

The Impact of Image-Difference Features on Perceived Image Differences

Jens Preiss (1), Ingmar Lissner (1), Philipp Urban (1), Matthias Scheller Lichtenauer (2)(3), Peter Zolliker (2);
(1) Institute of Printing Science and Technology, Technische Universität Darmstadt, Germany;
(2) Laboratory for Media Technology, Empa, Duebendorf, Switzerland;
(3) Friedrich-Schiller-Universität, Jena, Germany

Abstract

We discuss a few selected hypotheses on how the visual system judges differences of color images. We then derive five image-difference features from these hypotheses and address their relation to the visual processing. Three models are proposed to combine these features for the prediction of perceived image differences. The parameters of the image-difference features are optimized on human image-difference assessments.

For each model, we investigate the impact of individual features on the overall prediction performance. If chromatic features are combined with lightness-based features, the prediction accuracy on a test dataset is significantly higher than that of the SSIM index, which only operates on the achromatic component.

Introduction

A measure that accurately predicts perceived image differences would be useful for evaluating and optimizing image-processing methods, particularly in the areas of image compression, transmission, and reproduction. Various image-difference measures (IDMs) have been proposed in the literature. They are based upon hypotheses on distortions to which the human visual system (HVS) is particularly sensitive. Such hypotheses include pixelwise color differences, pixelwise differences of color attributes (i.e., lightness, chroma, and hue), differences in spatial changes of color attributes (gradients or variance), color differences within large areas of the same color [1], global and local contrast changes, changes of structural information [2], and differences in special regions of interest (edges, corners) [3].

Although the cortical mechanisms of image-difference perception are poorly understood, existing IDMs quite successfully predict perceived image differences for many individual distortions such as compression artifacts, noise, and blur. Most methods also incorporate hypotheses on how the HVS judges combinations of distortions that occur, for instance, in gamut mapping or tone mapping. Nevertheless, the prediction performance of IDMs for images with such complex distortions can be improved.

In this paper we analyze the influence of several hypotheses on the perceived overall image difference. Assumptions about how the HVS combines individual differences into an overall image difference are also investigated.

We utilize the recently presented image-difference framework [4] for our evaluations. In addition, we describe in detail how our IDMs are derived from the hypotheses.

The Image-Difference-Feature Framework

The visual mechanisms that contribute to the assessment of image differences are very complex. We attempt to model these mechanisms empirically as a combination of simple hypotheses on how the visual system performs such tasks. These hypothe-

ses are mathematically expressed as *image-difference features* (IDFs). Given two input images, our framework computes several IDFs, which are then combined to optimize the prediction of human image-difference assessments [4].

As a first step of our IDF computation we use an image-appearance model to normalize the images to specific viewing conditions (illuminant, luminance level, viewing distance). They are then transformed into a working color space with certain beneficial properties — for instance, Euclidean distances in this space should closely match perceived color differences. This ensures that image features, such as edges and gradients, are judged correctly by the subsequent feature-extraction routines.

In summary, image-difference features are computed as follows:

1. Image-appearance-model transformation

The input images are transformed by an image-appearance model to normalize them to specific viewing conditions. Here, we use S-CIELAB, which filters the images “to simulate the spatial blurring by the human visual system” [5].

2. Color-space transformation

The images are transformed into a working color space. In this paper, we use the LAB2000HL color space [6], which is both hue linear and highly perceptually uniform with respect to the CIEDE2000 color-difference formula [7].

3. Difference-map generation

Maps are generated that reflect differences between the images, e.g., gradient differences or color differences.

4. Characteristic-value computation

Each map is finally transformed into a characteristic value, e.g., the mean value of all pixels.

The resulting IDFs are combined into an image-difference measure (IDM) using, e.g., an additive combination model. The coefficients of the model can be optimized on an image database with image-difference assessments of human observers.

