

Reduced reference image quality assessment via Boltzmann machines

Citation for published version (APA):

Mocanu, D. C., Exarchakos, G., Bou Ammar, H., & Liotta, A. (2015). Reduced reference image quality assessment via Boltzmann machines. In *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM), 11-15 May 2015, Ottawa, Canada* (pp. 1278-1281). Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/INM.2015.7140481>

DOI:

[10.1109/INM.2015.7140481](https://doi.org/10.1109/INM.2015.7140481)

Document status and date:

Published: 01/01/2015

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Reduced Reference Image Quality Assessment via Boltzmann Machines

Decebal Constantin Mocanu^{*†}, Georgios Exarchakos^{*}, Haitham Bou Ammar[†] and Antonio Liotta^{*}

^{*}Department of Electrical Engineering

Eindhoven University of Technology, The Netherlands

[†]Department of computer and information science

University of Pennsylvania, USA

Abstract—Monitoring and controlling the user’s perceived quality, in modern video services is a challenging proposition, mainly due to the limitations of current Image Quality Assessment (IQA) algorithms. Subjective Quality of Experience (QoE) is widely used to get a right impression, but unfortunately this can not be used in real world scenarios. In general, objective QoE algorithms represent a good substitution for the subjective ones, and they are split in three main directions: Full Reference (FR), Reduced Reference (RR), and No Reference (NR). From these three, the RR IQA approach offers a practical solution to assess the quality of an impaired image due to the fact that just a small amount of information is needed from the original image. At the same time, keeping in mind that we need automated QoE algorithms which are context independent, in this paper we introduce a novel stochastic RR IQA metric to assess the quality of an image based on Deep Learning, namely Restricted Boltzmann Machine Similarity Measure (RBMSim). RBMSim was evaluated on two benchmarked image databases with subjective studies, against objective IQA algorithms. The results show that its performance is comparable, or even better in some cases, with widely known FR IQA methods.

Keywords—Quality of Experience, Reduced Reference Image Quality Assessment, Restricted Boltzmann Machine, Deep Learning, Similarity Measure

I. INTRODUCTION

To improve the Quality of Experience (QoE) [1], [2], [3],[4] for the increasing amount of multimedia network services, here is a need for automated algorithms capable to assess the quality of those services such as human do [5]. An important slice of these services is represented by video services in their various forms and applications such as digital cinema, internet videos, digital television, video teleconferencing, mobile broadcasting, video streaming and so on. To evaluate the quality of a video is a difficult problem [6], [7], and in practice we have to decompose it at a lower level and to evaluate individual frames (i.e. images). Furthermore, the field of Image Quality Assessment (IQA) is quite broad in scope as it also contains the application domains of medical imaging, computer vision, image acquisition, pattern recognition, etc. [8].

Currently, the quality assessment methods which reflect accurately the nature of human perception are the subjective ones [9], such as Difference Mean Opinion Scores

Pre-printed version.
Cite as: D.C. Mocanu, G. Exarchakos, H. Bou Ammar, A. Liotta.

Reduced Reference Image Quality Assessment via Boltzmann Machines, IEEE/IFIP IM, 11-15 May 2015, Ottawa (Canada)

(DMOS) [10]. Nevertheless, due to the need of using automated algorithms, capable of fast reaction in different real world scenarios a good substitution for the subjective methods are the objective IQA algorithms [11]. These are split in three main categories (i.e. FR IQA, NR IQA and RR IQA), each one with its own advantages and disadvantages. In the literature the most widely used methods are the FR IQA. They usually offer the best approximation for the image quality assessment, but at the same time they have a big downside, they need the original image, which makes them unsuitable for network applications [12]. Furthermore, the amount of research done in the area of the NR IQA and RR IQA algorithms is less. In the case of NR IQA the main advantage is that there is no need of the original image or any additional information from it. The evaluation of the impaired image is done based just on general statistics. Unfortunately, this type of methods, need knowledge about the type of distortion used (e.g. Gaussian blur) [13], and they are immature when there is no information about the distortions. This information is not easy to obtain in the case of network services, where the impairments can be packet loss or lossy compression, for example. By aiming to solve the previous issues, in this paper, we are focusing on RR IQA methods, which are the most suitable for the evaluation of image quality in network services. Furthermore, by trying to develop a general RR IQA framework which does not need any information about the impairments type and also to be general enough to cover different types of images, not just natural image scenes, we focus on Unsupervised Learning (UL) [14] methods. Our intuition is that by using UL, an automated algorithm would be able to extract by itself the most important features of an image and to use those features on the impaired image to assess its quality. Following this intuition, and by extending our previous research, where we applied similar ideas to assess the quality of 3D images [15] or to measure the similarity between different Markov Decision Process [16], in this paper we propose a novel RR IQA method. We dubbed this method Restricted Boltzmann Machine Similarity Measure (RBMSim), and it is an automated stochastic procedure capable to assess the quality of a distorted image, independently from the type of impairments or the type of image itself. To test the proposed approach, we have compared it with objective IQA algorithms on two subjective benchmark image databases [17], [18] which contains different types of distortions. The next step would be to test network impaired videos.

