Learning Local Distortion Visibility from Image Quality

N.K. Kottayil, Giuseppe Valenzise, Frédéric Dufaux, Irene Cheng

► To cite this version:

N.K. Kottayil, Giuseppe Valenzise, Frédéric Dufaux, Irene Cheng. Learning Local Distortion Visibility from Image Quality. IEEE International Conference on Image Processing (ICIP'2018), Oct 2018, Athens, Greece. 10.1109/icip.2018.8451732 . hal-01791412

HAL Id: hal-01791412

<https://hal.science/hal-01791412>

Submitted on 10 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LEARNING LOCAL DISTORTION VISIBILITY FROM IMAGE QUALITY

Navaneeth Kamballur Kottayil¹, Giuseppe Valenzise², Frederic Dufaux² and Irene Cheng¹

¹ MRC, University of Alberta, ² CNRS, CentraleSupélec, Université Paris-Sud.

ABSTRACT

Accurate prediction of local distortion visibility thresholds is critical in many image and video processing applications. Existing methods require an accurate modeling of the human visual system, and are derived through psychophysical experiments with simple, artificial stimuli. These approaches, however, are difficult to generalize to natural images with complex types of distortion. In this paper, we explore a different perspective, and we investigate whether it is possible to learn local distortion visibility from image quality scores. We propose a convolutional neural network based optimization framework to infer local detection thresholds in a distorted image. Our model is trained on multiple quality datasets, and the results are correlated with empirical visibility thresholds collected on complex stimuli in a recent study. Our results are comparable to state-of-the-art mathematical models that were trained on psychovisual data directly. This suggests that it is possible to predict psychophysical phenomena from visibility information embedded in image quality scores.

1. INTRODUCTION

Predicting the visibility of visual distortion is of paramount importance in a number of image processing applications, such as image compression [1, 2], watermarking [3] and quality assessment [4, 5]. Models of contrast detection and masking have been widely studied in the literature, e.g., [4, 6, 7]. These studies are based on an accurate modeling of the human visual system (HVS), including intra-ocular scattering, contrast sensitivity at different frequencies, luminance masking, intra- and inter-channel masking, etc.

Traditionally, measurements of distortion visibility are obtained through psychophysical experiments employing simple, artificial stimuli, such as Gabor patterns, sine-wave gratings, or wide-band noise [8, 6]. The simplicity and non-naturalness of the stimuli enable to describe them through well-defined features (e.g., frequency, brightness, etc.). These features are then used to derive mathematical models of distortion visibility. In order to understand how these observations generalize to the case of natural image masks, Alam *et al.* [9] have recently presented a dataset of local masking for natural scenes. Specifically, they use Gabor patterns as stimuli and natural images as masks to measure detection

thresholds – the minimum magnitude of the stimulus making it distinguishable from the background mask – for small image patches. Later, this data has been used to train a convolutional neural network in order to predict the detection threshold of a patch from its input pixel values [10].

Conventional approaches to model distortion visibility strongly rely upon psychophysical experiments that are, in their nature, based on a simplification of real-world conditions. For example, models that describe visibility of sine-wave gratings might be enough to predict visibility in DCT-based image compression; however, they are probably failing in modeling different or multiple concurrent artifacts. Furthermore, local visibility is influenced by surrounding regions, and is ultimately linked to image semantics. It is evident that modeling all these complex factors *only* through psychophysical experiments is unfeasible.

In this work, we explore a different perspective. Instead of learning distortion visibility directly from psychophysical data, we propose to learn it *indirectly*, leveraging the large availability of alternative, yet related, data: subjectively annotated image quality assessment (IQA) datasets. Image quality scores provide higher level information about the visual appearance of a picture, compared to psychophysical measurements. At the same time, they bring information about the visibility of distortion. Indeed, a common assumption in image quality assessment is that the perceived quality is directly related to the visibility of the error signal [11, 12]. In other words, the per pixel error is weighted locally by the ensemble of perceptual phenomena, such as contrast sensitivity and several forms of masking, which discount its visibility to the human visual system.