Distortions Affecting Image Differences

Hypotheses of Distortions

In the following, we present a list of low-level hypotheses investigated in this paper. We also explore their importance for the overall image-difference impression. More hypotheses, especially on high-level (semantic) image differences, can be found in the literature.

Hypothesis 1: *The HVS is sensitive to lightness, chroma, and hue differences.*

This, of course, includes color differences that are combinations of color-attribute differences. Gamut-mapping algorithms, for instance, incorporate this assumption [8].

Hypothesis 2: *The HVS is sensitive to achromatic contrast differences.*

Hypothesis 3: *The HVS is sensitive to achromatic structural differences.*

This is particularly important if the distorted image contains artifacts such as banding.

Hypothesis 4: *The HVS is more sensitive to contrast changes in low contrast image regions.*

This hypothesis is related to Weber’s law, stating that threshold difference is proportional to the magnitude of the stimulus (in our case contrast). This phenomenon can also be observed for suprathreshold contrast differences: if, for instance, the contrast of a distinct edge is noticeably reduced, the edge information is still present; the same contrast change applied to a weaker edge may cause it to vanish. The latter contrast change is likely to be considered more disturbing.

Hypothesis 5: *The HVS is not equally sensitive to differences described in hypothesis 1–4 on different spatial-frequency bands.*

For reasons of simplicity we do not consider this hypothesis here. It can be incorporated into the analysis by a multi-scale approach.

Derived Image-Difference Features

A widely used image-difference measure is the structural similarity (SSIM) index proposed by Wang et al. [2]. It incorporates hypotheses 1 (at least for luminance differences), 2, and 3 by applying three functions in a sliding window centered around corresponding pixels x and y of two images:

$$\text{luminance: } l(x, y, c_1) = \frac{2\mu_x\mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1}, \quad (1)$$

$$\text{contrast: } c(x, y, c_2) = \frac{2\sigma_x\sigma_y + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}, \quad (2)$$

$$\text{structure: } s(x, y, c_3) = \frac{\sigma_{xy} + c_3}{\sigma_x\sigma_y + c_3}, \quad (3)$$

where μ_x and μ_y denote the means, σ_x and σ_y the standard deviations, and σ_{xy} the mean-adjusted inner product of the pixel values within the window. The variables $c_i > 0$ serve to avoid instabilities for small denominators. The variable c_2 of $c(x, y, c_2)$ can be adjusted to account for the contrast-masking property of the HVS described in hypothesis 4.

These functions are combined using a factorial model, resulting in

$$\text{SSIM}(x, y) = l(x, y, c_1)^{\alpha_1} c(x, y, c_2)^{\alpha_2} s(x, y, c_3)^{\alpha_3}, \quad (4)$$

where c_1 , c_2 , and c_3 are adjusted to the working color space, and α_1 , α_2 , and α_3 may be used to weight the contribution of each function to the overall image-difference prediction.

The SSIM index is a simple IDM and incorporates many of the hypotheses described above. In addition, a multi-scale extension was proposed, which incorporates hypothesis 5 [9]. Due to its advantageous properties and its popularity we create image-difference features based on the terms of the SSIM index.

Comments on Image-Difference Features

Before we start our investigations, we address some issues of using the SSIM functions in our framework:

1. The SSIM index is designed for grayscale images and ignores color information. To fully incorporate hypothesis 1, ad-




Figure 1. Contour plot of the SSIM l -function for CIELAB lightness with $c_1 = 1$. For constant lightness differences, the function increases with increasing absolute lightness, e.g., $\Delta L^*(\mu_x = 10, \mu_y = 5) = \Delta L^*(\mu_x = 80, \mu_y = 75) = 5$, but $l(x', y') = 0.8$ and $l(x'', y'') = 0.99$.




Figure 2. The influence of c_2 on c -function-based difference maps.

ditional functions are required that include hue and chroma differences. Some attempts were already made [4, 10].

