INSTITUTE OF PHYSICS PUBLISHING NETWORK: COMPUTATION IN NEURAL SYSTEMS

Network: Comput. Neural Syst. **14** (2003) 321–333 PII: S0954-898X(03)60967-X

Modular organization of adaptive colouration in flounder and cuttlefish revealed by independent component analysis

J C Anderson[1], R J Baddeley[1], D Osorio[1], N Shashar[2], C W Tyler[3], V S Ramachandran[4], A C Crook[5] and R T Hanlon[6]

[1] School of Biological Sciences, Sussex University, Brighton, UK
[2] The Inter University Institute of Eilat, Eilat, Israel
[3] The Smith-Kettlewell Eye Research Institute, San Francisco, CA, USA
[4] Brain and Perception Laboratory, University of California, San Diego, CA, USA
[5] Department of Zoology and Animal Ecology, University College, Cork, Ireland
[6] Marine Resources Center, Marine Biological Laboratory, Woods Hole, MA, USA

Received 8 May 2002, in final form 31 October 2002
Published 11 April 2003
Online at stacks.iop.org/Network/14/321

Abstract

Flounders and cuttlefish have an impressive ability to change colouration, for camouflage and, in the case of cuttlefish, for communication. We pursue the hypothesis that these diverse patterns are created by combining a small number of distinct pattern modules. Independent component analysis (ICA) is a powerful tool for identifying independent sources of variation in linear mixtures of signals. Two versions of ICA are used, one assuming that sources have independence over time, and the other over space. These reveal the modularity of the skin colouration system, and suggest how the pattern modules are combined in specific behavioural contexts. ICA may therefore be a useful tool for studying animal camouflage and communication.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Animals and humans vary their appearance for a variety of purposes. Facial expressions, body postures and skin colouration patterns reflect basic emotional states, including fear, aggression and sexual arousal (Hinde 1970), while posture and pattern are also used for camouflage (Cott 1940). Some animals actively regulate their colouration, and of these, bottom-living marine animals, especially flounders (*Pleuronectidae*) and cuttlefish (*Sepiidae*), are masters (Hanlon and Messenger 1988, Ramachandran *et al* 1996). These animals change their appearance using chromatophores, small sacs of pigment in which the area of visible pigment can be modified. Flounders probably modify colouration mainly for camouflage (figures 1(a)–(d)),

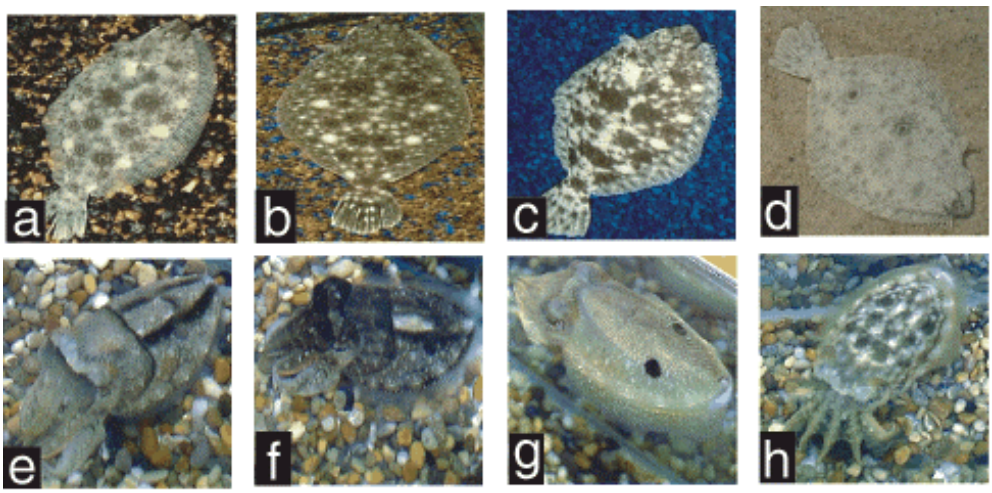

Figure 1. A selection of images (512 × 512 RGB; lossless TIFF) of a single summer flounder, *Paralichthys dentatus* ($n = 64$ images, age ∼8 months, body length ∼13 cm), and a single cuttlefish, *Sepia officinalis* ($n = 30$ images, age ∼6 months, mantle length ∼18 cm), were grabbed (MGI photosuite) from digital video (Sony VX-1000) of animals. (a)–(d) Flounder images were obtained after allowing the fish time to camouflage itself on a range of patterns created by mixing coloured gravels. (e)–(h) Images of cuttlefish; in (g) the animal gives an eye-spot signal in response to the threat of an approaching stick and (h) is the cuttlefish eating a crab.

while cuttlefish, in addition to camouflage, use visual patterns as a rich mode of expression for visual communication (figures 1(e)–(h); Hanlon and Messenger 1988).

Adaptive colouration has been studied by observation of skin surface markings in cuttlefish (Hanlon and Messenger 1988) and, for flounder, by principal component analysis (PCA) of body images (Ramachandran *et al* 1996). This previous work indicates that patterns are generated by a relatively small number of pattern modules (figure 2), but both these methods have weaknesses. Visual inspection is subject to the problems associated with subjective, non-quantitative analysis, including the possibility of bias from theoretical preconceptions. Factor analysis methods (and related methods such as PCA) potentially provide a quantitative solution to structure identification. Unfortunately, whilst factor analysis can find low dimensional descriptions of high dimensional data, there is no reason to believe that the dimensions identified have any biological relevance (for most factor analysis methods, any rotation of the factors identified has equal validity).

A recent development in automatic structure identification, independent component analysis (ICA), potentially provides a more powerful quantitative method to identify the biologically relevant dimensions in signalling systems. This paper describes the application of ICA to colouration patterns, and identifies the virtues and problems of this potentially very powerful method for studying communication in general and adaptive colouration of flounder and cuttlefish in particular.

