



**Politecnico  
di Torino**

**Politecnico di Torino**

Master's Degree in Biomedical Engineering

Master's Degree Thesis

**HISPACE: Histological Image  
Synthesis with Pattern And Control  
Engine**

**Candidate:**

Giulio CHIOSSO

**Supervisors:**

Prof. Massimo SALVI

Ing. Alen SHAHINI

Academic Year 2024/25



# Contents

<b>Abstract</b>	<b>2</b>
<b>1 Introduction</b>	<b>3</b>
<b>2 Generative models in digital pathology</b>	<b>5</b>
2.1 Current State-of-the-Art Techniques . . . . .	5
2.1.1 Conditional GAN Approaches . . . . .	7
2.1.2 Diffusion Models . . . . .	10
2.1.3 Other Generative Techniques . . . . .	10
<b>3 Methods</b>	<b>12</b>
3.1 Dataset . . . . .	12
3.1.1 MoNuSAC Dataset . . . . .	12
3.1.2 Steatosis Dataset . . . . .	13
3.2 Data preparation and feature extraction . . . . .	14
3.2.1 Coordinate and Centroid Processing . . . . .	14
3.2.2 Fourier Descriptors . . . . .	17
3.2.3 Feature Extraction and Clustering . . . . .	17
3.2.4 First Module Pipeline . . . . .	21
3.3 Generative Model . . . . .	22
3.3.1 Design Choices . . . . .	23
3.3.2 Generation Algorithm . . . . .	26

<b>4</b>	<b>Results</b>	<b>30</b>
4.1	Validation of the semantic content . . . . .	30
4.2	Image Synthesis from Semantic Masks . . . . .	33
4.3	Performances improvement in a segmentation task . . . . .	35
4.4	Expert-Based Qualitative evaluations . . . . .	36
<b>5</b>	<b>Conclusions</b>	<b>38</b>
5.1	Limits . . . . .	39
5.2	Future Applications . . . . .	40
	<b>Bibliography</b>	<b>42</b>

# Abstract

The growing integration of Digital Pathology (DP) within histological analysis has opened new perspectives in biomedical imaging and data-driven research. Whole Slide Imaging (WSI) technologies enable large-scale digitalization of tissue samples, allowing visualization, storage, and computational processing. However, the development of reliable data-driven methods remains limited by the lack of annotated dataset and especially due to the high inter-operator variability in the labelling process performed by pathologists. To address these challenges, this thesis proposes to invert the traditional paradigm by introducing a controlled generation framework for synthetic histological images with user-defined semantic content. The proposed system called HISPACE (Histological Image Synthesis with Pattern And Content Engine) is based on a hybrid pipeline, where semantic content generation is driven by explicit mathematical modelling, while the subsequent translation into photorealistic images is achieved through a GAN model. In the first stage, each cell contour is reconstructed using Fourier Descriptors (FDs), ensuring morphological coherence with real samples, while the overall tissue layout is modelled through Kernel Density Estimation (KDE) to reproduce realistic spatial patterns derived from real data distributions. The model allows direct user control over key semantic parameters – the number of cells, tissue density, and spatial configuration – enabling the creation of realistic and customizable histological images. The pipeline was developed and validated using two annotated datasets, MoNuSAC and a private dataset derived from liver WSI samples. The semantic generator was shown to produce realistic and cluster-consistent tissue layouts, while the integration with the SENSE model enabled the synthesis of high-quality histological images, as confirmed

by expert evaluation. The addition of HISPACE-generated samples improved the performance of a steatosis segmentation network, particularly in scenarios with extensive steatosis, which are poorly represented in standard training data. Overall, the results highlight HISPACE as a flexible and biologically plausible framework for controlled histological image synthesis with potential applications in computational modelling and pathology training.

# Chapter 1

## Introduction

Histology is a fundamental branch of biology which deals with microscopic analysis of biological tissues. The study of cellular organization and tissue architecture provides essential knowledge to describe pathological functioning. In a medical context, it assumes a central role regarding the recognition of pathological mutation allowing to improve diagnostic accuracy, therapeutic decision making and evaluation of surgical outcomes. Histopathology uses specialized techniques to study tissue alterations such as staining. This allows to highlight the tissue components including nucleus, cytoplasm, stroma, and more that otherwise could not have been seen under a microscope. The most widely used stain in medical diagnosis is hematoxylin-eosin stain (H&E) where hematoxylin stains cell nuclei a purplish blue and eosin stains the extracellular matrix and cytoplasm pink.

Nowadays histology has been progressively combined with Digital Pathology (DP), which merges computer tools and advanced imaging techniques with tissue analysis. Using high-resolution scanners, histological samples are converted into Whole Slide Images (WSI), granting visualization, storage and sharing of slides in digital format. This approach also leads to new opportunities with applications of image analysis algorithms and Machine Learning(ML)-based approaches.

In the histological field, the availability of annotated data represents a significant

challenge: whereas for some pathologies or tissues relatively large datasets exist; for others the material available is limited or poorly balanced. In addition, traditional data augmentation techniques generate images lacking realism or, conversely, images that are too similar to real ones. Hence, the need of developing new strategies in order to create synthetic data that not only improve samples heterogeneity but also maintain the authenticity of the biological information.

The target of this study is to develop an algorithm able of generating synthetic images with user-defined semantic content. This approach aims to combine precision and adaptability, operating non-deterministically while ensuring a high level of control. In the final work the user may select different cell classes, in order to create customized layout usable for experimenting or training processes. HISPACE can be used to create annotated datasets for training and validating ML-methods but also to train future pathologists, considering the intrinsic availability of the synthetic ground truth.



## Chapter 2

# Generative models in digital pathology

With the development of DP and so the ability to scan histological slides using high-resolution scanners to achieve WSI, it also starts to expand the Computational Pathology (CPath) that is a data-driven advanced application of digital pathology, applying automated analysis algorithms to extract quantitative information from the digital slides. The purpose is not only to support diagnosis, but also to lead to inferences of new biomarkers [1], to predict therapeutic outcomes [2] or to correlate histological data with genomic data [3]. However, as already specified, deep-learning models require datasets that are both very large and fully annotated and so generative models for synthetic images have become a useful and functional tool in histopathology. In addition, synthetic images can be employed for data augmentation, while simultaneously helping to mitigate privacy concerns and overcome ethical or legal barriers [4], [5].

### 2.1 Current State-of-the-Art Techniques

The generation of synthetic histological images through deep learning methods began with the introduction of Generative Adversarial Networks (GAN). GANs, initially proposed by Goodfellow et al. in 2014 [6], were rapidly adopted in digital pathology as a tool for the creation of realistic images, with applications ranging from color normalization [7], to segmentation tasks [8], up to data augmentation [9].

GANs works via an adversarial process, where two models are trained simultaneously: the generative model (G) that captures the data distribution, and the discriminative model (D) that assess the probability that a sample came from the training data rather than G. Assuming a latent variable  $z$  is extracted from a distribution  $p_z(z)$ . The generator  $G(z; \theta_g)$ , typically a deep neural network, converts this noise in a sample into the data space, approximating the real distribution  $p_{\text{data}}(z)$ . The discriminator  $D(x, \theta_d)$ , also a neural network, returns the probability that  $x$  belongs to the real dataset rather than being produced by G. Training consists of a min-max optimization problem, where D tries to maximize the ability to discriminate real and generated data and G tries to deceive it, reducing that difference:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))] \quad (2.1)$$

The opportunity of obtaining qualitatively convincing images, allows GANs to become the reference technique for studies about synthetic tissue generation [10]. In the following years, the original architecture has varied its specifications, creating models such as CycleGAN or StyleGAN, that allowed to extend practical applications and further improve image quality.

However, recently the focus has shifted towards Diffusion Models, introduced by Ho et al. in 2020 [11]. These models quickly gained popularity in the histological field because of their stability during training and their ability to generate samples close to the real distribution.

Diffusion models are probabilistic generative models based on the Markov chain. The model is defined in two opposite phases: the forward or diffusion process, where starting from real input, noise is gradually added to the data until signal is destroyed, resulting in a distribution that can be easily sampled; and the reverse process, where a neural network is trained to estimate the reversed transitions of Markov chain, in order to gradually remove noise and reconstruct the realistic data. Training is performed by optimizing the usual variational bound on negative log likelihood, which takes the

form:

$$\mathcal{L}(\theta) = \mathbb{E}_{q(x_{0:T})} \left[ -\log \frac{p_{\theta}(x_{0:T})}{q(x_{1:T} | x_0)} \right] \quad (2.2)$$

where  $q(x_{1:T} | x_0)$  defines the forward process and  $p_{\theta}(x_{0:T})$  represents the reverse process.

The most promising aspect of diffusion models, compared to GANs, lies in their greater training stability. Unlike GANs, they are not affected by mode collapse and can therefore explore the data distribution more efficiently [12]. This results in the generation of high-quality images with realistic textures and fine details.

