

# Contrast sensitivity functions in autoencoders

**Qiang Li**

Image Processing Lab, Parc Científic, Universitat de València, Spain



**Alex Gomez-Villa**

Computer Vision Center, Universitat Autònoma de Barcelona, Spain



**Marcelo Bertalmío**

Instituto de Óptica, Spanish National Research Council (CSIC), Spain



**Jesús Malo**

Image Processing Lab, Parc Científic, Universitat de València, Spain



Three decades ago, Atick et al. suggested that human frequency sensitivity may emerge from the enhancement required for a more efficient analysis of retinal images. Here we reassess the relevance of low-level vision tasks in the explanation of the contrast sensitivity functions (CSFs) in light of 1) the current trend of using artificial neural networks for studying vision, and 2) the current knowledge of retinal image representations. As a first contribution, we show that a very popular type of convolutional neural networks (CNNs), called autoencoders, may develop human-like CSFs in the spatiotemporal and chromatic dimensions when trained to perform some basic low-level vision tasks (like retinal noise and optical blur removal), but not others (like chromatic adaptation or pure reconstruction after simple bottlenecks). As an illustrative example, the best CNN (in the considered set of simple architectures for enhancement of the retinal signal) reproduces the CSFs with a root mean square error of 11% of the maximum sensitivity. As a second contribution, we provide experimental evidence of the fact that, for some functional goals (at low abstraction level), deeper CNNs that are better in reaching the quantitative goal are actually worse in replicating human-like phenomena (such as the CSFs). This low-level result (for the explored networks) is not necessarily in contradiction with other works that report advantages of deeper nets in modeling higher level vision goals. However, in line with a growing body of literature, our results suggests another word of caution about CNNs in vision science because the use of simplified units or unrealistic architectures in goal optimization may be a limitation for the modeling and understanding of human vision.

## Introduction

The human contrast sensitivity function (CSF) characterizes the psychophysical response to visual gratings of different frequency (Campbell & Robson, 1968). Filter characterizations in the Fourier domain are complete only for linear, shift-invariant systems. Human vision certainly is more complicated than that, however, this simple measure of the bandwidth of the system is still of paramount significance in biological vision: the CSF filter is an image-computable model that roughly describes the kind of visual information that is available for humans (Watson & Ahumada, 2016). Moreover, although it is defined for threshold conditions, there are many examples that illustrate the relevance of the CSF in more general situations (Watson et al., 1986; Watson & Malo, 2002; Watson & Ahumada, 2005), so it has shaped image engineering over decades (Mannos & Sakrison, 1974; Hunt, 1975; Wallace, 1992; Taubman & Marcellin, 2001). This theoretical and practical relevance motivated the measurement of CSFs, not only for spatial gratings (Campbell & Robson, 1968), but also for moving gratings (Kelly, 1979), chromatic gratings (Mullen, 1985), spatiotemporal chromatic gratings (Díez-Ajenjo et al., 2011), at different luminance levels (Wuerger et al., 2020), and for alternative basis of the image space (Malo et al., 1997).

## Principled explanations of the human CSFs

Of course, the psychophysical CSFs have physiological roots in the spatiotemporal bandwidths of the center-surround cells tuned to achromatic and chromatic stimuli (Enroth-Cugell & Robson, 1966;

Citation: Li, Q., Gomez-Villa, A., Bertalmío, M., & Malo, J. (2022). Contrast sensitivity functions in autoencoders. *Journal of Vision*, 22(6):8, 1–45, <https://doi.org/10.1167/jov.22.6.8>.



de Valois & Pease, 1971; Ingling & Martinez-Uriegas, 1983; Martinez-Uriegas, 1994; Cai et al., 1997; Reid & Shapley, 1992, 2002). However, the physiological basis of psychophysical phenomena does not explain the functional role (or *goal*) of the underlying computation (Marr & Poggio, 1976; Marr, 1982). The discussion about the goal of certain mechanism relies on deriving the biological behavior from a computational principle. In the specific case of the CSFs, the classical work of Atick et al. (Atick et al., 1992; Atick & Redlich, 1992; Atick, 2011) derived the spatiochromatic CSFs from the maximization of the information transferred from the input to the response of the system that, under certain conditions, is equivalent to optimal deblurring and denoising of the retinal signals. These classical explanations were based on clever observations about the 2nd order properties of natural images, but relied on linear filtering models. As a result, the consideration of more flexible (nonlinear) models could lead to a better fulfillment of the computational goal and, eventually to better explanations of the CSFs. A step forward in a more general (nonlinear) derivation of these phenomena from low-level principles was given by Karklin and Simoncelli (2011), where they obtained sensors with center-surround receptive fields optimizing the information transferred by a linear+nonlinear layer of neurons with noisy inputs. However, this work did not consider the chromatic or the temporal dimensions of the problem, and no explicit comparison with the psychophysical CSFs was done. Similarly, (Lindsey et al., 2019) also reproduced center-surround sensors close to the retina when training anatomically constrained artificial neural nets (in this case, training for a higher level task, such as object recognition). Again, these center-surround cells eventually would induce CSFs, but this was not analyzed in that paper.

## Emergence of CSFs in artificial neural networks

Automatic differentiation (Baydin et al., 2018) has simplified the search of computational principles in vision science because it allows the optimization of complex models according to different goals without the burden of obtaining the analytical derivatives of the goals w.r.t. the model parameters. Automatic differentiation is at the core of the current explosion of *deep learning* (Goodfellow et al., 2016). A full analytical description of the derivatives of realistic nonlinearities in visual neuroscience is certainly possible (Martinez et al., 2018), but the widespread availability of deep-learning tools for simplified neurons makes the exploration of these artificial architectures much easier. Conventional convolutional neural networks (CNNs) are too simplistic from the neuroscience perspective,<sup>1</sup> but the freedom to combine multiple of such simplified layers in any possible way may compensate this shortcoming. In the end, one has a flexible system that

can be optimized with automatic differentiation to fulfill whatever computational goal under consideration. As a result, deep learning models are becoming standard in visual neuroscience (Kriegeskorte, 2015; Yamins & DiCarlo, 2016; Cadena et al., 2019).

According to the above, the study of the CSF of artificial neural networks is interesting for two reasons: 1) CNNs are flexible and easily optimizable tools that may allow us to investigate principled explanations of the human CSFs with more generality than the classical methods considered, and 2) given the widespread use of CNNs in computer vision and their recent use in visual neuroscience, the eventual emergence of human-like sensitivities in these artificial systems has intrinsic interest.

Very recently, two groups have reported complementary results on the emergence of CSFs in deep networks: first, in order to explain the human-like nature of some of the brightness and color illusions in CNNs trained for *low-level* visual tasks found in Gomez-Villa et al. (2019), a novel eigenanalysis of the networks was proposed (Gomez-Villa et al., 2020). This analysis revealed the emergence of human-like chromatic channels and achromatic and chromatic CSFs in these channels. Then Akbarinia et al. (2021) have found that networks trained for *high-level* visual tasks, such as classification, also may develop an achromatic CSF, in this case not explicitly imposing low-level constraints.

## Contributions and scope of this work

- First, following Atick et al. (1992), Atick and Redlich (1992), Atick (2011), Karklin and Simoncelli (2011), we reconsider principled explanations of the CSFs from *low-level* visual tasks in light of new available methods: i) the current tools from deep-learning, and ii) the current knowledge of retinal image representations. We check the emergence of spatiotemporal chromatic CSFs in a wider range of low-level (goal/architecture) situations with more realistic inputs.

Regarding the retinal input, we use recent models of the human modulation transfer function (MTF) (Watson, 2013), and recent calibrated estimations of the noise in the cones (Esteve et al., 2020) obtained via the retina models implemented in ISETbio (Cottaris et al., 2019, 2020). In this way, here we generate realistic spatiotemporal noisy inputs to the visual pathway in a plausible representation: the cones of (Stockman & Sharpe, 2000) tuned to long, medium, and short wavelengths (LMS cones).

Regarding the deep learning tools, we use spatiotemporal extensions of the *convolutional autoencoders* used in our analysis of color illusions in CNNs (Gomez-Villa et al., 2019, 2020). We

elaborate on the proper determination of the CSF for convolutional autoencoders: instead of the linear characterization of the autoencoder used in Gomez-Villa et al. (2020), which hides its nonlinear nature into a single matrix, here we stimulate the networks with gratings of different contrast. In this way, the changes of the attenuation functions describe the nonlinearities of the system.

Regarding the architectures, in this work we focus on autoencoders that reconstruct the signal in the input domain as opposed to the consideration of more general architectures that encode the images into more abstract representations to achieve higher level tasks, such as classification. This limitation in scope is reasonable if one wants to model early vision stages like the lateral geniculate nucleus (LGN), which do not imply change of domain and may function according to error minimization and signal enhancement principles (Martinez-Otero et al., 2014). If the CSFs are related to the response of LGN neurons, as is usually assumed, autoencoders seem a reasonable computational framework to use.

Regarding the tasks, in this low-level context with autoencoder tools, we consider different visual tasks which may be implemented as early as in the retina-LGN path: a) the enhancement of the retinal signal (related to information maximization) when the input is subject to different degrees of degradation owing to different pupil diameters or different plausible levels of retinal noise, b) the compensation of changes in the spectral illumination of the scene in a reasonable range of color temperature, and c) the reconstruction of the signal when some information may be lost in eventual bottlenecks.

- Second, here we provide experimental evidence of “deeper CNNs are not necessarily better” (in representing this abstraction level). The bigger generality of flexible CNN models over fixed linear models is obvious, but one may ask: do more flexible architectures necessarily lead to more human CSFs? or does better accuracy in the goal imply more human CSFs and masking behavior? Consistently with previous results in low-level tasks (Flachot et al., 2020; Gomez-Villa et al., 2020), our CSF results presented also seem to favor shallow networks (in the explored range of architectures).

Our findings at this low level of abstraction complement other results where deeper architectures actually imply closer resemblance to human behavior (Yamins et al., 2014; Cichy et al., 2016; Cadena et al., 2019; Lindsey et al., 2019). But this is not contradictory because they refer to different abstraction levels (high-level object recognition vs. our low-level color constancy and error minimization goals).

The structure of this article is as follows. We used estimating contrast sensitivity in autoencoders that extends the theory proposed in Gomez-Villa et al. (2020) to obtain the CSFs of autoencoders with an analysis of the energy (or standard deviation) of the input and the output gratings. Experiments describe the considered low-level visual tasks (compensation of biodistortions, chromatic adaptation, and signal reconstruction after bottlenecks) and the setting of the numerical experiments. Results show the main empirical findings of the work: the emergence of human-like CSFs in the spatiotemporal and chromatic dimensions in shallow CNN autoencoders trained to minimize the distortion introduced by the optics of the eye and the noise in the cones. Finally, we discuss the implications of the empirical results: on the one hand, statements about the goal or organization principles are difficult to separate from the implementation because the final behavior very much depends on the algorithmic level (or selected architecture). On the other hand, special care has to be taken in using deep models in low-level vision science: their ability for function approximation may make them excel in the performance of a sensible score, but without the appropriate architecture constraints, this does not guarantee the similarity with humans. Appendix A provides details of the implementation of the models. Appendix B describes the image/video datasets to train the models and the sinusoids used to probe the networks. Appendix C illustrates the proper training and convergence of all the considered CNNs in all experiments. It shows the learning curves and explicit examples of the responses (reconstructed signals in test) for all the considered goal/architecture scenarios.

## Methods: Estimating contrast sensitivity in autoencoders

Here we consider different linear characterizations of the autoencoders including the eigenanalysis proposed in Gomez-Villa et al. (2020). That theory is extended with the explicit consideration of the image acquisition process in the human eye, which leads us to propose a procedure to estimate the autoencoder CSF that is more connected to the definition of the CSF in human observers.

### Autoencoders

Autoencoders are artificial networks that transform the signal into an inner representation through an *encoder*, and a *decoder* transforms this inner representation back in the input domain.

$$\mathbf{x} \xrightarrow{N_{\theta}(\mathbf{x})} \mathbf{y} \quad (1)$$

In Equation 1, we do not made explicit the encoding and decoding operations, that is,  $\mathbf{x}$  and  $\mathbf{y}$  are in the image space. In this work, we will not make any assumption on the nature of the inner representation of the autoencoder. This is because the basic goal function in autoencoders (reconstruction error) is defined in the image domain, shared by input and response. Moreover, with the appropriate stimuli, the CSF characterization can be defined in this image domain.

Following Gomez-Villa et al. (2019, 2020), we focus on convolutional autoencoders. We discuss and explore different computational goals, but, for now, let's consider that the parameters  $\theta$  are trained to compensate the blur and noise introduced in the signal by the image acquisition process. In this context, given a clean image,  $\mathbf{x}_c$ , the input to the neural system would be a distorted version:  $\mathbf{x} = H \cdot \mathbf{x}_c + \mathbf{n}_r$ , where  $H$  is a blurring operator related to the optics of the eye and  $\mathbf{n}_r$  is the noise associated with the response of the LMS photodetectors at the retina. Both  $H$  and  $\mathbf{n}_r$  are unknown to the neural system. The goal of the network at this early stage is inferring  $\mathbf{x}_c$  from  $\mathbf{x}$ . Accurate models of LGN cells show that this may be one of the goals of the biological processing after retinal detection (Martinez-Otero et al., 2014).

In the supervised learning setting of artificial neural networks, the parameters  $\theta$  of the network are selected so that the average reconstruction error  $\varepsilon = \|\mathbf{x}_c - N_\theta(\mathbf{x})\|_2$  is minimized over a set of training images (Goodfellow et al., 2016). In our case, we refer to the average reconstruction error as  $\varepsilon_{LMS}$  because the input signal is expressed in the LMS cone space (Stockman & Sharpe, 2000). Of course, supervised learning and parameter updates using backpropagation may not be biologically plausible (Lillicrap et al., 2020). However, our initial aim here is looking for statistical explanations of human frequency sensitivity and hence artificial neural networks can be seen as convenient tools to optimize the selected goal. With this focus on the goal, the specific learning algorithm is not as important as ensuring that the final network actually fulfills the goal. We will see that the situation may not be that simple because networks optimizing the same goal with equivalent performance may display human or non-human CSFs, depending on their architecture.

## Filter definition of the CSF and linearized autoencoders

The CSF describes the linear response of human viewers for low-contrast sinusoids (Campbell & Robson, 1968; Kelly, 1979; Mullen, 1985). In that linear setting, the CSF describes an input-output mapping where an input sinusoid of frequency  $f$ , the basis function  $\mathbf{b}^f$ , leads to an output,  $\mathbf{y}^f$ , with attenuated contrast (or attenuated standard deviation,  $\sigma$ ). The output standard deviation is given by,  $\sigma(\mathbf{y}^f) = \text{CSF}(f) \sigma(\mathbf{b}^f)$ . In the

case of humans the attenuation factor,  $\text{CSF}(f)$ , has to be obtained from contrast thresholds because there is no access to the output. However, for autoencoders, the computation of the output is straightforward. If the degradation of the acquisition is taken into account, the sinusoids,  $\mathbf{b}^f$ , used to simulate the measurement of the CSF have to undergo the degradation as well, and we should consider an *eye+network* system,  $S$ :

$$\mathbf{b}^f \xrightarrow{H \cdot \mathbf{b}^f + \mathbf{n}_r} \mathbf{b}_*^f \xrightarrow{N_\theta(\mathbf{b}_*^f)} \mathbf{y}^f \quad (2)$$

$S_\theta(\mathbf{b}^f)$

Therefore, one could check the attenuation factor by comparing the standard deviation of output and input:

$$\text{CSF}(f) = \frac{\sigma(S_\theta(\mathbf{b}^f))}{\sigma(\mathbf{b}^f)} = \frac{\sigma(N_\theta(H \cdot \mathbf{b}^f + \mathbf{n}_r))}{\sigma(\mathbf{b}^f)} \quad (3)$$

Note that the CSF ratio in Equation 3 (which uses degraded sinusoids to probe the network) is different from checking the Fourier response of the network, where one would use clean sinusoids at the input:

$$\mathcal{N}^F(f) = \frac{\sigma(N_\theta(\mathbf{b}^f))}{\sigma(\mathbf{b}^f)} \quad (4)$$

The relation of Equation 3 with the regular determination of the CSF in humans is illustrated in Figure 1. Of course, the ratio in Equation 3 should be computed for low-contrast sinusoids to keep parallelism with human CSF and keep the (eventually) nonlinear autoencoder in the low-energy range. For chromatic sinusoids the deviations have to be computed separately over the achromatic, red–green, and blue–yellow color channels (Mullen, 1985). In this work we use a classical opponent color space (Hurvich & Jameson, 1957) to generate achromatic and purely chromatic gratings and to decompose the corresponding responses.

Of course, plain attenuation for sinusoids in Equation 3 may not provide a full description of the action of nonlinear systems. In principle, it is not obvious why we should perform the analysis in a specific basis. Therefore, one should check to what extent waves are indeed eigenfunctions of the system.

A way to test this point is linearizing the response of the autoencoders in the low-contrast regime and check that it is shift invariant. Using a Taylor expansion, the response for low-contrast images can be approximated by the Jacobian around the origin (the zero-contrast image,  $\mathbf{0}$ , which is just a flat gray patch):

$$\begin{aligned} \mathbf{y} &= S_\theta(\mathbf{x}) \\ \mathbf{y}_{\text{low}} &= S_\theta(\mathbf{0} + \mathbf{x}_{\text{low}}) \\ &\approx S_\theta(\mathbf{0}) + \nabla_{\mathbf{x}} S_\theta(\mathbf{0}) \cdot \mathbf{x}_{\text{low}} \quad (5) \\ \mathbf{y}_{\text{low}} &\approx \nabla_{\mathbf{x}} S_\theta(\mathbf{0}) \cdot \mathbf{x}_{\text{low}} \end{aligned}$$

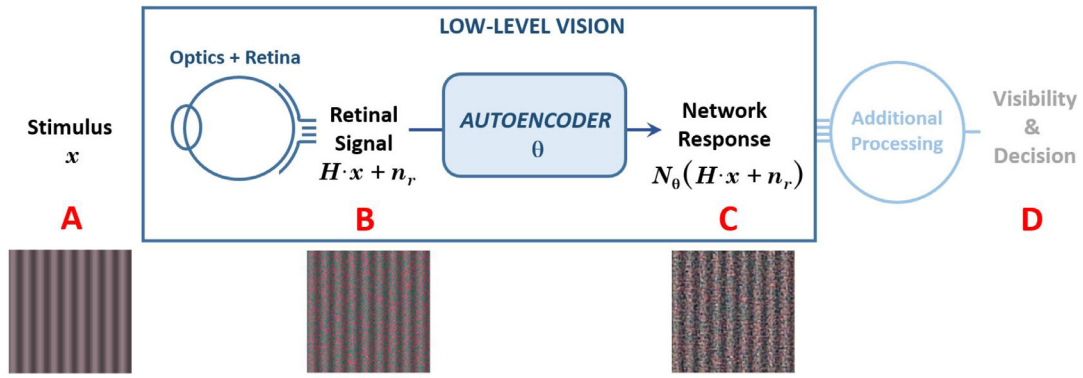


Figure 1. Definition of the CSF as a frequency-dependent attenuation factor in a system to develop low-level vision tasks. The diagram illustrates transforms in the visual signal from the input stimulus (A), the degraded signal owing to optical blur and retinal noise (B), the process of the early neural path where the output is still in a spatial LMS representation (C) modeled here by autoencoders, and additional mechanisms that compute a decision on the visibility (D). In the conventional view of the CSF as a filter, the process from A to C is assumed to be linear and (in humans) the visibility of gratings is assumed to be based on the amplitude of the response at the point C. In human psychophysics (with no access to C) the observer makes visibility decisions, and the attenuation factors are determined from the thresholds. When dealing with artificial systems we do have access to the response in C so we do not need to model the decision mechanism and we can simply estimate the CSF from the ratio in Equation 3.

where we assumed that the response for zero-contrast images is zero. If the behavior of the system at this low-energy regime is shift invariant, the Jacobian matrix can be diagonalized as  $\nabla_x S_\theta(\mathbf{0}) = B \cdot \lambda \cdot B^{-1}$ , with extended oscillatory basis functions in the columns of  $B$  (and rows of  $B^{-1}$ ). Fourier basis and cosine basis are examples of extended (nonlocal) oscillatory functions that diagonalize shift invariant systems. The reason for this result is equivalent to the emergence of cosine basis when computing the principal components of stationary signals (shift-invariant autocorrelation) (Clarke, 1981). As a result, the slope of the response for low-contrast sinusoids (the CSF) will be related to the eigendecomposition of the Jacobian of the system at  $\mathbf{0}$ . Let's compute the response for a sinusoid in this Taylor/Fourier setting to see the relation. A basis function  $\mathbf{b}^f$  with specific frequency  $f$  is orthogonal to all rows (sinusoids) in  $B^{-1}$  except that of the same frequency, that is,  $B^{-1} \cdot \mathbf{b}^f = \delta^{f'f}$ . And this delta selects the corresponding column (of frequency  $f$ ) among all the columns in the matrix  $B$ :

$$\begin{aligned} \mathbf{y}^f &= S_\theta(\mathbf{b}^f) \approx \nabla_x S_\theta(\mathbf{0}) \cdot \mathbf{b}^f \\ &\approx B \cdot \lambda \cdot B^{-1} \cdot \mathbf{b}^f \\ &\approx B \cdot \lambda \cdot \delta^{f'f} \\ &\approx \lambda_f \mathbf{b}^f \end{aligned} \quad (6)$$

So the slope of the response for basis functions of frequency  $f$  is  $\lambda_f$  (the corresponding eigenvalue of the Jacobian of the autoencoder). As a result, for systems with shift invariance in the low-contrast regime, the eigenvalues of the linear approximation of the system (eigenvalues of the Jacobian) are conceptually similar

to the CSF. A direct comparison of the eigenvalue spectrum with the CSF may not be simple because the eigenfunctions may differ from Fourier sinusoids. Examples of this include isotropic systems (with a constant sensitivity for certain  $|f|$  independent of orientation). In this case, the eigenbasis may be not sinusoids, but arbitrary linear combinations of sinusoids of the same frequency and different orientation.

Nevertheless, if the linearized version of the system (the Jacobian at  $\mathbf{0}$ ) is shift invariant, which can be seen from a convolutional structure in the Jacobian matrix, oscillatory waves are eigenfunctions of the system, and hence Equation 3 may provide a good description of the behavior of the system.

### Alternative linear characterizations of the autoencoders

A two-dimensional (2D) cartoon of the impact of the degradation and restoration processes in the probability density function (PDF) of the signal can illustrate alternative characterizations of the neural networks optimized to enhance the retinal signal (see Figure 2). In this diagram, two-pixel natural scenes (left) follow a PDF obtained from independent Student t-sources mixed by a matrix that introduces strong correlation between the luminance of the pixels. This kind of two-pixel representations is common to describe the statistics of natural images (Simoncelli & Olshausen, 2001), and mixtures of sparse components is a widely accepted model for natural scenes (Hyvärinen et al., 2009; van den Oord & Schrauwen, 2014; Malo, 2020),

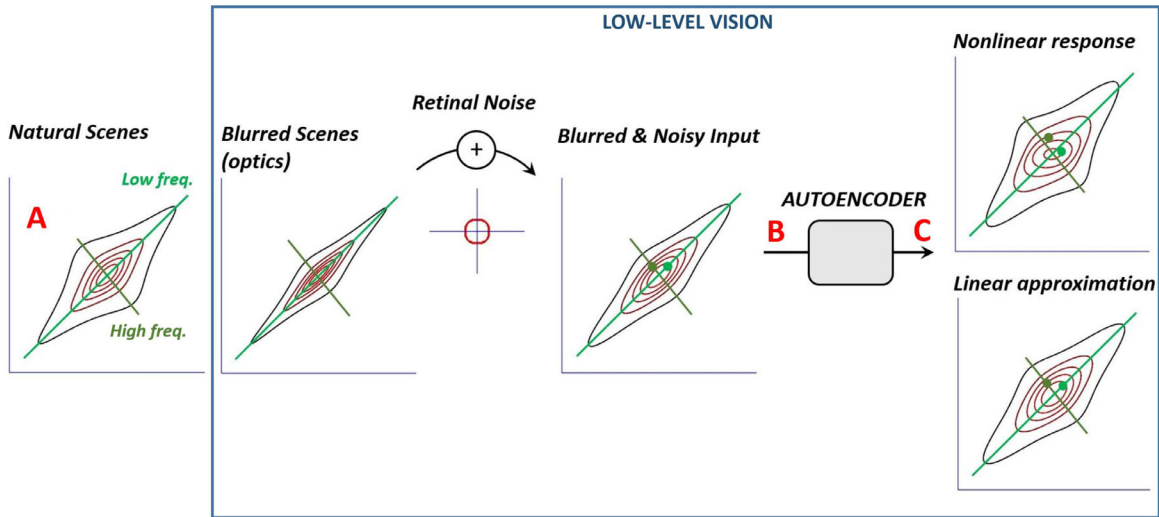


Figure 2. Degradation of the signal PDF and linear and nonlinear strategies to compensate the degradation. The axes of the plots represent the luminance in each of the photodetectors of two-pixel images as in [Simoncelli and Olshausen \(2001\)](#). (A) The PDF of natural scenes: marginal heavy tailed distributions of oscillatory functions mixed to have strong correlation in the pixel domain. Optical blur (second panel) implies contraction of the high frequencies. Additive retinal noise implies the convolution by the PDF of the noise leading to the PDF in (B). The solutions to the restoration problem at point C may include i) nonlinear transforms such as the one represented by the PDF at the right-top, and ii) linear transforms as the one at the right-bottom.

and appropriate enough for this illustration. In this diagram, the low-frequency direction corresponds with the main diagonal (where the two pixels have the same luminance) and the high-frequency direction is orthogonal (for images where one of the pixels is brighter than the other). The zero-contrast image is at the crossing point of the frequency axes.