ICA starts with the assumption that the observed data (in this case body patterns) are created from a (linear) mixture of different 'sources'. If the sources themselves are non-Gaussian distributed, ICA can estimate both the individual sources, and the transformation from the observed data to the underlying sources. The method works by using an implication of the central limit theorem: namely that a combination of any two non-Gaussian distributed variables tends to be more Gaussian than each individually. Therefore, the set of projections

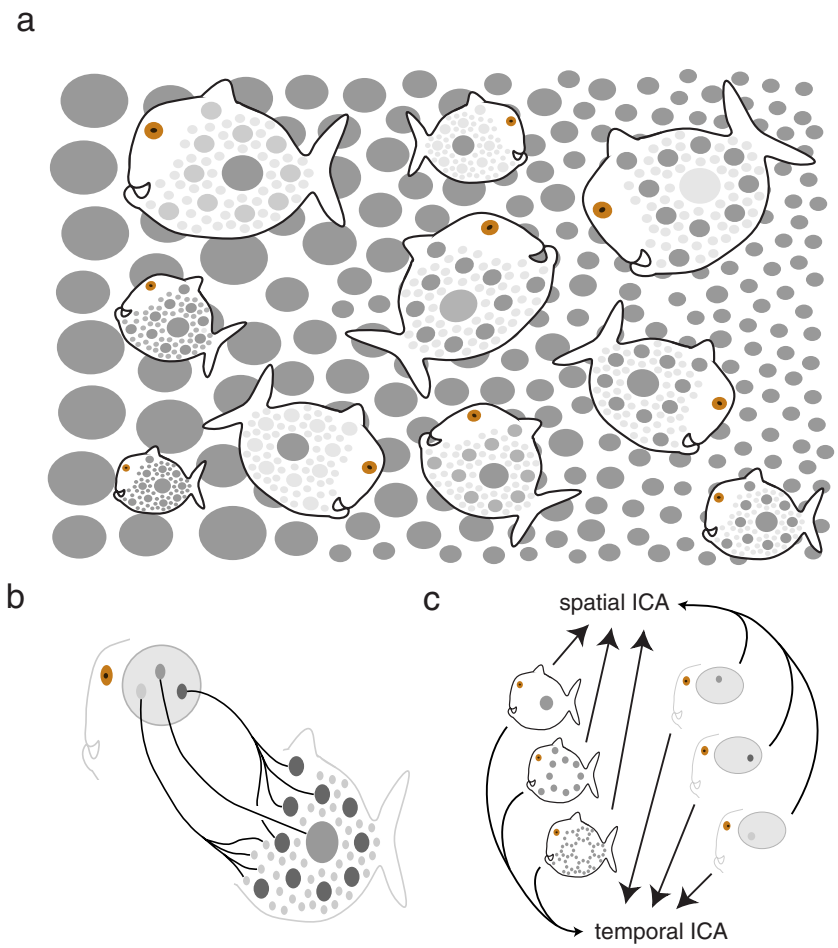

Figure 2. Diagrammatic illustration of how ICA can be used to analyse colouration patterns. (a) A school of fish camouflage themselves against a variable background by (b) independently adjusting the level of expression of three distinct groups of chromatophores (pattern modules), each via a separate neural pathway. Given a large collection of images, ICA provides identifies independent sources of variation, by finding linear projections of a set of data (images) that are maximally non-Gaussian. (c) ICA can be applied in two distinct ways, depending on what the underlying sources are assumed to represent. In spatial ICA (ICAs) sources are local features on the body surface. Alternatively, when using temporal ICA (ICAt), the sources are the expression strengths of the corresponding pattern modules.

of the data that are maximally non-Gaussian (by some appropriate measure) identifies the underlying sources. The procedure then is simple: start with a proposed set of 'features' within the data set; measure the non-Gaussianity of the projection of these features on to the data and then use some optimization algorithm to adaptively change the features to maximize this measure. If the assumption that the observed patterns consist of linearly combined, non-Gaussian sources is satisfied, this approach gives impressive results. For example, ICA has been successfully used to separate individual speakers from recorded mixtures (the cocktail party problem; Bell and Sejnowski 1995) and to interpret fMRI (McKeown *et al* 1998) and EEG data (Makeig *et al* 1996).

To apply ICA to a particular domain one has to decide what the underlying sources represent and how they are mixed (figure 2). There are two possible ways of understanding how patterns are generated. The first proposes that each image comprises the expression pattern modules. If each pattern module comprises a set of local features (splotches) on a uniform background, the distribution of intensities in these modules will be less Gaussian than any mixture of them.

Hanlon and Messenger (1988) proposed a set of local features (which they confusingly called chromatic 'components') on the body surface of the cuttlefish *Sepia officinalis*. We confirmed that if cuttlefish colouration patterns are created by linear combinations of such pattern modules, ICA could identify these modules. This was done by creating synthetic cuttlefish images as linear mixtures of pattern modules, consisting of blobs, lines and squares. Given that the distribution of intensities in each pattern module (the blobs, lines and squares) was non-Gaussian, ICA robustly extracted these individual modules from the synthetic images. This method, which assumes that each pattern module has a non-Gaussian distribution of intensities, is called spatial ICA (ICAs; Stone *et al* 1999).

The second regularity that can be exploited by ICA is the way in which the expression of the pattern modules varies over time. For instance a particular module may normally be absent, but occasionally expressed strongly. This will result in a non-Gaussian distribution of expression strengths, and in this case, the independence of the pattern module *expression* serves to identify the spatial structure of the pattern modules.

