Automatic Image Quality Analysis of Arbitrary Targets with Colorite

Henrik Johansson; National Library of Sweden

Abstract

Institutions are often reluctant to introduce color management because they believe that an introduction demands large changes to existing digitization workflows and results in lower production volumes. To make the introduction easier and to simplify image quality analysis, the Digital Production Division at the National Library of Sweden has developed Colorite. Colorite is a cross-platform software tool for automatic target-based analysis of image quality, designed with emphasis on both flexibility and efficiency. The flexibility enables institutions to incorporate Colorite into a wide range of production environments with minimal changes to existing workflows and the efficiency allows for large production volumes.

Colorite employs state-of-the-art image matching algorithms to allow for the use of both arbitrary targets and of arbitrary target placements. Colorite accurately determines the location of an arbitrary target in the image frame without any user intervention, regardless of the target's placement and orientation. Using the location of the target, Colorite automatically computes a comprehensive suite of established image quality metrics that are compared to pre-determined tolerances. The results from the comparison are presented individually for each metric and areas in which the quality needs to be improved are automatically pinpointed.

Introduction

To achieve consistent image quality, targets must both be included and analyzed in the digitization workflow. Several initiatives and guidelines for target-based analysis of image quality have matured during the last few years and they are now ready for large-scale adoption – the most prominent of these being FADGI [5] and Metamorfoze [4]. Despite these established guidelines, many institutions are reluctant to include targets and color management into their digitization workflows because color management is often believed to result in both lower production volumes and in large changes to existing digitization workflows. Furthermore, institutions tend to overestimate the technical knowledge needed to implement the guidelines and to adopt a color managed workflow. Hence, software tools that remedy these concerns and beliefs must be developed to simplify the implementation of target-based image quality analysis.

The Digital Production Division at the National Library of Sweden has developed Colorite, a flexible and efficient cross-platform software tool for automatic target-based analysis of image quality. Colorite is a major part of the division's ongoing effort to create a fully automated post-processing workflow. Colorite conforms to the image quality specifications and guidelines set forth in FADGI and in Metamorfoze.

Colorite employs state-of-the-art image matching algorithms to allow for the use of both arbitrary targets and arbitrary target placements. As a result, institutions using Colorite are not restricted to specific targets; they can choose freely among all available targets or even use targets of their own design. To minimize both the workload and the impact on existing workflows, Colorite automatically detects the location of the target in the image frame, regardless of the target's placement and its orientation. To our knowledge, the ability to use and locate arbitrary targets is unique for Colorite.

Colorite automatically computes a comprehensive suite of quality metrics, as defined in the FADGI and Metamorfoze guidelines. The results are presented individually for each metric and any area in which the quality needs to be improved is automatically pinpointed. Colorite uses XML to store comprehensive information about the employed target, the image quality tolerances, and the result of the image quality analysis. This information is suitable for long-term preservation to provide both security and traceability.

By using Colorite for image quality analysis, an institution only needs to make a few small modifications to its current workflow to get all the benefits of a color managed workflow. The impact on the digitization volumes will be minimal if color management is introduced together with Colorite while the volumes are likely to increase for institutions that already use color management.

Related work

There exist several commercial tools for image quality analysis, e.g. Image Engineering iQ-Analyzer [6], X-rite ColorChecker Passport Camera Calibration [12], Image Science Associates Golden Thread [8] and ColorGauge [7], Certifi Media Pedigree [3], and Imatest Master [9]. All of these tools are powerful and have numerous advanced features and configuration options. Unfortunately, for institutions that lack prior experience with color management, these features and configuration options can actually increase the resistance to introduce a color managed workflow.

Furthermore, several of the tools are more suitable to characterize image capture devices than to continuously evaluate the image quality in a production environment. These tools generally require large targets that are difficult to place efficiently together with an object during image capture.

All of the existing tools are limited to a single or at most a few pre-defined targets, reducing the flexibility. Except for iQ-Analyzer, the tools are only available for Windows, decreasing the flexibility even further.