Optical blur implies the attenuation of high-frequency versus low-frequency components and hence the contraction of the dataset as shown in the second panel. Assuming linear and noisy photoreceptors, the PDF of the retinal response results from the convolution of the PDF of the blurred images with the PDF of the noise (the function with circular support). The result (third panel) is the input to the autoencoder, whose goal is recovering the distribution at the first panel. Linear solutions are limited to global scaling of the domain (for instance, by inverting the contraction introduced by the blur), whereas nonlinear solutions may twist the domain in arbitrary ways.

In this setting, the computation of the CSF according to [Equation 3](#) means putting low-contrast sinusoids (e.g., the samples highlighted in green in the third panel) through the system, and checking the amplitude of the output (green dots at the panels at the right) over the directions of the input. This nonlinear example illustrates the fact that the behavior can be contrast dependent (see the different twist in the concentric contours). This graphical view illustrates the difference between three possible linear characterizations, with  $y = M \cdot x$ :

- **The optimal linear solution:** The matrix  $M$  that better relates the input  $x$  with the desired output  $x_c$ . This is the  $M$  that minimizes the expected value  $E\{|x_c - M \cdot x|_2\}$ . Assuming a representative set of  $N$  clean/distorted pairs stacked in the matrices  $X_c = [x_c^{(1)} x_c^{(2)} \dots x_c^{(N)}]$  and  $X = [x^{(1)} x^{(2)} \dots x^{(N)}]$ , the optimal solution in Euclidean terms is given by the pseudoinverse:

$$M = X_c \cdot X^\dagger \quad (7)$$

- **Globally linearized network:** The matrix  $M$  that better describes the nonlinear behavior of the network over the whole set of natural images. This is the  $M$  that minimizes  $E\{|S(x, \theta) - M \cdot x|_2\}$ . Assuming a representative set of  $N$  input/output pairs stacked in the matrices  $X = [x^{(1)} x^{(2)} \dots x^{(N)}]$ , and  $Y = [y^{(1)} y^{(2)} \dots y^{(N)}]$ , the solution is given by the pseudoinverse:

$$M = Y \cdot X^\dagger \quad (8)$$

- **Locally linearized network at 0:** The matrix  $M$  that better describes the nonlinear behavior of the network for low-contrast images. This is the  $M$  that minimizes  $E\{|\nabla_x S(\mathbf{0}, \theta) \cdot x_{low} - M \cdot x_{low}|_2\}$ . Of course, this could be empirically approximated by  $M = Y_{low} \cdot X_{low}^\dagger$ , but in this case the obvious exact solution is:

$$M = \nabla_x S_\theta(\mathbf{0}) \quad (9)$$

Although the optimal linear solution (or the optimal linear network) is a convenient reference to describe the problem, the other two options are different characterizations of the autoencoder. Equation 8 summarizes the behavior of the network in a single matrix, and Equation 9 is a description only valid around  $\mathbf{0}$ , and hence more closely connected to the low-contrast regime of the CSF. The eigenanalysis cited for  $\nabla_x S$  in Equation 6 can be applied for the three matrix characterizations, but it is important to note the differences between them.

The Jacobian of cascades of linear+nonlinear layers (as in autoencoders based on CNNs) can be obtained either analytically,<sup>2</sup> or it can be obtained via automatic differentiation or alternative methods based on system identification (Berardino et al., 2017). However, these procedures are tedious, so in Gomez-Villa et al. (2020), we took the more straightforward approach represented by Equation 8.

The different linear characterizations considered in this section and the diagram in Figure 2 illustrate that the behavior of a nonlinear autoencoder for high contrasts may be substantially different from the threshold behavior. Therefore, the attenuation of sinusoids by the linearized system (by the matrix Equation 8) will be compared with the result of Equation 3.

### Limitation of the proposed CSF definition in autoencoders

To maximize the equivalence to human CSFs, the proposed procedure (the ratio in Equation 3, which compares the signals at points C and A in Figure 1) implies the consideration of the retinal degradation process. This consideration of the retinal noise will be shown to improve similarity with human CSFs in the experiment 3, but it comes at a cost. Note that, even if the role of the autoencoder is compensating the retinal noise, complete removal is not possible. Therefore, there is some residual distortion in the response after the autoencoder. As a result, the standard deviation in the numerator of the proposed Equation 3 not only measures the contrast of the output grating, but also measures the energy of the residual noise. In this way, when the contrast of the sinusoids  $b^f$  is very small, as expected in threshold conditions, the standard deviation maybe measuring more the residual noise than the contrast of the output. The limitation of Equation 3 is that it has to be applied to sinusoids of relatively high contrasts so that the energy of the response coming from the sinusoid is bigger than the energy of the response coming from the noise.

One can overcome this limitation in two ways: 1) by computing the response many times for different noise evaluations and cancelling the residual noise by

averaging over the realizations, and 2) by using relatively high-contrast sinusoids so that the effect of the residual noise is negligible.

In this work (for computational convenience) we used the second approach: we probed the models with sinusoids with contrasts in the range [0.07, 0.6]. The lower limit is certainly higher than the minimum absolute threshold of the standard spatial observer (which is approximately 0.005) (Watson & Malo, 2002; Watson & Ahumada, 2005). Nevertheless, we choose this range for two reasons: first, 0.07 is the average of the threshold achromatic contrasts in the standard spatial observer, and second, we empirically checked that the effect of the noise was negligible above this value.

## Experiments

The Introduction raised questions on the role of low-level vision goals to explain the CSFs, the emergence of the CSFs in autoencoders working to solve these goals, and the eventual advantage of progressively more flexible models in explaining the CSFs. To address these issues in the more general spatiotemporal chromatic case, 1) we perform two extensive experiments (one with images, and one with video) to compensate biologically sensible degradation of the retinal signal (compensation of biodistortion), using a range of CNN architectures of different depth or flexibility, 2) we consider alternative low-level functional goals such as chromatic adaptation and the compensation of the effect of bottlenecks, 3) we consider different levels of biodistortion, chromatic shifts in different directions, and bottlenecks with different restrictions, and 4) we consider the consistency of the results under changes in the statistics of the signal. In this section, we describe the experimental setting of these simulations.

### Functional goals

Compensation of retinal biodistortion (biological blur and noise) consists of overcoming the degradation introduced in the acquisition of the visual signal. Specifically, the top panel in Figure 3 shows how a natural scene is degraded at the output of the retina according to the variations of the eye MTF for different pupil diameters (from top to bottom,  $d = 2$  mm,  $d = 4$  mm and  $d = 6$  mm), and a sensible range of Poisson retinal noise levels (from left to right, Fano factors  $F = 0.25$ ,  $F = 0.5$  and  $F = 1$ ). Variations of the MTF have been simulated with the expression in Watson (2013), and the noise in LMS sensors has been estimated in the discrete representation of the



Figure 3. Functional goals. Possible low-level goals of the autoencoders are compensating the following distortions in the visual input. *Top panel*: different levels of retinal biodistortion. *Bottom panel (first row)* Changes in the spectral illumination. *Bottom panel (second row)* Changes in illumination + retinal biodegradation. An alternative low-level goal is the reconstruction of the signal in presence of bottlenecks (as in the architectures considered below, Figure 4 right). Note: This selfie of the corresponding author was sent to the first author to test the models on a copyright-free image.

input digital image as in Esteve et al. (2020). In that work, noise was obtained by stimulating the ISETBio retinal model (Cottaris et al., 2019, 2020) with flat stimuli of controlled size and tristimulus values over short and long exposure times. Cartesian resampling of the random cone mosaic of the retinal model and

integration of the photocurrents over space/time reveals the effective Poisson nature of the noise (in the original LMS units) and allows the estimation of the effective Fano factor in the original discrete grid of the input image. In that way, we can easily generate calibrated noisy retinal images by adding this effective



Poisson noise in the LMS representation of the digital image. The illustrations in Figure 3 come from the transformation of the LMS tristimulus images into the RGB digital counts for proper display.

### Chromatic adaptation

Consists of the compensation of the deviations of the signal induced by the change of illuminant. The bottom panel shows how the image of a natural scene changes under changes in the shape of the spectral radiance of the illuminant. The change of illuminant in a digital image was simulated in this way: each pixel of the image was associated with a reflectance chosen from a large database of natural reflectances so that under an equienergetic illuminant led to the tristimulus values of the pixel. Then, a black-body radiator, which simulates natural ambient light along the day, was used to generate spectra of the same energy but different shape. From there, we could get versions of the scene under arbitrary color temperatures. This process is straightforward using the functionalities and databases of Colorlab (Malo & Luque, 2002). Of course, this process is just an approximation because it disregards the (unknown) geometry of the scene and assumes a flat Lambertian world. Nevertheless, as illustrated in the examples in Figure 3, it does a good qualitative job to generate controlled samples to check chromatic adaptation in large image databases.

### Compensation of chromatic shifts and biodistortions

The reason to consider this combination is that pure chromatic shifts with no additional distortion is not a realistic input for the visual pathway: the image acquisition front-end does exist and, hence, what we called biodistortion has to be taken into account. Such combination of distortions is illustrated by the second row of the bottom panel in Figure 3. Note that in the examples involving chromatic deviations (bottom panel) we introduced a flat-reflectance frame to help the models to cope with the chromatic adaptation<sup>3</sup>.

### Compensation of bottlenecks (pure reconstruction)

Consists of recovering the input after the signal has gone through a bottleneck. Examples of bottlenecks include the restriction of the spatial resolution or the restriction of the number of features (or channels) in the representation. Figure 4 (right panel) shows an illustrative range of architectures: from cases that expand the number of features (no bottleneck) to a variety of cases that introduce local pooling, reduce the number of features, or try to compensate the effect of spatial undersampling by increasing the number of features. Bottlenecks may imply severe information loss if the representation is not optimized. Therefore, pure

reconstruction of the signal in presence of bottlenecks is a sensible low-level goal to explore.

### Compensation of bottlenecks and biodistortion

As stated, errors in the acquisition front end (biodistortion) do exist, so their consideration together with bottleneck compensation makes the goal more realistic.

All in all, we explored nine levels of biodegradation, chromatic adaptation in the blueish and the reddish directions ( $T = 8600 K$  and  $T = 4400 K$ , respectively), and the combination of the central biodistortion with the considered chromatic deviations. We considered a pure reconstruction task with the eight bottleneck configurations in Figure 4, and the compensation of these bottlenecks was also combined with the central biodistortion case. The optical/retinal degradation in movies was applied in a frame-by-frame basis. No experiments involving chromatic adaptation or bottlenecks were done in movies, but only in natural and cartoon images.

The above computational goals are all measured in *distortion* terms, or how well the deviations  $\varepsilon_{LMS}$  were compensated. However, even within this low abstraction level, other computational goals could be considered together with the distortion, as for instance the information or the energy of the signal. In the experiments we restrict ourselves to the considered cases of distortion minimization and purely architectural bottlenecks. The discussion suggests how the goals considered here could be related or combined with other kind of goals or more general (energy or information) bottlenecks.

## Architectures

In this work, we consider 2D CNNs that act on spatiochromatic signals and three-dimensional (3D) CNNs that act on spatiotemporal signals (color images and color videos).

The set of explored architectures is shown in Figure 4. These are variations of the basic toy networks studied in Gomez-Villa et al. (2019, 2020): autoencoders with convolutional layers made of 8 feature maps with kernels of spatial size  $5 \times 5$  and sigmoids or rectified linear units (ReLU) as activation functions. From that starting point, here we consider a range of nets of increasing *depth* and flexibility: from the linear network in Equation 7 (as a convenient baseline reference of one layer with no flexibility), and CNNs with two layers to eight layers, both for the 2D and 3D cases. Moreover, we also consider a range of architectures with different bottlenecks, in this case, only 2D.

Of course, the range of possible architectures is virtually infinite and an exhaustive exploration of the architecture space is out of the scope of this work.




Figure 4. Architectures. (Left) Range of architectures of different *depth* following the basic structure of the nets studied in Gomez-Villa et al. (2019, 2020): three channels at the input and at the last (output) layer. In this work the input is in an LMS color representation. The rest of the layers have eight features with no undersampling or bottleneck, as represented by the eight blue lines of the same length. These architectures of increasing depth are used to study the compensation of the biodistortion in images and videos. The best of these architectures in terms of CSFs (which turns out to be the two layers example) is used with images to explore chromatic adaptation. (Right) Range of architectures used to illustrate the effect of *bottlenecks*. In this case, the inner layer in the four layer architecture at the left is systematically expanded (in A) or contracted (from C to G, either in the number of features or in the spatial resolution). See numbers and indications of pooling/upsampling and corresponding length of the layers-) to generate a range of bottlenecks. Finally, an illustrative U-Net with no residual connections is also considered in the experiments with bottlenecks.

However, note that the considered set of architectures of progressive flexibility and constraints is appropriate for the aim of this work for two reasons: 1) these architectures do a good job in fulfilling the goal so they are good examples to reason about systems that work according to the considered function, and 2) they display a range of flexibility and accuracy in the goal which is appropriate to illustrate the proposed questions (eventual emergence of the CSF and other nonlinearities, and qualitative effect in the CSF of increased flexibility and improvements in the goal accuracy).

The first point (the considered *toy* models do a reasonably good job in fulfilling the goal) is a technical issue that is demonstrated by the performance tables shown below and by the specific learning curves and reconstructions included in Appendix A. However, to put this quantitative performance in context, it is interesting to note that the retinal biodistortion is not an easy task to solve for general-purpose state-of-the-art image restoration CNNs. In particular, following Gomez-Villa et al. (2020), on top of the described toy networks, the computation of the CSFs of cutting-edge deeper models designed for restoration could be an illustrative limit to consider. However, we found that the combination of representative examples

of generic CNNs for denoising (Zhang et al., 2017; Soh & Cho, 2021) and deblurring (Tao et al., 2018), which gave excellent results with arbitrary Gaussian noise and blur in Gomez-Villa et al. (2020), is not satisfactory with biological distortion. In particular, generic enhancement algorithms did not produce better results than the considered simple architectures (specifically trained for this biodistortion).

Of course, this does not mean that the toy models used here are better than the state of the art, nor that state-of-the-art models are intrinsically unable to deal with this biological degradation. One could certainly fine-tune these deep architectures for the biodistortion and then get a better result than with the considered set of architectures, but that is not the goal of this work. The relevant argument in favor of the considered (toy) architectures for our purposes here is this: the fact that generic blind restoration CNNs need to be retrained to get better results than the proposed models means that these simple models can be considered as good (enough) examples of systems actually fulfilling the goal.

Regarding the second point (the considered set of architectures is good enough to illustrate interesting questions), consider that i) according to the results presented below (Results, Tables 1 and 4) the toy

nonlinear models decrease by up to 35% and 48% the error of the optimal linear solution in images and video, respectively, and ii) the best nonlinear model reduces the error of the shallower nonlinear model by 21% and 12% in images and video respectively.

In summary, the considered set of architectures (progressively deeper CNNs and a range of bottlenecks) does a reasonable job in optimizing the goals, and it is wide enough to illustrate changes in the achievement of the goals. As a result, the considered set of architectures is appropriate to address the questions raised in the introduction.

See [Appendix A](#) for implementation details. Data and code are available at <http://isp.uv.es/code/visioncolor/autoencoderCSF.html>

See [Appendix B](#) for details on the databases to generate the training stimuli and the stimuli used to probe networks.

## Assessing the quality of the CSF results

The CSFs defined for the autoencoders may be subject to two arbitrary scale factors. On the one hand, the response of the network could be multiplied by an arbitrary global scale factor and hence, the numerator in [Equation 3](#) (and the CSF amplitude) would be multiplied by this scale factor as well. We refer to this global scale factor in the amplitude as  $\alpha_{\text{CSF}}$ . On the other hand, the sampling assumptions (or assumptions on the extent of the signal, or the viewing distance) introduced in the description of the stimuli are arbitrary and they imply an arbitrary scaling in the frequency axis of our Fourier domains. We refer to this scale factor on frequency as  $\alpha_f$ .

The factor on amplitude is not a major problem: one network and a modified version with its outputs multiplied by  $\alpha_{\text{CSF}}$  are equivalent and their quality should be rated the same. The factor on frequency does not reduce the validity of the results either as long as it is moderate. Note that using the MTF expressions in [Watson \(2013\)](#), if the filter corresponding with a pupil of 3.5 mm is modified by applying  $\alpha_f = 0.75$  or  $\alpha_f = 4.5$ , the resulting MTF is similar to what would have been obtained with  $d = 2$  mm or  $d = 6$  mm, respectively. Therefore, as changes in the MTF (the only element where the scaling in frequency matters) are plausible if  $\alpha_f \in [0.75, 4.5]$ , one should also discount moderate variations of this factor when assessing the quality of the CSFs.

The similarity between the model and the human CSFs will be measured by the Euclidean distance between the CSF vectors, averaging over the frequency,  $f$ , and the chromatic channels,  $c$  (achromatic, red–green and yellow–blue), which will be referred to as:

$$\text{RMSE} = \left( \sum_{f,c} (\text{CSF}_c^{\text{scaled}}(f) - \text{CSF}_c^{\text{human}}(f))^2 \right)^{\frac{1}{2}}, \quad (10)$$

where the *scaled* attenuation factors of the model are related to the *raw* attenuation factors of the model as:

$$\text{CSF}^{\text{scaled}}(f) = \alpha_{\text{CSF}} \cdot \text{CSF}^{\text{raw}}(\alpha_f \cdot f) \quad (11)$$

In the following, we report the scaled CSFs together with the scaling factors that minimize the distance with human CSFs.

It is important to mention that the relative scaling between the CSFs in the three chromatic channels is a characteristic feature of a network (or system) and it should not be modified. Therefore, the same factors in [Equation 11](#) are applied to the three CSFs. With these considerations, the CSFs reported below represent the closest approximation the models may give to the human CSFs, and hence the comparison between them is fair.

The magnitude of the RMSE errors has to be understood in reference to the maximum value of the human sensitivity. As a convenient example to have in mind, RMSE – 22 corresponds with an average deviation of 10% of the scale of the human spatiotemporal CSF at every frequency and chromatic channel. This is because the maximum sensitivity is approximately 200 for stationary gratings and about 220 for moving gratings ([Watson & Malo, 2002](#); [Kelly, 1979](#)).

## List of experiments

The empirical exploration of the considered architectures consists of six experiments. Experiments 1 through 5 deal with spatiochromatic stimuli and 2D networks, and experiment 6 deals with spatiotemporal chromatic stimuli and 3D networks. As stated, the computational goals are measured by the Euclidean distance between the reconstructed image and the original image, referred to as  $\epsilon_{LMS}$ . The similarity with the human behavior is measured in terms of the Euclidean distance between the model CSFs and the human CSFs, that is, the RMSE defined in [Equation 10](#).

- **Experiment 1: Spatiochromatic CSFs from biodistortion compensation by a range of architectures.** This experiment is focused on the central degradation shown in the first panel of [Figure 3](#) ( $d = 4$  mm,  $F = 0.5$ ) and analyzes in detail the CSFs for nine architectures: the optimal linear network, and eight CNN architectures with two, four, six, and eight layers with either sigmoid or ReLU activations, all optimized according to this distortion–compensation goal. Once the architectures are properly trained (using  $20 \cdot 10^3$  images of the ImageNet database cited in [Appendix B](#),  $18 \cdot 10^3$  for training and  $2 \cdot 10^3$  for validation), we get the numerical performance

of the models in the independent test set of  $10^3$  images. The sizes of the train/validation/test sets are the same in all experiments with images, experiments 1 to 5. Throughout all the experiments, the performance is expressed as the average  $\varepsilon_{LMS}$  of the reconstruction in LMS space over 20 batches of 50 randomly chosen images per batch. The standard deviation over these 20 computations is also reported. The learning curves (train/validation) and the reconstructions of one representative test image are given in [Appendix C](#). Then, the CSFs (attenuation factors) of the trained models are computed according to the method described in the [Methods: Estimating contrast sensitivity in autoencoders](#) for gratings of different contrasts. The eventual variation of the attenuation reveals the nonlinear nature of the contrast response for gratings. In experiment 1, we also show the CSFs of the linear network and the linearized versions of the nonlinear networks introduced in the [Methods: Estimating contrast sensitivity in autoencoders](#). From the results of experiment 1, one of the nonlinear models is chosen as having representative resemblance with human behavior in terms of the CSFs (two layers with ReLU activation). Experiments 2, 3, and 4 further explore the behavior of this specific model in a number of conditions.

- **Experiment 2: Consistency of the CSFs from biodistortion compensation over a range of distortion levels.** This experiment is focused on the representative architecture selected after experiment 1, and checks its CSFs when trained for the nine different degradation levels considered in the first panel of [Figure 3](#).
- **Experiment 3: CSFs from chromatic adaptation and biodistortion compensation.** This experiment checks the CSFs of the representative architecture selected after experiment 1, when it is trained for i) the biodegradation compensation alone ii) the degradation compensation together with compensation of a bluish illuminant, iii) the degradation compensation together with compensation of a reddish illuminant, iv) pure compensation of a bluish illuminant, and v) pure compensation of a reddish illuminant. In the illustration of [Figure 3](#), these correspond with the five distorted versions closer to the clean image under equienergetic illuminant. As stated, the purely chromatic deviations are not realistic because they disregard the optics and retinal noise. However, they represent an illustrative reference. In the same vein, as a convenient reference, in this experiment we compute the CSF in two ways: a) the proposed (realistic) way, [Equation 3](#), by putting the clean gratings through the retinal degradation before entering the network, and b) the idealized

way, [Equation 4](#), in which we simply put the clean gratings through the considered network. This will stress the difference in the obtained CSFs when considering realistic spatial degradations or not.

- **Experiment 4: Consistency of the human/non-human CSFs under change in signal statistics.** Here we reconsider the chromatic adaptation and the degradation–compensation goals of experiment 3 now using stimuli of (apparently) quite different spatiochromatic statistics: the images from the Pink Panther cartoons. All the other settings remain the same as in experiment 3.
- **Experiment 5: CSFs from bottleneck–compensation and biodistortion compensation.** This experiment shows the CSFs of the systems that emerge from imposing pure reconstruction of the signal in presence of bottlenecks in the network (the eight examples in [Figure 4](#), right). Pure reconstruction is compared with the compensation of biodistortion in the same architectures. Given the similarity between activation options found in experiment 1, here we just explore the ReLU case.
- **Experiment 6: Spatiotemporal chromatic CSFs from biodistortion compensation by a range of architectures.** Here we check the fundamental findings of experiment 1 for spatiotemporal chromatic gratings on 3D networks optimized for degradation–compensation. Given the similarity between activation options found in experiment 1, here we just explore the sigmoid case. Therefore, we explored five architectures: the linear one and two, four, six, and eight layers with sigmoid. In this spatiotemporal case we used  $22 \cdot 10^3$  video patches in the learning ( $20 \cdot 10^3$  for training and  $2 \cdot 10^3$  for validation), and  $3 \cdot 10^3$  for test.

