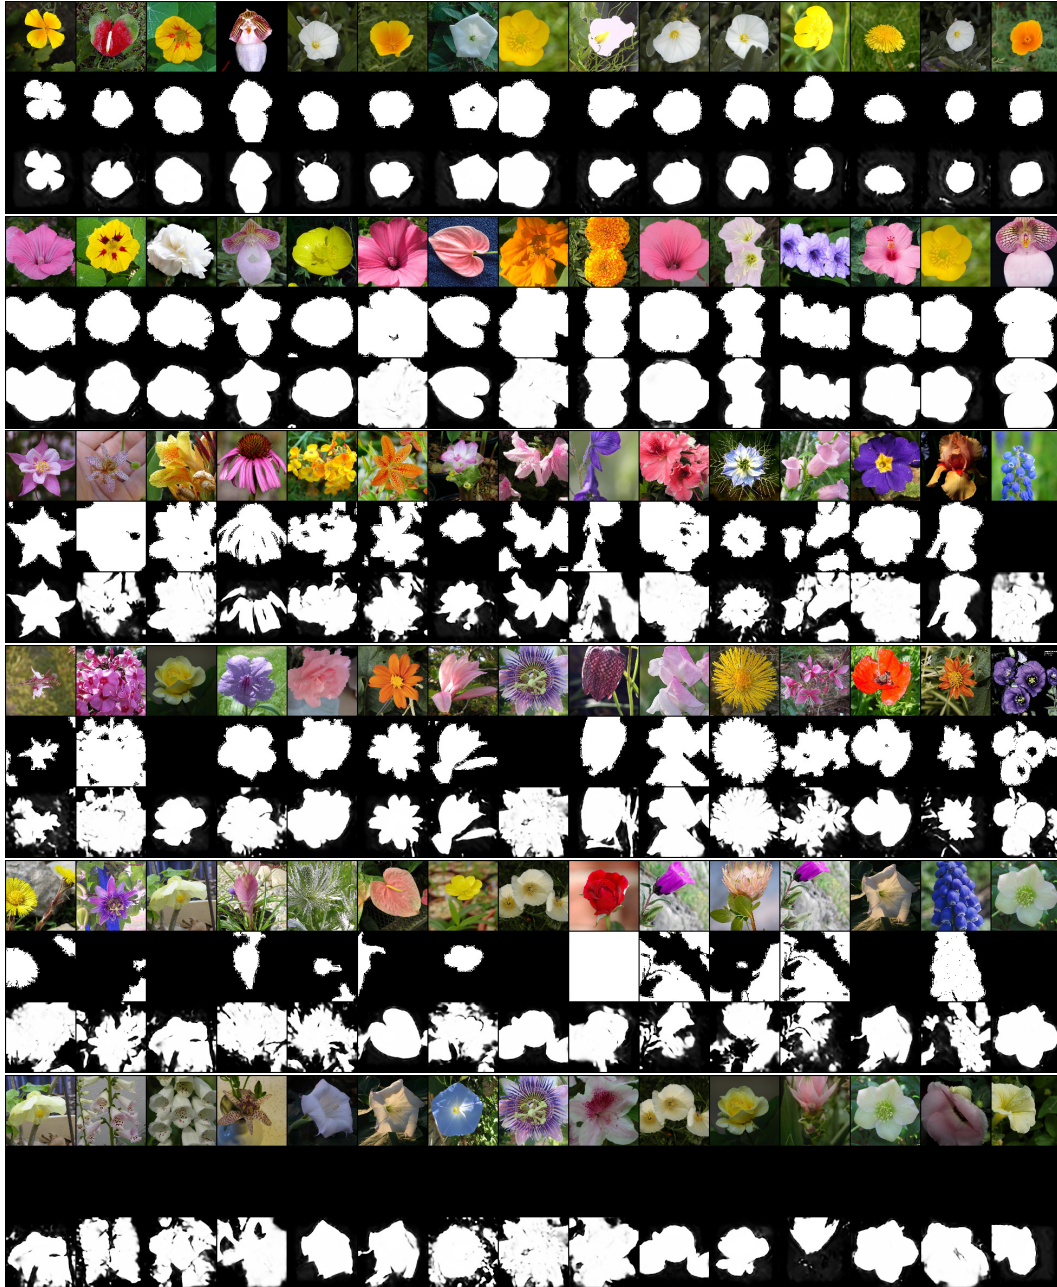


Figure 5: Masks obtained for images from the Flowers test set. Each bloc of three rows depict from top to bottom input image, ground truth and output of our model. Each bloc from top to bottom: 1) top masks according to accuracy 2) top masks according to IoU 3-4) randomly sampled masks 5) worst masks for accuracy 6) worst masks for IoU. Because the ground truth for Flowers were obtained via an automated process, our model actually provide better predictions in worst agreement case.