2. The l -function shown in Eq. (1) is designed for luminance. It is not suitable for judging lightness errors in a perceptually uniform color space, where Euclidean distances match perceived color differences. Lightness differences in darker image areas are considered more disturbing (smaller function values) than the same differences in lighter areas. This is illustrated in Fig. 1.

3. The contrast function c reflects the contrast-masking property of the HVS as described in hypothesis 4. The degree of masking must be adjusted to the working color space using parameter c_2 . Its impact on the difference map is illustrated in Fig. 2. Large values of c_2 emphasize differences of edges, where even large contrast differences are hardly noticeable. Small values of c_2 emphasize contrast differences in low-contrast areas, where they are most disturbing.

4. The factorial model used to combine the l -, c -, and s -functions assumes that the image-difference features are com-

bined multiplicatively by the HVS. Although this is mathematically convenient because it keeps the index in a predefined range, it does not necessarily reflect the actual visual processing. Another question is whether a pixelwise combination of difference maps and subsequent averaging — as suggested by the SSIM index — yields the best predictions. The reverse processing order, averaging the maps and subsequent combination, is also conceivable. Please note that all approaches are purely empirical; it is not yet understood how the HVS combines individual distortions into an overall image-difference impression.

Methodology

Considered Image-Difference Features

We incorporate each of the SSIM functions into a separate IDF. However, the issues stated in the previous section necessitate several modifications as described in the following.

The luminance function l must be replaced by a function that is invariant to the lightness level. We also need to consider chroma and hue differences, ΔC and ΔH , that are in a similar perceived range as lightness differences ΔL in a perceptually uniform color space (at least for small differences). We used a function proposed in [10] to incorporate these color-attribute differences:

$$l_A(x, y, c_A) = \frac{1}{c_A \cdot A(x, y)^2 + 1}, \quad (5)$$

where $A = \Delta L$, ΔC , or ΔH . We replaced the SSIM l -function by the $l_{\Delta L}$ -function and included two functions, $l_{\Delta C}$ and $l_{\Delta H}$, into separate IDFs. ΔC and ΔH are defined analogously to ΔC_{ab}^* and ΔH_{ab}^* (see [11]) using the coordinates of the LAB2000HL color space [6] instead of CIELAB.

The function in Eq. (5) operates on a perceptually uniform color space. Using the Weber-Fechner law and other reasonable assumptions, it can be derived from the function in Eq. (1), which is suitable for an intensity-linear space.

In conclusion, we investigate five image-difference features that correspond to hypotheses 1–4. Figure 3 shows how these features are extracted from two images.

Combining Image-Difference Features

In the following, X and Y are images compared by the IDM and c_i is the parameter that corresponds to IDF_i . In contrast to the SSIM index, the IDMs proposed here return zero for visual equivalence of the input images. Higher values correspond to larger predicted image differences. Note that visual equivalence does not necessarily mean that the images are equal; spatial sub-threshold distortions may have been removed by the initial normalization with an image-appearance model.

We investigate three different ways to combine the considered IDFs into an IDM:

1. Additive model:

$$IDM_{\text{Additive}}(X, Y) = \sum_{i=1}^5 (1 - IDF_i(X, Y, c_i)) \quad (6)$$

2. Factorial model:

$$IDM_{\text{Factorial}}(X, Y) = 1 - \prod_{i=1}^5 IDF_i(X, Y, c_i) \quad (7)$$

3. Hybrid model:

$$IDM_{\text{Hybrid}}(X, Y) = \left(1 - \prod_{i=1}^3 IDF_i(X, Y, c_i) \right) + (1 - IDF_4(X, Y, c_4)) + (1 - IDF_5(X, Y, c_5)) \quad (8)$$

The hybrid model assumes that lightness distortions are combined multiplicatively, whereas achromatic and chromatic distortions are combined additively.

