

Image Quality Assessment: Learning to Rank Image Distortion Level

Shira Faigenbaum-Golovin, Duke University, North Carolina, USA
Or Shimshi, Tel-Aviv, Israel

Abstract

Over the years, various algorithms were developed, attempting to imitate the Human Visual System (HVS), and evaluate the perceptual image quality. However, for certain image distortions, the functionality of the HVS continues to be an enigma, and echoing its behavior remains a challenge (especially for ill-defined distortions). In this paper, we learn to compare the image quality of two registered images, with respect to a chosen distortion. Our method takes advantage of the fact that at times, simulating image distortion and later evaluating its relative image quality, is easier than assessing its absolute value. Thus, given a pair of images, we look for an optimal dimensional reduction function that will map each image to a numerical score, so that the scores will reflect the image quality relation (i.e., a less distorted image will receive a lower score). We look for an optimal dimensional reduction mapping in the form of a Deep Neural Network which minimizes the violation of image quality order. Subsequently, we extend the method to order a set of images by utilizing the predicted level of the chosen distortion. We demonstrate the validity of our method on Latent Chromatic Aberration and Moiré distortions, on synthetic and real datasets.




Figure 1. Automatically ordered images (a)-(c) from the TID2013 [15] by their predicted Lateral Chromatic Aberration level (0, 2, 4 pixels accordingly) using the proposed method in this paper; the “<” sign indicates that the left image is less distorted than the one on the right; (d) a plot of the expected (the x-axis) versus the predicted image order (the y-axis).

1. Introduction

Evaluating the quality of an image is a vital task in the domain of image processing. This is crucial for measuring the performance of image processing algorithms, which may unintentionally damage image quality (e.g., a denoising algorithm can reduce the sharpness of the edges). While talking about image quality, often questions arise regarding an image’s quality assessment: like *What is a good image?* And also *Why is image quality assessment so difficult?* [20]. While the Human Visual System often provides an answer to these questions, it is not efficient and at times prone to subjective judgment.

In order to avoid repeated evaluation by image quality experts, an automatic image quality procedure is essential. Image quality (IQ) assessment can be addressed either by absolute or relative measure. In the first method, a number representing the IQ of a single image is computed. In the second method, given two images,

we indicate which image looks better. It should be noted that usually providing a relative measure is easier, as opposed to evaluating the quality of a single image and scoring it based on its defects. The reason for this is that there is no need to specify what image characteristics influenced the scoring.

Along the years there were various algorithms suggested to crack the enigma of Human Visual System (HVS), which constantly aids humans in this task [10, 13]. Those studies paved the way towards designing an objective procedure for image quality evaluation [2-3, 7, 13, 22]. However, measuring the perceptual image quality still remains a challenge. In recent years, the interest in image quality metrics was renewed, with the rise of Deep Neural Networks (DNN). The main idea of DNN is that an image quality metric can be defined and evaluated on images. Subsequently, a DNN is constructed to associate between the image and its calculated quality. Recently published papers [4, 6, 8-9, 12, 18, 21, 23] demonstrate the benefits of using DNN for IQ evaluation. These methods require the defining and evaluating an absolute image quality metric. This requires the designer to develop a metric and evaluate the distortion as a pre-processing step of DNN. However, the main challenge is that certain image distortions are ill-defined (though they are easy to acquire or simulate), and therefore, no metric exists to evaluate them.

In this paper, we address the question of image quality assessment by introducing a framework for learning *relative IQ*. We look on image distortion as the degradation of the ideal image or as a deviation from the “perfect” image. Thus, we propose a relative-order-preserving image quality, and bypasses the challenge of defining the desired distortion.

Our method was inspired by [19], in which the semantic hierarchy of words, and sentences was learned. Although this particular hierarchy is based on the hypernymies of words, for our purposes, the aim is to maintain the order of images based on their quality. In what follows, we introduce our relative IQ method (subsection 2.1) and later it is extended to rank a set of images (subsection 2.2). Then, in section 3, we describe the construction of the training, validation, and testing datasets. The validity of our method is demonstrated through the order-preserving dimension reduction of two distortions: Chromatic Aberration and Moire (section 4). The paper concludes with a discussion of future directions for methodological enhancements (section 5).

