MBS 94-28

A Formal Approach to Color Constancy: The Recovery of Surface and Light Source Spectral Properties Using Bilinear Models

Michael D'Zmura, Geoffrey Iverson

This technical report describes our recent work on the problem of color constancy (D'Zmura, 1992; D'Zmura & Iverson, 1993a,b, 1994; Iverson & D'Zmura, 1994). Our formal approach to the problem emphasizes the recovery of the spectral properties of light sources and surfaces that combine to produce the reflected lights that reach the eye. This approach relies on bilinear models of the visual response to reflected lights, and our first aim in this report, after a brief review of color constancy, is to describe such bilinear models.

The recovery of surface and light source spectral properties using bilinear models had been studied primarily in the case in which a single light shines on a collection of surfaces (Sallstrom, 1973; Brill, 1978, 1979; Buchsbaum, 1980; Maloney & Wandell, 1986). Yet lights and surfaces play symmetric roles in bilinear models, and the general case is one in which multiple lights shine sequentially on multiple surfaces (D'Zmura, 1992). each illuminant provides a different view of the surfaces, in analogy to work on structure from motion (Ullman, 1979a,b).

We worked to determine the circumstances under which recovery of spectral properties using bilinear models is possible. an important issue is uniqueness, which involves problems of identifiability similar to those met in other disciplines. We describe numerical and analytical methods for determining whether particular bilinear models, when presented a certain number of views of a certain number of surfaces, relate visual data uniquely to light and surface spectral properties. A unique relation between data and scenes is required for spectral recovery. We discuss uniqueness in the context of two stage recovery algorithms, in which the spectral properties of surfaces (or light sources) are recovered first and then used in a second stage to recover the spectral properties of light sources (or surfaces). Results support the intuition that information provided by additional views of surfaces lets a visual system determine a higher-dimensional, more accurate description of surface spectral properties (D'Zmura, 1992; D'Zmura & Iverson, 1993a,b; Iverson & D'Zmura, 1994).

We then turn to a general recovery algorithm that uses bilinear models to recover the spectral properties of lights and surfaces simultaneously (D'Zmura & Iverson, 1994). The algorithm is applicable to all problems handled by two-stage recovery algorithms but has a wider scope. In particular, for a p-chromatic visual system ($p = 3$ for a trichromatic visual system), there are a finite number of situations in which two-stage linear recovery algorithms may be applied, while there are an infinite number of situations in which general linear recovery may be applied. The report concludes with a discussion of how a visual system can use bilinear models to determine spectral descriptions of arbitrarily high dimension for lights and surfaces, when provided adequate information.