

ColorVideoVDP: A visual difference predictor for image, video and display distortions

RAFAŁ K. MANTIUK, University of Cambridge, UK
PARAM HANJI, University of Cambridge, UK
MALIHA ASHRAF, University of Cambridge, UK
YUTA ASANO, Reality Labs, USA
ALEXANDRE CHAPIRO, Reality Labs, USA

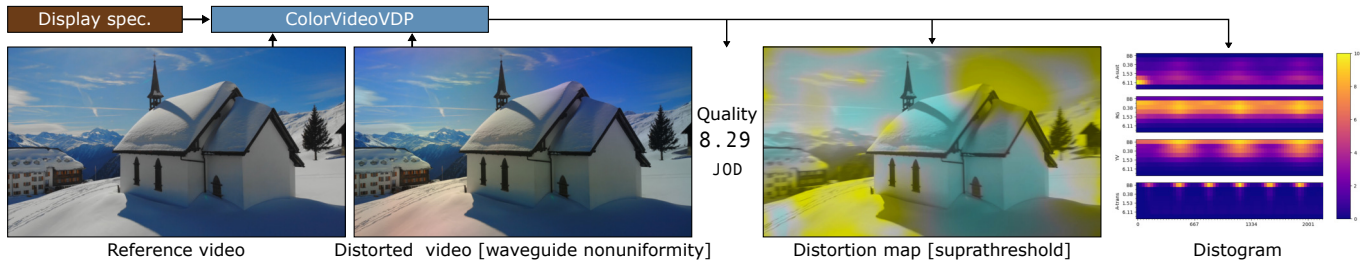


Fig. 1. ColorVideoVDP predicts the visibility of distortions for a pair of test and reference videos (or images) as seen on a display with a provided specification. The predictions are represented as a single quality value in Just-Objectionable-Difference (JOD) units, a distortion map video, and a *distogram*, which visualizes the distortions over time, separately for each channel and spatial frequency band. See Fig. 25 for the content attribution.

ColorVideoVDP is a video and image quality metric that models spatial and temporal aspects of vision for both luminance and color. The metric is built on novel psychophysical models of chromatic spatiotemporal contrast sensitivity and cross-channel contrast masking. It accounts for the viewing conditions, geometric, and photometric characteristics of the display. It was trained to predict common video-streaming distortions (e.g., video compression, rescaling, and transmission errors) and also 8 new distortion types related to AR/VR displays (e.g., light source and waveguide non-uniformities). To address the latter application, we collected our novel XR-Display-Artifact-Video quality dataset (XR-DAVID), comprised of 336 distorted videos. Extensive testing on XR-DAVID, as well as several datasets from the literature, indicate a significant gain in prediction performance compared to existing metrics. ColorVideoVDP opens the doors to many novel applications that require the joint automated spatiotemporal assessment of luminance and color distortions, including video streaming, display specification, and design, visual comparison of results, and perceptually-guided quality optimization. The code for the metric can be found at <https://github.com/gfxdisp/ColorVideoVDP>.

CCS Concepts: • **Computing methodologies** → **Perception**.

Authors' addresses: Rafał K. Mantiuk, rafal.mantiuk@cl.cam.ac.uk, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge, UK, CB3 0FD; Param Hanji, param.hanji@gmail.com, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge, UK, CB3 0FD; Maliha Ashraf, ma905@cam.ac.uk, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge, UK, CB3 0FD; Yuta Asano, yasano@meta.com, Reality Labs, , Redmond, USA, ; Alexandre Chapiro, alex@chapiro.net, Reality Labs, , Sunnyvale, USA, .

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '24 Technical Papers, July 28–Aug 1, 2024, Denver, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/3658144>

Additional Key Words and Phrases: image quality, video quality, visual difference predictor, contrast sensitivity, visual metric

ACM Reference Format:

Rafał K. Mantiuk, Param Hanji, Maliha Ashraf, Yuta Asano, and Alexandre Chapiro. 2024. ColorVideoVDP: A visual difference predictor for image, video and display distortions. In *SIGGRAPH 2024 Technical Papers (SIGGRAPH '24 Technical Papers)*, Jul 28– Aug 1, 2024, Denver, USA. ACM, New York, NY, USA, Article 129, 20 pages. <https://doi.org/10.1145/3658144>

1 INTRODUCTION

Evaluating the visual quality of displayed content is a perennial task in computer graphics and display engineering. The most direct route, involving visual appraisal by human observers, is often too costly and slow. Subjective studies may also be infeasible when a large trade-space of competing variables needs to be studied quickly to find optimal settings. In this case, automated metrics are of great importance as tool for evaluation and design, which are often employed as cost functions for optimization.

This need led to the creation of many general-purpose image and video metrics, but these techniques often ignore important aspects of human vision, such as color or temporal vision. This happens due to the inherent complexity of the visual system, which does not allow for holistic modelling. Further, accurate models typically rely on psychophysical data to make predictions, but due to the multi-dimensional nature of color, data on interactions between color and spatiotemporal characteristics of stimuli can be difficult to model.

Accurate reproduction of contrast and color is of central importance to the quality of displayed content. Achromatic artifacts stemming from graphics pipelines, such as visible blur or contrast loss, can be modelled by existing luminance-only metrics such as

SSIM [Wang et al. 2003] or FovVideoVDP [Mantiuk et al. 2021], but color ones, such as chroma subsampling, cannot. On the other hand, color difference formulas, such as the popular CIEDE2000 [CIE 2018; Sharma et al. 2005], do not model spatial or temporal aspects of vision, and as a consequence may ignore important aspects of an artifact, such as the spatial distribution or changes of color distortion over time. Notably, spatiotemporal color artifacts are especially problematic in modern display applications, in particular for emerging display technologies such as wide-color-gamut, virtual and augmented reality (XR) displays. The latter require novel architectural solutions, which lead to chromatic artifacts like color fringing caused by lens aberrations, or chromatic nonuniformity due to optical waveguides (see Fig. 1).

This work presents ColorVideoVDP, a full-reference quality metric that models spatiotemporal achromatic and chromatic vision. The metric is built on novel psychophysical models of chromatic spatiotemporal contrast sensitivity and cross-channel contrast masking. Thanks to its psychophysical foundations, the metric accounts for the physical specification of a display (size, resolution, color characteristic) and viewing conditions (viewing distance and ambient light). This is the first video and image quality metric that explicitly models human spatiotemporal and chromatic vision simultaneously and is capable of modeling XR display artifacts. Our work builds on FovVideoVDP [Mantiuk et al. 2021] borrowing its general processing pipeline. The main novel contribution is integrating the spatio-chromatic model of near- and supra-threshold contrast perception, which lets us quantify the visibility of both achromatic and chromatic distortions. As chromatic distortions tend to reside in lower spatial frequencies, ColorVideoVDP also improves the modeling of low-frequency differences. Compared to FovVideoVDP, ColorVideoVDP is fully differentiable¹ and much improves the accuracy of quality predictions on non-foveated image and video datasets.

The key challenge of developing any new quality metric is its effective calibration and robust validation. To that end, and to ensure that the metric can provide reliable predictions for display applications, we collected a new XR-Display-Artifact-Video quality dataset (XR-DAVID) with 8 common display artifacts (Sec. 4). To ensure the diversity of distortion and content types, we combined our new XR-DAVID dataset with a large HDR/SDR image dataset UPIQ [Mikhailiuk et al. 2022]. As the combined datasets consist of terabytes of data, calibration required a non-trivial mixture of end-to-end, and feature-space training. This effort allowed us to match and exceed the state-of-the-art results on unseen datasets (in a cross-dataset validation) and on the testing portion of the training datasets (cross-content validation). In Sec. 6, we demonstrate new applications of ColorVideoVDP, including analysis of chroma subsampling, display color tolerance specification, and quantifying observer metamerism variations on a target dataset.

Limitations. ColorVideoVDP lacks higher-level models of saliency or annoyance, resulting in lower accuracy when the semantic content has a strong influence on the quality judgments. It was not trained to predict accurate spatial distortion maps [Ye et al. 2019] (as

¹FovVideoVDP does not model the differences in low-frequency bands and, therefore, it cannot propagate low-frequency differences when used as a differentiable loss term.

no such data is available for video). ColorVideoVDP does not model the effect of glare (inter-ocular light scatter, found in HDR-VDP [Mantiuk et al. 2023]), gaze-contingent vision (found in FovVideoVDP [Mantiuk et al. 2021]), eye motion [Denes et al. 2020; Laird et al. 2006], or binocular vision [Didyk et al. 2011].

2 RELATED WORK

This section reviews the existing work that addresses the problem of predicting visible differences or quality in color images and video — the main focus of our metric. The representative examples of these methods are listed in Table 1. The table also specifies whether the metric attempts to model spatial vision, temporal vision, offers distinctive processing of color and whether it accounts for the colorimetric characteristic of the display, its resolution, and viewing distance. As shown in the table, no existing metric is capable of addressing these four important areas of image quality. Next, we review the metrics by the groups indicated in the table.

Color difference formulas. Perceived color differences for uniform patches can be predicted using one of the standard display formulas, such as CIE ΔE_{ab}^* or CIEDE2000 [CIE 2018; Sharma et al. 2005]. The standard CIE formulas, however, were not meant to predict differences for luminance below 1 cd/m² or above the illuminance of 1 000 lux [CIE 1993] (this translates to approximately 318 cd/m² for a Lambertian white surface). ΔE_{ITP} [ITU-R BT. 2124 2019] was proposed as a color difference formula for luminance levels found in wide-color-gamut high-dynamic-range television, and potentially intended for video content. It is unclear how to use the color difference formula with complex images and the color difference is typically computed per pixel and then averaged. Such treatment obviously ignores all spatial and temporal aspects of vision, which are partially addressed by the next group of metrics.

Spatial color difference formulas. The spatial component of vision was included in a spatial extension of the CIELAB difference formula — sCIELAB [Zhang and Wandell 1997]. The authors proposed to compute the CIE ΔE_{ab}^* color differences on images prefiltered by a contrast sensitivity function (CSF). Flip and HDR-Flip metrics [Andersson et al. 2020, 2021] improve on sCIELAB by employing a color difference formula that better quantifies large color differences. Both metrics also emphasize differences at edges, which tend to be more salient. Choudhury et al. proposed to compute color differences in the ITP color space [2021], which is more suitable for HDR color values. The strength of such spatial extensions is their simplicity. The main weakness is that such an application of the CSF is overly simplistic — it does not account for the changes in contrast sensitivity with luminance and does not account for supra-threshold vision (e.g. contrast masking and contrast constancy).

Image quality. The most popular image quality metrics, such as SSIM or MS-SSIM [Wang et al. 2003], do not attempt to explicitly model human vision, but, instead, they combine hand-crafted statistical measures that are likely to correlate with quality judgments. Although the early metrics, such as SSIM and MS-SSIM, operate only on the luma channel of the image, some later metrics, including FSIMc [Zhang et al. 2011] and VSI [Zhang et al. 2014], separate images into luma and two chroma channels, akin to the color space

Table 1. Quality/difference metrics and their capabilities. Columns indicate whether the metric models spatial vision, temporal vision, color vision, and whether it accounts for display geometry and photometry.

<i>Metric</i>	Spatial	Temporal	Color	Display model	Approach
PSNR	No	No	No	No	Signal quality
CIEDE2000 [CIE 2018]	No	No	Yes	Yes	Color difference formula
ΔE_{ITP} [ITU-R BT. 2124 2019]	No	No	Yes	Yes	Color difference formula
sCIELAB [Zhang and Wandell 1997]	Yes	No	Yes	Yes	CSF + Color difference formula
ΔE_{ITP}^S [Choudhury et al. 2021]	Yes	No	Yes	Yes	CSF + Color difference formula
MS-SSIM [Wang et al. 2003]	Yes	No	No	No	Multi-scale structural similarity
FSIMc [Zhang et al. 2011]	Yes	No	Yes	No	Similarity of phase congruency and gradients
VSI [Zhang et al. 2014]	Yes	No	Yes	No	Saliency + SSIM
LPIPS [Zhang et al. 2018]	Yes	No	Yes	No	Difference of CNN features
FLIP [Andersson et al. 2020]	Yes	No	Yes	No	CSF + Color difference + edge detectors
IQT [Cheon et al. 2021]	Yes	No	Yes	No	CNN features + transformer autoencoder
STRRED [Soundararajan and Bovik 2012]	Yes	Yes	No	No	Entropy differences in wavelet subbands
VMAF [Li et al. 2016a]	Yes	Yes	No	No	Features + SVR
FUNQUE [Venkataramanan et al. 2022]	Yes	Yes	No	No	Wavelet decomposition + features + SVR
HDR-VDP-3 [Mantiuk et al. 2023]	Yes	No	No	Yes	Psychophysical model
FovVideoVDP [Mantiuk et al. 2021]	Yes	Yes	No	Yes	Psychophysical model
ColorVideoVDP (ours)	Yes	Yes	Yes	Yes	Psychophysical model

transforms used in video compression. Those metrics, however, do not account for the geometry of a display (e.g. resolution, size, viewing distance), nor for its photometry (e.g. peak brightness, black level). The latter shortcoming can be addressed by employing the Perceptually Uniform (PU) transform [Mantiuk and Azimi 2021], which also extends these metrics to operate on high-dynamic-range images.

Zhang et al. [2018] observed that the activation layers of many convolution neural networks (CNNs) provide features that are well correlated with human judgments of image similarity. Their proposed metric, LPIPS, became very popular in computer vision, despite its underwhelming performance on image quality datasets [Ding et al. 2021]. Prashnani et al. [2018] proposed training a CNN on triplets of patches, two distorted and one reference, with supervision based on the Bradley-Terry pairwise comparison model. Akin to LPIPS training, their PieAPP metric was trained on a large dataset of patches with pairwise comparison labels. Cheon et al. [2021] combined CNN features with transformer-based embeddings to regress quality scores. Their IQT metric was ranked first among 13 participants in the NTIRE 2021 perceptual image quality assessment challenge.