### 2.1.1 Conditional GAN Approaches

While the original formulation of GANs enables the generation of realistic samples by learning the underlying data distribution, it does not provide explicit control over the properties of the synthesized images. In histopathology this lack of control represents a significant limit: it is usually necessary to produce images with specific morphological patterns, belonging to certain diagnostic classes or paired with annotations useful for segmentation tasks. Conditional GAN (cGAN) improve the standard model introducing an additional information that guide the generative process. In addition of the random latent vector, the generator and the discriminator receive a conditioning signal that can assume different forms: class labels, semantic masks or structured layout defined by the user. This approach allows to constrain the synthesis, generating images consistent with the input conditions.

#### Label-conditioned Synthesis

These generative models constrain the synthesis process through a global reference label, such as tissue type or pathology. This approach does not support explicit control over cellular distribution or histological architecture, generating images consistent with morphological characteristic of a given class, proving remarkably useful in data augmentation.

A representative example is the work of Falahkheirkhah et al. [13]. They presented a class-conditional GAN for the synthesis of prostate and colon histological images. The generator architecture combines residual blocks with conditional normalization and progressive upsampling steps, while the discriminator is based on pix2pixHD [14]. The training uses a combination of GAN loss and perceptual loss, ensuring the preservation of both visual quality and morphological coherence.

### **Layout-Conditioned Synthesis**

The generation process here is driven by an explicit object-level description of the histological architecture. The conditional input provided is a structural layout specifying position, shape and type of histological entities (cells, nuclei, glands and more). This concept grants a fine control of the tissue composition, generating conditions that are rare and difficult to find in real datasets.

SynCLay [15] framework introduces this type of conditioning and allows to construct realistic images from custom cellular layouts. The user defines the position and type of cells in a two-dimensional plane. The model learns to translate these layouts into realistic images of size 256x256 pixels, implementing direct control over the cellular distribution. Architecturally, SynCLay integrates a GAN network with specialized modules for nuclear segmentation and classification. The framework employs a pre-trained HoVer-Net [16] to simultaneously generate, together with the synthetic images, also the corresponding nuclear masks, improving morphologic consistency and supplying automatic annotations.

Also Deshpande et al. [17] proposed a framework able to generate realistic images of colorectal tissue, together with their respective glandular mask, starting from a gland layout defined by the user. Doing so it is possible to control fundamental parameters such as number, position and dimension of the gland. The generation process occurs in two main phases: firstly, the layout is translated into a synthetic glandular mask using a neural network that construct and combine individual masks for each gland, preserving the morphological variability by injecting gaussian noise; secondly the mask

is given as conditioning for an encoder-decoder network with skip-connections inspired from pix2pix that creates the realistic image. The training is adversarial and operates through three different discriminators: one for the mask, one for the image and the last specifically on the gland.

### **Mask-Conditioned Synthesis**

The generator receives as input a semantic map in which each pixel is labeled with a specific class. Rather than describing only the location and outline of the structures, the mask encodes a full-resolution map of tissue composition, capturing fine boundaries and subtle topology. The generator learns to translate this pixel-wise annotation into a photorealistic histological image, enabling accurate reproduction of complex microarchitecture and facilitating tasks such as segmentation to image translation or simulation of rare pathological patterns.

SENSE [18] is a framework designed to generate realistic histological images with precise control over cellular distribution. Once the synthetic semantic maps are created, the generator operates through a hybrid ViT-Pix2Pix to translate them into high-resolution images and the discriminator employs Patch Classifier to evaluate the authenticity of synthesized outputs.

DEPAS [1] introduces a mask generator trained from random noise using a DC-GAN architecture. The model generates high-resolution semantic maps that reproduce the plausible microarchitecture of various tissues, including skin, prostate, and lung. Using pix2pixHD, these maps are converted into histological images, enabling the control of cellular characteristics and type of staining.

Also PriorPath [19] fits into this class of models, here the approach is to apply a coarse-to-fine process in generating masks: starting from raw ones, obtained through morphological operations or manually drawn by pathologist, this model creates fine and detailed masks, closer to the real distribution of data. This structure allows to preserve the control over the layout and to overcome the problem of mode collapse, typical of GANs. PriorPath employs pix2pix and CycleGAN to translate the masks from coarse

to fine and uses pix2pixHD to generate the histological RGB images. Training takes place in an adversarial manner, utilizing multiscale discriminators to ensure global and local coherence.

### **2.1.2 Diffusion Models**

A first example of application of diffusion models in the histopathological field is the work of Moghadam et al. [20]. It investigates the ability of these type of architectures to generate realistic images of brain tumors. The model relies on a UNet denoiser enhanced with multi-resolution attention mechanisms that recognize both microscopic and macroscopic patterns. To improve image quality and avoid typical reconstruction artifacts, residual blocks derived from BigGAN are employed during downsample and upsample. Additional information is also provided as input through dedicated embedding: one for timestep for the neural network and one for the genotype. Training is carried out with a combination of score matching losses to effectively learn the inverse denoising process.

### **2.1.3 Other Generative Techniques**

Besides the most widespread approaches, that are GAN and Diffusion Models, literature reports alternative ones which despite being less popular, have found specific applications in histopathology. These include both probabilistic models such as Deep Convolutional Gaussian Mixture Models (DCGMM), and parametric models, based on mathematical and heuristic description of tissue morphology. Even if the photorealistic quality of generated images is usually lower compared to that of the aforementioned models, these approaches are particularly useful in application scenarios such as targeted data augmentation, creation of synthetic benchmarks with perfect annotations or study of specific morphological patterns.

Deep Convolutional Gaussian Mixture Model (DCGMM) combines in a single framework the potentials of convolutional neural networks and the probabilistic flexibility of Gaussian Mixture Models (GMM). Traditional GMMs describe data dis-

tribution as a combination of multiple multivariate gaussians and are usually employed in histopathology to normalize the H&E staining. DCGMM [21] differs by performing the estimation of the responsibility coefficients using a CNN fully-convolutional instead of the E-step in an EM-based optimization. By doing so, it is possible to extract structural information related to tissue morphology. Optimization occurs through end-to-end learning, maximizing the log-likelihood of the model using gradient descent. This approach shows considerable potential in specific application, such as stain normalization in histological dataset, a crucial step to ensure the generalizability of automated analysis algorithms.

An alternative way to achieve image synthesis are parametric models, which describe the biological tissue using a set of mathematical rules and distributions, explicitly defined starting from empirical observations on real samples. The generation is not driven by machine learning of hidden patterns, but by biologically and morphologically significant parameters estimated from experimental data and then used to reconstruct virtual structures. An example is proposed by Kovacheva et al. [22]. They developed a parametric model capable of synthetically reproducing the colon microenvironment in both healthy and pathological scenarios. The model is based on the explicit description of fundamental tissue structures such as crypts, lumen and goblet cells, each determined from parameters calculated on real H&E images. The model is based on the explicit description of fundamental tissue structures: colon crypts are simulated with shape and distribution that vary depending on the level of tumor differentiation, then the number of cells to be inserted is determined. The modelling of lumen and goblet cells employs procedural synthesis techniques and Voronoi diagrams to reproduce the observed distribution, and finally, parameters regarding the shape and texture of each cell are included. In doing so, there is an intrinsic explicit control over cellular density, overlap level, tumor grade and image resolution, granting flexibility and interpretability. In addition, the simulation of measurement errors contributes to make data even more similar to the real images, allowing the model to be useful as benchmark for segmentation, grading and normalization algorithms.

# Chapter 3

## Methods

### 3.1 Dataset

#### 3.1.1 MoNuSAC Dataset

The first dataset employed in this study is MoNuSAC (Multi-organ Nuclei Segmentation and Classification) [23], released for the MoNuSAC 2020 Challenge. It is a collection of histological images acquired from different organs and clinical centers, with the aim of providing a reference benchmark for the segmentation and classification of cell nuclei. The images are derived from H&E stained sections and include a wide heterogeneity in terms of tissues, pathological conditions, and laboratory variability, making the dataset particularly representative of real challenges in the histopathological domain. MoNuSAC consists of manual annotations at the single nucleus level, provided by experts, which allows not only to segment the nuclei but also to distinguish them into four main classes: epithelial cells, macrophages, lymphocytes, and neutrophils. Each image is therefore accompanied by multilabel segmentation masks that identify the position, shape, and type of cell uniquely. For this work, MoNuSAC was mainly employed for the design and model validation phase, as it offers a broad and diverse basis for testing preprocessing procedures, morphological and spatial descriptor extraction and clustering strategies. The size of the dataset, combined with the quality of the



annotations, makes it an ideal tool to justify and motivate the methodological choices adopted in the early stages of HISPACE, before moving to the use of the second dataset, specifically intended for the generation and evaluation of synthetic images.

### **3.1.2 Steatosis Dataset**

The second dataset used in this study consists of liver biopsy samples collected from a monocentric cohort at the *AOU Città della Salute e della Scienza di Torino Hospital* (Torino, Italy). The dataset includes consecutive liver needle biopsies obtained either at organ retrieval or at the end of the transplantation procedure. All specimens were processed following standard routine protocols of the Pathology Unit, yielding a total of 292 high-quality samples. The original H&E-stained glass slides were digitized to produce WSIs. Each digital slide then underwent stain normalization, a preprocessing step applied to reduce inter-samples color variability and improve the consistency of downstream image analysis. Semantic information is extracted using an automatic segmentation network specifically designed to detect steatosis. The liver biopsy dataset is primarily used in the evaluation phase of HISPACE, serving as an independent benchmark to assess both semantic-controlled content generation and translation into realistic histological images.