Since the characteristic values determined from the difference maps are combined, these IDF combinations differ from the combination used by the SSIM index. The SSIM index performs a pixelwise multiplication of the maps and calculates the characteristic value from the resulting map. Furthermore, we do not use any additional parameters to weight the individual IDFs. The SSIM index allows such adjustments using the parameters α_1 , α_2 , and α_3 as shown in Eq. (4).

Results and Discussion

For our evaluations, we used a combined dataset collected in ten gamut-mapping paired-comparison experiments [10, 12, 13, 14, 15]. This dataset contains 498 reference images (some of which show the same scene) and 3,891 distorted images corresponding to different gamut-mapping algorithms and gamut sizes. A total of 50,662 decisions were collected excluding ties.

Fitting the Model Parameters

Hit Rate

In each trial of a paired-comparison experiment two distorted images, i and j , and the corresponding reference image o are presented to an observer. The observer is asked to select the image $s(i, j, o) \in \{i, j\}$ that better represents the reference o . This selection can be compared with the prediction $p(i, j, o) \in \{i, j, o\}$ of an IDM, which is $p(i, j, o) = i$ if $IDM(i, o) < IDM(j, o)$, $p(i, j, o) = j$ if $IDM(j, o) < IDM(i, o)$, and $p(i, j, o) = o$ if $IDM(j, o) = IDM(i, o)$, as follows:

$$\delta(s(i, j, o), p(i, j, o)) = \begin{cases} 1, & s(i, j, o) = p(i, j, o) \\ 0, & \text{else} \end{cases} \quad (9)$$

The hit rate (HR) is defined as

$$HR = \frac{\sum_{(i, j, o, k) \in T} \delta(s_k(i, j, o), p(i, j, o))}{|T|}, \quad (10)$$

where T is the set of trial parameters (each element of T contains the image triplet (i, j, o) shown in trial k of the experiment), $s_k(i, j, o) \in \{i, j\}$ is the observer's decision in trial k , and $|T|$ is the number of elements in T . If the experiment allows ties, the corresponding elements are excluded from T . The higher the hit rate, the better the prediction accuracy of the investigated IDM is on the underlying visual data.

Hit rates allow a meaningful comparison of IDM predictions on a specific dataset and can therefore be used in objective functions for model fitting. However, hit rates do not provide sufficient information to judge the absolute prediction performance of IDMs. This is because the hit rate of a particular IDM strongly depends on the inter- and intraobserver variabilities of the underlying data, which makes it extremely difficult to compare hit rates across different datasets. We therefore employ the majority hit rate (MHR), which is the maximum achievable hit rate for a dataset. The MHR for our test dataset is 0.7741. It is only reached if the predictions of an IDM agree with the decisions of the majority of observers for any shown image triplet.

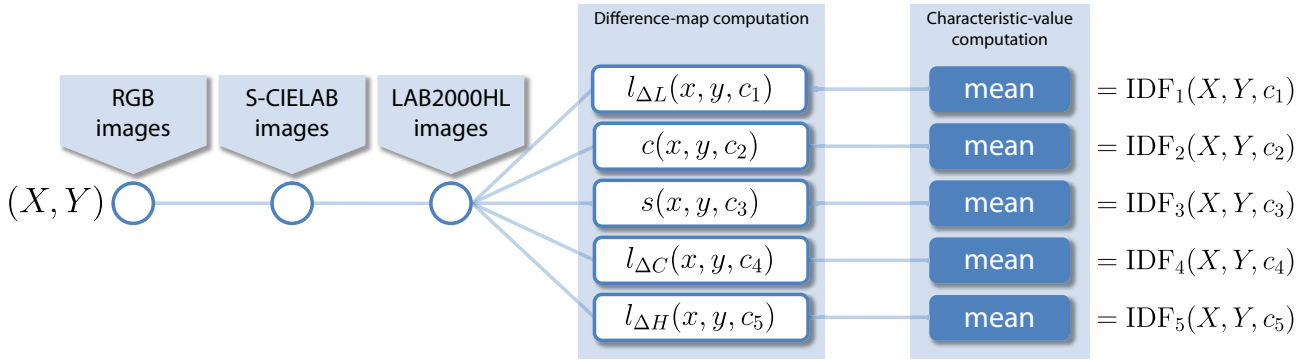


Figure 3. Investigated image-difference features (IDFs). S-CIELAB images result from spatial filtering as proposed by Zhang and Wandell [5].




