

CURRENCY DETECTION FOR VISUALLY IMPAIRED PEOPLE USING MACHINE LEARNING

Submitted in partial fulfillment of the requirements for the award of
Bachelor of Engineering degree in Computer Science and Engineering

by

**Maddi Lakshmi Sai Ratna Sri
(17CS0251)**



**DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
SCHOOL OF COMPUTING**

SATHYABAMA

**INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)**

**Accredited with Grade "A" by NAAC | 12B Status by UGC | Approved by AICTE
JEPPIAAR NAGAR, RAJIV GANDHI SALAI, CHENNAI - 600 119**

March - 2021



SATHYABAMA

INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)

Accredited with "A" grade by NAAC | 12B Status by UGC | Approved by
AICTE

Jeppiaar Nagar, Rajiv Gandhi Salai, Chennai – 600 119

www.sathyabama.ac.in

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING **BONAFIDE CERTIFICATE**

This is to certify that this Project Report is the bonafide work of Ms. Maddi Lakshmi Sai Ratna Sri (17CS0251) and Ms. Nithyashree.M.K(37110517) who carried out the project entitled " **CURRENCY DETECTION FOR VISUALLY IMPAIRED PEOPLE USING MACHINE LEARNING**" under my supervision from November 2020 to March 2021.

Internal Guide
(YOGITHA.R)

Head of the Department

Submitted for Viva voce Examination held on_____

Internal Examiner

External Examiner

DECLARATION

Maddi Lakshmi Sai Ratna Sri 17CS0251) declare that the ProjectReport entitled **CURRENCY DETECTION FOR VISUALLY IMPAIRED PEOPLE USING MACHINE LEARNING** done by me under the guidance of Ms. Yogitha.R.(Internal) and ____ (External) at Sathyabama Institute of Science and Technologyis submitted

in partial fulfillment of the requirements for the award of Bachelor ofEngineering / Technology degree in Computer Science and Engineering.

DATE:

PLACE:

SIGNATURE OF THE CANDIDATE

ACKNOWLEDGEMENT

I am pleased to acknowledge my sincere thanks to **Board of Management of SATHYABAMA** for their kind encouragement in doing this project and for completing it successfully. I am grateful to them.

I convey my thanks to **Dr. T.Sasikala M.E., Ph.D .**, Dean, School of Computing and **Dr.L.Lakshmanan M.E., Ph.D.,and Dr.S.Vigneshwari M.E., Ph.D.**, Heads of the Department of Computer Science and Engineering for providing me necessary support and details at the right time during the progressive reviews.

I would like to express my sincere and deep sense of gratitude to my Project Guide Ms.Yogitha.R for her valuable guidance, suggestions and constant encouragement paved way for the successful completion of my project work.

I wish to express my thanks to all Teaching and Non-teaching staff members of the **Department of Computer Science and Engineering** who were helpful in many ways for the completion of the project.

ABSTRACT

In the current world of cheating and other frauds, people find it very difficult to lead their lives. The same is even worse in the case of Blind or visually challenged. They face many more challenges in their daily life, especially while dealing with currency and other money related issues. In order to help blind people we are developing a project that helps them identify the denomination of the currency, as there is no brail mark on the currency note. Counterfeit note or fake note is another difficult task to identify, for both healthy and blind person. Hence, in addition to it we are incorporating fake note detection system, which helps every citizen from verge of being cheated. The currency recognition algorithm discussed in this paper using image processing is based on an ORB (Oriented fast and Rotated Brief). It is faster and also rotation invariant. The proposed algorithm shows 90% true recognition rate.

TABLE OF CONTENTS

S.NO	TITLE	PAGE.NO
	ABSTRACT	v
	LIST OF FIGURES	viii
1	INTRODUCTION	01
	1.1 ABOUT THE PROJECT	01
	1.2 DOMAIN OVERVIEW	02
2	LITERATURE REVIEW	04
3	AIM AND SCOPE	07
	3.1 AIM AND SCOPE	07
	3.2 OBJECTIVES	07
	3.3 EXISTING SYSTEM	07
	3.4 PROPOSED SYSTEM	07
4	SYSTEM IMPLEMENTATION & METHODOLOGY	09
	4.1 SYSTEM REQUIREMENTS	09
	4.1.1 HARDWARE USED	09
	4.1.2 SOFTWARE USED	09
	4.2 SYSTEM ARCHITECTURE	10
	4.3 MODULES	10
	4.3.1 IMAGE ACQUISTION	11
	4.3.2 PRE-PROCESSING	11
	4.3.3 THRESHOLDING	11
	4.3.4 EDGE DETECTION	12
	4.3.5 SEGEMENTATION	12
	4.4 METHODOLOGY	12
	4.4.1 IMAGE PROCESSING	12
	4.4.2 CONVOLUTION MATRIX	19
	4.4.3 DITHERING	23
	4.4.4 EDGE DETECTORS	23
	4.4.5 MASKING	33
	4.4.6 PIXELS	37
5	RESULTS AND DISCUSSION	46

6	CONCLUSION AND FUTURE ENHANCEMENT	49
	6.1 CONSLUSION	49
	6.2 FUTURE WORK	49
	REFERENCES	50

f

LIST OF FIGURES

FIGURE NO	FIGURE NAME	PAGE NO
4.1	System Architecture	10
4.2	Detailing on Image Processing	14
4.3	2×2 pixel area displaying one composite color	19
4.4	Convolution illustration	20
4.5	1 st & 2 nd derivative of an edge illustrated in 1D	24
4.6	Convolution kernel	29
4.7	Truth-tables for AND and NAND	30
4.8	Simple graylevel image	34
4.9	2 connected components based on 4 connectivity	39
4.10	Some example structuring elements	42
4.11	Mapping function for wrapping the pixel	44
4.12	Mapping function for saturating	44
5.1	Verification of Input	46
5.2	Comparision of Input with Trained Data	47
5.3	Output of the image	47
5.4	Input of a Fake note	48
5.5	Output of a Fake note	48

CHAPTER 1

INTRODUCTION

1.1 ABOUT THE PROJECT

The lives of visually impaired people are very miserable, they often face problems to survive in this world. Few people tend to take advantage of their situation and use them for their own needs. Blind people are betrayed easily by others. Thanks to the growing technology, that the blind are able to rely on it, instead of trusting an outsider.

We made one such effort in our project, it helps visually impaired people to identify denomination of currency with the help of VGG algorithm (a CNN algorithm) a Deep Learning mechanism. Also, the physical properties of the note, like checking if the note is torn or not are also included, to help them utmost. The output is in the form of audio such that, it wont be a problem for the blind in usage of the project.

Currency is one of the most important aspect in our day to day life, every small and big thing we do in our daily routine is linked to money. Being the most essential segment, many people tend to do forgery in the same and cheat other. Fake currency detection is a serious issue worldwide, affecting the economy of almost every country including India. Currency duplication also known as counterfeit currency is a vulnerable threat on economy. It is now a common phenomenon due to advanced printing and scanning technology. Normal people also face a problem in recognizing a counterfeit note, so the scope of the project is extended from visually impaired to almost everyone having a problem with identification of counterfeit notes. The possible solutions are to use either chemical properties of the currency or to use its physical appearance. The approach presented in this paper is based upon physical appearance of the Indian currency. Image processing algorithms have been adopted to extract the features such as security thread, intaglio printing (RBI logo) and identification mark, which have been adopted as security features of Indian currency. The output is in the form of audio that says "This is a fake note", the output is taken in the form of audio such that it can be helpful to visually impaired as well as normal people.

1.2 DOMAIN OVERVIEW

Machine learning is to predict the future from past data. Machine learning (ML) is a type of artificial intelligence (AI) that provides computers with the ability to learn without being explicitly programmed. Machine learning focuses on the development of Computer Programs that can change when exposed to new data and the basics of Machine Learning, implementation of a simple machine learning algorithm using python. Process of training and prediction involves use of specialized algorithms. It feed the training data to an algorithm, and the algorithm uses this training data to give predictions on a new test data. Machine learning can be roughly separated in to three categories. There are supervised learning, unsupervised learning and reinforcement learning. Supervised learning program is both given the input data and the corresponding labeling to learn data has to be labeled by a human being beforehand. Unsupervised learning is no labels. It provided to the learning algorithm. This algorithm has to figure out the clustering of the input data. Finally, Reinforcement learning dynamically interacts with its environment and it receives positive or negative feedback to improve its performance.