ICA applied in this context is called temporal ICA (ICAt; Stone *et al* 1999). ICAt works because any proposed pattern module that is a mixture of two genuine pattern modules will have a more Gaussian distribution of *expression* than each of the genuine pattern modules separately. The relationship between ICAt and ICAs, and its practical implications in the context of fMRI data, is discussed in more detail in Stone *et al* (1999).

These two regularities of non-Gaussian pattern structure, and non-Gaussian expression structure, provide two methods to find components of complex colouration patterns (ICAs and ICAt respectively). We find that both methods of ICA are able to reveal the basic modularity of adaptive camouflage in cuttlefish and flounder.

2. Methods

2.1. Warping

The analysis is based on comparing digital photographs of animals displaying a range of colouration patterns, and requires images to be compared point by point. As the study concerns spatial patterns, not colour, only the green channel of the RGB images was used. Given that pictures are taken from varying angles, and animals can distort their bodies, matching points on the body surfaces between images is a major technical problem. The solution is to warp images to a standard view by identifying a set of reference points on the surfaces of the body (figure 3). Meshes of 31 points were used on the flounder while 13 points could be reliably identified on the cuttlefish. The points form the apices of triangles, and bilinear interpolation transforms the pattern within each triangle to the appropriate reference triangle. As points lying outside the mesh grid could not be reliably identified, they were discarded (figures 3(b), (c)).

The quality of the results depends crucially on accurate warping, as fine details are lost through registration errors. There are two sources of registration errors. Firstly, human error derives from inconsistently identifying similar locations on the body surface across the image set. Stable anatomical features provide the most reliable markers. However, these are often too few to give a sufficiently detailed mesh. Identifying more subtle markings requires practice

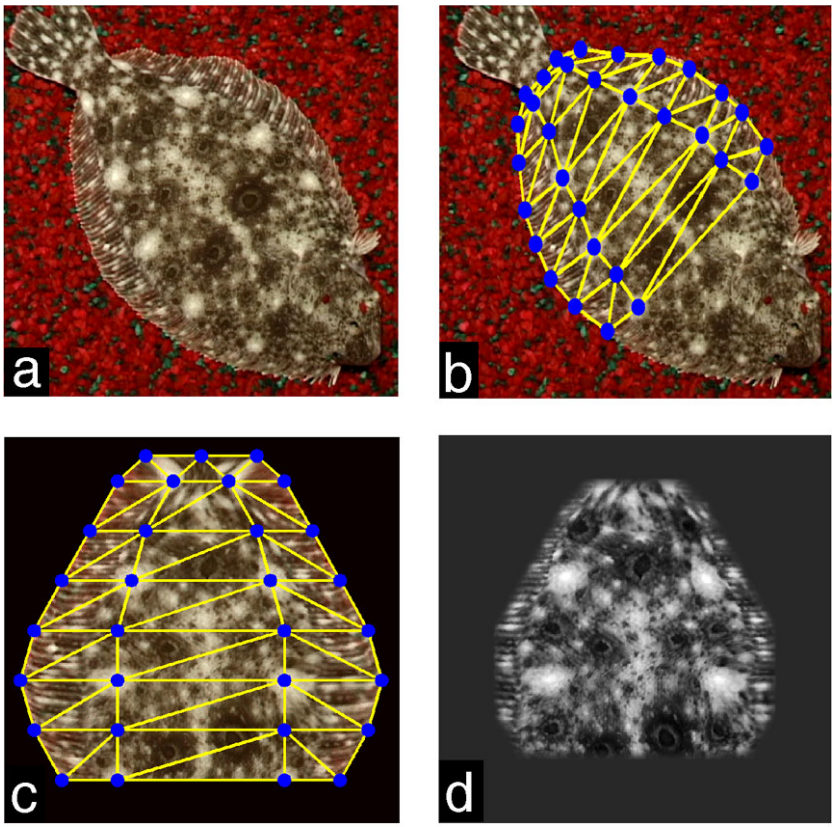

Figure 3. Image pre-processing, shown for the flounder. (a) As animals moved during filming, the position and the angle from which they were viewed changed. Views were transformed to a standard reference frame by image warping based on bilinear interpolation. (b), (c) A series of local affine transformations, defined by a triangular mesh grid consisting of 31 points on the skin of the flounder and 13 points on the cuttlefish, which could be relatively easily identified. (d) As points lying outside the mesh grid could not be reliably identified, each image was windowed using a mask defined by the mesh perimeter. Each image was normalized by subtracting the mean and dividing by the standard deviation in intensity. These values were retained for reconstructing images from the independent components.

but does improve with experience. We are currently exploring the use of tattooing (adding skin markings by dye injection) to provide more permanent mesh points.

The second source of registration error derives from the inability of the warping algorithm to correct the observed distortion. By using warping techniques borrowed from face recognition and fMRI we can be confident that systematic errors of this kind are minimized, given accurately identified registration points. The structural deformation we witnessed in the animals are certainly much less elaborate than those observed for the human face, and given accurately identified control points visual inspection suggests the results obtained were very good.

It is also worth noting that while accurate warping improves the quality of the results, we found that the artifacts caused by poor registration were quite easy to spot (see figure 6). Thus, the general findings of a particular study may not be severely compromised even if the investigator is unable to accurately align the image set.

2.2. Identifying pattern modules

After warping (figure 3), each square image matrix of intensity levels ($N \times N$ pixels) was converted to an image vector ($1 \times N^2$), and all these vectors were combined into a single compound image matrix ($n \times N^2$, where $n =$ total number of images). For spatial ICA (ICAs) each row corresponded to a single image, and each column to a fixed pixel location across all the images. For temporal ICA (ICAt) the compound image matrix was transposed. This means that fixed pixel locations across all the images now represent the mixtures. ICA was performed on this matrix using Hyvärinen and Oja's (1997) fixed-point algorithm (deflationary method using cubic and tanh non-linearity) implemented by Hugo Gavert ('fastICA', running under Matlab[1] v5.3;).

