Colour Constancy: Biologically-inspired Contrast Variant Pooling Mechanism

BMVC 2017 Submission # 331

Abstract

Pooling is an ubiquitous operation in image processing algorithms that allows higherlevel processes to collect relevant low-level features from a region of interest. Maxpooling is one of the most commonly used operators in the computational literature. However, it can lack robustness to outliers due to the fact that it relies merely on the peak of a function. Pooling mechanisms are also present in the primate visual cortex where neurons of higher cortical areas pool signals from lower ones. The receptive fields of these neurons have been shown to vary according to the contrast by aggregating signals over a larger region in the presence of low contrast stimuli. We hypothesise that this contrast-variant-pooling mechanism can address some of the shortcomings of maxpooling. We modelled this contrast variation through a histogram clipping in which the percentage of pooled signal is inversely proportional to the local contrast of an image. We tested our hypothesis by applying it to the phenomenon of colour constancy where a number of popular algorithms utilise a max-pooling step, e.g. White-Patch, Grey-Edge and Double-Opponency. For each of these methods we investigated the consequences of replacing their original max-pooling by the proposed contrast-variant-pooling. Our experiments on three colour constancy benchmark datasets suggest that previous results can significantly improve by adopting a contrast-variant-pooling mechanism.

⁰²⁷ 1 Introduction

000

007

011

Many computer vision frameworks contain a pooling stage that combine local responses at different spatial locations [I]. This operation is often a sum- or max-pooling mechanism implemented by a wide range of algorithms, *e.g.* in feature descriptors (such as SIFT [I] and HOG [II]) or convolutional neural networks [I], III). Choosing the correct pooling operator can make a great difference in the performance of a method [I]. Current standard pooling mechanisms lack the desired generalisation to find an equilibrium between frequently-occurring and rare but informative descriptors [I]. Therefore, many computer vision applications can benefit from a more dynamic pooling solution that takes into account the content of pooled signals.

Such pooling operators are commonly used in modelling the phenomenon of colour constancy (*i.e.* a visual effect that makes the perceived colour of surfaces remain approximately constant under changes in illumination [12, 26]) both in biological [2] and computational [33] solutions. Despite decades of research, colour constancy still remains as an open question [12], and solutions to this phenomenon are important from a practical point of view, *e.g.* camera manufacturers need to produce images of objects that appear the same

^{© 2017.} The copyright of this document resides with its authors.

J45 It may be distributed unchanged freely in print or electronic forms.

as the actual objects in order to satisfy customers, etc. Motivated by above, in this article we 046 propose a contrast-variant-pooling mechanism and investigate its feasibility in the context of 047 computational colour constancy. 048

1.1 Computational Models of Colour Constancy

Mathematically, the recovery of spectral reflectance from a scene illuminated by light of 052 unknown spectral irradiance is an ill-posed problem (it has infinite possible solutions). The 053 simplest and most popular solution has been to impose some arbitrary assumptions regarding 054 the scene illuminant or its chromatic content. Broadly speaking colour constancy algorithms 055 can be divided into two categories: (i) low-level driven, which reduce the problem to solv-056 ing a set of non-linear mathematical equations [21, 53, 54], and (ii) learning-based, which 057 train machine learning techniques on relevant image features [1], [1], [1]. Learning-based 058 approaches may offer the highest performance results, however, they have major setbacks which make them unsuitable in certain conditions: for example, (a) they rely heavily on 060 training data that is not easy to obtain for all possible situations, and (b) they are likely to 061 be slow [2] and unsuitable for deployment inside limited hardware. A large portion of lowlevel driven models can be summarised using a general Minkowski expression [12], 12] such 063 as: 064

$$L_{c}(p) = \left(\int_{\Omega} [f_{c}(x)]^{p} dx\right)^{\frac{1}{p}} = ke_{c},$$
(1) 06

where f(x) is the image's pixel value at the spatial coordinate $x \in \Omega$; *c* is one of the three 067 $\{R, G, B\}$ channels; *p* is the Minkowski norm; and *k* is a multiplicative constant chosen such 068 that the illuminant colour, *e*, is a unit vector. 069

Distinct values of Minkowski norm p results into different pooling mechanisms. Setting 070 p = 1 in Eq. 1 reproduces the well known Grey-World algorithm (sum-pooling), in which 071 it is assumed that all colours in the scene average to grey [**D**]. Setting $p = \infty$ replicates 072 the White-Patch algorithm (max-pooling), which assumes the brightest patch in the image 073 corresponds to the scene illuminant [**C**]. In general, it is challenging to automatically tune 074 p for every image and dataset. At the same time inaccurate p values may corrupt the results 075 noticeably [**C**].