Figure 4. Hit rate statistics for the three models. The blue and white bars represent $\pm 2\sigma$ intervals. The black lines indicate median hit rates.

The HR/MHR ratio is not affected by inter- and intraobserver variabilities and may be used for comparisons of IDMs across datasets.

Cross Validation

We used a cross-validation approach to optimize the parameters of our three models. The visual data were split into two sets of approximately equal size. One set was used for cross validation, the other was used for testing. All distorted images derived from the same reference image were assigned to the same set. We randomly extracted 20% of all reference images from the training set and adjusted the parameters of the investigated IDM to maximize its hit rate HR_{CV}^{OPT} on the corresponding trials. The optimized model was used to calculate the hit rate HR_{CV}^{TEST} on the remaining trials of the cross-validation dataset. This was repeated 81 times. We used the median parameters shown in Table 1 to evaluate the IDM on the test dataset.

	c_1	c_2	c_3	c_4	c_5
Additive	0.0012	0.1500	0.0875	0.0010	0.0375
Factorial	0.0014	0.1500	0.1000	0.0010	0.0375
Hybrid	0.0012	0.1500	0.0750	0.0010	0.0250

Table 1. Optimized parameters for the three models.

Figure 4 shows hit-rate statistics on the cross-validation dataset. Interestingly, all three models show similar performance. As the training hit rates HR_{CV}^{OPT} are not significantly better than the test hit rates HR_{CV}^{TEST} , the models do not overfit the data (see next subsection for significance tests).

Testing Significance

We are interested if differences of hit rates on the test dataset are significant, meaning that they are unlikely to be the result

of chance. To test this, we assume that each sample of IDM-predicted observer decisions follows a binomial distribution with unknown success probability. Please note that each such sample defines a hit rate. To judge whether two samples belong to different success probabilities, p_1 and p_2 , we calculate Yule's two-sample binomial confidence interval [16] for $|p_1 - p_2|$ using $\alpha = 5\%$. If the interval boundaries are positive we conclude that the two hit rates are significantly different. This approach does not consider any influence of inter- or intraobserver variabilities on the hit rate.

Testing the Models

We assessed the prediction performance of the IDMs using their hit rates on the test dataset. The corresponding hit rate of the SSIM index served as a reference. The SSIM index was calculated on RGB images using the default parameters given in [2]. Table 2 contains the resulting hit rates.

	SSIM		IDF-based		
	default	optim.	additive	factorial	hybrid
Hit rate	0.5992	0.6024	0.6319	0.6319	0.6340

Table 2. Hit rates of IDMs on the test dataset.

To analyze the impact of individual IDFs we removed IDF terms from the IDMs defined in Eqs. (6)–(8) and calculated their hit rates. Please note that the parameters c_i of the IDFs were optimized for the IDMs including all five IDFs. Figure 5 shows hit rates for all combinations of IDFs. Our findings can be summarized as follows:

1. The three investigated combination models do not have considerable influence on the hit rate.
2. The lightness-contrast-based IDF_2 is the most important IDF — in terms of achieved hit rate — of all IDMs with just one IDF. More generally, the most accurate among all IDMs with a certain number of IDFs *always* includes IDF_2 .
3. IDMs composed of two or more IDFs perform almost always better than those with only one IDF.
4. IDMs containing only the lightness-based IDFs (IDF_1 , IDF_2 , and IDF_3) perform significantly better than the default SSIM index. One possible reason is an intrinsic relationship between training and test data. To avoid such biased comparisons, we optimized the c_i parameters of the SSIM model on the training data using the cross-validation approach described above. The hit rate of the optimized SSIM index on the test data is not significantly different from that of the default SSIM index. It is still significantly lower than those of the IDMs with only lightness-based IDFs. Replacing the SSIM l -function with the function

