

How Multi-Illuminant Scenes Affect Automatic Colour Balancing

Liwen XU, Brian FUNT

School of Computing Science, Simon Fraser University

ABSTRACT

Many illumination-estimation methods are based on the assumption that the imaged scene is lit by a single source of illumination; however, this assumption is often violated in practice. We investigate the effect this has on a suite of illumination-estimation methods by manually sorting the Gehler et al. ColorChecker set of 568 images into the 310 of them that are approximately single-illuminant and the 258 that are clearly multiple-illuminant and comparing the performance of the various methods on the two sets. The Grayworld, Spatio-Spectral-Statistics and Thin-Plate-Spline methods are relatively unaffected, but the other methods are all affected to varying degrees.

Keywords

Colour balancing, illumination estimation, digital photography.

INTRODUCTION

The usual first step in automatic colour balancing of digital imagery is to estimate the chromaticity of the illumination. Although there are some recent exceptions (Beigpour 2014; Gijssenij 2012; Joze 2013), most illumination-estimation methods assume that the relative spectral power distribution of the illumination is constant throughout the scene. However, many scenes contain multiple illuminants with differing SPDs, and we investigate the effect this has on automatic colour balancing.

Somewhat surprisingly, the Gehler et al. (Gehler 2008) “Colorchecker” data set of 568 images, which is widely used in evaluating competing illumination-estimation methods, contains many images of multiple-illuminant scenes. For example, Figure 1 depicts an indoor scene that also includes a window through which daylight is clearly falling on the counter. Is the scene illuminant the light from inside the room or outside the window? Figure 2 shows an outdoor scene with at least three illuminant types: the cloudy sky, the shadowed areas, and the traffic light.

Each image in the Colorchecker dataset contains an Xrite/Macbeth ColorChecker, which is used to provide a ground-truth measure of the illumination’s ‘colour’. However, since many of the scenes do contain multiple illuminants, a single such measurement cannot possibly represent the colour of all the illuminants correctly, but rather must represent some sort of compromise. Whether the illumination-estimation method assumes there is a single illuminant or multiple illuminants, a single colorchecker cannot correctly represent the ground-truth illumination in a multi-illuminant scene. In this paper, we investigate how much of an effect this has on a representative set of illumination-estimation methods; namely, MaxRGB (Funt 2012), Grayworld (Weijer 2007), Shades-of-Gray (Finlayson 2004), Edge-based (Weijer 2007), N-jet (Gijssenij 2010), Thin-Plate-Spline (Shi 2011) and Spatio-Spectral Statistics (Chakrabarti 2012).

SCENE CLASSIFICATION

The Gehler et al. dataset (Gehler 2008) contains 568 images taken with two digital single lens reflex cameras, a Canon 5D and a Canon 1D. All images were saved in Canon RAW format. Each image contains an Xrite/Macbeth ColorChecker for reference. The image coordinates (measured by hand) of each Colorchecker square are provided with the dataset. In the tests below, we used the Shi et al. (Shi 2011) reprocessed version of the Gehler et al. data. The original dataset consists of non-linear TIFF images that were automatically generated from the RAW data. The reprocessed dataset contains PNG images that are linear and do not include any automatic white balancing or de-mosaicing.

We manually sorted the original 568 images into two groups according to whether the images were of single-illuminant or multiple-illuminant scenes. Sorting in this way is difficult because it can be hard to discern the nature of the scene illumination from the image. Figure 1 shows a typical case where the presence of multiple sources of illumination is very clear. Similarly, the traffic light in Figure 2 is an obvious additional illuminant. Figure 3 shows a situation in which it seems pretty clear that there is only a single illuminant. Figure 4 shows a somewhat ambiguous case, since there are areas in direct sun and others in shadow. The Colorchecker itself appears to be partly in sun and partly in shadow. There are also the clouds in the distance. However, this appears to be a typical outdoor scene basically dominated by sunlight/skylight and so we classified it as a single-illuminant scene. If we were to be any more strict in our interpretation of what constitutes a single-illuminant scene then almost the entire dataset would be classified as multiple-illuminant. Based on this type of analysis of each image, the 568 dataset is divided into 310 single-illuminant and 258 multiple-illuminant scenes. We denote the two images subsets as S (single) and M (multiple), and the full set of 568 images as F. The complete lists of image numbers for sets S and M are listed in the Appendix.



