Max-Planck-Institut für Mathematik in den Naturwissenschaften Leipzig

A Computational Model of Dysfunctional Facial Encoding in Congenital Prosopagnosia

by

Rainer Stollhoff, Ingo Kennerknecht, Tobias Elze, and Jürgen Jost

Preprint no.: 9 2011

A Computational Model of Dysfunctional Facial Encoding in Congenital Prosopagnosia

Rainer Stollhoff[1,*], Ingo Kennerknecht[3], Tobias Elze[1], Jürgen Jost[1,2],

1 Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

2 Santa Fe Institute for the Sciences of Complexity, Santa Fe, NM, USA

3 Institute of Human Genetics, Westfälische Wilhelms-Universität, Münster, Germany

∗ E-mail: rainer.stollhoff@mis.mpg.de

Keywords: Infomax; ICA; independent component analysis; face recognition; prosopagnosia; hereditary prosopagnosia; congenital prosopagnosia; faceblind

Abstract

Congenital prosopagnosia is a selective deficit in face identification that is present from birth. Previously, behavioral deficits in face recognition and differences in the neuroanatomical structure and functional activation of face processing areas have been documented mostly in separate studies. Here, we propose a neural network model of congenital prosopagnosia which relates behavioral and neuropsychological studies of prosopagnosia to theoretical models of information processing.

In this study we trained a neural network with two different algorithms to represent face images. First, we introduced a predisposition towards a decreased network connectivity implemented as a temporal independent component analysis (ICA). This predisposition induced a featural representation of faces in terms of isolated face parts. Second, we trained the network for optimal information encoding using spatial ICA, which led to holistic representations of faces. The network model was then tested empirically in an experiment with ten prosopagnosic and twenty age-matched controls. Participants had to discriminate between faces that were changed either according to the prosopagnosic model of featural representation or to the control model of holistic representation. Compared to controls prosopagnosic participants were impaired only in discriminating holistic changes of faces but showed no impairment in detecting featural changes.

In summary, the proposed model presents an empirically testable account of congenital prosopagnosia that links the critical features - a lack of holistic processing at the computational level and a sparse structural connectivity at the implementation level. More generally, our results point to structural differences in the network connectivity as the cause of the face processing deficit in congenital prosopagnosia.

1 Introduction

Faces are a special class of visual stimuli. They are rapidly detected in images and provide a multitude of different information important for social communication such as gaze direction, facial expressions, age, gender, and identity. Under normal conditions of cortical maturation and development, faces are processed in a distributed, hierarchical neural system of face perception (Kanwisher et al., 1997; Gauthier et al., 2000b; Hoffman and Haxby, 2000; Haxby et al., 2000). More specifically, facial identity is processed primarily along a ventral occipito-temporal stream with refined processing steps proceeding from the basic analysis of isolated facial features to the structural encoding of holistic, partially view-dependent individual face representations to the establishment of modality-independent personal recognition memories (Pourtois et al., 2005; Quiroga et al., 2005).

Prosopagnosia, colloquially also referred to as "face-blindness", is defined as a profound deficit in the specific task of face identification (Bodamer, 1947). This deficit can be either acquired due to brain damage (see e.g. Mazzucchi and Biber, 1983, for a review of 74 cases), or it is present from birth, i.e. congenital (Kress and Daum, 2003; Hasson et al., 2003; Behrmann and Avidan, 2005; Grueter et al., 2007; Kennerknecht et al., 2008b). Congenital prosopagnosia (CP) is highly familial (McConachie, 1976; De Haan, 1999; Duchaine and Nakayama, 2006; Kennerknecht et al., 2006, 2007, 2008b,a). We therefore coined the term hereditary prosopagnosia (HPA) which can be used synonymously to CP (Kennerknecht et al., 2006). Yet, when initially asked most index subjects are not aware of other impaired family members unless actively interviewed and probed into. Behavioral studies of congenital prosopagnosia have revealed a dissociation between face and object recognition deficits (Gauthier et al., 2004; Duchaine and Nakayama, 2005; Duchaine, 2006), between face detection and face recognition (Garrido et al., 2008), and between the processing of facial identity and facial expressions (Humphreys et al., 2007), either by testing single aspects in isolation or by conducting a battery of tests with the same participants (Behrmann et al., 2005; Le Grand et al., 2006; Schmalzl et al., 2008a; Garrido et al., 2009b).

The original symptomatic characterization of prosopagnosia by Bodamer (1947) clearly states what prosopagnosia is. But it only includes a vague specification of the processing differences underlying the deficit: “With unimpaired perception of the formal parts of physiognomies, the process of recognition fails” due to an inability to perceive “the structured picture making up an individual, personal whole” (translations taken from Ellis and Florence, 1990). Following up on this characterization of the processing deficits at the computational level we proposed that in CP the failure in integrating information is compensated by a strategy of serially processing informative face parts in isolation (Stollhoff et al., 2010). Such a serial, featural processing can explain the observation of more frequent eye-movements and dispersed gaze behavior in CP (Schwarzer et al., 2007; Schmalzl et al., 2008b), and an increase in inspection or reaction times (Behrmann et al., 2005; Stollhoff et al., 2010). As an intermediate step between processing faces via isolated face parts or as undifferentiated wholes, holistic encoding (Farah et al., 1998), deficits in processing changes in the configuration of features, i.e. the spatial arrangement of face parts, have been documented in cases of acquired prosopagnosia (Barton and Cherkasova, 2005; Barton et al., 2003, 2002).

So far, prosopagnosia has only been modeled in the acquired case where an existing, functional face recognition system suffers from an externally inflicted damage. The models of acquired prosopagnosia can be roughly divided into two classes: Conceptual models with a focus on neurophysiological correspondence between the functional deficit and the location of the lesion (Breen et al., 2000; Ellis and Lewis, 2001; Fox et al., 2008), and abstract, computational models with a focus on task differences in the required recognition or recall accuracy (Virasoro, 1988, 1989; Farah et al., 1993; Burton et al., 1999; Pessa et al., 1999; Zifan et al., 2007). In both classes of models a fully functional, mature system is degraded, e.g. by removing nodes or clipping connections. Evaluation of the model is then based on comparing the properties of the network before and after degradation either descriptively, analytically, or numerically using simulation studies. Here, a neural network model of congenital prosopagnosia (CP) is derived from formal considerations, implemented in an artificial neural network model of facial encoding, and tested empirically in experiments with prosopagnosic and control participants. The main accomplishment of the model is to provide a direct, testable link between the critical features of congenital prosopagnosia as the lack of holistic processing at the computational level and a reduced structural connectivity of face processing areas at the level of neuronal implementation.

1.1 Neuroanatomy of Face Processing

Under normal conditions the brain develops a specialized neural system for face recognition (de Haan et al., 2002; Scherf et al., 2007; Polk et al., 2007) which has been further differentiated into spatially segregated functional processing modules (Kanwisher et al., 1997; Gauthier et al., 2000c; Hoffman and Haxby, 2000; Haxby et al., 2000; Kawashima et al., 2000; Grill-Spector et al., 2004; Gauthier et al., 2005). Irrespective of the exact developmental processes underlying the functional specialization, damage inflicted to a specific cortical region can therefore lead to restricted deficits conditional on the interconnectedness and interdependence of the distributed processing (Damasio et al., 1990; De Renzi et al., 1991, 1994; Fox et al., 2008). More specifically, the behavioral heterogeneity in acquired prosopagnosia can largely be explained by differences in the extent and location of the brain damage causing the deficit (Damasio et al., 1990; De Renzi et al., 1991, 1994; Fox et al., 2008).

In contrast, functional imaging studies of CP have so far found no unequivocal evidence for activation differences in this region; neither using classical localizer paradigms (Hasson et al., 2003; Avidan et al., 2005) nor adaptation paradigms (Avidan et al., 2005; Avidan and Behrmann, 2009). First indications of structural neuroanatomical differences point to a volumetric reduction of the anterior fusiform gyrus (Behrmann et al., 2007) and the anterior inferior temporal lobe (Garrido et al., 2009b), regions involved in more associative and mnestic aspects of face recognition (Haxby et al., 2000). Analysis of a large group of CP participants, revealed diminished gray matter density in the lingual gyrus bilaterally, the right middle temporal gyrus and the dorsolateral prefrontal cortex (Dinkelacker et al., 2010). In a diffusion tensor imaging study, Thomas et al. (2009) reported a reduced structural connectivity in the ventral occipito-temporal white matter tracts, presumably involved in more apperceptive aspects of face recognition. In our model, these observations provide the rationale for implementing the structural differences in CP as a predisposition towards a reduced network connectivity between regions involved in the structural encoding of facial information.

1.2 Modelling Principles

Possible alterations in the process of structural encoding of face images will be studied in the framework of single-layer feedforward networks. The input units of the network register the stimulus which is then encoded into a sensory description based on the activation of the output units. In a single-layer feedforward network, the input stimulus $X = (X_1, \ldots, X_D)$ is mapped to the output activation $S = (S_1, \ldots, S_n)$ by

$$S_j = h\left(\sum_{d=1}^{D} w_{jd} X_d\right),$$

where $w_{jd} \in \mathbb{R}$ is the weight associated with the connection from the d^{th} input unit, X_d, to the j^{th} output unit, S_j, and h is the activation function, in matrix notation: $S = h(\mathbf{W}X)$. The network thus translates an observable input vector or stimulus X into an internal representation S, which can be used for further processing. While in this formulation the projection of input to output units involves only feedforward computations, training of the network, e.g. weight adaptation, often draws on information that is not available locally, e.g. the activations of the other output units.

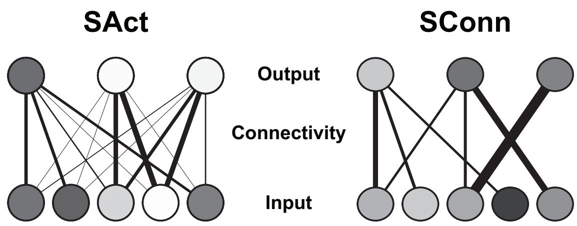

Figure 1. Schematic depiction of the two network models, where output units (top) are connected to input units (bottom). Line strength is chosen to represent connection strength and shading of the units represents activation (light = low activation, dark = high activation). The sparseness constraint is either placed on output activations in the SAct model (left) or on network connectivity in the SConn model (right).

For our models of functional and dysfunctional facial encoding we trained single-layer feedforward networks to represent a set of frontal face images under two different constraints (see figure 1). On the one hand, we introduced a constraint on the sparseness of the output unit activation which was implemented at the algorithmic level using spatial ICA. Sparseness in neural activation is commonly observed in the human cortex (Olshausen and Field, 1997) and is assumed to serve an important purpose in information processing (Field, 1994). We will call this model of functional encoding of facial identity the Sparse Activation (SAct) model. On the other hand, we introduced a structural constraint on sparse network connectivity between input and output units as the basis of a model of dysfunctional encoding in CP which was implemented using temporal ICA. We will call this model of dysfunctional encoding of facial identity the Sparse Connectivity (SConn) model.

1.3 Principles for the Experimental Validation

In general, it is difficult to derive and test quantitative predictions of computational models for actual human behavior. First, the kind of simple feedforward networks used here, operating with a very small number of units and connections, can only be a strong simplification of the actual neuronal architecture found in the human visual cortex. Second, the dataset of face images provided to train the models is vastly smaller than the dataset likely to be used by humans to evolve their face recognition capabilities. Third, human performance under tightly controlled experimental conditions is probably different from everyday performance. Particularly because experimental conditions often allow an extensive use of compensatory strategies like feature matching.

