

# PERCEPTNET: A HUMAN VISUAL SYSTEM INSPIRED NEURAL NETWORK FOR ESTIMATING PERCEPTUAL DISTANCE

Alexander Hepburn\* Valero Laparra† Jesús Malo† Ryan McConville\* Raul Santos-Rodriguez\*

\* Department of Engineering Mathematics, University of Bristol

† Image Processing Lab, Universitat de Valencia

## ABSTRACT

Traditionally, the vision community has devised algorithms to estimate the distance between an original image and images that have been subject to perturbations. Inspiration was usually taken from the human visual perceptual system and how the system processes different perturbations in order to replicate to what extent it determines our ability to judge image quality. While recent works have presented deep neural networks trained to predict human perceptual quality, very few borrow any intuitions from the human visual system. To address this, we present PerceptNet, a convolutional neural network where the architecture has been chosen to reflect the structure and various stages in the human visual system. We evaluate PerceptNet on various traditional perception datasets and note strong performance on a number of them as compared with traditional image quality metrics. We also show that including a nonlinearity inspired by the human visual system in classical deep neural networks architectures can increase their ability to judge perceptual similarity. Compared to similar deep learning methods, the performance is similar, although our network has a number of parameters that is several orders of magnitude less.

**Index Terms**— perceptual distance, human visual system, neural networks

## 1. INTRODUCTION

The human visual system’s ability to compare images through a variety perturbations is still unparalleled. Within machine learning and computer vision, *perception* has become increasingly relevant, and many metrics have attempted to capture characteristics of the human perceptual system in order to replicate the ability to perceive differences between images. Despite being a well established field, an often overlooked aspect is where the foundations of image processing originate; the human visual perceptual system. In the past, subjective image quality metrics were proposed following two principles: the *visibility of errors* derived from psychophysical models [1–4], or the preservation of perceptual *structural similarity* as in SSIM and variants [5–7]. More recent efforts has been focused on training neural networks to distinguish

between patches of images [8]. Although these networks are shown to be successful in datasets containing a wide variety of perturbations, the network structure takes no inspiration from the human visual system. These networks also contain millions of parameters and are often difficult to interpret. In fact, recently it has been shown that blind fitting of architectures which are not properly constrained may lead to failures in reproducing fundamental perceptual phenomena [9, 10].

We propose combining the recent methodologies from deep learning approaches and traditional image quality metrics by constructing a distance where the architecture takes inspiration from what we understand about the human visual system. In this paper we present PerceptNet, a carefully constructed network that has been trained on a limited set of perturbations and has an ability to generalise to perturbations in other datasets. PerceptNet outperforms classical measures and in traditional image quality databases and performs similarly to deep learning measures despite having two orders of magnitude less parameters.

## 2. RELATED WORK

Image quality metrics (IQMs) have long been relying on our understanding of the human visual perceptual system. Methods based on assessing the *visibility of errors* apply models of the psychophysical response to the original image and to the distorted image, and then compute Euclidean distances in the transformed domain. These response models have always been cascades of linear+nonlinear layers, mainly wavelet-like filters followed by divisive normalisation saturations. The difference between old implementations of this idea [1–4] and newer ones [11, 12] is the biological sophistication of the models and the way they are optimised. Models based on *structural similarity* such as SSIM [13] and its variants MS-SSIM and FSIM [5, 7] check the integrity of the statistics of the distorted image. However, it has been shown error visibility models may be as adaptive as structural similarity [4]. Therefore, current versions of error visibility models based on normalised laplacian pyramids (NLAPD) [11, 14] clearly outperform structural similarity methods whilst also shown to be effective at enforcing perceptual quality in image generation [15]. The linear+divisive normalisation layer can also


be formulated as a convolutional+nonlinear layer in CNNs, called Generalised Divisive Normalisation (GDN) [16].