## Results

Results in all the experiments have two parts: 1) the perception part, with the CSFs and the contrast responses of the networks, and 2) the technical part, with evidences of the convergence of the models, numerical performance in reconstruction, and visual examples of the performance in reconstructing images. The main text is focused on the perception part, while all the technical material is given in the [Appendix C](#).

### Experiment 1: Spatiochromatic CSFs from biodistortion compensation

[Figure 5](#) shows the achromatic and chromatic CSFs of the considered models (the linear solution and the eight CNNs) together with the human CSFs for convenient reference. The human data come from

	Comput. goal		CSFs		Comput. goal		CSFs	
	$\varepsilon_{LMS}$		RMSE		$\varepsilon_{LMS}$		RMSE	
Distortion	15.5 ± 0.2							
Linear net	<b>13.1 ± 0.1</b>		<b>24.4</b>					
CNNs	Sigmoid nonlinear	ReLU nonlinear	Sigmoid nonlinear	ReLU nonlinear	Sigmoid linearized	ReLU linearized	Sigmoid linearized	ReLU linearized
2 Layers	<b>10.3 ± 0.1</b>	<b>10.8 ± 0.1</b>	<b>26.9</b>	<b>24.5 ± 0.3</b>	12.5 ± 0.1	12.6 ± 0.2	24.1	22.8
4 Layers	<b>8.9 ± 0.1</b>	<b>9.1 ± 0.1</b>	<b>26.6</b>	<b>28.5</b>	12.5 ± 0.1	12.5 ± 0.1	23.2	23.1
6 Layers	<b>8.7 ± 0.1</b>	<b>8.5 ± 0.1</b>	<b>29.7 ± 0.7</b>	<b>33.1</b>	12.5 ± 0.2	12.5 ± 0.2	23.2	23.1
8 Layers	<b>8.9 ± 0.2</b>	<b>8.7 ± 0.1</b>	<b>31.2</b>	<b>31.6</b>	12.6 ± 0.1	12.7 ± 0.1	27.0	27.4

Table 1. Experiment 1: Emergence of CSFs from biodistortion compensation (computational goal and error in CSFs). The achievement of the computational goal is described by  $\varepsilon_{LMS}$  (error of the reconstructed signal in LMS) for batches of images of the independent test set (averages and standard deviations from 20 realizations with 50 images/batch). The distance between the CSFs of the networks and the human CSFs is measured by the RMSE between the functions, Equation 10. Uncertainty of the RMSE was estimated only in two cases (two layer ReLU and six layer sigmoid) and is represented here by the standard error of the corresponding means. It is interesting to note that the optimal linear solution (computed from the train set) has worse performance in the test set than the linearized versions of the networks. Numbers in bold style refer to nonlinear networks and numbers in regular style refer to linearized networks.

the achromatic standard spatial observer (Watson & Malo, 2002; Watson & Ahumada, 2005) and from the measurements in Mullen (1985). The plots for the nonlinear models show the attenuation factors (CSFs) for gratings of different contrast (dark to light colors mean lower to higher contrasts).

These plots include the RMSE measure of the difference of the artificial CSFs with the human CSFs. The insets also show the optimal values of the arbitrary scaling factors ( $\alpha_f$ ,  $\alpha_{CSF}$ ) applied to the axes of the raw CSFs of the network to minimize the distance with the human CSFs. Because these optimal scaling factors values were found exhaustively in all cases, the comparison of the final CSFs and RMSE values is fair.

Results show the emergence of a band-pass sensitivity in the achromatic channel and low-pass sensitivities in the chromatic channels. The bandwidth of the chromatic channels is always substantially narrower than the achromatic bandwidth. These properties are qualitatively in line with human behavior.

Shallower networks (either ReLU or sigmoid) display a greater resemblance with human CSFs. In particular, deeper nets introduce substantial distortion in the chromatic channels: note that the red–green channel is over attenuated (particularly for the eight layer architectures but also in the six layer cases). The RMSE scores summarize these differences and show that shallower nets (two and four layers) provide better explanations of the CSFs than deeper nets (six and eight layers).

Interestingly, the optimal linear solution (a single dense layer with identity activation) is the one that better reproduces the CSFs. However, by its linear nature, it cannot include contrast-dependent behavior.

In this regard, the shallow networks (two layers) display a consistent decay of the gain (attenuation factor) with contrast. This decay has an impact on the contrast response curves for gratings. The contrast response curves describe the evolution of the amplitude of the response to a grating as a function of the contrast of the grating. In humans, contrast response curves are increasing saturating functions both for achromatic gratings (Legge & Foley, 1980; Legge, 1981) and for chromatic gratings (Martinez-Uriegas, 1997). The decay found in two layer CNNs implies a saturation of the contrast response curves for these shallow CNNs, in line with human behavior. Figure 6 shows representative examples of these response curves: Although the two layer network (top row) consistently displays saturating behavior for every frequency, the deeper net (bottom row) shows non-human (linear or expansion) responses.

Finally, Figure 7 shows the CSFs corresponding with the linearized versions of the nonlinear CNNs, Equation 8. Of course, the linear approximations have contrast-independent behavior and hence the same CSF for all contrasts. The global linear approximations of the nonlinear models improve the resemblance of the CSFs with human behavior: the linearized shallow nets are closer to humans than the linear model, and linearization corrects over attenuation of the chromatic channels in the six layer models. However this increased similarity with human CSF comes at the cost of a significant drop in the performance (see the increase in  $\varepsilon_{LMS}$  error in Table 1). The linearization leads to rigid models that disregard the differences between the original nonlinear models and behave more similarly. In any case, linearization does not overcome the

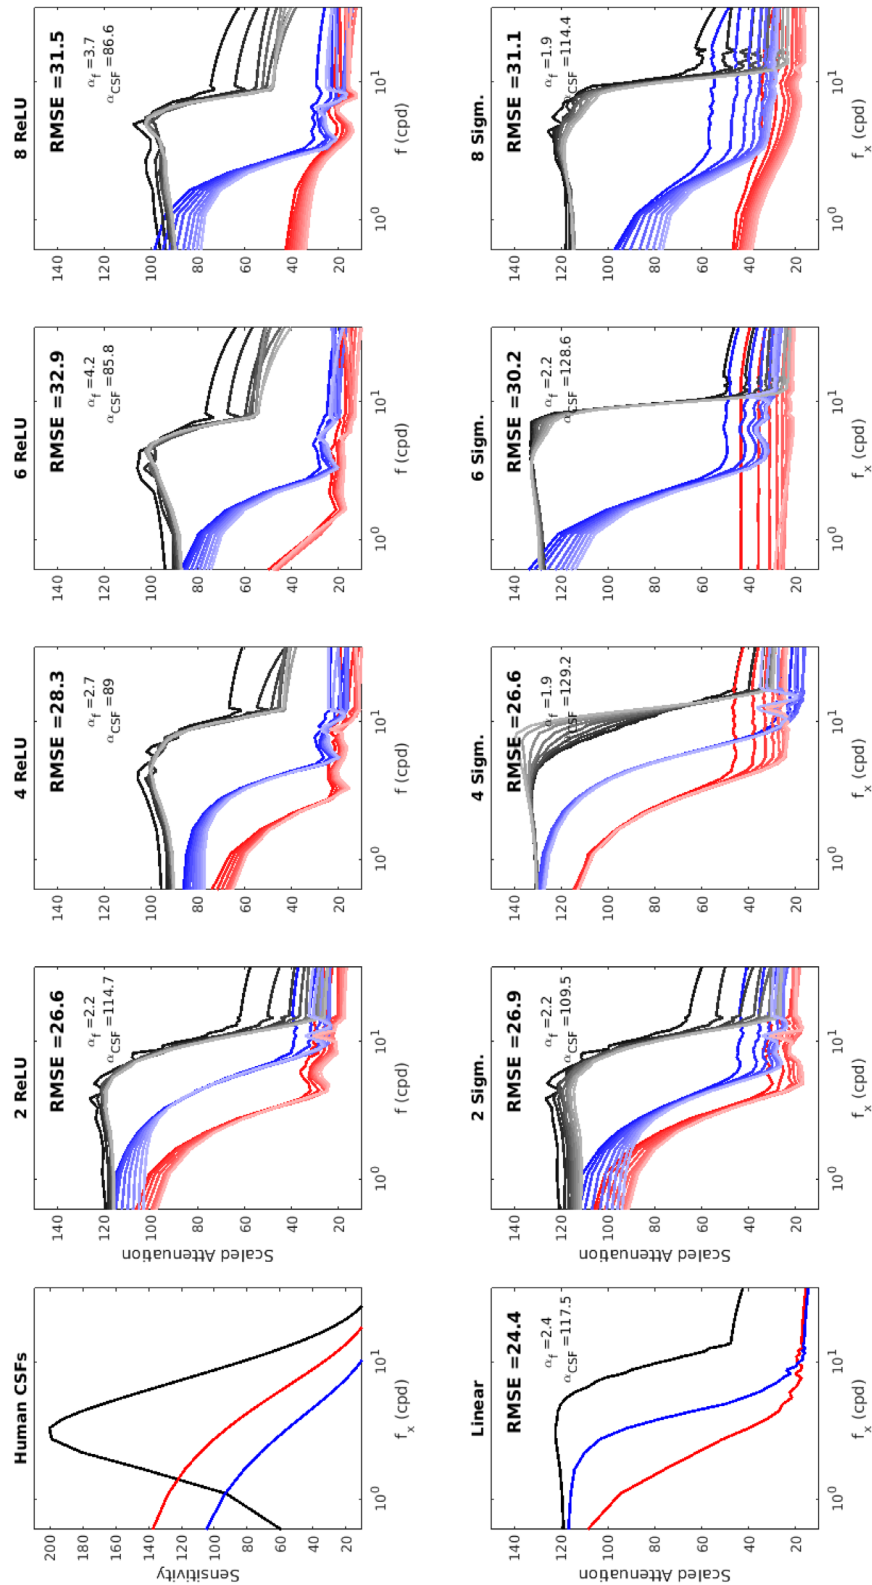


Figure 5. Experiment 1 (spatiochromatic CSFs from the compensation of biodistortion). Attenuation factors for gratings of different contrast for a range of CNN autoencoders trained for biodistortion compensation. Achromatic, blue, and red lines refer to the CSFs of the achromatic, red–green, and blue–yellow channels, respectively. Dark to light colors refer to progressively higher contrasts (evenly spaced in the range [0.07–0.6]). The human CSFs (top left) and the CSFs of the optimal linear solution (bottom left) are also shown as a convenient reference. RMSE measures describe the differences between the model and the human CSFs. The plots also display the values for the scaling factors in frequency and amplitude described in the text, Equation 11.

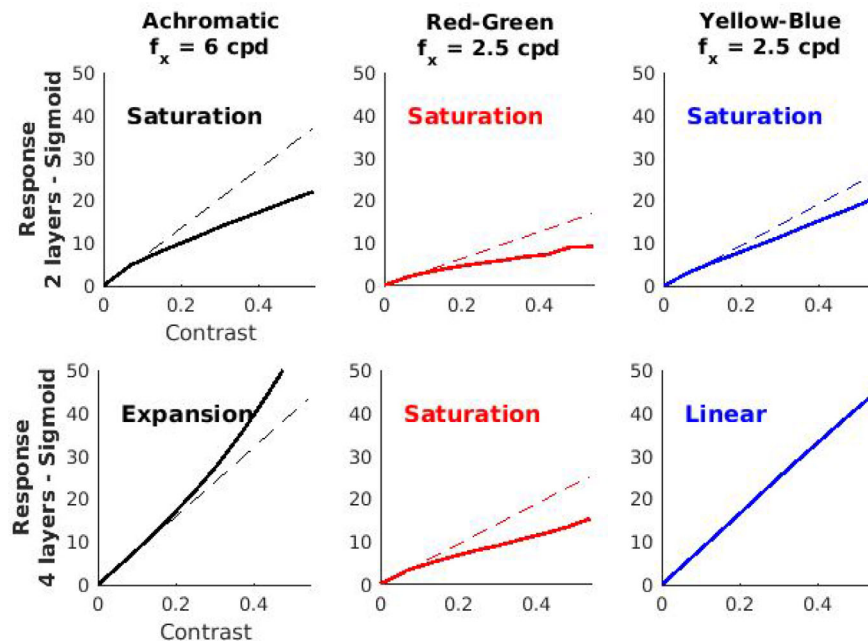


Figure 6. Experiment 1 (Illustrative contrast responses from the compensation of biodistortion). Representative examples of nonlinear responses for achromatic and chromatic sinusoids found. Saturation in these responses comes from the decay in the attenuation factors with contrast in Figure 5. Similarly, expansion comes from the increase in the attenuation factors. The linear behavior at the low-contrast regime has been plotted with dashed line as useful reference to highlight the nonlinear behavior. It is interesting to note that the saturating or expansive nature of the final contrast nonlinearity is a collective behavior that is not trivially attached to the specific (saturating or expansive) nonlinearities in individual neurons.

overattenuation of the red–green channel in the eight layer models.

Table 1 shows that while deeper networks are significantly better at fulfilling the computational goal (as expected from their increased flexibility), they are worse than shallow nets in reproducing the human behavior (as seen in Figures 5–7).

Progressive improvement in the goal for increasing depth is numerically substantial (and also visible in the reconstructed signals in Figure 14 in Appendix A) from two, four, to six layers, and the numerical performance stays (statistically) the same for eight layers. For this last case, there are chromatic issues in line with what was been found in the CSFs: the colorfulness of the reconstruction in Figure 14 is related to the relative gain of the chromatic channels. In particular, the consistent underestimation of the red–green CSF by the eight layer CNNs (either using ReLU or sigmoid activation) leads to low-saturation images. Interestingly, this effect is also visible in the reconstructed images coming from the linearized CNNs (Figure 15) and is consistent with the strong attenuation of the RG channel in the linearized eight layer architectures in Figure 7.

It is important to stress that the deviations in the chromatic CSFs in deep models do not come from not fulfilling the goal or having poor convergence in the training. First, all models (even the linear one) do reduce the error of the original retinal degradation

(Table 1) so they are fulfilling the goal. And second, the learning curves in the Appendix C (Figure 14) show that all models achieved a plateau in the training thus indicating proper convergence. Moreover, the asymptotic values achieved in the learning are consistent with the  $\varepsilon_{LMS}$  in test shown in Table 1.

As stated, RMSE errors in Table 1 have to be interpreted in terms of the scale of the human CSF. For example, the best and the worst CNNs (RMSE = 24.4 and RMSE = 33.1, respectively) have average deviations of 11% and 15% with regard to the maximum human sensitivity. Of course, a single figure of merit averaged over frequencies and chromatic channels may hide an uneven distribution of the errors. For instance, consider the specific six layer sigmoid CSF shown in Figure 5, which displays a clear over attenuation of the red–green channel. In that case, if the global RMSE = 30.2 is broken down into its chromatic components we have 29.0, 35.0, and 26.8 for the achromatic, red–green, and yellow–blue errors, which clearly point out that the biggest problem is in the red–green sensitivity. The same is true for the average over spatial frequencies: the global description does not stress the discrepancy in the low frequencies of the achromatic channel. That is why the (necessarily limited) description in the tables comes together with the explicit plots of the three CSFs for different contrasts.




Figure 7. Experiment 1 (spatiochromatic CSFs for linearized networks in biodistortion compensation). The human CSFs and the CSF of the optimal linear network are also included as reference.



Another important technical issue is the consistency of the CSFs over random initialization. This is easy to check by training a number of times the same architecture for the same computational task and over the same set of stimuli but from different initial values of the model parameters. Given the intensive computation required,<sup>4</sup> we checked this variability only in two illustrative models: one with reasonably human-like behavior (two layer ReLU), and another with less-human CSFs (six layer sigmoid). In these two models, we retrained them 20 additional times and recomputed the corresponding CSFs (results not plotted). In the six layer case, all the explored seeds lead to a flat red–green CSF of too low sensitivity (i.e., a non-human behavior), and in some cases even the blue–yellow sensitivity was strongly attenuated too. On the contrary, the two layer case systematically leads to better CSFs, as summarized by the RMSE in [Table 1](#), where the uncertainty is represented by the standard error of the mean. The shape of the sensitivities is pretty consistent in both cases, always better for the two layer case. At the same time, and not surprisingly, the six layer architectures systematically led to lower  $\varepsilon_{LMS}$  error. Only 1 of the 42 realizations (21 per model) led to a clear outlier (RMSE = 43.8 in the six layer case) and even for this CSF-outlier the distortion  $\varepsilon_{LMS}$  was not off the distribution. According to the observed consistency, the remaining 49 configurations of task/architecture in the work were studied using a single random initialization of the parameters.

The next experiments explore the consistency of the human-like behavior found in shallow autoencoders in a number scenarios. According to the results found in experiment 1, we select the two layer ReLU autoencoder as a representative example of shallow architecture with reasonable human-like behavior (RMSE of 11% of the maximum sensitivity) so we focus on this architecture in experiments 2 through 4.

## Experiment 2: Consistency of CSFs over a range of biodistortions

[Figure 8](#) shows the CSFs obtained when training the two layer ReLU net to compensate a range of retinal degradations (as described by the different pupil diameters and Fano factors). Learning curves that show the good convergence of the models and representative visual examples of the reconstructions are given in the [Appendix C](#) ([Figures 16](#) and [17](#)).

The results in [Figure 8](#) show that band pass / low-pass channels with distinct bandwidths consistently appear in all cases, and the RMSE with human CSF ( $25 \pm 2$ ), mean and standard deviation, stays in the low range of the values found in experiment 1.

Regarding the evolution of the CSFs with contrast, it is important to note that for some conditions (low blur and high noise) the gain in the achromatic CSF

increases with contrast, which is equivalent to contrast response curves that are not saturating.

## Experiment 3: CSFs from chromatic adaptation versus biodistortion compensation

[Figure 9](#) (top row) shows the CSFs emerging when the representative shallow network with human-like behavior in experiment 1 is trained for a range of alternative low-level tasks, some involving compensation of the retinal degradation (first, second, and third cases), and some others only involving chromatic adaptation (fourth and fifth). The corresponding learning curves for the models and visual examples of reconstruction are given in the [Appendix A](#) ([Figure 18](#)). [Table 2](#) (top) summarizes the numerical performance of the models in this experiment ( $\varepsilon_{LMS}$  for the computational goal, and RMSE for the CSFs).

First, let's focus on the case where the determination of the CSF faithfully follows [Equation 3](#) and hence we have a realistic eye+retina degradation (solid lines). Results show that only the cases where the task involves the biodegradation imply a clear difference in bandwidth between the achromatic and the chromatic channels. In the cases where there is only chromatic adaptation, the three CSFs are wider and of similar bandwidth. This behavior is clearly non-human, as confirmed by the RMSE measures in the fourth and fifth panels at the right.

Second, this difference is more clear in the idealized cases, [Equation 4](#), where clean sinusoids are used to determine the CSFs (dashed lines). In this situation, the CSFs of the purely chromatic goals are wider and flatter indicating that the networks are not performing any particular spatial modification in any chromatic channel. As a result, the RMSE values for the chromatic adaptation cases (light style numbers below the frequency axis) substantially increase indicating poorer description of human CSFs. In this regard, the errors for the cases in which the task involves biodegradation are lower, but they are even lower if the CSF is measured considering the realistic degradation in the input.

In summary, the results show two trends. On the one hand, human-like features emerge in the CSFs if the degradation–compensation task is considered, but they do not if only chromatic adaptation is considered. On the other hand, the CSFs are closer to human in RMSE if the determination takes the retinal degradation into account in the sinusoids.

Finally, there is an interesting human chromatic feature that is well-captured by all the CNN models that were trained for chromatic adaptation: all of them display a sort of Von-Kries modification of the red–green and yellow–blue channels. Note that, when the red illuminant has to be compensated (third and fifth cases), the red–green CSF is attenuated while the

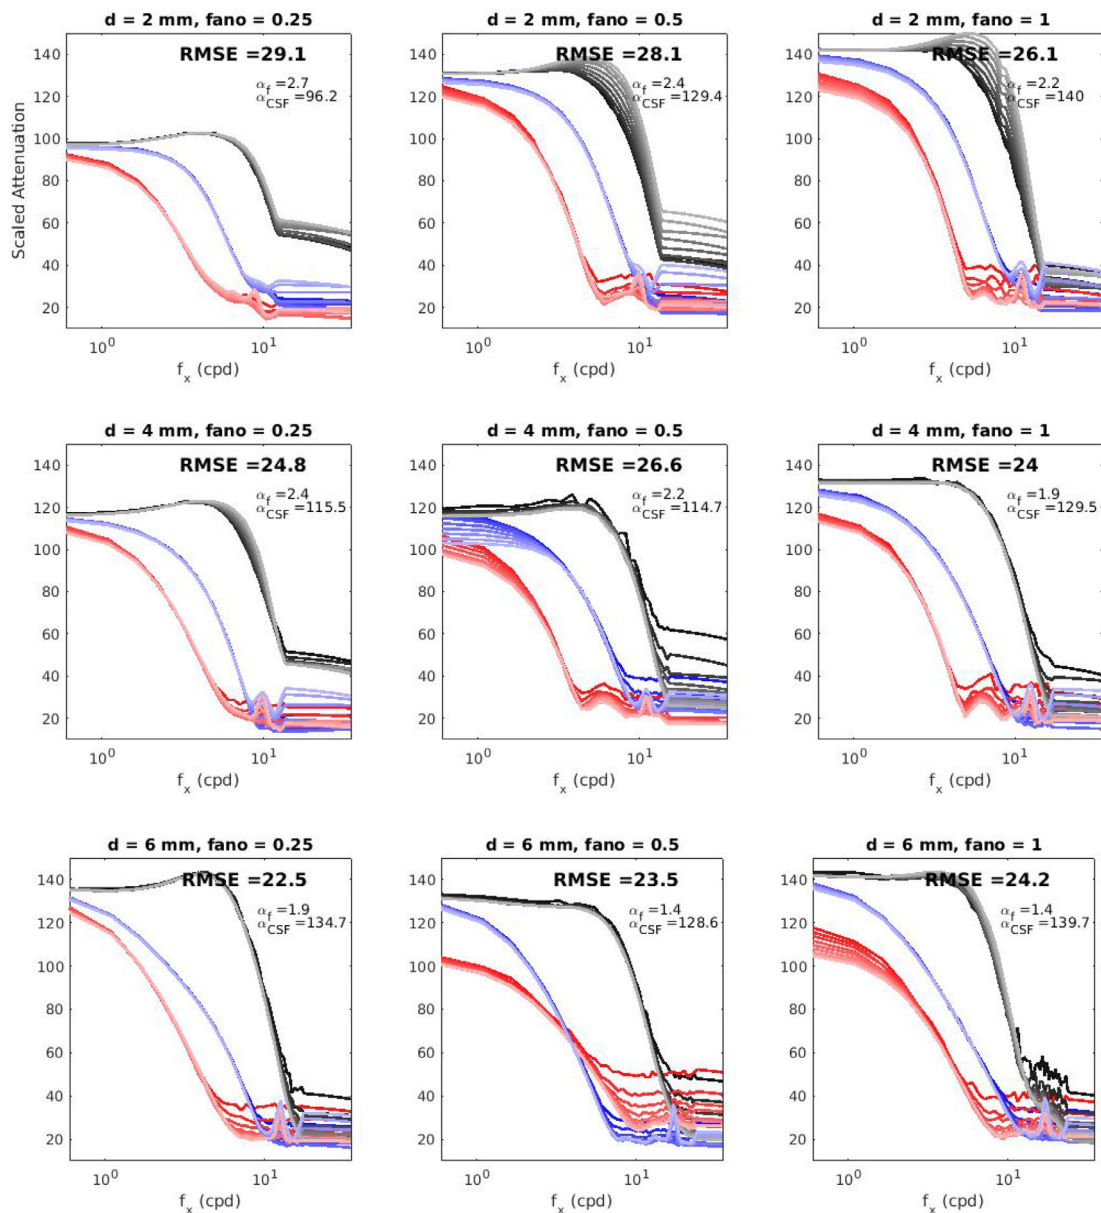


Figure 8. Experiment 2: Consistency of human-like result over a range of retinal degradations. Spatiochromatic CSFs of the two layer ReLU net for a range of retinal degradations. The RMSE distortion over the nine retinal conditions is  $25 \pm 2$  (mean and standard deviation). The line style conventions and meaning of the numbers is the same as in experiment 1.

blue–yellow CSF is boosted, and the other way around in in the compensation of a bluish illuminant (second and fourth cases, where the blue–yellow channel is attenuated).