The application of fastICA to a data set yields two separate matrices: a mixing matrix (a matrix transforming the unobserved sources to the observed data) and the loading matrix, which contains the projection of the original data on to the mixing matrix. For the case of ICAt, the mixing matrix carries the spatial details of the pattern modules we are interested in, while the loading matrix corresponds to the levels of expression of these modules within the image dataset. In ICAs this situation is reversed with the loading matrix conveying the spatial information whilst the mixing matrix describes the expression of the pattern modules in the original images. In this paper we are interested primarily in the spatial structure of the pattern modules. To minimize any confusion between the mixing matrix and the loading matrix we simply refer to the loading matrix revealed by ICAs as 'spatial components' and the mixing matrix revealed by ICAt as 'temporal components' as these are the respective matrices that contain the spatial structure of the pattern modules.

Prior to estimation of the independent components, the compound image matrix was projected on to its principal components and normalized (Bishop 1995) so the transformed mixtures were decorrelated and of unit variance. For large datasets, automatic methods for estimating the dimensionality of this 'whitened' subspace and the number of independent components to extract have been developed (Attias 1999). However, we have comparatively small datasets, and are faced 'with the curse of dimensionality' severely limiting the number of components that we can hope to find. An informal estimate of the dimensionality attainable can be generated by viewing ICA as a method of (projection pursuit) density estimation. The minimum number of data points required to estimate a density of dimension n, where to identify the deviation form normality for each dimension requires D data points, is D^n. We have 64 images for the flounder and 30 for the cuttlefish. Even the crudest measure of deviation from normality requires four data points, so in the case of ICAt we cannot realistically hope to recover more than three components. This constraint does not apply to ICAs where we are measuring deviations from normality in terms of intensity variations across the image, which are based upon about 512^2 data points. The number of components which can be reasonably estimated in this case is defined by the number of mixtures (images): 64 in the case of the flounder and 30 for the cuttlefish. Thus, attempts were made to estimate a larger number (16) of spatial components using ICAs (figure 6). However, for the purpose of comparing the two methods we used three dimensions for both principal component and independent component stages. At least for cuttlefish, it is certain that more exist (Hanlon and Messenger 1988), and should be discovered by applying the method to larger datasets.

Potentially, the pattern modules identified could be affected by the exact definition of deviations from normality. In practice, ICAs gave indistinguishable results when running the analysis with tanh or cubic non-linearity (except that ICAt sometimes failed to converge when using the cubic non-linearity).

[1] http://www.cis.hut.fi/projects/ica/fastICA/

3. Results

Figure 1 shows samples from 64 images of a single flounder, *Paralichthys dentatus*, on a range of artificial floors created by mixing different coloured gravel, and samples from 30 images of an immature cuttlefish, *Sepia officinalis*. The cuttlefish was also startled and observed hunting and feeding, generally encouraging it to display a range of patterns. Each set of images was warped (figure 3) to align the body plans of the animals, ensuring that we compared identical locations on the body surface. ICA was then performed on these two sets of images (figures 4, 5 and 6). For ICA to work the sources must be less Gaussian than the mixtures. This was indeed the case (table 1). ICAt reliably extracted sources with a higher kurtosis than the original mixtures. ICAs also gave similar and consistent results, except the kurtosis value for one of the extracted sources was lower than that of the original data.

The three spatial components revealed by applying ICAs to the set of flounder images account for roughly 90% of the variance in intensity across the image set (figure 4—second row). Temporal components revealed by ICAt (fourth row) account for slightly less, around 78% of the variance. By taking a linear weighted sum of the spatial or temporal components we can approximately reconstruct the original images. Reconstructions of three fish images, chosen because they express predominantly one of the components, are shown with the three original images placed centrally for comparison. Reconstructions using spatial components match the original images more closely than when using temporal components.

A similar pattern of results was found when applying ICA to cuttlefish (figure 5). Here, three spatial components account for 86% of the variance in intensity across the image set compared to 80% for the temporal components. A feature of the cuttlefish results is that while the first and second spatial components are asymmetric, only the first temporal component is markedly asymmetric. This asymmetry reflects the animal's capacity to give the illusion of spatial depth through shadowing by half shading the central square on its back (figure 1(f)). ICAs reveals two distinct motifs that may be combined to produce the effect of shading; ICAt however suggests the two halves of the square maybe under joint neural control, and are not modified independently. However, a larger set of images is needed to confirm this tentative conclusion.

The lower panels of figure 5 illustrate a cuttlefish displaying two different patterns. Those on the right show the animal startled by an approaching hand, where it predominantly expresses the third spatial component, which contributes relatively little to the patterns displayed in the left-hand panels.

While the number of components that can be realistically estimated using ICAt is severely limited by the number of images, in the case of ICAs it is reasonable to estimate as many components as there are images. Figure 6 displays the first 16 spatial components. Also shown are the variations in the level of expression of each component (left $>$ right, top $>$ bottom) together with their respective kurtosis values. It is immediately obvious that apart from two components in the flounder, and perhaps five or six in the cuttlefish, most do not contain localized features that relate to identifiable regions on the body surface. Thus, while in theory ICAs can glean more information than ICAt from small image datasets, if there are only a small number of independent spatial sources underlying the observed mixtures then much of this additional information represents little more than noise. This finding, particularly in the case of the flounder, illustrates one of the major difficulties associated with applying blind source separation techniques when the number of underlying sources is unknown and emphasizes the need for sensible *a priori* criteria for interpreting the validity of any results generated.

