

UNIVERSITAT DE BARCELONA

FUNDAMENTAL PRINCIPLES OF DATA SCIENCE MASTER'S
THESIS

SSSGAN:Satellite Style and Structure Generative Adversarial Networks

Author:

Emilio TYLSON BAIXAULI

Supervisor

Dr. Sergio ESCALERA

GUERRERO

Dr. Javier MARIN TUR

*A thesis submitted in partial fulfillment of the requirements
for the degree of MSc in Fundamental Principles of Data Science*

in the

Facultat de Matemàtiques i Informàtica

June 30, 2021

UNIVERSITAT DE BARCELONA

Abstract

Facultat de Matemàtiques i Informàtica

MSc

SSSGAN:Satellite Style and Structure Generative Adversarial Networks

by Emilio TYLSON BAIXAULI

This work presents Satellite Style and Structure Generative Adversarial Network (SSGAN), a generative model of high resolution satellite imagery to support image segmentation. Based on spatially adaptive denormalization modules (SPADE) that modulate the activations with respect to segmentation map structure in addition to global descriptor vectors that capture the semantic information in a vector with respect to Open Street Maps (OSM) classes, this model is able to produce consistent aerial imagery. By decoupling the generation of aerial images into a structure map and a carefully defined style vector, we were able to improve the realism and geo-diversity of the synthesis with respect to the state-of-the-art baseline. Therefore, the proposed model allows to control the generation not only with respect to the desired structure, but also with respect to a geographic area.

Acknowledgements

Throughout the writing of this master Thesis I have received a great deal of support and assistance.

I would first like to thank my two supervisors, Dr. Sergio Escalera who trusted in me and gave me the first opportunity to get inserted in the academic field and Dr. Javier Marin who supported me since the first time when I proposed to him this research. Your expertise was invaluable in formulating the research questions and methodology and innovation. The insightful and constant feedback pushed me to sharpen my thinking and brought my work to a higher level.

In addition, I would like to thank my life partner, Noelia. Nothing would have been possible without her companion and her constant support of countless days and nights of working. I would like to thank my parents, Claudio and Silvia. They made me understand effort and knowledge as virtues to constantly follow in life, and this work is a proof of these values.

Finally, I would like to acknowledge my friend and colleague Alejandro Hernandez from my master thesis grade. His suggestions were very valuable in the conformation of this thesis.

This work was supported by the European Regional Development Fund and the Spanish Government, Ministerio de Ciencia, Innovación y Universidades - Agencia Estatal de Investigación - RTC2019-007434-7.

Contents

Abstract	iii
Acknowledgements	v
Contents	1
1 Introduction	3
1.1 Introduction	3
1.2 Related Work	5
1.3 Problem Formulation	9
1.4 Research Questions	10
2 Datasets	11
2.1 Inria Aerial Image Labeling Dataset	11
2.2 Open Street Map (OSM)	11
3 Methods	13
3.1 SPADE	13
Spatially-Adaptive Denormalization	13
3.2 SSSGAN	16
3.3 Global Semantic Vector	17
3.4 Metrics	18
Frechet Inception	19
Sliced Wasserstein	19
4 Results	21
4.1 Implementation details	21
4.2 Quantitative Analysis	21
4.3 Qualitative Analysis	22
5 Conclusions	27
6 Future Works	29
6.1 Improve losses	29
6.2 Augment capabilities	29
Bibliography	31

Chapter 1

Introduction

1.1 Introduction

The commercialization and the advancement of the geospatial industry has led to an explosive amount of remote sensing data being collected to characterize our changing planet Earth. Public and private industries are taking advantage of this increasing availability of information in order to perform analytics and obtain more precise information about geographic areas in order to support decisions and automatize technology. Due to the increasingly revisiting frequency of satellites and a fine pixel resolution (up to 5cm per pixel), satellite imagery became of interest because computer vision algorithms can capture the presence of objects in an automatically and efficiently manner at large scale. Commonly studied computer vision tasks such as semantic and instance segmentation, object detection or height estimation are aimed to address problems such as land cover classification, precision agriculture, flood detection, building, road or car detection, that at the same time it helps to provide information of geographic zones that can improve agriculture, navigation, retail, smart city technologies, 3D precise world reconstruction or even assistance after natural disasters.

State-of-the-art methods comprise mostly on deep learning algorithms. With the presentation of AlexNet [14] in 2012 as winner of ImageNet LSVRC-2012 competition by a large margin, deep networks have dominated the scene of computer vision. Due to their large amount of parameters, they present a high complexity that make them require a high volume of data to correctly extract latent features from imagery, key to achieve outstanding results.

Particularly, in the field of geo-informatics and remote sensing, datasets are usually sparse, expensive and difficult to collect when it comes to tasks that require from high to very high resolution images (0.05 to 1 meters). To overcome this situation of scarcity of images, a commonly used technique is transfer learning. This approach for training consists of using pre-trained weights as a starting point in order to improve performance and decrease training time. Pre-training is done with a highly varied high-volume dataset, so the network can extract low-level features of the world. Then, this pre-trained model is trained again with a smaller task-specific available dataset that is known as fine-tuning. This tuning can be performed by a variety of strategies, that ranges from the most basic ones such as freezing most of the low level layers (layers that have learnt primitive low level features) and only tuning the shallow layers, to more complex schemes that apply different learning rates to different layers.



FIGURE 1.1: Synthetic images generated by SSSGAN

The idea is the model to take advantage of low-level extracted features to learn more easily task-specific features in the fine-tuning. Generally, public pre-trained models are trained in datasets such as ImageNet [3] or similar ones that consist of labeled images used in visual recognition tasks (ground level visualization). Those pre-trained models are applied in totally different domains obtaining an increment in performance with respect to training the network from scratch. For example ImageNet presents completely different visual features with respect to satellital images. Aerial-imagery contains the presence of high-frequency details and a background clutter that heavily depends on the environment, geographic zone, weather conditions, illumination, sensors and pixel resolutions. Those factors constitute a challenge itself for computer vision models to work well in a variety of cities, countries, regions, continents or even pixel resolutions.

The performance of algorithms varies markedly across geographies, image qualities and resolutions. Performance of a model applied in new areas depends, on one hand, on the target texture and topology related to cultural regions and countries [1]. Other crucial characteristics present in the image are the geographic location, weather and type of terrain. An image taken from a rural area totally differs from an urban area or from the coast. Even a specific rural area contains a different biome from a rural area of a different country/region. These points explain why it is really difficult to train a general deep network that works well with images of different locations. Additionally to the image content characteristics, there are image technical characteristics related with the methodology of extraction, such as type of sensor, radiometry, off nadir angle, and atmospheric conditions at the top layers of the atmosphere.

Supervised learning techniques that use deep networks are usually trained with a large number of classes that can go from tens to thousands of labels. Thus, labeling satellite imagery is a fundamental step in the training of deep networks. Depending on the quality of labels and resolution of the images the cost of annotating scenes varies. Generally the most quality satellite imagery labeling is performed by trained professionals with knowledge on GIS and geographic imagery, making this demanding annotation process slow and costly. Even the cost is tightly related

with the resolution of the images, as the spatial resolution of the image increases, the cost of annotation grows accordingly. This produces a scarcity of public datasets and a bias towards most developed urban regions that have enough resources to afford this data acquisition. Scientists should do a careful selection and analysis of the datasets before starting the data annotation phase and also they should put special attention to the quality of the labels.

When a study or a research presents a model claiming to efficiently extract, detect a specific target, it usually implies that they are presenting a model trained with a dataset with specific geographic, cultural and quality conditions that perform well. In order to overcome such necessity, one possibility can be to generate a large collection of diverse synthetic images with their corresponding labels. In this case, it would be necessary to contemplate the different characteristics mentioned before, so the resulting satellite images can augment efficiently in those desired directions.