Data scientists use many different kinds of machine learning algorithms to discover patterns in python that lead to actionable insights. At a high level, these different algorithms can be classified into two groups based on the way they “learn” about data to make predictions: supervised and unsupervised learning. Classification is the process of predicting the class of given data points. Classes are sometimes called as targets/ labels or categories. Classification predictive modeling is the task of approximating a mapping function from input variables(X) to discrete output variables(y). In machine learning and statistics, classification is a supervised learning approach in which the computer program learns from the data input given to it and then uses this learning to classify new observation. This data set may simply be bi-class (like identifying whether the person is male or female or that the mail is spam or non-spam) or it may be multi-class too. Some examples of classification problems are: speech recognition, handwriting recognition, bio metric 2 identification, document classification etc. Supervised Machine Learning is the majority of practical machine learning uses

supervised learning. Supervised learning is where we have input variables (X) and an output variable (y) and use an algorithm to learn the mapping function from the input to the output is $y = f(X)$.




Fig 1.1 Process of Machine learning

The goal is to approximate the mapping function so well that when you have new input data (X) that you can predict the output variables (y) for that data. Techniques of Supervised Machine Learning algorithms include logistic regression, multi-class classification, Decision Trees and support vector machines etc. Supervised learning requires that the data used to train the algorithm is already labelled with correct answers. Supervised learning problems can be further grouped into Classification problems. This problem has as goal the construction of a succinct model that can predict the value of the dependent attribute from the attribute variables. A classification model attempts to draw some conclusion from observed values. Given one or more inputs a classification model will try to predict the value of one or more outcomes. A classification problem is when the output variable is a category.

CHAPTER 2

LITERATURE SURVEY

Personalized End-to-End Mandarin Speech Synthesis using Small-sized Corpus - ChenHan Yuan (2020)

A start to finish technique is required for customized discourse combination framework by consolidating the plans to handle these issues and attempt to make the information assortment task feasible. An incorporated the phonetic/acoustic element extraction of the discourse corpus by receiving appropriate neural organizations can be utilized. Along these lines, the customary semantic element extraction module which depends on the master information to construct could be subbed.

Image Authentication based on Q-FMT via Blind Geometric Correction - Ram Kumar Karsh (2020)

The quaternion Fourier Mellin change (Q-FMT) is utilized in picture hashing, because of its characteristic properties of heartiness against turn, scaling, and interpretation. The info RGB shading picture is changed over into a quaternion area, which gives the spatial just as chromatic data.

Saliency Consistency-Based Image Re-Colorization for Color Blindness - Jinjiang Li (2020)

The use of picture recovery techniques to choose an enormous number of pictures and utilize co-saliency strategies to identify striking zones of standard shading pictures and partially blind reproduced pictures. As per the location results, a picture with a similar identification result was chosen as the reference picture. Grayscale the altogether changed picture, recolor the grayscale picture utilizing the reference picture.

Edge Sensitive Unsupervised Image-to-Image Translation - Ibrahim Batuhan Akkaya (2020)

The edge data which is determined by Sobel administrator is used for diminishing the curios. The edge-protecting misfortune work, in particular Sobel misfortune is acquainted with accomplish this objective which is defined as the L2 standard between the Sobel reactions of the first and the interpreted pictures.

Vocoder-free End-to-End Voice Conversion with Transformer Network - June-Woo Kim (2020)

Discourse generators with the MFB approach require an extra computationally costly vocoder for the preparation cycle. The pre-and post-handling required by the MFB and the vocoder isn't basic to change over human voices, since it is conceivable to utilize just the crude range to produce distinctive style of voices with intelligible articulation.

An Object Detection Technique For Blind People in Real-Time Using Deep Neural Network - Ashwani Kumar (2020)

Because of ongoing headway in profound learning with picture preparing, empower us to build up this model. The model takes shading pictures as information and train the mode until the mistake rate is less. To expand the computational exhibition of the model they have utilized single shot multi-box indicator calculation.

Tactile Graphics Exploration Studies Using Fingertip Tracking Based on Colour Markers Detection for Visually Impaired People - Muhammad Ikmal Hakim Shamsul Bahrin (2019)

To get visual data like located individuals seeing pictures, dazzle or outwardly disabled individuals need to utilize material designs materials and should use their touch faculties to peruse and investigate each data on them. In any case, there are constraints to decipher data on material designs with tangible discernment. The perspective on material designs is restricted to the size of an individual's fingertips.

A Hybrid Fake Banknote Detection model using OCR, Face Recognition and Hough Features - Adiba Zarin (2019)

Many note recognition frameworks are available in banks yet they are exorbitant and frequently mistaken. The fields of picture handling, neural organization and machine vision can possibly fundamentally beat this issue. This paper proposes a model contained Optical Character acknowledgment (OCR), Face Recognition and Hough change calculation. The microprinting, water-imprint, and bright lines highlights of Bangladeshi notes are separated for testing of certifiable notes.

Fake Colorized Image Detection with Channel-wise Convolution based Deep-learning Framework - Long Zhuo (2019)

WISERNet (Wider Separate-then-get-together Network), an as of late proposed profound learning based information driven shading picture steganalyzer in the field of phony colorized picture recognition. The factual irregularities presented by various programmed colorization techniques can be caught by cutting edge profound learning based information driven shading picture steganalyzers.

Gradient-Based Illumination Description for Image Forgery Detection - Falko Matern (2019)

The key idea is that the basic over an incline field of a thing shows the orientation of scene light in the image plane. As opposed to before 2-D lighting techniques, the proposed system is shockingly solid to changes in article tone and assortments in customer commitment, as it works generally speaking thing locale instead of item shapes.

CHAPTER 3

AIM AND SCOPE

3.1 AIM AND SCOPE

The aim of the project is to identify currency note with the help of it's length, breadth and other hidden features, not visible to human eye without help of an UV ray. Also to recognize if the note is original or fake.

3.2 OBJECTIVE

- To detect the denomination of the currency note – with help of its orientation and colour.
- To detect fake note alarm – to help the users from verge of being cheated.

3.3 EXISTING SYSTEM

In existing system, Image processing is being used with legacy version of machine learning algorithm. Device user friendliness is not good.

Detection of banknotes images done using sensor (sensor based) suffers from less accuracy than vision-based. The main shortfall of sensor-based system is the involvement of several electrical components and limitation of the sensor.

The specialists utilize a covering sensor to perceive shades of money note by then being managed by a microcontroller to see and send a revamp sound manual for a toll as sign. The banknote plan system in get money pictures utilizing one dimensional unmistakable light sensor.

3.4 PROPOSED SYSTEM

The proposed system is made using opencv tool which is light weight tool and python too so we get our output as soon as possible. Compared to existing system we increase time and accuracy because of shift algorithm it another name called pattern recognition. So, it takes each and every feature, compare input with trained data. The currency recognition system algorithm shows the appropriate preprocessing, feature extraction

via ORB & feature matching. With the help of confusion matrix, accuracy can be calculated. The result is integrated with the help of graphical user interface (GUI). This system can be implemented where money related transactions are involved and found to be reliable for visually disabled people.

CHAPTER 4

METHODOLOGY AND SYSTEM IMPLEMENTATION

4.1 SYSTEM REQUIREMENTS

4.1.1 SOFTWARE REQUIREMENTS

- Operating system: Windows 7/10.
- Coding Language: Python
- Opencv tool
- Matlab

Packages:

- Pyttsx3 package
- Matplotlib package

4.1.2 HARDWARE REQUIREMENTS

- System: Pentium Dual Core.
- Hard Disk: 120 GB.
- Monitor: 15" LED
- Ram: 4 GB.

4.2 SYSTEM ARCHITECTURE




Fig 4.1 System Architecture

4.3 MODULES

- Image Acquisition
- Pre-processing
- Edge detection
- Thresholding
- Segmentation

4.3.1 Image Acquisition

Image Acquisition is a technique to develop rough pictures received from cameras/sensors set on satellites for various applications. Various strategies have been made in Image Processing during the last decade. Processing structures are getting notable as a result of basic openness of stunning staff.

The inputted image under goes greyscale conversion using gaussian grey method or the gaussian blur method, this is done so, to identify any hidden features in the image. Then, the image is also converted into its negative, this enables the system to be more efficient and find even more hidden elements in the image.

4.3.2 Pre-processing

Pre-processing is a process where activities with pictures at the most insignificant level of impression of both information and yield are power pictures. The mark of pre-processing is development of the image data that suppresses unwanted turning. All the undesirable mutilation is taken out in this module.

Image blur and dilate image methods are used to remove unnecessary things(noise) from the image, it helps in image filtration.

4.3.3 Thresholding

Programmed thresholding is an extraordinary method to separate valuable data encoded into pixels while limiting foundation commotion. This is refined by using an input circle to upgrade the limit an incentive prior to changing over the first grayscale picture to double. The thought is to isolate the picture into two sections the foundation and frontal area.

Binary Thresholding is done to give outlines of the inputted image, the method is also called as masking. Adaptive thresholding is again done, to get accuracy.