Visual quality scores thus implicitly embed latent information on error visibility. In our recent work [13], we have proposed a deep convolutional neural network (CNN) architecture to disentangle the per pixel distortion and what we called the “perceptual resistance”, in the context of no-reference quality estimation of high dynamic range compressed pictures. Our results demonstrated that it is possible to effectively estimate these two terms over a broad range of qualities, starting from supra-threshold quality scores. In this paper, we investigate whether a similar approach can also be used to predict near-threshold visibility. To this end, we train our proposed system in a full-reference fashion, i.e., we assume that the error signal is known, as it is the case in most

IQA datasets. However, the inference step does not require the knowledge of the error, and can produce an estimate of local masking for any input image. Interestingly (and perhaps surprisingly), we find that perceptual scaling learned from image quality scores can predict the detection thresholds in [9] with similar accuracy as the CNN-based regressor in [10], although our model is learned on other datasets with different contents and several kinds of visual impairments. This makes the proposed approach potentially more general than previous work.

The rest of the paper is organized as follows. In Section 2, we describe our model to derive visibility threshold from image quality scores. In Section 3, we compare our predictions to state-of-the-art models on a local masking visibility dataset. Finally, we draw conclusions in Section 4.

2. PROPOSED MODEL

In this section we present our proposed model to estimate local distortion visibility thresholds. Specifically, we first discuss the assumptions of our model and provide a mathematical framework to split (supra-threshold) quality scores into per pixel error and a perceptual scaling term, which we show in Section 3 to be a good predictor of visibility thresholds in local masking. Afterwards, we describe how to implement this model using a deep convolutional neural network architecture.

2.1. Mathematical framework and assumptions

Let I_R and I_D be an original reference image and its distorted version, respectively, and $Q \in [0, 1]$ the quality score for I_D . Without loss of generality, we assume that Q is given as a differential mean opinion score (DMOS). We assume that we have access to the local quality $q(i, j)$ of an image patch $I_D(i, j)$ of size $N \times N$ pixels, centered at location (i, j) . The pooling process linking the quality of individual patches to the overall quality Q may depend on many factors, e.g., saliency [14]. Here, we assume that the local quality of an image block is the same as the global image quality score, similarly to the setting in [15]. While this is a strong assumption, which is often not met in practice, it has been proved to be accurate enough to predict image quality [15, 13].

In order to model local quality, we further assume that per pixel distortion in a patch is discounted by a perceptual weight, $T(i, j)$, that accounts for typical masking and visibility effects. Specifically, we measure pixel distortion through the average absolute error $E(i, j)$ of a patch, defined as:

$$E(i, j) = \frac{1}{N^2} \sum_{k=1}^{N^2} |I_D(i, j)_k - I_R(i, j)_k|, \quad (1)$$

where k is the pixel index in the patch. We then approximate local quality as a function of the error and the perceptual

Fig. 1: Neural network architecture for extracting contrast detection thresholds from IQA databases.

weight, that is:

$$q(i, j) \approx 1 - \exp \left(\left| \frac{\alpha \cdot E(i, j)}{T(i, j)} \right|^\beta \right), \quad (2)$$

where α and β are two model parameters.

Notice that this formulation has been previously used to model the probability of detecting localized noise distortion, e.g., in [11], but we introduce an additional scaling factor α . Eq. (2) is inspired by the common practice in vision science of expressing the magnitude of the error in multiples of the just-noticeable difference [2], although the relationship in this case is nonlinear. We therefore refer to $T(i, j)$ as the *local visibility threshold*, even if this is technically correct only for near-threshold distortion. We show in Section 3 that, for this latter case, T does indeed model local visibility. Moreover, the thresholds we estimate are not in the same scale as those obtained through psychophysical experiments. We compensate for this by means of the parameter α , which provides an additional degree of freedom in the optimization to push the predicted local quality as close as possible to ground truth. The value of the parameter β is generally found by matching the results of psychophysical experiments. In practice, we found that the choice of β does not significantly affect the performance of our model, and α alone provides already enough flexibility to minimize the loss function. Hence, we simply set $\beta = 1$ in Eq. (2).

2.2. Implementation

A scheme of the proposed model is shown in Figure 1. The input of our systems are overlapping patches of 32×32 pixels, similar to [15, 13]. The distortion, $E(i, j)$, is computed directly as in Eq. (1) during the training, where both I_D and I_R are available. The local visibility thresholds $T(i, j)$ are computed in a module that we name “P-net”, while an estimate of local quality, $\hat{q}(i, j)$ is obtained by implementing Eq. (2) in the “Mixing function” block. Notice that this structure is required for training the P-net, as T is considered a latent variable which depends implicitly on the observations of the input content and perceived quality. Instead, for inference the P-net is employed as a standalone block. Furthermore, we are generally interested in applying the learned P-net on the *original* pictures, rather than on the distorted ones. However, we found

that training the P-net with noisy versions of the image was more effective, as this increases the variability of input data, leading to improved generalization capabilities.

