



Original Articles

Understanding face familiarity

Robin S.S. Kramer^{a,b}, Andrew W. Young^a, A. Mike Burton^{a,*}^a Department of Psychology, University of York, UK^b School of Psychology, University of Lincoln, UK

ARTICLE INFO

Keywords:

Face recognition

Familiarity

Face matching

Face learning

ABSTRACT

It has been known for many years that identifying familiar faces is much easier than identifying unfamiliar faces, and that this familiar face advantage persists across a range of tasks. However, attempts to understand face familiarity have mostly used a binary contrast between ‘familiar’ and ‘unfamiliar’ faces, with no attempt to incorporate the vast range of familiarity we all experience. From family members to casual acquaintances and from personal to media exposure, familiarity is a more complex categorisation than is usually acknowledged. Here we model levels of familiarity using a generic statistical analysis (PCA combined with LDA) computed over some four thousand naturally occurring images that include a large variation in the numbers of images for each known person. Using a strong test of performance with entirely novel, untrained everyday images, we show that such a model can simulate widely documented effects of familiarity in face recognition and face matching, and offers a natural account of the internal feature advantage for familiar faces. Furthermore, as with human viewers, the benefits of familiarity seem to accrue from being able to extract consistent information across different photos of the same face. We argue that face familiarity is best understood as reflecting increasingly robust statistical descriptions of idiosyncratic within-person variability. Understanding how faces become familiar appears to rely on *both* bottom-up statistical image descriptions (modelled here with PCA), and top-down processes that cohere superficially different images of the same person (modelled here with LDA).

1. Introduction

The concept of familiarity is central to our understanding of face recognition. It has been known for many years that perception of familiar and unfamiliar faces differs in a number of ways (for reviews see Johnston & Edmonds, 2009; Young & Burton, 2017), and this point is emphasised in theoretical models (Bruce & Young, 1986; Burton, Bruce, & Hancock, 1999). For example, in studies of recognition memory, familiar faces are recognised faster and more accurately than unfamiliar faces (Ellis, Shepherd, & Davies, 1979; Klatzky & Forrest, 1984; Yarmey, 1971). This difference is not in any straightforward sense purely a memory effect, because in more recent studies of perceptual face *matching*, participants are again more accurate with familiar (compared to unfamiliar) faces, when judging whether two images depict the same person (e.g. Bruce, Henderson, Newman, & Burton, 2001; Bruce et al., 1999; Burton, Wilson, Cowan, & Bruce, 1999; Megreya & Burton, 2006; Megreya & Burton, 2008).

Despite these differences, our working definition of familiarity has been unsophisticated and our understanding of what happens when a face becomes increasingly familiar has been limited at best. Almost all studies compare unfamiliar, never previously seen, faces to highly

familiar people, often well-known celebrities. However, our daily experience tells us that familiarity is not simply a dichotomy. We all know many people with varying levels of familiarity, from members of our family encountered every day over long periods, to casual acquaintances perhaps seen occasionally on our route to work, or serving us in an infrequently-visited café. In this paper, we aim to capture familiarity in all its diversity. We present a model of face recognition which incorporates a large range of familiarity, and explore the consequences of increasing familiarity.

One key effect of familiarity is that it leads to generalisable representations for recognition. Early memory studies consistently showed that superficial image changes in pose, expression or lighting were detrimental to memory for unfamiliar faces, but had very little effect on familiar face memory (e.g. Bruce, 1982; Hill & Bruce, 1996; O’Toole, Edelman, & Bülthoff, 1998; Patterson & Baddeley, 1977). This has led to the idea that unfamiliar face processing is highly image-bound (Hancock, Bruce, & Burton, 2000; Megreya & Burton, 2006). In consequence, recognition declines as a function of differences between study and test photos (Beveridge et al., 2011; Estudillo & Bindemann, 2014), since representations of unfamiliar faces are tied to the specific images that were encountered. This image-dependence for unfamiliar

* Corresponding author at: Department of Psychology, University of York, York YO10 5DD, United Kingdom.
E-mail address: mike.burton@york.ac.uk (A.M. Burton).



Fig. 1. Unconstrained ambient images of the same person. Depicted variation is due to changes in pose, lighting, expression, age, camera settings, and so on. Image attributions from left to right: Eva Rinaldi (Own work) [CC BY-SA 2.0], Grant Brummett (Own work) [CC BY-SA 3.0], Gage Skidmore (Own work) [CC BY-SA 3.0], Eva Rinaldi (Own work) [CC BY-SA 2.0], Eva Rinaldi (Own work) [CC BY-SA 2.0].

faces seems to hold even after extensive training involving repeated exposure to a small number of different views of the same face (Liu, Bhuiyan, Ward, & Sui, 2009; Longmore, Liu, & Young, 2008). In such circumstances, particular training examples themselves become well-recognised, but show little generalisation to novel examples of the learned faces.