#### Experiment 4: Consistency of spatiochromatic CSFs under changes of signal statistics

Figure 9 (bottom row) shows the CSFs emerging when the representative shallow network with human-like behavior in experiment 1 is trained for the range of low-level tasks considered in experiment 3 optimizing the performances over cartoon images

(as opposed to regular photographic images). The corresponding learning curves for the models and visual examples of reconstruction are given in the Appendix C (Figure 19). Table 2 (bottom) summarizes the numerical performance of the models in this experiment.

The parallelism in the results of experiments 3 and 4 confirms the robustness of the behaviors shown in experiment 3 to certain changes in signal statistics. Note that this parallelism does not mean that the CSFs are independent of the signal statistics. It just means that they are invariant to this change of statistics. It is important to remark that the low-level statistics of these (apparently different) sources may not be that different. Colors of the Pink Panther images are

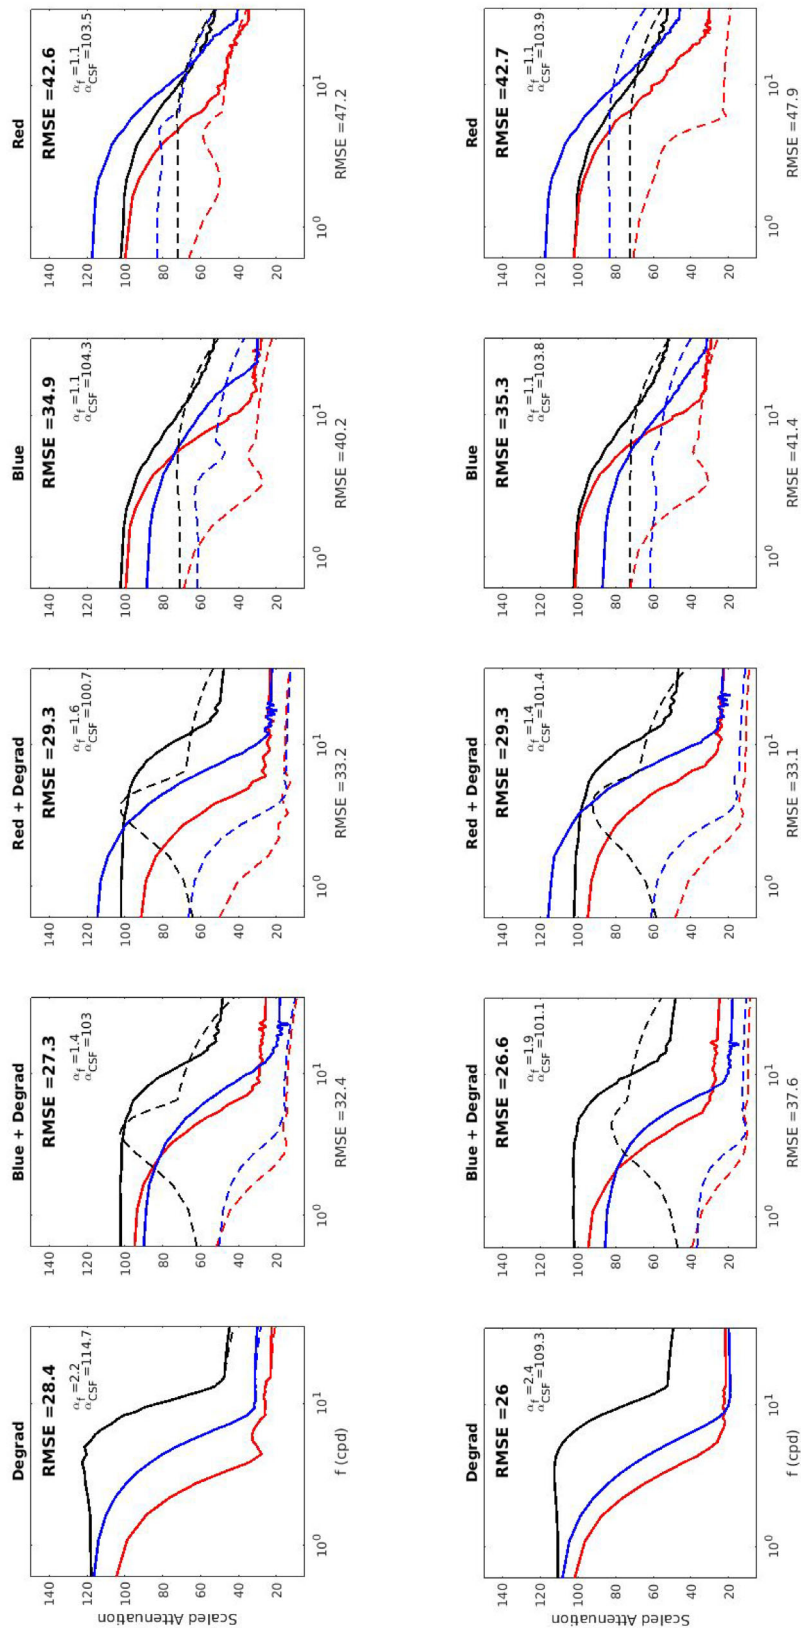


Figure 9. Experiment 3 (top row): CSFs from biodegradation compensation versus chromatic adaptation in natural images. From left to right, we consider the CSFs that emerge from (a) the compensation of biodegradation (left plot), (b) combinations of chromatic adaptation and degradation compensation (second and third plots), and (c) pure chromatic adaptation panels (fourth and fifth plots at the right). The cases in the columns first-to-third involve retinal degradation (with or without color adaptation) while the cases in the fourth and fifth columns only involve chromatic adaptation. For the sake of clarity, only low contrast results are shown. Solid lines

← correspond with CSFs determined using Equation 3 (where the input to the network includes realistic acquisition process). Dashed lines correspond with CSFs determined using Equation 4 (where the network is probed with clean sinusoids). The RMSE values in bold correspond with the CSFs determined in realistic conditions (curves in bold). The RMSE values displayed below the frequency axis correspond to the CSFs determined with clean sinusoids (dashed lines). In the cases involving chromatic adaptation the input sinusoids were color shifted according to the corresponding change in the illuminant. Experiment 4 (bottom row): Consistency of the results for different signal statistics (cartoon images). The computational goals are the same as above. The only difference is that the models are trained with cartoon images (from the Pink Panther show) as opposed to regular photographic images from ImageNet (used in experiments 1–3).

<b>Natural images</b>					
<b>Comput. goals</b>	Degradat. only	Degradat. + blue adapt.	Degradat. + red adapt.	Blue adapt.	Red adapt.
Original $\varepsilon_{LMS}$	12.1 ± 0.2	18.1 ± 0.2	15.2 ± 0.3	13.2 ± 0.2	10.2 ± 0.3
$\varepsilon_{LMS}$ after 2-layers ReLU	9.79 ± 0.08	8.7 ± 0.1	8.9 ± 0.1	1.92 ± 0.05	1.01 ± 0.01
<b>RMSE of CSFs</b>					
CSF <sub>realist.</sub>	28.4	27.3	29.3	34.9	42.6
CSF <sub>simplist.</sub>		32.4	33.2	40.2	47.2
<b>Cartoon Images</b>					
<b>Comput. Goals</b>	Degradat. Only	Degradat. + Blue Adapt.	Degradat. + Red Adapt.	Blue Adapt.	Red Adapt.
Original $\varepsilon_{LMS}$	12.9 ± 0.2	18.1 ± 0.1	15.6 ± 0.2	13.4 ± 0.3	9.1 ± 0.3
$\varepsilon_{LMS}$ after 2-layers ReLU	10.14 ± 0.08	9.5 ± 0.1	9.2 ± 0.1	1.34 ± 0.02	0.837 ± 0.06
<b>RMSE of CSFs</b>					
CSF <sub>realist.</sub>	26.0	26.6	29.3	35.3	42.7
CSF <sub>simplist.</sub>		37.7	33.1	41.4	47.9

Table 2. Experiment 3 (top). Compensation of biodegradation versus chromatic adaptation in natural images. Performance in the goals and eventual human-like CSFs are described by  $\varepsilon_{LMS}$  and RMSE, respectively. CSF<sub>realist.</sub> and CSF<sub>simplist.</sub> refer to the way the CSF is computed (taking into account or neglecting the retinal degradation in the sinusoidal stimuli). Experiment 4 (bottom). Consistency under change of image statistics. The considered goals and magnitudes have the same meaning as in experiment 3. The only difference is in the scenes used to train the models.

certainly more saturated, but beyond this obvious fact, other differences may be subtle. In particular, we took precautions to get frames from a 5-hour compilation where backgrounds around the whole chromaticity range appear not to bias the chromatic CSFs. Regarding the spatial content, note that there are plenty of edges of arbitrary orientations and also low-frequency transitions and shadows in the cartoons. More radical modifications of spatial information (e.g., edit the cartoon images to make them isoluminant; i.e., zero contrast in the achromatic channel) could lead to substantial variation of the CSFs, but the goal of this illustration is to point out the robustness of the result more than look for its limits.

### Experiment 5: CSFs from bottleneck compensation versus biodistortion compensation

Figure 10 shows the CSFs of the systems that emerge from the architectures with bottlenecks considered

in Figure 4 (right), when considering two different functional goals: 1) pure reconstruction of the input signal, that is, compensation of the information loss imposed by the bottleneck, and 2) compensation of the bottleneck together with compensation of the biodistortion. Table 3 summarizes the distortions in the CSFs, RMSE, and the performance in the reconstruction,  $\varepsilon_{LMS}$ .

The Appendix C, Figure 20, confirms that these architectures converged to a plateau of  $\varepsilon_{LMS}$ . Moreover, consistently with the data in Table 3, Figures 20 and 21 show that these systems achieve the computational goals to an extent that depends on the severity of the bottleneck in a very intuitive fashion (see comments in Appendix C).

More interesting is what happens to the emerging CSFs in Figure 10. In the absence of a bottleneck, pure reconstruction leads to wide filters equal in the three chromatic channels, a clearly non-human result with RMSE approximately 40 (architectures A and B). Similarly to pure chromatic adaptation, unconstrained pure reconstruction induces no spatial

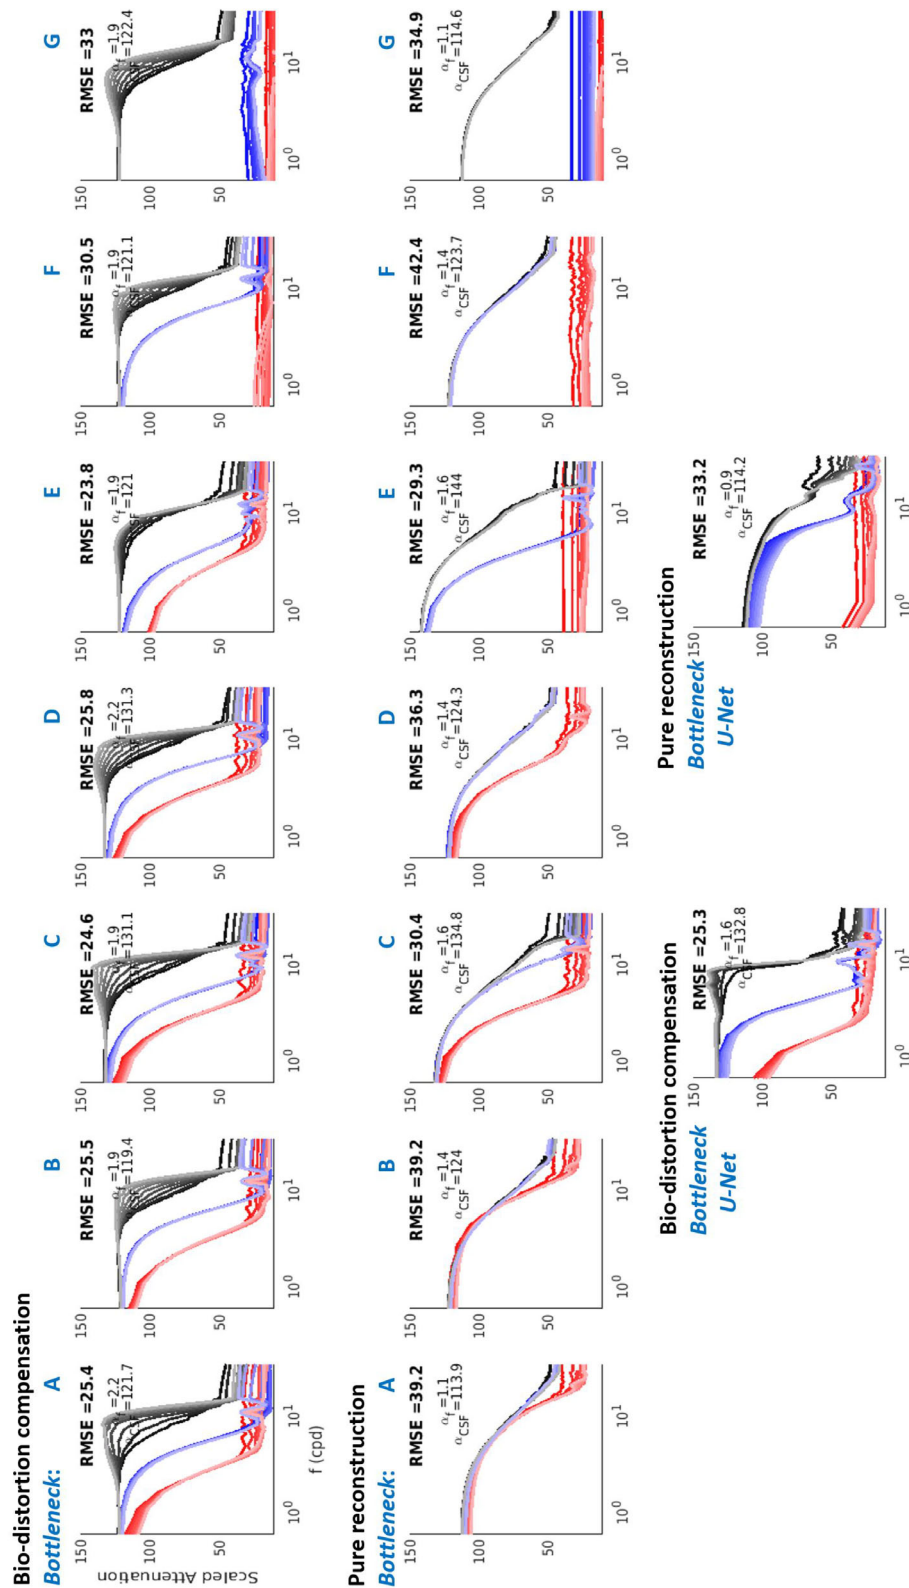


Figure 10. Experiment 5: CSFs from architectures with bottlenecks. Notation of the architectures with bottlenecks (letters in blue) refers to the diagrams shown in Figure 4. (Top row) CSFs emerging from bottleneck compensation and biodistortion degradation in progressively more restrictive bottlenecks in a four-layer architecture. (Middle row) CSFs emerging from pure reconstruction of the signal in the same architectures. (Bottom row) CSFs emerging in U-Nets from biodistortion compensation (left) or pure reconstruction (right).

	No bottleneck		Bottlenecks					
	A	B	C	D	E	F	G	U-Net
<b>Bio-Distort.</b>								
$\epsilon_{LMS}$	9.1 ± 0.3	9.2 ± 0.2	9.6 ± 0.1	9.8 ± 0.2	10.3 ± 0.4	11.2 ± 0.2	15.6 ± 0.5	12.3 ± 0.3
RMSE	25.4	25.5	24.6	25.8	23.8	30.5	33.0	25.3
<b>Pure Recons.</b>								
$\epsilon_{LMS}$	1.2 ± 0.1	2.2 ± 0.1	4.8 ± 0.3	3.0 ± 0.5	12.8 ± 0.7	5.6 ± 0.8	11.9 ± 0.5	12.2 ± 0.6
RMSE	39.2	39.2	30.4	36.3	29.3	42.4	34.9	33.2

Table 3. Experiment 5: Compensation of biodistortion versus pure reconstruction in networks with bottlenecks. The bottlenecks in the architectures A through G and U-Net are described in Figure 4 (right). Performance in the reconstruction goals is measured by the  $\epsilon_{LMS}$  error (in a test set) and the quality of the CSFs is given by the RMSE error. The considered biodistortion was the central case in Figure 3 with an original level of  $\epsilon_{LMS} = 15.5$ , which is reduced to the values reported in the first row after the application of the models. In the pure reconstruction case the original retinal distortion is  $\epsilon_{LMS} = 0$  so the errors reported in the 3rd row come from poor reconstruction or an incomplete compensation of the bottleneck.

selectivity and hence small similarity with human vision. Mild bottlenecks restricting the number of features and/or the spatial resolution do introduce differences in the bandwidth of the achromatic/chromatic channels, but the shape of the filters is far from human ( $RMSE \geq 30$  in architectures C and D). Then, more severe bottlenecks (architectures E–G and U-Net) quickly leads to over-attenuation of one or both chromatic CSFs (and hence non-human behavior with a RMSE of approximately 35 in these architectures for reconstruction). On the other hand, the very same architectures trained for the compensation of biodistortion lead to more human-like CSFs. See the band-pass/low-pass shape of the achromatic/chromatic CSFs and the RMSE of approximately 25, except for architectures F and G that overattenuate the chromatic CSFs but still preserve the band-pass nature of the achromatic channel. Better preservation of chromatic CSFs by the systems tuned to compensate the biodistortion is visually confirmed by the reconstructions of a representative image in Appendix C, Figure 21.

In summary, pure reconstruction with the explored bottlenecks induces a difference between the relative bandwidths of the achromatic and chromatic CSFs. However, the results become closer to human (both in the shape of the filters and in RMSE) when considering the compensation of the biodegradation of the retinal signal. And this resemblance remains even if the system is not constrained by a bottleneck.

## Experiment 6: Spatiochromatic–temporal CSFs from biodistortion compensation

Figure 11 shows the attenuation factors found for low-contrast moving sinusoids (both achromatic and chromatic) in the plane  $(f_x, f_t)$  for a range of 3D CNN autoencoders and for the optimal linear solution.

Experimental human CSFs for achromatic moving gratings (Kelly, 1979), and for chromatic moving gratings (Diez-Ajenjo et al., 2011) are also included as a useful reference. The learning curves for the models and visual examples of reconstructions are given in the Appendix C (Figure 22). Table 3 summarizes the numerical performance of the models in this experiment.

The CSF results show that the main feature of the spatiotemporal human window of visibility (its diamond shape), with smaller spatial bandwidth for higher temporal frequencies (or speeds) (Kelly, 1979; Watson, 2013) is reproduced by all the models as well as the substantially lower bandwidth of the chromatic channels, focused on very low spatiotemporal frequencies. The error of the best net is a RMSE of 17% of the maximum sensitivity.

Consistently with the results found in images (experiment 1), resemblance with human CSFs is bigger in shallower models (linear, two layers with a RMSE of approximately 17% or 18%, respectively) than in deeper models (six layers, and eight layers with a RMSE of approximately 22%) despite the performance of the deeper models in the goal is substantially better than the performance of the linear or the two layer model. The major differences are in the scaling of the chromatic CSFs: note that deeper models over attenuate the chromatic patterns. The RMSE measures confirm the superiority of the shallower solutions. For instance, note that the over attenuation of the red–green channel in the CNNs implies that the greenish hue of the background in the visual example of Figure 22 fades away, while it does not in the linear solution (which has obvious problems in other respects).

The linear solution cannot display a contrast dependent behavior, but the two layer architecture displays a consistent decay of the gain with contrast that is in line with the saturating nature of contrast response curves of humans for moving sinusoids

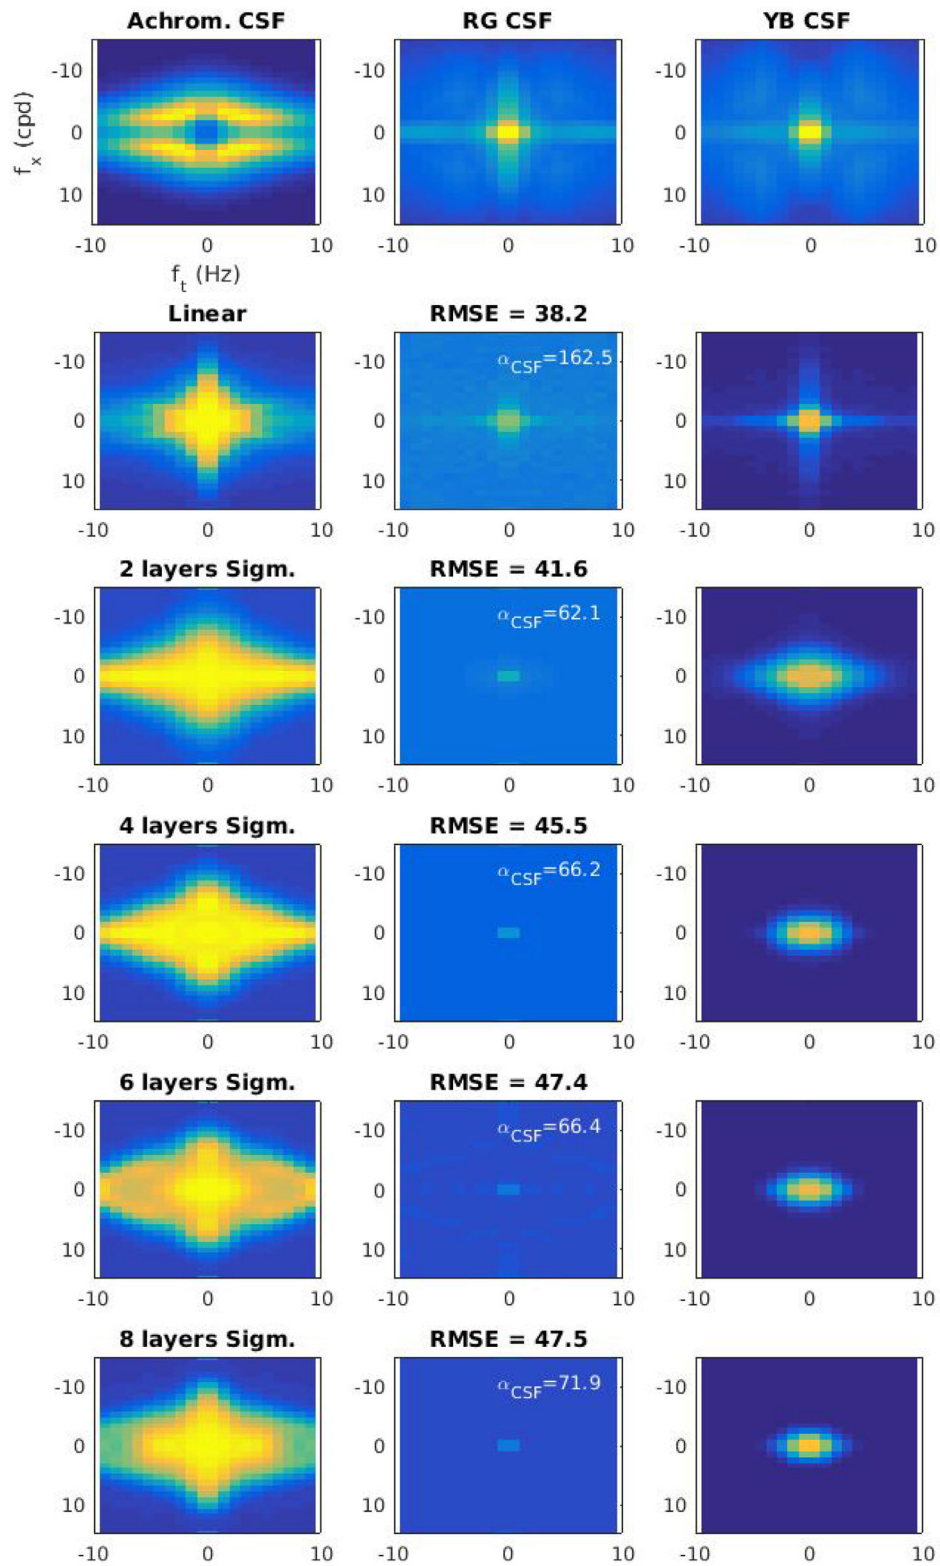


Figure 11. Experiment 6: Spatiotemporal chromatic CSFs from biodistortion compensation. The first row shows the human CSFs in  $(f_x, f_t)$ , and the following show the model CSFs. The RMSE numbers (average over channels) represent the distance between them. In this experiment,  $\alpha_f = 1$  in all cases.