In the experimental validation, we adopted an indirect way of testing the models by using the different model representations of faces to construct specific test stimuli. In a similarity judgement task, we assessed each participant's ability to discriminate between face images manipulated either according to the SAct model of functional encoding or according to the SConn model of dysfunctional encoding of facial identity. More specifically, we constructed morph series of face images where the direction of each morphs series was obtained directly from the model representations. In the test participants were presented with a display of three face images: An average face in the center and two different test faces on the left and right side. The test faces differed either only in the degree of morphing (experiment 1) or additionally in the direction of morphing (experiment 2).

Based on the hypothesis that the face recognition deficit in congenital prosopagnosia can be modeled by a sparseness constraint on network connectivity, it was assumed that congenital prosopagnosics would perform worse than controls in discriminating between manipulations according to the SAct model but would show no impairment with respect to SConn model.

2 Models of Facial Encoding

Single-layer feedforward networks have previously been used in representing face images. A prominent example is the "Eigenface" approach based on a principal component analysis (Turk and Pentland, 1991). In terms of neural networks, projection onto the first principal component can be implemented in a network with a single output neuron by Oja's rule (Oja, 1982); projection onto the first Q components in a network with Q output neurons by the Generalized Hebbian Algorithm (Sanger, 1989), a combination of Oja's Rule and a Gram-Schmidt orthogonalization process. A network performing PCA provides a representation in terms of uncorrelated output units. For non-Gaussian distributed inputs X, the requirement of uncorrelated representations in PCA can be strengthened by assuming independent representations, which is then referred to as independent component analysis or ICA (Hyvarinen et al., 2001). ICA is often applied in two different architectures depending on the task: Temporal ICA to model time series data and spatial ICA to model static, multidimensional data (Bartlett et al., 2002; Draper et al., 2003). In the context of face representations, one assumes that every given image X is the linear projection A of a set of source activations S, where the columns of A are called basis images. To roughly summarize the difference between the two architectures: Whereas spatial ICA assumes that the source activations S are independent and sparsely distributed, temporal ICA places a sparseness constraint on the pixel activations in the basis images (see table 1).

	spatial ICA	temporal ICA
observations	face images: $(x_1, \ldots, x_N)$	pixel sequences: $(p_1, \ldots, p_D) = (x_1, \ldots, x_N)^t$
generative model	$X = AS$	$X^t = (AS)^t$
sparse distribution	on sources S_j	on connectivity $A_i j$

Table 1. Differences between spatial and temporal ICA

Previos applications of ICA for face representation have revealed striking differences between spatial and temporal ICA (Bartlett et al., 2002; Bartlett, 2007; Draper et al., 2003). If the dataset consists of observations of different individuals and the task is to construct a representation separating the individuals, spatial ICA leads to a better performance (Draper et al., 2003). Intuitively, output units signal the presence of any individual, activations should be binary, and basis images are equal to the (mean-centered) observations of individual faces. More generally, output unit activations will be close to zero for most of the population samples and non-zero only in a local neighborhood around the original individual faces. In contrast, if observations are of different individuals and each individual is observed under various non-rigid transformations which have to be classified (e.g. different facial expressions), temporal ICA yields better representations (Draper et al., 2003). Intuitively, output unit activations correspond to different small manipulations of a face. Each manipulation only affects a small subset of the original features (sparseness constraint on basis images) and each transformation is modeled by the combination of several small manipulations. In facial expressions, the small manipulations could correspond to mimics affecting only parts of the face that can be activated in different expressions.

In general, ICA can be formulated in (at least) three different but equivalent ways: Either as a maximum likelihood estimation of source activations, or as the maximization of output entropy or equivalently the minimization of multiinformation in a feedforward network with an output nonlinearity (see Hyvarinen et al., 2001, for a more detailed treatment). In the context of feedforward networks, the assumptions of independent sparsely distributed source activations can be interpreted as a prior on output unit activations $S = h(\mathbf{W}X)$ that are independent and sparsely distributed. For example, Foeldiak (1990) introduced a direct Hebbian/anti-Hebbian learning rule to foster sparse representations in a single layer network with lateral connectivity. This was further elaborated by Olshausen and Field (1997), where an equivalence is established to the Infomax-ICA by Bell and Sejnowski (1995). To emphasize the aspect of sparse activations we will refer to the functional model of facial encoding as the SAct model. To implement the SAct model, we used a maximum likelihood implementation of spatial ICA and assumed a leptokurtic or super Gaussian distribution for the source activations S.

In contrast, in our implementation of a model of CP as a dysfunctional model of face encoding, we used a maximum likelihood implementation of temporal ICA. We again assumed a leptokurtic distribution for source activation which now corresponds to pixel activations in basis images. As will be shown later, placing a constraint on pixel activations in basis images A leads to similar network properties as placing a constraint on the network connectivity W. Intuitively, this can be explained by the equivalence of feedforward connectivity and basis images ($A = W^{-1}$). In keeping with the network interpretation, we will refer to this model as the *SConn model*.

In addition to these two pure models of sparse activations and sparse connectivity we constructed intermediate models based on a spatio-temporal ICA. Spatio-temporal ICA (Stone et al., 2002) combines spatial and temporal ICA in that it makes the same distributional assumptions on both the source activations, as in spatial ICA, and the basis images, as in temporal ICA. This is achieved by treating the weights (inverse basis images or forward connectivity) as random variables which follow a sparse prior distribution. This formulation changes the optimization problem from a univariate maximization of the source activation likelihood $P(\mathbf{S}|\mathbf{W})$ with fixed weights $\mathbf{W}$, as in spatial ICA, to maximizing the conditional posterior $P(\mathbf{W}|\mathbf{S}) \propto P(\mathbf{S}|\mathbf{W})P(\mathbf{W})$. To investigate the influence of the additional sparseness constraint on weights, we introduced a sparseness parameter α with higher values of α corresponding to more sparse weight distributions.

2.1 Materials and Methods

In our implementation of spatial, temporal, and spatio-temporal ICA we closely followed the setup of Bartlett et al. (2002), treating full images as pixel vectors and applying a PCA for dimensionality reduction and whitening prior to conducting the ICA itself. A dataset of 200 frontal face images was used to construct PCA and ICA representations. PCA was only used as a pre-processing step. ICA was applied to the PCA projections of the images as spatial, temporal or spatio-temporal ICA. Spatio-temporal ICA was implemented for five different sparseness constraints on the network connectivity.

2.1.1 Face Images and Image Pre-Processing

Images were taken from the publicly available Face Database of the MPI for Biological Cybernetics (http://faces.kyb.tuebingen.mpg.de/), which contains 200 head models, obtained by laser-scan (Cyberware TM) and slightly morphed to avoid resemblances to actual individuals (Troje and Bülthoff, 1996; Blanz and Vetter, 1999). In the simulations, only frontal views were used.

All images were converted to gray levels, trimmed to the facial outlines, normalized in size such that the face width x face height equaled 2628 pixels^2 and embedded into a 50x62 image by adding black borders. Size variations in width and height - as measured by the Fano factor $\frac{\mathrm{Var}(X)}{E[X]}$ - in the normalized images are about 75% of variations found in direct anthropometric measurements of the human population (Farkas, 1981). Images were taken as vectors of length 3100 which resulted in an image dataset $\mathbf{X} \in \mathbb{R}^{3100\times200}$. The dataset was centered to zero mean for every pixel.

To facilitate computations the image dataset was not used directly as input to the ICA, but only projections onto the first 50 principal components, i.e. eigenvectors of $\mathbf{X}\mathbf{X}^t$ were used as input ($\mathbf{U}_{50} := (\mathbf{V}_{50})^t\mathbf{X}$), leading to temporal ICA being applied to a dataset with a reduced number of observations ($\mathbf{Z}_T = \mathbf{V}_{50}^t \in \mathbb{R}^{50\times3100}$) and spatial ICA to a dataset with reduced dimensionality ($\mathbf{Z}_S = \mathbf{U}_{50} \in \mathbb{R}^{50\times200}$).

2.1.2 Implementation

A maximum likelihood implementation of ICA was used in the application to face images, with identical probability densities for source activations S_q

$$p_S(s) \propto (\cosh(s))^{-2},$$

where in the following the subscript for the density will be dropped for ease of notation. The empirical log-likelihood of source activations for given observation Z with $S = W\,Z$

$$l(\mathbf{S}|\mathbf{W}) = \sum_{n=1}^{N}\left[\sum_{q=1}^{Q}[\log\, p(\mathbf{W}_{q\cdot}z_n)] + \log\,|\det\, W|\right] \tag{1}$$

was maximized w.r.t. the weight matrix $\mathbf{W}$, where the z_n (and Z respectively) are either projections of the observations of individuals onto the first 50 principal components (spatial ICA or SAct) or the first 50 principal components themselves, i.e. as features (temporal ICA or SConn).

In terms of the original images x_n, for spatial ICA with $S_S = W_S\,X$, (1) can be expressed as

$$\sum_{n=1}^{N}\left[\sum_{q=1}^{Q}[\log\, p(\underbrace{(\mathbf{W}\mathbf{V}_{50}^{t})_{q\cdot}x_{n}}_{(S)_{qn}})] + \log\,|\det\, W|\right]$$

and temporal ICA with $S_T = W_T X$ as

$$\sum_{n=1}^{N}\left[\sum_{q=1}^{Q}[\log\, p(\underbrace{\mathbf{W}_{q\cdot}(\mathbf{V}_{50}^{t})_{\cdot n}}_{(W_T)_{qn}})] + \log\,|\det\, W|\right],$$

which shows that a sparse distribution p is assumed either for the source activations S_S (SAct) or for the weight matrix W_T, i.e. the connectivity between input and output nodes (SConn).

For spatio-temporal ICA (sICA$_\alpha$) non-uniform, sparse priors on the weight matrix were assumed to follow Laplace distributions with zero mean and variable variance, i.e.

$$\log(p_\alpha(w)) = -\alpha|w| - \log(\frac{2}{\alpha}),\ \text{for}\ \alpha > 0. \tag{2}$$

Using the approximation $\alpha|w| \approx \log(\cosh(\alpha w))$ the log of the posterior can be written as

$$\begin{aligned}\log(P_\alpha(\mathbf{W}|\mathbf{Z})) &= \sum_{q=1}^{Q}\left[\sum_{n=1}^{N}[\log p(\mathbf{W}_{q\cdot}z_n)] + \sum_{q'=1}^{Q}[\log p_\alpha(\mathbf{W}_{qq'})]\right] + N\log|\det\, W| + c(\alpha)\\ &= \sum_{q=1}^{Q}\left[\sum_{n=1}^{N}[\log p(\mathbf{W}_{q\cdot}z_n)] + \sum_{q'=1}^{Q}[\log p(\mathbf{W}_{q\cdot}\alpha e_{q'})]\right] + N\log|\det\, W| + c(\alpha),\end{aligned}$$

where the z_n or Z are the projections onto the first 50 principal components, analogous to spatial ICA, and the $e_{q'}$ is the q'-th canonical unit vector. This formulation allows a simpler algorithmic implementation of the prior assumptions in terms of additional (weighted) observations of unit vectors and allows to use the fast optimized algorithms available for spatial ICA (Hyvarinen et al., 2001).