A perceptual distance that is often used in the deep learning literature is the Learned Perceptual Image Patch Similarity (LPIPS) [8]. LPIPS utilises architectures trained to classify images on ImageNet [17] as feature extractors. A weighted importance vector is learned and a combination of spatial average and  $\ell_2$  distance is used to compute the perceptual distance. An alternative is to train using random initialisations and the Berkeley-Adobe Perceptual Patch Similarity (BAPPS) dataset. BAPPS contains traditional perceptual distortions, convolution-based perturbations and a combination of both. In traditional human judgement experiments, the observer is presented with an original image and two distorted images and is asked to select the distorted image that is most similar with the original. BAPPS contains only the fractional preference for each combination of two images and, as such, the distance output from the networks must be transformed into a preference score. A network  $\mathcal{G}$  containing two fully connected layers is used, which takes as input the distances from the original to both distorted images and outputs a predicted preference. Interestingly, it is known that convolution neural networks trained on ImageNet have a texture bias which contradicts what we know about the human visual perceptual system [9]. Similar departures from the desired perceptual behaviour have been also reported when the training set is not appropriate and the architecture is not properly constrained [10].

### 3. A PERCEPTUALLY CONSTRAINED ARCHITECTURE

To address the aforementioned shortcomings of current deep learning approaches, we devise an architecture, PerceptNet, for our proposed networks following the program suggested in [18]: a cascade of canonical *linear filters + divisive normalisation* layers that perform a series of perceptual operations in turn simulating the retina - LGN - V1 cortex pathway [12].

The architecture is depicted in Fig. 1. Firstly, we use GDN to learn Weber-like saturation [19] at the RGB retina. Then, we learn a linear transformation to an opponent colour space, analogue to the achromatic, red-green, yellow-blue colour representation in humans [19]. This linear transform is subsequently normalised again using GDN to learn a chromatic adaption process similar to Von-Kries [20]. Afterwards, spatial convolutions are allowed to learn center-surround filters as in LGN [21], which are known to have nonlinear GDN-like behaviour [3, 11, 12]. Finally, we include a new convolution+GDN stage to account for the wavelet-like filters at V1 cortex and the divisive normalisation [4, 22]. This domain replicates the representation at the end of the primary visual cortex, where most of the information is contained in various orientation sensitive edge detectors whilst preserving



**Fig. 1:** A diagram displaying the structure of PerceptNet. Each stage has a corresponding process in the human visual system: gamma correction  $\rightarrow$  opponent colour space  $\rightarrow$  Von Kries transform  $\rightarrow$  center-surround filters  $\rightarrow$  LGN normalisation  $\rightarrow$  Orientation sensitive and multiscale in V1  $\rightarrow$  divisive normalisation in V1.  $(a \times a)$  denotes the filter size used. gdn1 is applied independent to each channel.

a map of spatial information.

The network is trained to maximise the Pearson correlation,  $\rho$ , between the mean opinion score (MOS) and the  $\ell_2$  distance of the two images in the transformed domain:

$$\max_f \rho(\|f(x) - f(d(x))\|_2, y), \quad (1)$$

where  $f(x)$  is the transformation of the network (from RGB space to the more informative perceptual space),  $x$  is the reference image,  $d(x)$  is the distorted image and  $y$  is the corresponding MOS calculated from human observer experiments.

A number of properties are recognisable in the way that humans process images; one being a focus on medium frequency in the receptive field [3]. The contrast sensitivity function of the spatial standard observer (SSO) models this behaviour. The SSO model is used to judge perceptual distance between two contrast patterns and tends to also focus on medium frequencies as a result. Our network captures these characteristics as they are intrinsically linked to judging human perceptual distance.

### 4. EXPERIMENTS

It is our aim to find a representation that informs us of the overall perceptual quality, generalising to distortions not seen during the training phase. To this end, we use the TID2008 [23] dataset for training. It contains 17 distortions, with 1428 distorted images and corresponding MOS. Our code and models are publicly available<sup>1</sup>. We evaluate the network on multiple perceptual datasets; TID2013 [24], CSIQ [25], LIVE [26] and BAPPS [8]. A simple description of the datasets can be seen in Table 1. Both TID2013 and BAPPS contain distortions that are not present in the

<sup>1</sup><https://github.com/alexhepburn/perceptnet>


TID2008 dataset. The BAPPS dataset is slightly different in that it does not report MOS. The dataset also contains interesting distortions such as convolutional neural network (CNN) based distortions, while the test set contains distortions that are made using other deep learning algorithms, such as super-resolution, video deblurring and colourisation.