4.3.4 Edge detection

Edge detection is a picture processing method for finding the limits of articles inside pictures. It works by distinguishing discontinuities in brightness. In this module the edges of the note is distinguished utilizing the idea of progress in brightness.

Edge detection is done using sobel edge detection method, to improve accuracy it is again done using canny edge detection. The thresholded image output is fed into Laplacian edge method, to provide presicion.

4.3.5 Segmentation

The well known strategies utilized for picture segmentation are: thresholding strategy, edge location based methods, district based procedures, bunching based methods. Picture thresholding is a basic, yet compelling, method of apportioning a picture into a forefront and foundation. This picture examination method is a sort of picture segmentation that segregates protests by changing over grayscale pictures into paired pictures.

4.4 METHODOLOGY

4.4.1 IMAGE PROCESSING

Image Processing system includes treating images as two dimensional signals while applying already set signal processing methods to them.

It is among rapidly growing technologies today, with its applications in various aspects of a business. Image Processing forms core research area within engineering and computer science disciplines too.

Image processing basically includes the following three steps:

- Importing the image with optical scanner or by digital photography.
- Analyzing and manipulating the image which includes data compression and image enhancement and spotting patterns that are not to human eyes like satellite

photographs.

- Output is the last stage in which result can be altered image or report that is based on image analysis.

Purpose of Image processing

The purpose of image processing is divided into 5 groups.

They are:

1. Visualization - Observe the objects that are not visible.
2. Image sharpening and restoration - To create a better image.
3. Image retrieval - Seek for the image of interest.
4. Measurement of pattern – Measures various objects in an image.
5. Image Recognition – Distinguish the objects in an image.

Types of Image Processing

The two types of methods used for Image Processing are Analog and Digital Image Processing. Analog or visual techniques of image processing can be used for the hard copies like printouts and photographs. Image analysts use various fundamentals of interpretation while using these visual techniques. The image processing is not just confined to area that has to be studied but on knowledge of analyst. Association is another important tool in image processing through visual techniques. So analysts apply a combination of personal knowledge and collateral data to image processing.

Digital Processing techniques help in manipulation of the digital images by using computers. As raw data from imaging sensors from satellite platform contains deficiencies. To get over such flaws and to get originality of information, it has to undergo various phases of processing. The three general phases that all types of data have to undergo while using digital technique are Pre- processing, enhancement and display, information extraction.




Fig 4.2 Detailing on Image Processing

IMAGE PROCESSING CONCEPTS

Binary Images

Binary images are images whose pixels have only two possible intensity values. They are normally displayed as black and white. Numerically, the two values are often 0 for black, and either 1 or 255 for white.

Binary images are often produced by thresholding a grayscale or color image, in order to separate an object in the image from the background. The color of the object (usually

white) is referred to as the foreground color. The rest (usually black) is referred to as the background color. However, depending on the image which is to be thresholded, this polarity might be inverted, in which case the object is displayed with 0 and the background is with a non-zero value.

Some morphological operators assume a certain polarity of the binary input image so that if we process an image with inverse polarity the operator will have the opposite effect. For example, if we apply a closing operator to a black text on white background, the text will be opened.

Color Images

It is possible to construct (almost) all visible colors by combining the three primary colors red, green and blue, because the human eye has only three different color receptors, each of them sensible to one of the three colors. Different combinations in the stimulation of the receptors enable the human eye to distinguish approximately 350000 colors. A RGB color image is a multi-spectral image with one band for each color red, green and blue, thus producing a weighted combination of the three primary colors for each pixel.

A full 24-bit color image contains one 8-bit value for each color, thus being able to display $2^{24} = 16777216$ different colors.

However, it is computationally expensive and often not necessary to use the full 24-bit image to store the color for each pixel. Therefore, the color for each pixel is often encoded in a single byte, resulting in an 8-bit color image. The process of reducing the color representation from 24-bits to 8-bits, known as color quantization, restricts the number of possible colors to 256. However, there is normally no visible difference between a 24-color image and the same image displayed with 8 bits. An 8-bit color images are based on colormaps, which are look-up tables taking the 8-bit pixel value as index and providing an output value for each color.

8-bit Color Images

Full RGB color requires that the intensities of three color components be specified for each and every pixel. It is common for each component intensity to be stored as an 8-bit integer, and so each pixel requires 24 bits to completely and accurately specify its color. If this is done, then the image is known as a 24-bit color image. However there are two problems with this approach:

- Storing 24 bits for every pixel leads to very large image files that with current technology are cumbersome to store and manipulate. For instance a 24-bit 512×512 image takes up 750KB in uncompressed form.
- Many monitor displays use colormaps with 8-bit index numbers, meaning that they can only display 256 different colors at any one time. Thus it is often wasteful to store more than 256 different colors in an image anyway, since it will not be possible to display them all on screen.

Because of this, many image formats (e.g. 8-bit GIF and TIFF) use 8-bit colormaps to restrict the maximum number of different colors to 256. Using this method, it is only necessary to store an 8-bit index into the colormap for each pixel, rather than the full 24-bit color value. Thus 8-bit image formats consist of two parts: a colormap describing what colors are present in the image, and the array of index values for each pixel in the image.

When a 24-bit full color image is turned into an 8-bit image, it is usually necessary to throw away some of the colors, a process known as color quantization. This leads to some degradation in image quality, but in practice the observable effect can be quite small, and in any case, such degradation is inevitable if the image output device (e.g. screen or printer) is only capable of displaying 256 colors or less.

The use of 8-bit images with colormaps does lead to some problems in image processing. First of all, each image has to have its own colormap, and there is usually no guarantee that each image will have exactly the same colormap. Thus on 8-bit displays it is frequently impossible to correctly display two different color images that

have different colormaps at the same time. Note that in practice 8-bit images often use reduced size colormaps with less than 256 colors in order to avoid this problem.

Another problem occurs when the output image from an image processing operation contains different colors to the input image or images. This can occur very easily, as for instance when two color images are added together pixel-by-pixel. Since the output image contains different colors from the input images, it ideally needs a new colormap, different from those of the input images, and this involves further color quantization which will degrade the image quality. Hence the resulting output is usually only an approximation of the desired output. Repeated image processing operations will continually degrade the image colors. And of course we still have the problem that it is not possible to display the images simultaneously with each other on the same 8-bit display.

Because of these problems it is to be expected that as computer storage and processing power become cheaper, there will be a shift away from 8-bit images and towards full 24-bit image processing.

24-bit Color Images

Full RGB color requires that the intensities of three color components be specified for each and every pixel. It is common for each component intensity to be stored as an 8-bit integer, and so each pixel requires 24 bits to completely and accurately specify its color. Image formats that store a full 24 bits to describe the color of each and every pixel are therefore known as 24-bit color images.

Using 24 bits to encode color information allows $2^{24} = 16777216$ different colors to be represented, and this is sufficient to cover the full range of human color perception fairly well.

The term 24-bit is also used to describe monitor displays that use 24 bits per pixel in their display memories, and which are hence capable of displaying a full range of colors.

There are also some disadvantages to using 24-bit images. Perhaps the main one is that it requires three times as much memory, disk space and processing time to store and manipulate 24-bit color images as compared to 8-bit color images. In addition, there is often not much point in being able to store all those different colors if the final output device (e.g. screen or printer) can only actually produce a fraction of them. Since it is possible to use colormaps to produce 8-bit color images that look almost as good, at the time of writing 24-bit displays are relatively little used. However it is to be expected that as the technology becomes

Color Quantization

Color quantization is applied when the color information of an image is to be reduced. The most common case is when a 24-bit color image is transformed into an 8-bit color image.

Two decisions have to be made:

1. which colors of the larger color set remain in the new image, and
2. how are the discarded colors mapped to the remaining ones.

The simplest way to transform a 24-bit color image into 8 bits is to assign 3 bits to red and green and 2 bits to blue (blue has only 2 bits, because of the eye's lower sensitivity to this color). This enables us to display 8 different shades of red and green and 4 of blue. However, this method can yield only poor results. For example, an image might contain different shades of blue which are all clustered around a certain value such that only one shade of blue is used in the 8-bit image and the remaining three blues are not used.

Alternatively, since 8-bit color images are displayed using a colormap, we can assign any arbitrary color to each of the 256 8-bit values and we can define a separate colormap for each image. This enables us perform a color quantization adjusted to the data contained in the image. One common approach is the popularity algorithm, which creates a histogram of all colors and retains the 256 most frequent ones. Another

approach, known as the median-cut algorithm, yields even better results but also needs more computation time. This technique recursively fits a box around all colors used in the RGB colorspace which it splits at the median value of its longest side. The algorithm stops after 255 recursions. All colors in one box are mapped to the centroid of this box.