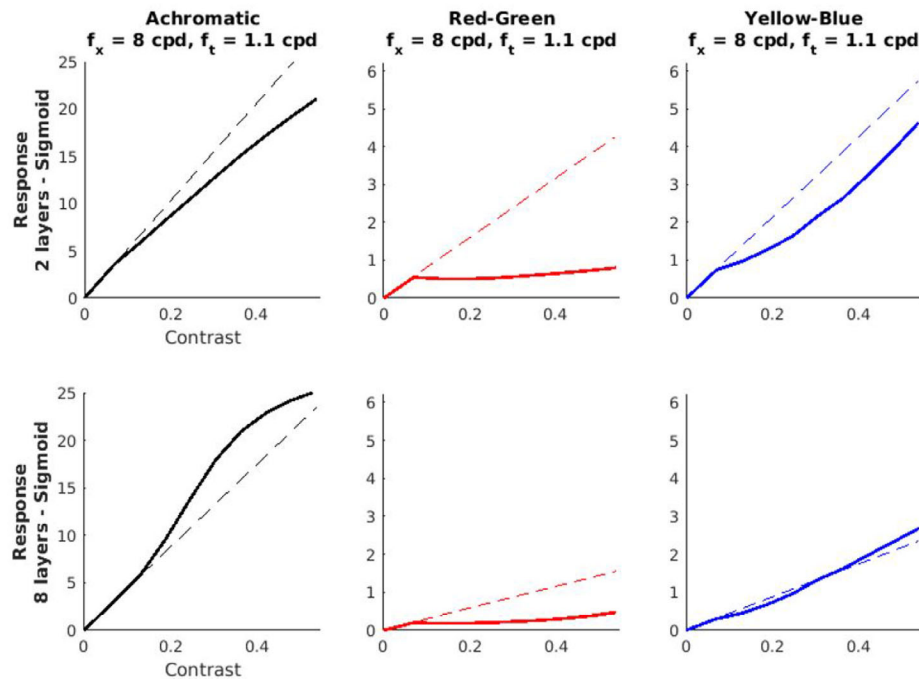


Figure 12. Experiment 6: Representative examples of nonlinear responses for spatiotemporal achromatic and chromatic gratings in shallow (top) and deeper (bottom) networks for biodistortion compensation.

	Comput. goal $\varepsilon_{LMS}$	CSFs RMSE
Distortion	$5.2 \pm 0.1$	
Linear Net	$3.5 \pm 0.2$	38.2
CNNs	Sigmoid nonlinear	Sigmoid nonlinear
2-Layers	$2.07 \pm 0.06$	41.6
4-Layers	$1.96 \pm 0.07$	45.5
6-Layers	$1.83 \pm 0.06$	47.4
8-Layers	$1.89 \pm 0.07$	47.5

Table 4. Experiment 6: Emergence of spatiotemporal chromatic CSF in 3D CNNs for compensation of biodistortion. The measures of the achievement of the goal  $\varepsilon_{LMS}$  and the distance with human behavior (RMSE difference between model and human CSF) have the same meaning as in the rest of experiments. The degradation included in the movies was the same as the considered for images (same pupil diameter and Fano factor). However, the numerical  $\varepsilon_{LMS}$  deviations turned out to be lower because the considered movies are *darker* and hence smaller LMS values lead to substantially smaller Poisson noise.

(Simoncelli & Heeger, 1998; Morgan et al., 2006). Figure 12 shows illustrative examples of these response curves: Although the two layer network (top row) consistently displays saturating behavior, the deeper net (bottom row) shows greater variability on the shape of the response.

As in the image case, the deviations in the chromatic CSFs in deep models do not come from not fulfilling the goal or having poor convergence in the training. First, all models (even the linear one) do reduce the error of the original retinal degradation so they are solving the computational problem. And second, the learning curves in the Appendix C (Figure 22) show that all models achieved a plateau in the training thus indicating proper convergence. Moreover, the asymptotic values achieved in the learning are consistent with the  $\varepsilon_{LMS}$  in test shown in Table 4.

## Discussion

### Summary of results

In these experiments, we trained a range of CNN autoencoders over natural scenes to solve different low-level vision goals: the compensation of retinal distortions, the compensation of changes in the illumination, the compensation of information loss after simple bottlenecks (or pure reconstruction after bottlenecks), and combinations of these.

Following the analysis of linearized networks presented in Gomez-Villa et al. (2020), it makes sense to stimulate these nets with achromatic, red–green and yellow–blue isoluminant sinusoids and moving sinusoids. The attenuation suffered by these gratings shows that:



- Human-like CSFs may emerge in systems that compensate retinal distortion: specifically, 2D shallow autoencoders trained to compensate retinal distortion display narrow low-pass behavior in the chromatic channels and wider band-pass behavior in the achromatic channel, so the shape and relative bandwidth of these artificial CSFs resemble those of humans (Figures 5, 7, and 8). Of course the match is not complete: the best CSFs obtained from the explored CNNs still deviate from human CSFs (RMSE of approximately 11% of the maximum sensitivity). Deeper autoencoders for the same goal also show CSFs with these basic shapes but the resemblance with human CSFs is consistently lower (RMSE of approximately 15% of the maximum sensitivity), particularly due to poor scaling of the chromatic CSFs (Figures 5 and 7).
- Artificial CSFs obtained from the compensation of retinal distortion differ from human CSFs in two qualitative aspects: a) The decay of network sensitivity found at low frequencies for achromatic gratings is not as big as in humans, and b) The relative amplitude of the red–green and the yellow–blue CSFs in the networks is inverted with regard to the humans. In our networks, the yellow–blue CSF is always bigger than the red–green CSF, and interestingly, this is pretty consistent over different architectures and datasets with different image statistics.
- Similar sensitivities consistently appear in shallow autoencoders for a range of levels in retinal distortions (Figure 8).
- Human-like CSFs with distinct bandwidths in achromatic/chromatic channels do not appear in pure chromatic adaptation tasks, but they do as soon as the retinal distortion compensation goal is considered (with or without chromatic adaptation). The compensation of chromatic shifts together with the compensation of biodistortion leads to systems in which the chromatic CSFs change their global gain similarly to a Von-Kries mechanism (Figure 9, top).
- CSFs emerging from chromatic adaptation and degradation compensation goals are similar for natural images and cartoon images (Figure 9, bottom).
- Pure reconstruction in architectures with a restrictive bottleneck induces changes in the relative bandwidths of the achromatic and chromatic CSFs with regard to trivial all-pass filters found in systems without bottleneck. However (in the explored cases), these CSFs are remarkably non-human. Interestingly, the very same architectures lead to more human-like CSFs as soon as the retinal distortion compensation goal is considered (Figure 10).

- The 3D autoencoders for retinal degradation compensation display a wide diamond-shaped achromatic bandwidth and very narrow chromatic bandwidths in the spatiotemporal Fourier domain, in parallel with humans. And this similarity is larger in the linear and shallow autoencoders (RMSE of approximately 17% of the maximum sensitivity) while it decays for deeper networks (RMSE of approximately 22% of the maximum sensitivity), again owing to poor scaling of the chromatic CSFs (Figure 11).
- The gain in shallow autoencoders decays with contrast and hence the contrast responses for gratings saturate with contrast. This happens both in the spatial and the spatiotemporal cases (Figures 6 and 12). This resembles contrast masking in humans. However, in deeper autoencoders this consistent saturation (and hence similarity with humans) is not found.

The emergence of human-like features in the CSFs (distinct bandwidth and shape of achromatic and chromatic channels) is related to the different properties of achromatic and chromatic patterns in visual scenes. The statistical unbalance towards achromatic patterns is known from long ago in terms of variance (Ruderman et al., 1998) and, more recently, it has been quantified in accurate information theoretic units (Malo, 2020). The eventual problems in preserving the saturation (or poor scaling of chromatic CSFs) in deeper models, do not come from training. Note that, according to the learning curves, all the models achieved proper convergence. On the contrary, the problems may come from the small (statistical) relevance of chromatic textures as opposed to the achromatic textures and the inability of deeper models to deal with this unbalance with a low-level  $\varepsilon_{LMS}$  goal: (too) flexible networks optimized to compensate the distortions focus (too much) on the spatial achromatic information to optimize the goal and are likely to distort chromatic information. The consequence is a negative impact on the chromatic CSFs. This does not seem to be a problem for more rigid shallower architectures and even the linear solution.

At this low abstraction level, where the minimization of distortion in LMS is simply connected with information maximization, and in the set of architectures considered, shallow networks seem more appropriate to explain the CSFs.

## Relation to other accounts of the CSFs

Our results revisit classical work on the statistical grounds of the CSFs (Atick et al., 1992; Atick & Redlich, 1992; Atick, 2011) in light of the new possibilities provided by automatic differentiation.

From the technical point of view, a number of assumptions that had to be done in the 1990s, either have been confirmed with the use of large data sets, or are not necessary with the use of more flexible models. In particular, regarding the signal, Atick et al. assumed translation invariance, independence between color and space-time, and second-order relations (autocorrelation with  $1/|f|^2$  decay). Moreover, regarding the model, they restricted themselves to linear solutions similar to Wiener filters. More recent studies with colorimetrically calibrated scenes (Gutmann et al., 2014) have confirmed the correctness of the color/space independence assumption. However, the focus on the power spectrum and the linear solutions has proven to be too limited for denoising (Gutiérrez et al., 2006). Adaptive (nonlinear) models that take into account additional features of the signal are required. Nevertheless, the nonlinear networks considered here turn out to be roughly translation invariant (Gomez-Villa et al., 2020), as expected from their convolutional architectures and the stationary nature of the problems they face. Another technical difference is in the formulation of the statistical goal: Atick et al. maximized the mutual information between the clean signal and the response  $I(\mathbf{x}_c, \mathbf{y})$ ; while here we minimized  $\varepsilon_{LMS}$  between the clean signal and the response,  $|\mathbf{X}_c - \mathbf{Y}|_2$ . These goals are exactly equivalent when the difference between clean signal and the response is Gaussian, which is not the case in general. However, note that these goals are always related because the limit  $|\mathbf{X}_c - \mathbf{Y}|_2 \rightarrow 0$  implies  $I(\mathbf{x}_c, \mathbf{y}) \rightarrow \infty$ . Beyond the spatiochromatic case, in our work we check the emergence of the CSFs with spatiotemporal signals, which was mentioned but not addressed by Atick et al. Finally, the consideration of nonlinear models allows us to show that the error-minimization goal may also lead to saturation of the contrast responses that, of course, was not possible in the linear framework of Atick et al.

As stated in the Introduction, other group has been working independently on the emergence of CSFs in artificial neural networks (Akbarinia et al., 2021). Their results are restricted to the spatial CSF (no chromatic nor spatiotemporal cases) and are based on networks trained for higher level goals, such as classification. Therefore, their results from higher level goals are complementary to ours, obtained from lower-level goals intended for the analysis of early visual stages such as the LGN. More generally, higher level goals such as classification performance may be an indirect way to impose preservation of colorfulness or a proper scaling of the chromatic CSFs. Although chromatic information may have small relevance to minimize  $\varepsilon_{LMS}$  in reconstruction, it may be more crucial for recognition.

Other works have obtained center surround sensors by optimizing a linear+nonlinear network with a

low-level infomax+energy goal (Karklin & Simoncelli, 2011), or deeper nets with higher level classification goals (Lindsey et al., 2019). These sensors could induce CSF-like bandwidths in the corresponding models but this aspect was not addressed in these works.

## Individual non-Euclidean distances from the optimization of average Euclidean distance

An interesting consequence of our low-level result is that the Euclidean measure,  $\varepsilon_{LMS}$ , averaged over the set of natural images leads to systems that measure individual differences in non-Euclidean ways. Note that in the systems that we trained given two input signals  $\mathbf{x}$  and  $\mathbf{x} + \Delta\mathbf{x}$ , the difference between the corresponding responses is  $\Delta\mathbf{y} \approx \mathbf{M} \cdot \Delta\mathbf{x}$ . As a network should assess the difference between the two signals from  $\Delta\mathbf{y}$ , the perceived difference for the system will depend on  $\mathbf{M}$ . Specifically (Epifanio et al., 2003; Laparra et al., 2010), the perceived distance for the system will depend on the metric,  $\mathbf{M}^T \cdot \mathbf{M}$ , and hence it will depend on  $\lambda^2$  (the eigenspectrum of  $\mathbf{M}$ ) or, as seen here, on the CSF<sup>2</sup>. This metric is non-Euclidean: for instance, high-frequency distortions will be less relevant for the network than medium or low-frequency distortions. Even though here we did not check the correlation between the image distortions perceived by humans and networks, the observed CSFs in the networks are consistent with the fact that the Euclidean distance between images at the retina is not a good representation of human distortion metrics (Wang & Bovik, 2009; Laparra et al., 2010; Hepburn et al., 2020).

The emergence of a non-Euclidean distance from the minimization of the average Euclidean distance over natural images is a counterintuitive consequence of the highly nonuniform distribution of natural scenes: distortions in less populated regions of the image space (e.g., in high-frequency directions or in chromatic channels) have to be under-rated to favor the average match to the data in highly populated or more informative regions (low-frequency, achromatic patterns).

Recent work on autoencoders with low-level rate-distortion constraints on natural images has shown the emergence of non-Euclidean distances correlated with human opinion of distortion (Hepburn et al., 2022). Human opinion of distortion is known to be strongly mediated by the CSFs, but the bandwidth of this autoencoder and its eventual similarity with the CSF was not explored in that work.

## Alternative low-level computational goals

Here we considered different low-level alternatives to the retinal signal enhancement goal proposed by Atick et al.: although our results are conclusive regarding the

role of chromatic adaptation, more research is definitely required about the relative relevance of bottlenecks in shaping the CSFs.

On the one hand, given the small role of spatial information in the changes of the LMS image purely owing to changes in illumination, it is not surprising that systems designed for pure chromatic adaptation have wide (all-pass) behavior in all channels (i.e., no spatial effect). As a consequence (as confirmed by our results), pure chromatic adaptation does not lead to CSFs with a human-like shape. This shape and relative bandwidths have to be related to other goals (e.g., the compensation of biodistortion). However, training for chromatic adaptation does introduce an important human-like behavior (which may not emerge from other tasks): it leads to adaptive global scaling of the red–green or yellow–blue CSFs. This effect is consistent with the observations done on spatiochromatic adaptation under changes in spectral illumination (Gutmann et al., 2014): the spatial structure of the receptive fields remains almost constant but their chromatic tuning basically changes according to Von-Kries adaptation.

On the other hand, as opposed to chromatic shifts, the spatial effect of bottlenecks is relevant. However, we only explored a small range of architectural bottlenecks: the toy examples of Figure 4-right. In this restricted set our results suggest that the compensation of the biodegradation at the retina plays a stronger role in the emergence of human-like CSFs than the consideration of the bottlenecks. However, bottlenecks in architectures C, D, and U-Net favor the emergence of nontrivial frequency selectivity. This would be consistent with Lindsey et al. (2019), who reported positive effects of bottlenecks in the emergence of center surround receptive fields. Nevertheless, the specific configuration of the bottleneck that maximizes the human nature of the CSFs and the relative role of bottlenecks in the compensation of the retinal distortion are interesting matters for further research.

More generally, other low-level goals could be considered together with the distortion, as for instance the information or the energy of the signal. Architectural bottlenecks considered here or in Lindsey et al. (2019) indirectly constrain the energy and the entropy of the signal by reducing the dimensionality of the signal. However, one could consider more general factors beyond the dimensionality as for instance the neural noise, the PDFs of signal and noise and the redundancy of the visual signals in the representation. In fact, transmitted information may be modulated by changes of the representation and by the amount of noise even without changes in the dimensionality (Malo, 2020).

In a separate study (Hepburn et al., 2022) we have shown that rate-distortion bottlenecks in autoencoders induce distance measures which are correlated with

subjective opinion of distortion. The autoencoders we presented here do not include constraints on information, but the emergence of a non-Euclidean metric depending on  $M$  (and hence on the CSFs) suggests that the distance will be correlated with human opinion in line with Hepburn et al. (2022).

Alternative low-level goals could include non-human retinal degradation. Other species have different optical quality and noise in their retinas may be substantially different. This may affect the kind of computations required to extract the appropriate information from this degraded input, and hence their corresponding CSFs.

All these issues, the specific impact of more sophisticated bottlenecks in the CSFs, which was not analyzed here or in Karklin and Simoncelli (2011), Lindsey et al. (2019), or Hepburn et al. (2022), the emergence of human-like image distortion measures from the enhancement of retinal signals, and the consideration of retinal degradation for other species, is a matter for future research.

## Goal and architecture are not independent

More important than the technical generalizations over Atick et al. (1992), Atick and Redlich (1992), and Atick (2011), is that the current freedom to explore different linear and nonlinear architectures stresses the relevance of the architectural constraints. The conventional interpretation of the efficient coding hypothesis (Barlow et al., 1961) is the following: obtaining human-like results from certain statistical goal seems to suggest that the human visual system has been shaped by this goal. However, it is important to realize that the results have been obtained via the optimization of certain model. In the case of Atick et al., it was a single model (the linear filter), but in our case here we tried a range of models (architectures). Because the results for the different architectures is markedly different, the conclusion can not be about the goal, but about specific combinations of goal and architecture. Our results are a specific illustration of the fact that the computational and the algorithmic levels of analysis of visual processing systems (Marr & Poggio, 1976; Marr, 1982) are not independent (Poggio, 2021). This dependence prevents about premature conclusions about the organization principles at the computational level if sensible architectures are not adopted.

## Beyond accuracy

Human-like CSFs are obtained for shallow autoencoders (two layers), or even linear networks, despite deeper architectures achieving similar or better performance in the goal. The previous literature has warned about the limitations of a single

accuracy/performance measure to identify human-like behavior. Achieving similar performance on a task does not guarantee that two models actually use the same strategy (Firestone, 2020). For instance, different strategies may become evident if performance degrades in different ways when changing the experimental setting (Wichmann et al., 2017; Geirhos et al., 2019). Therefore, additional checks different from the optimization goal have to be done in order to confirm the human-like behavior of a model. Examples include verifying additional psychophysics not included in the goal (Martinez et al., 2019), or disaggregating the results checking the consistency between model and humans in individual trials, not on averages over the data set (Geirhos et al., 2020).

In this complexity/accuracy discussion, it is important to stress that our results (shallower networks better reproduce the scale of human chromatic CSFs) is in line with the results of Gomez-Villa et al. (2020) and Flachot et al. (2020), which also show that shallower networks obtain more human-like colour representation. In a similar vein, although using higher level classification goals, Kubišius et al. (2019) show that lower performance networks may correlate better with human brain activity or psychophysics.

## Final remarks

In visual neuroscience, deep models are emerging as the new standard to reproduce the activity of visual areas under natural scene stimulation. On the one hand, conventional deep models driven by object recognition goals reproduce the response from V1 (Güçlü and van Gerven, 2015; Kriegeskorte, 2015), dorsal and ventral streams (Cichy et al., 2016), and IT (Cadieu et al., 2014; Yamins et al., 2014). On the other hand, deep networks are powerful enough to fit the mappings between stimuli and measured responses (Prenger et al., 2004; Antolík et al., 2016; Batty et al., 2017). These two approaches (goal-driven and measurement-driven deep models) have been thoroughly compared in V1 and were found to be superior to linear filter-banks and simple linear–nonlinear models (Cadena et al., 2019). However, more recently, the same team has shown that linear–nonlinear models with general divisive normalization make a significant step towards the performance of state-of-the-art CNN with interpretable parameters (Burg et al., 2021).

In our low-level goal-driven case, the emergence of human-like CSFs for certain CNN autoencoders generalizes in different ways previous statistical explanations of the CSF based on linear models (Atick & Redlich, 1992; Atick et al., 1992; Atick, 2011), and is consistent with optimizations of nonlinear encoders using alternative low-level (Karklin & Simoncelli, 2011; Hepburn et al., 2022) or higher level (Lindsey et al., 2019; Akbarinia et al., 2021) goals. However, we find a

strong dependence of the CSFs on the architecture with better results for shallower autoencoders (although they have similar or lower performance in the goal).

This is not in contradiction with the literature cited elsewhere in this article showing that deep networks with object recognition goals match very well higher visual areas. Note that the scope of our low-level goal is restricted to early visual stages (e.g., the retina–LGN path), and hence simpler architectures may be required there.

Beyond this difference in abstraction level, our results do illustrate the relevance of using appropriate architectures when checking a statistical goal. Following the move from conventional CNNs in Cadena et al. (2019) to more realistic divisive normalization models in Burg et al. (2021), we think that future goal-driven derivations of low-level visual psychophysics (e.g., pattern masking or perceptual distortion) should include more realistic architectures too, as opposed to conventional CNNs (although they may be flexible enough to fulfill the goal). Examples include divisive normalization with parametric interaction between features (Martinez et al., 2018, 2019) and generalizations of Wilson–Cowan interactions (Bertalmío et al., 2020). Learning frameworks with rate-distortion bottlenecks are already available (Ballé et al., 2017; Hepburn et al., 2022), and we advocate for the study of their artificial psychophysics using realistic and interpretable architectures.

*Keywords: spatiotemporal and chromatic contrast sensitivity, convolutional autoencoders, modulation transfer function, noisy cones, deblurring and denoising, chromatic adaptation, natural images, statistical goals, architectures*

## Acknowledgments

Partially funded by these grants from GVA/AEI/FEDER/EU: MICINN DPI2017-89867-C2-2-R, MICINN PID2020-118071GB-I00, and GVA Grisolia-P/2019/035 (for JM and QL), and MICINN PGC2018-099651-B-I00 (for A.G.V. and M.B.).

Commercial relationships: none.

Corresponding author: Jesus Malo.

Email: [jesus.malo@uv.es](mailto:jesus.malo@uv.es).

Address: Image Processing Lab, Building E4, Parc Científic de la Universitat de València, Carrer Catedràtic Escardino, 46980 Paterna, Valencia, Spain.

## Footnotes

<sup>1</sup>For example, neurons in conventional CNNs have fixed nonlinearities, as opposed to the known adaptive nature of real neurons (Wilson & Cowan, 1973; Carandini & Heeger, 2012).

<sup>2</sup>For optical blur where the linear operator  $H$  can be obtained from the MTF (Watson, 2013), and the retinal noise is Poisson,  $\mathbf{n}_r = F \cdot \mathbb{D}_{(|H \cdot x|^{\frac{1}{2}})} \cdot \mathbf{n}$ , where  $\mathbb{D}_v$  is a diagonal matrix with vector  $\mathbf{v}$  in the diagonal,  $F$  is the Fano factor, and  $\mathbf{n}$  is drawn from a unit-variance Gaussian (Esteve et al., 2020); the Jacobian in Equation 9, is  $\nabla_x S_\theta(\mathbf{0}) = \nabla_x N_\theta(\mathbf{0}) \cdot (I - \frac{F}{2} \cdot \mathbb{D}_{(|H \cdot \mathbf{0}|^{\frac{1}{2}})}) \cdot H$ , where the Jacobian of the network,  $\nabla_x N_\theta(\mathbf{0})$ , can be obtained analytically (Martinez et al., 2018).

<sup>3</sup>We prepared the samples that way before actually knowing how well the networks are able to cope with this distortion.

<sup>4</sup>In our computer cluster typical training of the 2D models takes approximately 10 to 20 hours.