## 3.2 Data preparation and feature extraction

The data preparation stage constitutes the first functional block of the pipeline. Starting from the original dataset of already segmented and annotated masks, HISPACE operates on two complementary workflows:

- **Cellular shape description:** For each cell, preserving the class distinction, the normalized contour is extracted and then a compact representation is calculated using Fourier coefficients. These descriptors allow for an accurate reconstruction of single-cell morphology during the generative phase.
- **Spatial pattern interpretation:** At the image level, features are computed to describe the relative arrangement of cell centroids, once again separated by class. These features capture the spatial organization of tissue and are used as input for an unsupervised clustering process via Agglomerative Clustering [24] to identify different types of cellular patterns.

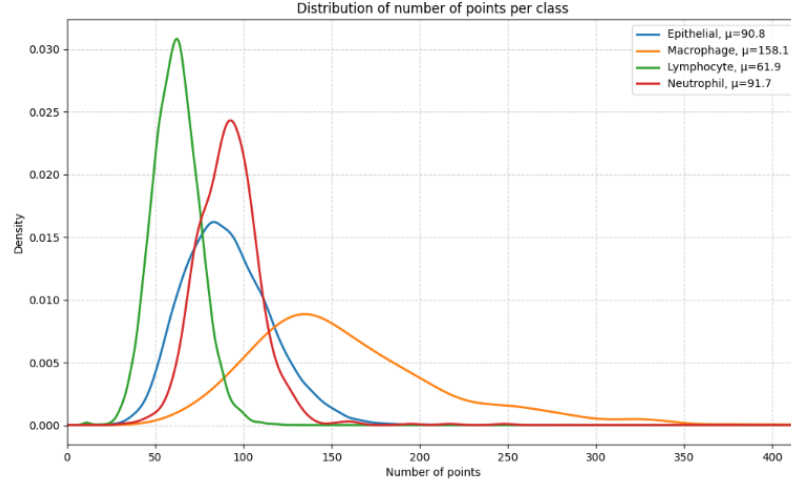
This double representation - morphological and spatial - provides a solid foundation for both statistical analysis and subsequent controlled generation of semantic masks. The user in fact will be able to specify the number of cells to be placed and the desired pattern for each class; the generator will reconstruct the single shapes starting from Fourier descriptors and place the cells accordingly to the distribution learned from the clustering process, granting a non-deterministic component of variability.

### 3.2.1 Coordinate and Centroid Processing

The data processing pipeline is designed to ensure uniform discretization of cellular contours by resampling all shapes with a fixed number of points, and to extract robust cell centroids.

To define the number of points to be used in the discretization, a preliminary analysis has been conducted on the MoNuSAC dataset. For each cell class (Epithelial, Lymphocyte, Macrophage, and Neutrophil) all available contours are considered, and the number of points in each shape is analyzed. The distribution of the number of

points is reported in the following graph.



**Figure 3.1:** Distribution of the number of points across all classes in MoNuSAC dataset.

This analysis enlightened the extent of intra-class variability and provided mean and standard deviation of the number of points per class. The mean value has been chosen to resample all shapes of the class in question, this allows to properly preserve the morphological details of the contour and to keep computational costs manageable and uniform for all classes.

Cell class	N
Epithelial	91
Macrophage	158
Neutrophil	92
Lymphocyte	62

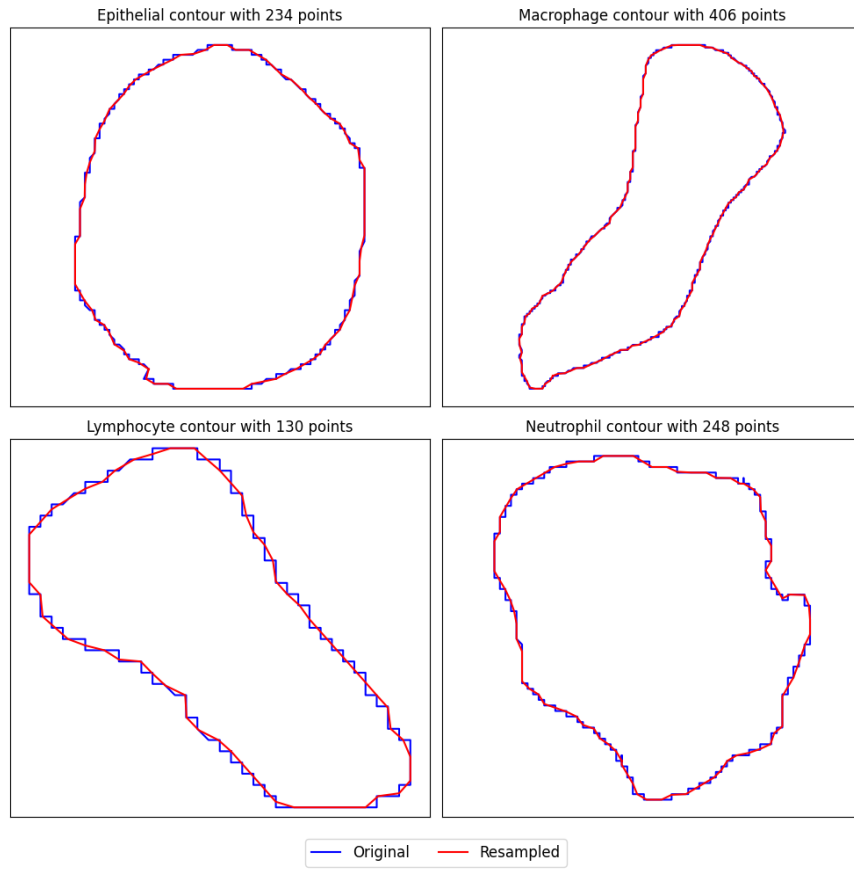
**Table 3.1:** Number of points selected for the resample (MoNuSAC).

Following the same procedure, the discretization analysis was repeated on the Steatosis dataset. The selected resampling values are summarized in Table 3.2. These values were adopted in all subsequent stages of the HISPACE pipeline and therefore constitute the effective configuration used to generate and evaluate the synthetic histological images presented in this work.

Cell class	N
Steatosis	46
Hepatocyte Nuclei	49
Other Nuclei	26
Artifact	31

**Table 3.2:** Number of points selected for the resample (Steatosis Dataset).

Fig. 3.2 illustrates borderline cases used to assess the validity of resampling in preserving cellular shape for the MoNuSAC dataset. These cases correspond to the contours at the margins of the distribution for each class shown in Fig 3.1



**Figure 3.2:** Down-sample of cellular contours at the edges of the distributions.

Therefore, the resample operation accurately maintains the original geometry even when applied to contours defined by a very different number of points from those selected in the tables 3.1 and 3.2.

For each segmented cell, its geometric centroid was calculated starting from the bi-dimensional contour represented as an array of coordinates  $(x_i, y_i)$ . The shape is considered a closed polygon, and so its centroid is determined as the center of mass. Results are then exported as a JSON file, divided by class and associated with relative metadata, to be used during the extraction of spatial features.

### 3.2.2 Fourier Descriptors

Let the contour be a closed curve defined by  $N$  points, parametrized as a complex function  $z_n = x_n + iy_n$  where  $(x_n, y_n)$  are the coordinates of each point relative to the origin. The Discrete Fourier Transform (DFT), implemented in *SciPy*, decomposes a finite sequence of samples into a sum of sinusoidal components at different frequencies, each defined by a complex coefficient encoding both amplitude and phase. Those coefficients are computed as follows, given the sequence of  $N$  complex points  $z_0, z_1, \dots, z_{N-1}$ :

$$a_k = \frac{1}{N} \sum_{n=0}^{N-1} z_n e^{-i2\pi kn/N}, \quad k = 0, 1, \dots, N-1$$

The Fourier Descriptors (FDs)  $a_k$  provide a compact representation of the contour geometry in terms of spatial frequency components. The full set of coefficients is stored in a *NumPy* array and constitutes the basis for the subsequent stages of the pipeline, most notably the synthesis of shapes.

### 3.2.3 Feature Extraction and Clustering

To compactly characterize the spatial arrangement of cells within each image, a set of eight features was computed using the centroid coordinates alone. This strategy aims to describe not only the average density, but also the degree of irregularity and the possible presence of periodic patterns in the distribution.

To evaluate the local density, for each centroid the  $k = 5$  closest neighbors were identified and their relative distances calculated. Considering the mean of these distances, the **Mean\_knn\_dist** parameter was estimated. This value provides an immediate indication of how strongly cells tend to cluster: reduced mean distances indicate more compact regions, whereas higher values correspond to more rarefied tissue.

To account for spatial heterogeneity, the centroid coordinates were used to generate a two-dimensional Voronoi diagram. In this representation, each cell is associated with the region of the plane closest to the corresponding centroid. The variance of the cell areas, defined as **Voronoi\_var**, serves as an indicator of the regularity of the cellular arrangement: reduced values denote a generally uniform distribution, while higher values highlight irregularities or local aggregations.