In this case we can use a fairly rudimentary definition to highlight components of potential interest. Those with highest and lowest kurtosis values are most likely to represent unique

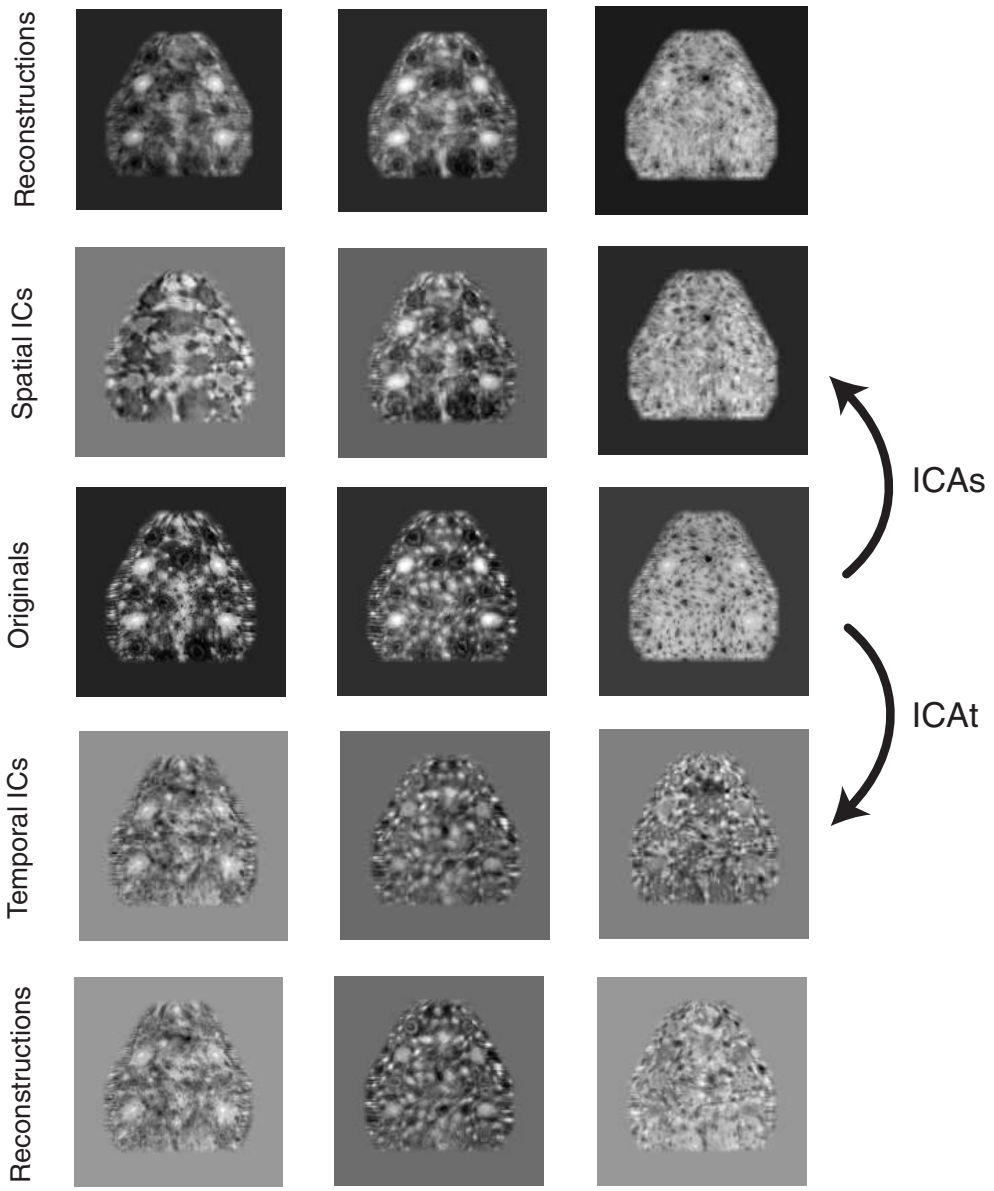

Figure 4. ICA of a flounder. *Second row*: three spatial components calculated by applying spatial ICA (ICAs) to 64 flounder images. Each component indicates regions on the flounder's skin, which undergo simultaneous changes in brightness independent of those indicated on the other components. Features on the spatial components correlate strongly with identifiable features on the original images (e.g. *third row*). *Fourth row*: temporal components calculated by applying temporal ICA (ICAt) to the same set of images. Features on these temporal components show some correlation with the fish's skin markings (e.g. centre panels) although clearly less than spatial components. *First row:* reconstructions of original images (centre panels) made by linearly mixing the three spatial components (second row). Reconstructions match the original images. *Fifth row:* images reconstructed by linearly mixing the three temporal components extracted using ICAt. These temporal components are worse at reconstructing the original images, but it is important to note that reconstruction is not the goal of ICA, which is to identify underlying pattern modules.