Image metrics are not meant to predict video quality, however, they can perform surprisingly well in this task. The typical route to employ image metrics to video is to average predictions across individual frames.

Video quality. Video quality metrics combine both spatial and temporal features. STRRED [Soundararajan and Bovik 2012] compares the per-frame entropy of fitted local distributions of wavelet coefficients. This entropy is then weighted by local spatial and temporal variance. VMAF [Li et al. 2016a] combines two spatial features (VIF and DLM) with a mean absolute difference of consecutive frames, and then maps those into quality scores using a support

vector regression (SVR). The success of VMAF motivated a series of fusion-based metrics, such as FUNQUE [Venkataramanan et al. 2022] and 3C-FUNQUE+ [Venkataramanan et al. 2023], which use wavelet transform modulated by the contrast sensitivity function as input to several individual predictors, which similarly as VMAF, are combined using SVR. However, all those metrics use simplified models of contrast sensitivity, which ignore the effect of luminance and do not model spatiotemporal vision. For example, the temporal processing of all of these state-of-the-art metrics considers just two consecutive frames, and it does not model the temporal characteristics of human vision. Our metric explicitly models achromatic and chromatic visual channels to address this shortcoming.

Visual Difference Predictors. The visual difference predictors (VDPs), such as DCTune [Watson 1993], VDP [Daly 1993], HDR-VDP-2 [Mantiuk et al. 2011] and HDR-VDP-3 [Mantiuk et al. 2023], explicitly model aspects of low-level human vision, such as contrast sensitivity and masking. The advantage of this approach is that these metrics are built on sound psychophysical models and generalize more easily to unseen conditions, such as displays of varying size, resolution, or peak luminance. While it is tempting to assume that this type of general modeling would perform worse or be more computationally expensive than metrics relying on hand-crafted features, as we will demonstrate in Sec. 5, this is not the case. Modern VDPs perform on par or better than metrics with hand-crafted features. When they are optimized to run on a GPU, they are as fast as feature-based metrics.

Our ColorVideoVDP metric takes inspiration from and is based on similar building blocks as FovVideoVDP [Mantiuk et al. 2021] but with several important improvements. First, ColorVideoVDP models the visibility of chromatic (color) differences by employing a novel spatiotemporal-chromatic contrast sensitivity function (castleCSF)

[Ashraf et al. 2024]. This involves modeling both spatial and temporal chromatic vision. Second, we account for supra-threshold color differences so that the perceived magnitude of achromatic and chromatic contrast is properly mapped to the metric response. Third, while FovVideoVDP is insensitive to low-frequency distortions (below 1 cpd²), ColorVideoVDP correctly accounts for those. This lets us account for low-frequency display artifacts (e.g., waveguide non-uniformity) and use ColorVideoVDP as a differentiable loss function, which can propagate both low and high-frequency differences. Fourth, ColorVideoVDP models within-channel and cross-channel masking, where each of the four modeled channels (two achromatic and two chromatic) can be masked by the combination of contrast in the other channels. Finally, ColorVideoVDP is trained and tested on multiple image and video datasets, both SDR and HDR, including a novel dataset of display artifacts described in Sec. 4. Unlike FovVideoVDP, we do not model foveated vision due to the increased overhead of modeling interactions between eccentricity and color. While modeling foveated vision is necessary in some specific usage scenarios, such as foveated rendering, it is less useful in most general-use cases, such as video streaming, display engineering, and non-foveated algorithm design.

3 COLOR VIDEO VISUAL DIFFERENCE PREDICTOR

Our goal is to create an image and video quality metric that accounts for color and spatiotemporal perception and has sound psychophysical foundations. The metric should rely on psychophysical models as they can help to extrapolate predictions to unseen conditions, such as different frame rates, spatial resolution, or absolute luminance levels. The metric should be able to predict a single-valued quality correlate that can help to both evaluate and optimize visual results. It should also produce spatial and temporal difference maps, which provide the visual explanation for the predicted quality correlate. Finally, the metric should be efficient to compute and fully differentiable so that it could be used as an optimization criterion.

The processing diagram of the metric is shown in Fig. 2. First, the test and reference content, typically stored as video or images, is transformed into colorimetric quantities of light emitted from a given display. Then, the frames are decomposed into color opponent channels, temporal channels and spatial frequency bands, mimicking the mechanisms of the visual system. The core component of the metric is the model of contrast sensitivity and masking, which computes a per-band visual difference between test and reference content. It relies on the near-threshold and supra-threshold models of contrast detection and discrimination. In the last two steps, the visual differences are pooled across the bands and channels and then the resulting visual difference value is regressed into an interpretable Just-Objectionable-Difference (JOD) scale. The following sections explain each step in detail.

3.1 Display model

The display model has two purposes: to convert spatial pixel coordinates to perceptually meaningful units of visual degrees, and to model the display’s photometric response in a given environment. We assume a flat panel display spanning a limited field of view,

²cpd – cycles per visual degree – the measure of spatial frequency.

so we can approximate the conversion from pixel coordinates to degrees in the visual field with a single constant:

$$n_{\text{ppd}} = \frac{\pi}{360 \arctan\left(\frac{0.5 d_{\text{width}}}{r_w d_v}\right)}, \quad (1)$$

where n_{ppd} is the effective display resolutions in pixels per degree, d_{width} is the width of the screen, r_w is the screen’s horizontal resolution in pixels and d_v is the viewing distance. The display width and viewing distance must be provided in the same units (e.g. meters). We rely on n_{ppd} later when expressing the spatial frequencies in cycles per visual degree. It should be noted that the above approximation is inaccurate for near-eye displays spanning a large field of view, and accurate geometric mapping should be used for these types of displays (see Equation 2 in [Mantiuk et al. 2021]).

The second responsibility of the display model is to convert pixel values represented in one of the standard color spaces into colorimetric quantities of light emitted from a given display. It accounts for the display’s peak luminance, color gamut, its black level, and ambient light reflected from the display. The display-encoded pixel values, $I_{\text{de},c}$ for color channel c ($c \in \{R, G, B\}$), are transformed into absolute linear colorimetric values:

$$I_{\text{lin},c}(\mathbf{x}) = \min\{(L_{\text{peak}} - L_{\text{black}})E(I_{\text{de},c}(\mathbf{x})) + L_{\text{black}}, L_{\text{peak}}\} + L_{\text{refl}}, \quad (2)$$

where L_{peak} is the peak luminance of the display and L_{black} is its black level. \mathbf{x} is used to denote spatial pixel coordinates throughout the paper. $E(\cdot)$ is the electro-optical transfer function (EOTF) of particular pixel coding, for example, the sRGB non-linearity (IEC 61966-2-1:1999) for standard dynamic range content and PQ (Perceptual Quantizer, SMPTE ST 2084) for high dynamic range content. Because the PQ EOTF transforms display-encoded values into absolute linear values (between 0.005 and 10 000), we do not multiply the EOTF by $(L_{\text{peak}} - L_{\text{black}})$ if PQ is used. We currently do not model tone-mapping, which is present in most displays with HDR capabilities, because it varies from one display to another. Instead, we clip the values at L_{peak} .

The amount of light reflected from the display, L_{refl} , is computed as:

$$L_{\text{refl}} = k_{\text{refl}} \frac{E_{\text{amb}}}{\pi}, \quad (3)$$

where k_{refl} is the reflectivity of the screen (typically 0.01–0.05 for glossy screens, 0.005–0.015 for matt screens) and E_{amb} is the ambient illumination in lux units. As the last step, the linear color values are converted into device-independent CIE XYZ color space. This conversion is standardized for popular color spaces, such as BT.709 or BT.2020, or alternatively can be computed for the primaries of a display.

3.2 Opponent color channels

The sensitivity to chromatic changes is typically explained for color modulations represented in a space that separates three cardinal mechanisms of human color vision: achromatic channel and two chromatic channels, the latter commonly known as red-green and violet-yellow [Stockman and Brainard 2010]. Here, we use the same color space as our contrast sensitivity function – the Derrington-Krauskopf-Lennie (DKL) colorspace [Derrington et al.




Fig. 2. Processing stages of ColorVideoVDP. Test and reference images are first processed by the same pipeline: the display model maps pixel values to linear color (CIE 1931 XYZ color space), linear color is transformed to the opponent color space (DKL), the achromatic channel is decomposed into sustained (continuous lines) and transient (dashed lines) temporal channels, then each of those is decomposed into multiple spatial bands (Laplacian pyramid). The decomposed video/image goes into the contrast sensitivity and masking models, explained in more detail in Figure 6. The result of the masking model is pooled across all spatial bands, temporal and color channels, and finally regressed to a JOD score.




Fig. 3. The frequency characteristic of the four temporal channels used in ColorVideoVDP.

1984]. DKL was selected for castleCSF as this color space is linearly related to cone responses, is well established in vision science, and much of the chromatic contrast detection data was collected in that space. The DKL space coordinates can be computed from the device-independent XYZ (provided by the display model) as:

$$\begin{bmatrix} I_{\text{ach}}(\mathbf{x}) \\ I_{\text{rg}}(\mathbf{x}) \\ I_{\text{vy}}(\mathbf{x}) \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 1 & -\frac{L_0}{M_0} & 0 \\ -1 & -1 & \frac{L_0+M_0}{S_0} \end{bmatrix} M_{\text{XYZ} \rightarrow \text{LMS}} \begin{bmatrix} I_X(\mathbf{x}) \\ I_Y(\mathbf{x}) \\ I_Z(\mathbf{x}) \end{bmatrix}, \quad (4)$$

where $M_{\text{XYZ} \rightarrow \text{LMS}}$ is a matrix converting CIE 1931 XYZ coordinates into the LMS cone responses³:

$$M_{\text{XYZ} \rightarrow \text{LMS}} = \begin{bmatrix} 0.187596 & 0.585169 & -0.026384 \\ -0.133397 & 0.405506 & 0.034502 \\ 0.000244 & -0.000543 & 0.019407 \end{bmatrix}. \quad (5)$$

L_0 , M_0 and S_0 specify chromaticity of the adapting color. Here, we assume adaptation to a D65 background: CIE 1931 $(x, y) = (0.3127, 0.3290)$, $(L_0, M_0, S_0) = (0.7399, 0.3201, 0.0208)$.

3.3 Temporal channels

Psychophysical masking experiments showed evidence that the information is processed in the visual system by separate temporal

channels. Two or three channels have been identified for the achromatic mechanism [Anderson and Burr 1985; Hess and Snowden 1992], and one or two channels for the chromatic mechanisms [Cass et al. 2009; McKeefry et al. 2001]. Here, we assume two temporal achromatic channels, as the third channel was observed only for low frequencies [Hess and Snowden 1992]. We assume just one temporal channel for the red-green and violet-yellow cardinal directions as the evidence for the second channel shows that it has a less prominent role [Cass et al. 2009]. Modeling fewer temporal channels also has the benefit of lower memory and computational cost.

The temporal channels of ColorVideoVDP are directly defined by castleCSF [Ashraf et al. 2024] (see next section). The frequency tuning of the channels can be seen in Fig. 3. To find digital filters, we perform a real-valued (symmetric) fast Fourier transform on the frequency space filters. We found that a filter support of 250 ms is sufficient to capture filter characteristics. Finally, the two achromatic and two chromatic channels are convolved with the digital filters along the time dimension. This splits the achromatic signal into sustained (low-pass) and transient (band-pass) channels. The chromatic channels are low-pass filtered and become insensitive to high-frequency chromatic flicker.

3.4 Contrast sensitivity

A contrast sensitivity function models our ability to detect patterns of different spatial and temporal frequency, size, and shown at different luminance levels. It is a cornerstone of ColorVideoVDP — it defines the temporal and chromatic channels and enables modeling of contrast masking, which is the key component of our metric. Since there is no contrast sensitivity model that could explain sensitivity to both chromatic modulations and different temporal frequencies, we have created our own model, named castleCSF. This CSF models color, area, spatial and temporal frequencies, luminance, and eccentricity. Because of its complexity, castleCSF is explained in detail in a separate paper [Ashraf et al. 2024]. Here, we summarize its key components.

castleCSF decomposes contrast into three cardinal directions of the DKL color space, the same as those we used to separate achromatic and chromatic channels in Sec. 3.2. The achromatic direction is split into sustained and transient channels, again the same as used by our metric. That lets us directly map the mechanism modeled in castleCSF to the channels in ColorVideoVDP.

³Our CSF is defined using CIE 2006 color matching functions while most of the content still relies on the CIE 1931 color matching functions. The matrix was derived to convert between the two using the spectral emission data for an LCD with an LED backlight.

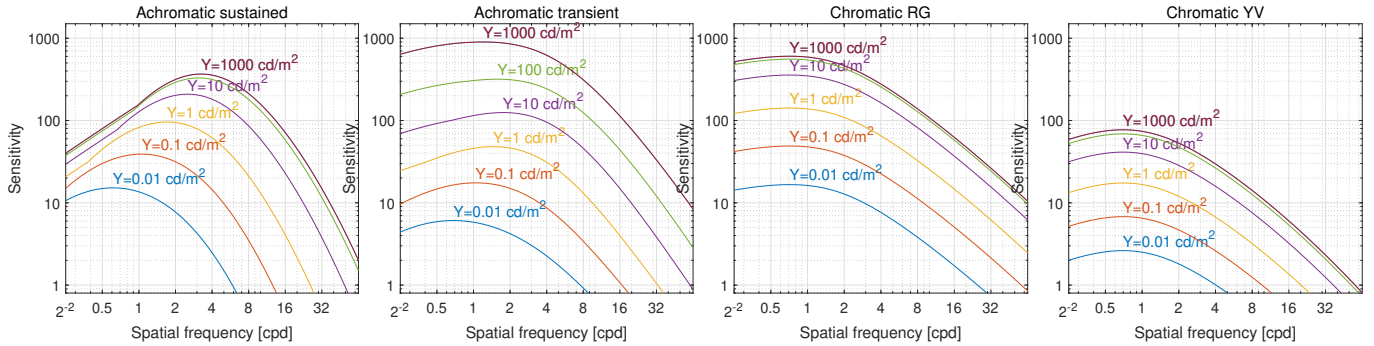


Fig. 4. castleCSF contrast sensitivity function for the four channels of ColorVideoVDP. The sensitivity is expressed in the cone-contrast units [Wuerger et al. 2020]. Note that the achromatic transient and chromatic RG channels appear to have higher sensitivity than the achromatic sustained channel. This is due to the scaling used in the DKL color space [Derrington et al. 1984] to represent chromatic contrast units. In practice, the transient and two chromatic channels are much less sensitive to patterns found in complex images.