The remainder of this paper is organized as follows. Section II, presents background knowledge on QoE assessment

for the benefit of the non-specialist reader. Section III, details the proposed method including all the mathematical details. Section IV, describes the experiments performed and reflects upon the attained results. Finally, Section V concludes and presents directions of future work.

II. BACKGROUND ON IMAGE QOE ASSESSMENT

The best methods to reflect the human perception are the subjective QoE methods, due to the fact that they involve humans to evaluate human-perception factors. However, it's not always possible to employ subjective tests, particularly when the application requires real-time feedback for the purpose of service monitoring and control. More commonly, subjective studies are successfully used to benchmark the results of the objective QoE methods. One of these is the Difference Mean Opinion Scores (DMOS) [9] which is a subjective QoE procedure to measure the image quality, known in the literature as hidden reference removal. The interested reader is referred to [10] for a more comprehensive discussion. Among all the objective QoE algorithms for image quality assessment we would like to enumerate here some of them: Peak signal-to-noise ratio (PSNR), Structural Similarity (SSIM), Multi-scale Structural Similarity (MS-SSIM), Visual Information Fidelity (VIF), Information Fidelity Criterion (IFC), Noise Quality Measure (NQM), Visual signal-to-noise ratio (VSNR), Weighted signal-to-noise ratio (WSNR), Blind/Referenceless Image Spatial QUality Evaluator (BRISQUE), and Most Apparent Distortion (MAD). The benchmark databases [17], [18], used by us in Section IV to assess the performance of RBMSim, provide the performance of all the aforementioned IQA algorithms.

III. RBMSIM

In this section, we introduce a novel stochastic Reduced Reference Image Quality Assessment method, so called Restricted Boltzmann Machines Similarity Measure (RBMSim). Firstly, we discuss the Restricted Boltzmann Machines [19] and their training method, and we end up by giving an intuition and detailing the RBMSim procedure.

A. Gaussian Bernoulli Restricted Boltzmann Machines

Restricted Boltzmann Machines (RBMs) [19] are energy-based models for unsupervised learning. These models are stochastic with stochastic nodes and layers, making them less vulnerable to local minima [20]. Furthermore, due to their architecture and neural configurations, RBMs and their variants possess excellent generalization capabilities [21], [22], [23].

Formally, a RBM consists of visible and hidden binary layers (i.e. \mathbf{v} and \mathbf{h} respectively), as shown in Fig. 1. The visible layer represents the data, while the hidden one increases the learning capacity by enlarging the class of distributions that can be represented to an arbitrary complexity. In their original form RBMs contain just binary units, which make them unusable for a large number of applications with real numbers. This problem was solved by Salakhutdinov & Hinton in 2006 [24], when based on Welling et al. work about Exponential Family Harmoniums [25], they introduced Gaussian Bernoulli Restricted Boltzmann (GRBM), where the binary units from the visible layer \mathbf{v} are replaced by linear units with Gaussian




Fig. 1. RBM general architecture, where $\mathbf{v}(i)$, $\mathbf{h}(j)$, and \mathbf{W} represent the visible neurons, hidden neurons and the undirected connections between the neurons from different layers.