The Minkowski framework can be generalised further by replacing f(x) in Eq. 1 with 077 its higher-order derivative [44]. These non-linear solutions are analogous to centre-surround 078 mechanisms of visual cortex [22], which is typically modelled by a Difference-of-Gaussians 079 (DoG) operators where a narrower, positive Gaussian plays the role of the "centre" and a broader, negative Gaussian plays the role of the "surround" [1, 5]. Recently, a few 081 biologically-inspired models of colour constancy grounded on DoG have offered promising results [2], solution their efficiency largely depends on finding an optimum pooling strategy for higher cortical areas. In short, pooling is a crucial component of many colour constancy models driven by low-level features (or even in deep-learning solutions [**B**]). In the primate visual systems the size of the receptive field varies according to the local contrast 086 of the light falling on it [], []] presenting a dynamic solution dependent on the region of 087 interest.

The low-level models mentioned above are related to the early stages of visual processing, *i.e.* the primary visual cortex (area V1), that are likely to be involved in colour constancy. Physiological measures suggest that although receptor fields triple in size from area V1 to operative field size modulation is similar, keeping the same optimized to the same optimized to the early stages of visual processing, *i.e.* the primary visual cortex (area V1), that are likely to be involved in colour constancy.

size dependency with contrast properties found in V1 [1]. This is consistent with the large 092 body of physiological and psychophysical literature highlighting the significance of contrast in the visual cortex. In computer vision, contrast-dependent models have also shown encouraging results in various applications such as visual attention [22], tone mapping [42]. boundary detection [2], to name a few. From these we can hypothesise the convenience and explanatory value of various "pooling strategies" such as those proposed by previous 097 colour constancy methods. In the rest of this work we will explore the advantages of replacing the different feed-forward (pooling) configurations of some successful colour constancy 099 models [21], 53, 54] by that of the visual system (as described by current neurophysiological 100 literature [3, [3]). Our aim in doing so is dual, in one hand we want to explore the technolog-101 ical possibilities of creating a more efficient algorithm and on the other hand we would like to test the idea that contrast-variant-pooling might play an important role in colour constancy.

1.2 Summary of Contributions

In the present article we propose a generic contrast-variant-pooling (CVP) mechanism that can replace standard sum- and max-pooling operators in a wide range of computer vision applications. CVP is based on local contrast and therefore it offers a dynamic solution that adapts the pooling mechanism according to the content of region of interest. Figure 1 illustrates the flowchart of CVP. We tested the feasibility of CVP in the context of colour constancy by substituting the pooling operation of four algorithms: (a) White-Patch [I], (b) first-order Grey-Edge [I], (c) second-order Grey-Edge [I], and (d) Double-Opponency [I] in three benchmark datasets. Results of our experiments show the quantitative and qualitative benefits of CVP over max-pooling.



Figure 1: Flowchart of the proposed contrast-variant-pooling (CVP) mechanism in the context of colour constancy. We implemented CVP through a *top-x-percentage-pooling*. Given an input image or a feature map: (i) value of the *x* is computed according to the inverse of local contrast for each channel, and (ii) we estimate the scene illuminant as the average value of the *top-x-percentage* of pooled pixels (those on the right side of the depicted dashed lines).

2 Model

2.1 **Max-pooling Colour Constancy**

141 One of the earliest computational models of colour constancy (White-Patch) is grounded on 142 the assumption that the brightest pixel in an image corresponds to a bright spot or specular 143 reflection containing all necessary information about the scene illuminant [1]. Mathemati-144 cally this is equivalent to a max-pooling operation over the intensity of all pixels: 145

$$L_c = \underset{x,y}{\arg\max} I_c(x,y), \tag{2} 14$$

147

179

180

138

140

148 where L represents the estimated illuminant at each chromatic channel $c \in \{R, G, B\}$; I is the 149 original image and (x, y) are the spatial coordinates in the image domain.