All above techniques restrict the number of displayed colors to 256. A technique of achieving additional colors is to apply a variation of half-toning used for gray scale images, thus increasing the color resolution at the cost of spatial resolution. The 256 values of the colormap are divided into four sections containing 64 different values of red, green, blue and white. As can be seen in Figure 1, a 2×2 pixel area is grouped together to represent one composite color, each of the four pixels displays either one of the primary colors or white. In this way, the number of possible colors is increased from 256 to 64^4 .




Fig 4.3: A 2×2 pixel area displaying one composite color

4.4.2 CONVOLUTION

Convolution is a simple mathematical operation which is fundamental to many common

image processing operators. Convolution provides a way of 'multiplying together' two arrays of numbers, generally of different sizes, but of the same dimensionality, to produce a third array of numbers of the same dimensionality. This can be used in image processing to implement operators whose output pixel values are simple linear combinations of certain input pixel values.

In an image processing context, one of the input arrays is normally just a graylevel image. The second array is usually much smaller, and is also two-dimensional (although it may be just a single pixel thick), and is known as the kernel. Figure 1 shows an example image and kernel that we will use to illustrate convolution.

I₁₁	I₁₂	I₁₃	I₁₄	I₁₅	I₁₆	I₁₇	I₁₈	I₁₉
I₂₁	I₂₂	I₂₃	I₂₄	I₂₅	I₂₆	I₂₇	I₂₈	I₂₉
I₃₁	I₃₂	I₃₃	I₃₄	I₃₅	I₃₆	I₃₇	I₃₈	I₃₉
I₄₁	I₄₂	I₄₃	I₄₄	I₄₅	I₄₆	I₄₇	I₄₈	I₄₉
I₅₁	I₅₂	I₅₃	I₅₄	I₅₅	I₅₆	I₅₇	I₅₈	I₅₉
I₆₁	I₆₂	I₆₃	I₆₄	I₆₅	I₆₆	I₆₇	I₆₈	I₆₉

K₁₁	K₁₂	K₁₃
K₂₁	K₂₂	K₂₃

Fig 4.4 Convolution illustration

Figure 4.4 An example small image (left) and kernel (right) to illustrate convolution. The labels within each grid square are used to identify each square.

The convolution is performed by sliding the kernel over the image, generally starting at the top left corner, so as to move the kernel through all the positions where the kernel fits entirely within the boundaries of the image. (Note that implementations differ in what they do at the edges of images, as explained below.) Each kernel position corresponds to a single output pixel, the value of which is calculated by multiplying together the

kernel value and the underlying image pixel value for each of the cells in the kernel, and then adding all these numbers together.

So, in our example, the value of the bottom right pixel in the output image will be given by:

$$O_{57} = I_{57}K_{11} + I_{58}K_{12} + I_{59}K_{13} + I_{67}K_{21} + I_{68}K_{22} + I_{69}K_{23}$$

If the image has M rows and N columns, and the kernel has m rows and n columns, then the size of the output image will have M - m + 1 rows, and N - n + 1 columns.

Mathematically we can write the convolution as:

$$O(i, j) = \sum_{k=1}^m \sum_{l=1}^n I(i + k - 1, j + l - 1)K(k, l)$$

where i runs from 1 to M - m + 1 and j runs from 1 to N - n + 1.

Note that many implementations of convolution produce a larger output image than this because they relax the constraint that the kernel can only be moved to positions where it fits entirely within the image. Instead, these implementations typically slide the kernel to all positions where just the top left corner of the kernel is within the image. Therefore the kernel 'overlaps' the image on the bottom and right edges. One advantage of this approach is that the output image is the same size as the input image. Unfortunately, in order to calculate the output pixel values for the bottom and right edges of the image, it is necessary to invent input pixel values for places where the kernel extends off the end of the image. Typically pixel values of zero are chosen for regions outside the true image, but this can often distort the output image at these places. Therefore in general if you are using a convolution implementation that does this, it is better to clip the image to remove these spurious regions. Removing n - 1 pixels from the right hand side and m - 1 pixels from the bottom will fix things.

Convolution can be used to implement many different operators, particularly spatial filters and feature detectors. Examples include Gaussian smoothing and the Sobel

edge detector.

Distance Metrics

It is often useful in image processing to be able to calculate the distance between two pixels in an image, but this is not as straightforward as it seems. The presence of the pixel grid makes several so-called distance metrics possible which often give different answers to each other for the distance between the same pair of points. We consider the three most important ones.

Euclidean Distance

This is the familiar straight line distance that most people are familiar with. If the two pixels that we are considering have coordinates (x_1, y_1) and (x_2, y_2) , then the Euclidean distance is given by:

$$D_{Euclid} = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$$

City Block Distance

Also known as the Manhattan distance. This metric assumes that in going from one pixel to the other it is only possible to travel directly along pixel grid lines. Diagonal moves are not allowed. Therefore the 'city block' distance is given by:

$$D_{City} = |x_2 - x_1| + |y_2 - y_1|$$

Chessboard Distance

This metric assumes that you can make moves on the pixel grid as if you were a King making moves in chess, *i.e.* a diagonal move counts the same as a horizontal move. This means that the metric is given by:

$$D_{Chess} = \max(|x_2 - x_1|, |y_2 - y_1|)$$

Note that the last two metrics are usually much faster to compute than the Euclidean metric and so are sometimes used where speed is critical but accuracy is not too

important.

4.4.3 DITHERING

Dithering is an image display technique that is useful for overcoming limited display resources. The word dither refers to a random or semi-random perturbation of the pixel values.

Two applications of this techniques are particularly useful:

● Low quantization display: When images are quantized to a few bits (e.g. 3) then only a limited number of graylevels are used in the display of the image. If the scene is smoothly shaded, then the image display will generate rather distinct boundaries around the edges of image regions when the original scene intensity moves from one quantization level to the next. To eliminate this effect, one dithering technique adds random noise (with a small range of values) to the input signal before quantization into the output range. This randomizes the quantization of the pixels at the original quantization boundary, and thus pixels make a more gradual transition from neighborhoods containing 100% of the first quantization level to neighborhoods containing 100% of the second quantization level.

● Limited color display: When fewer colors are able to be displayed (e.g. 256) than are present in the input image (e.g. 24 bit color), then patterns of adjacent pixels are used to simulate the appearance of the unrepresented colors.

4.4.4 EDGE DETECTORS

Edges are places in the image with strong intensity contrast. Since edges often occur at image locations representing object boundaries, edge detection is extensively used in image segmentation when we want to divide the image into areas corresponding to different objects. Representing an image by its edges has the further advantage that the amount of data is reduced significantly while retaining most of the image information.

Since edges consist of mainly high frequencies, we can, in theory, detect edges by applying a highpass frequency filter in the Fourier domain or by convolving the image with an appropriate kernel in the spatial domain. In practice, edge detection is performed in the spatial domain, because it is computationally less expensive and often yields better results.

Since edges correspond to strong illumination gradients, we can highlight them by calculating the derivatives of the image. This is illustrated for the one-dimensional case in Figure 1.




Fig 4.5: 1st and 2nd derivative of an edge illustrated in one dimension.

We can see that the position of the edge can be estimated with the maximum of the 1st derivative or with the zero-crossing of the 2nd derivative. Therefore we want to find a technique to calculate the derivative of a two-dimensional image. For a discrete one-dimensional function $f(i)$, the first derivative can be approximated by

$$\frac{df(i)}{d(i)} = f(i+1) - f(i)$$

Calculating this formula is equivalent to convolving the function with [-1 1]. Similarly the 2nd derivative can be estimated by convolving f(i) with [1 -2 1].

Different edge detection kernels which are based on the above formula enable us to calculate either the 1st or the 2nd derivative of a two-dimensional image. There are two common approaches to estimate the 1st derivative in a two-dimensional image, Prewitt compass edge detection and gradient edge detection.

Prewitt compass edge detection involves convolving the image with a set of (usually 8) kernels, each of which is sensitive to a different edge orientation. The kernel producing the maximum response at a pixel location determines the edge magnitude and orientation. Different sets of kernels might be used: examples include Prewitt, Sobel, Kirsch and Robinson kernels.

Gradient edge detection is the second and more widely used technique. Here, the image is convolved with only two kernels, one estimating the gradient in the x-direction, G_x , the other the gradient in the y-direction, G_y . The absolute gradient magnitude is then given by

$$|G| = \sqrt{G_x^2 + G_y^2}$$

and is often approximated with

$$|G| = |G_x| + |G_y|$$

In many implementations, the gradient magnitude is the only output of a gradient edge detector, however the edge orientation might be calculated with

The most common kernels used for the gradient edge detector are the Sobel, Roberts Cross and Prewitt operators.

