Question 1: Image Features (14 pts.)

a) **Trajectory fitting** We wish to track the trajectory of a projectile being shot in the night sky. The projectile is equipped with a blinking flashlight, such that it becomes visible at regular intervals over its trajectory. We only have a single long-exposure gray-scale image of resolution 800x800 pixels to perform the task. Our image, like the one in Figure 1a, exposed the image sensor to the sky for five minutes, so that all the blinks from the projectile are visible at once on the image, and form a parabolic trajectory. However, also stars are visible on the background and are not clearly distinguishable from the blinks, if not for their random scattering in the image. Our objective is to find the parameters of the parabola traced by the flying projectile.

i) The first task is to clean all of the disturbance from the clouds and the background, and decide for every pixel whether it **could** belong to a blink or not. Figure 1b shows an example of the desired output of this step. Note that at this step stars don't need to be removed, but clouds and the general background should be removed. To study the statistics of the image in Figure 1a, we collect an intensity histogram of its pixels, reported in Figure 1c. Note that such histogram is reported with the vertical axis in **logarithmic scale**, to make small counts more visible. Briefly describe a procedure that could perform the task of cleaning the night image using the histogram. Provide explicit values for every parameter that you may use in your procedure. **2 pts.**

ii) We now have the binarized image in Figure 1b, where each pixel is either 0 (black, background) or 1 (white, a blink or a star). We want to further summarize this representation into a matrix $B \in \mathbb{R}^{N \times 2}$ where N is the number of blobs (i.e. clusters of white pixels connected together) in the binary image, and every row $b_i \in \mathbb{R}^2$ represents the average pixel coordinates of the i-th blob. Describe a procedure that would produce matrix B from the binarized image. **3 pts.**

iii) Finally, we have the matrix B built in the previous exercise which lists N possible pixel locations where a blink could have occurred. Assuming that actual blinks are well aligned along a parabola, and stars are randomly scattered around the image, describe how to use a Hough transform to distinguish the blinks from the stars. In particular, you should describe:

- The dimensionality D of the necessary Hough transform.
- The parameterization of the necessary Hough transform.
- How a point $b_i \in \mathbb{R}^2$ is used to update the discrete Hough space.
- How to select the correct parameterization of the parabola.

5 pts.

b) **Correlation and convolution**

Let $K \in \mathbb{R}^{M \times M}$ be a kernel such that $K_{ij} \in \mathbb{R}$ and $i, j \in [-N, N] \subset \mathbb{Z}$, $M = 2N + 1$. We denote by $K * I$ the convolution of K on image I, and with $K \circ I$ the correlation of K on image I.

i) For which kernels K we have $K * I = K \circ I$ for every image I? Give an example of such kernel. **2 pts.**

ii) Let now K be defined by:

$$K_{ij} = \frac{1}{6\pi} e^{-\frac{i^2+j^2}{6}}$$

Name a use for this kernel, and show that it is separable. **2 pts.**

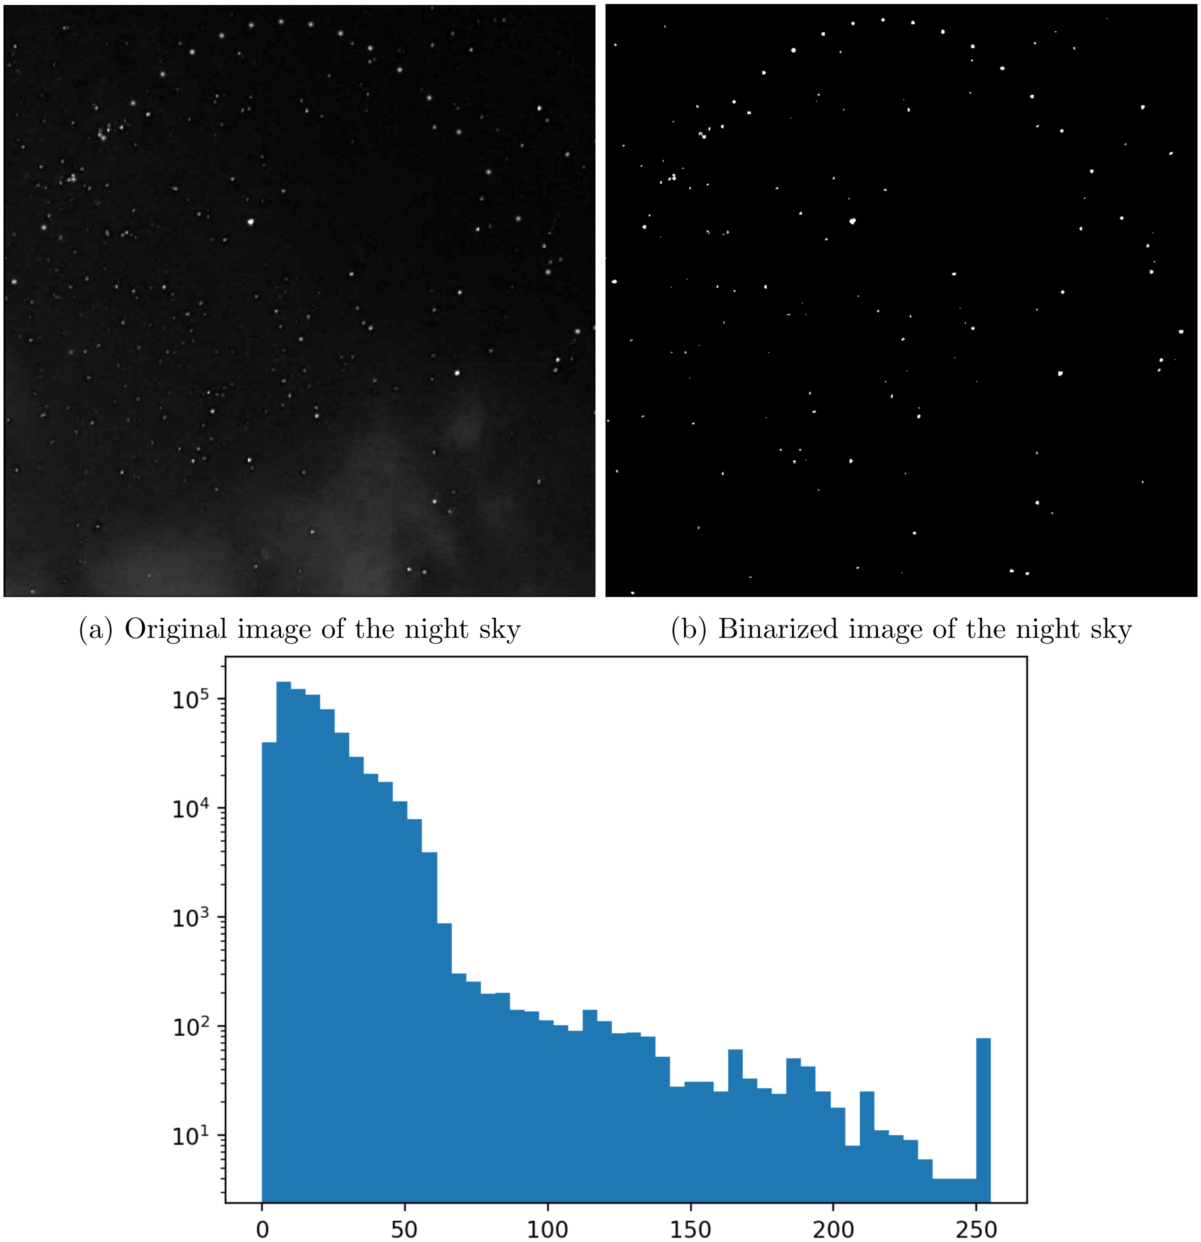

(a) Original image of the night sky

(b) Binarized image of the night sky

(c) Logarithmic histogram of Figure 1a

Figure 1: A blinking projectile over the night sky.