To find the detection threshold, castleCSF modulates (multiplies) the contrast associated with each mechanism by the sensitivity of that mechanism and then pools those to form the contrast energy. The detection threshold is assumed to be the contrast at which the energy is equal to 1. The sensitivity of each mechanism is modeled as a function of spatial and temporal frequency, area, background luminance and eccentricity, in a similar manner as for stelaCSF [Mantiuk et al. 2022]. castleCSF was optimized to predict the data from 19 contrast sensitivity datasets (10 achromatic, 6 chromatic and 3 mixed).

castleCSF predicts sensitivity as the inverse of cone contrast [Wuerger et al. 2020] while ColorVideoVDP operates on a contrast in the DKL color space. To obtain the sensitivity units consistent with our contrast definition, we use the contrast transformation procedure from [Kim et al. 2021] and explained in more detail in the supplementary. Because we do not account for foveated viewing in ColorVideoVDP, we assume that the eccentricity is 0. We found in ablations that the metric performs the best when the area is set to $a = \pi 1.5^2 \text{ deg}^2$. The resulting sensitivity functions are plotted in Fig. 4.

3.5 Multi-scale decomposition and color contrast

In addition to the temporally-tuned channels, psychophysical [Foley 1994; Stromeyer and Julesz 1972] and neuropsychological [De Valois et al. 1982] evidence proves the existence of channels that are tuned to the bands of spatial frequencies and orientations. We must account for such channels to model visual masking, as explained in the next section. Following FovVideoVDP [Mantiuk et al. 2021], we use the Laplacian pyramid [Burt and Adelson 1983] to decompose each of the four temporal channels (Y-sustained, Y-transient, RG, YV) into spatial-frequency selective bands. However, unlike FovVideoVDP, we consider low frequencies, including the base band (the coarsest low-pass band in the Laplacian pyramid). This change is due to finding that many display distortions can only be detected in the low-spatial frequency bands (see Fig. 1). We select the height of the pyramid so that the lowest peak frequency of the band-pass filter is greater than 0.2 cpd. Similarly as in FovVideoVDP, we do not

consider orientation-selective channels because of the substantial computational overhead of those.

The contrast at the spatial frequency band b , channel c and frame f is computed as:

$$C_{b,c,f}(\mathbf{x}) = \frac{\mathcal{L}_{b,c,f}(\mathbf{x})}{\uparrow \mathcal{G}_{b+1,S,f}(\mathbf{x})} = \frac{\mathcal{L}_{b,c,f}(\mathbf{x})}{L_{\text{bkg},b,f}(\mathbf{x})}, \quad (6)$$

where $\mathcal{L}_{b,c,f}$ represents the b -th band of the Laplacian pyramid, $\mathcal{G}_{b+1,S,f}$ is the $(b+1)$ band of the Gaussian pyramid and \uparrow is the upsampling (expand) operator. The Laplacian pyramid is created for each frame of each channel, obtained by temporal filtering (Sec. 3.3). The Laplacian pyramid coefficients are divided by the values from the Gaussian pyramid at the band $b+1$ (one lever coarser) of the sustained luminance channel (S). Note that the upsampled version of the Gaussian pyramid band $b+1$ is a by-product of computing a Laplacian pyramid so can be obtained at no computational cost. This formulation is similar to that of local band-limited contrast [Peli 1990], but here we extend it to both achromatic and chromatic channels. These contrast values are computed separately for the test and reference frames. The values in all bands except the low-pass (base-band) and the high-pass bands are multiplied by 2 to account for the reduction in amplitude [Mantiuk et al. 2021, Fig. 5].

3.6 Cross-channel contrast masking

The model of cross-channel contrast masking transforms local physical contrast differences between the test and reference frames into perceived differences — differences that are scaled by the local contrast visibility. It accounts for both contrast sensitivity, such as lower sensitivity to high spatial and temporal frequencies, and contrast masking. Masking accounts for differences being less likely to be noticed in heavily textured areas. The cross-channel component models that a strong contrast in one channel can reduce the visibility of contrast in another channel [Switkes et al. 1988], for example, contrast in the chromatic channels can mask contrast in achromatic channels. Finally, the masking model also needs to account for suprathreshold contrast perception — must match the perceived magnitude of contrast (changes) across luminance, frequencies, and the directions of color modulation.

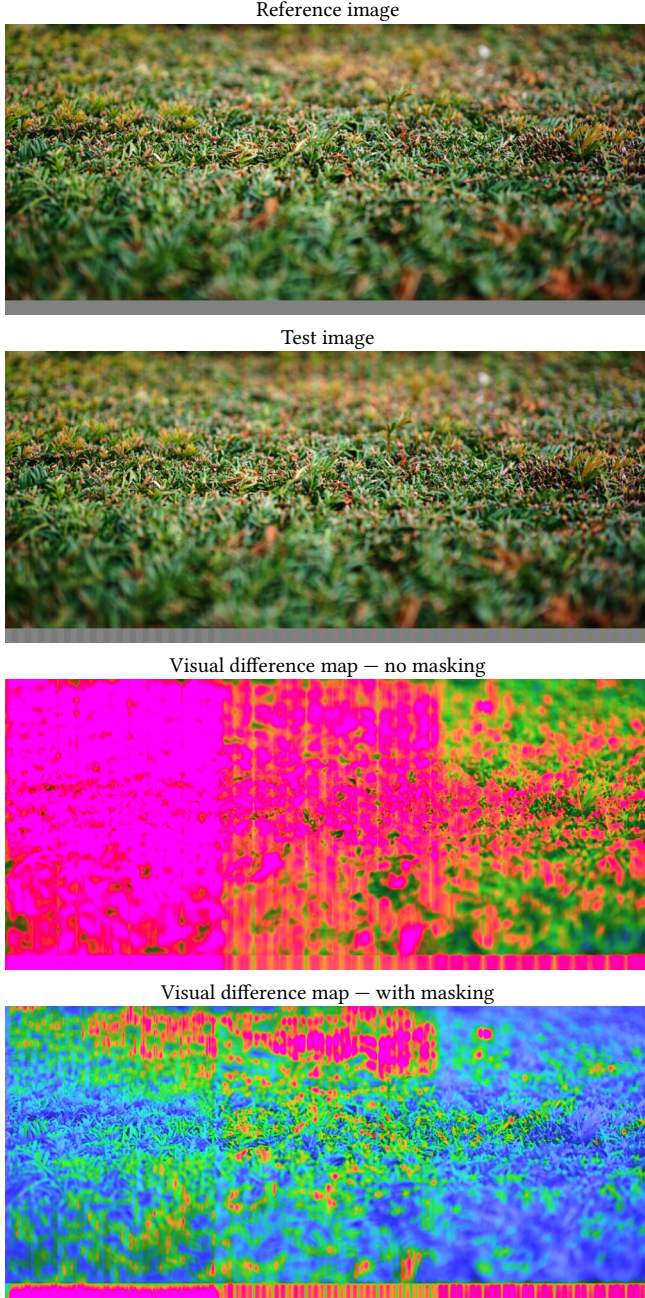


Fig. 5. Example of contrast masking. A sinusoidal grating of 4 cpd (when seen printed from 40 cm) was added to a reference image (top row) to obtain a distorted image (second row). The image was split into three parts and the grating was modulated along achromatic, red-green, and yellow-violet directions in each respective part. The third row shows the visual difference map generated by ColorVideoVDP without masking but still using the CSF. The map over-predicts the visibility in textured areas. The bottom row shows the prediction with the masking model. It is worth noting; although all three color directions are equally visible in the gray bar at the bottom of the image (where there is little masking), the red-green pattern is more visible in the textured area because it is weakly masked by the achromatic channel. ColorVideoVDP correctly predicts this phenomenon.

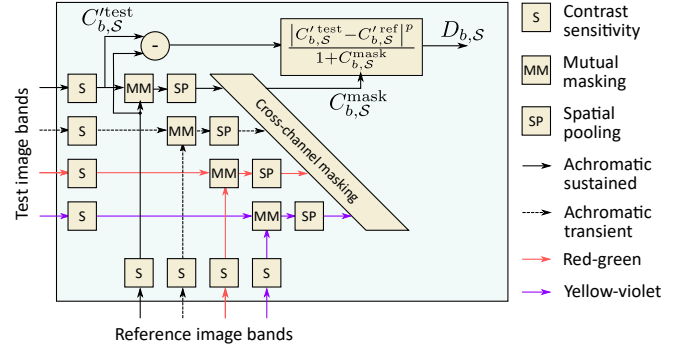


Fig. 6. Our masking model. Here, the resulting visual difference for the sustained channel, $D_{b,S}$, is visualized for a single spatial frequency band (the frame index f is omitted for clarity). Each band response is multiplied by the contrast sensitivity function (“S” boxes). The CSF-normalized band-responses are used to calculate the difference between test and reference frames. The masking signal is computed by first finding mutual masking between both channels (“MM” blocks), applying a spatial pooling in the local neighborhood of each pixel (“SP” blocks) and then combining the masking signal from multiple channels (cross-channel masking). The visual per-channel and per-band difference between the test and reference is calculated as the ratio of excitatory difference between the test and reference images, and the inhibitory masking signal, as shown in the equation in the box and in Eq. (9).

Contrast masking is the critical component of the metric that determines its performance, as we show later in the ablation studies (Sec. 5.2). To make the optimal choice, we analyzed and compared six models: two contrast encodings (additive and multiplicative) combined with three masking models. We tested the mutual masking model from the original VDP [Daly 1993], the contrast transducer proposed by Watson and Solomon [1997], and the contrast similarity formula used in SSIM and many other metrics. To keep this text concise, we describe below only the mutual masking model, which performed the best in our tests. We encourage the reader to refer to the supplementary materials with the detailed description and analysis of all the models.

Contrast masking shows different characteristics for different spatial frequency bands and color channels. However, it was shown that for both luminance [Daly 1993] and chromatic channels [Cass et al. 2009; Switkes et al. 1988] masking characteristics can be unified if both the test and masker contrast are multiplied by the contrast sensitivity function $S_{b,c}(\mathbf{x})$:

$$C'_{b,c,f}(\mathbf{x}) = C_{b,c,f}(\mathbf{x}) S_{b,c,f}(\mathbf{x}). \quad (7)$$

In our case, contrast sensitivity is provided by the castleCSF model:

$$S_{b,c,f}(\mathbf{x}) = s_{\text{corr}} s_{\text{ch},c} S_c(\rho_b, \omega_c, L_{\text{bkg},b,f}^{\text{ref}}(\mathbf{x})), \quad (8)$$

where ρ_b is the spatial frequency of band b in cycles per degree, ω_c is the peak temporal frequency of channel c , and $L_{\text{bkg},b,f}^{\text{ref}}(\mathbf{x})$ is the local background luminance for the reference image (see Eq. (6)). s_{corr} is a trainable parameter adjusting the absolute sensitivity of the model (all trained parameters can be found in Table 3). The additional benefit of the multiplicative contrast encoding from

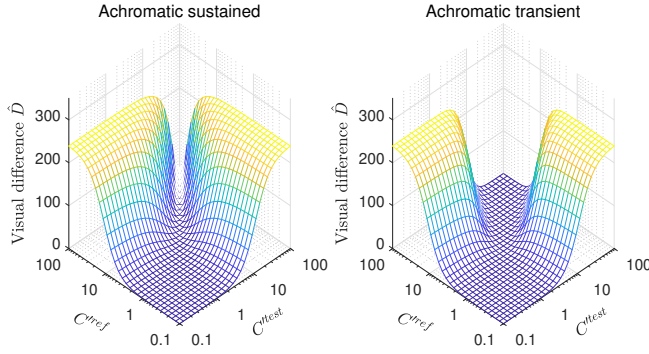


Fig. 7. The visual difference for a pair of test and reference contrast values (normalized by the CSF), as predicted by the masking model (see Eq. (9) and Eq. (13)). The visual difference D is zero when the test and reference contrasts are identical and increases as they start to differ. However, there is no increase for small contrast values due to contrast sensitivity.

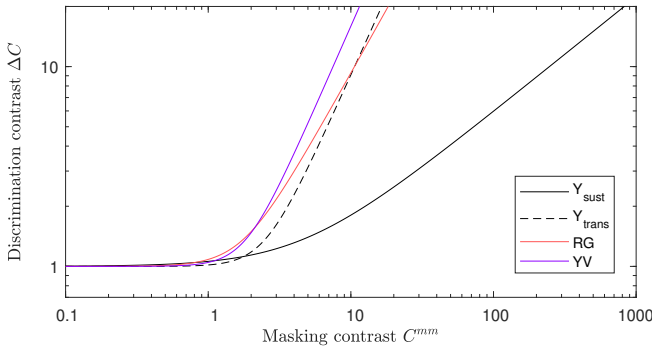


Fig. 8. Contrast discrimination functions for our masking model for the sustained and transient channels. The value on the y-axis denotes the difference in the normalized contrast ($\Delta C = |C'^{\text{test}} - C'^{\text{ref}}|$) required to obtain the response $D = 1$ (see Eq. (9)).

Eq. (7) above is that it helps in matching the perceived magnitude of suprathreshold contrast across color modulation directions [Switkes and Crognale 1999], luminance [Peli 1995] and to a lesser extent across spatial frequencies (see the full analysis in the supplementary). However, we found that CSF alone is too inaccurate to match suprathreshold contrast across color modulation directions, and we had to introduce the color matching correction factor $s_{\text{ch},c}$ based on the contrast matching data of Switkes and Crognale [1999]: $s_{\text{ch},c} = [1 \ 1 \ 1.7 \ 0.237]$ corresponding to achromatic sustained, transient, red-green and yellow-violet channels (see the full explanation in the supplementary).