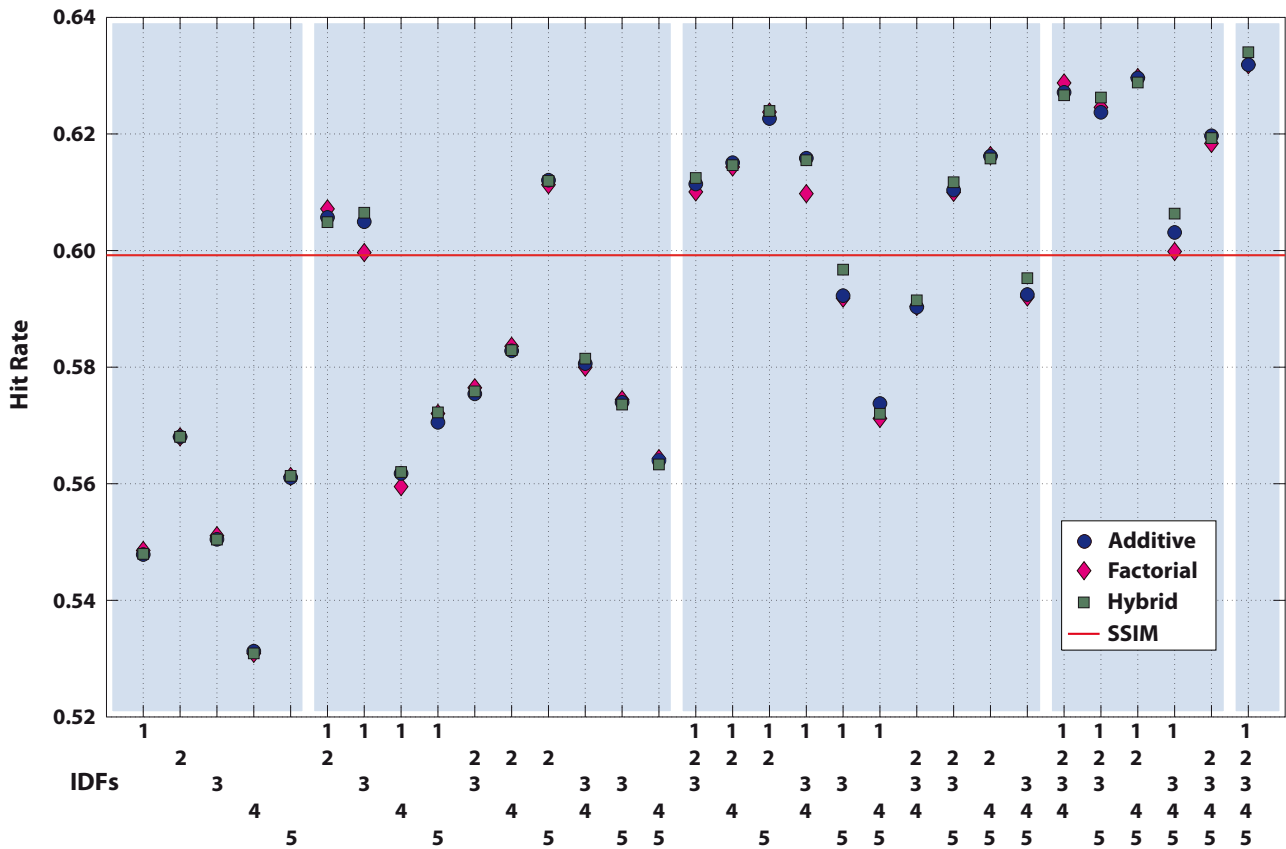


Figure 5. Hit rates on the test dataset of all possible IDF combinations incorporated into three different IDM models.