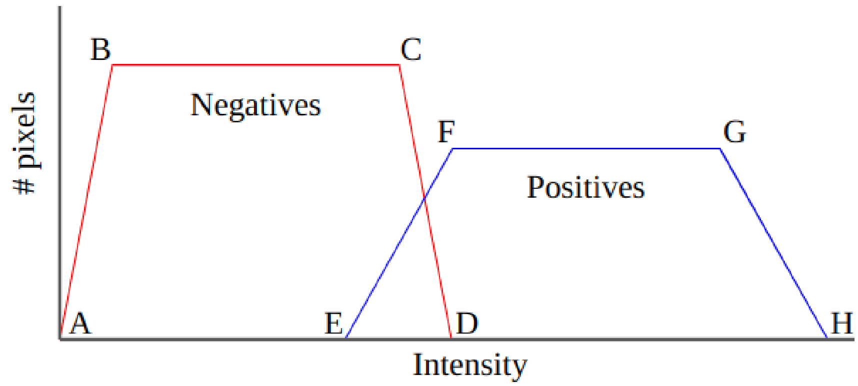

Figure 2: Separate histograms of Positives and Negatives. Refer to the text for coordinates of points A to H

Question 2: Basic image operations (9 pts.)

a) **Segmentation**

We wish to segment a grayscale image I with a simple thresholding approach. The intensity histogram of I separately for the positive class (in blue) and negative class (in red) is reported in Figure 2. The coordinates of points A to H are given: A(0, 0), B(20, 100), C(100, 100), D(120, 0), E(85, 0), F(120, 70), G(200, 70), H(235, 0). Points A to D define the histogram of the negative class, while points E to H define the histogram of the positive class. Let the threshold used be named τ:

i) What is the maximum True Positive Rate achievable? Give a value of τ that achieves it:

2 pts.

ii) What is the minimum False Positive Rate achievable? Give a value of τ that achieves it:

2 pts.

iii) What is the minimum number of incorrectly classified pixels? Give a value of τ that achieves it. Note: for this question you can assume that intensities and pixel counts are continuous quantities. **5 pts.**

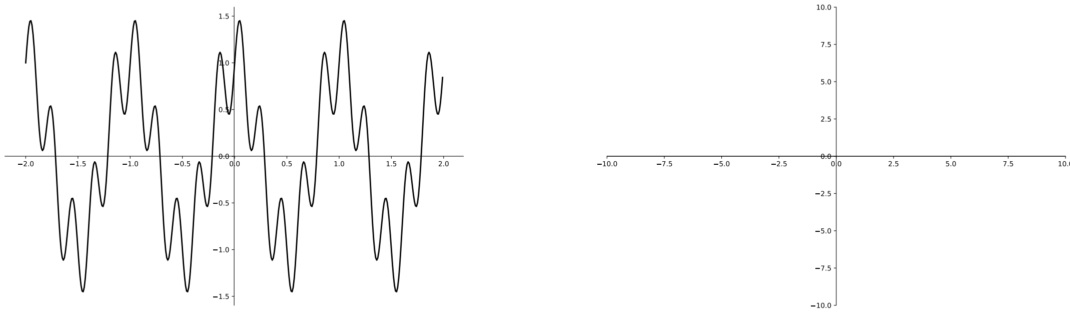

(a) (b)

Figure 3: Superposition of sinusoidal signals and the corresponding Fourier transform.

(a) (b)

Figure 4: Images with various frequencies.

Question 3: Fourier transform and filtering (12 pts.)

i) Figure 3a is showing a superposition of different sinusoidal signals. Draw the Fourier transform of this signal on Figure 3b. **2 pts.**

ii) Consider the two images of Figure 4: which one has the highest frequency components? Justify your answer. **1 pts.**

iii) Suppose that we want to downsample Image 4b by a factor of 8 and that we know that the maximal frequency of the Fourier transform of the image is $\frac{1}{4}$. Is it possible to directly subsample the image without artifacts? If not, what problem could arise and how would you solve it? Justify all your answers. **4 pts.**

iv) We now want to sharpen image 4a. What kind of filter would you use for that? Detail how you would apply it using the Fourier transform and frequency space. **3 pts.**

v) Among the following kernels, which one is the low-pass filter and which one the high-pass? **2 pts.**