Finally, the overall distribution of centroids was transformed into a continuous density map using kernel density estimation. From these maps, two types of information were extracted. The first is **Density\_entropy** which measures the level of uniformity: high entropy corresponds to a homogeneous arrangement, whereas lower values indicate marked concentrations. The second concerns the spectral structure of the distribution. By applying a two-dimensional Fourier transform to the density map, the five components with highest energy were identified and referred to as **FFT1**, **FFT2**, **FFT3**, **FFT4**, **FFT5**. These values reveal any periodicities or radial symmetries that are difficult to grasp with naked eye or other techniques.

At the end of the process, for each image an array with eight numerical descriptors was obtained. These values concisely yet informatively summarize the tissue geometry and are used in the morphological clustering stage to differentiate recurring cellular patterns that will be reproduced by the generator.

Feature	Description	Significance
<b>Mean_knn_dist</b>	Mean of euclidean distances to the nearest 5 neighbors	Local density: low = high density; high = low density
<b>Voronoi_var</b>	Variance of Voronoi regions areas	Uniformity of spatial arrangement: low = ordered, homogeneous tissue; high = spatial disorganization
<b>Density_entropy</b>	Shannon entropy calculated on KDE maps	Disorder: low = marked concentrations; high = uniform distribution
<b>FFT1-FFT5</b>	Spectral intensity of the dominant frequency components in the Fourier transform of the density map	FFT1: large-scale, slow patterns
		FFT2-3: intermediate patterns, local repetitiveness
		FFT4-5: fine patterns, micro-structures, small-scale regularity

**Table 3.3:** Features description

Once the eight features have been extracted for each image, the final step of this module consists in grouping cellular patterns through a clustering algorithm.

The Agglomerative Clustering method from the *sklearn.cluster* library was employed. It is a hierarchical technique that identifies latent structures in the data based on similarity criteria, without assuming specific shapes or prior distributions. For each class, the extracted features were first normalized using a StandardScaler and subsequently transformed via Uniform Manifold Approximation and Projection (UMAP). This non-linear transformation reduces the dimensionality of the feature space to five components while preserving local relationships between samples and facilitating the identification of more compact and interpretable morphological patterns.

Four experimental setups (Table 3.4) were defined, each specifying a different

number of clusters to be identified.

Setup	number of clusters
Setup1	3
Setup2	4
Setup3	5
Setup4	6

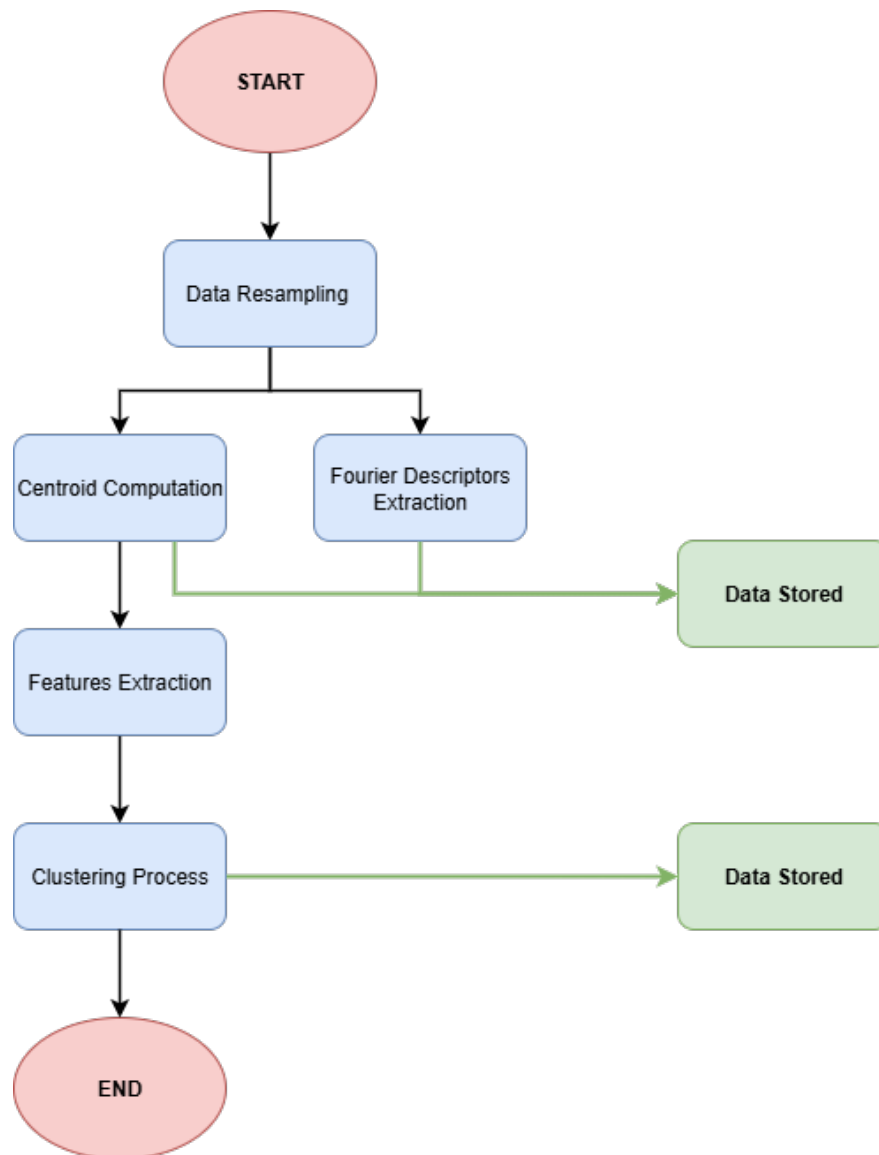
**Table 3.4:** Different setups for Agglomerative clustering

The results corresponding to each setups are stored, including both the image-level cluster assignments and the *silhouette score* used to evaluate cluster separability. In addition, a summary table reporting the relative cluster frequencies is provided to facilitate interpretation. The selection of the most suitable clustering setup is left to the user and depends on the specific objectives of the generative phase.



### 3.2.4 First Module Pipeline

To enhance the clarity of the described operations, Fig. 3.3 shows a flowchart that synthesized the data preparation and feature extraction phase. The diagram highlights the main steps of the pipeline, from the acquisition of cellular contours to the calculation of data useful for the generative program.



**Figure 3.3:** Flowchart for the first module pipeline.

### 3.3 Generative Model

The second module represents the actual generative model responsible for constructing synthetic images based on the information previously extracted. HISPACE aims to reproduce realistic cellular layouts while maintaining morphological and spatial coherence with real data, and allowing the user to control key parameters during the process. The user-defined parameters of the algorithm are specified for each class to be generated:

- Total number of cells
- Selected setup
- Selected cluster

This flexibility enables the generation of images that reflect different spatial layouts observed in real samples. As a result, it is possible to create various configurations such as healthy tissue, disorganized cellular structures or patterns associated with specific pathological conditions such as tumor alterations or steatosis. From an architectural perspective, the generative process is divided into two steps:

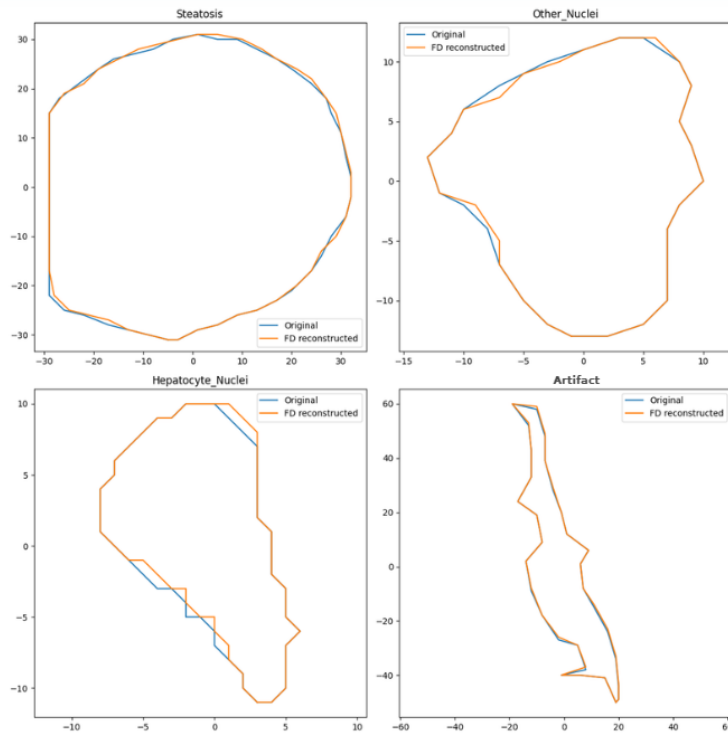
1. **Local morphology:** The reconstruction of each cell's contour is achieved through the Fourier Descriptors extracted in the previous module. This approach enables the generation of complex cellular shapes in a compact and controllable manner, ensuring both variability among shapes and consistency with real cells.
2. **Global topology:** The spatial distribution of cells in the resulting image is modeled using Kernel Density Estimation (KDE), computed on the centroids of real cells belonging to the selected cluster. This distribution acts as a probability map for generating new synthetic centroids, hence recreating spatial structures with characteristics similar to those observed in the original data.