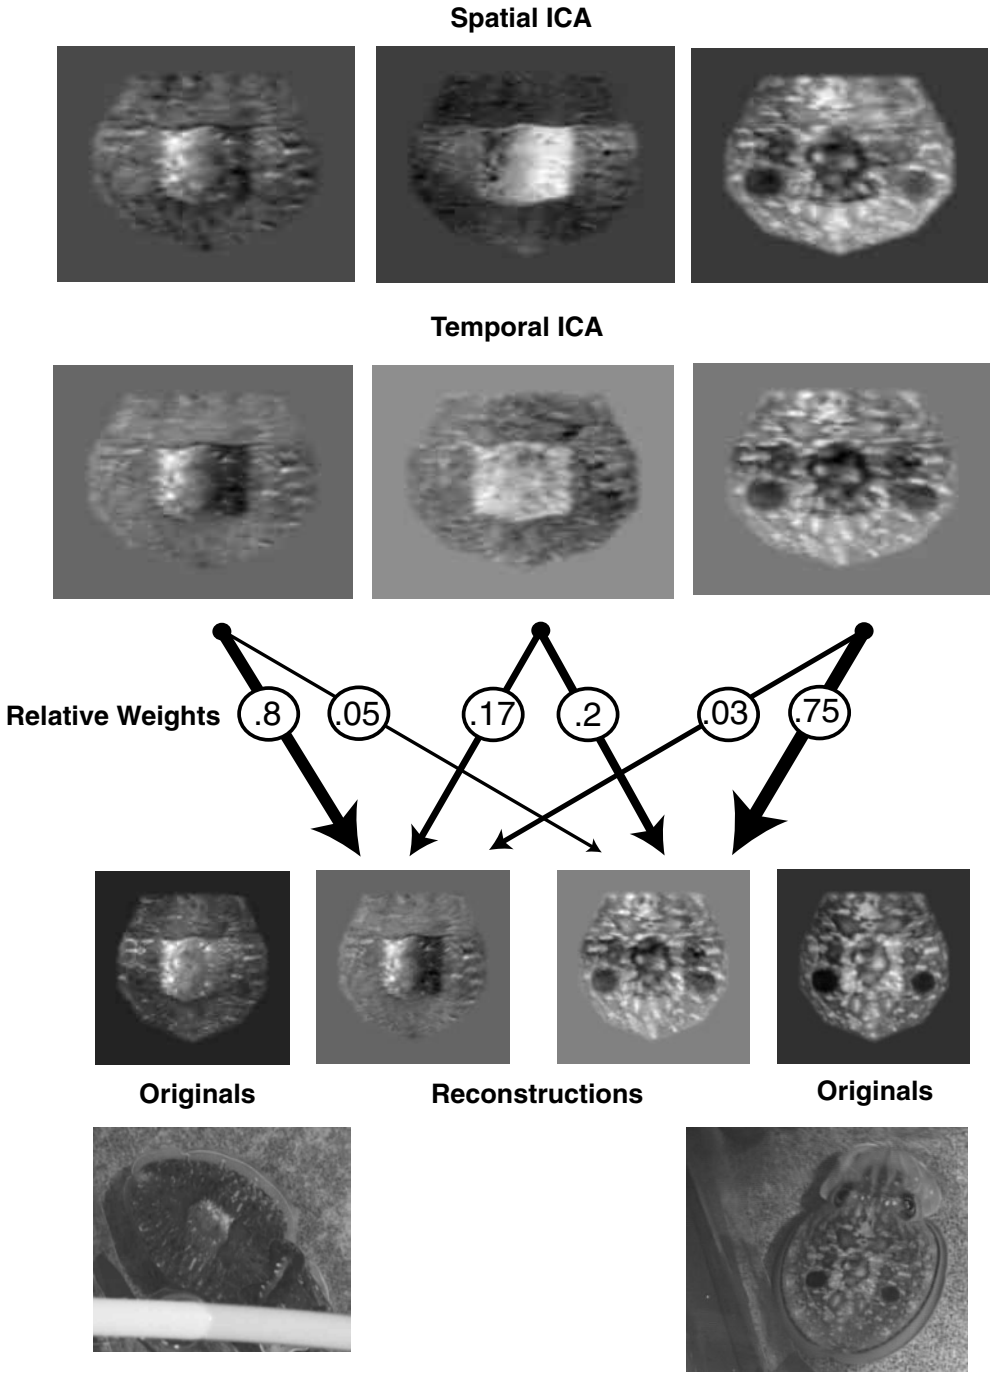

Figure 5. ICA of cuttlefish patterns. *Top row:* three spatial components calculated by applying ICAs to a set of 30 images of cuttlefish. *Second row:* temporal components extracted by applying ICAt to the same set of images. Arrows from these components point to reconstructions (*third row*) and indicate the relative weights of the three components while the cuttlefish is undergoing two distinct displays. Lower panels show original images.