noise. The hidden units \mathbf{h} remains binary. The total energy of GRBM is given by equation 1.

$$E(v, h) = - \sum_{i,j} \frac{v_i}{\sigma_i} h_j W_{ij} - \sum_i \frac{(v_i - a_i)^2}{2\sigma_i^2} - \sum_j h_j b_j \quad (1)$$

where W_{ij} denotes the connection between visible neuron i and hidden neuron j , a_i is the bias for visible neurons i and b_j is the bias for hidden neurons j . The term $\sum_{i,j} \frac{v_i}{\sigma_i} h_j W_{ij}$ represents the total energy between neurons from different layers, $\sum_i \frac{(v_i - a_i)^2}{2\sigma_i^2}$ represents the energy of the visible layer and $\sum_j h_j b_j$ the energy of the hidden layer. The inference for the hidden neuron j is done by sampling from a sigmoid function $sigm(b_j + \sum_i \frac{v_i}{\sigma_i} W_{ij})$. The inference for the visible unit i is done by sampling from a Gaussian distribution, defined as $\mathcal{N}(a_i + \sum_j h_j W_{ij}, \sigma_i^2)$. In order to maximize the

likelihood of the model, the gradients of the energy function with respect to the weights have to be calculated. Unfortunately, in all types of RBMs the maximum likelihood can not be straightforwardly applied due to intractability problems. To prevent these problems, Contrastive Divergence presented next, was introduced by Geoffrey Hinton in [26].

B. Contrastive Divergence (CD)

In Contrastive Divergence, learning follows the gradient of:

$$CD_n \propto D_{KL}(p_0(\mathbf{x}) || p_\infty(\mathbf{x})) - D_{KL}(p_n(\mathbf{x}) || p_\infty(\mathbf{x})) \quad (2)$$

where, $p_n(\cdot)$ is the resulting distribution of a Markov chain running for n steps. To find the update rules for the parameters of RBMs we have to calculate the derivatives of the energy function from Equation 1 with respect to those parameters (i.e. \mathbf{W} , \mathbf{a} and \mathbf{b}). Since the visible units are conditionally independent given the hidden units and vice versa, learning can be performed using one step Gibbs sampling, which is carried in two half-steps: (1) update all the hidden units, and (2) update all the visible units. Thus, in CD_n the weight updates are done as follows: $W_{ij}^{\tau+1} = W_{ij}^\tau + \alpha (\langle \langle h_j v_i \rangle_p(\mathbf{h} | \mathbf{v}; \mathbf{W}) \rangle_0 - \langle h_j v_i \rangle_n)$ where τ is the training epoch, α is the learning rate, and the subscript (n) indicates that the states are obtained after n iterations of Gibbs sampling from the Markov chain starting at $p_0(\cdot)$. For a more comprehensive discussion about RBM and CD, the interested reader is referred to [21].

C. RBMSim procedure

To sum up the previous two subsections, by initializing a RBM, after it is learned, with some testing data in the visible layer which follows a distribution close to the training data, that RBM will be capable to reconstruct those data with a very small error after a Gibbs sampling is done. Our intuition, go beyond this well known information and we made the assumption that the reconstruction error of the testing data is directly dependent upon the distance between the training and testing data.

More precisely, in the specific case of IQA we use a GRBM to discover the hidden features and model a reference image \mathcal{RI} , noted by $GRBM_{\mathcal{RI}}$. As a practical observation, all the pixels from \mathcal{RI} can be put directly on the visible layer \mathbf{v} of the GRBM with their RGB values, each RGB channel of each pixel being encoded by one neuron. To avoid a big number of weights the image can be split into a set of smaller images \mathcal{SI}_p with the next constraints: $\mathcal{SI}_p \in \mathcal{RI} ; \cup_p \mathcal{SI}_p = \mathcal{RI} ; \cap_p \mathcal{SI}_p = \emptyset$. The mean and the standard deviation for each RGB channel of each \mathcal{SI}_p sub-image can be used on the visible layer of the GRBM, each value being encoded by one visible neuron. Furthermore, to assess the quality of a distorted image \mathcal{DI} , or in other words how far away is \mathcal{DI} from the \mathcal{RI} , we can initialize the visible layer of the previously learned $GRBM_{\mathcal{RI}}$, with the sub-images from \mathcal{DI} in the same manner as before. Let us note the visible layer of the $GRBM_{\mathcal{RI}}$ model initialized with the \mathcal{DI} sub-images, $\mathbf{v}^{\mathcal{DI}}$. Afterwards, we perform the Gibbs sampling, i.e. firstly we infer the values of the hidden neurons, and secondly we infer (reconstruct) the values of the visible neurons based on the hidden ones. The last step involves measuring the distance between the values of the visible neurons after the reconstruction (noted with $\widehat{\mathbf{v}^{\mathcal{DI}}}$), and $\mathbf{v}^{\mathcal{DI}}$. This can be done using any distance measure such as Root Mean Square Error (RMSE). The RMSE value will reflect accurately the quality of \mathcal{DI} . Formally, the RBMSim metric is defined in Equation 3:

$$RBMSim(\mathcal{DI}) = \sqrt{\frac{1}{n_v} \sum_{i=1}^{n_v} (\widehat{\mathbf{v}_i^{\mathcal{DI}}} - \mathbf{v}_i^{\mathcal{DI}})^2} \quad (3)$$

where n_v is the number of visible neurons. If the RGB color model is used, $RBMSim(\mathcal{DI}) \in [0, 255]$, where 0 means that \mathcal{RI} and \mathcal{DI} are the same, and 255 means that \mathcal{RI} and \mathcal{DI} are completely different. Due to the fact that $RBMSim$ is a stochastic procedure the extreme values can never be reached. The amount of information introduced by the parameters of $GRBM_{\mathcal{RI}}$ needed to assess the quality of \mathcal{DI} , is still much smaller than the \mathcal{RI} itself making the approach suitable for a wide variety of applications. To clarify, on the client side, just the $GRBM_{\mathcal{RI}}$ model has to be transferred.

IV. EXPERIMENTS AND RESULTS

In this section, we compare the performance of RBMSim with objective IQA metrics on two subjective benchmark image databases, namely LIVE Multiply Distorted Image Quality Database (LIVE) [17] and CSIQ database [18]. Both databases have associated, subjective studies, using DMOS. For both databases, to bring the predictions of RBMSim on the same scale with DMOS values, we use the logistic transform recommended by the the Video Quality Experts Group [18].

The LIVE database contains fifteen reference images, which were corrupted using two studies with multiple distortions. In the first study, the images are first blurred and then compressed with a JPEG encoder. In the second study, the images are first blurred using different defocus techniques and then corrupted by white Gaussian noise. For each scenarios, for each reference image there are 15 types of impairments. The subjective DMOS study was done on 19 participants for the first scenario and 18 participants for the second scenario. Each image has a resolution of 1280x720 pixels. To keep the number of parameters for RBMSim low, we split each image in small images of 64x36 pixels, so in total we had 400 sub-images. For each sub-image, we took into consideration the mean and standard deviation for each RGB channel. So, in total we had 2400 (i.e. 400x3x2) visible neurons and 10 hidden neurons. The learning rate was set to 0.001. The computational time for one training epoch was 0.03 seconds. In total we trained the RBMSim model for 200 epochs, but as it can be seen in Fig. 2 the model converged after approximately 50 epochs. For both LIVE studies, RBMSim was capable to approximate very well the DMOS values, as it is depicted in Fig. 3. Besides that, it outperforms all the FR and NR IQA metrics from [17], sharing the first place with the best performer BRISQUE-2, as it can be seen in Table I.

TABLE I. RBMSIM SCORES IN CORRELATION WITH DMOS ON LIVE DATABASE. THE IQA METRICS USED IN [17] ARE PSNR, MS-SSIM, VIF, IFC, NQM, VSNR, WSNR, BRISQUE-1, AND BRISQUE-2.