Our model is trained to predict local quality $\hat{q}(i, j)$ using the ground-truth quality $q(i, j)$ as targets, by minimizing the following cost function:

$$J(i, j) = |q(i, j) - \hat{q}(i, j)|. \quad (3)$$

Notice that $\hat{q}(i, j)$ depends implicitly on the latent variables $T(i, j)$ through Eq. (2). Thus, when optimizing J , the visibility thresholds are adjusted in such a way to weigh the error coherently with the observed ground-truth quality.

For the architecture of P-net, we make use of a hand-crafted layer that we named as augmented input layer [13]. In this layer, in addition to the luminance values of the $N \times N$ block, we compute the mean, variance and Mean Subtracted Contrast Normalized (MSCN) image [16]. The latter is defined as:

$$MSCN(x, y) = \frac{I(x, y) - \mu_M[I(x, y)]}{\sigma_M[I(x, y)] + \epsilon}, \quad (4)$$

where (x, y) denotes the location of a pixel in the patch, $\mu_M[I(x, y)]$ is the mean and $\sigma_M[I(x, y)]$ the variance of the patch, computed by replacing every pixel (x, y) with the mean and variance, respectively, over a local Gaussian window of size $M \leq N$ ($M = 7$ in our case) around (x, y) . The regularization term ϵ is set to 0.01. This is followed by one convolutional layer with $32 \ 5 \times 5$ filters and a fully connected layer with 100 nodes. We use *relu* activation in all neurons. Dropout layers are used to prevent overfitting. The implementation of the network can be found at <https://github.com/kamballu/Deep-Visibility>.

3. RESULTS AND ANALYSIS

In order to assess how well the estimated distortion visibility predicts ground-truth data from psychophysical experiments, we test the proposed model on the dataset of local masking thresholds in [9]. This dataset collects measured threshold values for 1080 image patches of size 85×85 pixels, extracted from the CSIQ dataset [17]. The detection thresholds are reported in terms of root-mean-squared (RMS) contrast and expressed in decibels (dB).

To test the proposed model, we train it on three different datasets: the CSIQ dataset [17], containing 855 images, 6 types of distortions; the TID 2013 dataset [18], with 3000 images and 25 types of distortion; and the LIVE dataset [19], featuring 779 images and 5 types of distortion. In general, the thresholds $T(i, j)$ found with our model do not lie on the same scale as those in [9]. In addition, the P-net can be trained on different IQA datasets, and the interpretation of DMOS in each dataset depends on the experiment carried out to collect the data [20]. Following a typical protocol in the evaluation of quality metrics [21], we compensate for this mismatch

Fig. 2: Scatter plots of contrast detection thresholds derived using our method vs experimentally measured values. The system is trained on (a) LIVE dataset (b) TID 2013 dataset. The polynomial fitting line is shown in red. Red points correspond to patches whose luminance is outside the range of the training datasets, see Figure 3. r is the PLCC (after fitting) on the whole test set; r' is the PLCC excluding the red points.

by linearizing the predictions with respect to psychophysical groundtruth through a monotonic third-order polynomial fitting before evaluating their statistical accuracy.

3.1. Performance

A comparison of the predicted and ground-truth local visibility thresholds is illustrated in the scatter plots of Figure 2, where our P-net has been trained on the LIVE and TID 2013 datasets, respectively. Each point in the scatter plot represents a 85×85 patch of [9]. Since our predictor can produce per pixel estimates of distortion visibility (using overlapping patches), we decimate the maps produced by the P-net to match the resolution of the ground truth, using a simple averaging filter (see Figures 4 and 5).