After parameter estimation one obtains least-squares ICA/PCA reconstructions of the original images as

$$
\begin{aligned}
\mathrm{ICA}_T\text{:} \quad \hat{\mathbf{X}} &= \quad (\mathbf{S}_T)^t \left((\mathbf{A}_T)^t \, \mathbf{U}_{50}\right) \\
\mathrm{ICA}_S \text{ (sICA}_\alpha\text{):} \quad \hat{\mathbf{X}} &= (\mathbf{V}_{50}\, \mathbf{A}_S)\; \mathbf{S}_S
\end{aligned}
$$

In the models the first component includes the basis images (columns of $\mathbf{S}_T^t$ and $\mathbf{V}_{50}\mathbf{A}_S$ respectively) while image activations for single basis images are given in the columns of the second component. Note that the sparseness assumptions on $\mathbf{S}_T$ and $\mathbf{S}_S$ affect different components in the corresponding models.

The implementation of ICA was custom-written in R (R Development Core Team, 2009) using the natural gradient descent algorithm outlined in Hyvarinen et al. (2001).

2.1.3 Analysis of Face Representations

The resulting face representations were analyzed in terms of basis images, weight matrices and source activations for the given dataset.

Density functions are estimated from empirical data using a Gaussian kernel density estimator with automated bandwidth selection (Silverman, 1986). Correlations are Pearson product-moment correlations.

Since ICA models are only specified up to scalar multiplication, a standardization procedure was required to compare different representations. Empirical source activations were standardized by subtracting average activation and dividing by the empirical standard deviation. For base image activations and weight matrices activation profiles were calculated by subtracting background activation from basis image pixel activations (or weight vector strength respectively), taking the absolute values and scaling to unit sum.

2.2 Results

The independence and sparseness constraint on the source activations imposed in the SAct model are reflected in a leptokurtic distribution (figure 2 A) with negligible correlations

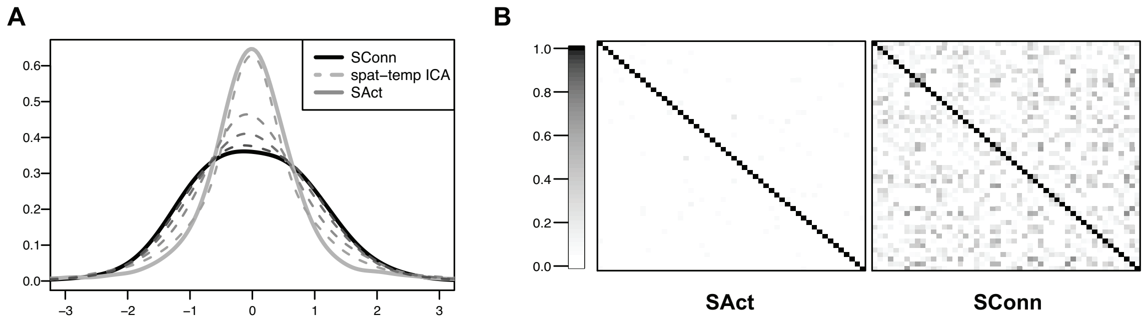

Figure 2. Source activations in SConn, SAct and spatio-temporal ICA representations of the face dataset. **(A)** In the SAct model a sparseness constraint is imposed on source activations. Increasing the sparseness constraint on network connectivity in spatio-temporal ICA (from $\alpha = 5$ to 10, 20, and 50, ordered from top to bottom) leads to more distributed activations as in the SConn model. **(B)** Source activations are uncorrelated in the SAct model but become correlated under the constraint of sparse connectivity in the SConn model.

among sources (figure 2 B). In contrast the SConn model imposed no constraints on source activations, which led to approximately normally distributed, but correlated activations. Activations in spatio-temporal ICA covered a wide range from almost normal distributions for high values of α, i.e. string sparseness constraint on the weight matrix, to highly peaked activations for low values of α.

A sample of base images obtained for the different face representations is shown in figure 3. Images were selected according using the same random index vector for all methods; column-wise similarities in the display are due to similarities in the methods and the fact that we always used exactly the same data. In the SAct model faces are represented by base images with a distributed activation pattern. Introducing an additional sparseness constraint on the weights in spatio-temporal ICA led to basis images in which activation is the more restricted to selected regions the higher the sparseness parameter α. For low values of α, the activation profiles exhibit a larger spread. For high values of α, the resulting basis images tend to be similar to the ones obtained in the SConn model. Although it was not specified in the assumptions of either the spatio-temporal ICA or the SConn model, pixels with high activations tend to cluster in specific regions.

Connectivity weights in the SConn model and in spatio-temporal ICA are more closely

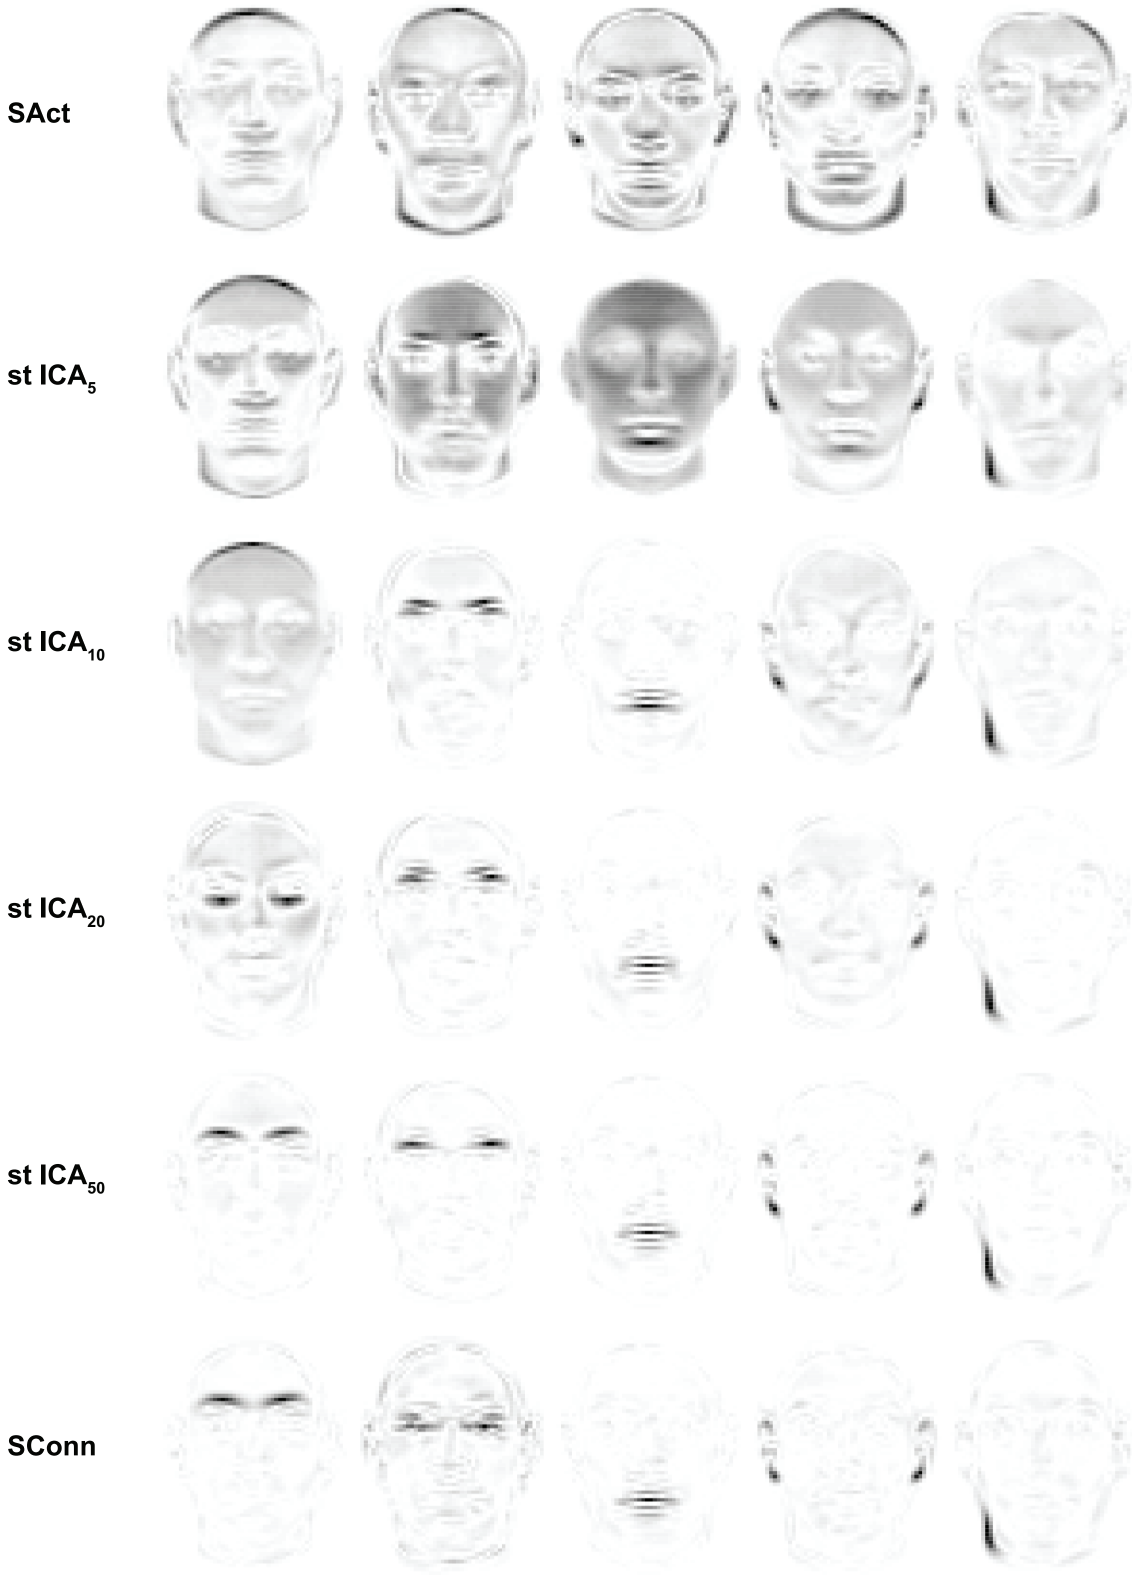

Figure 3. Graphical display of a randomly selected subset of base images (columns) for the different face representations (rows). Shown are pixel value activation profiles, with dark regions indicating areas where the basis image deviates from background.

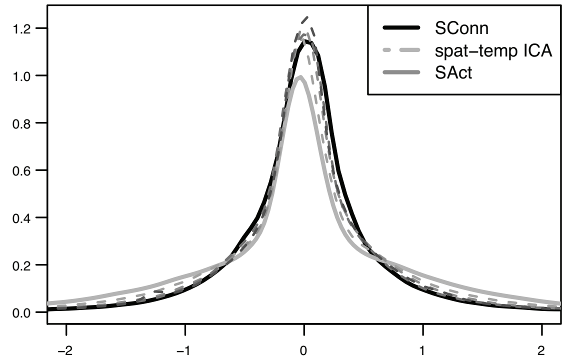

Figure 4. Distribution of connectivity weights vectors in SConn, SAct and spatio-temporal ICA models. The sparseness constraint on weight strength leads to more connections being closer to zero.

clustered around zero than in the SAct model (figure 4). Increasing the sparseness parameter in spatio-temporal ICA from low to high values (lines from top to bottom) increases the proportion of weights with values close to zero. Note that the approximation ($\alpha|w| \approx \log(\cosh(\alpha w))$) used in the derivation of the specific implementation of spatio-temporal ICA leads to more mass being concentrated close to zero compared to the SConn model. However this only slightly distorts the general similarity between the SConn model and spatio-temporal ICA with a large α on the one hand and the SAct model and spatio-temporal ICA with low αs on the other hand.