- shown in Eq. (5) seems to have a positive impact on the prediction accuracy.
5. Adding another IDF to an IDM does not negatively influence the prediction performance (except in one case).
 6. IDMs containing only the color-attribute differences (IDF₁, IDF₄, and IDF₅) perform significantly worse than all other IDMs composed of three IDFs.
 7. Adding the chroma-difference-based IDF₄ or the hue-difference-based IDF₅ to the lightness-based IDM composed of IDF₁, IDF₂, and IDF₃ increases the prediction performance significantly.
 8. The best investigated IDM achieves a hit rate of only 0.6340, i.e., there is still room for improvement to reach the majority hit rate MHR of 0.7741.

Please note that the visual data we used were collected to compare gamut-mapping algorithms (GMAs). GMAs commonly incorporate hypotheses on distortions to which the HVS is especially sensitive (e.g., hue shifts) and try to preserve the corresponding image attributes. Consequently, such distortions are less important for observer decisions in paired-comparison experiments. Furthermore, the studies compared only distorted images within the same gamut. Chroma distortions of such images are typically in the same range and therefore less crucial for the observers' decisions. This is why the relative importance of chroma differences obtained from an optimization is highly biased by the visual data. Figure 6 provides an example: although the chroma differences between the images are highly objectionable (see Fig. 2), the chroma-based difference map of IDF₄ shows very small intensities. This is due to a small parameter c_4 resulting from the optimization.

This simple example shows that the distortions assessed in

paired-comparison experiments must be uncorrelated to determine the actual impact of each IDF on image-difference perception. A visual database collected for such distortions is a prerequisite for optimizing IDMs that can reasonably be used as objective functions for gamut mapping.

Conclusions

We discussed several hypotheses on how the human visual system assesses image differences. We then proposed five image-difference features (IDFs), which are mathematical representations of these hypotheses. The IDFs were combined using three different models (additive, factorial, hybrid); their parameters were trained on an image database collected in gamut-mapping experiments. The image-difference predictions on a test set were compared with those of the SSIM index.

The predictions of our IDMs are significantly more accurate than those of the SSIM index, even if only lightness-based IDFs are used. Adding chroma- or hue-based IDFs (or both) further improves the predictions. Regardless of how many IDFs are combined, the highest accuracy is achieved if the lightness-contrast-based IDF is among them. Generally, the choice of combination model does not affect the prediction performance.

There is still room for improvement regarding the prediction performance: our best IDM achieves a hit rate of 63.4% on the test data, which is far below the maximum achievable hit rate of 77.4%.

As we only used visual data from gamut-mapping experiments, *all* distorted images compared by the observers had similar chroma differences to their respective reference images. As a consequence, the influence of the chroma-based IDF was reduced in the optimization. To avoid such bias, visual data with a greater variety of distortions should be used in the future.




Figure 6. Difference maps of the factorial IDM for the images shown in Fig. 2 (parameters from Table 1). Lighter areas correspond to greater differences.

Acknowledgements

This work would not have been possible without the financial support of the Deutsche Forschungsgemeinschaft (DFG).