We can observe from Figure 2 that the predicted thresholds capture relatively well the overall trends of the measures obtained from psychophysical experiments, even if they have been obtained by training on very different data (IQA scores) and using different source contents. This indicates that learning visibility thresholds from generic IQA datasets is feasible and can generalize sufficiently well. Nevertheless, we notice in Figure 2 that in some cases the predicted thresholds deviate significantly from the measured ones. Especially for the LIVE dataset, this degrades the performance, measured by the Pearson linear correlation coefficient (PLCC) $r = 0.56$. To further investigate this, we analyze the distribution of the average luminance intensity of patches in the TID and LIVE datasets compared to CSIQ in Figure 3. Notice that there is only a negligible amount of patches for the LIVE dataset in the intensity range $[0, 10]$ and $[250, 255]$. This lack of data can affect the performance of prediction in this specific intensity range. We verify this by highlighting in red those patches of the test set having luminance outside the interval $[10, 250]$ in Figure 2. We observe that these points correspond indeed to the outliers in the scatter plot. By removing these few very

Fig. 3: Distribution of average intensity of 32×32 patches in LIVE, CSIQ and TID datasets. The distributions of the three datasets overlap only in the intensity range $[10, 250]$.

Method	Training Data	RMSE
Watson <i>et al.</i>-KMF [6]	Pshyophysical visibility experiments	5.713
Watson <i>et al.</i> -JYS [6]		6.521
Teo & Heeger [22]		6.861
Chandler <i>et al.</i> [1]		6.879
Optimized GC [10]		5.192
Alam <i>et al.</i> CNN [10]		5.475
Proposed	CSIQ quality scores [17]	5.691
	LIVE quality scores [19]	5.991
	TID 2013 quality scores [18]	5.626

Table 1: Performance comparison between different algorithms. Highlighted are the best state-of-the-art methods (handcrafted and CNN-based), as well as our results.

dark or very bright patches (28 patches out of 1080), we observe that the PLCC increases to $r' = 0.68$.

We compare in Table 1 the performance of our method (trained with three different IQA datasets) with other predictors of local visibility thresholds proposed in the literature. The majority of these methods use handcrafted models directly derived from psychophysical experiment, while [10] employs a convolutional neural network. The performance criterion is root-mean-squared error (RMSE) between the ground-truth and predicted thresholds. We observe that our approach provides comparable or better results to state-of-the-art methods. It should be noticed that the two methods in [10] are trained on the same dataset [9] used for test, and thus their performance represents a sort of accuracy upper bound. Conversely, the proposed method achieve competitive results even if it is trained on the TID 2013 or LIVE datasets.

A qualitative illustration of the predicted thresholds is reported in Figure 4, where we also show per picture PLCC. We observe that the perceptual thresholds produced by our approach are intuitive, e.g., thresholds are lower for relatively smooth regions (sky) and higher for more complex regions (grass, leaves). There are also cases in which the prediction fails (examples in Figure 5). This happens in dark patches, as discussed above, where the training datasets do not offer sufficient samples to learn robustly the visibility thresholds.

Fig. 4: Examples of local visibility thresholds produced by our method trained on TID 2013. The estimated thresholds are decimated to match the resolution of the ground truth. Both sets of values are scaled to fit in the range $[0, 1]$ for visualization, with 0 (black) being the lowest threshold.

Fig. 5: Some failure cases of the proposed method trained on TID 2013. Often, these cases correspond to patches with low luminance, which are underrepresented in the training data.

4. CONCLUSION

We present a method to derive local visibility thresholds from image quality scores using a neural-network-based approach. We do not use the result of psychovisual experiments. Our experiments demonstrate that the latent information about distortion visibility carried by supra-threshold quality scores can be recovered and used to predict near-threshold local masking. One advantage of our approach, compared to models based on psychophysical data, is that it can leverage the larger availability of subjectively annotated image quality datasets. We plan to formalize further this approach in the future and apply it to objective image quality assessment.