In this work, we present Satellite Style and Structured Generative Adversarial Network (SSSGAN) to generate realistic synthetic imagery based on publicly available ground truth¹. Particularly, we propose the use of a conditional generative adversarial network (GAN) model capable of generating synthetic satellite images constrained by two components: a semantic style description of the scene in addition to a segmentation map that defines the structure of desired targets objects classes. By this way the structure and the style constraint are decoupled so the user can easily generate novel synthetic images by defining a segmentation mask of the desired foot print labels and then selecting the proportion of semantic classes expressed as number of a vector in addition to the selection of the region or city. With this generation rule the model can capture and express variability present in the satellite imagery while at the same time provides an easy-to-use generation mechanism with high expressiveness. The contribution of this thesis are:

- Development of a GAN model capable of producing highly diverse satellite imagery.
- Presentation of a semantic global vector descriptor dataset based on Open Street Maps (OSM). We analyse and categorize a set of 11 classes that semantically describes the visual features that are present in satellite imagery, leveraging the public description of this crowdsourced database.
- Evaluation and study that describes the different effects of the presented mechanisms.

1.2 Related Work

Synthetic image generation is an active researched topic in the field of computer vision. A vast variety of models have been developed in the past years since the presentation in 2014 of generative adversarial networks (GAN) [6]. Even though, before and after GANS there were a numerous of classical and deep learning methods, the increasing support and improvement on GAN models made this state-of-the-art technique reach outstanding results where the synthetic generated images are hardly distinguishable from the real ones.

As mentioned before, Generative Adversarial Networks (GANS) have stated the

¹Model and code available in https://bitbucket.org/ETyson/sssgan_tfm_ub/src/master/

baseline for deep generative learning. The model consists of two parts: a generator and a discriminator. The generator learns to generate synthetic realistic images while it is trying to fool the discriminator that is responsible to distinguish between real or fake generated images. This learning process consists of finding equilibrium in a two-player minimax game where each iteration the generator G gets better on capturing the real data distribution thanks to the feedback of the discriminator D that at the same time it is also learning important features that help to distinguish if the input image came from the training distribution or not. Mathematically, the generator G learns to map a latent random vector z to a generated sample tensor and tries to maximize the probability D making a mistake, that is to say minimize $\log(1 - D(G(z)))$. On the other hand, the opposite happens to D , it tries to maximize the probability of assigning the correct label $\log(D(x))$ and $\log(1 - D(G(z)))$, where x is a real image and z the latent vector

$$\min_G \max_D L_{\text{GAN}}(G, D) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(x)} [\log(1 - D(G(z)))] \quad (1.1)$$

In a slightly different point of view, this process can be seen as minimizing the distance between distributions. In other words, the generator tries to approximate to the real latent distribution of images by mapping from a completely random distribution. During the training process, the Jhensen-Shannon distance is applied, measuring how far the approximated distribution is from the real one. As it is optimizing the models using gradient descent, this gradient information is back-propagated to the generator. Despite the fact that they mathematically demonstrate that there is an unique solution where D outputs 0.5 for every output and G recovers the latent training data distribution, these models are unstable during training making it laborious to train. The problem arises due to the unfair competition between generator and discriminator generating mode collapse problems, discriminators shielding infinity predictions and generators producing blank images or always producing the same sample [22]. Moreover, the basic algorithm is capable of generating up to 64x64 images but runs into instabilities if the size is increased. Resolution of the generated image is an important topic to address since most of the geographic and visual properties are better expressed in high-resolution so it can be used in remote-sensing applications.

Having presented the cornerstone and basics of GANs, multiple models and different variations and flavours came up providing novel techniques, loss functions, layers or applications. Particularly, some studies such as DCGAN [21] that immediately came after the original GANs paper added convolutional neural network layer (CNN) in order to increase stability of synthetic image generation. Despite it proved to generate larger images of 128 x 128 pixels, studies such as [5] report not to be sufficient due to insufficient details in satellite images. They also include similar analysis to [6] about the input latent space, demonstrating that generators are capable of disentangling latent space dimensions by mapping particular dimensions to particular features of the generated images. Advanced techniques such as [22] provide new methods for training such as feature matching included in the loss, changing the objective of the loss function from maximizing the discriminator output, to reducing the distance between intermediate features maps of the discriminator extracted from real images and generator images. By doing this the generator is forced to generate

samples that produce the same feature maps in the discriminator as the real images, similar to perceptual losses [30]. Also, they further analyse the problem of mode collapse by proposing many strategies such as mini batch discriminator where the discriminator has information from other images included in the batch, they also propose historical averaging that adds weight to the costs and they even suggest a semi-supervised technique that trains the discriminator with labeled and unlabeled data.

Progressive Growing GAN (PGGAN) [13] propose a method that gradually trains the generator and the discriminator until they are capable of producing large resolution images of 512x512 and 1024x1024. Their method starts by training the generator on images of 4x4 pixels, and gradually adding new layers to double the generated resolution until it is capable of generating high-res images. In addition, they propose a couple of techniques that further stabilize the training and provide variation such as minibatch standard deviation layer at the end of discriminator helping it to compute statistics of the batch, they propose a weight initialization and a scaling factor during runtime, and inspired by [2] they implement a Wasserstein gradient penalty as loss function. They propose a novel metric called Sliced Wasserstein Distance (SWD) that allows to perform a multi scale statistical similarity between distributions of local real and fake images patches drawn from a Laplacian pyramid, providing a granular quantitative results at different scales of the generated image.

In addition to the generation of large images, researchers propose novel architectures for more complex applications such as image-to-image translation, mapping from an image to an output image (conditioned generation). Pix2Pix [11] and Pix2PixHD [28] are among the first to address both problems: the image-to image translation and high-resolution generation. [11] proposes PatchGAN discriminator that is highly involved in posterior GAN research. PatchGAN discriminator is applied in patches at different scales and then its outputs are averaged to produce one scalar. In combination with L1 loss that captures low-frequency information, this model, which uses fewer parameters, focuses on the high frequencies contained in each patch. Its successor, Pix2PixHD [28], is able to produce images up to 2048x1024 pixels by novel multi-scale generator and discriminator, and retaking ideas of [22] by adding perceptual pre-trained loss. Similar to [13], they divide the training in what they refer to as a coarse-to-fine generator. This generator G is divided into two U-Net models, global generator G_1 and local enhancer G_2 . First G_1 is trained in order to learn global characteristics at 1024x512 scale. In the second phase G_2 is added with the particularity that the encoder part is added at the beginning of G_1 and the decoder part is added at the end, leaving the G_1 in the middle. In this case D is divided into three PatchGAN that operate at different scales. The image is downsampled in order to generate a pyramid of three scales. Then each D_i operates at different scales with different receptive fields, the coarse scale with large receptive field leads to global consisting images while the finer scale leads to finer details. The final contribution is the instance level feature embedding, a mechanism to control the generation. First, they train an encoder to find a low-dimension feature vector that corresponds to a real image. Then, they train G and D with this vector and the instance map as the conditional input. After a K-means analysis to find the cluster descriptor of each feature, the user is able to control the generation in coordination of the interpretation that the G is assigned to each dimension.