2. Proposed Method

2.1 Ranking the Distortion of an Image Pair

First, we introduce the relative image quality measure in reference to the question “given two images, A, and B as well as image distortion d , is image B more distorted than image A with respect to d ”? We answer this question by defining image quality order with respect to the selected distortion. Subsequently, we look for a dimension reduction mapping from the image domain to a the

natural numbers, which maintain the IQ order. This mapping is later utilized to make order in an unseen image pair. In mathematical terms, we define the *IQ-order* of two given *Regions of Interest* (ROI) of an image as:

Definition 1. Let R_A and R_B be two registered ROIs, and let d be the distortion we would like to evaluate. Then, **ROI-IQ-order** is defined as $R_A <_d R_B$, meaning that ROI R_A is less distorted than ROI R_B , with respect to d .

Let us now extend this definition to the entire image:

Definition 2. Let A and B be two images, as well as the set of their registered ROIs $\{R_{A,i}\}, \{R_{B,i}\}$. If the following condition stands $\#\{i \mid R_{A,i} <_d R_{B,i}\} > \#\{i \mid R_{A,i} >_d R_{B,i}\}$ then we define **IQ-order** as $A <_d B$, i.e., image A is less distorted than image B with respect to the tested ROIs, and the selected distortion.

Next, we defined IQ-order-preserving mapping as:

Definition 3. Let A and B be two images, such that $A <_d B$ and let $S = \{(R_{A,i}, R_{B,i}) \mid R_{A,i} <_d R_{B,i} \}_{i=1..N}$ be their set of ordered registered ROI pairs, and $f: \mathbb{R}^{n \times m} \rightarrow \mathbb{R}$ a dimensional reduction function. We say that the mapping f is **IQ-order-preserving** if for any ROI pair in S , f is order-preserving. That is if $\forall i, R_{A,i} <_d R_{B,i} \Rightarrow f(R_{A,i}) < f(R_{B,i})$.

Now, the image quality question can be formulated with respect to the order-preserving mapping.

Problem definition: Let A, B be two images such that $A <_d B$ and also let $S = \{(R_{A,i}, R_{B,i}) \mid R_{A,i} <_d R_{B,i} \}_{i=1..N}$ be an ordered registered ROI pair set. The IQ-order-preserving mapping f , is found such that it will minimize the image quality order violation:

$$(1) \quad f = \underset{f: \mathbb{R}^{n \times m} \rightarrow \mathbb{R}}{\operatorname{argmin}} \frac{1}{N} \sum_{(R_{A,i}, R_{B,i}) \in S} E(R_{A,i}, R_{B,i}),$$

where the loss function, E , for an ordered pair (R_A, R_B) is defined as:

$$E(R_A, R_B) = \max(0, (f(R_A) + \epsilon) - f(R_B))^2.$$

Once the optimal f is found (see subsection 2.3 for details), given a pair of images we can find the less distorted image using definition 3.

2.2 Ranking the Distortion of Image Set

The definition of “order” on an image pair can be extended to rank image sets, utilizing their distortion levels. Once the mapping f , which minimizes equation (1), is found, it can be used to calculate the relative score of images. As a result, one can order a given image set with respect to the values of f . Specifically, for a set of images $\{A_j\}_{j=1..J}$, with the corresponding ROIs $R_{j,i}$, we calculate $f(R_{j,i})$. Subsequently, the ranking of this set with respect to a given distortion is achieved by ordering the values $f(R_{j,i})$ for each specific ROI index (i), and then calculating the median of the ranking across all the image patches (j).

We illustrate this procedure in the following example. We would like to rank four images using three patches with given predicted IQ values. Table 1 demonstrates all the steps of the process starting with (a) predicting the IQ values using some learned f (each

row in the matrix appearing in the first column corresponds to a different patch, and each column to corresponds to a different image), then (b) ranking their ROI’s and later (c) ranking the four images. We demonstrate this concept in a test-case scenario by ranking images with respect to Chromatic Aberration distortion in Figure 1. Images (a)-(c) are ordered according to their rank, and (d) is a plot of the expected ranking versus the predicted one (as can be seen the prediction, in this case, is perfect).

Table 1: Example of ranking four images, with respect to their three patches

Predicted IQ values	Calculated Patches Ranks	Calculated Images rank
[1 2 3 4; 4 8 9 12; 2 3 5 4]	[1 2 3 4; 1 2 3 4; 1 2 4 3]	[1 2 3 4]

2.3 Network Architecture

In this study we design the order-preserving mapping as a Deep Neural Network (DNN). It comprises of a Siamese network of a pair of ResNet architecture [11], each performing a dimension reduction. Later, a loss function, which maximizes the distance between mismatches of the dimension reduction via the equation (1), is calculated (the loss is also called squared negative smoothed hinge loss (SNSHL) [16]). See Table 2 for the detailed network architecture.

