Chapter 1

Introduction

Omnidirectional vision is the ability to see in all directions at the same time. Sensors that are able to achieve such omnidirectional vision (or at least come close to achieving it), offer several advantages to many areas of computer vision. One such area is that of tracking and surveillance, which benefits from the unobstructed views of the surroundings acquired by these sensors and allows for objects to be tracked simultaneously in different parts of the field-of-view without requiring camera motion.

There are several different types of omnidirectional sensors available, amongst them catadioptric systems, which use curved mirrors to achieve a hemispherical field-of-view or more. One of the advantages of these sensors is that of having a single viewpoint, which is often a requirement for many computer vision techniques. Catadioptric systems also present certain challenges because of their non-perspective projection and non-linear image geometry, and generally they have a low image resolution.

Tracking and surveillance is a very active field of research in computer vision, and many different methods have been studied. Normally, the first stage of the object tracking process consists of some form of motion detection, and in the case of stationary cameras, the preferred method is that of background subtraction, where motion detection is performed by detecting moving objects against a known background consisting of the stationary elements of the scene. The result of the motion detection phase is then used for locating objects and tracking them over time. The traditional way of performing this is to group the moving pixels of an image and then to match these pixel groups across frames, a process normally called 'blob' tracking. Recently, statistically based techniques are being used for object tracking because of their greater robustness to tracking problems such as object occlusion.

1.1 Aim

In this thesis, we propose to use an omnidirectional catadioptric camera for performing object tracking, taking advantage of the catadioptric camera's wide field-of-view. This allows objects to be potentially tracked from the time they first appear in the scene until they leave the scene – compared to limited field-of-view cameras, where objects, that are still present in the scene, are lost when they exit the field-of-view of the camera.

This thesis addresses the following issues:

1 Development of methods for **robust object tracking**. The methods must handle occlusion events and the presence of motion detection problems, like object fragmentation and shadows. For this issue, we investigate two different tracking approaches – a blob tracking method and a colour-based statistical tracking method.

2 A requirement of the tracking methods developed is that they should be **sufficiently generic to operate in both an indoor-based and outdoor-based environments**. These two types of environments normally provide quite diverse conditions from the point of view of object tracking. For this thesis, we use two datasets, which were created for the PETS workshops[1] to serve as standard test datasets for evaluating visual surveillance applications.

3 **Handling the non-linear image geometry** of the catadioptric camera. This includes presenting tracking results and images in a human-viewable form. A catadioptric camera with a single viewpoint allows for the creation of views as if generated by virtual perspective cameras.

   This issue also deals with the restrictions on the type of computer vision algorithms that can be used with omnidirectional images. Most of the techniques and algorithms in computer vision assume a perspective projection or linear image model, which are unsuitable in this case.

---

[1] Performance Evaluation of Tracking and Surveillance (PETS) workshops. More information on the datasets is given in §4.

1.2 Thesis Outline

Chapter 2 introduces the concept of omnidirectional vision, mentions the advantages obtained from having wide fields-of-view, and describes the different omnidirectional camera types currently available. This chapter defines the single viewpoint constraint as seeing the world from a single point in space and mentions why it is important for omnidirectional cameras to obey this constraint. The chapter concludes by mentioning that among the different types, catadioptric systems are the method of choice.

Chapter 3 describes catadioptric camera systems in more detail, and in particular examines the different mirror types that can be used. Of these, the paraboloidal mirror is found to be the best choice. This chapter also gives a review of current applications using catadioptric camera systems.

Chapter 4 introduces the application, called *OmniTracking*, developed for this thesis, describes the two different environments in which it is designed to work, and introduces the datasets used for testing the program. This chapter also discusses the advantages that the use of an omnidirectional camera provides to the application as well as the challenges it provides.

Chapter 5 examines the platform-specific issues that affect the application, such as operating system issues and thread control. It then introduces the OpenCV library used by the application and describes the processor's MMX technology. The chapter finishes off by describing in detail the structure of the application and the processing pipeline model adopted.

Chapter 6 introduces the equations that define the class of mirrors that generate a single viewpoint catadioptric system and derives the solution for the case of the paraboloidal mirror. Using these mirror equations, this chapter describes the idea of re-projecting rays from the omnidirectional image to generate virtual views. The chapter describes two virtual camera models (perspective and panoramic) and the related camera control mechanisms to allow the user to change the viewpoint of the virtual camera. These control mechanisms are also used by OmniTracking to auto-track targets.

Chapter 7 describes how catadioptric cameras can be calibrated and gives a review of the existing calibration methods. This is followed by a description of the calibration procedure adopted for the application of this thesis, together with the results obtained.

Chapter 8 starts by describing the different methods available for motion detection and then examines the background subtraction technique in more detail, in particular the different background models that can be used and how the background is updated. This chapter then describes the background subtraction method as implemented by the OmniTracking application and concludes with the results obtained.

Chapter 9 introduces the idea of tracking objects and lists the various problems that can be encountered during tracking. It provides an overview of the different tracking methods available in the literature and which of these are suitable to be used for omnidirectional images. The chapter finishes off by introducing the two methods adopted for this thesis.

Chapter 10 describes the implementation of the blob tracking method; mainly, the clustering method used for grouping connected components, the object features used for comparing objects and their related similarity measures. The chapter describes how objects are matched from frame to frame using a match scoring algorithm and temporal constraints and concludes with the results obtained from this method.

Chapter 11 describes the second tracking method, which is based on a probabilistic colour model expressed as a Gaussian mixture model. The theory underpinning mixture models and probabilistic tracking is described, together with how the Expectation-Maximisation (EM) algorithm can be used to learn the colour model. For updating the colour model at run-time, the chapter introduces a variant of the EM algorithm, called the Incremental EM algorithm. The chapter concludes with implementation details for improving the speed of the tracking algorithm and gives the results obtained.

Chapter 12 describes how the output from the tracking process can be used to estimate the 3D position of objects, given that objects move along a known ground-plane. The tracking results can also be used to automatically track objects using virtual cameras.

Chapter 13 concludes the thesis with a discussion of the results obtained and describes future work.