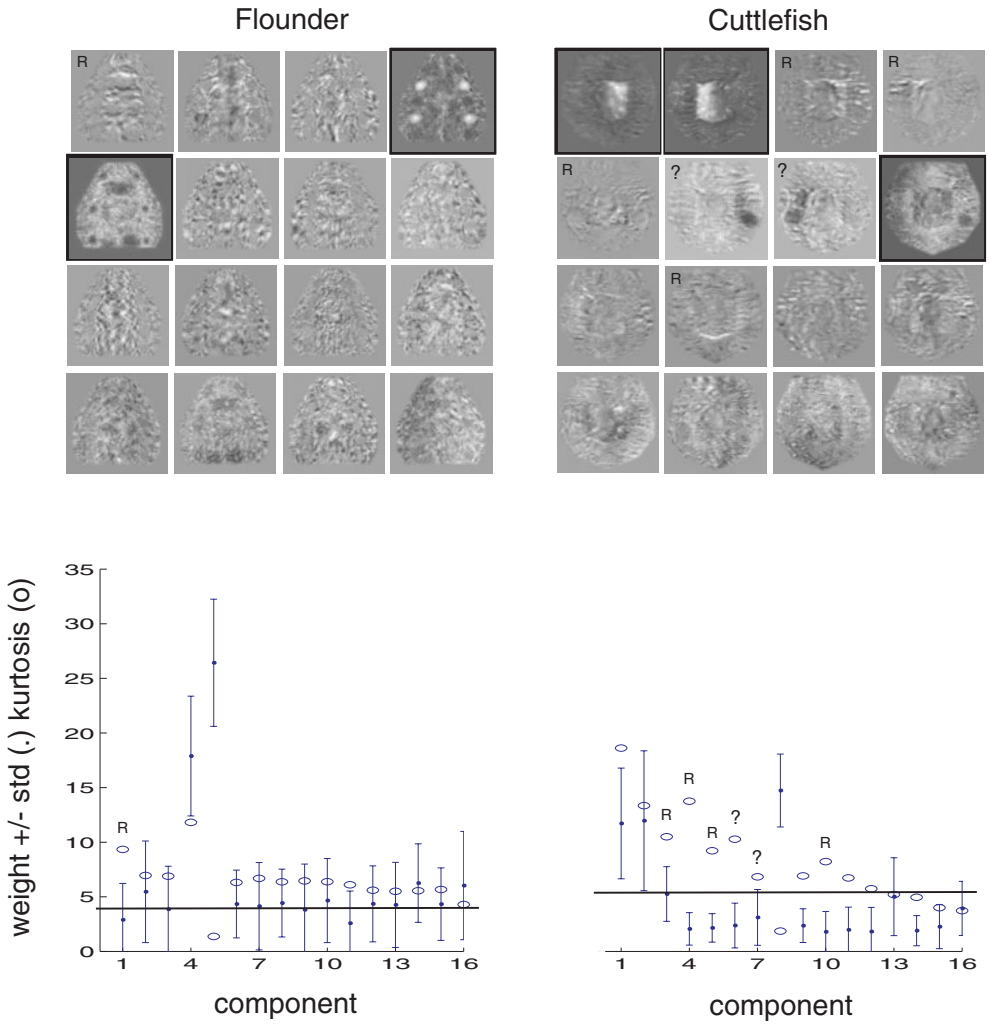

Figure 6. Spatial ICA of flounder and cuttlefish images. *Image panels*: 16 spatial components were estimated. *Graphs* indicate the mean and standard deviation in absolute level of expression of each of the components (number left to right, top to bottom) together with their kurtosis value. The bold line on each graph indicates the mean kurtosis value for the original images. Borders indicate those components which have particularly high or low kurtosis values and whose expression varies substantially over the image set. The components marked 'R' on the graphs most likely result from registration errors. The presence of independently driven eyespots on the cuttlefish (components marked '?') may be a result of trying to estimate too many sources (see the text for further discussion).

spatial sources as mixtures of these sources will, by the central limit theorem, have kurtosis values closer to the original mixtures (mean kurtosis values for the images (mixtures) are indicated by bold lines on the graphs in figure 6). In addition, those that are strongly expressed, or vary markedly in their level of expression over the image dataset, are likely to be of more perceptual relevance because they will generate bold visual stimuli. Significantly, the components which obviously satisfy both these criteria are broadly similar to those derived when constraining the projection to three dimensions. However, there are some exceptions. The third flounder component appears to be missing. This result suggests that our initial

Table 1. Kurtosis values for independent components estimated using ICAs and ICAt applied to images of flounder and cuttlefish.

	ICAs		ICAt	
Kurtosis (mean ± std)Gaussian = 3	Images (mixtures)	ICs (sources)	Images (mixtures)	ICs (sources)
Flounder	$3.7 \pm 1.3(64)$	10.7, 6.7, 1.8 (tanh/pow3)	2.14 ± 0.6 ($n \sim 512^2$)	3.6, 9.4, 7.5 (tanh)
Cuttlefish	$5.1 \pm 0.7(n = 30)$	15.3, 11.8, 2.7 (tanh/pow3)	$2.3 \pm 0.3(n \sim 512^2)$	3.3, 3.5, 3.9 (tanh)

search for three components may have been over-ambitious. In the case of the cuttlefish, the reverse appears to be true, with in excess of three components satisfying our selection criteria. However, two of these components contain localized features that do not obviously correlate with established body surface markings (Hanlon and Messenger 1988). These almost certainly derive from registration errors incurred during warping. Moreover, while these components have high kurtosis values they may still have been reasonably rejected on the basis of their relatively low level of expression across the image dataset. The capacity of the cuttlefish when compared to the flounder to dramatically distort its body together with the general difficulty in identifying consistent skin markings may have increased the frequency of these artifacts.

Another interesting feature of the less constrained spatial ICA is that cuttlefish components six and seven suggest that the eye-spots maybe independently modulated. Indeed, a small fraction of the original images that contained unilateral eye spot displays implied the same conclusion. This illustrates another potential pitfall when applying ICA. Underestimating the correct number of underlying generators may result in unique sources being extracted as linear combinations (i.e. the two eyespots combine to a single spatial feature when the dimensionality is constrained to three). As long as the combined sources have sufficiently high kurtosis values then ICA will group them as an apparently unique source. An analogous situation occurs during the blind source separation of sound signals using fewer microphones (mixtures) than sources. Here, for example, a combination of two instruments may be extracted as a single source. While results like these illustrate several potential weaknesses of ICA, such problems are inherent in the majority of interesting applications involving natural, rather than artificial signals.

4. Discussion

This paper introduces ICA as a quantitative method for describing animal colouration patterns that is potentially much more powerful than standard methods of factor analysis. ICA should be appropriate for behavioural work because, if the assumptions the algorithm makes hold for a given data set, it can uniquely identify the underlying 'generators' of the observed images. These may be related to neural control centres, motivational states, or perceptual mechanisms.

Although ICAs may better summarize intensity variations across the set of images, the main reason for using ICA is to identify independent sources of variation underlying the patterns. ICAt identifies such sources if the observed signals (i.e. colouration patterns) are constructed from linear combinations of pattern modules (figure 2), and the variation in expression of these modules across the set of images is non-Gaussian. The fact that here both methods give similar results (figure 4) is encouraging evidence that the components found are genuine pattern modules.

Identifying the specific body patterns that an animal displays in a given situation may help reveal the basic units employed for communication. Linear interplay of a small number of signal elements has been shown to form the basis for a range of communicative body patterns in several animals (e.g. Leyhausen 1956, Tembrock 1968). In cuttlefish three temporal independent components (figure 5) and at least three spatial components were identified (figures 4 and 6). Tracking the expression of these components over a wide range of clearly identified communication related behavioural situations could be useful in understanding the specific meaning they carry. Unfortunately, our limited database was inadequate for such a context related analysis.

