Computer Vision (CE316 and CE866): Concluding Remarks

Adrian F. Clark
⟨*alien@essex.ac.uk*⟩

The material presented in these lectures has given you an overview of computer vision but not really much depth. There are many techniques that we have not considered, either because they do not fall neatly into the lecture series or because they are too complicated to present in the time available; and I have been economical with the mathematical detail in much of what I have talked about — my aim has been to give you an understanding of how and why techniques work. Of course, there is an awful lot of exciting stuff that can be done using only the computer vision techniques we have considered. In particular, the machine learning approaches outlined in the last couple of lectures hold much promise — but this is an area that is evolving rapidly, so expect much of that to change in a year or two.

Whenever you look at the world around you, think of the cues you are using to work out an object's shape, or distance, or texture. Not until researchers have worked out how to emulate these processes and assimilate them in some system does computer vision stand any real chance of competing with the human visual system. Trying to get there is a fascinating problem.

Colophon

One of the sadnesses of modern desktop publishing is that authors give very little information regarding how their documents appear or were prepared; it is as though they take little pride in their work. This author's research area is image processing and computer vision, and one type of image processing is controlling just how the little dots of toner appear on a printed page. Some care has therefore been expended to make these lecture notes easy to read online and in printed form, and the author would welcome feedback where you do not find this to be the case.

This document was prepared using `Emacs`, the One True Editor, and typeset using pdfLaTeX. The basic document style is `tufte-book`, which supports the large number of marginal figures and notes that you see. The body font is Bitstream Charter and the sans serif font `Droid Sans`. The `microtype` package is used to enhance TeX's already good spacing algorithm. British hyphenation is used, thanks to the hard work of Dominik Wujastyk and Graham Toal. The graphs and many of the figures were prepared using TikZ; I have unashamedly stolen others from the Web (with acknowledgements, a good lesson for when you write your project final report). Hyperlinks within the document were generated automatically using the excellent `hyperref` package due to Sebastian Rahtz and Heiko Oberdiek.