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., . . . Zheng, X. (2016). TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv Comp. Sci.* <https://arxiv.org/abs/1603.04467>.
- Akbarinia, A., Morgenstern, Y., & Gegenfurtner, K. R. (2021). Contrast sensitivity is formed by visual experience and task demands. *Journal of Vision*, 21(9), 1996.
- Antolik, J., Hofer, S. B., Bednar, J. A., & Mrsic-Flogel, T. D. (2016). Model constrained by visual hierarchy improves prediction of neural responses to natural scenes. *PLoS Computational Biology*, 12(6), e1004927, doi:10.1371/journal.pcbi.1004927.
- Atick, J., Li, Z., & Redlich, A. (1992). Understanding retinal color coding from first principles. *Neural Computation*, 4(4), 559–572.
- Atick, J., & Redlich, A. (1992). What does the retina know about natural scenes? *Neural Computation*, 4(2), 196–210.
- Atick, J. J. (2011). Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 22, 4–44.
- Ballé, J., Laparra, V., & Simoncelli, E. P. (2017). End-to-end optimized image compression. *International Conference on Learning Representations (ICLR)*, <https://openreview.net/forum?id=rJxdQ3jeg>.
- Barlow, H. B. et al. (1961). Possible principles underlying the transformation of sensory messages. *Sensory Communication*, 1, 217–234.
- Batty, E., Merel, J., Brackbill, N., Heitman, A., Sher, A., Litke, A. M., . . . Paninski, L. (2017). Multilayer recurrent network models of primate retinal ganglion cell responses. *5th International Conference on Learning Representations (ICLR)*, <https://openreview.net/forum?id=9gmuVOIKfLa>.
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018). Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18(153), 1–43.
- Berardino, A., Ballé, J., Laparra, V., & Simoncelli, E. (2017). Eigen-distortions of hierarchical representations. In *Proceedings of the Neural Information Processing Systems*, 30, pp. 3533–3542, <https://papers.nips.cc/book/advances-in-neural-information-processing-systems-30-2017>.
- Bertalmío, M., Gomez-Villa, A., Martín, A., Vazquez, J., Kane, D., & Malo, J. (2020). Evidence for the intrinsically nonlinear nature of receptive fields in vision. *Scientific Reports*, 10, 16277.
- Burg, M., Cadena, S., Denfield, G., Walker, E., Tolia, A., & Bethge, M. et al. (2021). Learning divisive normalization in primary visual cortex. *PLoS Computational Biology*, 17(6), e1009028.
- Cadena, S., Denfield, G., Walker, E., Gatys, L., Tolia, A., & Bethge, M. et al. (2019). Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLoS Computational Biology*, 15(4), e1006897.
- Cadiou, C., Hong, H., Yamins, D., Pinto, N., Ardila, D., & Solomon, E. et al. (2014). Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS Computational Biology*, 10(12), e1003963.
- Cai, D., DeAngelis, G. C., & Freeman, R. D. (1997). Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *Journal of Neurophysiology*, 78(2), 1045–1061.
- Campbell, F., & Robson, J. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, 197, 551–566.
- Carandini, M., & Heeger, D. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62.
- Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, 6, 27755.
- Clarke, R. (1981). Relation between the Karhunen Loève and cosine transforms. *IEE Proceedings F Communications, Radar and Signal Processing*, 128, 359–361.
- Cottaris, N. P., Jiang, H., Ding, X., Wandell, B. A., & Brainard, D. H. (2019). A computational-observer model of spatial contrast sensitivity: Effects of wave-front-based optics, cone-mosaic structure, and inference engine. *Journal of Vision*, 19(4), 8.
- Cottaris, N. P., Wandell, B. A., Rieke, F., & Brainard, D. H. (2020). A computational observer model of spatial contrast sensitivity: Effects of photocurrent encoding, fixational eye movements, and inference engine. *Journal of Vision*, 20, 17.

- Donen, S. (1963). *Charade*. Universal Studios, CC Public Domain Mark 1.0, <https://archive.org/details/Charade19631280x696>.
- Díez-Ajenjo, M., Capilla, P., & Luque, M. J. (2011). Red-green vs. blue-yellow spatio-temporal contrast sensitivity across the visual field. *Journal of Modern Optics*, *58*, 1–13.
- Enroth-Cugell, C., & Robson, J. (1966). The contrast sensitivity of retinal ganglion cells on the cat. *Journal of Physiology (London)*, *187*, 516–552.
- Epifanio, I., Gutierrez, J., & Malo, J. (2003). Linear transform for simultaneous diagonalization of covariance and perceptual metric matrix in image coding. *Pattern Recognition*, *36*(8), 1799–1811.
- Esteve, J., Aguilar, G., Maertens, M., Wichmann, F., & Malo, J. (2020). Psychophysical estimation of early and late noise. *Arxiv: Quantitative Biology*, 1–15, <https://arxiv.org/abs/2012.06608>.
- Firestone, C. (2020). Performance vs. competence in human–machine comparisons. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(43), 26562–26571.
- Flachot, A., Akbarinia, A., Schütt, H. H., Fleming, R. W., Wichmann, F. A., & Gegenfurtner, K. R. (2020). Deep neural models for color discrimination and color constancy. *CoRR*, *abs/2012.14402*.
- Freleng, I. (1963). *The Pink Panther Show*. Los Angeles, CA, USA: DePatie-Freleng Enterprises (DFE Films), courtesy of Metro Goldwin Mayer.
- Geirhos, R., Meding, K., & Wichmann, F. A. (2020). Beyond accuracy: Quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. *arXiv: 2006.16736*.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2019). Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *International Conference on Learning Representations (ICLR)*, <https://openreview.net/forum?id=Bygh9j09KX>.
- Gomez-Villa, A., Martin, A., Vazquez, J., & Bertalmío, M. (2019). Convolutional neural networks can be deceived by visual illusions. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 19)* (pp. 12309–12317), doi:10.1109/CVPR.2019.01259.
- Gomez-Villa, A., Martin, A., Vazquez, J., Bertalmío, M., & Malo, J. (2020). Color illusions also deceive CNNs for low-level vision tasks: Analysis and implications. *Vision Research*, *176*, 156–174.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. Boston, MA, USA: MIT Press, <https://www.deeplearningbook.org>.
- Güçlü, U., & van Gerven, M. A. J. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, *35*(27), 10005–10014.
- Gutiérrez, J., Ferri, F. J., & Malo, J. (2006). Regularization operators for natural images based on nonlinear perception models. *IEEE Transactions on Image Processing*, *15*(1), 189–200.
- Gutmann, M. U., Laparra, V., Hyvärinen, A., & Malo, J. (2014). Spatio-chromatic adaptation via higher-order canonical correlation analysis of natural images. *PloS One*, *9*(2), e86481.
- Hepburn, A., Laparra, V., Malo, J., & Santos, R. (2020). Perceptnet: A human visual system inspired neural network for estimating perceptual distance. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (pp. 121–125), doi:10.1109/ICIP40778.2020.9190691.
- Hepburn, A., Laparra, V., Santos, R., Ballé, J., & Malo, J. (2022). On the relation between statistical learning and perceptual distances. *International Conference on Learning Representations, ICLR*, [https://openreview.net/forum?id=zXM0b4hi5\\_B](https://openreview.net/forum?id=zXM0b4hi5_B).
- Hunt, B. R. (1975). Digital image processing. *Proceedings of the IEEE*, *63*, 693–708.
- Hurvich, L., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, *64*(6), 384–404.
- Hyvärinen, A., Hurri, J., & Hoyer, P., 2009. *Natural image statistics: A probabilistic approach to early computational vision*. Heidelberg, Germany: Springer-Verlag.
- Ingling, C. R., & Martinez-Uriegas, E. (1983). The relationship between spectral sensitivity and spatial sensitivity for the primate r-g x-channel. *Vision Research*, *23*, 1495–1500.
- Karklin, Y., & Simoncelli, E. (2011). Efficient coding of natural images with a population of noisy linear-nonlinear neurons. In *Proceedings of the Advances in Neural Information Processing Systems*, *24*, <https://proceedings.neurips.cc/paper/2011/file/12e59a33dea1bf0630f46edfe13d6ea2-Paper.pdf>.
- Kelly, D. H. (1979). Motion and vision. ii. stabilized spatio-temporal threshold surface. *Journal of the Optical Society of America*, *69*(10), 1340–1349.
- Kingma, D. P., & Ba, J. (2017). Adam: A method for stochastic optimization. *arXiv:1412.6980*.
- Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, *1*(1), 417–446.
- Kubilius, J., Schrimpf, M., Kar, K., Rajalingham, R., Hong, H., Majaj, N., . . . DiCarlo, J. J. (2019). Brain-like object recognition with high-performing

- shallow recurrent ANNs. In *Proceedings Advances in Neural Information Processing Systems*, 32, <https://proceedings.neurips.cc/paper/2019/file/7813d1590d28a7dd372ad54b5d29d033-Paper.pdf>.
- Laparra, V., Muñoz, J., & Malo, J. (2010). Divisive normalization image quality metric revisited. *Journal of the Optical Society of America A*, 27(4), 852–864.
- Legge, G. E. (1981). A power law for contrast discrimination. *Vision Research*, 21(4), 457–467.
- Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, 70(12), 1458–1471.
- LeRoy, M. (1959). *The FBI story*. Los Angeles, CA, USA: Warner Brothers/Seven Arts.
- Lillicrap, T., Santoro, A., Marris, L., Akerman, C., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6), 335–346.
- Lindsey, J., Ocko, S. A., Ganguli, S., & Deny, S. (2019). The effects of neural resource constraints on early visual representations. *International Conference on Learning Representations, ICLR*, <https://openreview.net/forum?id=S1xq3oR5tQ>.
- Malo, J. (2020). Spatio-chromatic information available from different neural layers via gaussianization. *Journal of Mathematical Neuroscience*, 10(18), doi:10.1186/s13408-020-00095-8.
- Malo, J. (2022). Paraphrasing Magritte’s observation. *ArXiv: Computer Vision and Pattern Recognition*, 1–4, <https://arxiv.org/abs/2202.08103>.
- Malo, J., & Luque, M. (2002). ColorLab: A Matlab Toolbox for color science and calibrated color image processing. *Servei de Publicacions de la Universitat de Valencia*. Valencia, Spain, <https://isp.uv.es/code/visioncolor/colorlab.html>.
- Malo, J., Pons, A., Felipe, A., & Artigas, J. (1997). Characterization of the human visual system threshold performance by a weighting function in the gabor domain. *Journal of Modern Optics*, 44(1), 127–148.
- Mannos, J. L., & Sakrison, D. J. (1974). The effects of a visual fidelity criterion of the encoding of images. *IEEE Transactions on Information Theory*, 20, 525–536.
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman & Co.
- Marr, D., & Poggio, T. (1976). *From understanding computation to understanding neural circuitry (AI Memo No. AIM-357)*. Boston, MA, USA: MIT Libraries, <http://hdl.handle.net/1721.1/5782>.
- Martinez, M., Bertalmio, M., & Malo, J. (2019). In praise of artifice reloaded: Caution with natural image databases in modeling vision. *Frontiers in Neuroscience*, doi:10.3389/fnins.2019.00008.
- Martinez, M., Cyriac, P., Batard, T., Bertalmio, M., & Malo, J. (2018). Derivatives and inverse of cascaded linear+nonlinear neural models. *PLoS One*, 13(10), doi:10.1371/journal.pone.0201326.
- Martinez-Otero, L., Molano, M., Wang, X., Sommer, F., & Hirsch, J. (2014). Statistical wiring of thalamic receptive fields optimizes spatial sampling of the retinal image. *Neuron*, 81(4), 943–956.
- Martinez-Uriegas, E. (1994). Chromatic-achromatic multiplexing in human color vision. In D. H. Kelly (Ed.), (pp. 117–187). Boca Raton, FL, USA: CRC Press.
- Martinez-Uriegas, E. (1997). Color detection and color contrast discrimination thresholds. In *Proceedings of the OSA Annual Meeting ILS-XIII*, p. 81.
- Morgan, M., Chubb, C., & Solomon, J. (2006). Predicting the motion after-effect from sensitivity loss. *Vision Research*, 46(15), 2412–2420.
- Mullen, K. T. (1985). The CSF of human colour vision to red–green and yellow–blue chromatic gratings. *Journal of Physiology*, 359, 381–400.
- van den Oord, A., & Schrauwen, B. (2014). The student-t mixture as a natural image patch prior with application to image compression. *Journal of Machine Learning Research*, 15(60), 2061–2086.
- Poggio, T. (2021). *From Marr’s Vision to the problem of human intelligence (CBMM Memo No. 118)*. Boston, MA, USA: MIT Libraries, <https://dspace.mit.edu/handle/1721.1/131234>.
- Prenger, R., Wu, M., David, S., & Gallant, J. (2004). Nonlinear v1 responses to natural scenes revealed by neural network analysis. *Neural Networks*, 17(5-6), 663–79.
- Reid, R. C., & Shapley, R. (1992). Spatial structure of cone inputs to receptive fields in primate lateral geniculate nucleus. *Nature*, 356, 716–718.
- Reid, R. C., & Shapley, R. (2002). Space and time maps of cone photoreceptor signals in macaque lateral geniculate nucleus. *Journal of Neuroscience*, 22, 6158–6175.
- Ruderman, D. L., Cronin, T. W., & Chiao, C.-C. (1998). Statistics of cone responses to natural images: Implications for visual coding. *Journal of the Optical Society of America A*, 15(8), 2036–2045.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., & Ma, S. et al. (2015). Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115, 211–252.
- Simoncelli, E., & Heeger, D. (1998). A model of neuronal responses in visual area MT. *Vision Research*, 38(5), 743–761.

- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24(1), 1193–1216.
- Soh, J., & Cho, N. (2021). Deep universal blind image denoising. In *Proceedings International Conference on Pattern Recognition (ICPR)* (pp. 747–754), doi:10.1109/ICPR48806.2021.9412605.
- Stockman, A., & Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40(13), 1711–1737.
- Tao, X., Gao, H., Shen, X., Wang, J., & Jia, J. (2018). Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 8174–8182), doi:10.1109/CVPR.2018.00853.
- Taubman, D. S., & Marcellin, M. W. (2001). *Jpeg 2000: Image compression fundamentals, standards and practice*. Norwell, MA: Kluwer Academic Publishers.
- Valois, R. L. de, & Pease, P. L. (1971). Contours and contrast: Responses of monkey lateral geniculate nucleus cells to luminance and color figures. *Science*, 171, 694–696.
- Wallace, G. K. (1992). The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1), xviii–xxxiv, doi:10.1109/30.125072.
- Wang, Z., & Bovik, A. C. (2009). Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE Signal Processing Magazine*, 26(1), 98–117.
- Watson, A. B. (2013). High frame rates and human vision: A view through the window of visibility. *SMPTE Motion Imaging Journal*, 122(2), 18–32.
- Watson, A. B., & Ahumada, A. J. (2005). A standard model for foveal detection of spatial contrast. *Journal of Vision*, 5(9), 717–740.
- Watson, A. B., & Ahumada, A. J. (2016). The pyramid of visibility. In *Proceedings Human Vision and Electronic Imaging HVEI16*, 102, 1–6, doi:10.2352/ISSN.2470-1173.2016.16HVEI-102.
- Watson, A. B., Ahumada, A. J., & Farrell, J. E. (1986). Window of visibility: A psychophysical theory of fidelity in time-sampled visual motion displays. *Journal of The Optical Society of America A-optics Image Science and Vision*, 3, 300–307.
- Watson, A. B., & Malo, J. (2002). Video quality measures based on the standard spatial observer. In *Proceedings of the IEEE International Conference on Image Processing* (Vol. III, pp. 41–44), doi:10.1109/ICIP.2002.1038898.
- Welles, O. (1946). *The stranger*. New York, NY, USA: RKO Radio Pictures.
- Wichmann, F. A., Janssen, D. H. J., Geirhos, R., Aguilar, G., Schütt, H. H., Maertens, M., . . . Bethge, M. (2017). Methods and measurements to compare men against machines. *Electronic Imaging*, 10, 36–45.
- Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2), 55–80.
- Wuerger, S. M., Ashraf, M., Kim, M., Martinovic, J., Pérez-Ortiz, M., & Mantiuk, R. K. (2020). Spatio-chromatic contrast sensitivity under mesopic and photopic light levels. *Journal of Vision*, 20, 23, <https://doi.org/10.1167/jov.20.4.23>.
- Yamins, D., & DiCarlo, J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19, 356–365.
- Yamins, D., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United State of America*, 111, 8619–8624.
- Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. (2017). Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7), 3142–3155.

## Appendix A: Implementation details

As stated in the main text, all the models follow the the basic toy networks studied in Gomez-Villa et al. (2019, 2020): autoencoders with convolutional layers made of eight feature maps with kernels of spatial size  $5 \times 5$  and sigmoids or ReLUs as activation functions. As illustrated in Figure 4, the last reconstruction layer, has three features in every case (the three color channels) so that the input and output domains are the same. Following our purpose of using biologically plausible image representations the input to the networks and the output signals are expressed in the LMS color space (as opposed to generic RGB digital counts used in the cited references).

The spatiotemporal models follow the same spirit, in this case also including convolution in the temporal dimension: autoencoders with 3D convolutional layers made of eight feature maps with kernels of size  $5 \times 5 \times 5$  and sigmoid activation functions (we did not explore ReLU in videos because in images we found no qualitative difference between the ReLU and sigmoid results). As in the image case, the last (reconstruction)



layer for every architecture (two, four, six and eight layers) only has three feature maps (the LMS channels).

Implementation and training is done in the same way as in Gomez-Villa et al. (2019, 2020): mean squared error is used as loss function in all cases and all the models are implemented using Tensorflow (Abadi et al., 2016). We train our models using ADAM stochastic gradient descent (Kingma & Ba, 2017) with a batch size of 32 examples, momentum of 0.9, and a weight decay of 0.001. In principle, a standard early stopping criterion for convergence was used based on the number of iterations with no improvement in the validation set. However, to ensure appropriate convergence, we visualized the learning curves and we let the iteration continue until train and validation error reached a common plateau. All the learning curves are explicitly shown in the Appendix C, these show that all the CSF considered in the main text come from models with the proper convergence.

## Appendix B: Training stimuli and stimuli for CSF estimation

The natural stimuli to train the networks are regular photographic images from the same dataset used in Gomez-Villa et al. (2019, 2020): the Large Scale Visual Recognition Challenge, 2014 CLS-LOC validation dataset (which contains  $50 \cdot 10^3$  images), leaving  $10 \cdot 10^3$  images for validation purposes. This dataset is a subset of the whole ImageNet dataset (Russakovsky et al., 2015). The experiments with cartoon images were done using  $25 \cdot 10^3$  frames taken from The Pink Panther

Show (Freleng, 1963) reproduced with permission of the MGM. In every case we take  $128 \times 128$  images and assume a sampling frequency  $f_s = 70$  cpd, that is, we assume that the images subtend  $3.6^\circ$ .

The spatiotemporal models are trained over  $25 \cdot 10^3$  patches of size  $32 \times 32 \times 25$  from classical Hollywood films which are in public domain: the color movies *Charade* (Donen, 1963) and *The FBI story* (LeRoy, 1959), and the achromatic movie *The Stranger* (Welles, 1946). In all the video cases we assume a spatial sampling of 30 cpd and temporal sampling of 25 Hz, that is, we assume the patches subtend  $1.06^\circ$  and last for 1 second. These somewhat arbitrary selections of the sampling frequencies (or extent of the stimuli) have mild consequences on the quantitative evaluation of the CSFs as discussed elsewhere in this appendix.

The transform from digital counts to LMS tristimulus values was done assuming the primaries and gamma curves of a standard CRT display (Malo & Luque, 2002).

Regarding the stimuli for the estimation of the CSFs according to Equation 3, our  $\mathbf{b}^f$  are gratings in the classical opponent space of Hurvich and Jameson (1957). Figure 13 shows a representative subset of the gratings used to feed the networks for the estimation of the spatiochromatic CSFs. The justification of the use of these waves to probe the autoencoders follows the eigenanalysis of the linearized networks introduced in Gomez-Villa et al. (2020): the eigenfunctions of the matrices in Equation 8 were shown to be oscillating functions in space with chromatic variations in luminance and opponent red–green and yellow–blue directions. Consistently with Gomez-Villa et al. (2020) the corresponding spatiotemporal oscillations of increasing frequency for decreasing eigenvalue are




Figure 13. Representative spatiochromatic stimuli to feed the 2D networks. The 3D networks were probed with equivalent gratings moving at different speeds (see text).

obtained when the considered Jacobian corresponds to spatiotemporal autoencoders.

The full spatiochromatic set included gratings of 60 spatial frequencies linearly spaced in the range [0.5, 35] cpd (the Nyquist region assuming  $f_s = 70$  cpd), and nine contrasts linearly spaced in the range (0.07, 0.6). The average color was the white of the color system with  $30 \text{ cd/m}^2$ . The stimuli for the computation of the spatiotemporal CSFs were moving sinusoids with 16 spatial frequencies in the range of 0 to 15 cpd, 10 temporal frequencies in the range or 0 to 10 Hz, and 9 contrasts in the range of 0.07 to 0.6. The average color and luminance were the same as in the image case.

The lower limit of the explored contrast range comes from the limitation owing to noise discussed in Methods: Estimating contrast sensitivity in autoencoders. The upper contrast limit and the average luminance were selected to ensure that corrupted signals are reproducible in regular displays (which is the range of the scenes used in the training).

## Appendix C: Convergence and Performance of the models

To guarantee that the presented CSF results do not come from eventual training artifacts, this appendix illustrates the proper convergence and proper performance of all the considered CNNs in all the goal/architecture scenarios. For each considered case in the Experiments, we show the learning curves and explicit examples of the responses (reconstructed signals in test). Finally, as an illustrative example, we also show one extra case (for video, experiment 6) where convergence was not complete in one of the networks and the consequences in the reconstructions and in the CSFs.

### Experiment 1: Distortion compensation from a range of CNN architectures

Figure 14 shows the learning curves of all the CNN models used in experiment 1. Throughout the Appendix the gray/black curves refer to the  $\varepsilon_{LMS}$  distortion of the retinal signal in the LMS color space of Stockman and Sharpe (2000), which is constant along the learning. The cyan–blue curves show the evolution of the error in the response of the networks. The light color curve describes the error over image batches in the training phase while the dark color describes the same error in the validation set. The error of the response (solution) significantly drops below the error of the input signal (problem), thus indicating that the network is actually

achieving the functional goal it has been designed for. The plateau achieved by the blue curves (not only in training but, more significantly in validation) implies that a steady convergence was achieved and the resulting model is ready to be tested. Consistency between the train and validation sets is apparent from the parallel behavior of the light and dark curves. Performance tables in the main text (Table 1 for experiment 1) and performance in the visual examples shown here (Figure 14 for experiment 1) refer to an independent test set not used in the learning (training/validation) phase. In all CNNs used in experiment 1, the training has been done in a representative set because the errors in the independent test phase (Table 1 and Figure 14) are consistent with the asymptotic behavior of the learning curves shown in Figure 14. Figure 15 shows visual examples of the performance of the linearized versions of the nonlinear models in experiment 1. It is interesting to note that the optimal linear solution (computed from the train set) has worse behavior than the linearized versions of the networks (as also seen in Table 1).

### Experiment 2: Architecture trained on a range of distortion levels

Figure 16 shows the learning curves of the model trained in experiment 2 (two layer ReLU) for different levels of retinal degradation (noise/blur). Note that the specific cases where the iteration was stopped owing to the activation of the early stopping criterion (top left and bottom center), the convergence plateau was already reached. Figure 17 shows examples of the performance in test of the model considered in every training scenario considered in experiment 2.

### Experiment 3: Chromatic adaptation versus distortion compensation

Figure 18 (top) demonstrates that the model trained for the five computational goals considered in experiment 3 actually achieves the goals and has proper convergence. Figure 18 (bottom) shows visual examples of the performance.

### Experiment 4: Robustness under change of signal statistics

Figure 19 (top) demonstrates that the model trained for the five computational goals considered in experiment 4 actually achieves the goals and has proper convergence. Figure 19 (bottom) shows visual examples of the performance.