Image matching is limited in most tools. They tools generally require a tight crop of the reference image around the target and a nearly vertical or horizontal target orientation, restricting the target

Illustration 1. Different targets in the reference image are located by changing the target image. The target in the left image is custom made at the Division of Digital Production and based on the X-rite ColorChecker Mini. The target in the right image is the common Kodak Q-13. The images appear to be too dark because the viewer does not support color management.

placement. Often, manual involvement is also needed to locate the target. The only tool that seems to allow for arbitrary target placements is Certifi Media Pedigree.

Finally, except for Imatest Master, the commercial tools store the result of the image quality analysis in formats that are unsuitable for long-term preservation. The results from the analysis are normally exported as either plain text files or in proprietary file formats such as Microsoft Excel.

Design

To lower the barrier for implementation of target-based image quality analysis, Colorite is designed with emphasis on both flexibility and efficiency. The flexibility enables institutions to incorporate Colorite into a wide range of production environments with minimal changes to existing workflows. The efficiency allows for large production volumes, despite the addition of color management to the workflow.

In Colorite, flexibility is achieved by addressing a number of key operations in the workflow. First, Colorite allows for the use of arbitrary targets. Institutions are free to use any targets they like, even targets of their own design. Second, Colorite does not force its users to employ a specific set of quality metrics. While Colorite conforms to the FADGI and Metamorfoze guidelines, Colorite can also compute older RGB-based quality metrics. The user can always decide which quality metrics that should be computed and set the allowed tolerances for each metric. Third, Colorite accepts many common image formats, like TIFF, JPEG, JP2, and PNG. Fourth, Colorite is designed to run on multiple operating systems, allowing a user to choose their preferred computing platform. Finally, no expert knowledge in color management is needed to use Colorite, allowing for a wide range of potential users and operators.

The functionality of Colorite allows for efficient workflows and large production volumes. First, Colorite is designed to be fast. The complete run-time on a common computer is generally significantly less than 10 seconds; including the time to read a 60 MB TIFF-file, locating the target in the image, and computing all quality measurements. Second, Colorite has minimal impact on the image capture process because the software allows for fast and arbitrary target placements. Finally, Colorite can generally be executed without any explicit user configuration. At most, users need to spend a few seconds to specify the name and the location of a few files.

Beside the flexibility and the efficiency, a third area of great importance is long-term preservation. To allow for arbitrary use of an image file in the future, the results from the quality measurement should ideally be available for as long as the image exists. Colorite stores all information about both the target and the image quality in an intuitive XML-format. The XML-file is easy to include into SIPs (Submission Information Packet) or to store alongside the images.

Automatic target location

The core functionality of Colorite is the ability to automatically locate an arbitrary target in the image frame, regardless of the target's position and orientation. To locate the target, Colorite uses advanced image matching algorithms from the field of image analysis. An image matching algorithm that is considered for use in tools like Colorite should fulfil several important requirements. The algorithm must

- have a high accuracy and be able to identify an arbitrary target, regardless of any objects or shapes that might be present in the image

- be able to identify the target regardless of its position, size, and orientation in the image, i.e. the algorithm must be invariant to translation, scaling, and rotation
- be unaffected by badly exposed images, distortions, difficult lighting conditions, and noise
- be computational efficient to allow for large production volumes

Image matching is generally performed using two images. The first image – the target image – contains a cropped image of the target, or any other object that should be located. The second image – the reference image – can depict virtually anything, including the target object. Hence, to locate an arbitrary target in a reference image, a user only needs to supply an image of the target.

Traditionally, image matching has been performed by employing the target's appearance as a template. Using a sum-comparing metric, the location of the target is determined by testing all or a sample of the possible target locations within the reference image. A common method is to extract the edges in both the target image and in the reference image. A distance-based metric is computed using the edges in the target image and in the reference image. The edges of the target is then moved, scaled and rotated until the distance metric is minimized which implicates that the region in the reference image that is most similar to the target image has been found. One common algorithm that employs this methodology is hierarchical chamfer matching [2]. Algorithms that rely on the target's appearance generally produce accurate matchings for "well-behaved images", but the algorithms are unfortunately sensitive to noise, scaling, distortions, and clutter.