### **3.3.1 Design Choices**

The design choices were made to ensure that the output images achieve an effective balance between biological realism, controllability, and computational robustness. The objective is to reproduce the characteristic features of real histological tissues while remaining as faithful as possible to the input parameters.

#### **Morphological representation**

The adoption of Fourier Descriptors for contour representation of each cell in the pipeline was motivated by two main reasons. First, FDs allow extremely simple scale management: the contour size can be continuously and precisely adjusted by multiplying the coefficients by a scalar factor, thereby obtaining enlarged or reduced versions of the same shape without the need to recalculate or interpolate the contour points. Second, FDs provide a position-independent representation in space: the shape description no longer depends on the absolute coordinates of the image, making the shapes more suitable for subsequent geometric comparisons or manipulations. In Fig. 3.4 are shown several examples of cell shapes belonging to different classes, displayed in their original form alongside their contour reconstructions obtained through FDs.



**Figure 3.4:** Contour reconstruction using FDs.

## Spatial Modelling


To regulate the spatial arrangement of cells on the image plane, Kernel Density Estimation (KDE) was employed. This choice stems from the need to reconstruct cellular layouts that are not completely random but instead reflect spatial patterns statistically consistent with those observed in real data. KDE provides a continuous density function derived from the discrete positions of cell centroids belonging to single image. The model can then generate a probability map describing regions of higher or lower cell concentration, which serves as a basis for the controlled distribution of new synthetic centroids. This approach offers several distinctive advantages:

- **Preservation of spatial structure:** maintains the characteristic spatial relationships of the selected cluster, thereby preserving the global tissue organization.
- **Statistical variability:** Enables the generation of new configurations that are statistically similar but not identical to real data, resulting in variable yet biologically plausible outcomes.

- Control over dispersion: Provides direct control over the degree of spatial dispersion through the kernel bandwidth, a parameter that regulates the compactness or spread of the cellular layout.

The employment of a continuous density estimate, rather than discrete models or simple spatial jittering, has proven effective in maintaining topological coherence between regions of high cell density, while avoiding exact repetition of the original structures.

The effectiveness of this representation is illustrated by the density maps obtained for different clusters of the *Steatosis* cell class, shown in Fig. 3.5.



**Figure 3.5:** Comparison between heatmaps for different clusters

### Positioning constraints and geometric validation

To drive the placement of synthetic cells within the canvas, spatial constraints implemented in the *placing\_cell* function have been imposed; this ensures that the generative process avoids producing artificial or physically implausible layouts. In order for a new cell to be inserted, the reconstructed shape via FDs is placed at the location identified by the KDE distribution and then validated through a set of acceptance conditions, three main criteria are evaluated:

- Geometric contour validity: configurations exhibiting zero-area or degenerate contours are automatically discarded, ensuring the numerical consistency of the

representation.

- Tissue membership: if a tissue mask is available, each cell must be entirely contained within the biologically valid region. Actual coverage is assessed as the ratio between the total cell area and the portion that falls into the tissue mask. Placement is accepted if at least 95% of the total area is included in the mask.
- Absence of overlap: to prevent interference between adjacent cells, the intersection between the new cell and the already occupied area is computed. If the overlap exceeds 5%, the insertion is rejected and a new placement attempt is performed.

Once all these conditions are satisfied, the cell is inserted into the canvas and stained according to its corresponding class. By adopting this progressive validation strategy, a high level of spatial realism is guaranteed, excluding cells that are excessively overlapping or out of bounds. Overall, the model exploits these constraints to balance the flexibility introduced by probabilistic generation of the layout with the geometric rigor required to preserve the biological integrity of the synthesized tissue.

### 3.3.2 Generation Algorithm

The generative model is implemented as a modular pipeline, where each component sequentially contributes to the construction of the synthetic image. Every module performs a specific task, ranging from data loading to the final rendering phase.

The *Load\_cell\_data* function, operating on a class-by-class basis, opens and reads all the supporting files required by the generative process: the JSON file containing the centroids coordinates grouped by image, the Excel file with cluster allocation, and the NPY file storing the FDs related to the cellular contours. These data depend on user-defined input parameters, which are described in Chapter 3.3. In addition to loading the data, this function is responsible for selecting a reference cellular layout on which the target density heatmap is computed. This map is subsequently used as a probability mask for the positioning of new synthetic centroids. The selection of the reference

image is performed in a controlled yet non-deterministic manner: the chosen layout must belong to the user-defined combination of cluster and setup and must contain a number of cells between 80% and 140% of the target value. This tolerance range introduces an element of controlled variability that prevents the repetition of identical spatial patterns across different runs with the same input parameters.

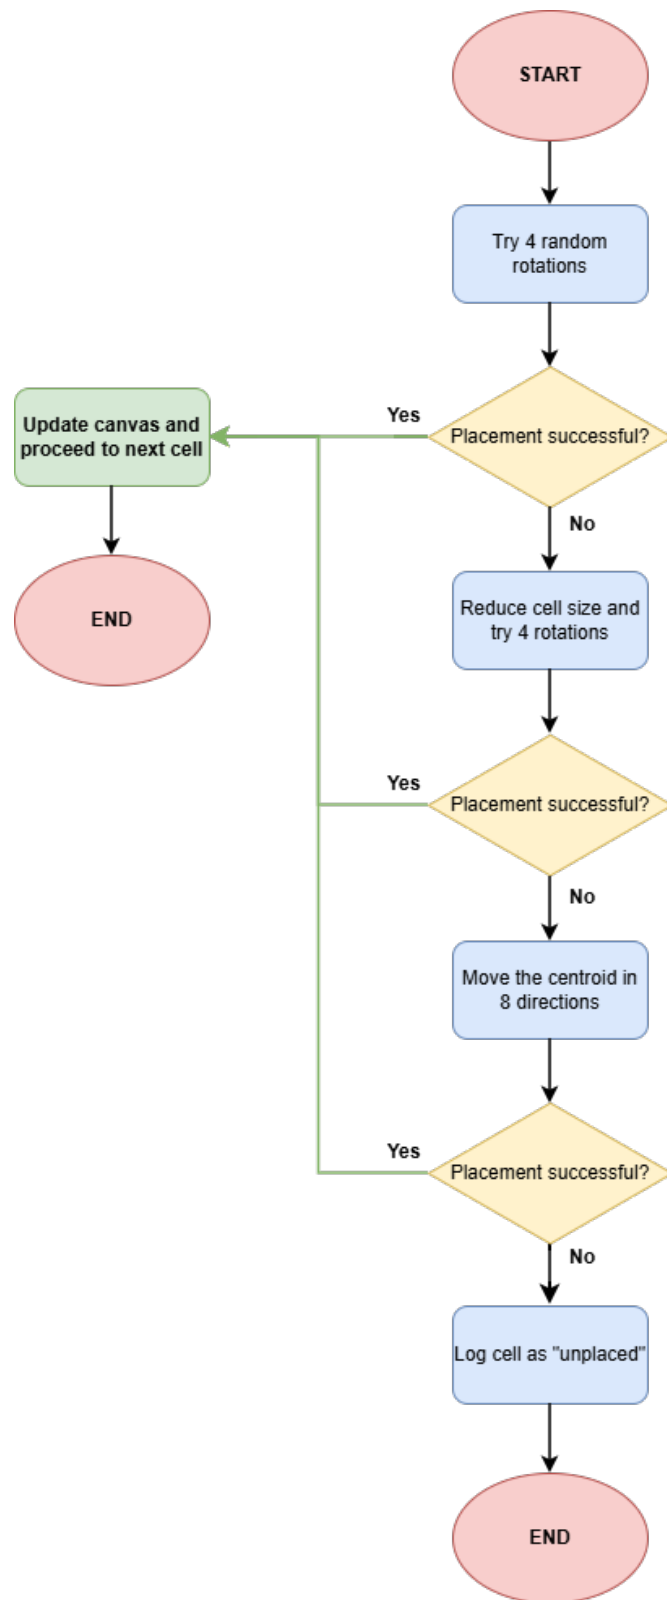
The progressive insertion of cells onto the target canvas is managed by the *Placing\_cell* function. This function receives as input the contour of the cell to be placed, the coordinates of its centroid, the main canvas, which is updated after each insertion, and the occupancy mask, which keeps track of the regions already covered. Once the contour is translated to the target position, the algorithm checks for overlaps with previously inserted cells. Specifically, it computes the percentage of shared area between the new cell and existing ones. If this value exceeds 5% of the total area of the new contour, the placement is rejected and the function returns negative; otherwise, the cell is accepted and drawn on the canvas colored according to its class, the occupancy mask is updated, and the function returns a positive result.