$$(a)\begin{bmatrix} 0.023 & 0.034 & 0.038 & 0.034 & 0.023 \\ 0.034 & 0.049 & 0.056 & 0.049 & 0.034 \\ 0.038 & 0.056 & 0.063 & 0.056 & 0.038 \\ 0.034 & 0.049 & 0.056 & 0.049 & 0.034 \\ 0.023 & 0.034 & 0.038 & 0.034 & 0.023 \end{bmatrix} \qquad (b)\begin{bmatrix} -0.0030 & -0.013 & -0.022 & -0.013 & -0.0030 \\ -0.013 & -0.060 & -0.098 & -0.060 & -0.013 \\ -0.022 & -0.098 & 0.84 & -0.098 & -0.022 \\ -0.013 & -0.060 & -0.098 & -0.060 & -0.013 \\ -0.0030 & -0.013 & -0.022 & -0.013 & -0.0030 \end{bmatrix}$$

Question 4: Optical Flow (11 pts.)

a) The Lucas-Kanade algorithm can be used to estimate the optical flow of an image sequence. Nevertheless, it does not work for all cases. State the 3 conditions that have to be fulfilled so that this method works. **3 pts.**

b) We consider a camera facing a planar scene along its Z axis, which points towards the scene. The camera captures images at a frame rate of 25 Hz and is a pinhole model with focal length $f = 500$ pixels. The camera plane is parallel to the scene and its projection equation relates a 3D point $\mathbf{P} = (X, Y, Z)$ to its pixel coordinates (x, y) by

$$x = \frac{fX}{Z}, \tag{1}$$

and similarly for y.

i) Express the 2D motion field $(\delta x, \delta y)$ (in pixels) of the 3D point $\mathbf{P}$ between two consecutive frames given its 3D velocity $\mathbf{V} = (V_x, V_y, V_z)$ in meters/second. Compute its value for $\mathbf{P} = (1, 2, 3)$ and $\mathbf{V} = (1, 0, 1)$. **5 pts.**

ii) What is the difference between this 2D motion field and the optical flow? **1 pt.**

iii) In which scenarios are the optical flow and the 2D motion field not equal? **2 pts.**

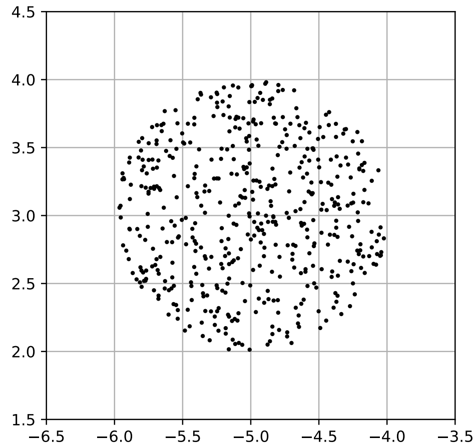

(a) Distribution A.

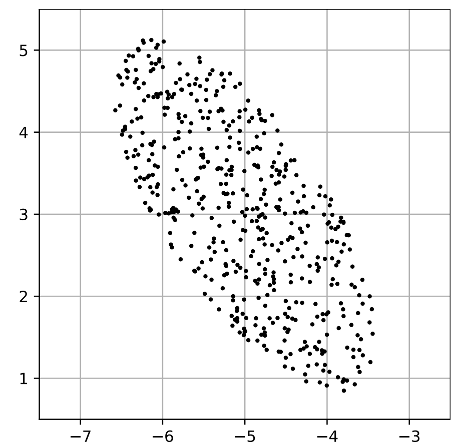

(b) Distribution B.

Figure 5: Illustrations for Question 5b.

Question 5: Principal Component Analysis (11 pts.)

a) We consider a dataset of 1000 faces images with size 40×40. Each array entry or number costs 1 unit of space. We wish to store this dataset given a budget of Z units of space. We choose to compress this dataset with PCA. Calculate the maximum number K of principal components allowed, taking into account all quantities required to decompress the dataset.

6 pts.

b) We consider two two-dimensional data distributions, A and B, shown in Figure 5.

   i) Let us apply PCA to each set of point. In each plot, draw the expected sample mean and the two principal components. Clearly indicate which one is the largest, or explain when this is not possible. **4 pts.**

ii) Given a new set of points, we wish to determine whether they were drawn from A or from B. Derive a simple test, with a unique threshold, based on the result of applying PCA to this point set. **1 pt.**