	Study 1		Study 2		Overall	
	SROCC	LCC	SROCC	LCC	SROCC	LCC
Lower score in [17]	0.66	0.75	0.21	0.38	0.42	0.50
Upper score in [17]	0.92	0.95	0.90	0.92	0.91	0.93
RBMSim	0.95	0.97	0.84	0.91	0.88	0.93

The CSIQ database [18] has 30 reference images, each of them having 6 types of impairments, with four or five different distortion levels. It has 5000 subjective ratings, made by 25 different observers, with DMOS. Each image has a resolution of 512x512 pixels. To test RBMSim, we split each image in squares of 32x32 pixels, obtaining in total 256 sub-images for each image. So, RBMSim has 1536 (i.e. 256x3x2) visible neurons and 10 hidden neurons. Fig. 4 shows the correlation of RBMSim scores with DMOS, and Table II compares the performance of RBMSim with the FR IQA metrics from [18] in terms of Pearson correlation coefficient (LCC), Spearman correlation coefficient (SROCC) and Outliers Ratio (OR). It is worth highlighting that RBMSim, even if it is outperformed by the FR IQA algorithms VIF and MAD, it outperforms widely used FR IQA metrics, e.g. PSNR, SSIM, on this database.

TABLE II. RBMSIM SCORES IN CORRELATION WITH DMOS ON CSIQ DATABASE. THE FR IQA METRICS USED IN [18] ARE PSNR, SSIM, MS-SSIM, NQM, VSN, VIF, AND MAD.

	SROCC	LCC	OR
Lower score in [18]	0.74	0.74	37%
Upper score in [18]	0.95	0.95	18%
RBMSim	0.85	0.85	26%

V. CONCLUSION

This paper introduced a novel stochastic RR IQA metric, namely RBMSim. Notably, it proposed the use of GRBM to

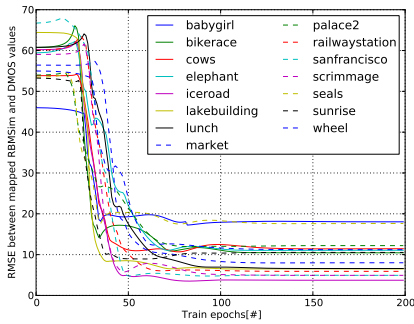


Fig. 2. The study of convergence for the learning phase of RBMSim, averaged over both LIVE studies and all distortions, for each reference image.




Fig. 3. DMOS against mapped RBMSim scores on LIVE. Each point represents a distorted image.




Fig. 4. DMOS against mapped RBMSim scores on CSIQ. Each point represents a distorted image.

measure the quality of distorted images. Its performance was evaluated on two subjective benchmarked image databases. The performance is at least comparable with widely known FR IQA metrics. Additionally, RBMSim shows a fast computational time, and is therefore suitable for on-line applications. As further directions, would be interesting to adapt RBMSim to videos by training RBMSim on all frames of a video, to be able to monitor in real-time the performance of network services from the user point of view.

ACKNOWLEDGMENT

This work has been partially supported by ARTEMIS project DEMANES (Design, Monitoring and Operation of Adaptive Networked Embedded Systems, contract 295372).

REFERENCES

- [1] S. Winkler and P. Mohandas, "The evolution of video quality measurement: From PSNR to hybrid metrics," *IEEE Transactions on Broadcasting*, vol. 54, no. 3, pp. 660–668, Sep. 2008.
- [2] V. Menkovski, G. Exarchakos, and A. Liotta, "The value of relative quality in video delivery," *J. Mob. Multimed.*, vol. 7, no. 3, pp. 151–162, Sep. 2011.
- [3] V. Menkovski, G. Exarchakos, A. Liotta, and A. C. Sanchez, "Quality of experience models for multimedia streaming," *IJMCMC*, vol. 2, no. 4, pp. 1–20, 2010.
- [4] V. Menkovski and A. Liotta, "Adaptive psychometric scaling for video quality assessment," *Image Commun.*, vol. 27, no. 8, pp. 788–799, Sep. 2012.
- [5] A. Liotta, "The cognitive NET is coming," *IEEE Spectrum*, vol. 50, no. 8, pp. 26–31, Aug. 2013.
- [6] D. Mocanu, A. Liotta, A. Ricci, M. Vega, and G. Exarchakos, "When does lower bitrate give higher quality in modern video services?" in *Network Operations and Management Symposium (NOMS), 2014 IEEE*, May 2014, pp. 1–5.
- [7] D. Mocanu, G. Santandrea, W. Cerroni, F. Callegati, and A. Liotta, "Network performance assessment with quality of experience benchmarks," in *Network and Service Management (CNSM), 2014 10th International Conference on*, Nov 2014, pp. 332–335.
- [8] Z. Wang, "Applications of objective image quality assessment methods [applications corner]," *Signal Processing Magazine, IEEE*, vol. 28, no. 6, pp. 137–142, 2011.
- [9] K. Seshadrinathan, R. Soundararajan, A. C. Bovik, and L. K. Cormack, "Study of subjective and objective quality assessment of video," *Trans. Img. Proc.*, vol. 19, no. 6, pp. 1427–1441, Jun. 2010.
- [10] M. Pinson and S. Wolf, "Comparing subjective video quality testing methodologies," in *SPIE Video Communications and Image Processing Conference*, 2003, pp. 8–11.