Once the encoded contrast C' is calculated, the visual difference between the bands is calculated as:

$$D_{b,c,f}(\mathbf{x}) = \frac{|C'_{b,c,f}{}^{\text{test}}(\mathbf{x}) - C'_{b,c,f}{}^{\text{ref}}(\mathbf{x})|^p}{1 + C_{b,c,f}^{\text{mask}}(\mathbf{x})}, \quad (9)$$

where p is a parameter of the model. The masking signal $C_{b,c}^{\text{mask}}$ combines local contrast across test and reference images, local spatial

neighborhood, and channels. First, similarly as in [Daly 1993, p.192], we calculate the mutual masking of test and reference bands (see also Fig. 6):

$$C_{b,c,f}^{\text{mm}}(\mathbf{x}) = \min \left\{ |C'_{b,c,f}{}^{\text{test}}(\mathbf{x})|, |C'_{b,c,f}{}^{\text{ref}}(\mathbf{x})| \right\}. \quad (10)$$

Then, the mutual masking signal is pooled in a small local neighborhood by convolving with a Gaussian kernel $g_{\sigma_{\text{sp}}}$, and is combined across channels, accounting for the cross-channel masking:

$$C_{b,c,f}^{\text{mask}}(\mathbf{x}) = \sum_i k_{i,c} ((C_{b,i,f}^{\text{mm}})^{q_c} * g_{\sigma_{\text{sp}}})(\mathbf{x}). \quad (11)$$

where $k_{i,c}$ is the cross-channel masking coefficient, describing the contribution of channel i to the masking signal of channel c . The weights for the trained model are shown in Fig. 9. The exponents q_c are model parameters set separately for each channel (two achromatic and two chromatic channels). One interpretation of Eq. (9) is that the division by the absolute amplitude of the neighbors reduces redundancy in natural images since neighboring pixels are highly correlated [Laparra et al. 2016]. Such normalization was shown to provide encoding that better correlates with our perception of differences.

The (low-pass) baseband of the Laplacian pyramid cannot be used to calculate contrast, as done in Eq. (6). Instead, the differences in the baseband are computed directly on the Gaussian pyramid coefficients, which are multiplied by the sensitivity:

$$D_{B,c,f}(\mathbf{x}) = k_{B,c} \left| \mathcal{G}_{B,c,f}^{\text{test}}(\mathbf{x}) - \mathcal{G}_{B,c,f}^{\text{ref}}(\mathbf{x}) \right| S_{B,c,f}(\mathbf{x}), \quad (12)$$

where B is the index of the baseband. Because baseband differences are in different units than those in other bands, we need to introduce a trainable scaling factor $k_{B,c}$, which varies across the channels.

One limitation of the original mutual masking model is that it results in excessively large contrast values when the sensitivity is high, and there is no masking signal (refer to the supplementary). We found it essential to restrict the maximum contrast in each band so that a few very large differences do not have an overwhelming impact on the predicted quality. We achieve that with a soft-clamping function

$$\hat{D}_{B,c,f}(\mathbf{x}) = \frac{k_C D_{B,c,f}(\mathbf{x})}{k_C + D_{B,c,f}(\mathbf{x})}, \quad (13)$$

where k_C is a trainable parameter. The equation accounts for the limited dynamic range of the retinal cells, which cannot encode large contrast values.

The shape of the masking function for the sustained and transient achromatic channels is plotted in Fig. 7. Another visualization of the masking model is shown Fig. 8 as a discrimination vs. masking plot. The plot shows that, as the mutual masking contrast C^{mm} increases, a higher difference between the test and reference bands is needed to trigger the same response ($D_{b,c,f}(\mathbf{x}) = 1$ in this example). Such an increase is more gradual for the sustained achromatic channel and more abrupt for the transient channel. The findings of Switkes et al. [1988] indicate that color can robustly mask luminance patterns, while luminance does not mask color but instead facilitates the detection of color patterns. Remarkably, the cross-channel masking weights obtained in metric fitting (Sec. 5) and shown in Fig. 9 confirm these findings – weights are very small in the 2nd and 3rd column of the top row, showing little influence of



Fig. 9. Weights of the cross-channel masking model — $k_{i,c}$. The values are normalized to one in each column to better show the contribution of each channel. The absolute values can be obtained by multiplying the cell value by a constant shown in the parenthesis at the bottom of each column. The matrix shows that both chromatic channels strongly mask the sustained achromatic channel (1st column), but the achromatic channels do not mask chromatic channels.

luminance (achromatic sustained channel) on masking of the color channels. Note that our mutual masking model cannot account for facilitation. An example of this effect is shown in Fig. 5, in which luminance contrast does not mask red-green chromatic contrast.

3.7 Pooling

Once the differences between the bands, $D_{b,c,f}$, are computed, they need to be pooled into a single quality correlate. We follow a similar strategy as done in FovVideoVDP and pool the differences across all the spatial dimensions (\mathbf{x}) in each band, across spatial frequency bands (b), across the channels (c), and finally across all the frames (f):

$$D_{\text{pooled}} = \frac{1}{F^{1/\beta_f}} \left\| \left\| w_c \left\| \frac{1}{N_b^{1/\beta_x}} \left\| \hat{D}_{b,c,f}(\mathbf{x}) \right\|_{\beta_x, \mathbf{x}} \right\|_{\beta_b, b} \right\|_{\beta_c, c} \right\|_{\beta_f, f}, \quad (14)$$

where $\|\cdot\|_{p,v}$ is a p -norm over the variable v :

$$\|f(v)\|_{p,v} = \left(\sum_v |f(v)|^p \right)^{1/p}. \quad (15)$$

w_c is the weight associated with each channel. The main advantage of pooling differences first across pixels is that we avoid the expensive step of the synthesis of the distortion map from the multiple levels of the Laplacian pyramid. It should be noted that the differences across pixels are normalized by the number of pixels in each band, N_b — we do not want the bands represented with more pixels in the pyramid to contribute more to image quality. The pooling exponents β_x and β_b were optimized in FovVideoVDP, but we found that such an optimization is unstable (because of exploding gradients) and unnecessary. Instead, we set $\beta_x = 2$ and $\beta_f = 2$ to represent the energy summation across the spatial dimensions and frames [Watson et al. 1983]. β_b and β_c are set to 4, which is representative of the summation across channels [Quick et al. 1978].

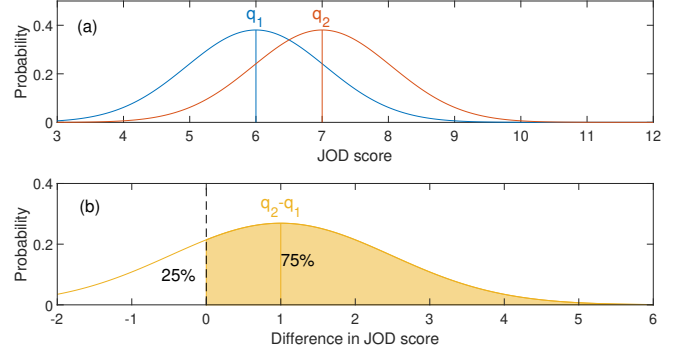


Fig. 10. Interpretation of the JOD scores. (a) JOD quality scores are normally distributed random variables — the reported JOD value is the mean of the distribution, which models that a certain portion of the population will rate the quality as higher or lower than the reported mean value. (b) The standard deviation of the distribution is selected in such a way that when the difference between the means is 1 JOD ($q_2 - q_1 = 1$), 75% of the population will select the condition with the higher JOD value as better. This allows us to interpret quality differences expressed in JOD units.

Using exponents greater than 1 roughly corresponds to a "winner-take-all" strategy, in which the strongest visual differences have the highest impact on the pooled value.

3.8 JOD regression

The visual difference correlate is regressed into interpretable units of just-objectionable-differences (JODs), using the same formula as FovVideoVDP:

$$Q_{\text{JOD}} = 10 - \alpha_{\text{JOD}} (D_{\text{pooled}})^{\beta_{\text{JOD}}}, \quad (16)$$

where α_{JOD} and β_{JOD} are tuned parameters. Here, 10 JOD represents the highest quality — when test and reference images are identical. The JOD units are scaled in terms of inter-observer variance — the drop in quality of 1 JOD means that 75% of observers will notice such a loss of quality in a pairwise comparison experiment. This concept is illustrated in Fig. 10.

3.9 Image quality

When predicting the quality of images, we use the same processing stages as for video but with two changes. First, we skip temporal decomposition and do not create the achromatic transient channel — all computations are performed on three channels. Second, we replace the temporal pooling from Eq. (14) with a single, trainable constant k_1 :

$$D_{\text{pooled}}^{\text{image}} = k_1 \left\| \left\| w_c \left\| \frac{1}{N_b^{1/\beta_x}} \left\| \hat{D}_{b,c,f}(\mathbf{x}) \right\|_{\beta_x, \mathbf{x}} \right\|_{\beta_b, b} \right\|_{\beta_c, c} \right\|. \quad (17)$$

k_1 could be interpreted as a fixation time. It lets us unify the quality between images and video.

3.10 Visualization — heatmaps and distograms

ColorVideoVDP offers two types of visualization that help to interpret the quality score. We can overlay a heatmap with per-pixel

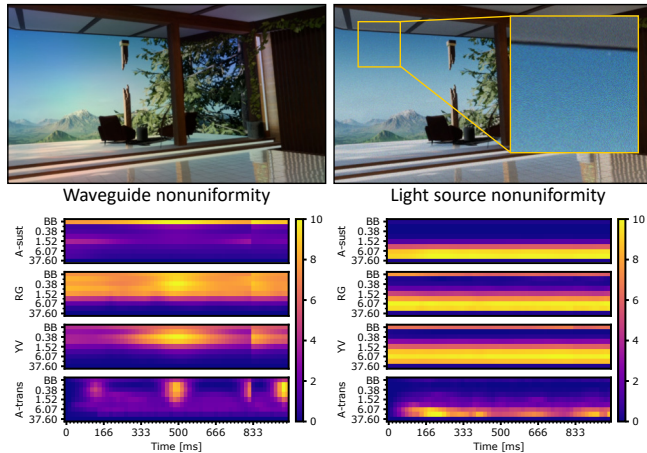


Fig. 11. Examples of *distogram* visualization of artifacts. The values on the y-axis denote spatial frequency in cpd. The higher (more yellowish) values denote stronger distortion. The low-frequency waveguide nonuniformity (left) is shown to be visible mostly in the base band (BB) and low frequency bands of both achromatic (A-sust) and chromatic (RG, YV) channels. The artifacts shows also as low-frequency flicker in the achromatic transient channel (A-trans). The (per-pixel) light source nonuniformity causes high-frequency shimmer, which is visible as high frequency distortions in the distogram on the right.

distortion intensity over the grayscale version of distorted content, as shown in Fig. 1. To obtain such a per-pixel distortion map, we pool the distortions across the channels

$$H_b(\mathbf{x}) = \left\| w_c \hat{D}_{b,c,f}(\mathbf{x}) \right\|_{\beta_{c,c}} \quad (18)$$

and perform the Laplacian pyramid synthesis step to obtain the map with visual difference correlates per pixel and per frame. Those can then be transformed into the JOD units using Eq. (16). It must be noted that the metric has not been trained to produce accurate heatmaps as there are no datasets that could be used for that purpose (such difference maps have been collected for images [Ye et al. 2019] but not for video). The heatmaps are meant to help interpret the single-valued JOD quality predictions.

Another visualization, which we name a *distogram*, lets us present video distortions within each channel and spatial frequency band, all in a single diagram. Two examples for two distortions from the XR-DAVID dataset are shown in Fig. 11. The examples show that depending on the characteristics of the distortion, the artifacts will show in either low- or high-frequency bands, in achromatic or chromatic channels. This visualization decomposes the distortions into their frequency components, visual channels, and their time series, which helps interpret and explain how those distortions contributed to the final JOD score.

3.11 Implementation details

ColorVideoVDP is implemented in PyTorch and optimized for fast execution on a GPU. To take advantage of the massive parallelism offered by the GPU, we load as many video frames as we can process into GPU memory. Then, the operations are executed in parallel on



Fig. 12. Physical hardware setup for the XR-DAVID data collection. Participants used a chinrest to fix their distance from the display at 0.73 m (two display heights) and used a 3-key keyboard to record their choices.

a set of test and reference frames, all four channels, and all pixels (all stored in a single tensor); only the pyramid bands are processed sequentially as each one has a different resolution. To avoid the computational overhead of castleCSF, the function is precomputed and stored as a set of 2D look-up-tables (LUTs) of luminance and spatial frequency, with a separate LUT for each channel (see Fig. 4). Thanks to these techniques, ColorVideoVDP has processing times comparable to the state-of-the-art metrics, while relying on a much more complex visual model. The timings can be found in the supplementary document. ColorVideoVDP is also fully differentiable, which lets us use it as an objective function in optimization, or calibrate its parameters, as explained in Sec. 5.

4 XR DISPLAY ARTIFACT VIDEO DATASET (XR-DAVID)

Calibrating a metric to real subjective data is an extremely important step. Notably, the datasets used for calibration must be representative of the types of artifacts that users will be applying the metric to. This can be challenging in cases where there is little prior art, for instance, while traditional artifacts like image and video compression are relatively well-studied, artifacts stemming from novel XR display architectures like waveguide nonuniformity are not.

To address this gap in quantitative data for XR distortions, we enumerated the most relevant artifacts for our study. We placed special emphasis on distortions that have a color or temporal component, as these are not well served by existing metrics and are the focus of ColorVideoVDP. We used this dataset of distortions to conduct a large-scale subjective study, collecting JOD quality scores in a new XR-DAVID video quality dataset⁴, which is suitable for calibration of ColorVideoVDP.