3 Experimental Validation

The main focus of the previous section was to show that placing different constraints on the network architecture leads to representation of faces with different functional properties. The translation from structural constraints to functional differences is essential in bridging the gap between the neuroanatomical and behavioral differences found in congenital prosopagnosia. More specifically, we propose that in congenital prosopagnosia an initial sparseness constraint on the neuronal connectivity leads to face representations based on local features (SConn); Such a representation is less sensitive to the global variations that are normally used to individuate face images (SAct).

Initially, each of the two models was trained to represent a set of frontal face images where each face image was represented as a linear combination of 50 basis images. Using these representation we constructed morph series of face images. Starting from an average face each morph series was constructed by moving along the direction provided by one of the basis images. In total, we chose six different basis images as morph directions, three basis images of the representation of the functional and three of the dysfunctional model. The transition stepsizes in each morph direction were normalized by the variance in the corresponding source activation. This standardization ensured that morph series, constructed using different representations and basis images, are comparable in their typicality, as measured by the distance from the average face.

To compare model representations with human representations, we applied a similarity judgment task. Presented with a display of three face images - an average face in the center and two different faces on the left and right side - the participants had to judge which of the two images on the sides is more similar to the image in the center. In the first experiment, the test faces on the left and right side were chosen out of the same morph series and differed only in the degree of morphing (experiment 1). In the second experiment, one face was chosen out of an SAct morph series and the other face out of a SConn series. Additionally, the morph degrees were varied for each of the two morph series independently.

Based on these similarity judgements we calculated participants discrimination information as the Kullback-Leibler divergence from a random judgement. In practice, participants obtained a high discrimination information if they consistently rated the morph closer to the average as being more similar to the average. In the second experiment we additionally investigated the frequency of judging the SConn manipulation being more similar to the average face than the SAct manipulation to calculate participants' response bias in favoring one manipulation over the other.

Based on the hypothesis that the face recognition deficit in congenital prosopagnosia can be modeled by a sparseness constraint on network connectivity, it was assumed that congenital prosopagnosics would perform similar to controls in isolated discriminations between SConn manipulations and worse for isolated discriminations between SAct manipulations (experiment 1). Furthermore, we expected that compared to controls CP participants would show no or a decreased bias for SAct manipulations in comparison to the SConn manipulations (experiment 2).

3.1 Materials and Methods

3.1.1 Face Images and Model Representations

Overall images were generated was analogously to the simulation study. Changes included an increase in the number of images and the dimensions of the images. Images were constructed using head models obtained by the MPI for Biological Cybernetics, Tübingen, Germany (Troje and Bülthoff, 1996; Blanz and Vetter, 1999). Based on a subset of 142 head models (75 female, 67 male), obtained by laser-scan (Cyberware TM), we added 50-50 morphs between every pair of female/male models to obtain an enlarged dataset of 5128 head models. From each model a frontal snapshot was taken as a face image.

All images were converted to gray levels, trimmed to the facial outlines, normalized in size such that the face width x face height equaled 11594 pixels^2 and embedded into a 100x150 image by adding black borders. Size variations in width and height - as measured by the Fano factor $\frac{\mathrm{Var}(X)}{E[X]}$ - in the normalized images were about 50% of variations found in direct anthropometric measurements of the human population (Farkas, 1981). Images were taken as vectors of length 15000 which resulted in an image dataset $\mathbf{X} \in \mathbb{R}^{15000 \times 5128}$. The dataset was centered to zero mean for every pixel.

To facilitate computations we performed pre-PCA projecting onto the first 142 principal components of the covariance matrix ($\mathbf{U}_{142} := (\mathbf{V}_{142})^t \mathbf{X}$), leading to temporal ICA being applied to a dataset with a reduced number of observations

($\mathbf{X}_T = (\mathbf{V}_{142})^t \in \mathbb{R}^{142\times 15000}$) and spatial ICA to a dataset with reduced dimensionality ($\mathbf{X}_S = \mathbf{U}_{142} \in \mathbb{R}^{142\times 5128}$).

For each ICA representation, out of a total of 142, three base images were chosen as directions of image manipulation: Three local manipulations obtained by temporal ICA, and three global manipulations obtained by spatial ICA. Selection was based on a graphical inspection of the base images to ensure that local manipulations in temporal ICA were distributed across facial regions (eyes, mouth, nose) and to exclude artifacts in face manipulations that are obvious to a human observer. For each of the six directions a series of six morphed images was constructed by pixel-wise linear addition of the basis image $b_{T,i}(i = 1, 2, 3)$ ($b_{S,i}$ for spatial ICA) onto the average face image x_0, i.e.

$$x_{T,i}(n) = x_0 + \frac{n}{2}\,\delta_{T,i}\,b_{T_i}. \tag{3}$$

For each direction stepsizes $\delta_{T,i}$ were taken as the standard deviation in source activation across all 5128 images.

The standardization lead to images that for each n have equal Mahalanobis distance to the mean image across dimensions. For a pair of observations x, y drawn from a sample with covariance matrix Σ the Mahalanobis distance d_M is defined by

$$d_M(x, y) := \left[(x - y)^T \Sigma (x - y)\right]^{\frac{1}{2}}$$

This ensured that for each direction there is a clear increase in distance between the average face and the constructed face. Also, if the Mahalanobis distance is used as a general measure of distance between any two source activation vectors, manipulations in different directions but of the same order, i.e. multiplicity of standard deviation, have exactly the same Mahalanobis distance to the average face.

While in the original images pixel (grayscale) values were confined to the interval $[0, 1]$

the linear manipulation of images in (3) may lead to pixel values outside the admissible range. Therefore first values below zero were set to zero and values above one to one and afterwards the mean value of each image was normalized to the mean value in the average face. For presentation in the experiments, all images were enlarged to 200x300 pixels by linear interpolation.

3.1.2 Experimental Setup

To study processing of the two different ICA representations two experiments were performed.

The first experiment tested discrimination accuracy for manipulations in a single ICA direction. On each trial participants were presented with a display of three face images: the average face in the center and two faces manipulated in one direction - the morph endpoint with a distance of ($n = 6$) and the morph at an intermediate distance ($1 \leq n \leq 5$) from the average face, which will be referred to as test face in the following. Participants had to indicate by mouse-clicks whether the image to the left or to the right was more similar to the average face image in the center. All combinations of six directions and five distance were presented ten times arranged into blocks of length 30. Order of the combinations and presentation position (left/right of average face) was randomized across blocks.

The second experiment tested similarity judgments across two different image manipulation directions. On each trial the participant was presented with a set of three images: the average face in the center and two faces manipulated according to two different manipulation directions. Again the task was to indicate which image was more similar to the average face image in the center. In each direction image manipulation was performed in six steps, which yields a total of (6*5)/2*6*6=540 trials. Trials were arranged in blocks of 15, such that each direction combination is contained in each block exactly once. The order of stepsizes over trials and presentation positions in each trial (to the left/right of the average face) was randomized. Here, only inter-model comparisons, i.e. trials with one SConn and one SAct direction are reported.

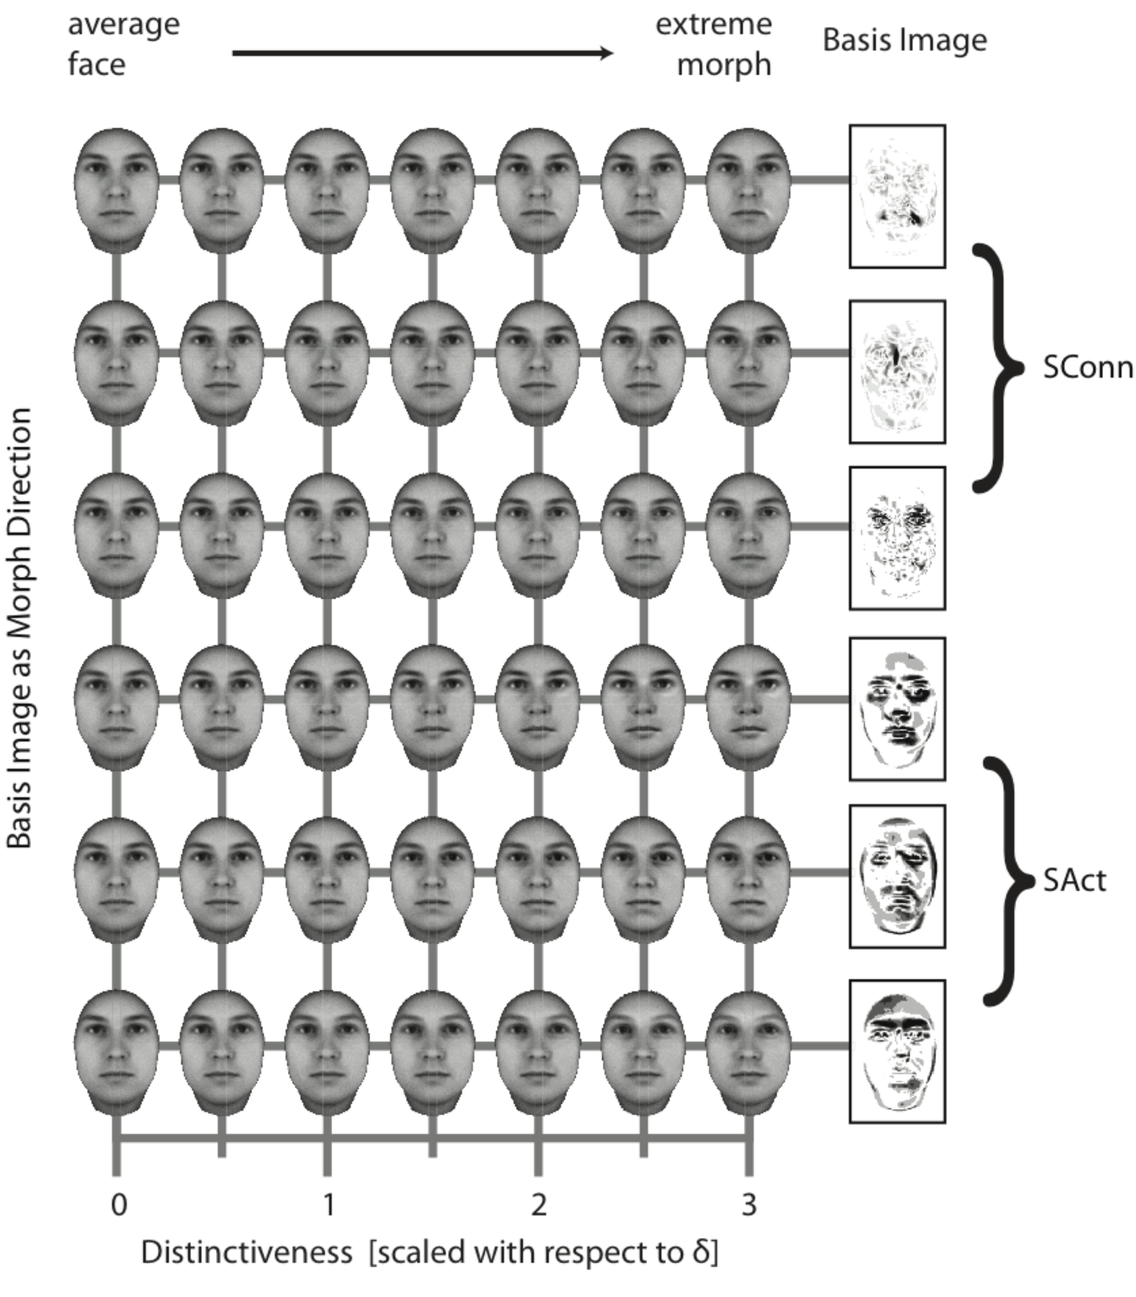

Figure 5. Face images were obtained by morphing into three SConn directions (upper half, with manipulations restricted to the mouth, nose and eye region) and into three SAct directions (lower half, all global manipulations). Each row contains the average face image on the left and morphed images in stepsizes of $\frac{1}{2}$ standard deviations. For small stepsizes changes are barely noticeable, while for large stepsizes changes become more visible especially for changes in SAct directions. Images on the right show the basis images which were used as morph directions; the darker a region is depicted the stronger the change in this region as compared to the average face. (absolute values, rescaled).