150 One important flaw of this simple approach is that a single bright pixel can misrepresent 151 the whole illuminant. Furthermore, the White-Patch algorithm may fail in the presence of noise or clipped pixels in the image due to the limitations of the max-pooling operator [22]. 153 One approach to address these issues is to account for a larger set of "white" points by 154 pooling a small percentage of the brightest pixels (e.g. the top 1%) [12], an operation referred 155 as top-x-percentage-pooling. In this manner the pooling mechanism is collectively computed considering a group of pixels rather than an single one. This small variant might be a crucial 156 factor in the estimation of the scene illuminant [5]. A similar mechanism has also been 157 deployed successfully in other applications such as shadow removal [16]. 158

In practice, given the chosen x-percentage, the top-x-percentage-pooling can be imple-159 mented through a histogram-based clipping mechanism [12, 16]. Let H be the histogram of 160 the input image I, and let $H_c(i)$ represents number of pixels in colour channel c with intensity 161 $i \in [0, \dots, K_c]$ (histogram's domain). The scene illuminant L_c is computed as: 162

$$L_{c} = \frac{1}{N_{c}} \sum_{i=1}^{K_{c}} i \cdot H_{c}(i), \qquad (3) \frac{163}{164}$$

where N_{k_c} is the total number of pixels within the intensity range k_c to K_c . The values of 167 $k_{c=\{R,G,B\}}$ are determined from the chosen *x*-percentage such that:

$$N_{k_c} = \sum_{i=k_c}^{K_c} H_c(i) = x_c \cdot P,$$
(4)
169
170
171

where P is the total number of pixels in the image and x_c is the chosen percentage for colour 173 channel c. Within this formulation it is very cumbersome to define a universal, optimally-174 fixed percentage of pixels to be pooled $[\square]$ and consequently the free variable x requires 175 specific tuning for each image or dataset individually. Naturally this limitation restricts the 176 usability of the *top-x-percentage-pooling* operator. In the following sections we show how 177 to automatically compute this percentage, based on the local contrast of each image. 178

2.2 Pooling Mechanisms in the Visual Cortex

181 We know from the physiology of the cerebral cortex that neurons of higher cortical areas pool information from lower areas over increasingly larger image regions. Although the exact pooling mechanism is yet to be discovered, "winner-takes-all" and "sparse coding" 183 kurtotical behaviour are common to many groups of neurons all over the visual cortex [1, 53], and it is conceivable that a mechanism analogous to max-pooling might be present within the cortical layers. Indeed such behaviour has been discovered across a population of cells in the cat visual cortex [51] and the activation level of cells with max-like behaviour was reported to vary depending on the contrast of visual stimuli.

Results reported by [1] suggest an inverse relationship between the contrast of a stim-189 ulus and the percentage of the signal pooled. When pooling neurons were exposed to low 190 contrast stimuli their responses shifted slightly away from pure max-pooling (selecting the 191 highest activation response within a region) towards integrating over a larger number of 192 highly activated neurons. In the language of computer vision, this can be regarded as top-193 *x-percentage-pooling*, where x assumes a smaller value in high contrast and a larger value 194 in low contrast. Interestingly, the pooling of those neurons remained always much closer 195 to max-pooling than to the linear integration of all neurons (sum-pooling) [1]. Mathematically, this can be interpreted as having a very small (top-x-percentage-pooling) x value. 197

It does not come as a great surprise that the pooling mechanism in the visual cortex depends of the stimulus' contrast. There is a large body of physiological studies showing that receptive fields (RF) of neurons are contrast variant (for a comprehensive review refer to [**I**]). Quantitative results suggest that RFs in visual area one (V1) of awake macaques double their size when measured at low contrast [**II**]. Similar expansions have also been reported for the RFs of neurons in extrastriate areas, such as V2 and V4. This implies that a typical neuron in higher cortical areas that normally pool responses from its preceding areas over about three neighbouring spatial locations [**II**] can access a substantially larger region to pool from in the presence of low contrast stimuli . This is in line with the reported pooling mechanism in the cat visual cortex [**II**].