We will be comparing PerceptNet with several baselines, namely, the  $\ell_2$  distance between reference and distorted image, traditional IQMs like SSIM [6], FSIM [7] and MSSIM [5]. We also compare against the NLAPD proposed in [11], but we replace the divisive normalisation step and each stage in the pyramid with a generalised divisive normalisation process, where the parameters are optimised using the TID2008 dataset. The main reference algorithm for deep learning architectures is the LPIPS measure proposed in [8]. Zhang *et al.* found that LPIPS AlexNet initialised from scratch trained on the BAPPS dataset performed best on the BAPPS test subset. For LPIPS measures, *scratch* denotes that the network was trained from random initialisation, and *tune* indicates that the network was pretrained on a dataset and fine-tuned to the BAPPS dataset. We also train AlexNet on perceptual datasets as a feature extractor. Importantly, it should be noted that LPIPS AlexNet requires 24.7m parameters whereas PerceptNet has 36.3k parameters. ImageNet contains millions of images compared to traditional perceptual datasets, which usually have thousands or hundreds of examples. When we train AlexNet on the traditional perceptual datasets such as TID2008, we use the feature extractor section of the network and disregard the classification section. We train the network using the same correlation loss in Eq. 1. When comparing with LPIPS, it is important to provide comparisons using the test subset of the dataset it was trained on – the BAPPS dataset. Although the LPIPS measures are trained using the  $G$  network that transforms two distances to a fractional preference, when evaluating the measure, only the main network is used. One measure for evaluating IQMs is two-alternative forced choice (2AFC). This is the percentage of images where the image closest in distance to the reference using the specific measure agrees with the majority of human voters.

Dataset	Number of Samples	Number of Distortions
TID2008 [23] Train	1428	17
TID2008 [23] Test	272	17
TID2013 [24]	3000	24
CSIQ [25]	899	6
LIVE [26]	982	5
BAPPS [8] Train	151.4k	425
BAPPS [8] Test	36.3k	425+

**Table 1:** An overview of the datasets used in the paper.

Table 2 contains Pearson and Spearman correlations be-



**Fig. 2:** Receptive field in the Fourier domain for channels that maximise  $\ell_2$  distance between reference and an achromatic version JPEG2000 transmission error distorted image (from Fig. 3).

tween MOS and distance measures on various datasets. The higher the correlations, the more aligned the distance is to how humans judge the distorted images. PerceptNet outperforms all algorithms, including AlexNet with ReLU or GDN on the TID2008 Test and LIVE datasets. AlexNet is the best performing algorithm on a dataset with unseen distortions, such as TID2013 but the number of parameters in AlexNet is two orders of magnitude greater than PerceptNet. The networks trained using BAPPS (LPIPS and PerceptNet tune and scratch) perform similar to traditional IQMs. Replacing the ReLU activation with GDN layers in AlexNet improves correlations on the TID2008 Test and LIVE datasets.

Table 3 contains 2AFC scores for combinations of networks and datasets. Networks trained on perceptual datasets perform poorly on the BAPPS dataset and networks trained on the BAPPS dataset perform poorly on the perceptual datasets containing less distortions (table 2). Training AlexNet and PerceptNet on TID2008, and evaluating on BAPPS, leads to similar results despite AlexNet having a larger number of parameters. Using LPIPS but replacing AlexNet with PerceptNet leads to slightly worse, but still similar performance, when the pretrained networks tuned on BAPPS. Training from scratch on BAPPS, PerceptNet outperforms AlexNet it is able to better generalise to other perceptual datasets.

Fig. 3 shows an example of the output from PerceptNet for a reference and distorted image. Each channel has been scaled to the domain  $[0, 255]$ . The main difference between the channels are where the distortions have taken place. Fig. 2 shows the receptive field in the Fourier domain for the corresponding channels (88 and 64). The fields resemble the Contrast Sensitivity Function of the Spatial Standard Observer in that the channels focus on the mid-frequencies, where humans have maximum sensitivity [3].