During the last ten years, image matching has increasingly focused on algorithms that use features for the matching instead of the appearance of the target. These new algorithms describe an image by a set of features, e.g. corners, blobs, lines, and high contrast areas. Feature-based algorithms match the feature description for the target image with feature subsets for the reference image until the most similar subset in the reference images is located. Because the feature-based algorithms generally are invariant to scale, orientation, affine distortion, and partially invariant to illumination changes, they can often accurately identify a target with an arbitrary orientation, even among clutter and under partial occlusion.

The most well-known feature-based algorithm is SIFT (Scale-invariant feature transform) [11]. In SIFT, the features of an image are stored in feature vectors. The image matching is based on the Euclidian distance between the feature vectors of the target and the reference image. From the full set of matches, subsets of features that agree on location, scale, and orientation are identified and filtered out. The subsets are subject to further computations to verify their similarity and to discard outliers. Finally, the probability that a particular feature subset in the reference image corresponds to the target is computed, given their similarity and an estimated number of false matches.

Building on the ideas of SIFT, a significantly faster and slightly more accurate algorithm was developed. The new algorithm, SURF (Speeded-Up Robust Features), uses an intermediate image representation that significantly reduces the number of computations compared to SIFT [1]. SURF also introduced a new and more accurate feature detector based on Hessians and the use of 2D Haar wavelets to describe the features.

Illustration 2. The SURF image matching algorithm often manages to locate the target even in cluttered images and under partial occlusion. Compared to Illustration 1, the image matching accuracy is slightly reduced.

SURF is considered to be the state-of-the-art algorithm for feature matching and its properties – high computational efficiency, high accuracy, and insensitivity to noise, translation, scaling and distortion – makes it the ideal algorithm for Colorite.

Image quality metrics

Colorite computes a number of metrics to assert the image quality. If the real-world size of the target is known, Colorite will also compute the theoretical resolution of the reference image. A prerequisite for the computation of most quality metrics is that the image uses the LAB color space. The quality metrics are listed below.

	Metric
Exposure	ΔL or ΔE
Colour accuracy	ΔE
Colour cast	ΔC
Tonal separation	Gain modulation (gm)
Noise	Standard deviation of sample

Table 1. Image quality metrics computed by Colorite.

The accuracy of the exposure is measured using

$$\Delta L = \sqrt{(L_1 - L_2)^2}$$

where the subscript "1" denotes the real-world tonal values on the target's grayscale patches while the subscript "2" denotes the tonal values in the reference image. ΔL is computed for each patch, as the average for all patches, and as the maximum value for a single patch.

Colorite uses the CIE76 definition of ΔE for color accuracy.

$$\Delta E = \sqrt{(L_1 - L_2)^2 + (A_1 - A_2)^2 + (B_1 - B_2)^2}$$

where the subscript "1" denotes the real-world color values on the target's color patches while the subscript "2" denotes the color values in the reference image. ΔE is computed for each patch, as the average for all patches, and as the maximum value for a single patch. It is possible to also use ΔE as a metric for the exposure instead of ΔL. Since ΔE includes color, the tolerances must in these cases be slightly increased compared to ΔL.

The ΔE-metric exists in several versions. The chosen CIE76 definition of ΔE is easy to implement, but more importantly, it is easy and intuitive to understand for users without prior experience of color management. Research has shown that the newer definitions of ΔE (CIE94 and CIEDE2000) are slightly more accurate for perceived color differences [10], but the improvements do not currently justify the work and higher complexity associated with their implementation in Colorite.

The color cast in the reference image is measured using

$$\Delta C = \sqrt{(A_1 - A_2)^2 + (B_1 - B_2)^2}$$

where the subscript "1" denotes the real-world color values on the target's gray scale patches while the subscript "2" denotes the color values in the reference image. The color cast is computed individually for each patch.

The gain Modulation is calculated using two separate measurements of ΔL, where each measurement only includes the tonal values from either real-world measurements or from the reference image.

$$gm = \frac{\sqrt{(L_1^1 - L_2^1)^2}}{\sqrt{(L_1^2 - L_2^2)^2}}$$

where the subscript denotes the used patches, the elevated digit "1" denotes the reference image, and the elevated digit "2" denotes the real-world measurements. Hence, above the fractional line, ΔL is computed using tonal values from the reference image; below the fractional line, ΔL is computed using tonal values from real-world measurements. The Metamorfoze guidelines specify which combinations of patches that preferably should be selected [4]. To comply with Metamorfoze, Colorite uses information contained in the target's configuration file to automatically assert which patches that should be selected to compute the gain modulation.