4.1 Experimental setup

Physical Setting. We selected an Eizo CG3146 professional reference monitor⁵ as our test vehicle. This 31.1" diagonal display has a resolution of 4096×2160 , and a contrast of 1,000,000:1 claimed by the manufacturer. The monitor was set to a maximum luminance

⁴The XR-DAVID dataset can be downloaded from <https://doi.org/10.17863/CAM.108210>

⁵Eizo CG3146, see <https://www.eizo.com/products/coloredge/cg3146/> for detail

of 300 cd/m², with sRGB EOTF, P3 color primaries, and a 60 Hz refresh rate. This display has a built-in colorimeter, which allowed for high-precision daily calibration to ensure accuracy as the study progressed. The distance of the observer was controlled using a chin rest, which was placed so that the effective resolution of the display was 77 pixels per visual degree (see Fig. 12).

Participants. We conducted the study using paid external participants. After two pilots (N=5 each), 77 naïve users took part in 1 hour long sessions of our main study, which included training, data gathering, and a short break half-way through. All participants signed informed consent forms, and the study was approved by an external ethics committee. Participants were screened via an Ishihara color test [Hardy et al. 1945] to ensure normal color vision.

Reference videos. We selected 14 high-quality videos following practical considerations. Selected thumbnails of scenes used in the study are shown in the supplementary document. Our references included videos spanning real, rendered, and productivity content, which are typical for AR/VR applications.

Experimental procedure. Participants performed a side-by-side pairwise comparison task answering which of the two versions of a video was less distorted. This method was preferred to alternatives (e.g., direct ratings on the mean-opinion-score scale) as pairwise comparisons have been shown to produce more accurate results [Zerman et al. 2018] by simplifying the task in each trial. Active sampling using the *ASAP* method [Mikhailiuk et al. 2021] was used to optimize information gain from each comparison, making it possible to explore a large number of artifacts without prohibitively increasing the number of pairwise comparisons. We allowed comparisons across different distortion types and levels but not across different video contents. A training session preceded the main experiment. In that session, the participants compared each reference video to its distorted version. A different distortion type was used in each trial. The training session was meant to familiarize the participants with the reference content and distortions. The results of the main experiment were scaled to a unified perceptual scale in just-objectionable-difference units (JODs) using the *pwcmp* software suite [Perez-Ortiz et al. 2019].

4.2 Distortions

Each reference video was distorted by one of 8 artifacts, each of which was applied at one of 3 different strengths. All 8 distortions at level 3 (most distorted) are illustrated in Figure 25. We discuss each artifact in detail below.

Spatiotemporal dithering. Display systems could be constrained such that bit depths of color channels are reduced. In such scenarios, dithering could be performed both spatially and temporally to improve image quality. This artifact simulates blue noise mask dithering per color channel and per frame for a given bit depth. The bit depths were 6, 5, and 4 for artifact levels 1, 2, and 3, respectively.

Light source nonuniformity (LSNU). Display light sources like μ OLED and μ LED tend to exhibit high-frequency spatial nonuniformity, as each pixel consists of a separate light source module, which may vary in its light output. This artifact was simulated by




Fig. 13. MTF Curves corresponding to the *Blur* artifact at level 1, 2, and 3.

assuming different levels of variations in pixel intensities for each color channel, resulting in spatial and color artifacts. Each pixel intensity was randomly modulated per color channel by up to 12%, 16%, and 18% (values corresponding to twice the standard deviation, and half the limit at which truncation occurred) using a Gaussian distribution for artifact levels 1, 2, and 3, respectively. The modulation ratios were kept constant for each pixel across all frames. Simulations were done in linear color space.

Blur (MTF degradation). Optical components such as lenses and waveguides in AR and VR displays will degrade the system MTF, resulting in blurry images. The artifact was simulated by applying a representative point spread function (PSF) with various severities as illustrated in Fig. 13. The same PSF was applied to all color channels. In addition to varying PSFs, the lateral chromatic aberration was simulated for artifact level 2 and 3 by shrinking the green channel frame by two pixels.

Reduced contrast. It is often challenging to achieve good contrast in optical see-through displays, particularly in bright environments. As a first-order approximation, the artifact was simulated in linear space such that contrast was reduced by increasing the minimum value of input videos while maintaining the maximum value as is. The contrast, defined as max/min, was reduced to 100:1, 50:1, and 30:1 for artifact levels 1, 2, and 3 respectively.

Waveguide nonuniformity (WGNU). AR displays with diffractive waveguides such as Microsoft HoloLens, Magic Leap, and Snap Spectacles [Ooi and Dingliana 2022] exhibit a characteristic spatially-varying color nonuniformity. This nonuniformity is typically low-frequency (less than 1 cycle per degree). In addition, it is heavily dependent on pupil position within the eye box, which can vary depending on the user’s fixation and eye movements. In particular, if a user keeps their gaze fixed on an object in the AR content while moving their head (engaged in what is typically termed a vestibulo-ocular reflex or VOR movement), they may notice its color shifting as the position of the eye changes in relation to the waveguide. In order to simulate the artifact, first, two nonuniformity patterns were obtained from empirical waveguide transmission data. Second, base variation maps were generated for the two patterns by cleaning noise in data and normalizing variations across channels. The variations were also scaled such that they appear to have similar magnitude across the two patterns. Finally, the amplitude of variations in the base map was varied to produce different artifact levels. The multipliers of the amplitude variations were 0.15, 0.3, and 0.4




Fig. 14. Examples of waveguide nonuniformity level 3 at different frames, using *Business* scene (cropped for better visibility). nonuniformity changes from one pattern (left) to another (right).

for artifact levels 1, 2, and 3 respectively. To simulate changes with eye position, the nonuniformity pattern transitions from the first base pattern to the second and back over the course of the video. This transition is repeated 4 times for level 2, and 7 times for level 3. Fig. 14 illustrates this temporally varying waveguide nonuniformity artifact.

Dynamic correction error (DCE). In AR displays suffering from waveguide nonuniformity across different exit pupil positions, as described above, a dynamic correction algorithm can be implemented to invert the distortion, reducing spatial nonuniformity. This algorithm would be dependent on gaze location, thus relying on estimated positions from eye trackers, which may be inaccurate. Namely, inaccurate eye tracker readings will cause a wrong estimate of the pupil position and temporal color artifacts (by either not correcting perfectly for distortions or performing an imprecise correction for a distortion pattern that is not present). To simulate this, the precision error was randomly generated per frame from a Gaussian distribution with a standard deviation of 0.27° , 0.39° , and 0.58° for artifact levels 1, 2, and 3 respectively. Simulated precision errors are in alignment with existing VR headsets currently in the market [Schuetz and Fiehler 2022].

Color fringes. An image’s color channels may be optically misaligned due to various reasons such as mechanical shifts of optical components in head-mounted displays, imperfect fabrication processes, and thermal loads. As a first-order approximation, the artifact was simulated by shifting each color channel by variable amounts in 2D dimensions globally. For artifact level 1, the red channel frame was shifted by $[0, 0.5]$ pixels in $[x,y]$ direction while the blue channel frame was shifted by $[-0.5, 0]$ pixels. For artifact level 2, the red channel frame was shifted by $[1, 0.5]$ pixels while the blue channel frame was shifted by $[0.5, -1]$ pixels. For artifact level 3, the red channel frame was shifted by $[1.5, 1]$ pixels while the blue channel frame was shifted by $[1, -1]$ pixels. The green channel frame was kept unshifted for all the artifact levels.

Chroma subsampling. Chroma subsampling is a common technique used in video compression and signal transmission to reduce spatial resolutions of chroma channels while maintaining the overall image quality. The technique takes advantage of a much-reduced sensitivity of the visual system to high frequencies modulated along chromatic (isoluminant) color directions. Spatial resolutions of chroma channels were reduced by $1/4$, $1/8$, and $1/12$ for artifact level 1, 2, and 3 respectively.




Fig. 15. The results of our display artifact study aggregated over all users and scenes are shown. Quality scores are expressed in JODs, counting down from 10 by convention. Values above the nominal reference are the result of chance selection for artifacts that appear very subtle. Error bars represent 95% confidence intervals, obtained via bootstrapping.

Results. The results of our study are shown in Fig. 15 in aggregate form over all scenes, and in full in Figure 2 in the supplementary. Fig. 15 shows a good range of quality levels, with some distorted videos almost indistinguishable from the reference. Such conditions are useful to test whether a metric can disregard invisible distortions. Having all artifacts graded on a single linear perceptual JOD scale allows for direct comparison, and reveals interesting interactions between individual videos and artifacts types. For instance, while blur was graded as very disturbing across all forms of content, waveguide nonuniformity was fairly visible for a web-browsing scenario (*‘wiki’*), but almost invisible in a dark scene (*‘bonfire’*) (see Figure 14 in the supplementary).

5 METRIC VALIDATION

A typical validation of a quality metric involves reporting correlation values for each individual dataset. While we still report such correlation values (in the supplementary HTML reports), here, we undertake a more challenging task of predicting absolute quality in JOD units, which could generalize across datasets scaled in such units. Therefore, unlike most work in this area, we do not train ColorVideoVDP individually for each dataset, but instead, we train a single version of the metric on multiple datasets with the goal of generalizing to new (unseen) data.

Datasets. We split the datasets into those used solely for testing and those used for both training and testing. The five selected datasets are listed in Table 2. The two datasets used for both training and testing are XR-DAVID, explained in detail in Sec. 4, and UPIQ [Mikhailiuk et al. 2022]. We selected the UPIQ dataset because it contains a large collection of both SDR and HDR images (over 4000) and is scaled in JOD units, similar to XR-DAVID. To test our metric on unseen datasets (cross-dataset validation), we chose LIVE HDR

Table 2. Datasets used for training and testing.

Dataset	Used for	Type	Scenes	Conditions	Distortions
XR-DAVID	Train & test	SDR video	14	336	8 display artifacts
UPIQ [Mikhailiuk et al. 2022]	Train & test	SDR/HDR images	84	4159	34 distortion types
LIVE HDR [Shang et al. 2022]	Test	HDR video	21	210	H.265, bicubic upscaling
LIVE VQA [Seshadrinathan et al. 2010]	Test	HDR video	10	150	H.264, MPEG-2, transmission
KADID-10k [Lin et al. 2019]	Test	SDR images	81	10125	25 distortion types

because it is a modern dataset representative of video streaming applications, KADID-10k because of its size (over 10k images), and LIVE VQA because it is widely used to test video quality metrics. Because KADID-10k and a part of UPIQ were collected in uncontrolled crowdsourcing experiments, we had to approximate the display specification and viewing conditions for those. We avoided video datasets that were collected in uncontrolled conditions (e.g., KonViD-1k).

Training and testing sets. The two datasets used for both training and testing were split into 7 parts: 5 parts were used for training, and 2 parts were used for testing. Each scene is present in only one part so that no scene is shared across the sets. 7 parts were selected because 7 is the common denominator for the number of unique scenes in XR-DAVID and LIVE HDR (see Table 2).

Training. Training a video quality metric is problematic as a very large amount of data is used to predict a single quality value. For example, 60 frames of 4K video requires the processing of almost 500 million pixels to infer just a single quality score. This makes gradient computation (through backpropagation) infeasible because of the memory requirements. Previous work dealt with this problem in several ways: some metrics were designed to extract low dimensional features from a video and then train a regression mapping those features to quality scores [Li et al. 2016b]. This approach, however, does not allow training or tuning the feature extraction stage. Another group of methods operated on patches of limited resolution (e.g. 64×64), assuming that the quality is the same across the entire image [Prashnani et al. 2018]. This assumption, however, can be easily proved wrong for localized distortions or for content in which the effect of contrast masking varies across an image. Other metrics used numerical gradient computation [Mantiuk et al. 2021], which, however, becomes too expensive when optimizing a large number of parameters.

We used a mixture of feature-based and end-to-end training. The parameters that are introduced after pooling across all pixels in a frame (such as JOD regression parameters) can be easily optimized by pre-computing pooled values/features ($\|\hat{D}_{b,c,f}(x)\|_{\beta_x, x}$ from Eq. (14)) and then optimizing the stages of the metric that follow the pooling stage. This approach not only significantly accelerates the training process due to the reduced memory footprint of pooled features compared to full videos, but it also lets us operate on much larger batches. We found that large batches, with smoother gradients, are required for stable training of the pooling and regression parameters.

Table 3. The trained parameters of the ColorVideoVDP. The parameters of the cross-channel masking can be found in Fig. 9.

Model component	Parameters
Contrast sensitivity	$S_{\text{corr}} = 0.9683$
Masking	$p = 2.264, q_c = [1.303, 2.889, 3.681, 3.589],$ $k_C = 366.6, \sigma_{\text{sp}} = 3,$ $k_{B,c} = [0.003633, 1.663, 4.119, 25.26]$
Pooling	$k_I = 0.5779, w_c = [1, 1, 1, 0.8081]$
JOD regression	$\alpha_{\text{JOD}} = 0.04396, \beta_{\text{JOD}} = 0.9302$

For the end-to-end training, we computed the full analytical gradient of the remaining parameters using two techniques. Firstly, we utilized gradient checkpointing [Chen et al. 2016] in PyTorch, by rerunning a forward pass for each checkpointed segment during backward propagation. This technique trades some speed for reduced memory requirements. We inserted a checkpoint after each block of frames, where the block size was determined based on the available GPU memory. Secondly, during training, we randomly sampled 0.5-second-long sequences from each video clip (full sequences were used for testing). This approach improved training time and introduced a form of data augmentation. The feature-based and end-to-end training were run in an interleaved manner, with 50 epochs of (fast) feature-based training followed by a single epoch of end-to-end training. This lets us jointly train all the parameters of the metric. The trained parameters of the metric can be found in Table 3.

5.1 Comparison with other metrics

We compare the performance of ColorVideoVDP with several state-of-the-art metrics, listed in Table 1.