Table 2: Architecture of the IQ order-preserving network. Building blocks are shown in brackets (and consists of three consequent ReLU’s), with the numbers of blocks stacked. The network input is two concatenated color patches of 32×32.

Layer name	Output size	Order-preserving Net104-layer	
slice		slice point 3	
conv1	16 × 16 × 2	7 × 7,64, stride 2	7 × 7,64, stride 2
conv2.x	8 × 8 × 2	$\begin{bmatrix} 1 \times 1,64 \\ 3 \times 3,64 \\ 1 \times 1,256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1,64 \\ 3 \times 3,64 \\ 1 \times 1,256 \end{bmatrix} \times 3$
conv3.x	4 × 4 × 2	$\begin{bmatrix} 1 \times 1,128 \\ 3 \times 3,128 \\ 1 \times 1,512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1,128 \\ 3 \times 3,128 \\ 1 \times 1,512 \end{bmatrix} \times 4$
conv4.x	2 × 2 × 2	$\begin{bmatrix} 1 \times 1,256 \\ 3 \times 3,256 \\ 1 \times 1,1024 \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1,256 \\ 3 \times 3,256 \\ 1 \times 1,1024 \end{bmatrix} \times 6$
conv5.x	1 × 1 × 2	$\begin{bmatrix} 1 \times 1,512 \\ 3 \times 3,512 \\ 1 \times 1,2048 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1,512 \\ 3 \times 3,512 \\ 1 \times 1,2048 \end{bmatrix} \times 3$
conv1	1 × 1 × 2	1 × 1,1, stride 2	1 × 1,1, stride 2
Squared negative smoothed hinge loss	1 × 1 × 1	1 × 1,1, stride 2	1 × 1,1, stride 2

2.4 Accuracy Evaluation of the Predicted Order

Measuring predicted order accuracy can be separated into two different scenarios: (a) for a pair of images, and (b) for an image set. While the first one can be evaluated as the True Positive (TP) percentage, the accuracy of the second is performed using Spearman's correlation coefficient [14], which is widely utilized to detect trends in data in light of reference information. Thus, given two measurements vectors x , y , and their corresponding ranks, r_x and r_y , Spearman's correlation coefficient is calculated as

$$\rho = \frac{cov(r_x, r_y)}{\sigma_{r_x} \sigma_{r_y}},$$

where σ_{r_x} , σ_{r_y} are the standard deviations of the rank variables.

Subsequently, for patch i , we apply Spearman's correlation of the ranks of the predicted values $f(R_{j,i})$ as well as a monotonically increasing sequence with equal length (resulting in a correlation coefficient ρ_i). The accuracy of the predicted image set ranking is the median of ρ_i across all images patched. In the example above, the correlation coefficients of the patches ranking (middle column in Table 1) are [1 1 0.8], with $median(\rho_i) = 1$. Therefore, we conclude that there is a monotonicity trend in the data, and that the predicted rank is perfect.

3. Database Creation

Our experimental flow consisted of the following steps: (a) acquiring a dataset of images for training and validation; (b) pairing or simulating images of the same scene, each corresponding to different levels of distortion, and accompanied with a predefined IQ-order; (c) extracting ROIs which contain the desired distortion; (d) learning the order-preserving dimensional reduction function, f ; and (e) executing the method on various datasets. The general flow is illustrated in Figure 2.




Figure 2. Our experimental flow consisting of (a) acquiring a database of images with a chosen distortion (subsection 3.1); (b) creating pairs of images which each corresponding to different levels of the chosen distortion (subsection 3.2); (c) finding areas which correspond to high values of the distortion and cropping ROIs of the size 32x32 (subsection 3.3); (d) learning the order-preserving dimension reduction; and (e) executing it on new datasets.

3.1 Database Acquisition

We tested our methodology on learning to rank the Lateral Chromatic Aberration and Moire distortions. Since no dataset of pairs of images accompanied by the level of distortion was available for training-validation-testing purposes, we created our own dataset both by simulating a distortion on existing images (for Lateral Chromatic Aberration) and by creating synthetic images (for Moire). In the first scenario, we took an existing set of images (an ImageNet dataset [5]) as a baseline and created a pair of images for each image in the dataset, each with a random level of distortion. For the Moire case, we created images with repetitive pattern and then produced an image pair with different distortion levels (see below). After the model was trained, we tested it on a set of real

images from the TE42.v2 chart (designed and produced by Image Engineering [1]).