The noise in the image is measured by computing the standard deviation of the tonal and color values included in the samples. To minimize the influence of dust and hot/dead pixels, all measurements are performed using a sample of 5x5 pixels.

Implementation

Colorite is mainly written in Java to be platform independent. To both minimize the dependence on commercial vendors and allow for the widest possible distribution, all external libraries and software employed by Colorite are platform independent, open-source, and free to use.

To locate the target in the reference image, an implementation of the SURF image matching algorithm from the OpenCV framework is used. Binaries compiled with OpenCV are platform specific, which makes it necessary to compile a separate version of the image matching component for each platform. Fortunately, only minor changes are needed to compile the component for multiple platforms. To reduce the run-time of the image matching, the size of the reference image is reduced internally within Colorite. The amount of reduction depends on the size of the reference image.

The color and tonal values are measured at the center of each patch in the reference image. To determine the center of a patch, the properties of the target is stored in a configuration file using XML. The file contains the relative distance from the target's upper left corner to the center of each patch, with regard to the total size of the target (example: if the target has a real-world size of 10 cm x 2 cm and the center of a patch is located 3 cm to the right and 1 cm below the upper-left corner, the values 0.3 and 0.5 are stored in the configuration file).

The image matching component determines the positions of the target's corners in the reference image. To pinpoint the center of a patch, Colorite uses the location of the corners and the relative distance from the target's upper-left corner to the center of the patch.

ImageMagick is used for all measurements of color and tonal values. By default, each sample is measured as the average of 5x5 pixels. Color and tonal measurements in ImageMagick is currently limited to images with a color depth of 8-bit/channel. Images stored with 16-bit/channel have to be converted into 8-bit color/channel before any measurements can be made.

To create and parse XML, Colorite uses functionality embedded in the framework libxml2.

Interface and usage

To allow for users without prior experience of color management, Colorite's graphical user interface is designed to be intuitive and easy to use. Information that is generally static, e.g. data about the target and the allowed tolerances for the image quality metrics, is stored in configuration files using XML. The use of configuration files dramatically reduces the number of configuration options needed in the interface. Instead of being overwhelmed by a large number of choices, the user only has to select the configuration files that correspond to the current job. The number of available configuration files will most likely be limited; most institutions will only need two or three configuration files for their preferred image quality tolerances and a single configuration file for each of their targets. In the rare cases where the image quality tolerances has to be changed or when a new

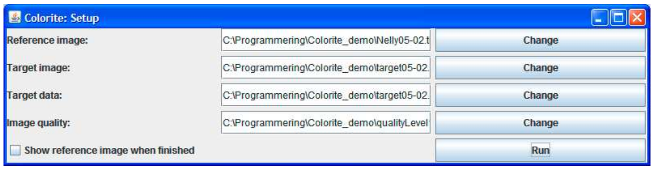

Illustration 3. The simple and intuitive user interface. Note that the user only has to specify at most four configuration files.

target is brought into use, it is easy to modify or create the corresponding configuration file with a simple text editor.

Default names for the configuration files, the reference image, and the target image are set in advance by the user. For the most common case – when no changes has to be made to either the configuration files or the image file names – the user only has to click the “Run”-button to start an image quality analysis.

To verify that the estimated target location in the reference image is correct, the user has the option to display an image where the target’s outline is superimposed on top of the reference image. This image will also highlight the estimated center of each patch.

After the execution, the user is presented with a comprehensive results window that makes it easy determine if the image quality is acceptable. At the top of the window, the used configuration files and the name of both the reference image and the target are shown. If computed, the resolution of the image is displayed at the bottom of the result window.

In the results window, both the reference color value and the measured color value are presented for each patch. Each of the image quality metrics are displayed alongside the color values for the corresponding patch. To easily assert the outcome of the analysis, each image quality metric that meets its specified tolerance is highlighted with the text “pass” in green. Metrics that fail to meet its specified tolerance are marked with a red “fail”.