References

- [1] G. Hong and M. R. Luo. Perceptually based colour difference for complex images. In *Proceedings of the 9th Congress of the International Colour Association*, pages 618–621, Rochester, NY, 2001.
- [2] 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*, 13(4):600–612, 2004.
- [3] P.-L. Sun and J. Morovic. What differences do observers see in colour image reproduction experiments? In *First European Conference on Color in Graphics, Imaging, and Vision*, pages 181–186, Poitiers, France, 2002.
- [4] I. Lissner, J. Preiss, and P. Urban. Predicting image differences based on image-difference features. In *19th Color and Imaging Conference*, pages 23–28, San Jose, CA, 2011.
- [5] X. Zhang and B. A. Wandell. A spatial extension of CIELAB for digital color image reproduction. *Society for Information Display Symposium Technical Digest*, 27:731–734, 1996.
- [6] I. Lissner and P. Urban. Toward a unified color space for perception-based image processing. *IEEE Transactions on Image Processing*, 21(3):1153–1168, 2012.
- [7] CIE Publication No. 142. Improvement to industrial colour-difference evaluation. Technical report, Central Bureau of the CIE, Vienna, Austria, 2001.
- [8] J. Morović. *Color Gamut Mapping*. John Wiley & Sons, 2008.
- [9] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multi-scale structural similarity for image quality assessment. In *Proceedings of the 37th IEEE Asilomar Conference on Signals, Systems and Computers*, pages 1398–1402, Pacific Grove, CA, 2003.
- [10] P. Zolliker, Z. Barańczuk, and J. Giesen. Image fusion for optimizing gamut mapping. In *19th Color and Imaging Conference*, pages 109–114, San Jose, CA, 2011.
- [11] R. S. Berns. *Billmeyer and Saltzman's Principles of Color Technology*. John Wiley & Sons, New York, 3rd edition, 2000.
- [12] Z. Barańczuk, P. Zolliker, and J. Giesen. Image-individualized gamut mapping algorithms. *Journal of Imaging Science and Technology*, 54(3):030201–(7), 2010.
- [13] P. Zolliker and K. Simon. Retaining local image information in gamut mapping algorithms. *IEEE Transactions on Image Processing*, 16(3):664–672, 2007.
- [14] J. Giesen, E. Schuberth, K. Simon, P. Zolliker, and O. Zweifel. Image-dependent gamut mapping as optimization problem. *IEEE Transactions on Image Processing*, 16(10):2401–2410, 2007.
- [15] F. Dugay, I. Farup, and J. Y. Hardeberg. Perceptual evaluation of color gamut mapping algorithms. *Color Research & Application*, 33(6):470–476, 2008.
- [16] L. Brown and X. Li. Confidence intervals for two sample binomial distribution. *Journal of Statistical Planning and Inference*, 130(1–2):359–375, 2005.

Author Biographies

Jens PREISS received his diploma in Physics (equivalent to a M.S.) from the University of Freiburg (Germany) in 2010. He is currently a research assistant at the Institute of Printing Science and Technology, Technische Universität Darmstadt (Germany), where he works as a doctoral candidate in the area of color and imaging science.

Ingmar LISSNER received his degree in Computer Science and Engineering from the Hamburg University of Technology (Germany) in 2009. He is currently working toward the PhD degree with the Institute of Printing Science and Technology, Technische Universität Darmstadt (Germany). His research interests include color perception, uniform color spaces, and image-difference measures for color images.

Philipp URBAN has been head of an Emmy-Noether research group at the Technische Universität Darmstadt (Germany) since 2009. His research focuses on color science and spectral imaging. From 2006–2008 he was a visiting scientist at the RIT Munsell Color Science Laboratory. He holds a MS in mathematics from the University of Hamburg and a PhD from the Hamburg University of Technology (Germany).

Matthias SCHELLER LICHTENAUER studied Computer Science at ETH Zürich, receiving his Master of Science in 2008. He then joined the Laboratory of Media Technology at Empa where he is researching the subject of design and analysis of psychometric measurements. He is also a PhD candidate in the group of Joachim Giesen at the Friedrich-Schiller-University in Jena (Germany).

Peter ZOLLIKER studied Physics at ETH Zürich and received his PhD in Crystallography from the University of Geneva in 1987. After his postdoc position at the Brookhaven National Laboratory in New York, he joined Gretag Imaging in 1988. Since 2003 he is working at Empa where he is engaged in color management and statistical analysis.