Figure 1 Example of a multiple-illuminant scene with light coming both from the room and window.



Figure 2 Example of a multiple-illuminant scene containing a visible light source.



Figure 3 Example of a clearly single-illuminant scene.



Figure 4 Example of a somewhat ambiguous scene with sunlight and shadow but classified as single-illuminant nonetheless.

COMPARATIVE PERFORMANCE ON SINGLE-VERSUS MULTIPLE-ILLUMINANT SCENES

We evaluate the illumination-estimation performance of all the methods separately on subset S, subset M, and the complete set F. The illumination-estimation methods are MaxRGB (Funt 2012), Gray-World (Weijer 2007), Shades of Gray (Finlayson 2004), Edge Based (Weijer 2007), N-jet (Gijssenij 2010), TPS (Shi 2011) and Spatio-Spectral Statistics (Chakrabarti 2012). The image pixels occupied by the Colorchecker in each image areas are replaced with zeros for the tests. These methods all estimate the rg-chromaticity of the illumination. The error in a given estimate is measured relative to the measured ground-truth illumination chromaticity. The error is evaluated in terms of the angular difference in degrees between the two chromaticities after each chromaticity is converted to a 3-vector as (r, g, 1-r-g). The overall accuracy across a given test set of images is reported in terms of the mean, median, RMS and maximum errors.

Table 1: Comparative illumination-estimation performance evaluated in terms of angular error. MaxP (MaxRGB w/o preprocessing), MaxM (MaxRGB after median filtering, GW (Grayworld), EB (Edge-Based, first and second order), 1-jet (Gamut mapping), 2-jet (Gamut mapping), SSS-ML (Spatio-Spectral Statistics with maximum likelihood, SSS-GP (Spatio-Spectral Statistics with general prior). Md (Median), Mn (Mean).

	Set S				Set M				Complete Dataset			
	Md	Mn	RMS	Max	Md	Mn	RMS	Max	Md	Mn	RMS	Max
MaxP	4.6	7.3	9.9	27	16	14	16	50	9.1	10	13	50
MaxM	3.1	5.3	7.8	26	8.8	10	13	42	4.7	7.7	11	42
GW	4.0	5.5	7.3	25	3.3	3.8	4.7	15	3.6	4.8	6.2	25
SoG (norm=6)	2.9	4.8	6.8	23	7.4	8.4	10	36	4.5	6.4	8.7	36
EB1 (norm=6)	2.6	4.6	6.8	26	7.2	9.0	12	38	3.8	6.6	9.4	38
EB1 (norm=1)	3.8	4.7	5.6	17	3.4	4.1	4.9	16	3.6	4.4	5.3	17
EB2	2.7	5.0	7.2	28	8.9	9.9	13	47	4.4	7.2	10	47
1-jet 3-fold	2.8	4.4	6.2	24	6.0	7.7	9.9	32	4.1	5.9	8.1	32
2-jet 3-fold	2.9	4.3	6.1	21	5.9	7.7	9.8	32	4.2	5.9	8.0	32
TPS 3-fold	2.6	3.5	4.5	17	3.0	3.6	4.5	15	2.7	3.5	4.5	17
SSS-ML	2.9	3.7	4.8	22	3.1	3.7	4.5	15	3.0	3.7	4.7	22
SSS-GP	2.9	3.6	4.7	22	3.0	3.6	4.4	15	3.0	3.6	4.6	22

The MaxRGB tests include MaxRGB without preprocessing and MaxRGB with median filtering MaxM (Funt 2012). The N-jet algorithms are tested using threefold cross-validation. Because the images are from two different cameras and the training is specific to each camera, we train and test on the images from each camera separately and then combine the results. TPS is also evaluated using threefold cross-validation.