All results from the image quality analysis are stored in a file using an intuitive XML schema. This file is suitable for long-term storage alongside the image files to allow for arbitrary use of the images in the future.

Initial experiences

Work has begun to integrate Colorite into the workflows at the Digital Production Division at The National Library of Sweden. Colorite is also being evaluated by several independent institutions, both in Sweden and in other countries. Initial experiences show that it is easy to integrate Colorite into existing workflows. For color managed workflows, it is generally straightforward to replace current methods or tools for image quality analysis with Colorite. For institutions that does not yet use color management, Colorite can easily be introduced in the workflow, preferably immediately after the image capture. None of the institution that evaluates Colorite has needed more than a few hours to install and start using Colorite for automatic image quality measurements.

Colorite results in minimal overhead, especially for workflows where the reference images are always given the same filename. For these cases, no configuration has to be performed and Colorite can be executed immediately after the image capture is finished. Even for workflows where different configurations and image quality requirements are regularly used, the additional overhead to select the configuration files is negligible.


Colorite: Results
Target: C:\Programmering\Colorite_demo\target05-02.xml
Image quality: C:\Programmering\Colorite_demo\qualityLevel1.xml
Reference Image: C:\Programmering\Colorite_demo\Nelly05-02.tif

Patch	Reference value	Measured value	Noise	Delta-E	Outcome	Delta-L	Outcome	Delta-C	Outcome	Gain Modulation	Value	Outcome
1:	(20.4, 0.0, -0.6)	(19.9, -0.1, -0.8)	1.0	0.55	Pass	0.5	Pass	0.22	Pass	L*95-L*80:	1.03	Pass
2:	(36.2, -0.5, -0.5)	(35.8, -0.8, -0.4)	0.0	0.51	Pass	0.4	Pass	0.32	Pass	L*80-L*20:	1.01	Pass
3:	(50.4, -0.5, 0.3)	(50.3, -1.5, 0.7)	0.0	1.08	Pass	0.1	Pass	1.08	Pass			
4:	(66.7, -0.8, -0.2)	(67.0, -2.1, -0.2)	0.0	1.33	Pass	0.3	Pass	1.3	Pass			
5:	(81.3, -0.5, 0.0)	(81.3, -3.0, -0.1)	1.0	2.5	Pass	0.0	Pass	2.5	Fail			
6:	(95.9, -0.6, 2.2)	(96.3, -1.6, 1.0)	0.0	1.61	Pass	0.4	Pass	1.56	Pass			
7:	(29.6, 12.2, -48.9)	(28.8, 12.0, -48.7)	1.0	0.85	Pass							
8:	(55.3, -37.7, 32.1)	(55.3, -38.4, 32.6)	1.0	0.86	Pass							
9:	(42.0, 50.2, 27.0)	(42.3, 50.8, 27.7)	3.0	0.97	Pass							
10:	(81.2, 4.9, 79.8)	(81.8, 4.1, 82.7)	6.0	3.07	Pass							
11:	(52.0, 49.6, -13.6)	(51.8, 49.9, -14.2)	0.0	0.7	Pass							
12:	(50.5, -28.5, -28.0)	(51.8, -29.6, -28.4)	0.0	1.75	Pass							
Mean:				1.37	Pass							

Image resolution: 189.46 ppi

Illustration 4. The results window.

Furthermore, Colorite's short run-time has minimal impact on the digitization throughput. For workflows where the image quality has previously been analyzed manually, Colorite is likely to increase the digitization throughput.

The ability to automatically locate the target in the reference image has been found to be remarkably accurate. The placement and the orientation of the target have never affected the accuracy of the matching. However, if the reference image is small compared to the total size of the image, the accuracy might be reduced if the size of the reference image is decreased too much within Colorite.

The noise metric is surprisingly useful as a measurement of the state of the target because noise levels are likely to be substantially increased for dirty or damaged patches. Hence, a high noise level for a patch can indicate that the target is in need of cleaning or replacement. A good example is the high noise level of patch number 10 in illustration 4, as a closer examination revealed a small but distinct smudge on the patch.