CycleGAN [31] proposes a model that learns to translate an image from one source

domain to a target domain, distressing the necessity of having two paired source and target datasets. This is done by adding an inverse mapping model in the loss that reverts the first transformation applied to the input, called cycle consistency. Additionally they reuse PatchGAN [11] as a discriminator. They conclude that by applying the cycle loop in addition to PatchGAN they are able to reach higher image sizes. PSGAN Progressive Structured GAN [7] is a work that adds conditionally to PGGAN. Their network is able to generate high-resolution anime characters by providing the skeleton structure of the character as an input. They take up the progressing growth by imposing the skeleton map at different scale levels while the generator and the discriminator are growing. StyleGAN [12] is GAN designed for style transfer purposes that can deal with higher resolutions and control the generation by learning high-level attributes and stochastic variations, allowing to control the style of synthesis. They use a progressive training in conjunction with Adaptive Instance normalization layers and Wassertein gradient penalty in addition to original GAN loss. This adapted generator learns a latent space domain and how to control features at different scales. Perceptual Adversarial network, PAN, [27] is a general framework that is also capable of performing high-resolution image-to-image translation. Their proposal also relies on feature matching of the D encouraging the generated images to have similar high-level features as the real ones while at the same time they use the output of D as the classical GAN loss.

Finally, we describe SPADE [20], a model that generates photorealistic imagery given a semantic map. They propose a spatially adaptive denormalization module (SPADE module), a conditional normalization layer that uses the input segmentation map to modulate the activation of the normalization layer at different scales of the generation. They demonstrate that batch normalization layers drown the signal, so they de-normalize the signal at each scale level by using SPADE layers. These layers consist simply of a convolutional layer that extracts the features of the input map and then learn by other two convolutional layers the scaling parameter at each spatial position the scale and bias according to the input map structure. By the addition of this simple modulation and residual blocks, they obtain consistent local semantic image translations that outperform previous models such as pix2pixHD and at the same time they remove the necessity of using an encoder-decoder network. They also comment that taking a progressive growing approach makes no difference in their technique. As a discriminator they reuse the multi-scale PatchGAN [28] with the last term replaced by Hinge loss.

In the field of remote sensing, there are not many studies focused general on image-to-image translation using GAN. In [5] they described the process of applying PGAN to synthetically generate satellite images of rivers and the necessity of high-resolution image generation for remote sensing applications that can capture particular high-frequency details of this kind of image that we have mentioned at the beginning of this work. Most of the work that uses GAN to remote sensing applications is done for cloud removal [25] or super resolution applications with GAN [29] and without GAN [23] that put special emphasis in the usage of dense skip or residual connections to propagate high-frequency signals that is particularly present on this kind of images. Works such as [24] evaluated models trained with synthetic images and demonstrate the improvement of adding them, but they do not delve into synthetic image generation techniques.

At the moment of this work, there are no vast formal studies specifically applied

to the image-to-image translation of generating satellite images conditioned into the segmentation map. Despite there are works that conduct similar tasks [31] [11], they rely on general translating satellite footprint to real images as an usage example rather than conducting a complete study of these challenging tasks. It is important to remark that there are a couple of companies, such as *OneView.ai*². These companies base their entire business model on providing synthetic image generation service for enriching training datasets by including in their pipeline their own developed GAN model to generate synthetic images from small datasets.

1.3 Problem Formulation

Before going deeper with more complex concepts and ideas, we first provide a high-level introduction about the principal ideas around this work. Let's start by considering we have $C = [0, \dots, K + 1]$ that represents K possible classes and 0 for the background. Let $m \in L^{H \times W}$ be a segmentation map, a matrix where in each position (x, y) contains a $k \in L$ the index of a class and H and W is the height and width of the image. Let $s = (v : r)$ a the $V + R$ -dimensional semantic global vector. Where each dimension of the first V -dimensional represents a proportion of one of V semantic global classes. The remaining R -dimensional vector is a categorical (one-hot encoding) vector that represents the categorical class of the region. In this way each scene is represented by a matrix M and a vector s . We present a deep neural network G that is capable of generating a satellite image I by receiving as an input M and s . Each pixel position (i, j) of the resulting I corresponds to the label of position (x, y) of m . Particularly, in this work we simplify the problem by choosing one class segmentation map despite it could be easily adapted to more classes. We decided M to be a building footprint map due to dataset availability and it was more than enough to validate the model and demonstrate the simplicity of generation. For the first V -dimensional part of the semantic global vector, we defined carefully 11 classes that express the amount of visual cues, land use and styles relative to classes such as forest, industrial, road, etc (in following sections we will explain more in detail). We selected 4 cities, with remarked style, cultural and geographic properties for the second categorical R -dimensional part of the vector. We ended up with a model that given a binary M mask with the shape and position of the buildings and a global semantic vector s that defines content related to style such as amount of roads, forest, industrial land use zone, etc and the city/region, it is capable of generating a satellite image that contains all the stylish visual cues in addition to buildings with the exact same position and shape as defined by the mask. Having this control mechanism a user can define its own segmentation mask, or even can modify the region or the amount of semantic class for a same mask helping it to efficiently augment a dataset with varied region/culture synthetic satellite imagery.

Finally, the model consists of a generator G of a GAN that is modified from a SPADE model [20] and a discriminator D Figure. 1.2. Mask m and vector s are passed to generator G for generating a synthetic scene to fool the discriminator that is responsible to discern between synthetic images and real ones.

²<https://one-view.ai/>

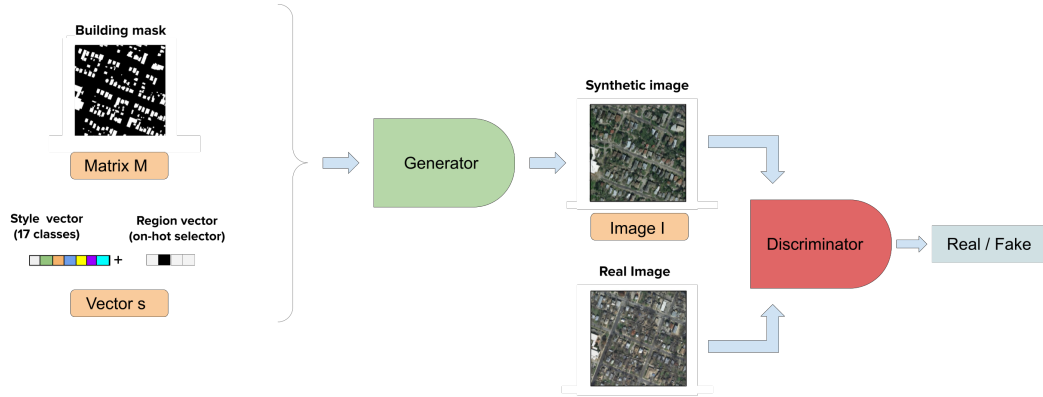


FIGURE 1.2: SPADE high level diagram. The generator takes the building footprint mask (m matrix) and the semantic global vector. It generates the synthetic image I and it is passed to the discriminator to determine if it is fake or real.

1.4 Research Questions

The objective of this work is to propose a simple mechanism that leverages the information of public geographic databases to enhance the geographic-properties of a synthetic image generated via a GAN model. While defining this mechanism, we wanted to evaluate if that enhancement would help to enrich satellite synthetic generation with finer details and properties using a simple representation such as a 17-dimensional vector. Therefore, the main research question of this thesis is:

How can a GAN model be modified to accept rich style satellite specific properties while at the same time this information comes in a small-dimensional representation ?

Subsequent questions may also be responded:

- How to leverage public annotation resources such as Open Street Maps to provide style information ?
- How to define visual distinct land cover properties ?
- Is the prior knowledge of region and style improving expressiveness of the GAN model ?

Chapter 2

Datasets

In this section we will talk about datasets and data sources we used for training the GAN model and for the development of the semantic global vector descriptor.

