


5 Computer Vision (jgd1000)

- (a) Inferring a 3D object shape from shading variation across a surface depends on assumptions about how Lambertian or how specular each area is. For a surface reflectance map $\phi(i, e, g)$ having a mixed form,

$$\phi(i, e, g) = \frac{s(n + 1)(2 \cos(i) \cos(e) - \cos(g))^n}{2} + (1 - s) \cos(i)$$


give a range of values for s and n that would arise for: (i) a matte surface, and (ii) a glossy surface. What form of reflectance map $\phi(i, e, g)$ describes (iii) a mirror, and what form describes (iv) the lunar surface? (v) Why is $\phi(i, e, g)$ as specified above sometimes called the “Face Powder Equation”? (vi) How does the lunar form of $\phi(i, e, g)$ explain why the full moon looks like a flat 2D penny in the sky, rather than a 3D sphere like a ping-pong ball? [8 marks]

- (b) A breakthrough in face recognition accuracy arose when machine learning on big datasets minimised a loss function involving terms like $\| f(x_i^a) - f(x_i^p) \|^2$ and $\| f(x_i^a) - f(x_i^n) \|^2$ on triples of embeddings for $f(x_i^a)$ (anchor faces), $f(x_i^p)$ (positive examples: same face), and $f(x_i^n)$ (negative examples: different faces).



This approach treats false matches and failures-to-match as equally bad errors. But their costs are vastly different for a 1-to-1 face verification system (that just makes a ‘yes/no’ decision), versus a face identification system that may need to search a database the size of an entire nation, returning an actual identity. Propose a parameterised loss function for an algorithm that can be tuned for the different costs of the two error types, false matches and failures-to-match. Explain how its parameter(s) should reflect the numbers of potential false match collisions that must be avoided in a large-scale search. [6 marks]

- (c) A surprising aspect of human vision is the prevalence of quite striking illusions, which cannot be defeated even by being aware of them. Are visual illusions “bugs”, or “features” that should be built into computer vision algorithms? Consider in your answer both the tiling illusion (in which all horizontal lines really are parallel), and the hollow mask illusion below (in which the face always appears convex even when the mask is concave in presentation).



[6 marks]