2.3 Contrast Variant Pooling

In order to model this contrast variant pooling mechanism, we first computed local contrast C of the input image I at every pixel location by means of its local standard deviation

212 213

214 215

223

$$C_{c}(x,y;\sigma) = \sqrt{\frac{1}{\#\mathcal{N}_{\sigma}(x,y)}} \sum_{(x',y')\in\mathcal{N}_{\sigma}(x,y)} \left(I_{c}(x',y') - \mu(x,y)\right)^{2}$$
(5)

where *c* indexes each colour channel; (x, y) are the spatial coordinates of a pixel; μ is an isotropic average kernel with size σ ; and $\mathcal{N}_{\sigma}(x, y)$ represents the neighbourhood centred on pixel (x, y) of radius σ .

To simulate this inverse relation between stimulus contrast and percentage of signal pooled [51] in the *top-x-percentage-pooling* operator, we determined the percentage x_c in Eq. 4 as the average of inverse local contrast signals:

$$x_c = \frac{1}{P} \sum_{x,y} \frac{1}{C_c(x,y;\sigma)}, \ \forall (x,y) \in \Omega,$$
(6)

where C_c is computed from Eq. 5, and Ω is the spatial image domain. In this fashion, instead of defining a fix percentage of signal (pixels) to be pooled (as in [II]), we chose an adaptive percentage according to the contrast of image. In term of colour constancy, this effectively relates the number of pooled pixels to compute the scene illuminant to the average contrast of an image. We illustrated this mechanism of contrast-variant pooling in the central panel of Figure 1, where red, green and blue signals correspond to each chromatic channel. Pixels 230 on the right side of the dashed lines (k_c) are pooled. In this example contrast is higher for the 231 red signal and therefore a smaller percentage of cells are pooled in the red channel. 232

Bearing in mind that "contrast" is just a fraction in the range [0, 1] – with 0 characterising 233 an absolutely uniform area and 1 representing points with the highest contrast, *e.g.* edges – 234 we can predict that the percentage *x* will be a very small number for natural images where 235 homogeneous regions are likely to form the majority of the scene. Consequently in our *top-*236 *x-percentage-pooling* operator we always pool a small percentage. This is in agreement with 237 observations of [53] which indicate that such pooling mechanism is always much closer to max-pooling than to sum-pooling. 239

2.4 Generalisation to Other Colour Constancy Models

There is a number of colour constancy models in the literature which are driven by low-level features that require a pooling mechanism on top of their computed feature maps in order to estimate the scene illuminant. In the Double-Opponency [21] algorithm this feature map is computed by convolving a colour-opponent representation of the image with a DoG kernel followed by a max-pooling operation. In the Grey-Edge [22] model, higher-order derivatives of the image are calculated through its convolution with the first- and second-order derivative of a Gaussian kernel. This is complemented by a Minkowski summation, which oscillates between sum- and max-pooling depending on its norm.

Similar to the White-Patch algorithm, the pooling mechanism of these models can also 250 be replaced with our *top-x-percentage-pooling* operator, where x is computed according to 251 the local contrast of image as explained above. The only difference is that instead of pooling 252 from an intensity image (as in case of the White-Patch algorithm), Double-Opponency and 253 Grey-Edge pool over their respective feature maps. This means that Eq. 3 receives a feature 254 map M instead of an intensity image I as input.

3 Experiments and Results

We investigated the efficiency of the proposed contrast-variant-pooling (CVP) framework when applied to four different colour constancy algorithms whose source code is publicly available: White-Patch [53], First-order Grey-Edge [54], Second-order Grey-Edge [54], and Double-Opponency [72]. We simply replaced their max-pooling operator with our proposed pooling mechanism. To evaluate each method we used the recovery angular error as defined by:

$$\boldsymbol{\varepsilon}^{\circ}(\boldsymbol{e}_{e}, \boldsymbol{e}_{t}) = \arccos\left(\frac{\langle \boldsymbol{e}_{e}, \boldsymbol{e}_{t} \rangle}{\|\boldsymbol{e}_{e}\|\|\boldsymbol{e}_{t}\|}\right),\tag{7}$$

where $\langle . \rangle$ is the dot product of the estimated illuminant e_e and the ground truth e_t ; and 268 $\|.\|$ represents the Euclidean norm of a vector. It is worth mentioning that this error mea-269 sure might not to correspond precisely to observer preferences [13], however, it is the most 270 commonly used comparative measure in the literature. We also computed the reproduction 271 angular error [13] in all experiments (due to lack of space these results are not reported here). 272 Readers are encouraged to check the accompanying supplementary materials.