2.1 Inria Aerial Image Labeling Dataset

Inria Aerial Image Labeling Dataset (Inria) [16] is a high-resolution dataset designed for pixel-wise building segmentation (Figure 2.1). It consists of high-resolution objectified color imagery with spatial resolution of 0.3m/pixel that covers 810km² of 5 cities (in the training dataset):

- Vienna, Austria
- Lienz, Austria
- Chicago, USA
- Kitsap county of Washington, USA
- Austin, Texas, USA

Segmentation maps are binary images where a 1 in position (i, j) means that pixel belongs to a building and 0 is that it belongs to the background class. This dataset became of our interest that besides containing the structure segmentation map of buildings, its images cover a large variety of dissimilar urban and not urban areas with different types of population, culture, urbanisation, ranging from highly urbanized Austin, Texas to the rural Tyrol region in Austria. The dataset was designed with the objective of evaluating the generalization capabilities of training in a region and extending it to images with varying illumination, urban landscape and time of the year. As we were interested only in the labeled images, we discarded the test set, and focused on the before mentioned cities. In consequence our dataset consists of 45 images of 3000x3000 pixels.

2.2 Open Street Map (OSM)

OpenStreetMap (OSM) [19], created in 2005, is an open and collaborative database that provides geodata and geo-annotations. Basically, it consists of a free editable map of the world that allows its more than two million users to annotate or to provide collected data to enrich the OSM geo-information database. Its data primarily consist of annotations at multiple semantic levels that are expressed in keys (categories) and values. Under each key they provide finer grained information in different formats depending on the object of the annotation. For example, they provide

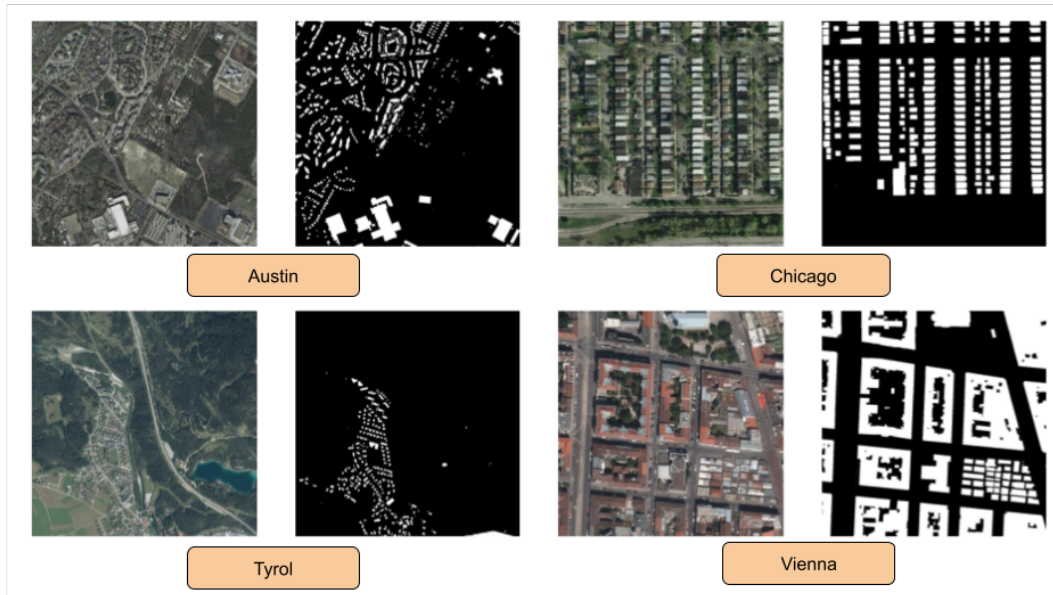


FIGURE 2.1: Inria building datasets sample [15]

annotations of land use that describe the human usage of an area as a polygon in a geojson. Another example is the annotation of roads, they structure the road network as a graph. There are many ways of accessing its data such as an API, or dedicated public or private geo-servers that digest and renderize the data. In our case we decided to use a public open source server that renders and compiles all the interested information for a specific area.

Therefore, we decided to download the render for each of the images using a rasterized tile server *rasterized tile server*¹ that provides cartographic style guidelines (Figure 3.5). As we have the source code of the server, we have the mapping between pixel colour and category. We ended up listing more than 200 categories present in the render and we were capable of reducing it to only 11 classes for the global semantic vector. We will explain this procedure more in detail in the following section.

¹<https://github.com/gravitystorm/openstreetmap-carto/>

Chapter 3

Methods

This section explains the methods used in this study. We will start from a more detailed analysis of the baseline model SPADE [20]. Then, we will delineate the proposed architecture modifications in order to develop SSSGAN. Next, we describe the creation of the global semantic vector. Finally, we present the metrics we used for evaluation.

3.1 SPADE

As previously explained, SPADE[20] proposed a conditional GAN architecture capable of generating high-resolution photorealistic images from a semantic segmentation map. They stated that generally image-to-image GANS receive the input at the beginning of the network, and consecutive convolutions and normalizations tend to wash away semantic and structural information, producing blurry and not aligned images. They propose to modulate the signal of the segmentation map at different scales of the network, producing better fidelity and alignment with the input layouts. In the following subsections we will explain different key contributions of the proposed model.

Spatially-Adaptive Denormalization

Spatially-Adaptive Layer is the novel contribution of this work. They demonstrated that spatial semantic information is washed away due to sequences of convolutions and batch normalization layers [10]. In order to avoid this, they propose to add these SPADE blocks that denormalize the signal in function of the semantic map input, helping to preserve semantic spatial awareness such as semantic style and shape. Let $m \in \mathbb{L}^{H \times W}$ be the segmentation mask whereas H and W is the height and width, and \mathbb{L} is a set of labels that refers to each class. Let h^i be the activation of i -th layer of a CNN. Let C^i , H^i and W^i the channels, height and width of the i -th layer. Assuming, that the batch normalization layer is applied channel wise, and obtain μ_c^i and σ_c^i for each channel $c \in C^i$ and i -th layer. The SPADE layer denormalization operation could be expressed as follows, if we consider $y \in H^i$, $x \in W^i$ and $n \in N$ be the batch size:

$$\gamma_{c,y,x}^i(m) \frac{h_{n,c,y,x}^i - \mu_c^i}{\sigma_c^i} + \beta_{c,y,x}^i(m) \quad (3.1)$$

Where μ_c^i and σ_c^i are the batch normalization parameter computed channel-wis for the batch N :

$$\mu_c^i = \frac{1}{NH^iW^i} \sum_{n,y,x} h_{n,c,y,x}^i \quad (3.2)$$

$$\sigma_c^i = \frac{1}{NH^iW^i} \sum_{n,y,x} (h_{n,c,y,x}^i - (\mu_c^i)^2) \quad (3.3)$$

The role of the SPADE layer is to learn the scale $\gamma_{c,y,x}^i(m)$ and bias term $\beta_{c,y,x}^i(m)$ with respect to the mask m , what they call modulation parameters Fig 3.1. What is interesting to put special emphasis on is that modulation parameters depend on the location (x, y) , thus it provides spatial awareness. This spatial awareness is what it differentiates this modulation with respect to batch normalization that does not consider spatial information. Those modulation parameters are expressed as a functional, because the SPADE layer passes m through a series of two convolutional layers in order to learn these spatially aware parameters. The structure of layers can be seen in Figure 3.1




FIGURE 3.1: (a) SPADE block internal architecture. (b) SPADE Residual block (SPADE ResBlk)

Having defined the SPADE block, the authors reformulate the common generator architecture that uses encoder-decoder architectures [11] and [28]. They remove the encoder layer since the mask is not fed in the beginning of the architecture. They decided to downsample the segmentation at different scales, and fed them via SPADE blocks after each batch normalization. Then they divided the network into four upscaling segments, where the last one generates an image with the size of the mask. Each segment that defines a scale level is composed of convolutional and upscaling layers followed by SPADE residual blocks. Each SPADE residual block consists of two consecutive blocks of SPADE layers (that ingest segmentation masks that have the same dimensions as the assigned to the SPADE residual block), followed by RELU activation layer and a 3x3 convolution Figure. 3.2. By this way, they removed the encoder, and ingested information of the shape and structure of the map at each scale, obtaining a lightweight generator with fewer parameters.