As our datasets include both SDR and HDR content, we need to ensure that those are handled correctly by all the metrics. For metrics that do not work with colorimetric data and instead operate on display-encoded (SDR) pixel values, we ran the evaluations using the original SDR pixel values. In the case of HDR content, we employed the perceptually-uniform transform (PU21) [Mantiuk and Azimi 2021] to encode the pixel values. The metrics that operate on colorimetric data (FovVideoVDP, HDR-VDP-3) were supplied the absolute luminance values, computed by the display model (Sec. 3.1). The metrics that account for the viewing distance or the resolution in pixels-per-degree (e.g., HDR-VDP-3, HDR-FLIP) were supplied with the correct values.

Because UPIQ, XR-DAVID and LIVEHDR are scaled in the same JOD units, we fit a single logistic function to map metric prediction to JODs. As neither KADID-10k nor LIVEVQA is scaled in JOD units, we fit a logistic function separately for each metric and dataset to map metric prediction to the subjective scores.

The results, shown in Fig. 16, indicate a substantial gain in performance of ColorVideoVDP over the second-best metric – VMAF. Image metrics that consider color – FSIMc and VSI – performed better than expected on XR-DAVID, but rather poorly on LIVEHDR video dataset. The color difference formulas and their spatial extensions performed worse than image and quality metrics. The non-reference metrics – NIQE, PIQE, and BRISQUE – show a very small correlation with subjective judgments. Other performance indices (SROCC and PLCC) and detailed results can be found in the HTML reports included in the supplementary. Overall, while it is possible to find a metric that performs well on a selected dataset, such as STRRED on LIVEVQA (LIVEVQA was used to calibrate and test STRRED), ColorVideoVDP offers good performance across all datasets with a wide variety of distortions and content. This could be explained by no other metric offering the same set of capabilities as ColorVideoVDP; only our metric models spatiotemporal color vision and accounts for the display model.

5.2 Ablations

We grouped ablation studies into those used to determine suitable contrast encoding and masking models, and into those used to test the importance of each component of the metric. The metric parameters were refitted for each ablation, as explained above. First, we tested a combination of two contrast encodings and three masking functions, all explained in detail in the supplementary document. The results of those ablations, shown in Fig. 17, clearly indicate that the mutual masking model with multiplicative contrast encoding offers much better performance than the alternatives. The results indicate that the multiplicative contrast encoding is necessary for mutual masking and transducer models. As mentioned in Sec. 3.6, such encoding can unify results across different spatial frequencies [Daly 1993] and color directions [Cass et al. 2009; Switkes et al. 1988]. Although the contrast transducer is one of the best-established masking models [Watson and Solomon 1997], which accounts for facilitation and performs well on selected datasets [Alam et al. 2014], it did not perform well as a part of the quality metric. The similarity formula exhibits masking properties and is used in many metrics, such as (MS-)SSIM [Wang et al. 2003], but it did not result in acceptable performance when integrated into our metric.

Second, we tested the importance of each component of our metric, namely:

- w/o temporal channel – ignored the contribution of the achromatic transient channel;
- w/o cross-channel masking – masking was allowed within each channel, but not across the channels (matrix $k_{i,c}$ from Eq. (11) was diagonal).
- w/o chromatic channels – ignored the contribution of the two chromatic channels;




Fig. 16. ColorVideoVDP compared with existing metrics, listed in Table 1. The prediction errors are reported for testing datasets (LIVEVQA, LIVEHDR and KADID) and the test portion of UPIQ and XR-DAVID. The error bars denote 95% confidence intervals. Other performance metrics (SROCC and PLCC) and detailed reports can be found in the supplementary HTML report.

- w/o masking model – disabled the masking model but used the contrast encoding from Eq. (7) and contrast saturation from Eq. (13);
- w/o CSF – the contrast sensitivity function did not vary with luminance or spatial frequency, but varied across the channels;
- w/o contrast saturation – disabled the contrast saturation formula from Eq. (13).

The results of these ablations, shown in Fig. 18, demonstrate that the contrast saturation formula is critical for the mutual masking model used in our metric. The metric also performs poorly if the contrast encoding is not modulated by the CSF (see “w/o CSF” in Fig. 18),

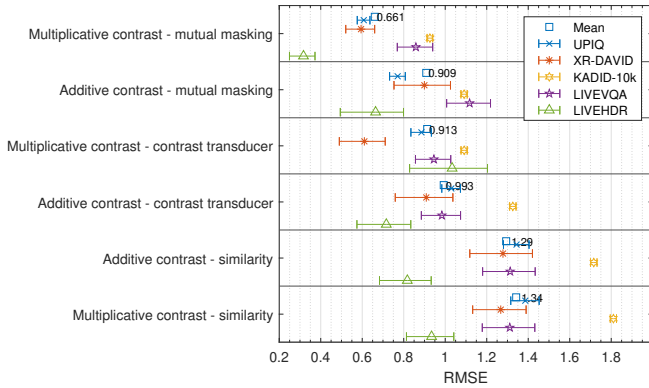


Fig. 17. Ablation studies on the contrast encoding and masking models used in ColorVideoVDP. The notation is the same as in Fig. 16. The masking models used in the ablations are explained in the supplementary.

or lacks the masking model. Although color is often regarded as a less critical aspect of quality assessment (especially in the context of video coding), here we show that it is important for the datasets that contain color distortions, such as XR-DAVID, UPIQ, or KADID-10k. The cross-channel masking is an important but subtle effect, which results in gains mostly for the XR-DAVID dataset. While the influence of the temporal channel may seem small in the results (see “w/o temporal channel” in Fig. 18), this is due to the gain only being observable for video datasets with temporal distortions, such as XR-DAVID.

Finally, we tested the importance of the XR-DAVID dataset when training ColorVideoVDP. We retrained ColorVideoVDP, but used UPIQ and LIVE-HDR as training sets (instead of UPIQ and XR-DAVID). The quality scores of LIVE-HDR were scaled in JOD units, as explained in the supplementary. The results of this training, shown as “Base model trained on LIVE-HDR” in Fig. 18, demonstrate that the XR-DAVID dataset is essential for training color video metrics. Video compression datasets, such as LIVEHDR, lack the variety of both temporal and color distortions, which is a necessary component for the calibration of video metrics. The correlation coefficients for all ablations and detailed reports can be found in the supplementary HTML report.

5.3 Synthetic tests

Validation datasets and ablations may not capture the edge cases for which a metric may perform differently than expected. To examine these, we created 14 sets of synthetic test and reference pairs, containing contrast, masking, flicker patterns, and typical distortions. One such example, shown in Fig. 19, demonstrates the metric’s ability to estimate the magnitude of supra-threshold contrast correctly. The lines shown in the plots connect contrast across three color directions that match in perceived magnitude, according to the data of Switkes and Crognale [1999]. The figure shows that while ColorVideoVDP can correctly predict suprathreshold contrast, HDR-FLIP [Andersson et al. 2021] overpredicts the contrast in the red-green and yellow-violet directions. The extensive report for all

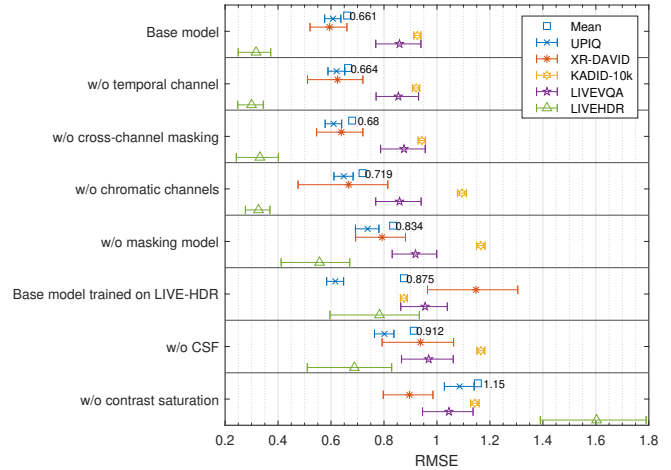


Fig. 18. Ablation studies on ColorVideoVDP. The top row shows the base model. Refer to the text for details.

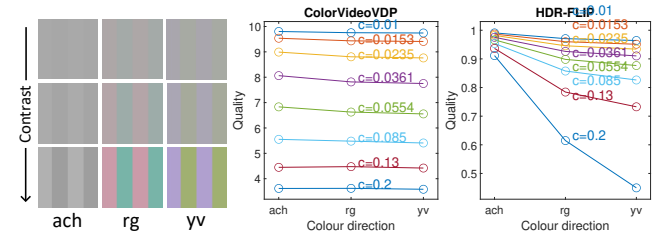


Fig. 19. Testing of supra-threshold contrast matching across color directions. The lines in the plots represent metric responses for the contrast modulated along one of the cardinal directions of the DKL color space (examples shown on the left). The magnitude of the contrast was selected to match the data of Switkes and Crognale [1999]. If a metric correctly predicts the magnitude of chromatic and achromatic contrast, the plots should form horizontal lines. The values for FLIP have been shown as 1-(FLIP mean difference) to make them comparable with other metrics. Note that the perceived magnitudes depend on the spatial frequency and, therefore, the viewing distance.

14 sets and multiple metrics can be found in the supplementary HTML report.

6 APPLICATIONS

As a general-purpose image and video difference metric, our metric can be used for a range of standard applications, such as optimization of video streaming. In this section, we present three proof-of-concept use cases, which go beyond such standard applications of video metrics.

6.1 Chroma subsampling

Chroma subsampling is a popular compression technique in which chroma channels are encoded with a lower resolution than luma. This works well in practice due to the lower sensitivity of our chromatic vision to high frequencies. The visibility of chroma subsampling artifacts cannot be easily predicted with pixel-wise color difference formulas, such as CIEDE2000, as they are unaware of image

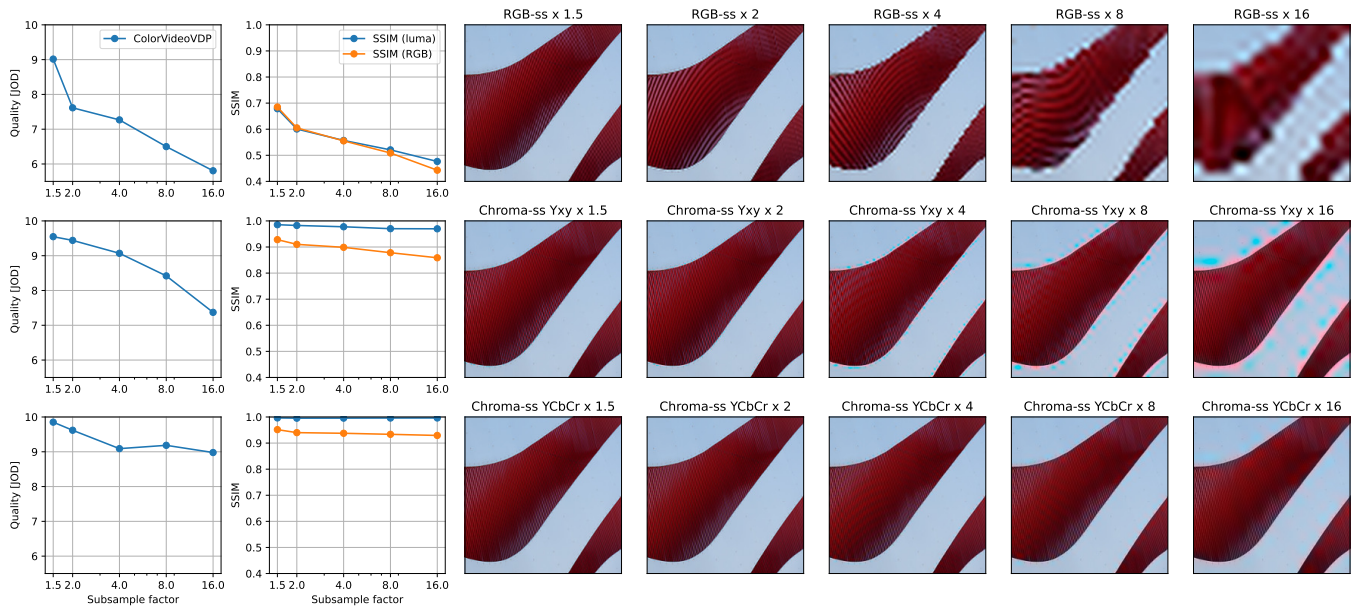


Fig. 20. The comparison of quality loss due to subsampling of RGB channels (top row), chroma channels of Yxy colorspace (middle row), and chroma channels of YCbCr color space (bottom row). The plots show the predictions of ColorVideoVDP (1st column) and SSIM (2nd column), the latter computed either on luma or RGB channels. SSIM fails to predict substantial quality loss when the chroma channels are severely subsampled.

structure. Spatial metrics, such as SSIM or even FSIMc, do not operate on color spaces that could properly isolate chromatic and achromatic mechanisms. This is shown in an example in Fig. 20, in which SSIM fails to predict a substantial loss of quality at high chromatic subsampling rates, even if the metric is computed on the RGB channels (the original SSIM operates only on luma). The SSIM (RGB) predictions for $\times 16$ subsampling of xy channels (last column, second row) indicates only a moderate loss of quality (0.87), much lower than even $\times 1.5$ RGB subsampling (third column, top row, 0.68), which poorly correlates with the perceived level of distortions. ColorVideoVDP provides easy-to-interpret predictions, which correspond well with the perceived image quality. It shows, for example, the strength of the YCbCr space, which can well balance subsampling distortions between the achromatic and chromatic channels (the chromatic plane of YCbCr is not isoluminant).

6.2 Optimization

ColorVideoVDP is fully differentiable, and it can be potentially used as an image loss function when training deep neural networks. However, a good quality metric may not serve as a good loss function [Mustafa et al. 2022], in particular when the features used by a metric are not surjective [Ding et al. 2021]. Moreover, complex functions can make the landscape of the loss highly non-convex and impede convergence. For that reason, a popular "perceptual loss" [Johnson et al. 2016] must be combined with L1 or L2 to ensure convergence.