5. REFERENCES

- [1] D. M. Chandler, M. D. Gaubatz, and S. S. Hemami, "A Patch-Based Structural Masking Model with an Application to Compression," *EURASIP Journal on Image and Video Processing*, vol. 2009, no. 1, pp. 1–22, 2009.
- [2] A. B. Watson, "DCTune: A technique for visual optimization of DCT quantization matrices for individual images."
- [3] A. Koz and A. Alatan, "Oblivious Spatio-Temporal Watermarking of Digital Video by Exploiting the Human Visual System," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 18, no. 3, pp. 326–337, 3 2008.
- [4] S. Daly, "Visible differences predictor: an algorithm for the assessment of image fidelity," in *SPIE/IS&T 1992 Symposium on Electronic Imaging: Science and Technology*, 1992, pp. 2–15.
- [5] D. Chandler and S. Hemami, "VSNR: A wavelet-based visual signal-to-noise ratio for natural images," *Image Processing, IEEE Transactions on*, vol. 16, no. 9, pp. 2284–2298, 2007.
- [6] A. B. Watson and J. A. Solomon, "Model of visual contrast gain control and pattern masking," *Journal of the Optical Society of America. A, Optics, image science, and vision*, vol. 14, no. 9, pp. 2379–91, 9 1997.
- [7] T. Carney, C. W. Tyler, A. B. Watson, W. Makous, B. Beutter, C.-C. Chen, A. M. Norcia, and S. A. Klein, "Modelfest: year one results and plans for future years," B. E. Rogowitz and T. N. Pappas, Eds., vol. 3959. International Society for Optics and Photonics, 6 2000, pp. 140–151.
- [8] G. E. Legge and J. M. Foley, "Contrast masking in human vision," *Journal of the Optical Society of America*, vol. 70, no. 12, pp. 1458–71, 12 1980.
- [9] M. M. Alam, K. P. Vilankar, D. J. Field, and D. M. Chandler, "Local masking in natural images: A database and analysis," *Journal of Vision*, vol. 14, no. 8, pp. 22–22, 7 2014.
- [10] M. M. Alam, P. Patil, M. T. Hagan, and D. M. Chandler, "A computational model for predicting local distortion visibility via convolutional neural network trained on natural scenes," in *2015 IEEE International Conference on Image Processing (ICIP)*. IEEE, 9 2015, pp. 3967–3971.
- [11] T. Zhu and L. Karam, "A no-reference objective image quality metric based on perceptually weighted local noise," *EURASIP Journal on Image and Video Processing*, vol. 2014, no. 1, pp. 1–8, 2014.
- [12] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image Quality Assessment: From Error Visibility to Structural Similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [13] N. K. Kottayil, G. Valenzise, F. Dufaux, and I. Cheng, "Blind Quality Estimation by Disentangling Perceptual and Noisy Features in High Dynamic Range Images," *IEEE Transactions on Image Processing*, vol. 27, no. 3, pp. 1512–1525, 3 2018.
- [14] A. K. Moorthy and A. C. Bovik, "Visual Importance Pooling for Image Quality Assessment," *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 2, pp. 193–201, 4 2009.
- [15] L. Kang, P. Ye, Y. Li, and D. Doermann, "Convolutional neural networks for no-reference image quality assessment," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1733–1740.
- [16] A. Moorthy and A. Bovik, "Blind image quality assessment: From natural scene statistics to perceptual quality," *Image Processing, IEEE Transactions on*, vol. 20, no. 12, pp. 3350–3364, 2011.
- [17] E. Larson and D. Chandler, "Most apparent distortion: a dual strategy for full-reference image quality assessment," in *Proc. SPIE*, vol. 7242, 2009.
- [18] N. Ponomarenko, L. Jin, O. Ieremeiev, V. Lukin, K. Egiazarian, J. Astola, B. Vozel, K. Chehdi, M. Carli, F. Battisti, and Others, "Image database TID2013: Peculiarities, results and perspectives," *Signal Processing: Image Communication*, vol. 30, pp. 57–77, 2015.
- [19] H. R. Sheikh, M. F. Sabir, and A. C. Bovik, "A Statistical Evaluation of Recent Full Reference Quality Assessment Algorithms," *IEEE Transactions on Image Processing*, vol. 15, no. 11, pp. 3440–3451, 2006.
- [20] E. Zerman, V. Hulusic, G. Valenzise, R. Mantiuk, and F. Dufaux, "The relation between MOS and pairwise comparisons and the importance of cross-content comparisons," in *Human Vision and Electronic Imaging Conference, IS&T International Symposium on Electronic Imaging (EI 2018)*, 2018.
- [21] ITU-T, "Methods, metrics and procedures for statistical evaluation, qualification and comparison of objective quality prediction models," ITU-T Recommendation P.1401, 7 2012.
- [22] P. Teo and D. Heeger, "Perceptual image distortion," in *Proceedings of 1st International Conference on Image Processing*, vol. 2. IEEE Comput. Soc. Press, pp. 982–986.