As a discriminator, they decided to use Pix2PixHD multiscale PatchGAN discriminator [28]. The task of differentiating high-resolution real images from fake ones represents a special challenge for D , since it needs to have a large receptive field that would increase network complexity. To address this problem, they used three identical PatchGAN discriminators at three different scales (factor of 2) D_1, D_2

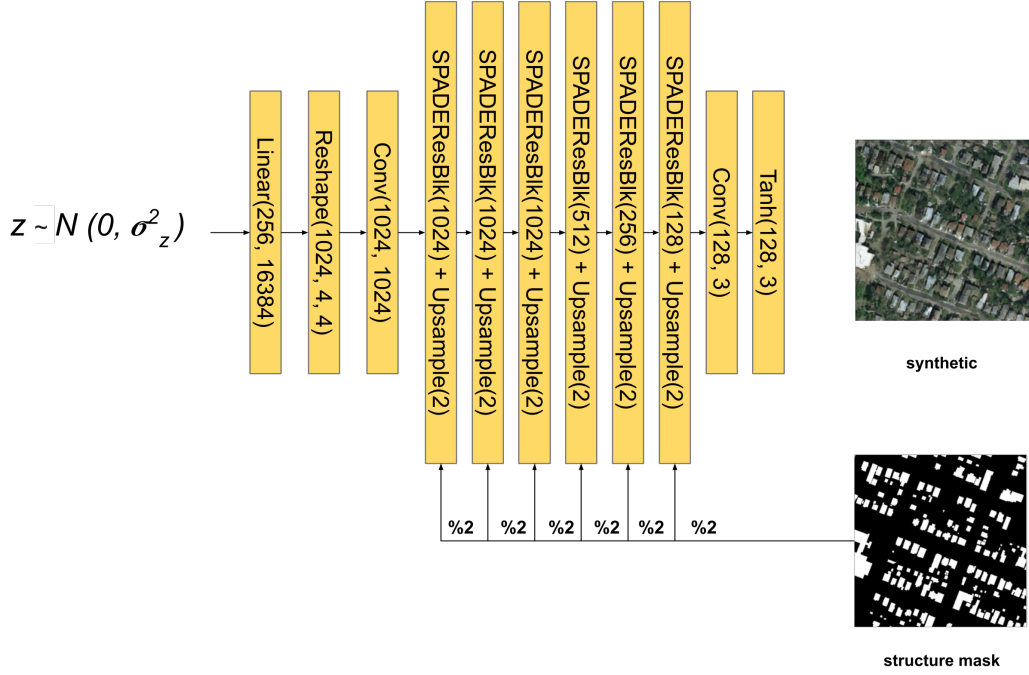


FIGURE 3.2: SPADE main architecture

and D_3 . The one that operates at the coarsest scale has larger global knowledge of the image while the one that operates at the finest scale forces the generator to produce finer details. Hence the loss function is the following, where k refers to the index of the three different scales:

$$\min_G \max_D L(G, D) = \sum_{k=1,2,3} L_{\text{GAN}}(G, D_k) \quad (3.4)$$

Another particularity is that they did not use the classical GAN loss function. Instead they used the least squared loss [17] term modification in addition to Hinge loss [18] that demonstrated to provide a more stable training, and avoiding vanishing gradient problems provided by the usage of the logistic function. Therefore, their adapted loss function is shown as follows:

$$L_D(G, D) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\min(0, -1 + D(x))] + \mathbb{E}_{z \sim p(z)} [\min(0, -1 - D(G(z)))] \quad (3.5)$$

$$L_G(G, D) = \mathbb{E}_{z \sim p(z)} [D(G(z))] \quad (3.6)$$

Additionally they used feature matching loss functions that we will not use in our experiments.

Finally, PatchGAN [11] is the lightweight discriminator network that is used at each scale figure 3.3. It was developed with the idea that the discriminator focuses on high-frequency details while L_1 focuses on low-frequencies. In consequence, they restricted the discriminator to look at particular $N \times N$ patches to decide if it is real or fake. Consequently, the discriminator is convolved through the image, by averaging its prediction of each $N \times N$ patch into a single scalar. This allows the discriminator to have fewer parameters and focus on granular details and composition of the generated image. In fewer words this discriminator is a simple ensemble of lightweight

discriminators that reduce the input to a unique output that defines the probability of being real or fake. The authors interpreted this loss as a texture/style loss.




FIGURE 3.3: PatchGAN diagram [4]. The entire discriminator is applied to $N \times N$ patches. Then, the model is convolved over the image and average their results in order to obtain a single scalar

3.2 SSSGAN

Having studied the principal component of SPADE in detail, we were able to spot its weak points for being used in our study. The key idea of SPADE is to provide spatial semantic modulation through the SPADE layers. That property is useful to guarantee spatial consistency in the synthesis related to the structural segmentation map, that in our case is the building footprint, it does not apply to the global semantic vector. Our objective is to ingest global style parameters through the easy-to-generate global semantic vector, that allows the user to define the presence of semantic classes while avoiding the necessity of generating a mask with particular location of these classes. As the semantic vector does not have spatial applicability it can not be concatenated neither fed through the SPADE layer. On the other hand, we can think of this vector as a human-interpretable and already disentangled latent space. Hence we force the network to adapt this vector as latent space.

We replace the latent random space generator of the SPADE model for a sequence of layers that receives the global semantic vector as an input (Figure 3.4). In order to ingest this information, we first generate the global vector by concatenating the first V classes, the 17 visual classes and the one hot encoding vector that defines the region or area (R -dimensional). The vector goes through three consecutive multi layer perceptron (MLP) blocks of 256, 1024 and 16384 neurons followed by an activation function. The resulting activations are reshaped to an $1024 \times 4 \times 4$ activation volume tensor. That volume is passed to a convolutional layer and a batch normalization layer. The output is then passed to a SPADE layer that modulates this global style information with respect to the structure map. [20] suggest that the style information tends to be washed away as the network goes deeper. In consequence, we decided to add skip connections between each of the scale blocks in channel-wise concatenation, similar to DenseNet [9]. In this way each scale block can receive the collective knowledge of previous stages, allowing the flow of the original style information. At the same time it allows to divide information in the way the SPADE block can focus on high-frequency spatial details, extremely important in aerial images, while the skip branch allows the flow of style and low-frequency information [23] [23]. In addition it is added reduction blocks (colored in green in Figure 3.4) that reduce the channel dimension, that is increased by the concatenation. This helps to stack more layers for the dense connections without a significant increment of memory. Thus it

is extremely important to add those layers. Besides all of that, this structure helps to establish the training process because the dense connections also allow the gradient to be easily propagated to the lower layers, even allowing deeper networks structures. This dense connection is applied by passing the volume input of each scale block with the output volume of the SPADE layer block. As the concatenation increases the channel (hence the complexity of the model) a 1×1 convolution layer is applied to reduce the volume.