- [11] A. K. Moorthy and A. C. Bovik, "Visual quality assessment algorithms: What does the future hold?" *Multimedia Tools Appl.*, vol. 51, no. 2, pp. 675–696, Jan. 2011.
- [12] D. M. Chandler, "Seven challenges in image quality assessment: Past, present, and future research," *ISRN Signal Processing*, 2013.
- [13] A. Rehman and Z. Wang, "Reduced-reference image quality assessment by structural similarity estimation," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3378–3389, 2012.
- [14] G. Hinton and T. Sejnowski, *Unsupervised Learning: Foundations of Neural Computation*, ser. A Bradford Book. MIT Press, 1999.
- [15] D. C. Mocanu, G. Exarchakos, and A. Liotta, "Deep learning for objective quality assessment of 3d images," in *Image Processing (ICIP), 2014 IEEE International Conference on*, Oct 2014, pp. 758–762.
- [16] H. B. Ammar, E. Eaton, M. E. Taylor, D. C. Mocanu, K. Driessens, G. Weiss, and K. Tuyls, "An automated measure of mdp similarity for transfer in reinforcement learning," in *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [17] D. Jayaraman, A. Mittal, A. Moorthy, and A. Bovik, "Objective quality assessment of multiply distorted images," in *Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on*, 2012, pp. 1693–1697.
- [18] E. C. Larson and D. M. Chandler, "Most apparent distortion: full-reference image quality assessment and the role of strategy," *J. Electronic Imaging*, vol. 19, no. 1, p. 011006, 2010.
- [19] P. Smolensky, "Information processing in dynamical systems: Foundations of harmony theory," in *Parallel Distributed Processing: Volume 1: Foundations*, D. E. Rumelhart, J. L. McClelland et al., Eds. Cambridge: MIT Press, 1987, pp. 194–281.
- [20] G. W. Taylor, G. E. Hinton, and S. T. Roweis, "Two distributed-state models for generating high-dimensional time series," *Journal of Machine Learning Research*, vol. 12, pp. 1025–1068, 2011.
- [21] Y. Bengio, "Learning deep architectures for ai," *Found. Trends Mach. Learn.*, vol. 2, no. 1, pp. 1–127, Jan. 2009.
- [22] D. C. Mocanu, H. B. Ammar, D. Lowet, K. Driessens, A. Liotta, G. Weiss, and K. Tuyls, "Factored four way conditional restricted boltzmann machines for activity recognition," *Pattern Recognition Letters*, 2015.
- [23] E. Mocanu, D. Mocanu, H. Ammar, Z. Zivkovic, A. Liotta, and E. Smirnov, "Inexpensive user tracking using boltzmann machines," in *Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on*, Oct 2014, pp. 1–6.
- [24] G. E. Hinton and R. R. Salakhutdinov, "Reducing the Dimensionality of Data with Neural Networks," *Science*, vol. 313, no. 5786, pp. 504–507, Jul. 2006.
- [25] M. Welling, M. Rosen-Zvi, and G. E. Hinton, "Exponential family harmoniums with an application to information retrieval," in *NIPS*, 2004.
- [26] G. E. Hinton, "Training Products of Experts by Minimizing Contrastive Divergence," *Neural Computation*, vol. 14, no. 8, pp. 1771–1800, Aug. 2002.