After having calculated the magnitude of the 1st derivative, we now have to identify

those pixels corresponding to an edge. The easiest way is to threshold the gradient image, assuming that all pixels having a local gradient above the threshold must represent an edge. An alternative technique is to look for local maxima in the gradient image, thus producing one pixel wide edges. A more sophisticated technique is used by the Canny edge detector. It first applies a gradient edge detector to the image and then finds the edge pixels using non-maximal suppression and hysteresis tracking.

An operator based on the 2nd derivative of an image is the Marr edge detector, also known as zero crossing detector. Here, the 2nd derivative is calculated using a Laplacian of Gaussian (LoG) filter. The Laplacian has the advantage that it is an isotropic measure of the 2nd derivative of an image, i.e. the edge magnitude is obtained independently from the edge orientation by convolving the image with only one kernel. The edge positions are then given by the zero-crossings in the LoG image. The scale of the edges which are to be detected can be controlled by changing the variance of the Gaussian.

A general problem for edge detection is its sensitivity to noise, the reason being that calculating the derivative in the spatial domain corresponds to accentuating high frequencies and hence magnifying noise. This problem is addressed in the Canny and Marr operators by convolving the image with a smoothing operator (Gaussian) before calculating the derivative.

Frequency Domain

For simplicity, assume that the image I being considered is formed by projection from scene S (which might be a two- or three-dimensional scene, etc.).

The frequency domain is a space in which each image value at image position F represents the amount that the intensity values in image I vary over a specific distance related to F . In the frequency domain, changes in image position correspond to changes in the spatial frequency, (or the rate at which image intensity values) are changing in the spatial domain image I .

For example, suppose that there is the value 20 at the point that represents the frequency 0.1 (or 1 period every 10 pixels). This means that in the corresponding spatial domain image the intensity values vary from dark to light and back to dark over a distance of 10 pixels, and that the contrast between the lightest and darkest is 40 gray levels (2 times 20).

The spatial frequency domain is interesting because: 1) it may make explicit periodic relationships in the spatial domain, and 2) some image processing operators are more efficient or indeed only practical when applied in the frequency domain.

In most cases, the Fourier Transform is used to convert images from the spatial domain into the frequency domain and vice-versa.

A related term used in this context is spatial frequency, which refers to the (inverse of the) periodicity with which the image intensity values change. Image features with high spatial frequency (such as edges) are those that change greatly in intensity over short image distances.

Grayscale Images

A grayscale (or graylevel) image is simply one in which the only colors are shades of gray. The reason for differentiating such images from any other sort of color image is that less information needs to be provided for each pixel. In fact a 'gray' color is one in which the red, green and blue components all have equal intensity in RGB space, and so it is only necessary to specify a single intensity value for each pixel, as opposed to the three intensities needed to specify each pixel in a full color image.

Often, the grayscale intensity is stored as an 8-bit integer giving 256 possible different shades of gray from black to white. If the levels are evenly spaced then the difference between successive graylevels is significantly better than the graylevel resolving power of the human eye.

Grayscale images are very common, in part because much of today's display and image capture hardware can only support 8-bit images. In addition, grayscale images

are entirely sufficient for many tasks and so there is no need to use more complicated and harder-to-process color images.

Image Editing Software

There is a huge variety of software for manipulating images in various ways. Much of this software can be grouped under the heading image processing software, and the bulk of this reference is concerned with that group.

Another very important category is what we call image editing software. This group includes painting programs, graphic art packages and so on. They are often useful in conjunction with image processing software packages, in situations where direct immediate interaction with an image is the easiest way of achieving something. For instance, if a region of an image is to be masked out for subsequent image processing, it may be easiest to create the mask using an art package by directly drawing on top of the original image. The mask used in the description of the AND operator was created this way for instance. Art packages also often allow the user to move sections of the images around and brighten or darken selected regions interactively. Few dedicated image processing packages offer the same flexibility and ease of use in this respect.

Idempotence

Some operators have the special property that applying them more than once to the same image produces no further change after the first application. Such operators are said to be idempotent. Examples include the morphological operators opening and closing.

Isotropic Operators

An isotropic operator in an image processing context is one which applies equally well in all directions in an image, with no particular sensitivity or bias towards one particular set of directions (e.g. compass directions). A typical example is the zero crossing edge detector which responds equally well to edges in any orientation. Another example is Gaussian smoothing. It should be borne in mind that although an

operator might be isotropic in theory, the actual implementation of it for use on a discrete pixel grid may not be perfectly isotropic. ***An example of this is a Gaussian smoothing filter with very small standard deviation on a square grid.***

Kernel

A kernel is a (usually) smallish matrix of numbers that is used in image convolutions. Differently sized kernels containing different patterns of numbers give rise to different results under convolution. For instance, Figure 1 shows a 3×3 kernel that implements a mean filter.

1	1	1	Set of coordinate points = { (-1, -1), (0, -1), (1, -1), (-1, 0), (0, 0), (1, 0), (-1, 1), (0, 1), (1, 1) }
1	1	1	
1	1	1	

Fig 4.6: Convolution kernel for a mean filter with 3×3 neighborhood.

The word 'kernel' is also commonly used as a synonym for 'structuring element', which is a similar object used in mathematical morphology. A structuring element differs from a kernel in that it also has a specified origin. This sense of the word 'kernel' is not used in HIPR.

Logical Operators

Logical operators are generally derived from Boolean algebra, which is a mathematical way of manipulating the truth values of concepts in an abstract way without bothering about what the concepts actually mean. The truth value of a concept in Boolean value can have just one of two possible values: true or false. Boolean algebra allows you to represent things like:

The block is both red and large

by something like:

$A \text{ AND } B$

where A represents 'The block is red', and B represents 'The block is large'. Now each of these sub-phrases has its own truth value in any given situation: each sub-phrase is either true or false. Moreover, the entire composite phrase also has a truth value: it is true if both of the sub-phrases are true, and false in any other case. We can write this AND combination rule (and its dual operation NAND) using a truth-table as shown in Figure 1, in which we conventionally represent true by 1, and false by zero.

A	B	Q
0	0	0
0	1	0
1	0	0
1	1	1

AND

A	B	Q
0	0	1
0	1	1
1	0	1
1	1	0

NAND

Fig 4.7: Truth-tables for AND and NAND

The left hand table shows each of the possible combinations of truth values of A and B, and the the resulting truth value of A AND B. Similar truth-tables can be set up for the other logical operators: NAND, OR, NOR, XOR, XNOR and NOT.

Turning now to an image processing context, the pixel values in a binary image, which are either 0 or 1, can be interpreted as truth values as above. Using this convention we can carry out logical operations on images simply by applying the truth-table combination rules to the pixel values from a pair of input images (or a single input image in the case of NOT). Normally, corresponding pixels from each of two identically sized

binary input images are compared to produce the output image, which is another binary image of the same size. As with other image arithmetic operations, it is also possible to logically combine a single input image with a constant logical value, in which case each pixel in the input image is compared to the same constant in order to produce the corresponding output pixel. See the individual logical operator descriptions for examples of these operations.

Logical operations can also be carried out on images with integer pixel values. In this extension the logical operations are normally carried out in bitwise fashion on binary representations of those integers, comparing corresponding bits with corresponding bits to produce the output pixel value. For instance, suppose that we wish to XOR the integers 47 and 255 together using 8-bit integers. 47 is 00101111 in binary and 255 is 11111111. XORing these together in bitwise fashion, we have 11010000 in binary or 208 in decimal.

Note that not all implementations of logical operators work in such bitwise fashion. For instance some will treat zero as false and any non-zero value as true and will then apply the conventional 1-bit logical functions to derive the output image. The output may be a simple binary image itself, or it may be a graylevel image formed perhaps by multiplying what would be the binary output image (containing 0's and 1's) with one of the input images.

Look-up Tables and Colormaps

Look-Up Tables or LUTs are fundamental to many aspects of image processing. An LUT is simply a table of cross-references linking index numbers to output values. The most common use is to determine the colors and intensity values with which a particular image will be displayed, and in this context the LUT is often called simply a colormap.

The idea behind the colormap is that instead of storing a definite color for each pixel in an image, for instance in 24-bit RGB format, each pixel's value is instead treated as an index number into the colormap. When the image is to be displayed or otherwise processed, the colormap is used to look up the actual colors corresponding to each

index number. Typically, the output values stored in the LUT would be RGB color values.

There are two main advantages to doing things this way. Firstly, the index number can be made to use fewer bits than the output value in order to save storage space. For instance an 8-bit index number can be used to look up a 24-bit RGB color value in the LUT. Since only the 8-bit index number needs to be stored for each pixel, such 8-bit color images take up less space than a full 24-bit image of the same size. Of course the image can only contain 256 different colors (the number of entries in an 8-bit LUT), but this is sufficient for many applications and usually the observable image degradation is small.