We tested our model on three benchmark datasets¹, (i) SFU Lab (321 images) [**D**], (ii) $\frac{1}{274}$

241

Colour Checker (568 images) [22], and (iii) Grey Ball (11,346 images) [21]. Table 1 shows the best median and trimean angular errors for each method (these metrics were proposed by [23] and [23] respectively to evaluate colour constancy algorithms since they are robust to outliers). Mean angular errors are reported in the supplementary materials.

281			SFU Lab [5]		Colour Checker [Grey Ball [
282	Method		Median	Trimean	Median	Trimean	Median	Trimean
283	White-Patch	[33]	6.44	7.51	5.68	6.35	6.00	6.50
284	Grey-Edge 1st-order	[44]	3.52	4.02	3.72	4.76	5.01	5.80
285	Grey-Edge 2 nd -order	[44]	3.22	3.65	4.27	5.19	5.72	6.39
286	Double-Opponency	[2]]	2.37	3.27	3.46	4.38	4.62	5.28
287	CVP White-Patch		2.99	3.42	2.97	3.45	5.02	5.15
288	CVP Grey-Edge 1st-o	order	3.29	3.72	2.48	2.79	4.70	5.17
289	CVP Grey-Edge 2 nd -	order	2.85	3.13	2.59	2.93	4.65	5.05
290	CVP Double-Oppone	ency	2.02	2.56	2.39	2.84	4.00	4.24

Table 1: Recovery angular errors of four methods with max-pooling and contrast-variantpooling (CVP) on three benchmark datasets. Lower figures indicate better performance.

Figure 2 illustrates three exemplary results of the proposed contrast-variant-pooling (CVP) operator for two colour constancy models: White-Patch and first-order Grey-Edge derivative. Qualitatively, we can observe that CVP does a better job than max-pooling in estimating the scene illuminant. This is also confirmed quantitatively for the angular errors, shown at the right bottom side of each computed output.



Figure 2: Qualitative results of White-Patch (WP) and the first-order Grey-Edge (GE1) under max- and contrast-variant-pooling (CVP). Angular errors are indicated on the right bottom corner. Images are from SFU Lab, Colour Checker and Grey Ball dataset respectively.

³¹⁵³¹⁶ 3.1 Influence of Free Parameters

294

295

298

310

311

314

For each free variable of tested models we compared performance of max- to contrastvariant-pooling. White-Patch does not have any free variable, therefore it is exempted from this analysis. In Figure 3 we have reported impact of different σ s (receptive filed size) on Double-Opponency algorithm for the best and the worst results obtained by free variable *k* in each dataset (results for all *k*s are available in the supplementary material). We can observe that almost in all cases contrast-variant-pooling outperforms max-pooling. The improvement 322 is more tangible for the Colour Checker and Grey Ball datasets and in low σ s. 323



free variables of Double-Opponency [22] algorithm (k and σ).

Figure 4 illustrates the impact of different σ s (Gaussian size) on the first- and secondorder Grey-Edge algorithm. We can observe similar patterns as with Double-Opponency 377 (contrast-variant-pooling outperforms max-pooling practically in all cases). This improvement is more significant for low σ s, for the Colour Checker dataset and for the secondorder derivative. It must be noted here that our objective is to study the performance of 340 max-pooling and CVP on top of the Grey-Edge algorithm. In this respect, *CVP Grey-Edge* 341 angular errors are on par with the best reported results for max-pooling Grey-Edge using 342 Minkowski norm optimization for each dataset [12], With the important caveat that CVP has 343 no extra variables to be tuned, whereas in the Minkowski norm optimisation the value of p 344 must be hand-picked for each dataset. 345



Figure 4: Comparison of max- and contrast-variant-pooling for free variable σ of Grey-Edge [12] algorithm (both first- and second-order derivatives). 356

357

362 363

364

From Figures 3 and 4 we can observe that the greatest improvement occurs in the Colour 358 Checker dataset. We speculate that one of the reasons for this is the larger range of intensity 359 values in the Colour Checker dataset (16-bit) in comparison to the other two datasets that 360 contain 8-bit images, therefore, an inaccurate max-pooling is greatly penalised. 361