Figure 6. Participants were shown three facial images on a black background and had to decide whether the test image on the left or on the right was closer to the image in the center (average face). **(A)** In the first experiment, test images were chosen as extreme and intermediate values of the same direction (left image: extreme at 3σ; right image: intermediate at σ)). **(B)** In the second experiment, test images were chosen from different directions and morph levels (left image: SAct, 3σ; right image: SConn, 3σ).

The two experiments were actually administered in reverse order, i.e. inter-model comparison first, to prevent participants from familiarizing with the single morphs directions in each model before having to make inter-model comparisons.

Presentation was on a Toshiba Satellite Pro 6100 with an Nvidia GeForce4 420 Go graphics adapter and a 15” (28.7 cm x 21.5 cm) TFT display with a resolution of 1024 x 768 pixels. Images subtended 5.6 cm x 8.4 cm, which corresponds to a visual angle of 5.6° x 8.4° at the initial seating distance of 1 m. Participants were free to take as much time for their decision as they wanted. All experiments were run using the visual psychophysics software FlashDot (Elze, 2009).

3.1.3 Participants

A total of 30 subjects participated in the experiments, 10 participants with congenital prosopagnosia (CPs) and 20 age- and mostly gender-matched controls.

All of the 30 participants, successfully completed both experiments. Post-hoc analysis of the data revealed that one control participant responded randomly with a response that was unrelated to experimental conditions and reaction times between 227 ms and 702 ms, well below the observed minimum reaction time of 769 ms among all other participants. This participant was thus excluded from further analysis.

Diagnosis of CP was based on a semi-structured interview which includes questions on everyday-problems with face and object recognition, mental imagery and avoidance strategies (Grueter et al., 2007; Kennerknecht et al., 2008b). Overall, most CPs had normal or borderline normal basic-level object recognition abilities as measured by BORB tests 6,7,10,13 (Riddoch and Humphreys, 1993) and VOSP tests 2,4,6 (Warrington and James, 1991). For CPs the face recognition deficits were confirmed in a series of experiments testing short-term (Stollhoff et al., 2010) and long-term (Stollhoff et al., 2011) recognition performance of faces and objects.

All CPs and controls provided written informed consent before participation. The study was approved by the ethical committee of the University of Münster, Germany, protocol No 3XKenn2.

3.1.4 Statistical Analysis

In general, if participants base their judgement of similarity on perceived distances between the images and the average face, the fraction of times the left image is considered more similar than the right image can be modeled as some function of the distance ratio, e.g.

$$F_{\text{left}}(x_{\text{left}}, x_{\text{right}}) = g\left(\frac{d(x_o, x_{\text{left}}) + \eta_1}{d(x_0, x_{\text{right}}) + \eta_2}\right), \tag{4}$$

where g is a monotonically decreasing function $d(x, y)$ is a similarity measure and $\eta_i (i = 1, 2)$ are random variables interpretable as noise terms.

In the first experiment responses of participants were either defined as correct, if they selected the image with the lower distance to the average face, or false, otherwise. To assess accuracy in processing manipulations in a single morph direction, we used the Kullback-Leibler (KL) divergence of the observed fraction of correct answers p from a uniform distribution (50-50 decision). In this context Kullback-Leibler divergence can be interpreted as the expected discrimination information for discriminating individual performance, p, from chance performance.

$$D_{KL}(p||\tfrac{1}{2}) := p\,\log 2p + (1-p)\,\log 2(1-p) \tag{5}$$

The Kullback-Leibler divergence was calculated for each of the six morphing directions and averaged for each standardized distance across all SAct or SConn directions respectively.

In similarity judgments across different morphing directions, as measured in the second experiment, it is difficult to define correct and false answers. We calculated the fraction of trials were SConn are rated closer to the average face than SAct manipulations as a measure of participant's bias between the two model representations. Additionally we again calculated the Kullback-Leibler divergence from a uniform distribution as a measure of precision of participants' decision criteria. Both the bias and the Kullback-Leibler divergence were calculated separately for each pair of morphing directions and then averaged across all comparisons of SAct and SConn comparisons.

In testing group differences between controls and CPs, we used model based comparisons based on families of nested (generalized) linear mixed models (GLMMs, see e.g. Tuerlinckx et al., 2006). This approach allows to include structural differences between participants (e.g. age) as well as individual variations (random effects) in the statistical model. For each analysis we fitted a nullmodel which included participants age as well as the degree(s) of morphing in the test images (σ or σ_{SAct} and σ_{SConn} resp.) as fixed effects (with parameters β, β_{SAct}, and β_{SConn} resp.), and participants' identity as a random effect. In the second experiment we also included the squared distance between the degrees of morphing in the SConn and the SAct direction ($\Delta(\sigma)$ or Δ with parameter β_{Δ}) in our analysis of differences in Kullback-Leibler divergence.

Based on the nullmodel a nested family of GLMMs was fitted by including first a main effect of group differences ($\beta_{CP,0}$), and second interaction effects, i.e. different slopes for the influence of the fixed effects, (e.g. $\beta_{CP,SAct}$). In this nested family we selected the best fitting model using the Bayesian Information Criterion (Karabatsos, 2006). To provide interpretable parameter estimates, we also calculated Bayesian maximum posterior estimates ($\hat{\beta}$) and highest posterior density intervals with 95% support ($\mathrm{HPDI}_{95\%}$) for group differences in main and/or interaction effects in the selected model.

Data analysis and statistical testing was done in the statistical programming language R (R Development Core Team, 2009). Fitting of generalized linear mixed models (GLMMs) was done using the R packages `lme4` (Bates and Maechler, 2009) and `MCMCglmm` (Hadfield, 2009). The algorithms used in `lme4` as well as the model based comparisons conducted here, are described by the main contributor to the `lme4` package in more detail in Faraway (2006). As prior distributions for the Bayesian model fitting we used a multivariate normal distribution with zero mean and a diagonal covariance matrix with large variances (order of 10^{10}) for fixed effects and an inverse Wishart distribution with degrees of freedom equal to one and the inverse scale equal to the unconditional variance of the response variable.

3.2 Results

3.2.1 Discrimination along Single Directions

Overall participants were able to discriminate between intermediate and extreme manipulations along a single morph direction (see figure 7). Discrimination information is on average lower in CPs compared to controls for SAct manipulations ($\hat{\beta}_{CP,0} = -0.19$, $\mathrm{HPDI}_{95\%} = [-0.32, -0.04]$) but not for SConn manipulations ($\hat{\beta}_{CP,0} = -0.06$, $\mathrm{HPDI}_{95\%} = [-0.15, 0.11]$). There's no discernable difference in the influence of morphing strength between the two groups.

3.2.2 Discrimination between Different Directions

In general, the weaker the SConn manipulation and the stronger the SAct manipulation, the more likely participants chose the SConn test image as being more similar to the

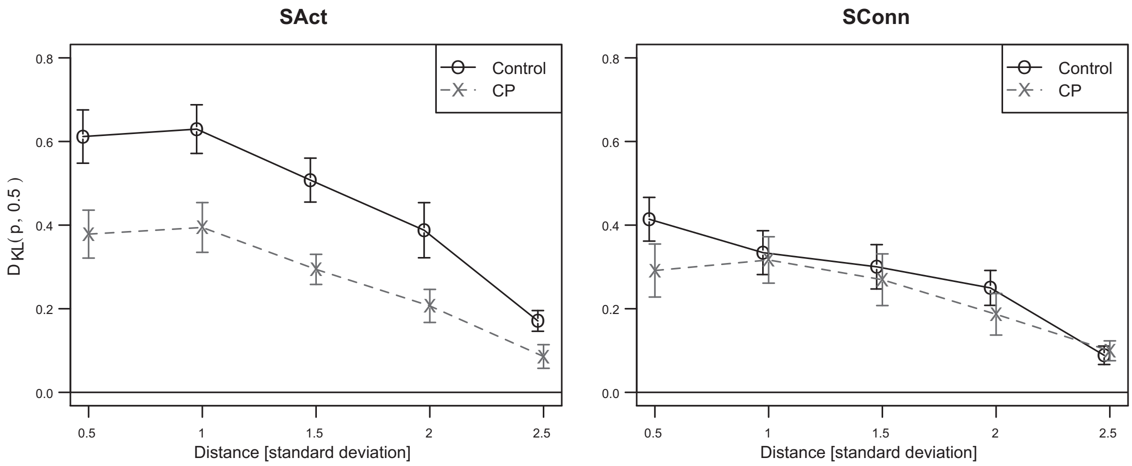

Figure 7. Discrimination information (KL divergence) for manipulations in SAct (left) and SConn (right) directions. Compared to controls (black, solid) CPs (light gray, dashed) are less able to discriminate manipulations according to SAct but there's no difference in manipulations according to SConn. (Error bars = SEM).

average. Compared to controls, CPs showed a main effect of a small but nonsignificant overall response bias in regarding SAct manipulations as less distinctive, i.e. more similar to the average, ($\hat{\beta}_{CP,0} = 0.05$, $\mathrm{HPDI}_{95\%} = [-0.03, 0.13]$).

With respect to the influence of morphing strength, controls had a small bias in regarding manipulations in SAct directions as more distinct than SConn manipulations even if they are matched in morphing strength ($\hat{\beta}_{SAct} = 0.09$, $\hat{\beta}_{SConn} = -0.06$, in standardized units of σ). Compared to controls, CPs showed less sensitivity for SAct manipulations ($\hat{\beta}_{CP,SAct} = -0.05$, $\mathrm{HPDI}_{95\%} = [-0.07, -0.02]$, standardized units) but the same sensitivity for SConn manipulations ($\hat{\beta}_{CP,SAct} = 0.01$, $\mathrm{HPDI}_{95\%} = [-0.02, 0.02]$, standardized units).