Secondly the use of a color table allows the user to experiment easily with different color labeling schemes for an image.

One disadvantage of using a colormap is that it introduces additional complexity into an image format. It is usually necessary for each image to carry around its own colormap, and this LUT must be continually consulted whenever the image is displayed or processed.

Another problem is that in order to convert from a full color image to (say) an 8-bit color image using a color image, it is usually necessary to throw away many of the original colors, a process known as color quantization. This process is lossy, and hence the image quality is degraded during the quantization process. Additionally, when performing further image processing on such images, it is frequently necessary to generate a new colormap for the new images, which involves further color quantization, and hence further image degradation.

As well as their use in colormaps, LUTs are often used to remap the pixel values within an image. This is the basis of many common image processing point operations such as thresholding, gamma correction and contrast stretching. The process is often referred to as anamorphosis.

4.4.5 MASKING

A mask is a binary image consisting of zero- and non-zero values. If a mask is applied to another binary or to a grayscale image of the same size, all pixels which are zero in the mask are set to zero in the output image. All others remain unchanged.

Masking can be implemented either using pixel multiplication or logical AND, the latter in general being faster.

Masking is often used to restrict a point or arithmetic operator to an area defined by the mask. We can, for example, accomplish this by first masking the desired area in the input image and processing it with the operator, then masking the original input image with the inverted mask to obtain the unprocessed area of the image and finally recombining the two partial images using image addition. An example can be seen in the worksheet on the logical AND operator. In some image processing packages, a mask can directly be defined as an optional input to a point operator, so that automatically the operator is only applied to the pixels defined by the mask .

Mathematical Morphology

The field of mathematical morphology contributes a wide range of operators to image processing, all based around a few simple mathematical concepts from set theory. The operators are particularly useful for the analysis of binary images and common usages include edge detection, noise removal, image enhancement and image segmentation.

The two most basic operations in mathematical morphology are erosion and dilation. Both of these operators take two pieces of data as input: an image to be eroded or dilated, and a structuring element (also known as a kernel). The two pieces of input data are each treated as representing sets of coordinates in a way that is slightly different for binary and grayscale images.

For a binary image, white pixels are normally taken to represent foreground regions, while black pixels denote background. (Note that in some implementations this convention is reversed, and so it is very important to set up input images with the

correct polarity for the implementation being used). Then the set of coordinates corresponding to that image is simply the set of two-dimensional Euclidean coordinates of all the foreground pixels in the image, with an origin normally taken in one of the corners so that all coordinates have positive elements.

For a grayscale image, the intensity value is taken to represent height above a base plane, so that the grayscale image represents a surface in three-dimensional Euclidean space. Figure 1 shows such a surface. Then the set of coordinates associated with this image surface is simply the set of three-dimensional Euclidean coordinates of all the points within this surface and also all points below the surface, down to the base plane. Note that even when we are only considering points with integer coordinates, this is a lot of points, so usually algorithms are employed that do not need to consider all the points.




Fig 4.8: Simple graylevel image and the corresponding surface in image space

The structuring element is already just a set of point coordinates (although it is often represented as a binary image). It differs from the input image coordinate set in that it is normally much smaller, and its coordinate origin is often not in a corner, so that some coordinate elements will have negative values. Note that in many implementations of morphological operators, the structuring element is assumed to be a particular shape (e.g. a 3×3 square) and so is hardwired into the algorithm.

Binary morphology can be seen as a special case of graylevel morphology in which the

input image has only two graylevels at values 0 and 1.

Erosion and dilation work (at least conceptually) by translating the structuring element to various points in the input image, and examining the intersection between the translated kernel coordinates and the input image coordinates. For instance, in the case of erosion, the output coordinate set consists of just those points to which the origin of the structuring element can be translated, while the element still remains entirely 'within' the input image.

Virtually all other mathematical morphology operators can be defined in terms of combinations of erosion and dilation along with set operators such as intersection and union. Some of the more important are opening, closing and skeletonization.

Multi-spectral Images

A multi-spectral image is a collection of several monochrome images of the same scene, each of them taken with a different sensor. Each image is referred to as a band. A well known multi-spectral (or multi-band image) is a RGB color image, consisting of a red, a green and a blue image, each of them taken with a sensor sensitive to a different wavelength. In image processing, multi-spectral images are most commonly used for Remote Sensing applications. Satellites usually take several images from frequency bands in the visual and non-visual range. Landsat 5, for example, produces 7 band images with the wavelength of the bands being between 450 and 1250 nm.

All the standard single-band image processing operators can also be applied to multi-spectral images by processing each band separately. For example, a multi-spectral image can be edge detected by finding the edges in each band and then ORing the three edge images together. However, we would obtain more reliable edges, if we associate a pixel with an edge based on its properties in all three bands and not only in one.

To fully exploit the additional information which is contained in the multiple bands, we should consider the images as one multi-spectral image rather than as a set of

monochrome graylevel images. For an image with k bands, we can then describe the brightness of each pixel as a point in a k -dimensional space represented by a vector of length k .

Special techniques exist to process multi-spectral images. For example, to classify a pixel as belonging to one particular region, its intensities in the different bands are said to form a feature vector describing its location in the k -dimensional feature space. The simplest way to define a class is to choose a upper and lower threshold for each band, thus producing a k -dimensional 'hyper-cube' in the feature space. Only if the feature vector of a pixel points to a location within this cube, is the pixel classified as belonging to this class. A more sophisticated classification method is described in the corresponding worksheet.

The disadvantage of multi-spectral images is that, since we have to process additional data, the required computation time and memory increase significantly. However, since the speed of the hardware will increase and the costs for memory will decrease in the future, it can be expected that multi-spectral images will become more important in many fields of computer vision.

Non-linear Filtering

Suppose that an image processing operator F acting on two input images A and B produces output images C and D respectively. If the operator F is linear, then

$$F(a \times A + b \times B) = a \times C + b \times D$$

where a and b are constants. In practice, this means that each pixel in the output of a linear operator is the weighted sum of a set of pixels in the input image.

By contrast, non-linear operators are all the other operators. For example, the threshold operator is non-linear, because individually, corresponding pixels in the two images A and B may be below the threshold, whereas the pixel obtained by adding A and B may be above threshold. Similarly, an absolute value operation is non-linear:

$$|-1 + 1| \neq |-1| + |1|$$

as is the exponential operator:

$$\exp(1 + 1) \neq \exp(1) + \exp(1)$$

4.4.6 PIXELS

In order for any digital computer processing to be carried out on an image, it must first be stored within the computer in a suitable form that can be manipulated by a computer program. The most practical way of doing this is to divide the image up into a collection of discrete (and usually small) cells, which are known as pixels. Most commonly, the image is divided up into a rectangular grid of pixels, so that each pixel is itself a small rectangle. Once this has been done, each pixel is given a pixel value that represents the color of that pixel. It is assumed that the whole pixel is the same color, and so any color variation that did exist within the area of the pixel before the image was discretized is lost. However, if the area of each pixel is very small, then the discrete nature of the image is often not visible to the human eye.

Other pixel shapes and formations can be used, most notably the hexagonal grid, in which each pixel is a small hexagon. This has some advantages in image processing, including the fact that pixel connectivity is less ambiguously defined than with a square grid, but hexagonal grids are not widely used. Part of the reason is that many image capture systems (e.g. most CCD cameras and scanners) intrinsically discretize the captured image into a rectangular grid in the first instance.

Pixel Connectivity

The notation of pixel connectivity describes a relation between two or more pixels. For two pixels to be connected they have to fulfill certain conditions on the pixel brightness and spatial adjacency.

First, in order for two pixels to be considered connected, their pixel values must both be from the same set of values V . For a grayscale image, V might be any range of graylevels, e.g. $V=\{22,23,\dots,40\}$, for a binary image we simply have $V=\{1\}$.

To formulate the adjacency criterion for connectivity, we first introduce the notation of neighborhood. For a pixel p with the coordinates (x,y) the set of pixels given by:

$$N_4(p) = \{(x+1, y), (x-1, y), (x, y+1), (x, y-1)\}$$

is called its 4-neighbors. Its 8-neighbors are defined as

$$N_8(p) = N_4 \cup \{(x+1, y+1), (x+1, y-1), (x-1, y+1), (x-1, y-1)\}$$

From this we can infer the definition for 4- and 8-connectivity:

Two pixels p and q , both having values from a set V are 4-connected if q is from the set $N_4(p)$ and 8-connected if q is from $N_8(p)$.

General connectivity can either be based on 4- or 8-connectivity; for the following discussion we use 4-connectivity.