3.2 Distortion Simulation

We verified the validity of the proposed order-preserving method on Lateral Chromatic Aberration (LCA) and Moire distortions (Figure 3). The LCA distortion appears when the color convergence point is not unique (which stems from the failure of the lens to focus). This effect is seen as a blur or "rainbow" edge in areas of contrast. The LCA dataset was constructed by distorting the ImageNet dataset [5], where the RGB channels of each image were changed with a random shift of size, i.e., $\sim U(1,5)$ pixel in one of the square's diagonal directions.

The Moire distortion (or aliasing) is an effect that causes different signals to become indistinguishable when sampled [17]. It occurs when repetitive patterns of high spatial frequencies exist, which are sampled at different frequencies. Since natural images usually do not depict a constant frequency which could train the Moire distortion, we had to create a synthetic image dataset with constant high-frequency patterns. Our dataset contained the following simulated repetitive patterns (i.e., resolution bars, resolution net, Siemens-star, resolution wedges, concentric rings). In order to simulate the Moire effect, we used an image resize with bicubic interpolation without an antialiasing option along with a randomly sampled resize factor $\sim U(1.5,10)$.




Figure 3. Left: Image with Chromatic Aberration distortion from the TID2013 dataset [15]; right: Resolution chart with Moire effect. Marked in red are the areas with the desired distortion.

3.3 ROIs Extraction

Once IQ-order-preserving image pair $\langle A, B \rangle$ was created (by applying two random levels of the chosen distortion), the IQ-order-preserving ROIs (size 32x32) were extracted. The ROI pairs $\langle R_{A,i}, R_{B,i} \rangle$ were chosen as those corresponding to the maximal values on the error map: $ErrMap = |A - B|$. We chose patches with a sufficient amount of distortion (we chose an ROI pair with $ErrMap$ larger than 0.025% of the image area). Thus, the ROI location selection is content dependent, and may be effected by the image saliency.

As a result, our constructed dataset for each distortion consisted of about six million image patches.

4. Experimental Results

We demonstrate the validity of the proposed method through learning two order preserving mappings, each corresponding to a different distortion (either the LCA or Moire distortions). We trained a DNN, described in subsection 2.3, using a dataset of ordered image pairs (discussed in section 3). The optimal ordering mapping was later used to (a) predict the IQ order of image pairs, and then (b) to rank registered sets of images from new datasets.

While the training and validation of the DNN performed well on synthetic data (Table 3), a question that remained was *what the accuracy of the model is when applied to real life images*. Unfortunately, a real-life dataset of distortions, with image rank, was not available. Therefore, we acquired a new dataset, which consisted of two tests sets of images of the TE42v2 chart [1]. One set depicted the TE42 chart in various LCA levels, and the other contained images with different Moire effects (about 13 images of each set were used for each distortion). Subsequently, the images were ranked by an independent image quality expert using his HVS. Later, we sampled different patches from the images using Monte Carlo to enrich our dataset. In order to simulate different rank tests, each pair was randomly cropped to a size of 150x150 pixels with the desired distortion. For each cropped pair, we predicted the IQ metric and ranked the cropped patches. The results are summarized below (see also Table 3).

We also used another dataset (entitled the TID2013 dataset [15]) to evaluate the performance of image set ranking. This dataset consisted of various images, each subjected to a set of image distortions at various levels.

4.1 Accuracy of Ranking Image Pair

Our trained models gave results of 97% and 94% TP for the simulated test data of the LCA and Moire distortions, respectively. Later, we used the images of the TE42 chart for evaluating the order prediction of real-life scenarios. With the 13 images, we simulated 150 different Monte Carlo experiments by sampling **pairs** of images from the test set, which contained the required distortion. For each cropped pair, we predicted the IQ metric and ranked the cropped patches. This experiment resulted in 80% and 85% TP for the LCA and Moire, respectively. This result indicates that, although the model was trained on synthetic data, it also performs well on real life images.