For both groups the functional relationship between the two different morphing strengths (σ_{SAct} and σ_{SConn}) and the discrimination information resembled the shape of a valley. The region with a low discrimination information, i.e. test image pairings for which participants responded indifferently, is close by and parallel to the diagonal. And as the difference in morphing strength between the two test images increases, so does participants discrimination information. Compared to controls, CPs showed a decreased sensitivity to

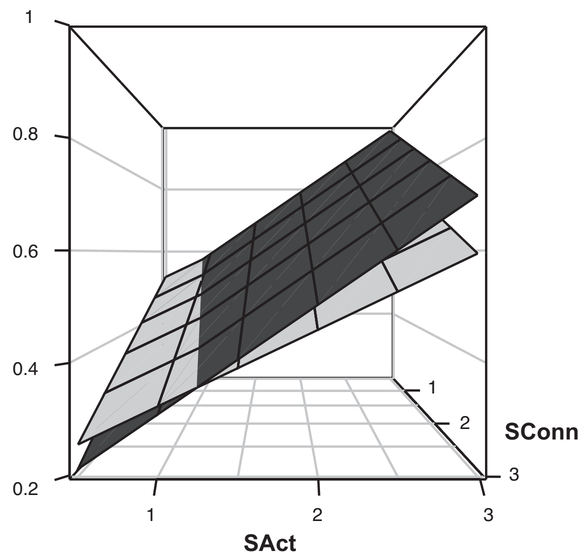

Figure 8. Fraction of responses rating the SConn test image as more similar to the average face than the SAct manipulation, as a function of morphing strength for controls (black) and CPs (light gray). Compared to controls, CPs have a similar sensitivity to SConn manipulations but are less sensitive to SAct manipulations (plane less tilted in the SAct direction).

differences in morphing strength ($\hat{\beta}_{CP,\Delta} = -0.005$, $\mathrm{HPDI}_{95\%} = [-0.008, -0.001]$, in units of σ^2). Summarily, while controls covered a large spectrum of responses, from close to uniform up to reliable preferences for one direction over another, CPs tended to make less sharp, more noisy discriminations.

4 Discussion

4.1 Summary

Two models for the encoding of facial information were proposed. The first model, SAct, was derived as a model of functional encoding maximizing the information about facial identity encoded. The second model, SConn, was proposed as a model of dysfunctional encoding in congenital prosopagnosia where facial encoding is constrained by a predisposition on reduced network connectivity.

Comparing the results on model representations, obtained in the first part of the study, with psychophysical findings on human face processing, the SAct encoding can be char-

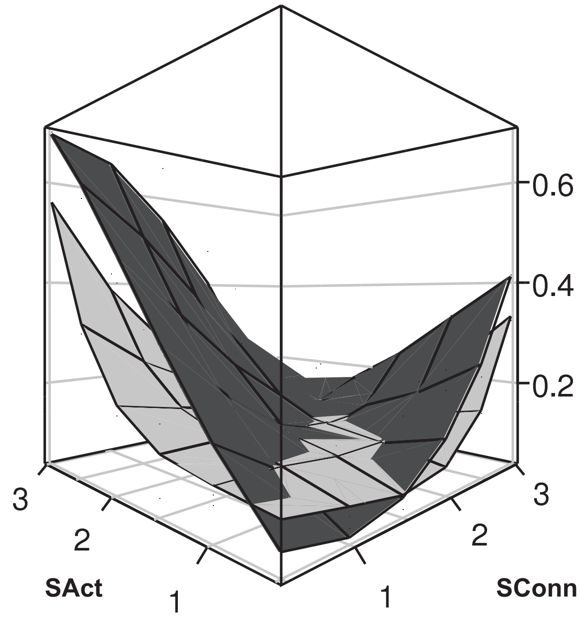

Figure 9. Discrimination information as a function of morphing strength in SConn or SAct directions for controls (black) and CPs (light gray). Overall CPs make less sharp and more noisy discriminations (less curved response plane). Also, the region of indifferent responses is more stretched out around the diagonal with equal morphing strength in both directions.

acterized as “holistic” processing. Pixel, i.e. input, activations are integrated across the whole face and encoded in sparse output unit activations. In contrast to the unconstrained SAct network, the face representations of the SConn model replaces this “holistic” information encoding of faces with a localized, “featural” representation in terms of face parts. Such a “featural” representation aligns with current models on CP face processing (Stollhoff et al., 2010).

In the second part of the study, an experimental validation of the model was conducted with a group of 10 CP participants and 20 age-matched controls. Participants had to judge similarity between an average face and a set of test face images manipulated according to either the SAct or the SConn model. In comparison to controls, participants with CP showed clear deficits in discriminating between face representations of the SAct model, but no differences with respect to SConn representations.

Taken together, the proposed models of functional and dysfunctional processing are substantiated by behavioral data in two ways:

- The representation of faces provided by the SAct model of functional processing aligns well with known aspects of human face processing.

- In a direct behavioral test, controls and participants with CP differ in discrimination abilities only for manipulations in the SAct model, but not in the proposed SConn model of CP face encoding.

More generally, our results reveal how a structural constraint on network connectivity can lead to a featural representation of facial information that is inappropriate to identify faces. The assumption of reduced network connectivity implemented in the SConn model has a direct neuropsychological equivalent in terms of reduced synaptic connectivity in the face processing areas of the inferotemporal cortex in CP (Thomas et al., 2009). By incorporating this assumption into a model of facial encoding, the observed structural neurophysiological differences can be related to functional deficits in the computational processing of facial information observed in CP. In contrast, if the assumption of a constrained network connectivity is replaced with a functional constraint on optimal information transfer, implemented in the SAct model, this leads to a holistic representation of faces at the computational level. Such a representation is in line with psychophysical studies on normal face processing (e.g. Farah et al., 1998).

In general, the proposed model - a simple linear feedforward network operating with a very small number of units and connections - can only be a strong simplification of the actual neuronal architecture found in the human visual cortex. Yet, even though the proposed ICA models are unlikely to be actually implemented in the brain as such, they still seem to capture the critical feature of congenital prosopagnosia not only at the computational but also at the implementation level. Algorithmically, the ICA used in this study is normally not implemented in a neural network architecture. But it can be shown, that it is in fact related to training a single layer network with lateral connectivity according to a neurobiologically inspired Hebbian/anti-Hebbian learning rule to foster sparse representations (Foeldiak, 1990). This association has been elaborated in more detail by Olshausen and Field (1997), where an equivalence is established to the Infomax-

ICA by Bell and Sejnowski (1995) used in this study.

4.2 Specificity of the Deficits

If one characterizes prosopagnosia as a specific instance of an identification agnosia, there are two related concepts: First, an identification of people by other sensory modalities. Second, the visual identification of other object classes. In the first case, an analog to prosopagnosia has been found in the domain of vocal identification, called phonagnosia. Similar to prosopagnosia, phonagnosia has been observed in both forms: Acquired phonagnosia after damages to the inferior and lateral parietal regions of the right hemisphere (Van Lancker et al., 1988, 1989) as well as developmental phonagnosia without an apparent cause (Garrido et al., 2009a). Due to inherent differences between auditory and visual processing it is difficult to judge, whether the results found in this study could be generalized to phonagnosia.

In the second case, although identification tasks might appear in other areas of visual expertise than face recognition, it is hard to find examples of visual identification tasks that occur as frequent and - as an area of visual expertise - as widespread as facial identification. One possible example might be the identification of specific landmarks, e.g. famous or familiar places. As a specific form of a topographical disorientation, landmark agnosia (Aguirre and D'Esposito, 1999) has been observed after lesions of the lingual and fusiform gyrus (Takahashi and Kawamura, 2002) and often co-occurs with prosopagnosia (Landis et al., 1986). Although faces and landmarks are processed in separate cortical regions these regions are in close spatial proximity, which is remarkable given the differences between faces and landmarks both in physical appearance as well as social relevance. However, based on the results of this study it is conceivable, that these two cortical areas should share a similar neuroanatomical structure, characterized by a high degree of structural connectivity. In this sense, spatial proximity would be a by-product of structural similarity.

A direct experimental investigation of processing differences in congenital prosopagnosia in the identification of non-face images, similar to the one conducted here for face images, would be of great interest. However there are at least two difficulties in designing such an experiment. Firstly, as mentioned above it is difficult to find an area in which humans normally develop a level of expertise comparable to faces. Secondly, while a lot of progress has been made regarding computational models of face recognition, it is yet an open question how best to represent images of non-face objects. As a consequence, a major technical problem would be the alignment of images with respect to variations in scale, rotation, luminance, etc. Yet, theoretical studies contrasting face recognition with general visual expertise, e.g. in dog experts or ornithologists, have argued for similarities at a computational level (Palmeri et al., 2004; Tong et al., 2008). Also, neuropsychological studies of non-face expert processing show a recruitment of classical face processing regions (Gauthier et al., 2000a; Bukach et al., 2006).

4.3 Structural Encoding in Hierarchical Networks

Appearance-based models in their simplest form operate directly on the pixel intensity values in a digitized image. An example is the "Eigenface" approach (Turk and Pentland, 1991) based on a principal component analysis of the pixel intensity matrix. However, there are some limitations to the applicability of appearance-based methods in general object recognition. Most notably, they are highly susceptible even to small changes induced by transformations of the input, e.g. translation of the image by a single pixel. This is a serious problem, as the pixel-intensity variation induced by transformations across images of the same exemplar can be larger than the variation across images taken from different exemplars (Ullman, 1997). Thus, more complicated models of face recognition are not direct appearance-based methods, but include (hierarchical) combinations of feature- or fragment-based image representation (Wiskott et al., 1997; Parga and Rolls, 1998; Ullman and Sali, 2000; Ullman, 2007; Serre et al., 2007; Wallis et al., 2008).

The first layer of hierarchically structured neural networks usually involves some form of local integration of information and is modeled by a layer of topographically arranged filters with properties similar to those observed in simple and/or complex cells in the human primary visual cortex (Lee, 1996; Bell and Sejnowski, 1997; Olshausen and Field, 1997; Hoyer and Hyvarinen, 2002). The second layer then integrates this information into features, which possibly extend over a larger part of the visual field and encode more complex information. Further layers perform subsequent integration and can achieve a (position-)invariant representation of different object classes (Bartlett and Sejnowski, 1998; Parga and Rolls, 1998; Deco, 2004). A decomposition into parts (features) and their respective spatial positions (configuration) is certainly appropriate for the representation of object classes. With respect to face identification such a decomposition has several shortcomings. More specifically, it is unclear, whether the decomposition of a face image into parts or features (e.g. eyes, nose) and their respective spatial positions is appropriate in the case of an individual face where these parts always appear in exactly the same combination and in exactly the same spatial configuration. In such a task a holistic representation of faces as an individual, personal whole seems more appropriate.

Nonetheless, a certain degree of hierarchical organisation is presumably beneficial even for the holistic processing of faces. For example, a hierarchical organisation could be used to engage in selective attention to different stimulus characteristics. In experiments on facial identification, participants rely more on the information contained in vertical spatial relations, e.g. eye-height (Goffaux, 2008). This is precisely what would be expected for an optimal observer if stimulus transformations more often disturb horizontal information but leave the vertical information intact (Stollhoff, 2010). Such a preference towards vertical information can also explain the findings that, if face stimuli are manipulated by in-depth rotations, participants show a better performance forrotations around the vertical axis (look left or right) as compared to the horizontal axis (look up or down; Favelle et al., 2007). Thus, a further development of the model introduced here, could include a preliminary layer of filtering according to orientation and spatial frequency. And only the filtered image would then be processed holistically. However, given the specificity of the deficit in congenital prosopagnosia it is unclear, whether this introduction of additional model complexity could lead to a better characterization of the deficit.