Future work

Using Colorite, institutions can easily incorporate color management into their digitization workflows. If an image has a too low quality, Colorite will automatically pinpoint the areas in which the quality needs to be improved. However, to be truly useful, especially for institutions that lack prior experience of color management, tools like Colorite should also automatically correct images where the quality is too low. If the automatic correction should fail, the tools should preferably give the user advice on how to manually correct the images or how to adjust the image capture equipment before re-digitization.

The Digital Production Division at the National Library of Sweden has recently initiated the development of Colorite+. For images that do not meet the image quality tolerances, as defined by the digitizing institution, Colorite+ will attempt to automatically correct both the tonal values and the colors. Colorite+ will use the provided image quality metrics to compute translations within the LAB color space. After these translations have been applied to the image, the quality metrics will be re-computed to assert if the quality has been sufficiently improved. Images that still have a too low quality after the correction will be marked for re-digitization. The development of Colorite+ will focus on image corrections, but the ultimate goal is to also provide advice on how to manually correct the images and how to adjust the image capture equipment. It is currently too early to estimate the efficiency the computed translations and to tell when a first version will be available.

Although development has already started on Colorite+, several improvements can also be made to the original version of Colorite. The quality metrics should preferably be expanded to also include both the optical resolution and the uniformity of the lighting. The resolution can be measured using the Spatial Frequency Response (SFR) and the lighting uniformity can be determined by the use of either multiple reference images or by using a large target with multiple measurement regions. The inclusion of SFR and lighting uniformity in Colorite provides no major technical difficulties. However, because resources at the Digital Production Division are scarce, it is uncertain if these features can be added in the near future, especially since they are more suitable for device characterization than for analysis of image quality in a production environment.

To remedy the rare cases where the image matching accuracy is decreased due to a too large reduction of the reference image size, additional logic will be included in Colorite. To reduce the run-time, a parallelization of the image matching step is being considered.

Work is also performed to integrate Colorite with a production database and workflow tool that is being developed at the Digital Production Division. When the integration is completed, the results from the image quality analysis will be stored in the data base and allow a user to easily analyze and compare the performance of different equipments and set-ups over longer periods of time. The integration will also result in the ability to analyze how the image quality for specific equipment changes over time.

References

[1] Herbert Bay, Andreas Ess, Tinne Tuytelaars, Luc Van Gool, "SURF: Speeded Up Robust Features", Computer Vision and Image Understanding (CVIU), Vol. 110, No. 3 (2008).

[2] Gunilla Borgefors, "Hierarchical Chamfer Matching: A Parametric Edge Matching Algorithm", IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 6, no. 10 (1988).

[3] Certifi Media, Certifi Media Pedigree QP, http://www.certifi-media.com (2011).

[4] Hans van Dormolen, Metamorfoze Preservation Imaging Guidelines, http://www.metamorfoze.nl (2010).

[5] Federal Agencies Still Image Digitization Working Group, Guidelines: Technical Guidelines for Digitizing Cultural Heritage Materials, http://www.digitizationguidelines.gov (2010).

[6] Image Engineering, Image Engineering iQ-Analyzer, http://www.image-engineering.de (2011).

[7] Image Science Associates, Color Gauge, http://www.imagescienceassociates.com/ (2011).

[8] Image Science Associates, Golden Thread, http://www.imagescienceassociates.com/ (2011).

[9] Imatest, Imatest Master, http://www.imatest.com (2011).

[10] Henri Kivinen, Mikko Nuutinen, and Pirkko Oittinen, Comparison of Colour Difference Methods for Natural Images, Proceedings of the Fourth European Conference on Colour in Graphics, Imaging and MCS/10 Vision 12th International Symposium on Multispectral Colour Science (2010).

[11] David Lowe, Object Recognition from Local Scale-Invariant Features, Proceedings of the International Conference on Computer Vision (1999).

[12] X-Rite, X-Rite ColorChecker Passport Camera Calibration, http://www.xrite.com (2011).

Author Biography

Henrik Johansson received both his MS in Engineering Physics (2003) and his PhD in Scientific Computing (2009) from Uppsala University, Sweden. He is currently working as a Production Developer at the Digital Production Division at the National Library of Sweden where he is responsible for the library's digitization development. Henrik is on the Swedish Standards Institute committee for imaging and a member of IS&T.