In marked contrast to unfamiliar face recognition, recognition of highly familiar faces is very robust. We can tolerate severe image degradation (Bruce et al., 2001; Burton et al., 1999) and considerable image distortion (Hole, George, Eaves, & Rasek, 2002) with very little effect on our ability to recognise the people we know. Why might this be? One proposal that lies at the heart of the approach we develop here is that our exposure to familiar faces has itself been highly diverse, including the very wide variability in the appearance of any particular individual that arises under everyday conditions (Burton, 2013; Jenkins & Burton, 2011; Jenkins, White, Van Montfort, & Burton, 2011). To illustrate this point, consider Fig. 1, comprising five photos of the actor Hugh Jackman. These pictures vary due to characteristics of the person (e.g. age, hairstyle, weight), the pose and facial expression, the image capture conditions (e.g. lighting, viewpoint) and the capture device (e.g. perspective settings, exposure levels). The images are therefore superficially very different in a way that is typical of everyday, ambient images (Burton, Jenkins, & Schweinberger, 2011). However, despite this diversity, a viewer familiar with the actor can recognise Hugh Jackman easily in all the photos. Our proposal in earlier work has been that this is because we have already encountered his face in a wide range of conditions, allowing us to have built up a representation of him which *includes* information about the ways in which his face can vary.

The nature of face representations has, of course, been a long-standing concern. In particular, many researchers have asked how it might be possible to build a representation that can be accessed when presented with *any* recognisable instance of a particular face (Bruce & Young, 1986; Eger, Schweinberger, Dolan, & Henson, 2005). Most conceptions, until recently, have emphasised what might potentially be common to all images of a person. For example, the most widely used idea involves the second-order configuration of distances between facial features (Carey & Diamond, 1977), though this is now known to run into both empirical and conceptual difficulties (Burton, Schweinberger, Jenkins, & Kaufmann, 2015; Maurer, Le Grand, & Mondloch, 2002). Alternatively, it has been pointed out that there might be common texture patterns across the face that can be captured through image averaging (Burton, Jenkins, Hancock, & White, 2005). Such approaches imply, at least implicitly, that familiarisation results in higher fidelity representations which can become sufficiently refined to be recruited when recognising a novel image of a known person. By focusing on what might be common to all views of the same face, research in this tradition thus often treats within-person variability – the extent to which the same face can look different – as noise. Typical experimental approaches in consequence tend to use highly controlled stimuli in

which images of different people are taken under very similar conditions (lighting, pose, expression, camera).

The approach used here represents a break from this tradition. We have recently followed an important insight of Bruce (1994) and suggested that, rather than being irrelevant noise, within-person variability can actually *assist* in finding information that is diagnostic of individual identity (Burton, Kramer, Ritchie, & Jenkins, 2016). This is because statistical analysis of multiple images of the same person shows that within-person variability is, to some extent, idiosyncratic. So, the ways in which one face varies are different from the ways in which another varies. Under this proposal, it is important to sample widely over different, naturally occurring images of someone in order to become familiar with that person – because part of familiarisation is learning that person's unique variability.

This proposal that variability is central to creating effective representations of face identities is gaining experimental support. For example, participants learn a face more effectively when exposed to greater variation in the images they see (Menon, White, & Kemp, 2015a; Murphy, Ipser, Gaigg, & Cook, 2015; Ritchie & Burton, 2017). So, while traditional approaches to face learning emphasise image-independent factors such as *duration* of exposure (Read, Vokey, & Hammersley, 1990; Reynolds & Pezdek, 1992), this may not be so critical as the image-dependent *type* of exposure, and especially the *range* of exposure. Likewise, if people have idiosyncratic facial variability, then we would expect any training on a particular face to have rather limited generalisability to other faces. Once again, this is borne out by experiments studying training in face recognition. Facial learning can be enhanced by various training regimes, but the benefits accrue only to those faces encountered, and do not generalise to others (Dowsett, Sandford, & Burton, 2016; Hussain, Sekuler, & Bennett, 2009).

Renewed interest in face learning, as described above, highlights the fact that we need a better understanding of familiarity. Studies manipulating levels of familiarity do so, almost exclusively, through a binary categorisation of faces as ‘familiar’ or ‘unfamiliar’, and tests of learning tend to dichotomise responses as ‘seen’ or ‘unseen’. An exception is a series of experiments by Clutterbuck and Johnston (2002, 2004, 2005) who show that pairwise matching – i.e. the ability to match two different images of a face – varies relatively smoothly with levels of familiarity. Nevertheless, for the most part, familiarity is treated in the research literature as a discontinuous variable with only two states.

In this paper, we take the important step of examining familiarity as a multi-valued function. We present a development of a previously implemented computational model (Kramer, Young, Day, & Burton, 2017a) using minimal assumptions and standard image analysis techniques involving a combination of Principal Components Analysis (PCA) and Linear Discriminant Analysis (LDA). While this approach has already been shown to simulate the specific property of image invariant familiar face recognition (Kramer, Young, Day & Burton, 2017), these

Download English Version:

<https://daneshyari.com/en/article/7285526>

Download Persian Version:

<https://daneshyari.com/article/7285526>

[Daneshyari.com](https://daneshyari.com)