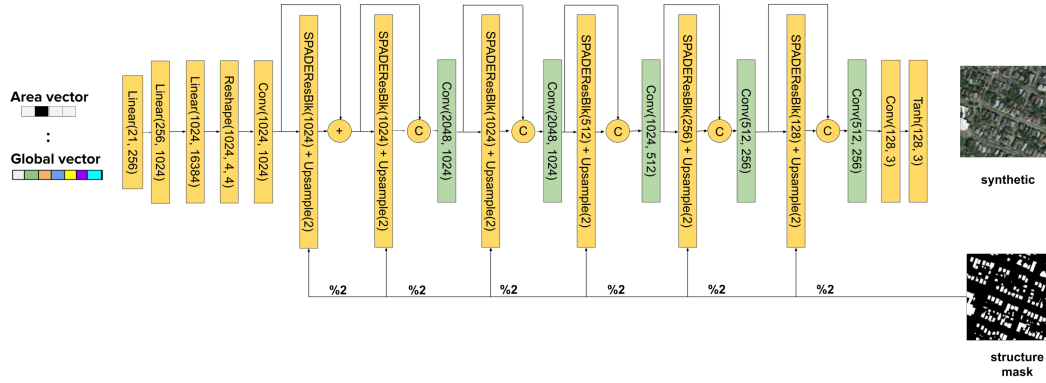


FIGURE 3.4: SSSGAN architecture. Every Linear layer is followed implicitly by Relu activation. The structure mask is downsampled to its half resolution each time it is fed to a SPADEResBlock (referenced by ‘%2’ label)

3.3 Global Semantic Vector

In this section we describe the creation of the global semantic vector. The key idea is to obtain a semantic description of the image that may help the generator to distinguish and pay attention to key properties present on the satellite image and it also helps to modify the image generation. In order to obtain a description, the principal idea of this thesis is that it can be easily generated from the OSM tags. This crowd-source tagged map is available publicly and it offers tags named by category for land use, roads, places, services. First, we download tags related to the area of interest: Chicago, Vienna, Austin, Kitsap and Tyrol. These tags come in multiple formats, for example, land use is defined by polygons while roads are defined as a graph. We obtained more than 150 values so we decided to rasterize these tags and then to define the value that corresponds to each pixel. After that process, we analysed the results and we found different problems regarding the labels. The first problem is that urban zones were more densely and finely detailed tagged than urban zones. For example, Vienna had much more detail in tags that even individual trees were tagged (Figure 3.5 (a)). While in the Tyrol region there were zones that were not even tagged. The second and more important problem was that there was no homogeneous definition of one tag in the same region or image. For example, in Chicago there were zones tagged as residential, while at the other side of the road which has the same visual appearance it was tagged as land (Figure 3.5 (b)). Moreover, we noticed that all images of Kitsap were not annotated at all, there were roads and residential zones that were missed (Figure 3.5 (c)). Finally, we come up with similar

conclusions to [1], a work that only used land use information. Labels refer to human activities that are performed in specific zones. Those activities sometimes may be expressed with different visual characteristics at ground level, but from the aerial point of view those zones do not contain visual representative features. The clear example is the distinction between commercial and retail. The official definition in OSM is ambiguous, commercial refers to areas for commercial purposes, while retail is for zones where there are shops. Besides this ambiguity in definition, both areas express buildings with flat grey roofs in the aerial perspective.



FIGURE 3.5: (a) Detailed annotation of the urban area of Vienna (b) Residential area from Chicago, from one side of the annotated road is defined as residential and the other side is not annotated, despite both belongs to the same visual residential cues (c) Area of Kitsap without annotation.

Having studied in detail all of these problems, we decided to perform a manual inspection of the data and we defined a series of conventions that help to aid the previously mentioned problems. The principal idea is to create a vector in an automatic way that digests all semantic visual information so it facilitates the model to put attention on particular visual characteristics. For that reason we decided to group all of these categories in 17 classes that have a clear visual representation despite the use. In that case, classes such as commercial and retail will constitute the same class, since while we were doing the manual visual inspection we decided that those classes are visually indistinguishable. We manually corrected zones that were not labeled and we defined a unique label for ambiguous zones, fixing the problem with residential and land labeled zones. Finally, we decided to remove images from the Kitsap region from the dataset due to the scarcity in label information.

At the end of this process we ended with the 17 classes expressed in Table 3.1. In order to compute the vector, we defined an index or position to each class in the vector. Having this grouping rule of classes we processed each image by counting the amount of pixels that belong to each class (taking into account the priority of the class) and then we normalized the vector to sum 1, obtaining a distribution of classes.

More conclusions were obtained from this analysis that also coincide with the ones expressed in [1]. A specific land use such as residential or commercial varies in visual characteristics from region to region due to architectural and cultural factors. In order to help the network to distinguish these cultural properties, and at the same time controlling the generation, we added to this vector a one hot encoding selector that defines the region: Chicago, Austin, Vienna or Tyrol.

3.4 Metrics

We decided to employ two state-of-the-art perceptual metrics used in [20] [13]. Since there is no ground truth, the quality of generated images is difficult to evaluate.

TABLE 3.1: CNN Performance Long Format

Semantic category	OSM tag	Index
grass	grass, heath, golf_course, farmland, ...	1
forest	forest, forest-text, orchard, scrub, ...	2
residential	residential, residential-line, land-color, ...	3
commercial	commercial, commercial-line, retail, ...	4
industrial	industrial, industrial-line, wastewater_plant, ...	5
parking	garages, parking, ...	6
construction	construction, construction_2, built-up-z12, quarry, ...	7
sports	pitch	8
highway	motorway, primary, secondary, ...	9
rail	motorway, primary, secondary, ...	10
road	living_street, residential, ...	11
footway	footway, pedestrian, ...	12
religious	cemetery, religious, ...	13
motorway	motorway, trunk, ...	14
water	water, ...	15
allotments	allotments, ...	16
block	block, ...	17

Perceptual metrics try to provide a quantitative answer of how close the generator managed to understand and reproduce the target distribution of real images. The following metrics provide a scalar that represents distance between distributions, and indirectly they are accessing how perceptually close are the generated images to the real ones.

Frechet Inception

Frechet Inception Distance (FID) [8] [20] is commonly used in GAN works for measuring their image quality generation. It is a distance that measures the distance between synthetically generated images and the real distribution. Its value refers to how similar two sets of images are in terms of vector features extracted by Inception V3 model [26] trained for classification. Each image is passed through the Inception V3, and the last pooling layer prior to the output classification is extracted obtaining a feature vector of 2048 activations. These vectors are summarized as a multivariate Gaussian, computing the mean and covariance of each dimension for each image in each group. Hence, a multivariate Gaussian is obtained for each group, real and synthetic images. The resulting Frechet distance between these Gaussian distributions is the resulting score for FID. A lower score means that the two distributions are close, the generator has managed to efficiently emulate the latent real distribution.

Sliced Wasserstein

Sliced Wasserstein Distance (SWD) is an efficient approximation to the earth mover distance between distributions. Briefly speaking, despite being computationally inefficient, earth mover distance provides the vertical distance difference between distribution, giving an idea of the differences between densities. [13] comment that metrics such as MS-SSIM is useful for detecting coarse errors such as mode collapse, but fails to detect fine-detailed variations in color and textures. Consequently, they

propose to build a Laplacian pyramid for each of the real and generated images, from 16x16 pixel and doubling resolution until the pyramid reaches the original dimensions. Basically each level of the pyramid is a downsampled version of the upper level. This pyramid was constructed having in mind that a perfect generator will synthesize similar images structures at different scales. Then, they select 16384 images for each distribution and extract 128 patches of 7x7 with 3 RGB channels (descriptors) for each Laplacian level. This process ends up with 2.1 M of descriptors for each distribution. Each patch is normalized with respect to each color channel mean and standard deviation. After that Sliced Wasserstein Distance is applied to both sets, real and generated. Lowering the distance, means that patches between both distributions are statistically similar.