DISCUSSION

The results in Table 1 show that the effect of multiple scene illuminants on illumination-estimation performance varies substantially across the various methods. MaxRGB is strongly influenced by the presence of multiple illuminants. Whether MaxRGB includes image preprocessing or not, the presence of multiple illuminants seriously influences the results. As Table 1 shows, the error is approximately tripled. For example, the median error for the MaxRGB variant MaxM rises from 3.1 to 8.8 degrees for the change from subset S to subset M. In other words, MaxRGB works very well when the single-illuminant assumption holds, but fails when it is violated. Since MaxRGB is based on estimating the maximum value in each of the R, G, B channels, it is particularly vulnerable to the presence of light sources such as the traffic light in Figure 2.

Interestingly, Grayworld's performance appears to be unaffected by the presence of multiple illuminants since the median angular error on sets S, M and F is 4.0, 3.3, and 3.6, respectively. Although its overall performance is poorer than several of the other methods, it has the advantage of being stable. The Shades of Gray approach is controlled by the choice of the Minkowski norm to vary between the extremes of Grayworld and MaxRGB. A norm of 6 has been reported to work well (Finlayson 2004). As a compromise between Grayworld and MaxRGB, however, its performance is then affected by the presence of multiple illuminants, with a median angular error of 2.9, 7.4 and 4.5, on sets S, M and F, respectively.

Just as Shades of Gray is more sensitive to the presence of multiple illuminants than Grayworld, the Edge-Based method using norm = 6 is more sensitive than the Edge-Based method using norm = 1. For the norm = 1 case, Edge-Based is simply averaging derivatives within each RGB channel instead of the RGB values themselves. With norm = 6, the Edge-Based method weights the large derivatives, which are likely to arise from illumination boundaries in multiple-illuminant scenes, more heavily thus leading to a concomitant increase in angular error. The performance of the 1-jet and 2-jet Gamut Mapping methods also degrades when the single-illuminant assumption is violated. For 1-jet, the median angular on S of 2.8 degrees increases to 6.0 degrees for M. Gamut mapping assumes that the gamut of RGBs and the gamuts of their derivatives are limited by the illuminant. In the presence of multiple illuminants these image gamuts expand such that the constraint used to estimate the illuminant is no longer as strong or accurate.

The learning-based methods appear to account for the presence of multiple illuminants quite well. The performance of the Spatio-Spectral Statistics methods tends to be very good and quite unaffected by multiple illuminants. With median angular errors of 2.6 (set S), 3.0 (set M) and 2.7 (set F), the Thin-Plate Spline method (TPS) is both the least affected by multiple illuminants and also attains the minimum error of all the methods on each of the three datasets.

ACKNOWLEDGMENT

Funding was provided by the Natural Sciences and Engineering Research Council of Canada

REFERENCES

- Beigpour, S., C. Riess, J. van d Weijer, and E. Angelopoulou. 2014. Multi-illuminant estimation with conditional random fields. *IEEE Trans. on Image Processing* 23(1): 83-96.
- Chakrabarti, A. K. Hirakawa, and T. Zickler. 2012. Color constancy with spatio-spectral statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(8): 1509-1519. Accompanying computer implementation downloaded from <http://vision.seas.harvard.edu/colorconstancy/>. Accessed February 23, 2015.
- Finlayson, G. and E. Trezzi. 2004. Shades of gray and colour constancy. In *Proc. IS&T/SID 12th Color Imaging Conf.* Springfield, VA, 37-41.
- Funt, B., and L. Shi. 2012. MaxRGB reconsidered. *Journal of Imaging Science and Technology* 56(2): 020501-1-020501-10.
- Gehler, P., C. Rother, A. Blake, T. Minka, and T. Sharp. 2008. Bayesian color constancy revisited. In *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition*. Anchorage, Alaska, 1-8.
- Gijssenij, A., T. Gevers, and J. van deWeijer. 2010. Generalized gamut mapping using image derivative structures for color constancy. *Int. J. Computer Vision*. 86, 127-139. Accompanying computer implementation downloaded from <http://colorconstancy.com/>. Accessed February 27, 2015.
- Gijssenij, A., R. Lu, and T. Gevers. 2012. Color constancy for multiple light sources. *IEEE Trans. on Image Processing* 21(2): 697-707.
- Joze, H.R.V. and Mark S. Drew. 2013. Exemplar-based colour constancy and multiple illumination. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 36(5): 860-873.
- Shi, L., W. Xiong, and B. Funt. 2011. Illumination estimation via thin-plate spline interpolation. *J. Opt. Soc. Am. A* 28(5): 940-948. Accompanying computer implementation downloaded from <http://www.cs.sfu.ca/~colour/code/>. Accessed February 25, 2015.
- Shi, L. and B. Funt. Re-processed version of the Gehler color constancy dataset of 568 images. Downloaded from http://www.cs.sfu.ca/~colour/data/shi_gehler/. Accessed Sept. 2014.
- Weijer, J. van de, T. Gevers, and A. Gijssenij. 2007. Edge-based color constancy. *IEEE Trans. Image Processing*. 16(9): 2207-2214. Accompanying computer implementation downloaded from <http://lear.inrialpes.fr/people/vandeweijer/software>. Accessed February 25, 2015.