A pixel p is connected to a pixel q if p is 4-connected to q or if p is 4-connected to a third pixel which itself is connected to q . Or, in other words, two pixels q and p are connected if there is a path from p and q on which each pixel is 4-connected to the next one.

A set of pixels in an image which are all connected to each other is called a connected component. Finding all connected components in an image and marking each of them with a distinctive label is called connected component labeling.

An example of a binary image with two connected components which are based on 4-connectivity can be seen in Figure 1. If the connectivity were based on 8-neighbors, the two connected components would merge into one.




Fig 4.9: Two connected components based on 4-connectivity.

Pixel Values

Each of the pixels that represents an image stored inside a computer has a pixel value which describes how bright that pixel is, and/or what color it should be. In the simplest case of binary images, the pixel value is a 1-bit number indicating either foreground or background. For a grayscale images, the pixel value is a single number that represents the brightness of the pixel. The most common pixel format is the byte image, where this number is stored as an 8-bit integer giving a range of possible values from 0 to 255. Typically zero is taken to be black, and 255 is taken to be white. Values in between make up the different shades of gray.

To represent color images, separate red, green and blue components must be specified for each pixel (assuming an RGB colorspace), and so the pixel 'value' is actually a vector of three numbers. Often the three different components are stored as three separate 'grayscale' images known as color planes (one for each of red, green and blue), which have to be recombined when displaying or processing.

Multi-spectral images can contain even more than three components for each pixel, and by extension these are stored in the same kind of way, as a vector pixel value, or as separate color planes.

The actual grayscale or color component intensities for each pixel may not actually be stored explicitly. Often, all that is stored for each pixel is an index into a colormap in

which the actual intensity or colors can be looked up.

Although simple 8-bit integers or vectors of 8-bit integers are the most common sorts of pixel values used, some image formats support different types of value, for instance 32-bit signed integers or floating point values. Such values are extremely useful in image processing as they allow processing to be carried out on the image where the resulting pixel values are not necessarily 8-bit integers. If this approach is used then it is usually necessary to set up a colormap which relates particular ranges of pixel values to particular displayed colors.

Primary Colors

It is a useful fact that the huge variety of colors that can be perceived by humans can all be produced simply by adding together appropriate amounts of red, blue and green colors. These colors are known as the primary colors. Thus in most image processing applications, colors are represented by specifying separate intensity values for red, green and blue components. This representation is commonly referred to as RGB.

The primary color phenomenon results from the fact that humans have three different sorts of color receptors in their retinas which are each most sensitive to different visible light wavelengths.

The primary colors used in painting (red, yellow and blue) are different. When paints are mixed, the 'addition' of a new color paint actually subtracts wavelengths from the reflected visible light.

RGB and Colorspaces

A color perceived by the human eye can be defined by a linear combination of the three primary colors red, green and blue. These three colors form the basis for the RGB-colorspace. Hence, each perceivable color can be defined by a vector in the three-dimensional colorspace. The intensity is given by the length of the vector, and the actual color by the two angles describing the orientation of the vector in the colorspace.

The RGB-space can also be transformed into other coordinate systems, which might be more useful for some applications. One common basis for the color space is IHS. In this coordinate system, a color is described by its intensity, hue (average wavelength) and saturation (the amount of white in the color). This color space makes it easier to directly derive the intensity and color of perceived light and is therefore more likely to be used by human beings.

Spatial Domain

For simplicity, assume that the image I being considered is formed by projection from scene S (which might be a two- or three-dimensional scene, etc.).

The spatial domain is the normal image space, in which a change in position in I directly projects to a change in position in S . Distances in I (in pixels) correspond to real distances (e.g. in meters) in S .

This concept is used most often when discussing the frequency with which image values change, that is, over how many pixels does a cycle of periodically repeating intensity variations occur. One would refer to the number of pixels over which a pattern repeats (its periodicity) in the spatial domain.

In most cases, the Fourier Transform will be used to convert images from the spatial domain into the frequency domain.

A related term used in this context is spatial frequency, which refers to the (inverse of the) periodicity with which the image intensity values change. Image features with high spatial frequency (such as edges) are those that change greatly in intensity over short image distances.

Another term used in this context is spatial derivative, which refers to how much the image intensity values change per change in image position.

Structuring Elements

The field of mathematical morphology provides a number of important image processing operations, including erosion, dilation, opening and closing. All these morphological operators take two pieces of data as input. One is the input image, which may be either binary or grayscale for most of the operators. The other is the structuring element. It is this that determines the precise details of the effect of the operator on the image.

The structuring element is sometimes called the kernel, but we reserve that term for the similar objects used in convolutions.

The structuring element consists of a pattern specified as the coordinates of a number of discrete points relative to some origin. Normally cartesian coordinates are used and so a convenient way of representing the element is as a small image on a rectangular grid. Figure 1 shows a number of different structuring elements of various sizes. In each case the origin is marked by a ring around that point. The origin does not have to be in the center of the structuring element, but often it is. As suggested by the figure, structuring elements that fit into a 3×3 grid with its origin at the center are the most commonly seen type.




Fig 4.10: Some example structuring elements.

Note that each point in the structuring element may have a value. In the simplest structuring elements used with binary images for operations such as erosion, the

elements only have one value, conveniently represented as a one. More complicated elements, such as those used with thinning or grayscale morphological operations, may have other pixel values.

An important point to note is that although a rectangular grid is used to represent the structuring element, not every point in that grid is part of the structuring element in general. Hence the elements shown in Figure 1 contain some blanks. In many texts, these blanks are represented as zeros, but this can be confusing and so we avoid it here.

When a morphological operation is carried out, the origin of the structuring element is typically translated to each pixel position in the image in turn, and then the points within the translated structuring element are compared with the underlying image pixel values. The details of this comparison, and the effect of the outcome depend on which morphological operator is being used.

Wrapping and Saturation

If an image is represented in a byte or integer pixel format, the maximum pixel value is limited by the number of bits used for the representation, e.g. the pixel values of a 8-bit image are limited to 255.

However, many image processing operations produce output values which are likely to exceed the given maximum value. In such cases, we have to decide how to handle this pixel overflow.

One possibility is to wrap around the overflowing pixel values. This means that if a value is greater than the possible maximum, we subtract the pixel value range so that the value starts again from the possible minimum value. Figure 1 shows the mapping function for wrapping the output values of some operation into an 8-bit format.




Fig 4.11: Mapping function for wrapping the pixel values of an 8-bit image.

Another possibility is to set all overflowing pixels to the maximum possible values --- an effect known as saturation. The corresponding mapping function for an 8-bit image can be seen in Figure 2.




Fig 4.12: Mapping function for saturating an 8-bit image.

If only a few pixels in the image exceed the maximum value it is often better to apply the latter technique, especially if we use the image for display purposes. However, by setting all overflowing pixels to the same value we lose an essential amount of information. In the worst case, when all pixels exceed the maximum value, this would lead to an image of constant pixel values. Wrapping around overflowing pixel retains

the differences between values. On the other hand, it might cause the problem that pixel values passing the maximum 'jump' from the maximum to the minimum value. Examples for both techniques can be seen in the worksheets of various point operators.

If possible, it is easiest to change the image format, for example to float format, so that all pixel values can be represented. However, we should keep in mind that this implies an increase in processing time and memory.

CHAPTER 5

RESULTS AND DISCUSSION

The image is inputted into the code in form of jpg or any other acceptable format, the image is stored inside a folder within the project.

The image is then compared with the all the images in the data set, then the image in the data set that has more matches is shown. The comparison lines are also showed. The line which is closer to greener shade, has more similarity. The line whose color is close to red shade has lesser similarity.




Fig 5.1: Verification of Input




Fig 5.2: Comparison of Input with Trained Data




Fig 5.3: Output of the image




Fig 5.4: Input of a Fake note




Fig 5.5: Output of a Fake note

CHAPTER 6

CONCLUSION AND FUTURE WORK

6.1 CONCLUSION

Identification of currency note for visually challenged people have been made easy, along with detection of fake note, which have been a severe problem in today's world have been solved.

To get the best of accuracy, two to three methods are implemented on each module to attain maximum efficiency.

6.2 FUTUREWORK

The project can be further developed into a real time application or a camera feature add on, that takes a picture of the note and can determine the denomination of the note and to identify the real note. The application can be made as a offline application, that can work without the help of internet.