Finally, the *Generate\_cell\_layouts* function integrates all the previously described modules to output a semantic map with user-defined characteristics. The process begins by filling the entire canvas with a neutral gray tone, which serves as the background color used to identify the tissue. Next, the function calls *Load\_cell\_data* to read all supporting files and user selections, and to select the reference layout used as mask. A density probability map is then computed in the form of a heatmap indicating the most probable positions of cell centroids. Through a brief optimization cycle, the algorithm tries to extract the best set of points by minimizing the loss with respect to the target layout. At the end of this stage, the model has access to the complete set of cell positions to be used in the subsequent insertion phase. During the placement, the algorithm operates sequentially, placing one cell at a time. For each cell a contour is randomly extracted from the file containing the FDs, reconstructed and passed to the *Placing\_cell* function for insertion. Four different rotation angles are first tested. If the placement succeeds, the process moves to the next cell; otherwise, the algorithm

continues by applying a set of descending scale factors, specifically [0.8, 0.6, 0.4, 0.2], to progressively reduce the contour size and attempt insertion again while respecting the overlap constraint. This resizing step is governed by the previously extracted parameters related to the minimum area thresholds defined for each class. A resized contour is accepted only if its area remains greater than the 5th percentile threshold for its class. In case of failure, the function repeats the placement attempt by testing multiple orientations again in search of a viable position. If all the previous attempts fail, the algorithm applies a final strategy by shifting the centroid used for placement with an increasing offset in each of the eight principal directions. If after this stage, the cell still cannot be positioned, the program logs the event and notifies the user that the specific cell could not be inserted. The canvas is then saved as a PNG file with an incremental identifier within the folder "Testing". Metadata are also embedded into the final PNG image, allowing users to retrieve class-specific information, such as which reference image was used to compute the KDE density map, how many cells were requested, and how many were successfully inserted.

To facilitate understanding of the logic adopted by the model during the cell insertion phase, Fig. 3.6 provides a graphical summary of the operational flow implemented within the *Generate\_cell\_layout* function. The diagram highlights the logical decision-making mechanisms that govern the sequential placement of each cell onto the canvas.





**Figure 3.6:** Placing\_cell pipeline

# Chapter 4

## Results

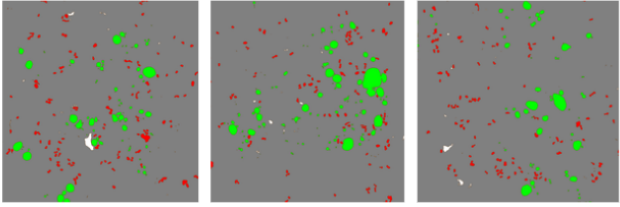
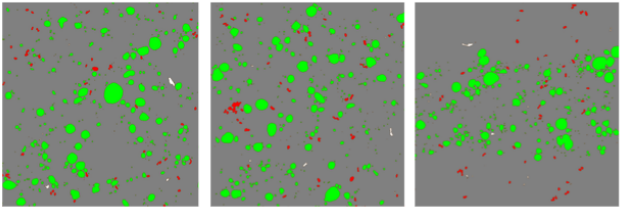
This chapter presents the results obtained with the HISPACE model, with the aim of evaluating three key aspects: (i) the model’s ability to generate controllable and coherent semantic masks, (ii) the quality and realism of the synthetic images produced by the GAN from the generated semantic content assessed by a pathologist, and (iii) the practical utility of these synthetic images, assessed through clinical evaluation and data-augmentation experiments in a steatosis segmentation task. These results allow to assess both the methodological soundness of the model and its potential applicability in pathological and computational contexts.

### 4.1 Validation of the semantic content

The first evaluation focuses on validating the semantic content generated by HISPACE, verifying whether the produced masks satisfy the two key properties that the architecture was designed to guarantee: non-deterministic variability and cluster-dependent spatial topology. Both aspects are essential to ensure that the synthetic masks are realistic, biologically plausible, and sufficiently diverse for downstream applications such as image synthesis and data augmentation.

## Variability and non-determinism

To assess the generative variability of the semantic masks, multiple samples were generated using the same configuration range, allowing the model to operate under its intrinsic stochastic components. The masks shown in Table 4.1, despite similar input, differ in their centroid distribution, local clustering, and overall arrangement.

Input			Output		
Class	# Cells	Cluster			
Steatosis	50-130	0			
Hepatocyte	120	0			
Artifacts	30	0			
Steatosis	350-450	2			
Hepatocyte	50	0			
Artifacts	30	0			

**Table 4.1:** Semantic content generated using different sets of input parameters: tissue is shown in gray, steatosis in green, hepatocyte nuclei in red, and artifacts in white.

The first source of variability originates from the construction of the KDE map used to sample synthetic centroids. Although the density map always reflects the spatial pattern of the selected cluster, the reference image from which KDE is computed is chosen at random among a pool of real images belonging to that cluster and with a similar number of cells. Given the large number of available images, the probability of selecting the same reference more than once is extremely low. This mechanism ensures that HISPACE reproduces the statistical structure of the cluster rather than replicating a specific real tissue layout.

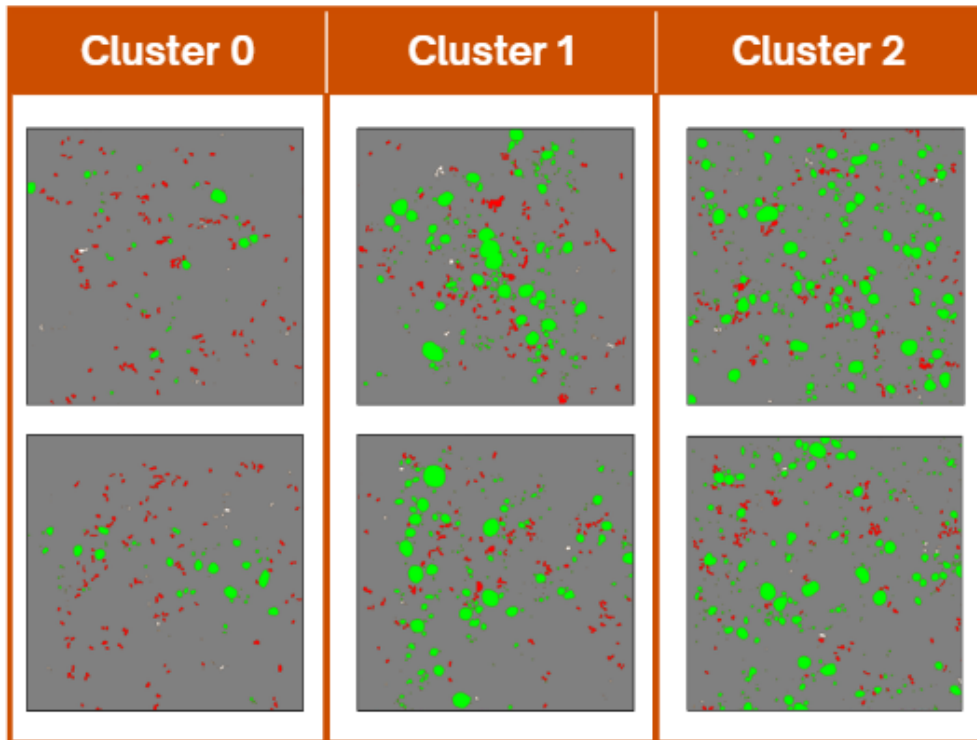
A second source of variability is the stochastic sampling of centroids from the KDE distribution, performed through an optimization-guided process that seeks positions consistent with the reference layout while allowing deviations and local perturbations. Even when two generations were to rely on the same image layout, this sampling

stage alone would yield substantially different centroid configurations.

Additional variability comes from the random selection of the cell contours from the FDs repository, as well as from the random placement attempts used during the insertion process. Finally, the number of cells requested in the test shown here was sampled from a fixed interval, so each mask shares the same generation conditions but not the exact numerical input, further contributing to natural variability.

### Sensitivity to the cluster selection

A second set of experiments was conducted to assess how the spatial arrangement of the synthetic masks changes when varying the cluster assigned to a specific class. In this analysis, the input parameters for Artifact and Hepatocyte Nuclei were fixed throughout all generations, while the cluster selected for the Steatosis class is varied. Table 4.2 presents the generated synthetic masks. The distribution of the steatosis changes significantly, reflecting the spatial topology learned from the real samples belonging to each cluster.



**Table 4.2:** Differences in the semantic content of the Steatosis class associated with different clusters.

The differences observed in the generated masks are consistent with the statistical properties of the clusters identified during the feature analysis:

- Cluster 0 groups images characterized by a low local density and a generally disorganized spatial arrangement. Cells tend to appear isolated or weakly clustered, without forming any recognizable large-scale pattern. The number of cells is relatively small and it contributes to the diffusely scattered appearance of the generated mask. This cluster is suitable for modelling early-stage or mild tissue involvement where no cohesive patterns are detected.

- Cluster 1 contains images with high local density, where cells frequently appear in compact groups. The feature profile indicates the presence of incipient large-scale patterns, as highlighted by moderate spectral components in the FFT features and a more regular local neighborhood structure.

- Cluster 2 is characterized by the highest degree of local density and by a more organized global topology. The structures of this group exhibit large regions of uniform high concentration as demonstrated also by the large number of cells. This last group corresponds to high-density, strongly structured steatosis patterns, resembling advanced pathological states.