4.4 Structural Constraints as a General Principle

By formulating constraints that act on the formation of face representations in neural networks, we formulated a model for the influence of congenital differences on developmental trajectories in CP.

Generalizing from our model, sparseness constraints or limitations could be implemented in a neural network at various levels. Examples include:

Sparse coding: Enforcing a sparsely distributed output activation maximizes information transmission through the network (Bell and Sejnowski, 1995) and increases storage capacity and response specificity (Field, 1994). This was the basis of the functional SAct model proposed here.

Sparse forward connectivity: Restricting the number and strengths of inputs for each output unit leads to a more featural description in terms of a parts decomposition. Such a constraint was the basis of the proposed SConn model of CP.

Sparse recurrent connectivity: The coupling strength of the recurrent connectivity influences the convergence rate as well as the capacity to recognize individual faces across environmental changes (Parga and Rolls, 1998).

Limited units: Storage capacity in neural networks increases with the number of units (e.g. linearly for Hopfield networks (Hopfield, 1982)). Limiting the number of units directly affects the number of individual faces that can be memorized.

Recent neuroanatomical studies lend support to this more general hypothesis that the face recognition deficit in CP can be explained by the influence of different sparseness constraints acting on the human system for face processing:

- In a sample of CP participants, Behrmann et al. (2007) observed a reduction in the cortical gray matter volume (i.e. number of units) in anterior inferotemporal areas which presumably are involved in memory storage. Moreover, the individual volumetric reductions were correlated with a decreased performance in a famous face test.

- Thomas et al. (2009) observed a reduction in the structural connectivity in the ventral occipito-temporal visual cortex in a sample of CP participants, both in the forward connectivity of lower visual areas to associative areas of the human face processing system as well as the recurrent connectivity inside these areas. The reduction was correlated with a decreased performance in face recognition.

- Investigating age-related differences in face processing, Thomas et al. (2008) found a robust correlation between a reduction in the integrity of fiber tracts connecting face processing areas and a decline in the ability to discriminate faces.

But the observation of structural differences in a mature state doesn't necessarily establish a cause for the deficit. It is also possible that the structural differences are merely an adaptive consequence of other, so far unobserved, differences causing the deficit.

4.5 Deficits at Higher Levels of a Hierarchical Network

In this study, a detailed model for CP was formulated for deficits in encoding frontal views of faces. As a first extension of the model, a deficit in the generalization across different views could be formulated based on the idea of identifying appropriate constraints on the structural connectivity.

Associating images taken from different viewpoints with the same exemplar can be accomplished in an attractor network. For example, Parga and Rolls (1998) studied a neurobiologically plausible recurrent network and showed that the capacity for view-invariant recognition depends on the interplay of the number of exemplars to be stored, the number of units available and the coupling strength of the recurrent connectivity. These observations motivate a generalization of the model for the encoding of facial information in congenital prosopagnosia proposed above. First, reduced recurrent connectivity in associative networks of the anterior inferotemporal cortex could lead to deficits in associating view-dependent holistic representations across changes in viewpoint. Second, a reduction in the number of units in the associative network (cf. a decreased cortical volume) could limit holistic processing of faces to a small number of individuals and/or views.

References

Aguirre, G. and D'Esposito, M. (1999). Topographical disorientation: a synthesis and taxonomy. *Brain*, 122:1613–1628.

Avidan, G. and Behrmann, M. (2009). Functional MRI Reveals Compromised Neural Integrity of the Face Processing Network in Congenital Prosopagnosia. *Curr Biol*, 19(13).

Avidan, G., Hasson, U., Malach, R., and Behrmann, M. (2005). Detailed exploration of face-related processing in congenital prosopagnosia: 2. Functional neuroimaging findings. *J Cognitive Neurosci*, 17(7):1150–1167.

Bartlett, M. S. (2007). Information maximization in face processing. *Neurocomputing*, 70(13-15):2204–2217.

Bartlett, M. S., Movellan, J. R., and Sejnowski, T. J. (2002). Face recognition by independent component analysis. *IEEE T Neural Networ*, 13(6):1450–64.

Bartlett, M. S. and Sejnowski, T. J. (1998). Learning viewpoint-invariant face representations from visual experience in an attractor network. *Network-Computation In Neural Systems*, 9(3):399–417.

Barton, J. J. S. and Cherkasova, M. (2005). Impaired spatial coding within objects but not between objects in prosopagnosia. *Neurology*, 65(2):270–274.

Barton, J. J. S., Press, D. Z., Keenan, J. P., and O'Connor, M. (2002). Lesions of the fusiform face area impair perception of facial configuration in prosopagnosia. *Neurology*, 58(1):71–8.

Barton, J. J. S., Zhao, J., and Keenan, J. P. (2003). Perception of global facial geometry in the inversion effect and prosopagnosia. *Neuropsychologia*, 41(12):1703–11.

Bates, D. and Maechler, M. (2009). *lme4: Linear mixed-effects models using S4 classes.* R package version 0.999375-31.

Behrmann, M. and Avidan, G. (2005). Congenital prosopagnosia: face-blind from birth. *Trends in Cognitive Sciences*, 9(4):180–7.

Behrmann, M., Avidan, G., Gao, F., and Black, S. (2007). Structural imaging reveals anatomical alterations in inferotemporal cortex in congenital prosopagnosia. *Cereb Cortex*, 17(10):2354–63.

Behrmann, M., Avidan, G., Marotta, J., and Kimchi, R. (2005). Detailed exploration of face-related processing in congenital prosopagnosia: 1. Behavioral findings. *J Cognitive Neurosci*, 17(7):1130–1149.

Bell, A. and Sejnowski, T. J. (1995). An Information Maximization Approach to Blind Separation and Blind Deconvolution. *Neural Computation*, 7(6):1129–1159.

Bell, A. and Sejnowski, T. J. (1997). The "independent components" of natural scenes are edge filters. *Vision Res*, 37(23):3327–3338.

Blanz, V. and Vetter, T. (1999). A Morphable Model For The Synthesis Of 3D Faces. *SIGGRAPH'99 Conference Proceedings*, pages 187–194.

Bodamer, J. (1947). Die Prosop-Agnosie. *Arch Psychiatr Nervenkr*, 179:6–53.

Breen, N., Caine, D., and Coltheart, M. (2000). Models of face recognition and delusional misidentification: A critical review. *Cognitive Neuropsychology*, 17(1-3):55–71.

Bukach, C. M., Bub, D. N., Gauthier, I., and Tarr, M. J. (2006). Perceptual expertise effects are not all or none: spatially limited perceptual expertise for faces in a case of prosopagnosia. *J Cognitive Neurosci*, 18(1):48–63.

Burton, A. M., Bruce, V., and Hancock, P. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, 23(1):1–31.

Damasio, A., Tranel, D., and Damasio, H. (1990). Face Agnosia and the Neural Substrates of Memory. *Annu Rev Neurosci*, 13:89–109.

De Haan, E. (1999). A familial factor in the development of face recognition deficits. *J Clin Exp Neuropsyc*, 21(3):312–315.

de Haan, M., Humphreys, K., and Johnson, M. H. (2002). Developing a brain specialized for face perception: A converging methods approach. *Dev. Psychobiol.*, 40(3):200–212.

De Renzi, E., Faglioni, P., Grossi, D., and Nichelli, P. (1991). Apperceptive and associative forms of prosopagnosia. *Cortex*, 27(2):213–21.

De Renzi, E., Perani, D., Carlesimo, G., Silveri, M., and Fazio, F. (1994). Prosopagnosia can be Associated with Damage Confined to the Right-Hemisphere - An MRI and PET Study and a Review of the Literature. *Neuropsychologia*, 32(8):893–902.

Deco, G. (2004). A Neurodynamical cortical model of visual attention and invariant object recognition. *Vision Res*, 44(6):621–642.

Dinkelacker, V., Grueter, M., Klaver, P., and Grueter, T. (2010). Congenital Prosopagnosia: Multistage Anatomical and Functional Deficits in Face Processing Circuitry. *Journal of Neurology*, (accepted with minor revisions).

Draper, B., Baek, K., Bartlett, M. S., and Beveridge, J. (2003). Recognizing faces with PCA and ICA. *Computer Vision and Image Understanding*, 91(1-2):115–137.

Duchaine, B. C. (2006). Prosopagnosia as an impairment to face-specific mechanisms: Elimination of the alternative hypotheses in a developmental case. *Cognitive Neuropsychology*, 23(5):714–747.

Duchaine, B. C. and Nakayama, K. (2005). Dissociations of Face and Object Recognition in Developmental Prosopagnosia. *J Cognitive Neurosci*, 17(2):249–261.

Duchaine, B. C. and Nakayama, K. (2006). Developmental prosopagnosia: a window to content-specific face processing. *Current Opinion in Neurobiology*, 16(2):166–73.

Ellis, H. and Florence, M. (1990). Bodamer's (1947) paper on prosopagnosia. *Cognitive Neuropsychology*, 7(2):81–105.

Ellis, H. D. and Lewis, M. (2001). Capgras delusion: a window on face recognition. *Trends in Cognitive Sciences*, 5(4):149–156.

Elze, T. (2009). FlashDot - A platform independent experiment generator for visual psychophysics. *JOV*, 9(14):58–58.

Farah, M. J., O'Reilly, R. C., and Vecera, S. P. (1993). Dissociated overt and covert recognition as an emergent property of a lesioned neural network. *Psychol Rev*, 100(4):571–88.

Farah, M. J., Wilson, K. D., Drain, M., and Tanaka, J. N. (1998). What is "special" about face perception? *Psychol Rev*, 105(3):482–98.

Faraway, J. J. (2006). *Extending the linear model with R: generalized linear, mixed effects and nonparametric regression models*. Chapman & Hall, FL.

Farkas, L. G. (1981). *Anthropometry of the Head and Face in Medicine*. Elsevier, New York.

Favelle, S. K., Palmisano, S., and Maloney, R. T. (2007). Things are looking up: Differential decline in face recognition following pitch and yaw rotation. *Perception*, 36(9):1334–1352.

Field, D. (1994). What is the Goal of Sensory Coding. *Neural Computation*, 6(4):559–601.

Foeldiak, P. (1990). Forming Sparse Representations by Local Anti-Hebbian Learning. *Biol Cybern*, 64(2):165–170.

Fox, C. J., Iaria, G., and Barton, J. J. S. (2008). Disconnection in prosopagnosia and face processing. *Cortex*, 44(8):996–1009.

Garrido, L., Duchaine, B. C., and Nakayama, K. (2008). Face detection in normal and prosopagnosic individuals. *J Neuropsychol*, 2:119–140.

Garrido, L., Eisner, F., McGettigan, C., Stewart, L., Sauter, D., Hanley, J. R., Schweinberger, S. R., Warren, J. D., and Duchaine, B. C. (2009a). Developmental phonagnosia: A selective deficit of vocal identity recognition. *Neuropsychologia*, 47(1):123–131.

Garrido, L., Furl, N., Draganski, B., Weiskopf, N., Stevens, J., Tan, G. C.-Y., Driver, J., Dolan, R. J., and Duchaine, B. C. (2009b). Voxel-based morphometry reveals reduced grey matter volume in the temporal cortex of developmental prosopagnosics. *Brain*, 132:3443–3455.