## 5. CONCLUSION


We describe a transformation inspired by various stages in the human visual perceptual system that can accurately pre-

Method	Trained On	Pearson Correlation (Spearman Correlation) with MOS			
		TID2008 Test	TID2013	CSIQ	LIVE
$\ell_2$		0.38 (0.53)	0.60 (0.69)	0.70 (0.81)	0.69 (0.94)
SSIM		0.51 (0.53)	0.62 (0.60)	0.77 (0.84)	0.84 (0.95)
MS-SSIM		0.78 (0.80)	0.78 (0.80)	0.81 (0.91)	0.77 (0.97)
FSIM <sub>c</sub>		0.68 (0.72)	0.67 (0.73)	0.63 (0.86)	0.80 (0.95)
NLAPD (with GDN)	TID2008	0.81 (0.82)	0.82 (0.81)	0.90 (0.92)	0.88 (0.96)
AlexNet (with ReLU)	TID2008	0.89 (0.89)	0.93 (0.91)	0.95 (0.95)	0.88 (0.94)
AlexNet (with GDN)	TID2008	0.91 (0.91)	0.92 (0.91)	0.94 (0.95)	0.93 (0.95)
PerceptNet	TID2008	0.93 (0.93)	0.90 (0.87)	0.94 (0.96)	0.95 (0.98)
LPIPS AlexNet (tune)	ImageNet + BAPPS	0.74 (0.75)	0.76 (0.76)	0.88 (0.93)	0.85 (0.96)
LPIPS AlexNet (scratch)	BAPPS	0.47 (0.47)	0.58 (0.57)	0.72 (0.80)	0.77 (0.89)
PerceptNet (tune)	TID2008 + BAPPS	0.67 (0.72)	0.75 (0.76)	0.81 (0.88)	0.85 (0.94)
PerceptNet (scratch)	BAPPS	0.56 (0.67)	0.67 (0.72)	0.77 (0.84)	0.80 (0.93)

**Table 2:** Traditional IQMs and state-of-the-art approaches evaluated on a variety of datasets. We report the Pearson and Spearman correlations between distances obtained using these methods and the MOS. For methods that are feature extractors (AlexNet, PerceptNet) we took the  $\ell_2$  distance between features obtained using the reference and distorted images.

Method	Trained On	2AFC Accuracy (%)						
		Average	Trad- itional	CNN Based	Super Res	Video Deblur	Colour- isation	Frame Interp
LPIPS AlexNet (tune)	ImageNet + BAPPS	69.7	77.7	83.5	69.1	60.5	64.8	62.9
LPIPS AlexNet (scratch)	BAPPS	70.2	77.6	82.8	71.1	61.0	65.6	63.3
LPIPS PerceptNet (tune)	TID2008 + BAPPS	67.8	69.4	81.3	70.6	60.9	61.9	62.6
LPIPS PerceptNet (scratch)	BAPPS	69.2	75.3	82.5	71.3	61.4	63.6	63.2
AlexNet	TID2008	63.2	56.1	77.4	66.1	58.6	61.6	56.2
PerceptNet	TID2008	64.9	58.1	80.5	68.3	59.6	61.6	58.2

**Table 3:** Two-alternative forced choice (2AFC) accuracy scores for various architectures, all evaluated on the BAPPS [8] dataset. The accuracy is the percentage of samples that the method agreed with the majority of human observers.



**Fig. 3:** Difference in the output of the network for a reference image and distortion image. The channels shown are those that are the maximum in  $\ell_2$  distance between the outputs. Each difference in channels was scaled to  $[0, 255]$ . The image is from the TID2008 test set and the distortion is the maximum magnitude for JPEG2000 transmission errors.

dict human perceived distance when images are subject to a number of distortions. This transformation is implemented as a deep neural network. We show that this network can generalise to datasets to more distortions than are present in the training set. It clearly performs better than traditional image quality metrics. Although it has two orders of magnitude less parameters, its performance is similar to the AlexNet net-

work. Visualising the output of the transformation shows that the perceptual space (output) contains a number of desirable properties that are thought to be present in the human visual system. We also show that substituting ReLU layers by GDN layers (inspired by the human visual system) in AlexNet increases its ability to judge perceptual similarity.