4.2 Accuracy of Ranking Image Set

We later turned to ranking sets of images using the methodology described in 2.2. Our first test was performed on the TID2013 dataset [15], using images created with chromatic aberrations (distortion marked as No. 23 in the dataset). The dataset contained images with chromatic aberration Levels 1–5; unfortunately, the levels of LCA used for the TID2013 creation were not specified. Based on our examination, Level 5 of LCA in TID2013 was more than 5 pixels (the value used in our training set) and therefore we did not use it in the ranking procedure.¹ But since the images were labeled according to distortion levels, we could check our prediction’s accuracy against that of the expected one. The median ranking for the 25 image sets resulted in an accuracy correlation of $\rho = 1$. Examples of order-preserving ranking can be seen in Figure 1 and Figure 4. We saw that the distortion level of the images with the airplane (Figure 1) and the woman were predicted perfectly, whereas for the parrot, the scoring was good, with images (c)-(e) receiving the same score.




Figure 4. Images (a)-(e) ordered by the rank of the predicted IQ measure from the TID2013 dataset and (f) inspected ROIs, chosen via the error map described in subsection 3.3; (g) a graph with the expected (x-axis) versus predicted rank values (y-axis); the accuracy correlation of those sets is $\rho = 0.8,1$ (from the top to bottom rows).

Subsequently, we tested our ranking methodology on the ordered images depicting the TE42.v2 chart (discussed above). We simulated 150 image sets by sampling quadruplets of images from the original dataset and cropping them randomly to a size of 150x150 pixels, depicting the desired distortion. For each cropped quadruplet, we predicted the IQ metric and ranked the cropped patches. The experiment resulted in a median accuracy of $\rho = 0.7$ and $\rho = 0.8$ for LCA and Moire. We provide two examples of the order prediction for the two distortions in Figure 5. We see that, for the LCA, the predicted order was almost perfect and that for Moire, the order was flawless. This example demonstrates that the proposed methodology also works well on real life images.

Table 3. Summary of order-preserving image quality assessment tests for LCA and Moire distortions, added artificially or newly acquired in real images.

Test	LCA	Moire
Synthetic data, test set (pairwise ordering %TP)	97%	94%
TE42 chart (pairwise ordering %TP)	80%	85%
TID2013 (image set ordering ρ)	1	N/A
TE42 chart (image set ordering ρ)	0.7	0.8




Figure 5. Example of quadruplets, ordered by rank (a-d); ROIs used for ranking (e); and a graph with the expectation (x-axis) versus the median predicted rank values (ty-axis) (f). In the first row, LCA distortion was measured with a result of 0.9 accuracy; in the second row, Moire was assessed with an accuracy of 1.

5. Summary and Future Directions

Assessing image quality is a key problem when evaluating image processing algorithms. The challenge can be addressed either by absolute or relative measurements. As the absolute metric is sometimes ill-defined, and can be harder to implement, a relative

¹ It should be noted that the ROIs were selected by calculating the error map as described above.

metric can solve the problem. In this paper we suggested a relative method, which looks for an optimal mapping that maintains the order of pair of images. Namely, given a pair of images, the mapping returns a pair of scalars that are ordered based on IQ. Subsequently, we proposed extending the mechanism for ranking a set of registered images. The ranking was performed by ordering the images by the found function values.

We demonstrated the validity of our method by constructing a Deep Neural Network and testing it on two distortions: Chromatic Aberration and Moire. The test's accuracy on synthetic data as well as real data showed satisfactory results. Our test demonstrates that even though the training was performed on synthetic data, the results achieved on real data was satisfactory. Our method paves the way towards learning to measure a wide range of image distortions.

Acknowledgments

We would like to thank Dr. Shay Maymon and Dmitry Paus for their insightful and valuable comments and suggestions. Shira is grateful to the Eric and Wendy Schmidt Fund for Strategic Innovation, and to the Zuckerman-CHE STEM Program for supporting her research.