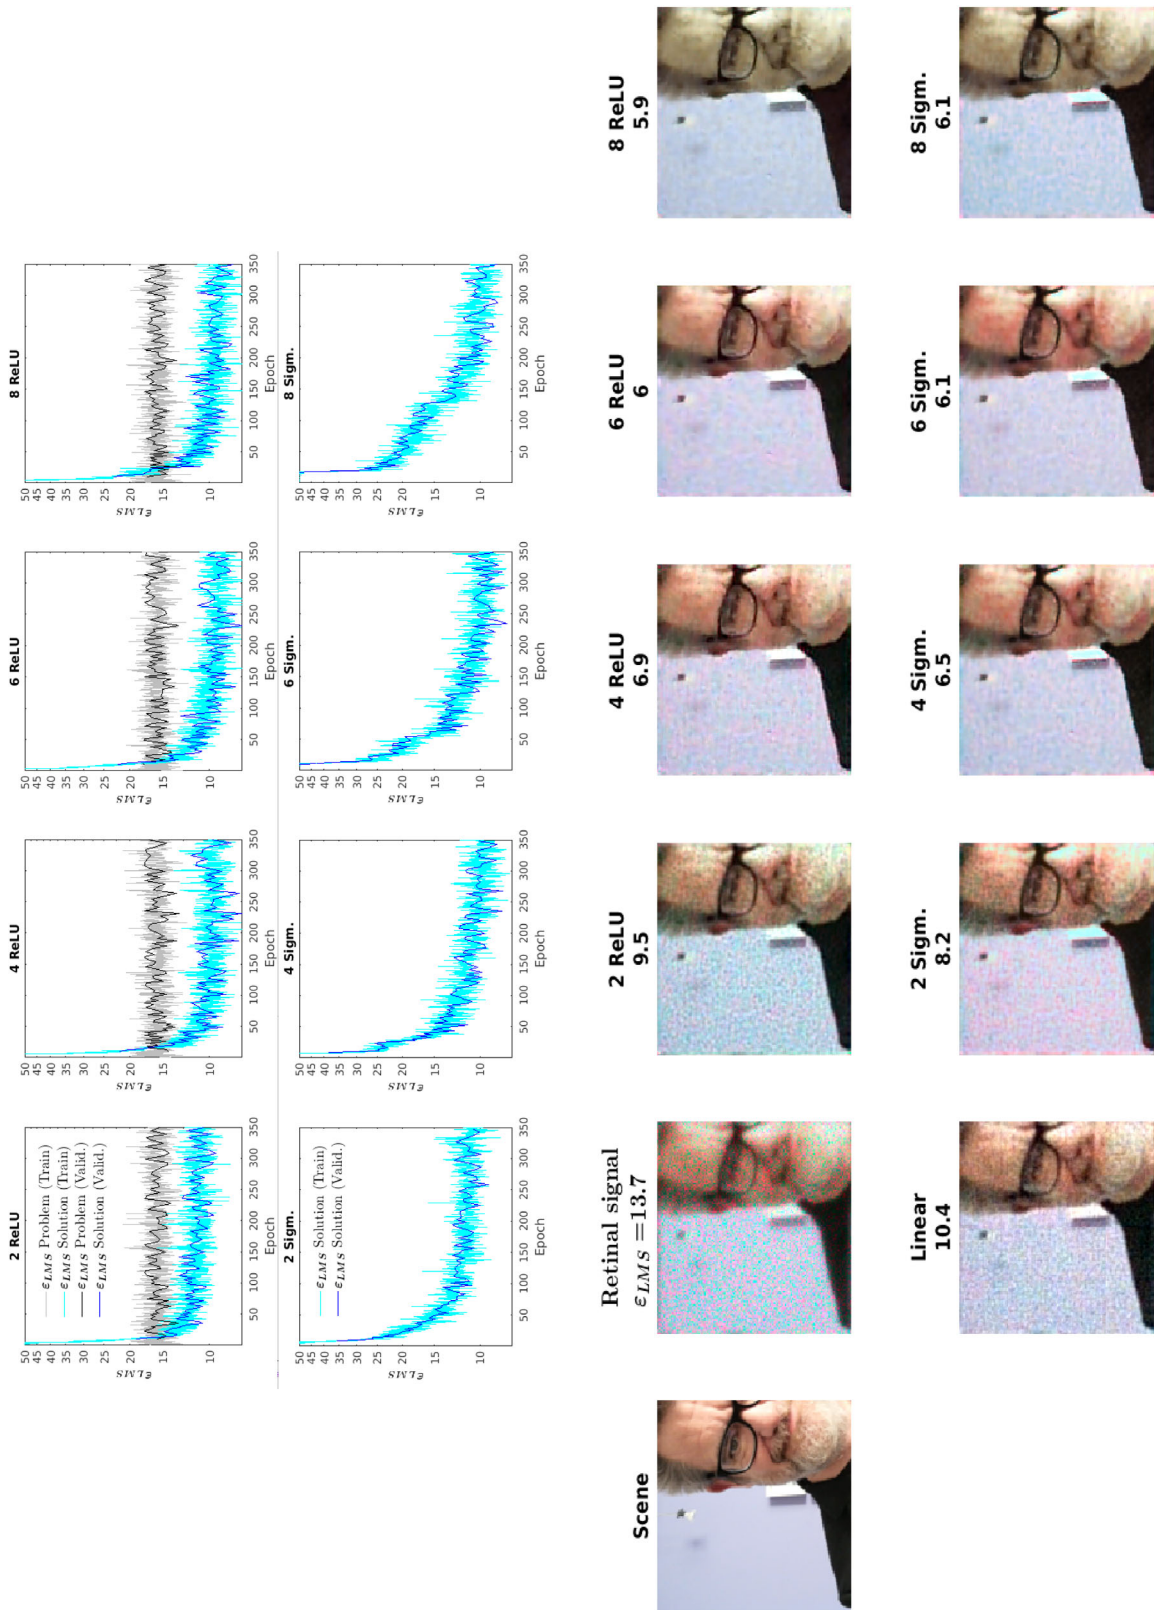


Figure 14. Experiment 1: Convergence and trained models. Learning curves (training/validation) and examples of visual performance (test) of all the models trained in experiment 1. The distortion  $\epsilon_{LMS}$ -Problem refers to the original degradation of the images (previous to the application of the net). This distortion describes how difficult the compensation problem is. The distortion  $\epsilon_{LMS}$ -solution refers to the degradation remaining in the signal after the application of the net. It describes how close the output is to the ideal result. Performance numbers have been truncated to the significance of the standard deviation in the test set.

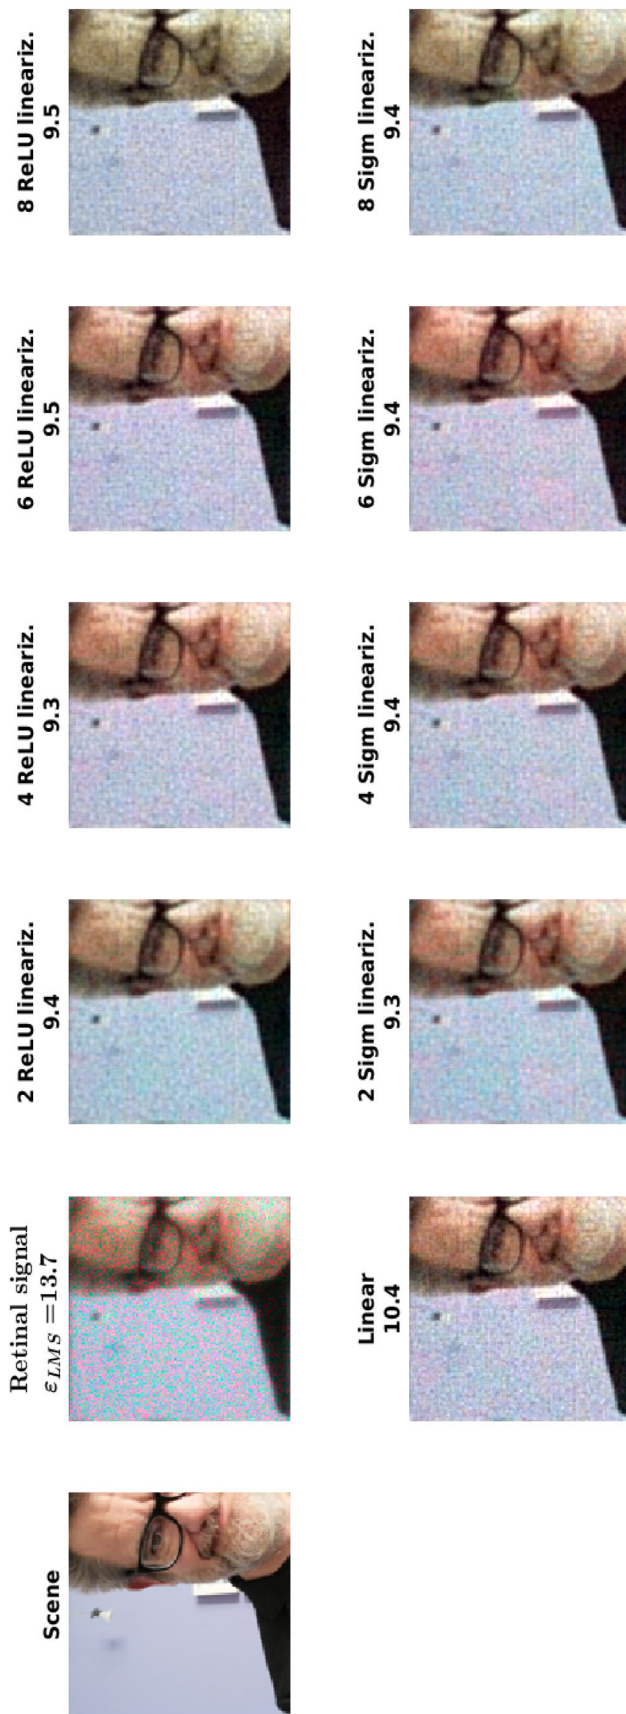


Figure 15. Experiment 1: linearized models. Examples of visual performance (in test) for the linearized CNNs of experiment 1.

## Experiment 5: CSFs from bottleneck compensation and biodistortion compensation


Figure 20 demonstrates that the models trained for the goals considered in experiment 5 have proper convergence and actually achieve the goals within the constraints imposed by the bottlenecks of progressive severity. Figure 21 shows visual examples of the performance.

The reconstruction error  $\epsilon_{LMS}$  behaves quite intuitively (see Figure 20 here and Table 3 in the main text): On the one hand, in the cases that involve biodistortion all the architectures do reduce the original value of  $\epsilon_{LMS}$  except the architecture  $G$ , which has a single feature in its bottleneck. On the other hand, the pure reconstructions cases introduce negligible distortion  $\epsilon_{LMS}$  when the inner representation does not restrict spatial resolution nor number of features (the no-bottleneck cases A and B). And, as expected, more severe bottlenecks imply higher  $\epsilon_{LMS}$ .

## Experiment 6: Spatiotemporal temporal CSFs

The main text includes the CSF results from a range of models trained in Charade (1963). Figure 22 shows the regular evidences on convergence (top) and visual performance in test (bottom) shown for the other Experiments.

In this Appendix, we also include a replication of experiment 6 trained on a movie with higher spatial resolution (*The FBI Story*, 1959). We give the corresponding learning curves (Figure 23) and CSFs (Figure 24). This is interesting for two reasons: 1) it confirms the superiority of shallower nets even for different resolution, and 2) it shows an example of failure in convergence (see that the eight layer model in Figure 23 got stuck in a local minimum (with poor performance) and this has consequences in the complete loss of chromatic information (frame not shown because of copyright issues). Also interesting is the fact that models with four or six layers (which converged as well as the two layer model), substantially over attenuate the red–green channel with the corresponding yellowish–bluish look of the reconstruction and the corresponding impact on the relative scaling of channels in the CSFs, which is not the case for the linear and the 2-layer solutions. This is consistent in the other image/video examples in the main text.



	Degradation			2-layer ReLU			CSF		
	$\epsilon_{LMS}$			$\epsilon_{LMS}$			RMSE		
	F=0.25	F=0.5	F=1	F=0.25	F=0.5	F=1	F=0.25	F=0.5	F=1
d=2	10.5±0.1	12.9±0.1	15.8±0.1	7.55±0.07	8.79±0.08	10.4±0.1	29.1	28.1	26.1
d=4	13.5±0.2	15.5±0.2	18.1±0.1	8.9±0.1	9.1±0.1	13.1±0.1	24.8	24.5	24.0
d=6	17.7±0.2	19.4±0.2	21.6±0.2	11.8±0.2	15.2±0.2	15.1±0.2	22.5	23.5	24.2

Figure 16. Experiment 2: Convergence and performance. Learning curves (training/validation) and numerical performance (in test and in the reproduction of the CSFs) of the CNN model trained in all conditions considered in experiment 2.

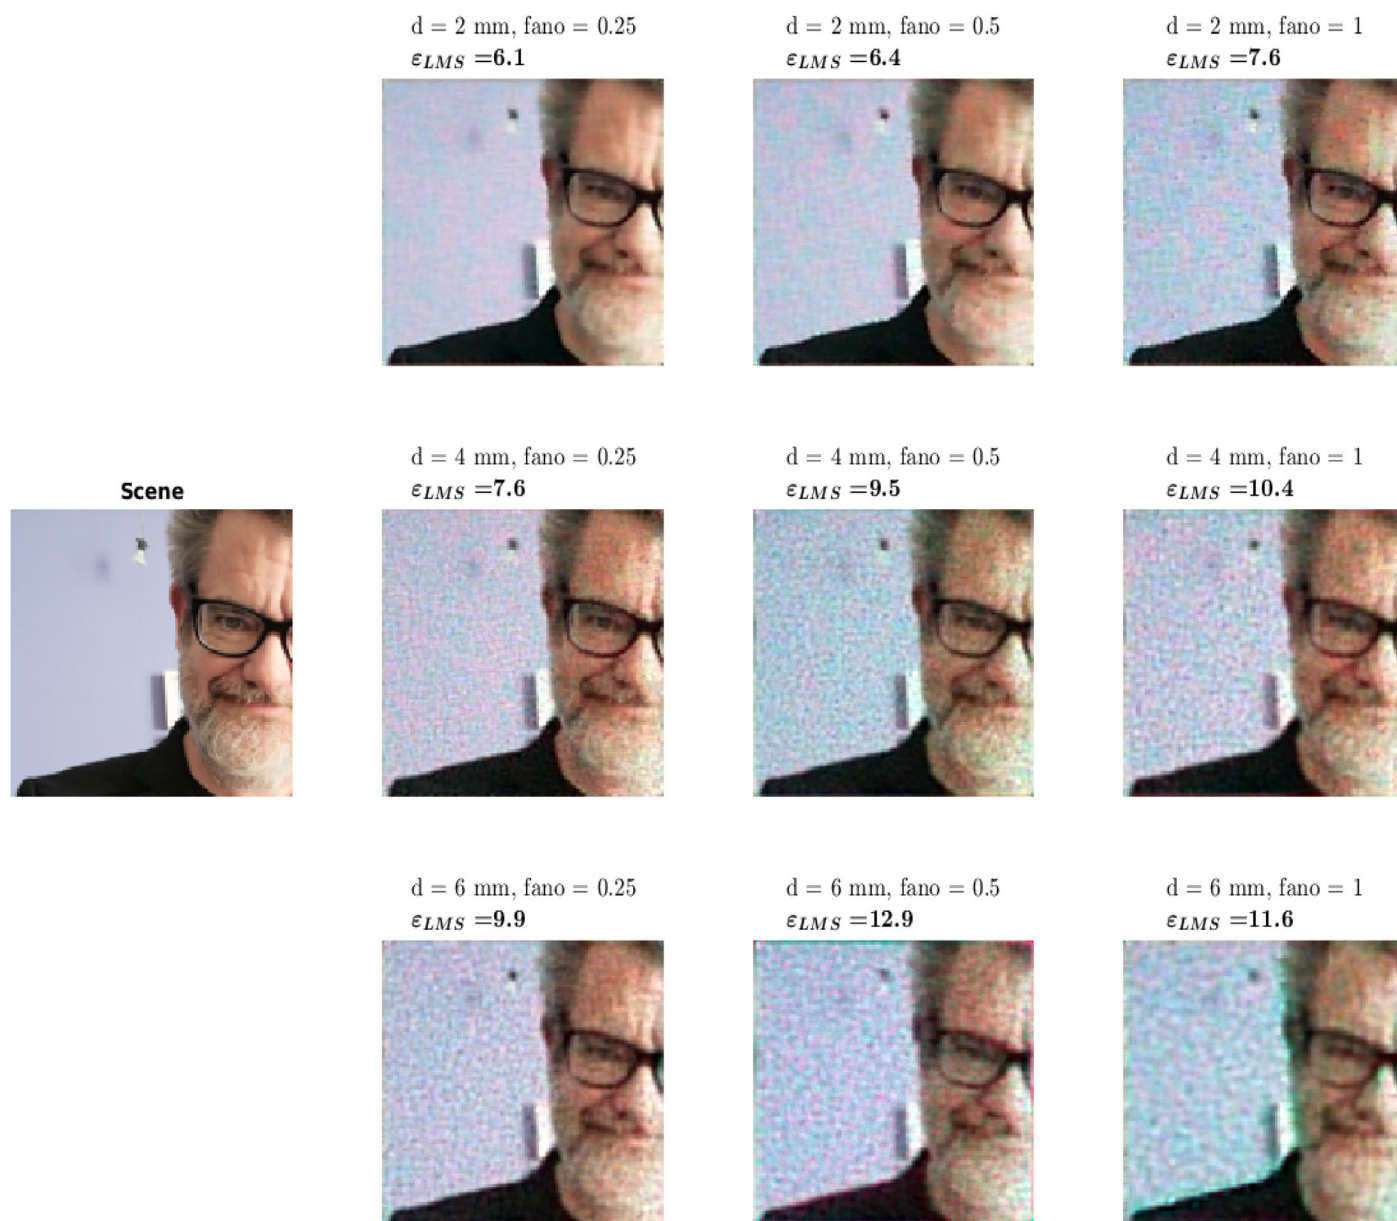


Figure 17. Experiment 2: Visual performance. Examples of reconstruction (in test) for the CNN in all the degradation scenarios considered in experiment 2.

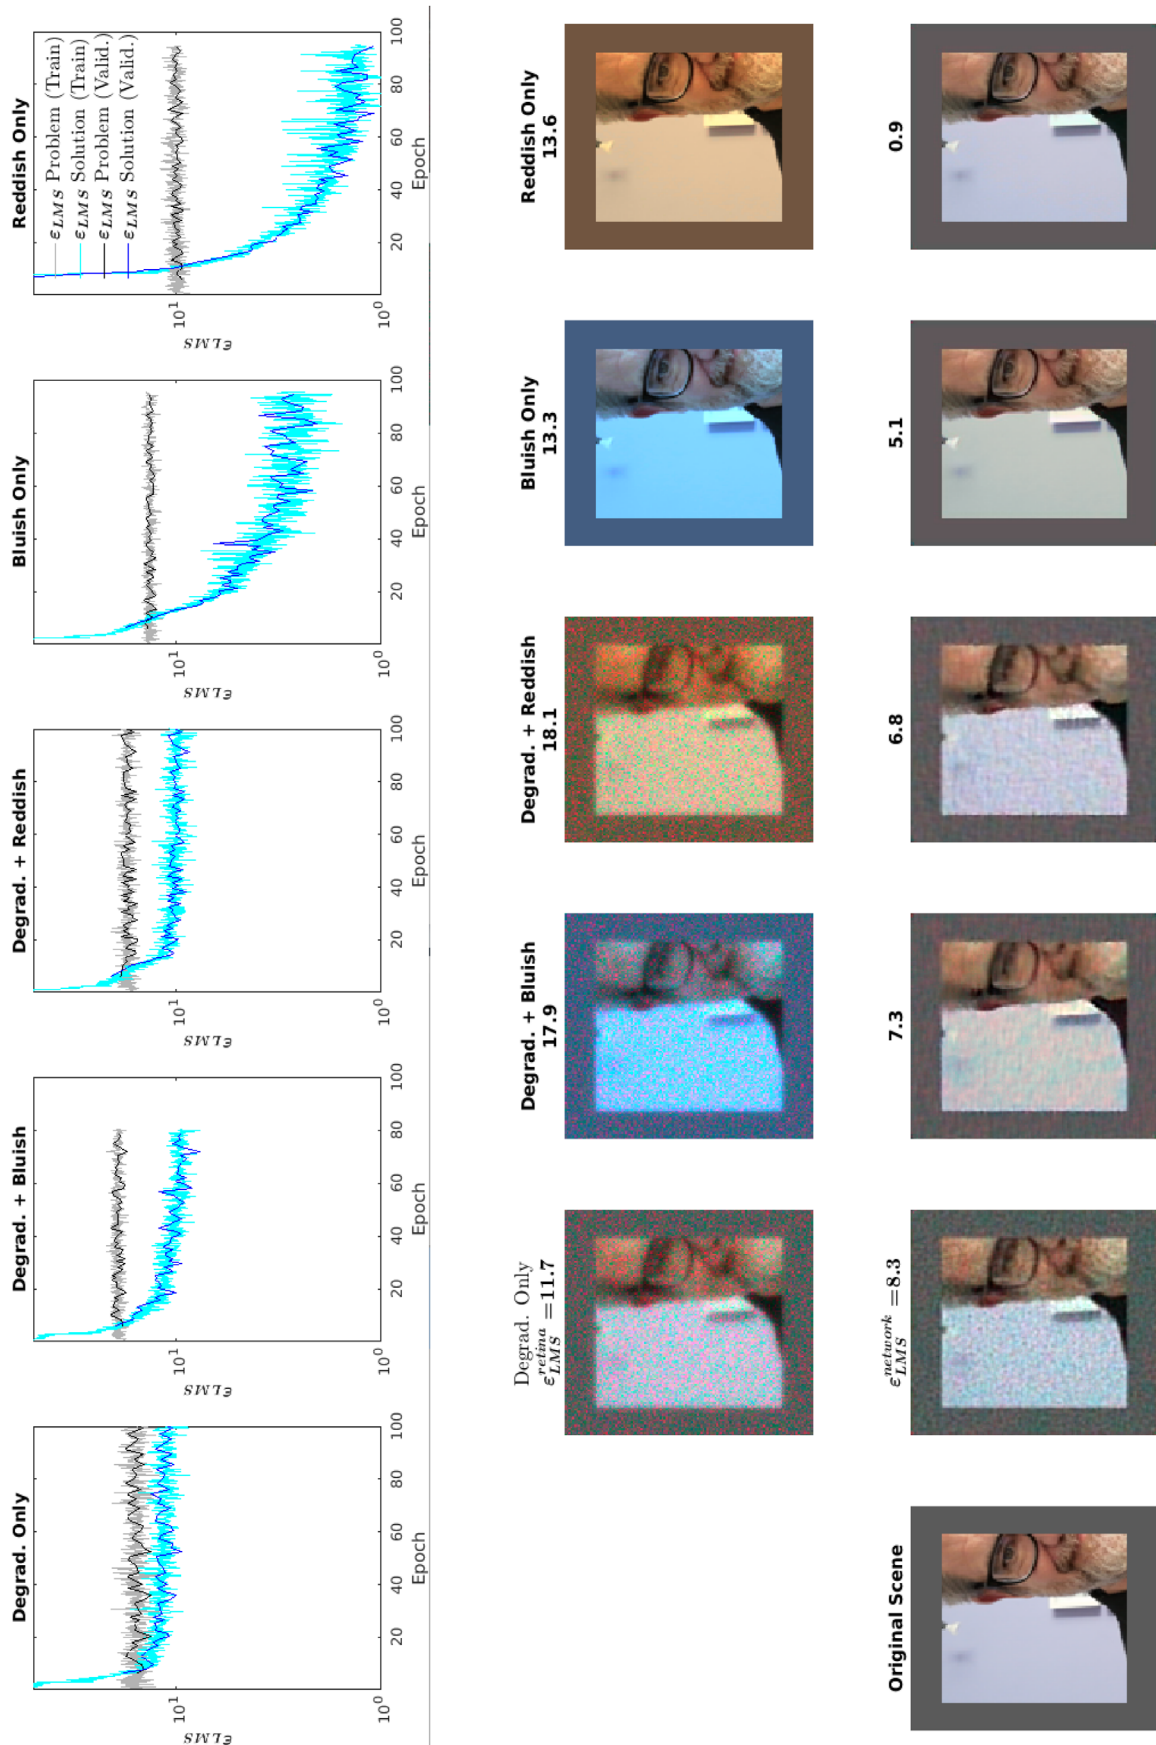


Figure 18. Experiment 3: Convergence and visual performance: Top, learning curves (train/validation) for the considered architecture in the different goals. Bottom: Visual example (test) for the CNNs in experiment 3 (natural images).