Therefore, this metric provides a granular quality description at each scale. Patches at 16x16 similarity indicate if the sets are similar in large-scale structures, while larger scale provides more information of finer details, color or textures similarities and pixel-level properties.

Chapter 4

Results

In this section we show quantitative and qualitative results using INRIA dataset along with our global semantic vector descriptor. We start in section 4.1 by describing the set up of the experiment. In section 4.2 we show the quantitative results by performing a simple ablation study. Finally, in section 4.3 we present some qualitative results, by showing how a change in the global vector changes the style of synthesised images.

4.1 Implementation details

Original SPADE was trained on an NVIDIA DGX1 with eight 32GB V100 GPUs [20]. In our case, we train our network with our network with eight NVIDIA 1080Ti of 11 GB each one. This difference in terms of computational resources made us reduce the batch down to 24 images, instead of 96. Usually training with larger batch sizes should help stabilize the training and produce better results. Regardless of this aspect, we show our approach is able to improve the generation's expressiveness and variety with respect to the baseline. Altogether, while changing the style and domain of the generated images.

We applied a learning rate of 0.0002 to both, the generator and the discriminator. We used ADAM optimizer with $\beta_1 = 0$ and $\beta_2 = 0.9$. Additionally we applied data a few data augmentation that consisted of simply random 180 and 90 rotations. Original images were cropped to 256x256 patches with an overlap of 128 pixels that provide more variability. We trained each network for the same amount of 50 epochs.

4.2 Quantitative Analysis

We trained the original SPADE implementation as a baseline, that we used for referencing any quantitative improvement provided by our proposal. Then, we decided to evaluate our main architecture by using only the global semantic vector. Finally, we conducted the full approach that uses the global semantic vector and the dense connections scheme. We applied our two aforementioned metrics, Frechet (FID) Inception Metric and Sliced Wasserstein distance (SWD) for obtaining quantitative results. Table 4.1 shows the comparison results between different versions of the model. By a great margin, we can appreciate that the full implementation of SSSGAN, that uses the complete global semantic vector, outperforms the original baseline. The model could reduce by more than a half almost all the metrics. The reduction of the FID from 53.1909 to 22.358 suggests that the generation was more close to estimate the latent distribution of the real images than the original baseline in general and global features. Moreover, provides a more granular and detailed

perspective about the generator performance at different scales. SSSGAN was able to reduce by a 56% at the original scale, an impressive 76.5% at 128x128 scale, a 67.6% at a 64x64 scale, a 64.3% at a 32x32 scale and a 45.8% at 16x16 scale the SWD score. Our hypothesis is that by forcing the generator to understand the already disentangled space for humans, we are providing more prior knowledge about the real distribution of the real images. During the training process the generator can assign a correlation between the presence of particular features of the image and an increment of the global vector value. In this way the generator could produce more variable synthetic images generations and it could capture finer details structures at different scales. The generator not only reduces each metric, it could reach a constant performance almost in every scale, by learning how to generate closer to reality scale specific features.

Intermediate results that use only the semantic vector suggest that approach provides variability to the image generation. Even though the absence of dense connections reduced considerably every score, the signal of the style that is fed into the beginning of the networks gets washed out by consistently activations modulation performed by SPADE blocks, that modulates activation only with respect to the structure of the buildings. The addition of dense connections before the modulation helps to propagate the style signal efficiently to each of the scales.

TABLE 4.1: Performance different SSSGAN versions with respect to SPADE baseline

model	FID	SWD-256	SWD-128	SWD-64	SWD-32	SWD-16
baseline	53.1909	338.5295	474.0834	486.1723	474.0834	665.2925
semantic(ours)	38.1688	207.6403	206.2924	241.2935	232.0195	404.9646
semantic+dense(ours)	22.358	148.9343	111.8898	153.751	173.8637	355.164

TABLE 4.2: model comparison. baseline refers to the original SPADE implementation [20]. While "semantic" refers to the SSSGAN with only style vector and "semantic+dense" is the full model SSSGAN with global semantic vector and dense connection structure

4.3 Qualitative Analysis

In this section, we show a comparison between SPADE baseline, SSSGAN with only semantic information, and the full version of SSSGAN. From Figure 4.1 to Figure 4.8 we see those networks compared in addition to the segmentation building mask, the semantic map to provide an idea of the proportion of the semantic classes and the three most influential classes of the semantic global vector. Qualitatively speaking the full version of SPADE was able to perform a simple relation between shape of buildings with region in order to generate more consistent scenes. For example, when it is presented a structure mask with the characteristic shape of Vienna's building (Fig. 4.8), SSSGAN can infer from the shape of building besides the information of the global vector the region of the intended image and generate region-specific features of the region like tree shapes, illumination, and the characteristic orange roof. Most of the time SPADE generated flat surfaces with absence of fine details,

textures and illumination Fig.4.1, Fig.4.2, Fig.4.6 or Fig.4.7. Another aspect to remark related to the style of the region is that SSSGAN was able to remarkably capture Tyrol style images (Fig.4.6 and Fig.4.7) with large light green meadows, tress, illumination and roads.



FIGURE 4.1: Visual comparison of Austin area. Mask of building footprint and main semantic classes of the vector are shown as reference.



FIGURE 4.2: Visual comparison of Austin area. Mask of building footprint and main semantic classes of the vector are shown as reference.



FIGURE 4.3: Visual comparison of Austin area. Mask of building footprint and main semantic classes of the vector are shown as reference.

Generally speaking, SSSGAN demonstrated its vast ability to capture style and context related to each of the four regions. For example, in contrast to the baseline, SSSGAN was able to produce detailed grass style of Tyrol and differentiate subtle trees properties of Austin and Chicago. In general, visual inspection of the generated images suggest that SSSGAN was able to capture railway track, roads and even consistent generation of cars like in Fig. 4.1.

Another remarkable point is the consistent shadowing of the scenes, it can be appreciated in every scene that the network is able to generate consistent shadows among every salient feature such as trees or buildings. Finally, we can see networks have difficulties in generating long rectified lines. The reason is that building mask



FIGURE 4.4: Visual comparison of Chicago area. Mask of building footprint and main semantic classes of the vector are shown as reference.

contains imperfect annotated boundaries that the networks reproduce and the adversarial learning procedure does not detect and therefore do not know how to overcome.



FIGURE 4.5: Visual comparison of Chicago area. Mask of building footprint and main semantic classes of the vector are shown as reference.



FIGURE 4.6: Visual comparison of Tyrol area. Mask of building footprint and main semantic classes of the vector are shown as reference.




FIGURE 4.7: Visual comparison of Tyrol area. Mask of building footprint and main semantic classes of the vector are shown as reference.

Finally, we show the generation capabilities in Fig. 4.10, where we increment the presence of four categories while diminishing the others in a Chicago building footprint. At the same time, we show how this generation mechanism is expressed



FIGURE 4.8: Visual comparison of Vienna area. Mask of building footprint and main semantic classes of the vector are shown as reference.



FIGURE 4.9: Finer observation of the increment of grass class and forest class

in each region by changing the one hot encoded area vector. Efficiently, we see how each row contains a global style color palette related to the region. For example, the row of the Tyrol region in Fig. 4.10 presents a global greenish style that is common in that region, while the row of Chicago presents brownish and diminished colors. The increment of forest efficiently Fig. 4.9 increases the presence of trees while the increment of industrial category tends to generate grey flat roofs over the buildings. It is important to remark that the style of the semantic category is captured, despite it does not show enough realism due to incompatibilities of building shapes with this specific style. For instance, when increasing industrial over a mask of residential houses of Chicago, the network is able to detect buildings and provide them a grey tonality, but is not providing finer details to these roofs because it is not relating the shape and dimensions of that building with respect to the increased style. Nevertheless, we can efficiently corroborate changes in style and textures by manipulating the semantic global vector.