We found that when ColorVideoVDP is used as an image loss function (without L1 or L2 term) in an image reconstruction task (see [Ding et al. 2021, Sec. 3]), it converges well when the reconstructed image is close to the reference. However, the convergence is worse when the test image is affected by random noise or is far from the

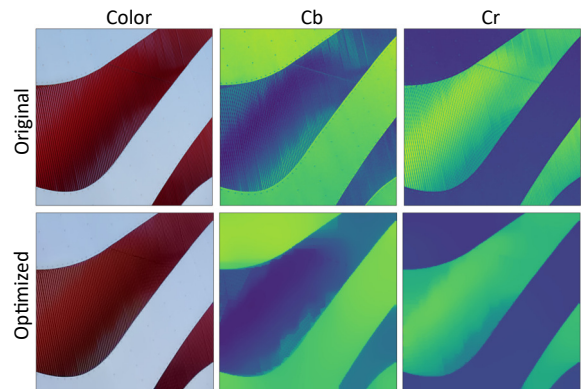


Fig. 21. Adaptive chroma subsampling removes high frequencies from chromatic Cb and Cr planes (2nd and 3rd column) of the YCbCr color space without affecting the appearance of the color image (1st column). Such adaptive subsampling can effectively improve image compression without the loss of chromatic details.

reference. This problem is alleviated when ColorVideoVDP is used when optimizing low-dimensional problems, such as the selection of encoding parameters in video streaming. In the next section, we show an example of ColorVideoVDP used as a loss function in a high-dimensional problem.

6.3 Adaptive chroma subsampling

Traditional chroma subsampling, as demonstrated in Sec. 6.1, globally reduces the resolution (sampling rate) of the chromatic planes

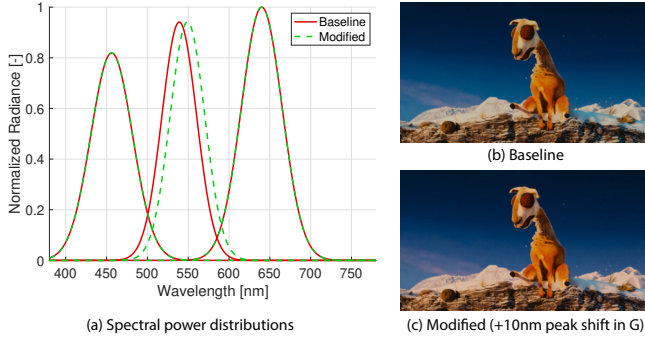


Fig. 22. (a) Spectral power distributions for baseline (red solid line) and modified curves (green dashed line), (b) Proof image shown on display with baseline primaries, (c) Proof image shown on display with modified primaries. © caminandes.com

regardless of image content. However, the visibility of chromatic distortion will strongly depend on image content, and therefore, we argue that chroma subsampling should be adaptive and vary across an image. Here, we show that we can use ColorVideoVDP as a differential loss to adaptively remove high frequencies from the image chroma channels while accounting for all factors that contribute to the visibility of the introduced changes. To do this, we optimize for the color image I :

$$\operatorname{argmin}_I \|I - I_{\text{org}}\|_{\text{cvvdp}} + \lambda \|\nabla I_{\text{Cb}}\|_1 + \lambda \|\nabla I_{\text{Cr}}\|_1 \quad (19)$$

where I_{org} is the original image, ∇I_{Cb} and ∇I_{Cr} are the gradients of the Cb and Cr planes of the YCbCr color space and λ is a regularization constant. $\|\cdot\|_{\text{cvvdp}}$ denotes ColorVideoVDP used as a differentiable loss and $\|\cdot\|_1$ is an L1 norm. Fig. 21 shows an example of such adaptive chroma subsampling. In this particular example, we could reduce the size of a PNG file by 10% with no impact on image quality.

6.4 Setting display color tolerance specifications

In a display manufacturing setting, it is common for display primaries' spectra to differ from the desired target values between individual units and suppliers. Manufacturers typically set specifications for how much each primary can deviate from the ideal target. Traditionally, it has been characterized in wavelength shifts, changes in chromaticity and luminance, or color difference metrics such as ΔE_{ab}^* and CIEDE2000. ColorVideoVDP can be used to set these specifications in an interpretable manner (in terms of JODs) with respect to sample content, as our metric is aware of image structure.

As an example use case, assume a VR display is being characterized in a factory calibration setting. We simulate the reference by generating a synthetic set of R, G, and B primaries producing a P3 gamut using Gaussian curves, as depicted by the red solid line in Fig. 22(a). Next, we simulate possible spectral primary deviations by varying the spectral peak and full width at half maximum (FWHM) for each primary to examine tolerances (green dashed line in Fig. 22(a)). A test image rendered with both baseline and distorted

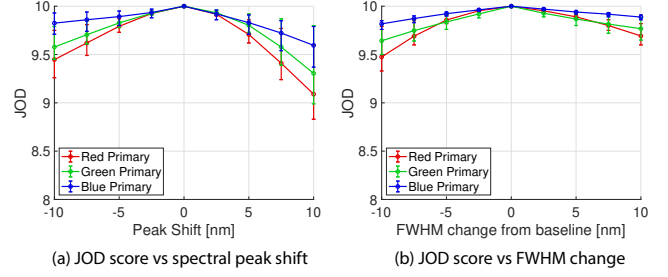


Fig. 23. (a) JOD scores vs spectral peak shift for R, G, B primaries. (b) JOD scores vs. FWHM change from baseline for R, G, B primaries. Error bars show maximum and minimum JOD scores among 4 scenes for each case.

primaries is shown in Fig. 22 (b) and (c), respectively. Finally, JODs scores were calculated between images with baseline and modified primaries for each spectral peak shift and FWHM change for each of four test images (taken from *Caminandes*, *Icons*, *Panel*, and *VR*, see the supplementary document).

The results are shown in Fig. 23. Note that JOD scores change nontrivially depending on the direction of the peak shift, FWHM magnitude, and type of primary (R, G, or B). In our imaginary example, factory tolerance specifications could then be derived by setting an acceptable limit for Δ JOD deviation from the baseline via a psychophysical or "golden eye" study. Further investigation can be conducted by looking into individual images for worst performers, and future tolerances can be easily tightened by reducing the Δ JOD specification.

6.5 Observer metamerism and variability

Observer metamerism has been an issue for wide-color-gamut displays. As wide color gamut displays typically have spectrally narrower primaries, it is increasingly likely that, when calibrated for a standard observer, they will appear color-inaccurate for individuals who deviate from this profile [Bodner et al. 2018; Hung 2019]. The severity of this observer metamerism depends on both the content being shown and the display characteristics, and can be evaluated by using ColorVideoVDP.

To demonstrate this, we calculated JOD scores for the 'Wiki' scene, simulating three different display primary spectral profiles and 11 different observer functions. The evaluated displays were Sony BVM32 CRT (100% of sRGB), Eizo ColorEdge 3145 (99% P3), and laser primaries (100% ITU-R BT.2020 color gamuts). Laser spectra were synthetically generated. Spectral power distributions for all primaries are shown in Fig. 24 (top). To simulate individual differences in color vision, we used Asano and Fairchild's 10 categorical observers [2020], as well as the 1931 2° standard observer used for reference. These 10-observer functions were derived as representative means for a color-normal population based on an individual colorimetric observer model [Asano et al. 2016]. The results are shown in Fig. 24 (bottom). As expected, the narrower spectra incur a larger metamerism error for non-standard observers, which is reflected in the higher Δ JOD values. ColorVideoVDP can be used to estimate the risk of metamerism error for populations on content, and to evaluate or optimize solutions for this problem.

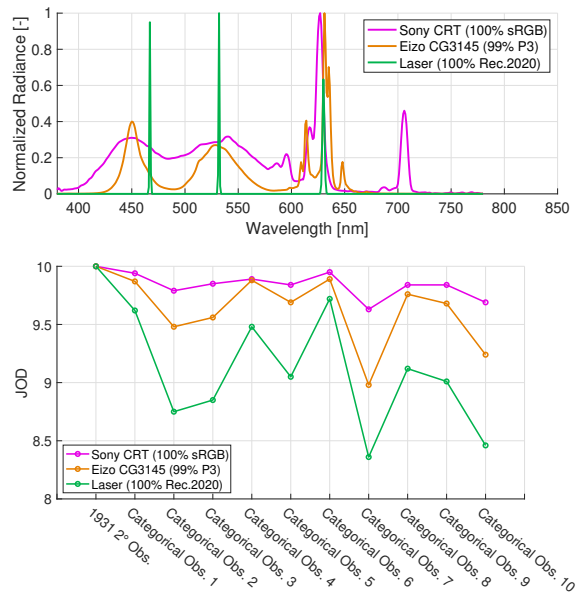


Fig. 24. (Top) Spectral power distributions for three types of displays used to simulate observer metamerism (see Sec. 6.5). (Bottom) JOD scores for all observers relative to 1931 2° standard observer for the ‘Wiki’ scene, for three display types.

7 CONCLUSIONS

In this work we introduce ColorVideoVDP, a general-purpose image and video metric that models several challenging aspects of vision simultaneously. Notably, our algorithm is calibrated in psychophysical JOD units, models color, high-dynamic-range, and spatiotemporal aspects of vision. Our metric is explainable, and each component used to build our pipeline is based on proven psychophysical models of the human visual system.

An important aspect of our metric is its extensive calibration. Our metric is simultaneously calibrated on 3 large datasets containing a variety of distortion types, which included our novel XR-DAVID video quality dataset, containing a range of display hardware-oriented artifacts critically important to the development of future displays. We demonstrated that our metric generalizes well to unseen datasets.

Finally, ColorVideoVDP is efficiently implemented to run on a GPU. This aspect is increasingly important in the modern display landscape, as resolution, frame rate, field of view, and bit-depth continue to increase, adding complexity to analyzed content. It is also fully differentiable, making it possible to use it as a loss function in optimization problems.

ACKNOWLEDGMENTS

We thank Dounia Hammou for her insightful comments. We thank Cameron Wood, Eloise Moore, Helen Ayele, Sydney Gregory and Ka Yan Wat for help running the user study. Thanks go to Curtis Torrey and Xin Li for assistance with experimental hardware, and to Jacqueline Beyrouthy for help with logistics.

REFERENCES

- M. M. Alam, K. P. Vilankar, David J Field, and Damon M Chandler. 2014. Local masking in natural images: A database and analysis. *Journal of Vision* 14, 8 (July 2014), 22–22. <https://doi.org/10.1167/14.8.22> Citation Key: Alam2014.
- Stephen J. Anderson and David C. Burr. 1985. Spatial and temporal selectivity of the human motion detection system. *Vision Research* 25, 8 (jan 1985), 1147–1154. [https://doi.org/10.1016/0042-6989\(85\)90104-X](https://doi.org/10.1016/0042-6989(85)90104-X)
- Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proc. of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (aug 2020), 1–23. <https://doi.org/10.1145/3406183>
- Pontus Andersson, Jim Nilsson, Peter Shirley, and Tomas Akenine-Möller. 2021. Visualizing Errors in Rendered High Dynamic Range Images. In *Eurographics Short Papers*. <https://doi.org/10.2312/egs.20211015>
- Yuta Asano and Mark D Fairchild. 2020. Categorical observers for metamerism. *Color Research & Application* 45, 4 (2020), 576–585.
- Yuta Asano, Mark D Fairchild, and Laurent Blondé. 2016. Individual colorimetric observer model. *PloS one* 11, 2 (2016), e0145671.
- Maliha Ashraf, Rafał K. Mantiuk, Alexandre Chapiro, and Sophie Wuerger. 2024. castleCSF — A Contrast Sensitivity Function of Color, Area, Spatio-Temporal frequency, Luminance and Eccentricity. *Journal of Vision* 24 (2024), 5. <https://doi.org/10.1167/jov.24.4.5>
- Ben Bodner, Neil Robinson, Robin Atkins, and Scott Daly. 2018. 78-1: Correcting Metameric Failure of Wide Color Gamut Displays. In *SID Symposium Digest of Technical Papers*, Vol. 49. Wiley Online Library, 1040–1043.
- P. Burt and E. Adelson. 1983. The Laplacian Pyramid as a Compact Image Code. *IEEE Transactions on Communications* 31, 4 (apr 1983), 532–540. <https://doi.org/10.1109/TCOM.1983.1095851>
- John Cass, C. W. G. Clifford, David Alais, and Branka Spehar. 2009. Temporal structure of chromatic channels revealed through masking. *Journal of Vision* 9, 5 (may 2009), 17–17. <https://doi.org/10.1167/9.5.17>
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174* (2016).
- Manri Cheon, Sung-Jun Yoon, Byungyeon Kang, and Junwoo Lee. 2021. Perceptual Image Quality Assessment with Transformers. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, Nashville, TN, USA, 433–442. <https://doi.org/10.1109/CVPRW53098.2021.00054>
- Anustup Choudhury, Robert Wanat, Jaclyn Pytlarz, and Scott Daly. 2021. Image quality evaluation for high dynamic range and wide color gamut applications using visual spatial processing of color differences. *Color Research & Application* 46, 1 (feb 2021), 46–64. <https://doi.org/10.1002/col.22588>
- CIE. 1993. *Parametric effects in colour-difference evaluation*. Technical Report. CIE 101-1993.
- CIE. 2018. CIE 015: 2018 Colorimetry. (2018).
- S.J. Daly. 1993. Visible differences predictor: an algorithm for the assessment of image fidelity. In *Digital Images and Human Vision*, Andrew B. Watson (Ed.). Vol. 1666. MIT Press, 179–206. <https://doi.org/10.1117/12.135952>
- R.L. De Valois, D.G. Albrecht, and L.G. Thorell. 1982. Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research* 22, 5 (1982), 545–559. [https://doi.org/10.1016/0042-6989\(82\)90113-4](https://doi.org/10.1016/0042-6989(82)90113-4)
- Gyorgy Denes, Akshay Jindal, Aliaksei Mikhailiuk, and Rafał K. Mantiuk. 2020. A perceptual model of motion quality for rendering with adaptive refresh-rate and resolution. *ACM Transactions on Graphics* 39, 4 (jul 2020), 133. <https://doi.org/10.1145/3386569.3392411>
- A M Derrington, J Krauskopf, and P Lennie. 1984. Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of Physiology* 357, 1 (dec 1984), 241–265. <https://doi.org/10.1113/jphysiol.1984.sp015499>
- Piotr Didyk, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, and Hans-peter Seidel. 2011. A perceptual model for disparity. *ACM Transactions on Graphics* 30, 4 (jul 2011), 1. <https://doi.org/10.1145/2010324.1964991>
- Keyan Ding, Kede Ma, Shiqi Wang, and Eero P. Simoncelli. 2021. Comparison of Full-Reference Image Quality Models for Optimization of Image Processing Systems. *International Journal of Computer Vision* 129, 4 (apr 2021), 1258–1281. <https://doi.org/10.1007/s11263-020-01419-7> arXiv:2005.01338
- John M. Foley. 1994. Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America A* 11, 6 (jun 1994), 1710. <https://doi.org/10.1364/JOSAA.11.001710>
- LeGrand H Hardy, Gertrude Rand, and M Catherine Rittler. 1945. Tests for the detection and analysis of color-blindness. I. The Ishihara test: An evaluation. *JOSA* 35, 4 (1945), 268–275.
- R.F. Hess and R.J. Snowden. 1992. Temporal properties of human visual filters: number, shapes and spatial covariation. *Vision Research* 32, 1 (jan 1992), 47–59. [https://doi.org/10.1016/0042-6989\(92\)90112-V](https://doi.org/10.1016/0042-6989(92)90112-V)
- Po-Chieh Hung. 2019. 61-3: Invited paper: CIE activities on wide colour gamut and high dynamic range imaging. In *SID Symposium Digest of Technical Papers*, Vol. 50. Wiley Online Library, 866–869.