ICA contrasts with the alternative method of PCA, which simply captures as much variance as possible, regardless of how sources are combined. In particular, temporal independent components may correspond to discrete behavioural units, represented by modular arrays of chromatophores under joint neural control (figure 2). The partial correlation between local features contained within the estimated pattern modules and colour patterns on the original fish images supports this conclusion. Ramachandran *et al* (1996) suggested that the tropical flounder *Bothus ocellatus* combines three chromatophore channels each tuned to a specific range of spatial frequencies, which together match the broadband spectra of seafloor textures. In *Paralichthys* we found no simple frequency dependent segregation. These early results suggest we can explore the detailed spatial structure of the chromatophore channels.

What general lessons can we draw from our results? The first is the importance of accurate registration. Considerable effort was invested into making registration as accurate as possible, but we still found evidence in the results for registration artifacts. If analysing vocal communication it would be vital to normalize for equivalent variations such as vocal track length. The second lesson is the importance of using large data sets. The maximum possible number of components that could be identified with ICAt is severely constrained by the number of example images (the number of images required increases exponentially with the number of dimensions to be identified). This will also be true of other quantitative methods. We are therefore generating a much larger data set in the hope of finding more detailed structure.

Of equal importance is some kind of 'sanity' check for any method. A quantitative method makes (often strong) assumptions about the data. In the case of ICA, it is assumed that the data are a linear combination of non-Gaussian sources, and this method is only useful if the assumption is approximately true. Here we used three checks: we confirmed that at least some projections had highly non-kurtotic distributions as required by the method. Though obvious, this check (equivalent to checking for normality in factor analysis) is often overlooked. The second check was that both ICAs and ICAt gave similar results. If the images were not created by combining different modules, there is no reason why the two methods will give similar results. Though there were differences (the symmetrical/asymmetrical eye spots), the results were similar enough to be encouraging. Lastly we checked that the expression of the identified pattern modules actually varied between images. Registration errors can result in highly kurtotic distributions with low variance. Requiring any behaviourally relevant components to have high contribution to the overall variance alleviates this problem. If ICA identifies a representation that only describes a small percentage of the variation, then the results must again be treated with extreme caution. In our case this was not a problem with ICAs accounting for 90/86% of the variance, and ICAt accounting for 78/80% (in contrast to PCA, there is no *a priori* reason that the identified structure should account for a large amount of the variance).

Methods of quantifying observations underpin our understanding of animal behaviour. Complex behaviours might often be succinctly summarized in terms of mixtures of basic modules (Hinde 1970). Unfortunately, in most cases these components usually mix nonlinearly,

making it difficult to identify them, and hence to analyse the behaviour quantitatively. Flounder and cuttlefish colouration patterns are a form of behavioural expression in which the components are directly visible and potentially combine linearly. As ICA identifies underlying components in linear combinations of patterns, it is well suited to the problem, and an appropriate tool for analysing this type of visual display. However, as with many powerful linear techniques, ICA needs to be applied cautiously if it is to provide useful insight.

Acknowledgment

We thank Tom Cronin for financial support and encouragement.

References

Attias H 1999 Inferring parameters and structure of latent variable models by variational Bayes *Proc. 15th Conf. on Uncertainty in Artificial Intelligence* ed K B Laskey and H Prade (San Francisco, CA: Morgan Kaufmann)

Bell A J and Sejnowski T J 1995 An information-maximisation approach to blind separation and blind deconvolution *Neural Comput.* **7** 1129–59

Bishop C M 1995 *Neural Networks for Pattern Recognition* (Oxford: Clarendon)

Comon P 1994 Independent component analysis—a new concept? *Signal Process.* **36** 287–314

Cott H B 1940 *Adaptive Coloration in Animals* (London: Methuen)

Hanlon R T and Messenger J B 1988 Adaptive coloration in young cuttlefish (*Sepia officinalis* L.): the morphology and development of body patterns and their relation to behaviour *Phil. Trans. R. Soc.* B **320** 437–87

Hanlon R T and Messenger J B 1996 *Cephalopod Behavior* (Cambridge: Cambridge University Press)

Hinde R A 1970 *Animal Behaviour* 2nd edn (New York: McGraw-Hill) ch 16

Hyvärinen A and Oja E 1997 A fast fixed-point algorithm for independent component analysis *Neural Comput.* **9** 1483–92

Leyhausen P 1956 Verhaltensstudien bei Katzen *Z. Teirpsychol. Beiheft* **2** 1–100

Makeig S, Bell A J, Jung T-P and Sejnowski T J 1996 Independent component analysis of EEG data *Adv. Neural Inf. Process. Syst.* **8** 145–51

McKeown M J, Makeig S, Brown G G, Jung T-P, Kindermann S S, Bell A J and Sejnowski T J 1998 Analysis of fMRI data by decomposition into independent spatial components *Hum. Brain Mapping* **6** 1–31

Ramachandran V S, Tyler C W, Gregory R L, Rogers-Ramachandran D, Duensing S, Pillsbury C and Ramachandran C 1996 Rapid adaptive camouflage in tropical flounders *Nature* **379** 815–18

Stone J V, Porrill J, Buchel C and Friston K 1999 Spatial, temporal and spatiotemporal independent component analysis of fMRI data *Proc. 18th Leeds Statistical Research Workshop on Spatial–Temporal Modelling and its Applications*

Tembrock G 1968 Land mammals *Animal Communication* ed T A Sebeok (Bloomington, IN: Indiana University Press) ch 16 pp 338–404

van Hateren J H and van der Schaaf A 1998 Independent component filters of natural images compared with simple cells in primary visual cortex *Proc. R. Soc.* B **265** 359–66