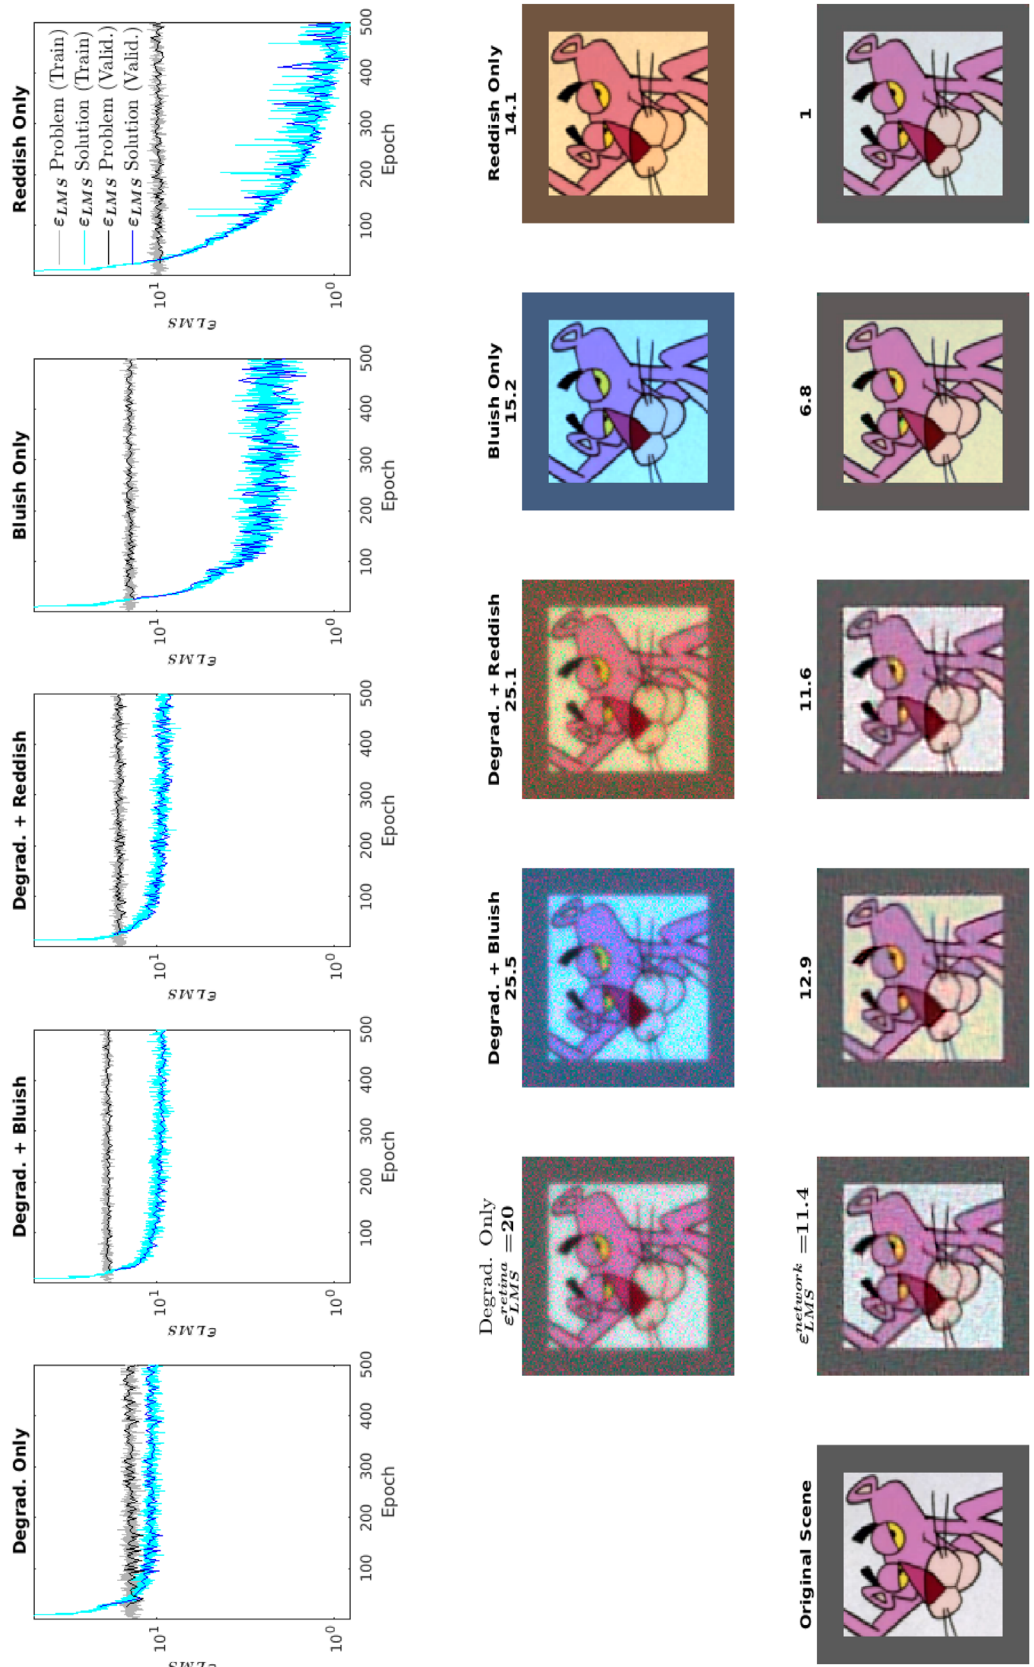


Figure 19. Experiment 4: Convergence and visual performance: Top, learning curves (train/validation) for the considered architecture in the different goals. Bottom: Visual example (test) for the CNNs in experiment 4 (cartoon images). Original image from *The Pink Panther Show* (Freleng, 1963) courtesy of MGM. Similar images (Malo, 2022) lead to equivalent performance in the networks.



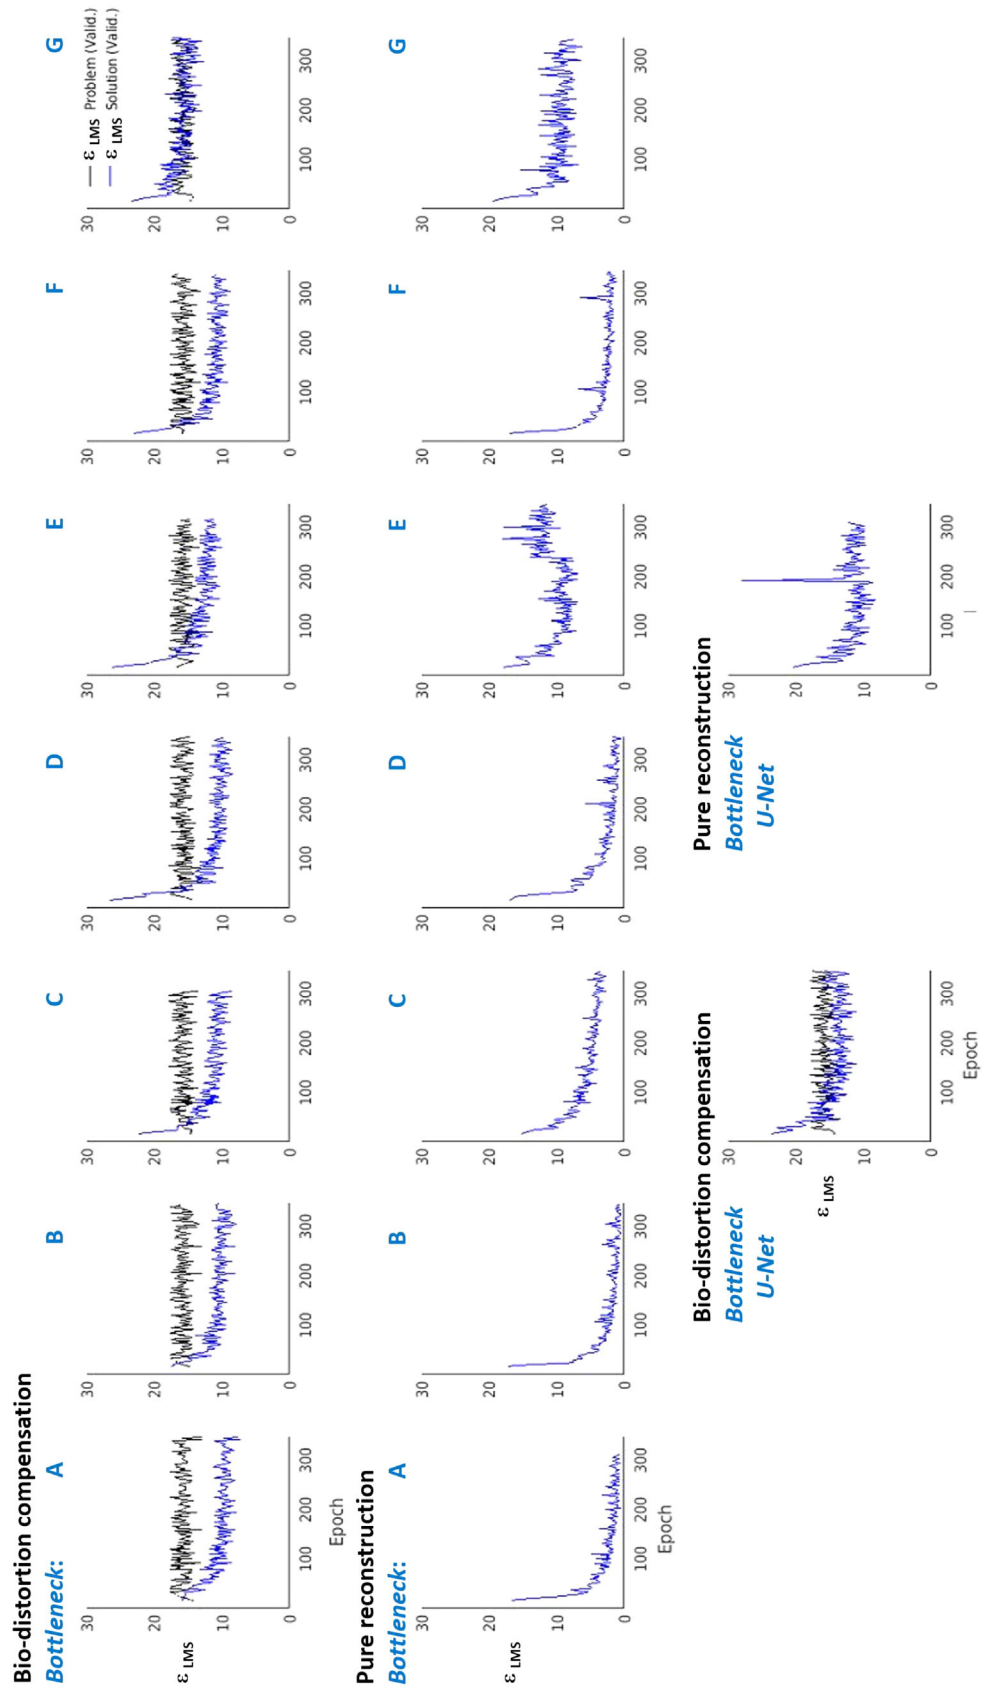


Figure 20. Experiment 5: Bottleneck compensation and biodistortion compensation. Learning curves (train/validation) for the considered architectures in reconstruction with compensation of biodistortion (top row and bottom row, left) and pure reconstruction (middle row and bottom row, right). The cases including biodistortion show the original  $\epsilon_{LMS}$  of the problem. See Figure 4 for the structure of the architectures referred by the letters in blue.



Figure 21. Experiment 5: Bottleneck compensation and biodistortion compensation. Examples of visual performance (in test) for the CNNs of experiment 5. See Figure 4 for the structure of the architectures referred by the letters in blue.

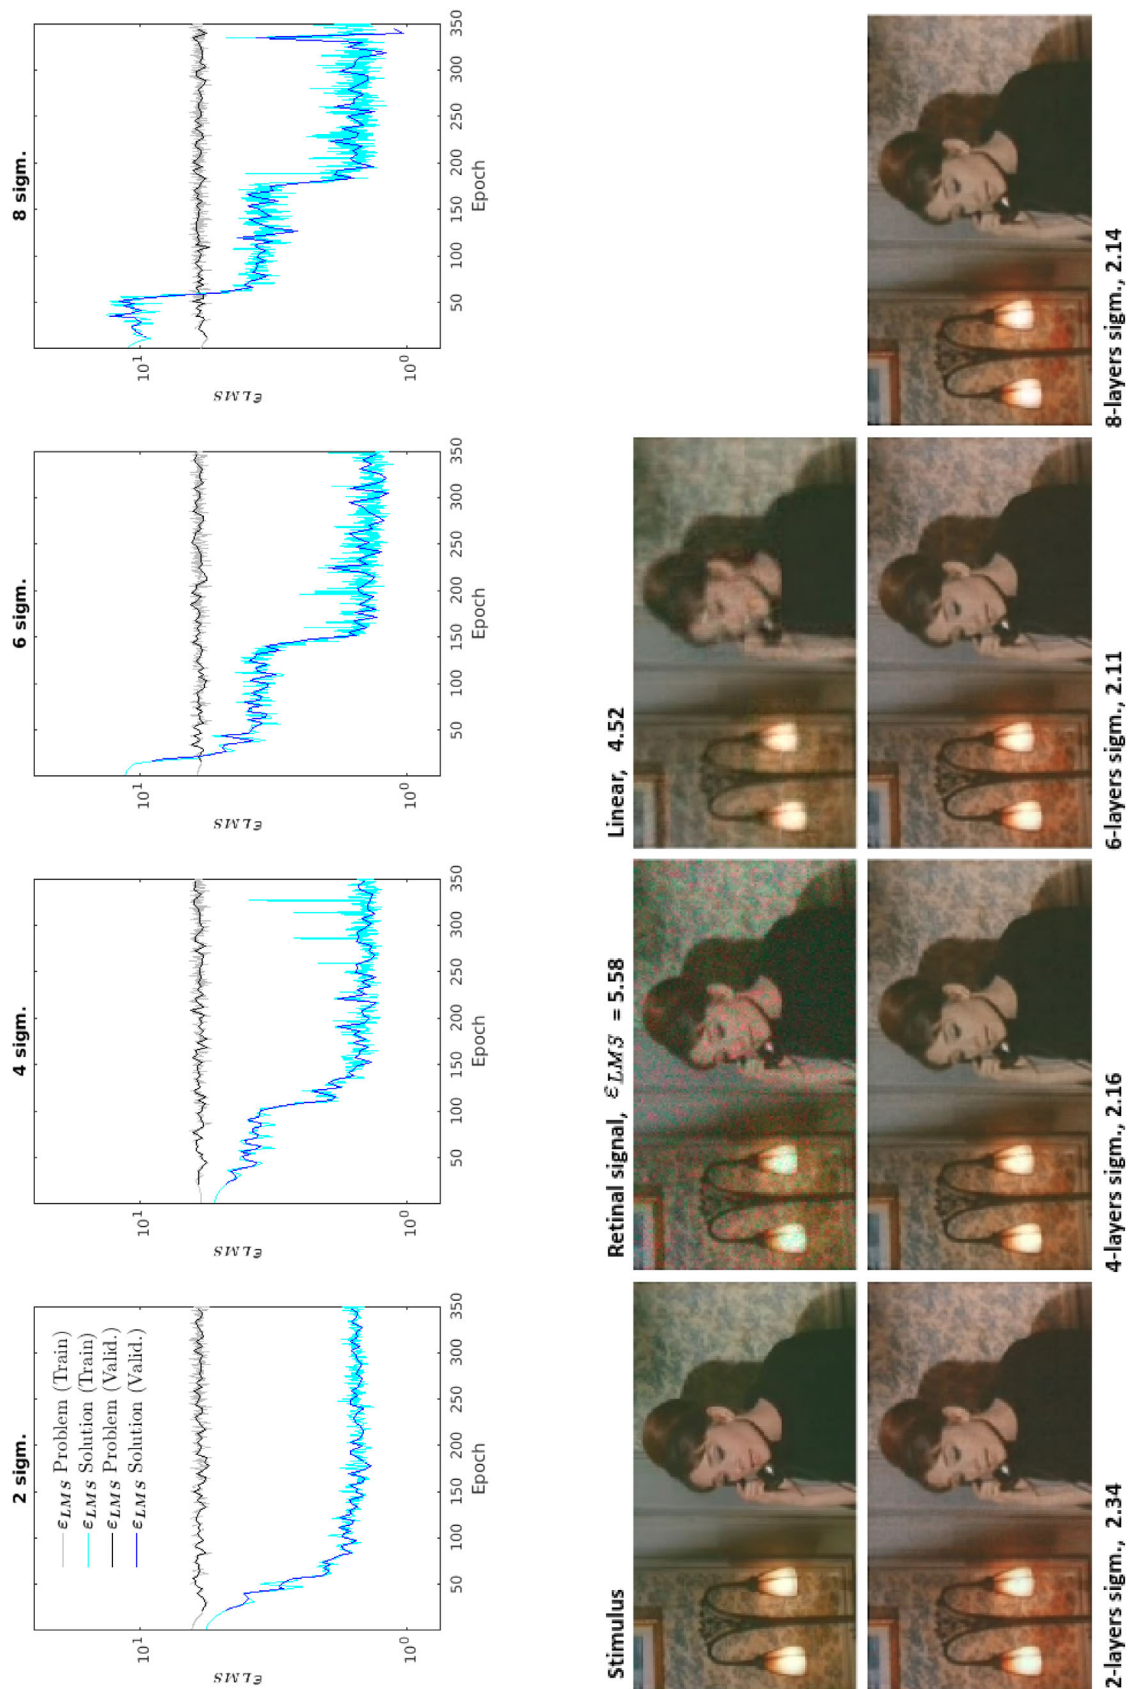



Figure 22. Experiment 6 (spatiotemporal chromatic CSFs in main text): Convergence and visual performance. Top, learning curves for the considered architectures. Visual example (test) for the linear solution and the CNNs (*Charade*, low-resolution movie). The original frame comes from the film *Charade* (Donen, 1963), which is in the public domain.



←

*The FBI story*; LeRoy, 1959). All models converge except the 8-layer architecture. Visual results are not shown due to copyright issues. Interested readers can obtain these specific results from the authors. The local minimum in which the eight layer architecture was trapped has consequences in the complete loss of chromatic information (frame not shown because of copyright issues). Also interesting is the fact that models with four and six layers (which converged as well as the two layer model), substantially over attenuate the red–green channel with the corresponding yellowish–bluish look of the reconstruction and the corresponding impact on the relative scaling of channels in the CSFs, which is not the case for the linear and the two layer solutions.

Figure 23. Extended experiment 6: Spatiotemporal chromatic CSFs with different training set (not shown in the main text): Convergence. Figure shows the learning curves for the considered architectures on a higher resolution movie (the film

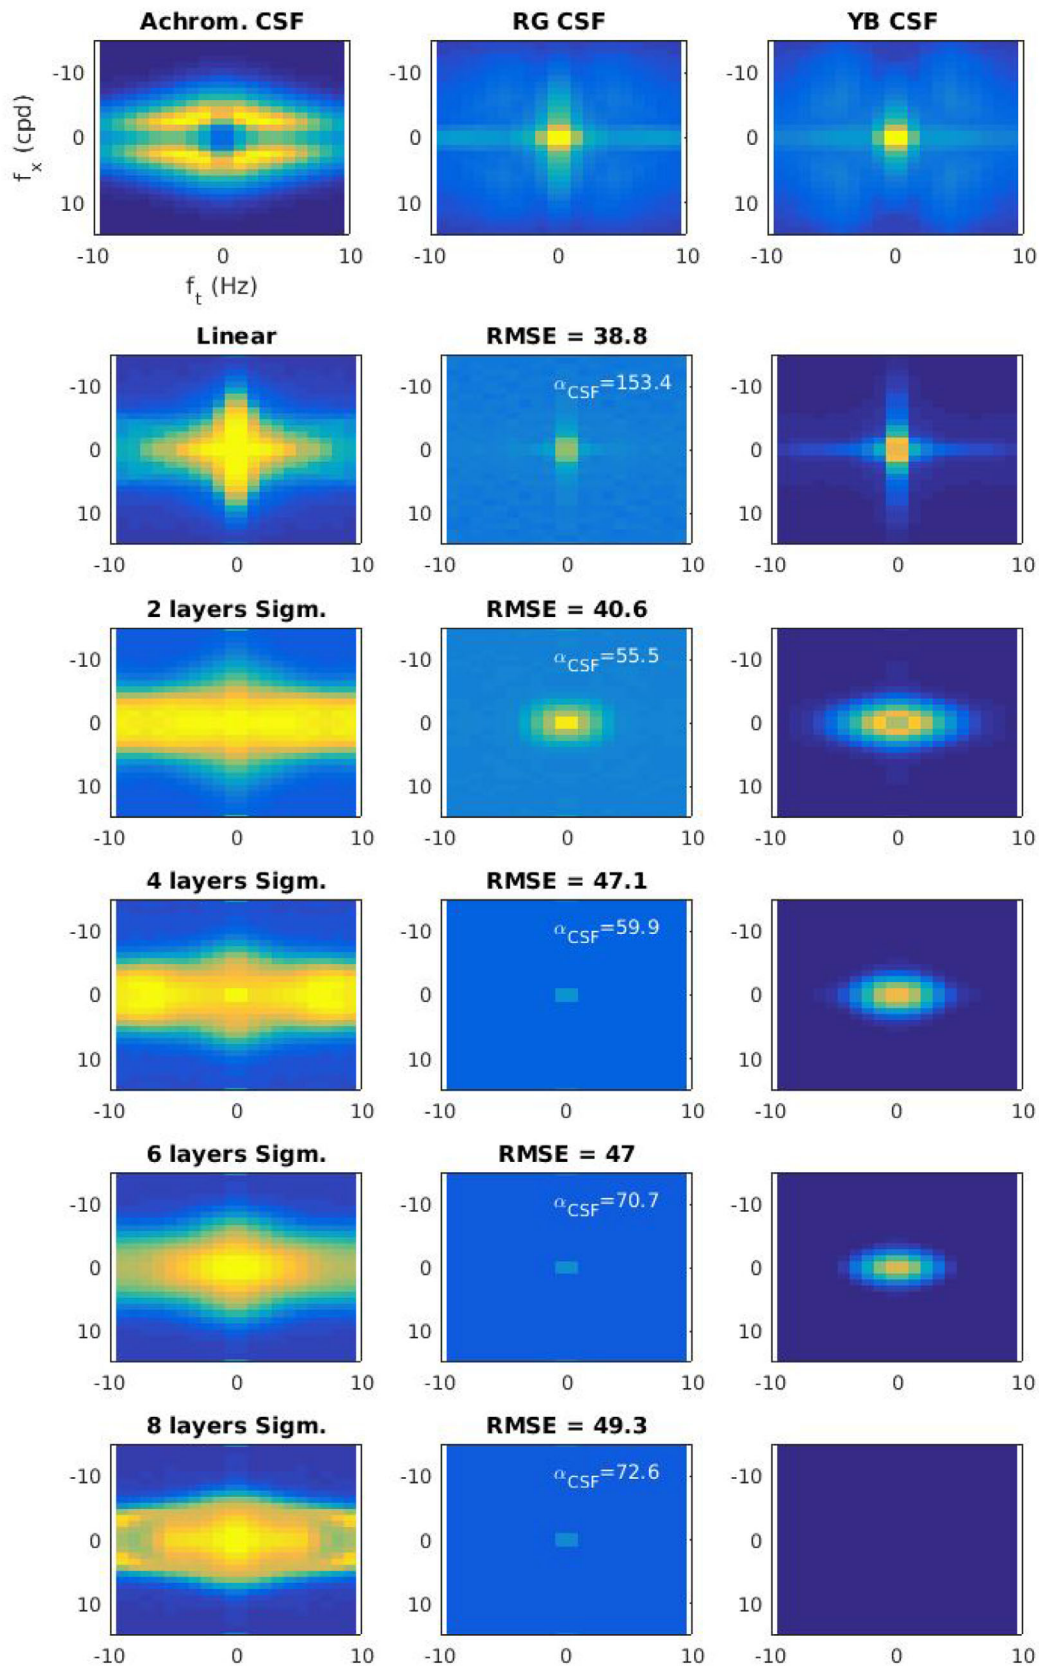


Figure 24. Extended experiment 6 (spatiotemporal chromatic CSFs of a high-resolution movie not shown in the main text). Note that in this case (see Figure 23) all models converged except the eight layer architecture, which totally removes the yellow–blue channel and almost removed the red–green channel.