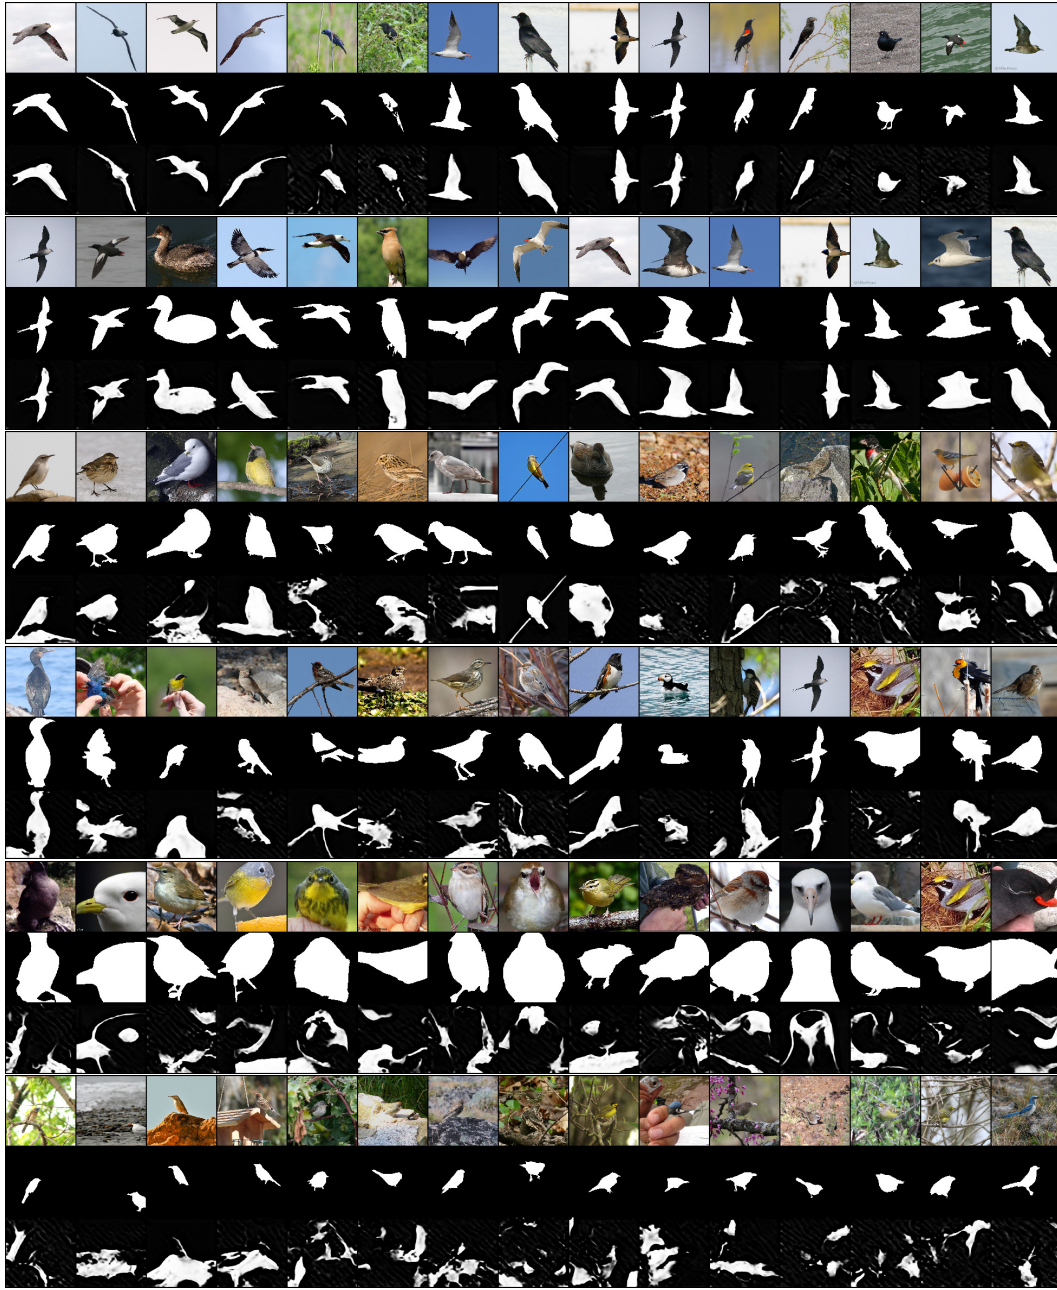


Figure 6: Masks obtained for images from the CUB test set. Each bloc of three rows depict from top to bottom input image, ground truth and output of our model. Each bloc from top to bottom: 1) top masks according to accuracy 2) top