3.2 Discussion

We would like to emphasise that the objective of this article is not to improve state-of-the-art ³⁶⁵ in colour constancy, but to show that contrast-variant-pooling (CVP) almost always produces ³⁶⁶ improvements over max-pooling. Surprisingly, the results we obtained are even competitive ³⁶⁷

with the state-of-the-art. For instance in the SFU Lab dataset, the lowest reported angu-369 lar error is 2.1 (obtained by an Intersection-based Gamut algorithm [2]). This means that CVP Double-Opponency with an angular error of 2.0 outperforms the state-of-the-art in this 370 dataset. In the Colour Checker and Grey Ball datasets there are a few learning-based mod-371 els (e.g. Exemplar-based method [23]) that obtain lower angular errors than CVP Double-372 *Opponency*, nevertheless our results are comparable with theirs. 373

Physiological evidence besides, the better performance of CVP can be explained intu-374 itively by the fact that max-pooling relies merely on the peak of a function (or a region of 375 interest), whereas in our model, pooling is defined collectively based on a number of ele-376 ments near the maximum. Consequently those peaks that are outliers and likely caused by 377 noise get normalised by other pooled elements. The rationale within our model is to pool a 378 larger percentage at low contrast since in those conditions, peaks are not informative on their 379 own, whereas at high contrast peaks are likely to be more informative and other irrelevant details must be removed (therefore a smaller percentage is pooled). 381

Although the importance of choosing an appropriate pooling type has been demonstrated both experimentally $[\Sigma X, \Box T]$, and theoretically $[\Omega]$, current standard pooling mechanisms lack the desired generalisation []]. We believe that contrast-variant-pooling can offer a more dynamic and general solution. In this article we evaluated CVP on the colour constancy phenomenon as a proof-of-concept, however our formulation of CVP is generic (and based on local contrast) and in principle can be applied to a wider range of computer vision 387 algorithms, such as deep-learning, where pooling is a decisive factor [1].

In our implementation of CVP we approximated local contrast through local standard deviation (see Eq. 5). There are at least two factors that require a more profound analysis: (i) incorporating more sophisticated models of contrast perception [22] accounting for extrema in human contrast sensitivity; and (ii) analysing the role of kernel size in the computation of local contrast.

Conclusion 4

396

401

In this article we present a novel biologically-inspired contrast-variant-pooling (CVP) mechanism grounded in the physiology of the visual cortex. Our main contribution can be summarised as linking the percentage of pooled signal to the local contrast of the stimuli, *i.e.* 400 pooling a larger percentage at low contrast and a smaller percentage at high contrast. Our CVP operator remains always closer to max-pooling rather than to sum-pooling since natural 402 images generally contain more homogeneous areas than abrupt discontinuities.

403 We tested the efficiency of our CVP model in the context of colour constancy by replac-404 ing the max-pooling operator of four algorithms with the proposed pooling. After that, we 405 conducted experiments in three benchmark datasets. Our results show that contrast-variant-406 pooling outperforms the commonly used max-pooling operator nearly in all cases. This can 407 be explained by the fact that our model allows more informative peaks to be pooled while 408 suppressing less informative peaks and outliers.

409 We argue that the proposed CVP is more generic and offers a dynamic (and automatic) 410 solution based on the image's local contrast, thus its application can be extended to a wider 411 range of computer vision frameworks. This opens a multitude of possibilities for future lines 412 of research and it remains an open question whether our model can reproduce its excellent 413 results in other domains as well.