FIGURE 4.10: Original building footprint is from Chicago. Each row shows the generation for that Chicago footprint mask in different regions. First column uses the original global semantic vector. Second column the grass category is increased. Third column forest class is increased. Finally, the fourth column industrial class is increased.

Chapter 5

Conclusions

Global high resolution images with corresponding ground truth are difficult to obtain due to the infrastructure and cost required to acquire and label them, respectively. In order to overcome this issue, we present a novel method, SSSGAN, that integrates a mechanism capable of generating realistic satellite images, improving the semantic features generation, by leveraging publicly available crowd sourced data from OSM. These static annotations, that purely describe a scene, can be used to enhance satellite image generation by encoding it in the global semantic vector. We also demonstrate that the use of this vector, in addition to the architecture proposed in this work, permits SSSGAN to effectively increase the expressiveness capabilities of the GAN model. In the first place, we manage to outperform the SPADE model in terms of FID and SWD metrics, meaning the generator was able to better approximate the latent real distribution of real images. By evaluating the SWD metric at multiple scales, we further show the consistent increment in terms of diversity at different scale levels of the generation, from fine to coarse details. In the qualitative analysis, we perform a visual comparison between the baseline and our model, comparing the increment in diversity and region-culture styles. We finish our analysis by showing the effectiveness of manipulating the global semantic vector. This brings to light the vast potential of the proposed approach. We hope this work will encourage future synthetic satellite image generation studies that will help to a better understanding of our planet.

Chapter 6

Future Works

Finally, we will talk about different approaches that we had to exclude from this thesis due to time constraints and we will leave it for future work.

6.1 Improve losses

We took out from this work the usage of feature matching or structural losses. We talked that the addition of feature matching losses improves the image quality because it forces the generator to produce similar deep learning features with respect to the real images. There are some state-of-the-art differentiable perceptual losses like LPIPS [30] that can be further tuned for the domain of satellite imagery that can improve quality of generation. We also thought in creating our own specific loss based on deep learning detectors already trained on satellite imagery, that can translate its knowledge about dynamic objects present in the images (e.g cars) through gradient backpropagation. We believe this is one of the main improvements that were left to SSSGAN.

6.2 Augment capabilities

We think that further studies can be conducted related to the augmentation capabilities of this network. It would be interesting to understand how other networks behave by training with synthetic augmented images. For example, further studies could contemplate the performance change of a simple UNet baseline network trained to detect building footprints. The study should contemplate different mixture proportions of synthetic and real images, and also evaluate other strategies such as pre-training with only synthetic images and then fine tuning with real ones. More evaluations could be conducted regarding the style and region interpolation. It would be interesting to train a model with a diverse variety of regions, and if the region capabilities of this SSSGAN could help to improve its performance by ingesting domain altered synthetic images.

Bibliography

- [1] Adrian Albert, Jasleen Kaur, and Marta C Gonzalez. “Using convolutional networks and satellite imagery to identify patterns in urban environments at a large scale”. In: *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 2017, pp. 1357–1366.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. “Wasserstein generative adversarial networks”. In: *International conference on machine learning*. PMLR. 2017, pp. 214–223.
- [3] Jia Deng et al. “Imagenet: A large-scale hierarchical image database”. In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee. 2009, pp. 248–255.
- [4] Thittaporn Ganokratanaa, Supavadee Aramvith, and Nicu Sebe. “Unsupervised anomaly detection and localization based on deep spatiotemporal translation network”. In: *IEEE Access* 8 (2020), pp. 50312–50329.
- [5] Akshat Gautam, Muhammed Sit, and Ibrahim Demir. “Realistic River Image Synthesis using Deep Generative Adversarial Networks”. In: *arXiv preprint arXiv:2003.00826* (2020).
- [6] Ian J Goodfellow et al. “Generative adversarial networks”. In: *arXiv preprint arXiv:1406.2661* (2014).
- [7] Koichi Hamada et al. “Full-body high-resolution anime generation with progressive structure-conditional generative adversarial networks”. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. 2018, pp. 0–0.
- [8] Martin Heusel et al. “Gans trained by a two time-scale update rule converge to a local nash equilibrium”. In: *arXiv preprint arXiv:1706.08500* (2017).
- [9] Gao Huang et al. “Densely connected convolutional networks”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 4700–4708.
- [10] Sergey Ioffe and Christian Szegedy. “Batch normalization: Accelerating deep network training by reducing internal covariate shift”. In: *International conference on machine learning*. PMLR. 2015, pp. 448–456.
- [11] Phillip Isola et al. “Image-to-image translation with conditional adversarial networks”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 1125–1134.
- [12] Tero Karras, Samuli Laine, and Timo Aila. “A style-based generator architecture for generative adversarial networks”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 4401–4410.
- [13] Tero Karras et al. “Progressive growing of gans for improved quality, stability, and variation”. In: *arXiv preprint arXiv:1710.10196* (2017).

- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. "Imagenet classification with deep convolutional neural networks". In: *Advances in neural information processing systems* 25 (2012), pp. 1097–1105.
- [15] Jingjing Ma et al. "Building extraction of aerial images by a global and multi-scale encoder-decoder network". In: *Remote Sensing* 12.15 (2020), p. 2350.
- [16] Emmanuel Maggiori et al. "Can Semantic Labeling Methods Generalize to Any City? The Inria Aerial Image Labeling Benchmark". In: *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE. 2017.
- [17] Xudong Mao et al. "Least squares generative adversarial networks". In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 2794–2802.
- [18] Takeru Miyato et al. "Spectral normalization for generative adversarial networks". In: *arXiv preprint arXiv:1802.05957* (2018).
- [19] OpenStreetMap contributors. *Planet dump* retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org>. 2021.
- [20] Taesung Park et al. "Semantic image synthesis with spatially-adaptive normalization". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 2337–2346.
- [21] Alec Radford, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks". In: *arXiv preprint arXiv:1511.06434* (2015).
- [22] Tim Salimans et al. "Improved techniques for training gans". In: *arXiv preprint arXiv:1606.03498* (2016).
- [23] Francesco Salvetti et al. "Multi-image Super Resolution of Remotely Sensed Images using Residual Feature Attention Deep Neural Networks". In: *arXiv preprint arXiv:2007.03107* (2020).
- [24] Jacob Shermeyer et al. "Rareplanes: Synthetic data takes flight". In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2021, pp. 207–217.
- [25] Praveer Singh and Nikos Komodakis. "Cloud-gan: Cloud removal for sentinel-2 imagery using a cyclic consistent generative adversarial networks". In: *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*. IEEE. 2018, pp. 1772–1775.
- [26] Christian Szegedy et al. "Rethinking the inception architecture for computer vision". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 2818–2826.
- [27] Chaoyue Wang et al. "Perceptual adversarial networks for image-to-image transformation". In: *IEEE Transactions on Image Processing* 27.8 (2018), pp. 4066–4079.
- [28] Ting-Chun Wang et al. "High-resolution image synthesis and semantic manipulation with conditional gans". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 8798–8807.
- [29] Zhongyuan Wang et al. "Ultra-dense GAN for satellite imagery super-resolution". In: *Neurocomputing* 398 (2020), pp. 328–337.
- [30] Richard Zhang et al. "The unreasonable effectiveness of deep features as a perceptual metric". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, pp. 586–595.

-
- [31] Jun-Yan Zhu et al. “Unpaired image-to-image translation using cycle-consistent adversarial networks”. In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 2223–2232.