- ITU-R BT. 2124. 2019. *Objective metric for the assessment of the potential visibility of colour differences in television*. Technical Report.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*. Lecture Notes in Computer Science, Vol. 9906. Springer International Publishing, Cham, 694–711. https://doi.org/10.1007/978-3-319-46475-6_43
- Minjung Kim, Maryam Azimi, and Rafal K. Mantiuk. 2021. Color Threshold Functions: Application of Contrast Sensitivity Functions in Standard and High Dynamic Range Color Spaces. *Electronic Imaging* 33, 11 (jan 2021), 153–1–153–7. <https://doi.org/10.2352/ISSN.2470-1173.2021.11.HVEI-153>
- Justin Laird, Mitchell Rosen, Jeff Pelz, Ethan Montag, and Scott Daly. 2006. Spatio-velocity CSF as a function of retinal velocity using unstabilized stimuli. In *SPIE 6057, Human Vision and Electronic Imaging XI*.
- Valero Laparra, Johannes Ballé, Alexander Berardino, and Eero P. Simoncelli. 2016. Perceptual image quality assessment using a normalized Laplacian pyramid. *Electronic Imaging* 28, 16 (2016), 1–6. <https://doi.org/10.2352/ISSN.2470-1173.2016.16.HVEI-103>
- Zhi Li, Anne Aaron, Ioannis Katsavounidis, Anush Moorthy, and Megha Manohara. 2016a. *Toward A Practical Perceptual Video Quality Metric*. Technical Report. The NETFLIX Tech Blog. <https://netflixtechblog.com/toward-a-practical-perceptual-video-quality-metric-653f208b9652>
- Zhi Li, Anne Aaron, Ioannis Katsavounidis, Anush Moorthy, and Megha Manohara. 2016b. *Toward A Practical Perceptual Video Quality Metric*. <https://netflixtechblog.com/toward-a-practical-perceptual-video-quality-metric-653f208b9652>
- Hanhe Lin, Vlad Hosu, and Dietmar Saupe. 2019. KADID-10k: A Large-scale Artificially Distorted IQA Database. In *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)*, Vol. 161. IEEE, 1–3. <https://doi.org/10.1109/QoMEX.2019.8743252>
- Rafal K. Mantiuk, Maliha Ashraf, and Alexandre Chapiro. 2022. StelaCSF: A Unified Model of Contrast Sensitivity as the Function of Spatio-Temporal Frequency, Eccentricity, Luminance and Area. *ACM Trans. Graph.* 41, 4, Article 145 (Jul 2022), 16 pages. <https://doi.org/10.1145/3528223.3530115>
- Rafal K. Mantiuk and Maryam Azimi. 2021. PU21: A novel perceptually uniform encoding for adapting existing quality metrics for HDR. In *2021 Picture Coding Symposium (PCS)*. IEEE, 1–5. <https://doi.org/10.1109/PCS50896.2021.9477471>
- Rafal K. Mantiuk, Gyorgy Denes, Alexandre Chapiro, Anton Kaplanyan, Gizem Rufo, Romain Bachy, Trisha Lian, and Anjul Patney. 2021. FovVideoVDP: A visible difference predictor for wide field-of-view video. *ACM Transaction on Graphics* 40, 4 (2021), 49. <https://doi.org/10.1145/3450626.3459831>
- Rafal K. Mantiuk, Dounia Hammou, and Param Hanji. 2023. HDR-VDP-3: A multi-metric for predicting image differences, quality and contrast distortions in high dynamic range and regular content. (apr 2023). arXiv:2304.13625 <http://arxiv.org/abs/2304.13625>
- Rafal K. Mantiuk, Kil Joong Kim, Allan G. Rempel, and Wolfgang Heidrich. 2011. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics* 30, 4 (July 2011), 1–14.
- DJ McKeefry, IJ Murray, and JJ Kulikowski. 2001. Red–green and blue–yellow mechanisms are matched in sensitivity for temporal and spatial modulation. *Vision Research* 41, 2 (jan 2001), 245–255. [https://doi.org/10.1016/S0042-6989\(00\)00247-9](https://doi.org/10.1016/S0042-6989(00)00247-9)
- Aliaksei Mikhailiuk, Maria Perez-Ortiz, Dingcheng Yue, Wilson Suen, and Rafal Mantiuk. 2022. Consolidated Dataset and Metrics for High-Dynamic-Range Image Quality. *IEEE Transactions on Multimedia* 24 (2022), 2125–2138. <https://doi.org/10.1109/TMM.2021.3076298>
- Aliaksei Mikhailiuk, Clifford Wilmot, Maria Perez-Ortiz, Dingcheng Yue, and Rafal K Mantiuk. 2021. Active sampling for pairwise comparisons via approximate message passing and information gain maximization. In *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2559–2566.
- Aamir Mustafa, Aliaksei Mikhailiuk, Dan Andrei Iliescu, Varun Babbar, and Rafal K. Mantiuk. 2022. Training a Task-Specific Image Reconstruction Loss. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Waikoloa, HI, USA, 21–30. <https://doi.org/10.1109/WACV51458.2022.00010>
- Chun Wei Ooi and John Dingliana. 2022. Color LightField: Estimation Of View-point Dependant Color Dispersion In Waveguide Display. In *SIGGRAPH Asia Posters*. 1–2.
- Eli Peli. 1990. Contrast in complex images. *Journal of the Optical Society of America A* 7, 10 (oct 1990), 2032–40. <https://doi.org/10.1364/JOSAA.7.002032>
- Eli Peli. 1995. Suprathreshold contrast perception across differences in mean luminance: effects of stimulus size, dichoptic presentation, and length of adaptation. *JOSA A* 12, 5 (1995), 817. <https://doi.org/10.1364/JOSAA.12.000817>
- Maria Perez-Ortiz, Aliaksei Mikhailiuk, Emin Zerman, Vedad Hulusic, Giuseppe Valenzise, and Rafal K Mantiuk. 2019. From pairwise comparisons and rating to a unified quality scale. *IEEE Transactions on Image Processing* 29 (2019), 1139–1151.
- Ekta Prashnani, Hong Cai, Yasamin Mostofi, and Pradeep Sen. 2018. PieAPP: Perceptual Image-Error Assessment Through Pairwise Preference. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 1808–1817. <https://doi.org/10.1109/CVPR.2018.00194>
- R. Frank Quick, W. W. Mullins, and T. A. Reichert. 1978. Spatial summation effects on two-component grating thresholds. *Journal of the Optical Society of America* 68, 1 (jan 1978), 116. <https://doi.org/10.1364/JOSA.68.000116>
- Immo Schuetz and Katja Fiehler. 2022. Eye tracking in virtual reality: Vive pro eye spatial accuracy, precision, and calibration reliability. *Journal of Eye Movement Research* 15, 3 (2022).
- Kalpana Seshadrinathan, Rajiv Soundararajan, Alan Conrad Bovik, and Lawrence K. Cormack. 2010. Study of Subjective and Objective Quality Assessment of Video. *IEEE Transactions on Image Processing* 19, 6 (jun 2010), 1427–1441. <https://doi.org/10.1109/TIP.2010.2042111>
- Zaixi Shang, Joshua P. Ebenezzer, Alan C. Bovik, Yongjun Wu, Hai Wei, and Sriram Sethuraman. 2022. Subjective Assessment Of High Dynamic Range Videos Under Different Ambient Conditions. In *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE, 786–790. <https://doi.org/10.1109/ICIP46576.2022.9897940>
- Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1 (feb 2005), 21–30. <https://doi.org/10.1002/col.20070>
- Rajiv Soundararajan and Alan C Bovik. 2012. Video quality assessment by reduced reference spatio-temporal entropic differencing. *IEEE Transactions on Circuits and Systems for Video Technology* 23, 4 (2012), 684–694.
- Andrew Stockman and David H. Brainard. 2010. Color vision mechanisms. In *OSA handbook of optics*. 11.
- C. F. Stromeyer and B. Julesz. 1972. Spatial-Frequency Masking in Vision: Critical Bands and Spread of Masking. *Journal of the Optical Society of America* 62, 10 (oct 1972), 1221. <https://doi.org/10.1364/JOSA.62.001221>
- Eugene Switkes, Arthur Bradley, and Karen K. De Valois. 1988. Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings. *Journal of the Optical Society of America A* 5, 7 (1988), 1149. <https://doi.org/10.1364/josaa.5.001149>
- Eugene Switkes and Michael A. Crognale. 1999. Comparison of color and luminance contrast: apples versus oranges? *Vision Research* 39, 10 (may 1999), 1823–1831. [https://doi.org/10.1016/S0042-6989\(98\)00219-3](https://doi.org/10.1016/S0042-6989(98)00219-3)
- Abhinav K. Venkataramanan, Cosmin Stejerean, and Alan C. Bovik. 2022. Funque: Fusion of Unified Quality Evaluators. In *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE, Bordeaux, France, 2147–2151. <https://doi.org/10.1109/ICIP46576.2022.9897312>
- Abhinav K. Venkataramanan, Cosmin Stejerean, Ioannis Katsavounidis, and Alan C. Bovik. 2023. One Transform To Compute Them All: Efficient Fusion-Based Full-Reference Video Quality Assessment. arXiv:2304.03412 (Nov. 2023). <http://arxiv.org/abs/2304.03412> arXiv:2304.03412 [eess].
- Z Wang, E.P. Simoncelli, and A.C. Bovik. 2003. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*. IEEE, 1398–1402. <https://doi.org/10.1109/ACSSC.2003.1292216>
- AB Watson and JA Solomon. 1997. Model of visual contrast gain control and pattern masking. *Journal of the Optical Society of America A* 14, 9 (1997), 2379–2391. <http://www.opticsinfobase.org/abstract.cfm?URI=josaa-14-9-2379>
- Andrew B Watson. 1993. DC Tune: a technique for visual optimization of DCT quantization matrices for individual images. In *Society for Information Display Digest of Technical Papers XXIV*. 946–949.
- Andrew B. Watson, H. B. Barlow, and John G. Robson. 1983. What does the eye see best? *Nature* 302, 5907 (mar 1983), 419–422. <https://doi.org/10.1038/302419a0>
- Sophie Wuergler, Maliha Ashraf, Minjung Kim, Jasna Martinovic, Maria Pérez-Ortiz, and Rafal K. Mantiuk. 2020. Spatio-chromatic contrast sensitivity under mesopic and photopic light levels. *Journal of Vision* 20, 4 (April 2020), 23.
- Nanyang Ye, Krzysztof Wolski, and Rafal K. Mantiuk. 2019. Predicting Visible Image Differences Under Varying Display Brightness and Viewing Distance. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 5429–5437. <https://doi.org/10.1109/CVPR.2019.00558>
- Emin Zerman, Vedad Hulusic, Giuseppe Valenzise, Rafal K Mantiuk, and Frédéric Dufaux. 2018. The relation between MOS and pairwise comparisons and the importance of cross-content comparisons. *Electronic Imaging* 2018, 14 (2018), 1–6.
- Lin Zhang, Ying Shen, and Hongyu Li. 2014. VSI: A Visual Saliency-Induced Index for Perceptual Image Quality Assessment. *IEEE Transactions on Image Processing* 23, 10 (oct 2014), 4270–4281. <https://doi.org/10.1109/TIP.2014.2346028>
- Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. 2011. FSM: A Feature Similarity Index for Image Quality Assessment. *IEEE Transactions on Image Processing* 20, 8 (aug 2011), 2378–2386. <https://doi.org/10.1109/TIP.2011.2109730>
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*. 586–595. <https://doi.org/10.1109/CVPR.2018.00068>
- X. Zhang and B. A. Wandell. 1997. A spatial extension of CIELAB for digital color-image reproduction. *Journal of the Society for Information Display* 5, 1 (1997), 61. <https://doi.org/10.1889/1.1985127>




Fig. 25. All 8 distortions as used in the XR-DAVID study at level 3 (most distorted). The undistorted reference is shown in the center for comparison. An inset of the *Snow* scene is shown to enhance artifact visibility. Note that DCE (Dynamic Correction Error) does not exhibit much visible distortion as the artifact is mostly present in the temporal domain. “A Drone Shot of a Church in Belalp” by SwissHumanity Stories under a Creative Commons license.