Unsupervised Recognition of Salient Colour for Real-Time Image Processing

David Budden[1,2] and Alexandre Mendes[3]

[1] National ICT Australia (NICTA), Victoria Research Lab
[2] Department of Electrical and Electronic Engineering
The University of Melbourne, Parkville, VIC 3052, Australia.
david.budden@unimelb.edu.au
[3] School of Electrical Engineering and Computer Science
The University of Newcastle, Callaghan, NSW 2308, Australia.
alexandre.mendes@newcastle.edu.au

Abstract. Humans have the subconscious ability to create simple abstractions from observations of their physical environment. The ability to consider the colour of an object in terms of "red" or "blue", rather than spatial distributions of reflected light wavelengths, is vital in processing and communicating information about important features within our local environment. The real-time identification of such features in image processing necessitates the software implementation of such a process; segmenting an image into regions of salient colour, and in doing so reducing the information stored and processed from 3-dimensional pixel values to a simple colour class label. This paper details a method by which colour segmentation may be performed offline and stored in a static look-up table, allowing for constant time dimensionality reduction in an arbitrary environment of coloured features. The machine learning framework requires no human supervision, and its performance is evaluated in terms of feature classification performance within a RoboCup robot soccer environment. The developed system is demonstrated to yield an 8% improvement over slower traditional methods of manual colour mapping.

Keywords: Computer vision, colour vision, robotics, RoboCup, LUT generation

1 Introduction

Szeliski describes image segmentation simply as the task of finding groups of pixels that "go together" [9]. This is an abstract notion that corresponds with an inherently subconscious human process: the ability to look at an image and identify salient features, such as a person, landmark or household item. Hundreds of algorithms exist for image segmentation, with most relying on the assignment of a *feature vector* (containing spatial and/or colour information) to each pixel. For real-time applications, the dimensionality of this feature vector is commonly reduced to contain only colour information, with each pixel assigned a colour *class label* corresponding with higher level notions of colour (such as "red" or