APPENDIX

Image numbers of single-illuminant scenes: 1 2 3 4 7 10 11 12 14 17 18 19 21 23 27 32 33
34 35 36 37 38 39 40 41 44 45 46 47 48 49 50 52 54 56 57 58 59 63 64 65 66 67 68 69 70
71 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 99 100 101 102 103
104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 121 122 126 127 128
129 130 131 132 134 135 136 137 139 140 141 142 143 144 147 149 150 151 152 153 154
155 156 157 158 159 160 161 162 163 164 165 166 168 169 170 171 172 173 174 176 177
178 179 180 185 186 187 188 189 191 193 196 197 199 200 203 204 206 212 219 224 225
226 227 228 229 230 232 233 234 235 248 249 250 251 253 254 255 256 257 258 259 260
261 262 264 265 266 267 268 269 270 271 273 274 275 281 282 283 285 286 287 288 289
291 292 293 295 298 299 300 301 302 303 304 305 306 307 308 309 310 342 343 344 353
355 358 359 362 364 365 367 372 377 381 384 386 394 395 396 398 401 402 407 410 411
414 416 417 418 419 420 422 423 424 425 427 428 429 430 431 432 433 434 436 437 439
441 444 449 451 453 454 455 456 458 459 461 462 463 466 473 474 475 478 479 480 482
484 487 489 490 491 493 494 496 500 502 503 504 516 522 528 530 536 537 538 541 543
546 549 560 561 562 564

Image numbers of multiple-illuminant scenes: 5 6 8 9 13 15 16 20 22 24 25 26 28 29 30 31
42 43 51 53 55 60 61 62 72 73 74 75 97 98 120 123 124 125 133 138 145 146 148 167 175
181 182 183 184 190 192 194 195 198 201 202 205 207 208 209 210 211 213 214 215 216
217 218 220 221 222 223 231 236 237 238 239 240 241 242 243 244 245 246 247 252 263
272 276 277 278 279 280 284 290 294 296 297 311 312 313 314 315 316 317 318 319 320
321 322 323 324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340 341
345 346 347 348 349 350 351 352 354 356 357 360 361 363 366 368 369 370 371 373 374
375 376 378 379 380 382 383 385 387 388 389 390 391 392 393 397 399 400 403 404 405
406 408 409 412 413 415 421 426 435 438 440 442 443 445 446 447 448 450 452 457 460
464 465 467 468 469 470 471 472 476 477 481 483 485 486 488 492 495 497 498 499 501
505 506 507 508 509 510 511 512 513 514 515 517 518 519 520 521 523 524 525 526 527
529 531 532 533 534 535 539 540 542 544 545 547 548 550 551 552 553 554 555 556 557
558 559 563 565 566 567 568.

*Address: Brian Funt, School of Computing Science, Simon Fraser University,
8888 University Drive, Burnaby, British Columbia V5A 1S6, Canada
E-mails: funt@sfu.ca, xuliwenx@gmail.com*