10 AUTHOR(S): BIOLOGICALLY-INSPIRED CONTRAST VARIANT POOLING MECHANISM

Re	ferences	414
[1]	Vivek Agarwal, Andrei V Gribok, and Mongi A Abidi. Machine learning approach to color constancy. <i>Neural Networks</i> , 20(5):559–563, 2007.	415 416 417
[2]	Arash Akbarinia and C. Alejandro Parraga. Biologically plausible boundary detec- tion. In <i>Proceedings of the British Machine Vision Conference (BMVC)</i> . BMVA Press, September 2016.	418 419 420
[3]	A Angelucci and S Shushruth. Beyond the classical receptive field: Surround modula- tion in primary visual cortex. <i>The new visual neurosciences</i> , pages 425–444, 2013.	422 423
[4]	Kobus Barnard. Improvements to gamut mapping colour constancy algorithms. In <i>Computer Vision–ECCV 2000</i> , pages 390–403. Springer, 2000.	424 425 426
[5]	Kobus Barnard, Lindsay Martin, Brian Funt, and Adam Coath. A data set for color research. <i>Color Research & Application</i> , 27(3):147–151, 2002.	427 428
[6]	Jonathan T Barron. Convolutional color constancy. In <i>Proceedings of the IEEE Inter-</i> national Conference on Computer Vision, pages 379–387, 2015.	429 430 431
[7]	Y-Lan Boureau, Jean Ponce, and Yann LeCun. A theoretical analysis of feature pooling in visual recognition. In <i>Proceedings of the 27th international conference on machine learning (ICML-10)</i> , pages 111–118, 2010.	432 433 434
[8]	Gershon Buchsbaum. A spatial processor model for object colour perception. <i>Journal of the Franklin institute</i> , 310(1):1–26, 1980.	435 436 437
[9]	Matteo Carandini and David J Heeger. Normalization as a canonical neural computa- tion. <i>Nature Reviews Neuroscience</i> , 13(1):51–62, 2012.	438 439
[10]	Florian Ciurea and Brian Funt. A large image database for color constancy research. In <i>Color and Imaging Conference</i> , volume 2003, pages 160–164. Society for Imaging Science and Technology, 2003.	440 441 442 443
[11]	Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In <i>Computer Vision and Pattern Recognition</i> , 2005. <i>CVPR</i> 2005. <i>IEEE Computer Society Conference on</i> , volume 1, pages 886–893. IEEE, 2005.	444 445 446
[12]	Marc Ebner. Color constancy, volume 6. John Wiley & Sons, 2007.	448
[13]	Christina Enroth-Cugell and John G Robson. The contrast sensitivity of retinal gan- glion cells of the cat. <i>The Journal of physiology</i> , 187(3):517–552, 1966.	449 450 451
[14]	Graham D Finlayson and Elisabetta Trezzi. Shades of gray and colour constancy. In <i>Color and Imaging Conference</i> , volume 2004, pages 37–41. Society for Imaging Science and Technology, 2004.	452 453 454
[15]	Graham D Finlayson and Roshanak Zakizadeh. Reproduction angular error: An improved performance metric for illuminant estimation. <i>perception</i> , 310(1):1–26, 2014.	455 456 457
[16]	Graham D Finlayson, Steven D Hordley, and Mark S Drew. Removing shadows from images. In <i>Computer Vision–ECCV 2002</i> , pages 823–836. Springer, 2002.	458 459

- [17] David A Forsyth. A novel algorithm for color constancy. *International Journal of Computer Vision*, 5(1):5–35, 1990.
- [18] David H Foster. Color constancy. Vision research, 51(7):674–700, 2011.
- [19] Brian Funt and Weihua Xiong. Estimating illumination chromaticity via support vector regression. In *Color and Imaging Conference*, volume 2004, pages 47–52. Society for Imaging Science and Technology, 2004.
- [20] Brian Funt, Kobus Barnard, and Lindsay Martin. Is machine colour constancy good
 enough? In *Computer Vision–ECCV'98*, pages 445–459. Springer, 1998.
- 470

480

494

- [21] Shao-Bing Gao, Kai-Fu Yang, Chao-Yi Li, and Yong-Jie Li. Color constancy using
 double-opponency. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*,
 37(10):1973–1985, 2015.
- 474 [22] Arjan Gijsenij and Theo Gevers. Color constancy using natural image statistics and scene semantics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33 (4):687–698, 2011.
- 478 [23] Arjan Gijsenij, Theo Gevers, and Marcel P Lucassen. Perceptual analysis of distance
 479 measures for color constancy algorithms. *JOSA A*, 26(10):2243–2256, 2009.
- [24] Andrew M Haun and Eli Peli. Perceived contrast in complex images. *Journal of vision*, 13(13):3–3, 2013.
- [25] Steven D Hordley and Graham D Finlayson. Reevaluation of color constancy algorithm
 performance. *JOSA A*, 23(5):1008–1020, 2006.
- [26] Paul M Hubel. The perception of color at dawn and dusk. *Journal of Imaging Science and Technology*, 44(4):371–375, 2000.
- [27] Laurent Itti and Christof Koch. Computational modelling of visual attention. *Nature reviews neuroscience*, 2(3):194–203, 2001.
- [28] Kevin Jarrett, Koray Kavukcuoglu, Yann LeCun, et al. What is the best multi-stage architecture for object recognition? In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2146–2153. IEEE, 2009.
- [29] Hamid Reza Vaezi Joze and Mark S Drew. Exemplar-based color constancy and multiple illumination. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*,
 36(5):860–873, 2014.
- [30] Hamid Reza Vaezi Joze, Mark S Drew, Graham D Finlayson, and Perla Aurora Troncoso Rey. The role of bright pixels in illumination estimation. In *Color and Imaging Conference*, volume 2012, pages 41–46. Society for Imaging Science and Technology, 2012.
- [31] Ilan Lampl, David Ferster, Tomaso Poggio, and Maximilian Riesenhuber. Intracellular
 measurements of spatial integration and the max operation in complex cells of the cat
 primary visual cortex. *Journal of neurophysiology*, 92(5):2704–2713, 2004.