REFERENCES

- [1] “ Personalized End-to-End Mandarin Speech Synthesis using Small-sized Corpus”
- ChenHan Yuan (2020)
- [2] “Image Authentication based on Q-FMT via Blind Geometric Correction” - Ram
Kumar Karsh (2020)
- [3] “Saliency Consistency-Based Image Re-Colorization for Color Blindness” - Jinjiang
Li (2020)
- [4] “Edge Sensitive Unsupervised Image-to-Image Translation” - Ibrahim Batuhan
Akkaya (2020)
- [5] “Vocoder-free End-to-End Voice Conversion with Transformer Network” - June-
Woo Kim (2020)
- [6] “An Object Detection Technique For Blind People in Real-Time Using Deep Neural
Network” - Ashwani Kumar (2020)
- [7] “Tactile Graphics Exploration Studies Using Fingertip Tracking Based on Colour
Markers Detection for Visually Impaired People” - Muhammad Ikmal Hakim Shamsul
Bahrin (2019)
- [8] “A Hybrid Fake Banknote Detection model using OCR, Face Recognition and
Hough Features” - Adiba Zarin (2019)
- [9] “Fake Colorized Image Detection with Channel-wise Convolution based Deep-
learning Framework” - Long Zhuo (2019)
- [10] “Gradient-Based Illumination Description for Image Forgery Detection” - Falko
Matern (2019)

SOURCE CODE

detect.py:

```
from utils import *
from matplotlib import pyplot as plt
import pyttsx3
max_val = 8
max_pt = -1
max_kp = 0

orb = cv2.ORB_create()
# orb is an alternative to SIFT

#test_img = read_img('files/test_100_2.jpg')
test_img = read_img('files/test_50_1.jpg')
#test_img = read_img('files/download.png')
#test_img = read_img('files/test_100_3.jpg')
#test_img = read_img('files/test_20_4.jpg')
display('original', test_img)

# resizing must be dynamic
original = resize_img(test_img, 0.4)
#display('original', original)

# keypoints and descriptors
# (kp1, des1) = orb.detectAndCompute(test_img, None)
(kp1, des1) = orb.detectAndCompute(test_img, None)

training_set = ['files/20.jpg', 'files/50.jpg', 'files/100.jpg',
                'files/500.jpg']

for i in range(0, len(training_set)):
    # train image
    train_img = cv2.imread(training_set[i])

    (kp2, des2) = orb.detectAndCompute(train_img, None)

    # brute force matcher
    bf = cv2.BFMatcher()
    all_matches = bf.knnMatch(des1, des2, k=2)

    good = []
    # give an arbitrary number -> 0.789
    # if good -> append to list of good matches
    for (m, n) in all_matches:
        if m.distance < 0.789 * n.distance:
            good.append([m])

    if len(good) > max_val:
        max_val = len(good)
        max_pt = i
        max_kp = kp2

    print(i, ' ', training_set[i], ' ', len(good))
```

```

if max_val != 8:
    print(training_set[max_pt])
    print('good matches ', max_val)
    train_img = cv2.imread(training_set[max_pt])
    img3 = cv2.drawMatchesKnn(test_img, kp1, train_img, max_kp, good, 4)

    note = str(training_set[max_pt])[6:-4]
    print('\nDetected denomination: Rs. ', note)
    engine = pyttsx3.init()
    engine.say(note)
    engine.runAndWait()
    (plt.imshow(img3), plt.show())

else:
    print('No Matches')
    engine = pyttsx3.init()
    engine.say("this is fake note")
    engine.runAndWait()

```

utils.py:

```

import cv2
import math
import numpy as np
import matplotlib.pyplot as plt
from pprint import pprint

# read image as is
def read_img(file_name):
    img = cv2.imread(file_name)
    return img

# resize image with fixed aspect ratio
def resize_img(image, scale):
    res = cv2.resize(image, None, fx=scale, fy=scale, interpolation =
cv2.INTER_AREA)
    return res

# convert image to grayscale
def img_to_gray(image):
    img_gray = cv2.cvtColor(image, cv2.COLOR_RGB2GRAY)
    return img_gray

# gaussian blurred grayscale
def img_to_gaussian_gray(image):
    img_gray = cv2.GaussianBlur(img_to_gray(image), (5, 5), 0)
    return img_gray

# convert image to negative
def img_to_neg(image):

```

```

    img_neg = 255 - image
    return img_neg

# binarize (threshold)
# retval not used currently
def binary_thresh(image, threshold):
    retval, img_thresh = cv2.threshold(image, threshold, 255,
cv2.THRESH_BINARY)
    return img_thresh

# NO IDEA HOW THIS WORKS
def adaptive_thresh(image):
    img_thresh = cv2.adaptiveThreshold(image, 255,
cv2.ADAPTIVE_THRESH_GAUSSIAN_C, cv2.THRESH_BINARY, 11, 8)
    # cv2.adaptiveThreshold(src, maxValue, adaptiveMethod, thresholdType,
blockSize, C[, dst]) → dsta
    return img_thresh

# sobel edge operator
def sobel_edge(image, align):
    img_horiz = cv2.Sobel(image, cv2.CV_8U, 0, 1)
    img_vert = cv2.Sobel(image, cv2.CV_8U, 1, 0)
    if align == 'h':
        return img_horiz
    elif align == 'v':
        return img_vert
    else:
        print('use h or v')

# sobel edge x + y
def sobel_edge2(image):
    # ksize = size of extended sobel kernel
    grad_x = cv2.Sobel(image, cv2.CV_16S, 1, 0, ksize=3, borderType =
cv2.BORDER_DEFAULT)
    grad_y = cv2.Sobel(image, cv2.CV_16S, 0, 1, ksize=3, borderType =
cv2.BORDER_DEFAULT)

    abs_grad_x = cv2.convertScaleAbs(grad_x)
    abs_grad_y = cv2.convertScaleAbs(grad_y)

    dst = cv2.addWeighted(abs_grad_x, 0.5, abs_grad_y, 0.5, 0)
    return dst

# canny edge operator
def canny_edge(image, block_size, ksize):
    # block_size => Neighborhood size
    # ksize => Aperture parameter for the Sobel operator

    # 350, 350 => for smaller 500
    # 720, 350 => Devnagari 500, Reserve bank of India

    img = cv2.Canny(image, block_size, ksize)

```

```

    # dilate to fill up the numbers
    #img = cv2.dilate(img, None)
    return img

# laplacian edge
def laplacian_edge(image):
    # good for text
    img = cv2.Laplacian(image, cv2.CV_8U)
    return img

# detect countours
def find_contours(image):
    (_, contours, _) = cv2.findContours(image, cv2.RETR_LIST,
cv2.CHAIN_APPROX_SIMPLE)
    contours = sorted(contours, key = cv2.contourArea, reverse = True)[:5]
    return contours

# median blur
def median_blur(image):
    blurred_img = cv2.medianBlur(image, 3)
    return blurred_img

# dialte image to close lines
def dilate_img(image):
    img = cv2.dilate(image, np.ones((5,5), np.uint8))
    return img

# erode image
def close(image):
    img = cv2.Canny(image, 75, 300)
    img = cv2.dilate(img, None)
    img = cv2.erode(img, None)
    return img

def harris_edge(image):
    img_gray = np.float32(image)

    corners = cv2.goodFeaturesToTrack(img_gray, 4, 0.03, 200, None, None,
2,useHarrisDetector=True, k=0.04)
    corners = np.int0(corners)

    for corner in corners:
        x, y = corner.ravel()
        cv2.circle(image, (x, y), 3, 255, -1)
    return image

# calculate histogram
def histogram(image):
    hist = cv2.calcHist([image], [0], None, [256], [0, 256])
    # cv2.calcHist(images, channels, mask, histSize, ranges[, hist[,

```

```

accumulate]])
    plt.plot(hist)
    plt.show()

# fast fourier transform
def fourier(image):
    f = np.fft.fft2(image)
    fshift = np.fft.fftshift(f)
    magnitude_spectrum = 20 * np.log(np.abs(fshift))

    plt.subplot(121), plt.imshow(image, cmap='gray')
    plt.title('Input Image'), plt.xticks([]), plt.yticks([])

    plt.subplot(122), plt.imshow(magnitude_spectrum, cmap='gray')
    plt.title('FFT'), plt.xticks([]), plt.yticks([])

    plt.show()

# calculate scale and fit into display
def display(window_name, image):
    screen_res = 1440, 900 # MacBook Air

    scale_width = screen_res[0] / image.shape[1]
    scale_height = screen_res[1] / image.shape[0]
    scale = min(scale_width, scale_height)
    window_width = int(image.shape[1] * scale)
    window_height = int(image.shape[0] * scale)

    # rescale the resolution of the window
    cv2.namedWindow(window_name, cv2.WINDOW_NORMAL)
    cv2.resizeWindow(window_name, window_width, window_height)

    # display image
    cv2.imshow(window_name, image)

    # wait for any key to quit the program
    cv2.waitKey(0)
    cv2.destroyAllWindows()

```