References

- [1] TE42. <https://www.image-engineering.de/products/charts/all/425-te42>. Accessed: 2021-09-12
- [2] B. Smith, E.H., Maggard, E., Line, S., Shaw, M., "Quantifying print quality for practice," in NIP & Digital Fabrication Conference, vol. 2015, pp. 157-162. Society for Imaging Science and Technology, 2015.
- [3] C.J. van den Branden Lambrecht, "A working spatio-temporal model of the human visual system for image restoration and quality assessment applications," in IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings, vol. 4, pp. 2291-2294, 1996.
- [4] A. Chetouani, "Image quality assessment without reference by mixing deep learning-based features," in IEEE International Conference on Multimedia and Expo (ICME), pp. 1-6, 2020.
- [5] J. Deng, W. Dong, R. Socher, L.J. Li, K. Li, L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 248-255, 2009.
- [6] K. Ding, K. Ma, S. Wang, E.P. Simoncelli, "Image quality assessment: Unifying structure and texture similarity," arXiv preprint arXiv:2004.07728, 2020.
- [7] S. Faigenbaum, A. Shaus, B. Sober, E. Turkel, E. Piasetzky, "Evaluating glyph binarizations based on their properties," in Proceedings of the 2013 ACM symposium on Document engineering, pp. 127-130, 2013.
- [8] F. Gao, D. Tao, X. Gao, X. Li, "Learning to rank for blind image quality assessment," IEEE transactions on neural networks and learning systems, vol. 26, no. 10, pp. 2275-2290, 2015.
- [9] F. Gao, Y. Wang, P. Li, M. Tan, J. Yu, Y. Zhu, "Deepsim: Deep similarity for image quality assessment," Neurocomputing vol. 257, pp. 104-114, 2017.
- [10] D.J. Granrath, "The role of human visual models in image processing," Proceedings of the IEEE vol. 69, no. 5, pp. 552-561, 1981.
- [11] K. He, X. Zhang, S. Ren, J. Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770-778, 2016.
- [12] W. Hou, X. Gao, D. Tao, X. Li, "Blind image quality assessment via deep learning," IEEE transactions on neural networks and learning systems vol. 26, no. 6, pp. 1275-1286, 2015.
- [13] J. Kleinmann, J., Wueller, D., "Investigation of two methods to quantify noise in digital images based on the perception of the human eye," in Image Quality and System Performance IV, vol. 6494, International Society for Optics and Photonics, 2007.
- [14] E. Lehmann, H. Dabrera, Nonparametrics: Statistical methods based on ranks, Holden-day inc. San Francisco, 1975.
- [15] N. Ponomarenko, L. Jin, O. Ieremeiev, V. Lukin, K. Egiazarian, J. Astola, B. Vozel, K. Chehdi, M. Carli, F. Battisti, C.C. J. Kuo, "Image database TID2013: Peculiarities, results and perspectives," Signal Processing: Image Communication vol. 30, pp. 57-77, 2015.
- [16] J.D. Rennie, N. Srebro, "Loss functions for preference levels: Regression with discrete ordered labels," In Proceedings of the IJCAI multidisciplinary workshop on advances in preference handling, pp. 180-186. Kluwer Norwell, MA, 2005.
- [17] D.N. Sidorov, A.C. Kokaram, "Suppression of moire patterns via spectral analysis," in Visual Communications and Image Processing, vol. 4671, pp. 895-907. International Society for Optics and Photonics, 2002.
- [18] H. Talebi, P. Milanfar, "Nima: Neural image assessment," IEEE Trans. on Image Processing vol. 27, no. 8, 3998-4011, 2018.
- [19] I. Vendrov, R. Kiros, S. Fidler, R. Urtasun, "Order-embeddings of images and language," arXiv preprint arXiv:1511.06361, 2015.
- [20] Z. Wang, A.C. Bovik, L. Lu, "Why is image quality assessment so difficult?," in IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), vol. 4, pp. 3313-3316, 2002.
- [21] L. Xu, J. Li, W. Lin, Y. Zhang, Y. Zhang, Y. Yan, "Pairwise comparison and rank learning for image quality assessment," Displays vol 44, pp. 21-26, 2016.
- [22] Y. Zhang, D.M. Chandler, "Opinion-unaware blind quality assessment of multiply and singly distorted images via distortion parameter estimation," IEEE Trans. on Image Processing, vol. 27, no. 11, pp. 5433-5448, 2018.
- [23] H. Zhu, L. Li, J. Wu, W. Dong, G. Shi, "MetaQA: Deep meta-learning for no-reference image quality assessment," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020.

Author Biography

Shira Faigenbaum-Golovin received her BSc and MSc in Mathematics from Tel Aviv University, Israel (in 2006 and 2013, respectively). In 2021 she received her PhD in Applied Mathematics from the Tel-Aviv University. Since then, she is a Phillip Griffiths Assistant Research Professor at Duke University, in the mathematics department, as well as at the Rhodes Interdisciplinary Initiative. Her research strives at developing machine learning and image processing algorithms and designing approximation tools for manifolds in high-dimensions.

Or Shimshi received his BSc in electrical engineering from the Tel Aviv University (2004). Since then, he has worked in applied research positions at multiple companies, such as Samsung, Intel, Qualcomm, and Citi. His work has mainly focused on Machine Learning, Computer Vision, and NLP.