This demonstrates that a change of the desired cluster, results in a radically different global spatial organization of the semantic content.

## 4.2 Image Synthesis from Semantic Masks

The semantic masks generated by HISPACE were translated into RGB histological images using the SENSE generative model [18]. SENSE is a conditional GAN trained following a Pix2Pix scheme, but it integrates a transformer-based generator (Restormer) capable of capturing long-range contextual dependencies. This design is particularly well suited for the HISPACE pipeline: the mask provides the spatial arrangement, while the GAN focuses exclusively on learning the appearance model (color, staining, and nuclear texture). The generator improves the reconstruction of heterogeneous

nuclear morphology and context-dependent color variations by attending to both local and global structures. In parallel, the PatchGAN discriminator stabilizes training and encourages the generation of high-frequencies details, helping preserve sharp edges and realistic H&E textures. SENSE was originally trained on real semantic-image pairs, ensuring that the mapping remains biologically plausible even when the semantic content is produced by a synthetic generator.

Below in Fig. 4.1 are shown examples of semantic masks generated by HISPACE alongside their corresponding synthetic histological images. Compared to the semantic visualizations, the generated images accurately reflect the spatial configuration defined in the masks while adding colour, texture, and fine-grained histological details. The semantic masks were generated by explicitly including all cell classes available in the dataset - Steatosis, Hepatocyte Nuclei, Other Nuclei and Artifact - unlike previous demonstrations, which focused on a reduced subset of classes. This allows the synthesized images to capture the full cellular composition of the tissue and to more closely resemble real histological samples.



**Figure 4.1:** Semantic masks and relative generated images. The "Other\_Nuclei" class is represented in blue.

### 4.3 Performances improvement in a segmentation task

To evaluate the practical usefulness of the synthetic images generated by HISPACE, a steatosis segmentation experiment was conducted using a U-Net based architecture. The model was first trained on the original dataset and evaluated on a fixed test set, producing the baseline metrics reported in the first row of Table 4.3. A conventional data augmentation scheme (random rotations and translations) was then applied during training, resulting in a modest improvement across all metrics. Finally, a third training setup was performed by enriching the training set with synthetic images characterized by a high steatosis content generated through the complete HISPACE-SENSE pipeline. These additional samples were designed to compensate for underrepresented scenarios in the real dataset. The augmentation led to a clear improvement in segmentation performances on the same test set.

Architecture	Training method	Sensitivity	Specificity	Dice Score
Unet	Baseline	0.733±0.138	0.993±0.118	0.852±0.106
	+ 100% Data Augmentation	0.741±0.147	0.994±0.112	0.857±0.110
	+ 200% Data Augmentation	0.743±0.141	0.994±0.110	0.858±0.108
	<b>+ HISPACE (200% data augmentation)</b>	<b>0.874±0.094</b>	<b>0.995±0.106</b>	<b>0.908±0.070</b>

**Table 4.3:** Performances improvement of a segmentation task when trained with specific images generated by HISPACE.

The results indicate that the inclusion of HISPACE-generated images substantially increases the model’s ability to correctly identify steatosis regions. Sensitivity improves markedly, indicating fewer false negatives and a better capacity to detect steatosis. The specificity remains high and even increases slightly, showing that the introduction of synthetic samples does not cause the network to over-segment or introduce false positives. Most importantly, the Dice score displays a notable improvement, reflecting a more accurate spatial agreement between the predictions and the ground truth.

Overall, these findings demonstrate that the synthetic data produced by HISPACE can effectively complement conventional augmentation techniques and significantly improve segmentation performances in underrepresented scenarios, such as rare pathological patterns or, in this case, tissue regions with high steatosis concentrations, which are typically harder for the network to learn due to their limited presence in the training data.

## 4.4 Expert-Based Qualitative evaluations

To further assess the plausibility of the synthetic images produced through the HISPACE pipeline, a specialized pathologist conducted a qualitative evaluation. The expert was asked to assign two independent scores (on a 1 to 5 scale) to each of the 20 generated images, focusing on two distinct aspects:

- **Visual realism of the generated image.** This score refers exclusively to the appearance of the synthesized RGB image. The pathologist was instructed to evaluate the hue and colour consistency with H&E staining, the texture realism, the coherence of local and global patterns, and the absence of artifacts or visually implausible regions.
- **Realism of the semantic content.** The second score refers instead to the biological plausibility of the underlying tissue composition generated by HISPACE. For this evaluation, the expert was asked to judge whether the spatial distribution of steatosis, nuclei and artifacts reflects patterns observed in real liver tissue; whether the number, size and arrangement of the architecture are compatible with real pathological scenarios; and whether any unlikely or impossible configurations appear.

The distributions of the assigned scores are summarized in Table 4.4.



	Scores				
	1	2	3	4	5
Realism of the Image	-	-	4	7	9
Realism of the semantic content	-	1	5	7	7

**Table 4.4:** Qualitative scores assessed by a expert pathologist.

The results indicate that the majority of samples (16/20) were judged to have high or very high visual realism (scores  $\geq 4$ ), confirming the effectiveness in reproducing histological textures, staining patterns and structural details.

The realism of the semantic content, although slightly more variable, also shows strong overall plausibility: 14 out of 20 images received a score  $\geq 4$ , suggesting that the spatial arrangements produced by HISPACE are generally consistent with real tissue morphology. A small subset of masks received an intermediate or low score (3, 2), typically associated with configurations where steatosis density or spatial organization deviated from common pathological patterns. These cases highlight the value of expert review and suggest potential directions for refining the semantic generator.

Overall, the expert evaluation confirms that the HISPACE framework can generate images that are both visually realistic and semantically plausible, supporting its applicability not only in downstream computational tasks but also as a resource for training and supporting pathologists, particularly in scenarios where controlled or rare tissue configurations are needed.

# Chapter 5

## Conclusions

This thesis work presented HISPACE, a novel framework for the controlled generation of histological semantic content and its translation into realistic synthetic images. The central objective of the model was to design, implement, and validate a methodology capable of producing biologically plausible tissue configurations while maintaining explicit user-level control over the desired structures.

The proposed architecture introduces a generative process that combines cluster-informed spatial modelling with stochastic sampling and contour-based shape reconstruction. This design enables HISPACE to generate masks that are both structurally coherent and inherently variable, capturing the diversity observed in real tissue. The semantic generator operates independently from the image synthesis model, resulting in a modular pipeline that separates the modelling of spatial semantics from the modelling of visual appearance.

The experimental evaluation confirms the effectiveness of this approach. First, the analysis of the generated semantic masks demonstrates that HISPACE produces non-deterministic yet cluster consistent spatial layouts, successfully reproducing a wide range of topological patterns associated with different degree of steatosis. Second, the integration with SENSE shows that these semantic structures can be translated into high-quality synthetic histological images, exhibiting realistic colour distributions, textures,

and morphological details. Third, the downstream segmentation experiment highlights the practical value of the synthetic data: the inclusion of HISPACE generated samples leads to a substantial improvement in the performances of a U-Net model, designed to segment steatosis. Finally, the expert-based evaluation provides an external validation of the realism of both the synthesized images and the underlying semantic content, further supporting the credibility of the framework.

Overall, the results indicate that HISPACE constitutes a flexible and effective tool for the generation of synthetic histological data. By combining semantic interpretability, biological plausibility, and visual realism, the framework offers a promising foundation for applications in computational pathology as well as for educational and training scenarios.

## 5.1 Limits

Despite the encouraging results, the HISPACE framework presents several limitations that should be taken into account when interpreting the findings and considering future extensions.

The generative process relies on iterative placement attempts for each cell, constrained by collision rules and spatial feasibility. In particular high-density configurations, a cell may fail to be inserted after a fixed number of attempts, and as a consequence, the final number of cells included in the mask may slightly deviate from the requested input. Although these variations are minimal and fully traceable, they introduce a degree of uncertainty that may be undesirable in scenarios that request strict quantitative control.

The semantic variability of the generated masks is strongly influenced by the clustering performed on real samples. While HISPACE successfully reproduces the spatial characteristics of each cluster, the interpretation and utility of such clusters depend on the dataset, the annotation quality, and the intended downstream application. Different clustering setups may result in different semantic partitions, requiring user to

carefully evaluate which configuration best suits their specific generative or analytical goals. This dependency limits the portability of a clustering setup across datasets or tissue type without an appropriate validation step.

A structural limitation of the current framework is that each class is generated independently, without modelling interactions or spatial dependencies between multiple classes. As a result, the final semantic mask does not enforce any multi-class constraints, such as class-specific adjacency rules, exclusion zones, or co-occurrence patterns that may be characteristic of real tissue. This combination of classes may lead to less realistic configurations, and many generation attempts may be necessary to obtain a fully coherent mask in which all classes exhibit a mutually consistent spatial layout.