Gauthier, I., Behrmann, M., and Tarr, M. J. (2004). Are Greebles like faces? Using the neuropsychological exception to test the rule. *Neuropsychologia*, 42(14):1961–1970.

Gauthier, I., Curby, K. M., Skudlarski, P., and Epstein, R. A. (2005). Individual differences in FFA activity suggest independent processing at different spatial scales. *Cognitive, affective & behavioral neuroscience*, 5(2):222–34.

Gauthier, I., Skudlarski, P., Gore, J. C., and Anderson, A. W. (2000a). Expertise for cars and birds recruits brain areas involved in face recognition. *Nat Neurosci*, 3(2):191.

Gauthier, I., Tarr, M. J., Moylan, J., Skudlarski, P., Gore, J. C., and Anderson, A. W. (2000b). The fusiform "face area" is part of a network that processes faces at the individual level. *J Cognitive Neurosci*, 12(3):495–504.

Gauthier, I., Tarr, M. J., Moylan, J., Skudlarski, P., Gore, J. C., and Anderson, A. W. (2000c). The fusiform "face area" is part of a network that processes faces at the individual level. *J Cognitive Neurosci*, 12(3):495–504.

Goffaux, V. (2008). The various phases of faces. *Perception*, 37:35–35.

Grill-Spector, K., Knouf, N., and Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nat Neurosci*, 7(5):555–562.

Grueter, M., Grueter, T., Bell, V., Horst, J., Laskowski, W., Sperling, K., Halligan, P. W., Ellis, H. D., and Kennerknecht, I. (2007). Hereditary Prosopagnosia: the First Case Series. *Cortex*, 43(6):734–749.

Hadfield, J. (2009). *MCMC methods for Multi-response Generalised Linear Mixed Models: The MCMCglmm R Package*. R package version 1.13.

Hasson, U., Avidan, G., Deouell, L. Y., Bentin, S., and Malach, R. (2003). Face-selective activation in a congenital prosopagnosic subject. *J Cognitive Neurosci*, 15(3):419–31.

Haxby, J. V., Hoffman, E., and Gobbini, M. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6):223–233.

Hoffman, E. and Haxby, J. V. (2000). Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nat Neurosci*, 3(1):80–84.

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *PNAS*, 79(8):2554–2558.

Hoyer, P. O. and Hyvarinen, A. (2002). A multi-layer sparse coding network learns contour coding from natural images. *Vision Res*, 42(12):1593–1605.

Humphreys, K., Avidan, G., and Behrmann, M. (2007). A detailed investigation of facial expression processing in congenital prosopagnosia as compared to acquired prosopagnosia. *Experimental Brain Research*, 176(2):356–373.

Hyvarinen, A., Karhunen, J., and Oja, E. (2001). *Independent component analysis*. Wiley, NY.

Kanwisher, N., McDermott, J., and Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci*, 17(11):4302–11.

Karabatsos, G. (2006). Bayesian nonparametric model selection and model testing. *J Mathematical Psychology*, 50(2):123–148.

Kawashima, R., Sato, N., Nakamura, A., Sugiura, M., Kato, T., Hatano, K., Ito, K., Fukuda, H., Schormann, T., and Zilles, K. (2000). Functional delineation of the human occipito-temporal areas related to face and scene processing. A PET study. *Brain*, 123(9):1903–12.

Kennerknecht, I., Grueter, T., Welling, B., Wentzek, S., Horst, J., Edwards, S., and Grueter, M. (2006). First report of prevalence of non-syndromic hereditary prosopagnosia (HPA). *Am. J. Med. Genet.*, 140A(15):1617–1622.

Kennerknecht, I., Ho, N. Y., and Wong, V. C. N. (2008a). Prevalence of hereditary prosopagnosia (HPA) in Hong Kong Chinese population. *Am. J. Med. Genet.*, 146A(22):2863–70.

Kennerknecht, I., Pluempe, N., and Welling, B. (2008b). Congenital prosopagnosia - a common hereditary cognitive dysfunction in humans. *Front Biosci*, 13:3150–3158.

Kennerknecht, I., Plümpe, N., Edwards, S., and Raman, R. (2007). Hereditary prosopagnosia (HPA): the first report outside the Caucasian population. *J Hum Genet*, 52(3):230–6.

Kress, T. and Daum, I. (2003). Developmental prosopagnosia: a review. *Behavioural neurology*, 14(3-4):109–21.

Landis, T., Cummings, J. L., Benson, D. F., and Palmer, E. P. (1986). Loss of topographic familiarity. An environmental agnosia. *Arch Neurol*, 43(2):132–6.

Le Grand, R., Cooper, P. A., Mondloch, C. J., Lewis, T. L., Sagiv, N., de Gelder, B., and Maurer, D. (2006). What aspects of face processing are impaired in developmental prosopagnosia? *Brain and Cognition*, 61(2):139–58.

Lee, T. S. (1996). Image representation using 2D gabor wavelets. *IEEE*, 18(10):959–971.

Mazzucchi, A. and Biber, C. (1983). Is prosopagnosia more frequent in males than in females? *Cortex*, 19(4):509–16.

McConachie, H. R. (1976). Developmental prosopagnosia. A single case report. *Cortex*, 12(1):76–82.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273.

Olshausen, B. and Field, D. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Res*, 37(23):3311–3325.

Palmeri, T., Wong, A., and Gauthier, I. (2004). Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences*, 8(8):378–386.

Parga, N. and Rolls, E. (1998). Transform-invariant recognition by association in a recurrent network. *Neural Computation*, 10(6):1507–1525.

Pessa, E., Bandinelli, P., and Penna, M. (1999). Simulating prosopagnosia through a lesion of lateral connections in a feed-forward neural network. *Ital J Neurol Sci*, 20(1):29–36.

Polk, T. A., Park, J., Smith, M. R., and Park, D. C. (2007). Nature versus nurture in ventral visual cortex: A functional magnetic resonance Imaging study of twins. *J Neurosci*, 27(51):13921–13925.

Pourtois, G., Schwartz, S., Seghier, M., Lazeyras, F., and Vuilleumier, P. (2005). Portraits or people? Distinct representations of face identity in the human visual cortex. *J Cognitive Neurosci*, 17(7):1043–1057.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107.

R Development Core Team (2009). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Riddoch, J. M. and Humphreys, G. W. (1993). *BORB: Birmingham Object Recognition Battery*. Psychology Press, NY.

Sanger, T. (1989). Optimal Unsupervised Learning in a Single-layer linear Feedforward Neural Network. *Neural Networks*, 2(6):459–473.

Scherf, K. S., Behrmann, M., Humphreys, K., and Luna, B. (2007). Visual category-selectivity for faces, places and objects emerges along different developmental trajectories. *Developmental Science*, 10(4):F15–30.

Schmalzl, L., Palermo, R., and Coltheart, M. (2008a). Cognitive heterogeneity in genetically based prosopagnosia: A family study. *J Neuropsychol*, 2:99–117.

Schmalzl, L., Palermo, R., Green, M., Brunsdon, R., and Coltheart, M. (2008b). Training of familiar face recognition and visual scan paths for faces in a child with congenital prosopagnosia. *Cognitive Neuropsychology*, 25(5):704–729.

Schwarzer, G., Huber, S., Grueter, M., Grueter, T., Gross, C., Hipfel, M., and Kennerknecht, I. (2007). Gaze behaviour in hereditary prosopagnosia. *Psychological Research*, 71(5):583–90.

Serre, T., Oliva, A., and Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *PNAS*, 104(15):6424.

Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, FL.

Stollhoff, R. (2010). Representing facial information using Gabor wavelets. *Perception*, 39(ECVP Abstract Supplement):127.

Stollhoff, R., Jost, J., Elze, T., and Kennerknecht, I. (2010). The Early Time Course of Compensatory Face Processing in Congenital Prosopagnosia. *PLoS ONE*, 5(7).

Stollhoff, R., Jost, J., Elze, T., and Kennerknecht, I. (2011). Deficits in long-term recognition memory reveal dissociated subtypes in congenital prosopagnosia. *PLoS ONE*, 1(6).

Stone, J., Porrill, J., Porter, N., and Wilkinson, I. (2002). Spatiotemporal independent component analysis of event-related fMRI data using skewed probability density functions. *NeuroImage*, 15(2):407–421.

Takahashi, N. and Kawamura, M. (2002). Pure topographical disorientation - The anatomical basis of landmark agnosia. *Cortex*, 38(5):717–725.

Thomas, C., Avidan, G., Humphreys, K., Jung, K. J., Gao, F., and Behrmann, M. (2009). Reduced structural connectivity in ventral visual cortex in congenital prosopagnosia. *Nat Neurosci*, 12(1):29–31.

Thomas, C., Moya, L., Avidan, G., Humphreys, K., Jung, K. J., Peterson, M. A., and Behrmann, M. (2008). Reduction in white matter connectivity, revealed by diffusion tensor imaging, may account for age-related changes in face perception. *J Cognitive Neurosci*, 20(2):268–284.

Tong, M., Joyce, C., and Cottrell, G. W. (2008). Why is the fusiform face area recruited for novel categories of expertise? A neurocomputational investigation. *Brain Research*, 1202:14–24.

Troje, N. F. and Bülthoff, H. H. (1996). Face recognition under varying poses: the role of texture and shape. *Vision Res*, 36(12):1761–71.

Tuerlinckx, F., Rijmen, F., Verbeke, G., and De Boeck, P. (2006). Statistical inference in generalized linear mixed models: a review. *The British journal of mathematical and statistical psychology*, 59(2):225–55.

Turk, M. and Pentland, A. (1991). Eigenfaces for recognition. *J Cognitive Neurosci*, 3(1):71–86.

Ullman, S. (1997). *High-level Vision: Object Recognition and Visual Cognition.* MIT Press, MA.

Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences*, 11(2):58–64.

Ullman, S. and Sali, E. (2000). Object classification using a fragment-based representation. *Biologically Motivated Computer Vision*, 1811:73–87.

Van Lancker, D. R., Cummings, J. L., Kreiman, J., and Dobkin, B. H. (1988). Phonagnosia: a dissociation between familiar and unfamiliar voices. *Cortex*, 24(2):195–209.

Van Lancker, D. R., Kreiman, J., and Cummings, J. L. (1989). Voice perception deficits: neuroanatomical correlates of phonagnosia. *J Clin Exp Neuropsyc*, 11(5):665–74.

Virasoro, M. (1988). The Effect of Synapses Destruction on Categorization by Neural Networks. *Europhys Lett*, 7(4):293–298.

Virasoro, M. (1989). Categorization in Neural Networks and Prosopagnosia. *Phys Rep*, 184(2-4):301–306.

Wallis, G., Siebeck, U. E., Swann, K., Blanz, V., and Bülthoff, H. H. (2008). The prototype effect revisited: Evidence for an abstract feature model of face recognition. *JOV*, 8(3):20.

Warrington, E. K. and James, M. (1991). *The Visual Object and Space Perception Battery.* Harcourt Assessment, London.

Wiskott, L., Fellous, J., Kruger, N., and von der Malsburg, C. (1997). Face recognition by elastic bunch graph matching. *IEEE*, 19(7):775–779.

Zifan, A., Gharibzadeh, S., and Moradi, M. H. (2007). Could dynamic attractors explain associative prosopagnosia? *Medical Hypotheses*, 68(6):1399–1405.