12 AUTHOR(S): BIOLOGICALLY-INSPIRED CONTRAST VARIANT POOLING MECHANISM

[32]	Edwin H Land. An alternative technique for the computation of the designator in the retinex theory of color vision. <i>Proceedings of the National Academy of Sciences</i> , 83 (10):3078–3080, 1986.	506 507 508
[33]	Edwin H Land et al. The retinex theory of color vision. Citeseer, 1977.	509
[34]	Yann LeCun, Bernhard E Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne E Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In <i>Advances in neural information processing systems</i> , pages 396–404, 1990.	510 511 512 513 514
[35]	David G Lowe. Distinctive image features from scale-invariant keypoints. <i>International journal of computer vision</i> , 60(2):91–110, 2004.	515 516
[36]	David Marr and Ellen Hildreth. Theory of edge detection. <i>Proceedings of the Royal Society of London B: Biological Sciences</i> , 207(1167):187–217, 1980.	517 518 519
[37]	Naila Murray and Florent Perronnin. Generalized max pooling. In <i>Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition</i> , pages 2473–2480, 2014.	520 521 522
[38]	Bruno A Olshausen et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. <i>Nature</i> , 381(6583):607–609, 1996.	523 524
[39]	C Alejandro Parraga and Arash Akbarinia. Colour constancy as a product of dynamic centre-surround adaptation. <i>Journal of Vision</i> , 16(12):214–214, 2016.	525 526 527
[40]	Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. In <i>ACM Transactions on Graphics (TOG)</i> , volume 21, pages 267–276. ACM, 2002.	528 529 530
[41]	Dominik Scherer, Andreas Müller, and Sven Behnke. Evaluation of pooling operations in convolutional architectures for object recognition. <i>Artificial Neural Networks–ICANN 2010</i> , pages 92–101, 2010.	531 532 533 534
[42]	L. Shi and B. Funt. Re-processed version of the gehler color constancy dataset of 568 images. http://www.cs.sfu.ca/~colour/data/.	535 536
[43]	S Shushruth, Jennifer M Ichida, Jonathan B Levitt, and Alessandra Angelucci. Com- parison of spatial summation properties of neurons in macaque v1 and v2. <i>Journal of</i> <i>neurophysiology</i> , 102(4):2069–2083, 2009.	537 538 539 540
[44]	Joost Van De Weijer, Theo Gevers, and Arjan Gijsenij. Edge-based color constancy. <i>IEEE Transactions on image processing</i> , 16(9):2207–2214, 2007.	541 542
[45]	Javier Vazquez-Corral, C Párraga, Ramon Baldrich, and Maria Vanrell. Color con- stancy algorithms: Psychophysical evaluation on a new dataset. <i>Journal of Imaging</i> <i>Science and Technology</i> , 53(3):31105–1, 2009.	543 544 545
[46]	HR Wilson and F Wilkinson. Configural pooling in the ventral pathway. <i>The new visual neurosciences</i> , pages 617–626, 2014.	546 547 548
[47]	Jianchao Yang, Kai Yu, Yihong Gong, and Thomas Huang. Linear spatial pyramid matching using sparse coding for image classification. In <i>Computer Vision and Pattern Recognition</i> , 2009. CVPR 2009. IEEE Conference on, pages 1794–1801. IEEE, 2009.	549 550 551