## 5.2 Future Applications

Building on the results presented in this thesis, several promising research directions and practical applications emerge for the HISPACE framework:

1. **Controlled and pathology-specific data augmentation:** The experiments conducted in this work demonstrate that samples generated by HISPACE can significantly improve segmentation performance in high-steatosis scenarios, which are typically underrepresented in real datasets. A natural extension is to leverage the same approach to perform targeted augmentation for a broader range of histological structures and pathological conditions, such as inflammatory infiltrates, fibrosis patterns, or rare morphological variants. Unlike generic augmentation strategies, HISPACE allows for fine-grained control of the semantic configuration, enabling the generation of synthetic samples tailored to specific morphological distributions or clinically relevant edge cases.
2. **Synthetic dataset generation for training and education:** Given the high visual and semantic realism of the generated images, HISPACE can support the creation of fully synthetic and intrinsically annotated histological datasets for computational modelling and medical training. Such data can be especially

valuable for training pathologists on rare or infrequent patterns, or providing consistent ground truth labels that are otherwise costly or time-consuming to obtain.

3. **Towards multi-scale semantic modelling of tissue architecture:** A further direction involves extending HISPACE from cell-level generative modelling to multi-scale semantic representation that include higher-level structures such as lobular organization, vascular patterns, connective tissue architecture, or pathological macro-regions. Such a multi-scale extension would enable the simulation of complex tissue interactions, distinctions between healthy and pathological spatial organization, and more realistic modelling of disease progression.

# Bibliography

- [1] Ariel Larey et al. “Harnessing artificial intelligence to infer novel spatial biomarkers for the diagnosis of eosinophilic esophagitis”. In: *Frontiers in Medicine* 9 (Oct. 2022). ISSN: 2296-858X. DOI: 10.3389/fmed.2022.950728. URL: <http://dx.doi.org/10.3389/fmed.2022.950728>.
- [2] Andrew H. Song et al. “Artificial intelligence for digital and computational pathology”. In: *Nature Reviews Bioengineering* 1.12 (Oct. 2023), pp. 930–949. ISSN: 2731-6092. DOI: 10.1038/s44222-023-00096-8. URL: <http://dx.doi.org/10.1038/s44222-023-00096-8>.
- [3] Richard J. Chen et al. “Pan-cancer integrative histology-genomic analysis via multimodal deep learning”. In: *Cancer Cell* 40.8 (Aug. 2022), 865–878.e6. ISSN: 1535-6108. DOI: 10.1016/j.ccell.2022.07.004. URL: <http://dx.doi.org/10.1016/j.ccell.2022.07.004>.
- [4] Jeremias Krause et al. “Deep learning detects genetic alterations in cancer histology generated by adversarial networks”. In: *The Journal of Pathology* (Feb. 2021). ISSN: 1096-9896. DOI: 10.1002/path.5638. URL: <http://dx.doi.org/10.1002/path.5638>.
- [5] W. Nicholson Price and I. Glenn Cohen. “Privacy in the age of medical big data”. In: *Nature Medicine* 25.1 (Jan. 2019), pp. 37–43. ISSN: 1546-170X. DOI: 10.1038/s41591-018-0272-7. URL: <http://dx.doi.org/10.1038/s41591-018-0272-7>.
- [6] Ian J. Goodfellow et al. *Generative Adversarial Networks*. 2014. DOI: 10.48550/ARXIV.1406.2661. URL: <https://arxiv.org/abs/1406.2661>.

- [7] Thomas de Bel et al. *Stain-transforming cycle-consistent generative adversarial networks for improved segmentation of renal histopathology*. 2019. URL: <https://openreview.net/forum?id=BkxJkgSlx4>.
- [8] Du Wang et al. “Adversarial neural networks for basal membrane segmentation of microinvasive cervix carcinoma in histopathology images”. In: *2017 International conference on machine learning and cybernetics (ICMLC)*. Vol. 2. IEEE. 2017, pp. 385–389.
- [9] Jerry Wei et al. “Generative image translation for data augmentation in colorectal histopathology images”. In: *Proceedings of machine learning research* 116 (2019), p. 10.
- [10] Maximilian E. Tschuchnig, Gertie J. Oostingh, and Michael Gadermayr. “Generative Adversarial Networks in Digital Pathology: A Survey on Trends and Future Potential”. In: *Patterns* 1.6 (Sept. 2020), p. 100089. ISSN: 2666-3899. DOI: 10.1016/j.patter.2020.100089. URL: <http://dx.doi.org/10.1016/j.patter.2020.100089>.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denoising Diffusion Probabilistic Models*. 2020. DOI: 10.48550/ARXIV.2006.11239. URL: <https://arxiv.org/abs/2006.11239>.
- [12] Sofoklis Katakis et al. “Generation of Musculoskeletal Ultrasound Images with Diffusion Models”. In: *BioMedInformatics* 3.2 (May 2023), pp. 405–421. ISSN: 2673-7426. DOI: 10.3390/biomedinformatics3020027. URL: <http://dx.doi.org/10.3390/biomedinformatics3020027>.
- [13] Kianoush Falahkheirkhah et al. “Deepfake Histologic Images for Enhancing Digital Pathology”. In: *Laboratory Investigation* 103.1 (Jan. 2023), p. 100006. ISSN: 0023-6837. DOI: 10.1016/j.labinv.2022.100006. URL: <http://dx.doi.org/10.1016/j.labinv.2022.100006>.
- [14] Ting-Chun Wang et al. *High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs*. 2017. DOI: 10.48550/ARXIV.1711.11585. URL: <https://arxiv.org/abs/1711.11585>.

- [15] Srijay Deshpande et al. “SynCLay: Interactive synthesis of histology images from bespoke cellular layouts”. In: *Medical Image Analysis* 91 (Jan. 2024), p. 102995. ISSN: 1361-8415. DOI: [10.1016/j.media.2023.102995](https://doi.org/10.1016/j.media.2023.102995). URL: <http://dx.doi.org/10.1016/j.media.2023.102995>.
- [16] Simon Graham et al. “Hover-Net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images”. In: *Medical Image Analysis* 58 (Dec. 2019), p. 101563. ISSN: 1361-8415. DOI: [10.1016/j.media.2019.101563](https://doi.org/10.1016/j.media.2019.101563). URL: <http://dx.doi.org/10.1016/j.media.2019.101563>.
- [17] Srijay Deshpande, Fayyaz Minhas, and Nasir Rajpoot. *Synthesis of Annotated Colorectal Cancer Tissue Images from Gland Layout*. 2023. DOI: [10.48550/ARXIV.2305.05006](https://doi.org/10.48550/ARXIV.2305.05006). URL: <https://arxiv.org/abs/2305.05006>.
- [18] Alen Shahini et al. “Semantic-driven synthesis of histological images with controllable cellular distributions”. In: *Computer Methods and Programs in Biomedicine* 261 (Apr. 2025), p. 108621. ISSN: 0169-2607. DOI: [10.1016/j.cmpb.2025.108621](https://doi.org/10.1016/j.cmpb.2025.108621). URL: <http://dx.doi.org/10.1016/j.cmpb.2025.108621>.
- [19] Nati Daniel et al. *PriorPath: Coarse-To-Fine Approach for Controlled De-Novo Pathology Semantic Masks Generation*. 2024. DOI: [10.48550/ARXIV.2411.16515](https://doi.org/10.48550/ARXIV.2411.16515). URL: <https://arxiv.org/abs/2411.16515>.
- [20] Puria Azadi Moghadam et al. *A Morphology Focused Diffusion Probabilistic Model for Synthesis of Histopathology Images*. 2022. DOI: [10.48550/ARXIV.2209.13167](https://doi.org/10.48550/ARXIV.2209.13167). URL: <https://arxiv.org/abs/2209.13167>.
- [21] Farhad Ghazvinian Zanjani, Svitlana Zinger, and Peter H. N. de With. “Deep Convolutional Gaussian Mixture Model for Stain-Color Normalization of Histopathological Images”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. Springer International Publishing, 2018, pp. 274–282. ISBN: 9783030009342. DOI: [10.1007/978-3-030-00934-2\\_31](https://doi.org/10.1007/978-3-030-00934-2_31). URL: [http://dx.doi.org/10.1007/978-3-030-00934-2\\_31](http://dx.doi.org/10.1007/978-3-030-00934-2_31).



- [22] Violeta N. Kovacheva, David Snead, and Nasir M. Rajpoot. “A model of the spatial tumour heterogeneity in colorectal adenocarcinoma tissue”. In: *BMC Bioinformatics* 17.1 (June 2016). ISSN: 1471-2105. DOI: 10.1186/s12859-016-1126-2. URL: <http://dx.doi.org/10.1186/s12859-016-1126-2>.
- [23] Ruchika Verma et al. “MoNuSAC2020: A Multi-Organ Nuclei Segmentation and Classification Challenge”. In: *IEEE Transactions on Medical Imaging* 40.12 (2021), pp. 3413–3423. DOI: 10.1109/TMI.2021.3085712.
- [24] Marcel R. Ackermann et al. “Analysis of Agglomerative Clustering”. In: *Algorithmica* 69.1 (Dec. 2012), pp. 184–215. ISSN: 1432-0541. DOI: 10.1007/s00453-012-9717-4. URL: <http://dx.doi.org/10.1007/s00453-012-9717-4>.