## 6. REFERENCES

- [1] A.B. Watson, *Digital images and human vision*, Cambridge, Mass. : MIT Press, 1993.
- [2] P. Teo and D. Heeger, “Perceptual image distortion,” in *IEEE Int. Conf. Im. Proc.*, 1994, vol. 2, pp. 982–986.
- [3] A.B. Watson and J. Malo, “Video quality measures based on the standard spatial observer,” in *Proc. IEEE Int. Conf. Im. Proc.*, 2002, vol. 3, pp. 41–44.
- [4] V. Laparra, J. Muñoz-Marí, and J. Malo, “Divisive normalization image quality metric revisited,” *J. Opt. Soc. Am. A*, vol. 27, no. 4, pp. 852–864, Apr 2010.
- [5] Z. Wang, E. P. Simoncelli, and A. C. Bovik, “Multiscale structural similarity for image quality assessment,” in *Asilomar Conference on Signals, Systems & Computers*. Ieee, 2003, vol. 2, pp. 1398–1402.
- [6] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al., “Image quality assessment: from error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [7] L. Zhang, L. Zhang, X. Mou, and D. Zhang, “Fsim: A feature similarity index for image quality assessment,” *IEEE transactions on Image Processing*, vol. 20, no. 8, pp. 2378–2386, 2011.
- [8] R. Zhang, P. Isola, A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE CVPR*, 2018, pp. 586–595.
- [9] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, “Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness,” *arXiv preprint arXiv:1811.12231*, 2018.
- [10] M. Martinez, M. Bertalmío, and J. Malo, “In praise of artifice reloaded: Caution with image databases in modeling vision,” *Front. Neurosci.*, vol. 13, pp. 1–17, 2019.
- [11] V. Laparra, J. Ballé, A. Berardino, and Simoncelli E P, “Perceptual image quality assessment using a normalized laplacian pyramid,” *Electronic Imaging*, vol. 2016, no. 16, pp. 1–6, 2016.
- [12] M. Martinez, P. Cyriac, T. Batard, M. Bertalmío, and J. Malo, “Derivatives and inverse of cascaded linear+nonlinear neural models,” *PLoS ONE*, vol. 13, no. 10, pp. e0201326, 2018.
- [13] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.
- [14] V. Laparra, A. Berardino, J. Ballé, and E. P. Simoncelli, “Perceptually optimized image rendering,” *Journal Optical Society of America, A*, 2017.
- [15] A. Hepburn, V. Laparra, R. McConville, and R. Santos-Rodriguez, “Enforcing perceptual consistency on generative adversarial networks by using the normalised laplacian pyramid distance,” *arXiv preprint arXiv:1908.04347*, 2019.
- [16] J. Ballé, V. Laparra, and E. P. Simoncelli, “Density modeling of images using a generalized normalization transformation,” *arXiv preprint arXiv:1511.06281*, 2015.
- [17] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [18] M. Carandini and D. Heeger, “Normalization as a canonical neural computation.,” *Nature Reviews. Neurosci.*, vol. 13, no. 1, pp. 51–62, 2012.
- [19] M.D Fairchild, *Color Appearance Models, 2nd Ed.*, Wiley-IS&T, Chichester, UK, 2005.
- [20] A.B. Abrams, J.M. Hillis, and D.H. Brainard, “The relation between color discrimination and color constancy,” *Neural Comp.*, vol. 19, no. 10, pp. 2610–2637, 2007.
- [21] R. Shapley and M.J. Hawken, “Color in the cortex: single- and double-opponent cells,” *Vision Research*, vol. 51, no. 7, pp. 701 – 717, 2011.
- [22] A. B. Watson and J. A. Solomon, “Model of visual contrast gain control and pattern masking,” *JOSA A*, vol. 14, no. 9, pp. 2379–2391, 1997.
- [23] N. Ponomarenko, V. Lukin, A. Zelensky, K. Egiazarian, M. Carli, and F.a Battisti, “Tid2008-a database for evaluation of full-reference visual quality assessment metrics,” *Advances of Modern Radioelectronics*, vol. 10, no. 4, pp. 30–45, 2009.
- [24] N. Ponomarenko, O. Ieremeiev, V. Lukin, K. Egiazarian, L. Jin, J. Astola, B. Vozel, K. Chehdi, M. Carli, F. Battisti, et al., “Color image database tid2013: Peculiarities and preliminary results,” in *European workshop on visual information processing (EUVIP)*. IEEE, 2013, pp. 106–111.
- [25] E. C. Larson and D. M. Chandler, “Most apparent distortion: full-reference image quality assessment and the role of strategy,” *Journal of Electronic Imaging*, vol. 19, no. 1, pp. 011006, 2010.
- [26] H. R. Sheikh, Z. Wang, L. Cormack, and A. C. Bovik, “Live image quality assessment database release 2,” <http://live.ece.utexas.edu/research/quality>, 2005.