masks according to IoU 3-4) randomly sampled masks 5) worst masks for accuracy 6) worst masks for IoU.



Figure 7: More randomly sampled output masks on Flower dataset.

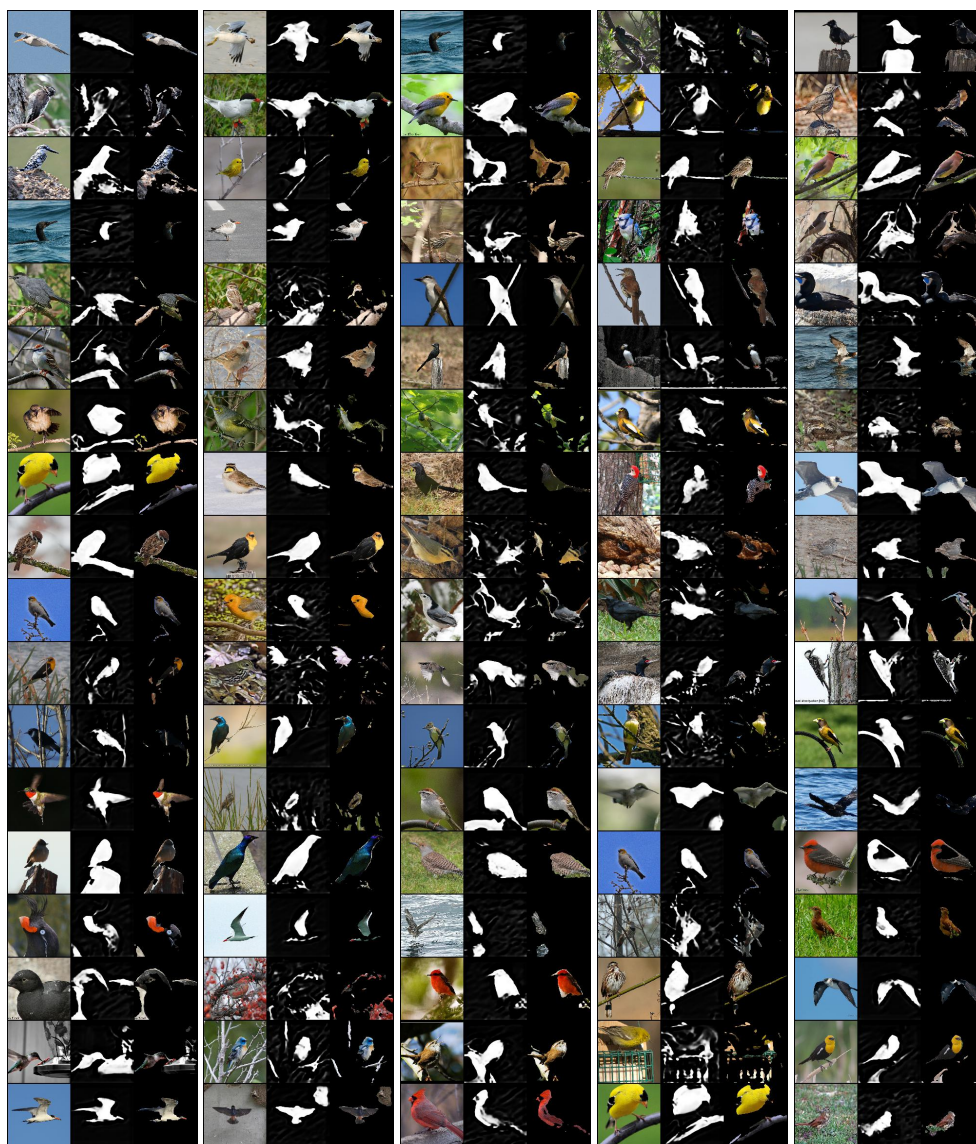


Figure 8: More randomly sampled output masks on CUB dataset.