Digital Colour Processing

Brian V. Funt

School of Computing Science
Simon Fraser University, Burnaby, B.C.
Canada V5A 1S6
funt@sfu.ca

Abstract

The members of the Computational Vision Laboratory at Simon Fraser University have been studying colour for over a decade. I describe some of the main colour issues, the progress we have made in understanding them and the application of our methods to colour-based object recognition and digital photography.

Introduction

Brockton Point in Stanley Park is one of my favourite spots in Vancouver, the world for that matter, so I've decided to stop there and write this paper on my laptop. It's a beautiful sunny day, blue water, blue sky, fresh white snow on the mountains with a few fluffy clouds clinging to them. I've taken a short walk to the water, inhaled the salt air, and now back in my car, I open my laptop. I can hardly see the characters on the screen; perhaps this wasn't such a good plan after all! My eyes have adapted to the bright sunlight, so the screen which normally seems quite bright to me, is now very dim. I notice also that the white background that I'm typing against looks yellowish, almost a tinge of orange. A few minutes later as my eyes adapt to the relative darkness of my car's interior, the screen now appears a bit brighter, but it's still yellow-orange. "Why?" I ask myself.

It's questions like this that have kept me fascinated with colour since 1984. What is colour? How do we see colour? Why do we see colour—what's its use to us? What's the relationship between the physics of reflection and our perception of colour? How does the context in which we see a colour change its appearance? Why does knowledge not affect our colour perception (e.g., I know my screen is white, so why now does it still look yellow-orange to me despite this knowledge?) Why does the colour clothes look change from the store to home? Why do colours not look even more different than they do given how different the colours of the illuminating lights often are?

I cannot answer all these questions, but I would like to describe some of the results we have obtained in my lab at Simon Fraser University over the past several years. I will be describing work done jointly with my recent students: Graham Finlayson, Janet Dueck, Kobus Barnard, Vlad Cardei and Subho Chatterjee, Louis Brassard and my research assistant, Michael Brockington. This work builds on earlier work with Ian Harder, Brigitte Dorner, Mark Drew and Jian Ho.

I come to the field of colour from an Artificial Intelligence background, so I naturally take the view that colour perception can be explained and modelled as a computational process. I'm interested in understanding how people think and perceive and using that understanding to produce better colour. To me, constructing computational models of colour perception is a part of artificial intelligence because there is no way to measure colour that does not relate to human perception. Colour is a perceptual quantity, not a physical one.

What we speak of when we speak about colour is our experience, not a physical phenomenon. Of course there are underlying physical phenomena creating our experience, but we can experience the same physical phenomenon differently under different circumstances, very much like the way a word can have different meanings in different contexts. Similarly we may experience two different physical phenomena as appearing the same. Hence in general, there is no one-to-one mapping between the spectrum of the light reflected into our eye's (the underlying physical phenomenon) and the colour we will perceive it to be.

The Main Problem of Colour Perception

To explain colour perception, we must explain how it is that we see colours as relatively stable despite changes in the incident illumination. I make the assumption that colour—like the rest of visual perception—is there to give us information about the world, the surface properties of objects in particular, and so the stability and reliability of the information is important. The problem of colour stability arises because the light reaching our eyes from an object is the product of the object's surface reflectance and the spectrum of the light illuminating the object. We do not have direct access to the properties of the incident light, so somehow we must estimate them from the light we receive from the object. To make matters worse, our eyes only measure the spectrum at extremely low resolution.

A typical spectrometer's resolution is 1 nanometer, which means that it yields 301 samples in the visible range of 400-700nm. If we consider this to be 301 "pixels" then the eye is, if we picture the spectrum as a 1-dimensional 'image', zooming out and reducing it