Moving Image Quality Analysis to the Cloud

Tim Zaman, Picturae BV, The Netherlands

Abstract

In this paper a novel and modern method of image quality analysis is presented; Delta. We are moving quality analysis to the cloud, where the platform can be accessed from any computer with an Internet connection. You can simply drag and drop the digital image(s) that include(s) the test target(s) on the website. It supports many target makes and models, opaque as well as transparent. The server will analyze the test target and report the results immediately, quantifying quality in terms of the guideline of your preference. The software is made to automatically detect which test target you have uploaded, and how to analyze it. The results are all stored in databases and available for later reference.

Test target serial number readout is also incorporated. This means that when a QR-code is put on or near the test target, color values can be checked in reference to its own true values, instead of the values of the average test target. Because this whole system is web-based, we can take full advantage of its ease of use, accessibility, storage, databases, visualization and functionality. Verifying quality now literally takes less than a minute. Delta has been in use internally for one year, and has so far checked over 10.000 images. Now we are going to open up the system to anyone with a test target, in anticipation that it will bring us all closer to the ultimate digitization efficiency, quality and consistency.

Introduction

When we analyze and quantify digital image quality, we can choose out of a vast variety of different test targets, software to read these targets, and guidelines. Any sensible combination of the previous should eventually lead to the same characterization of the image quality. As a large digitalization company, Picturae has to fit the various quality demands of its customers, requesting various test targets and guidelines. Verifying image quality can be cumbersome, while its quantification is mostly objective. From 15 minutes per verification involving multiple employees (operator, datamanager, qualitymanager) two years ago, verification now takes less than a minute by only the operator himself. The new work flow is simple: open your web-browser, and drop your target image on it. The server will check the target to the desired guidelines that are set for the project he is going to work for, and display the results.

Generally, the best existing image quality analysis (QA) software today are licensed per installation, static and stand-alone tools that can usually only handle one type of target. It requires skill to operate and interpret its results, with quality demands being project-specific. One source of error in quantifying color accuracy is when you test your target in reference to the value of the average target. This error can never be accounted for if the target can not be uniquely identified. Therefore, one has to extract the targets serial number, and check it against its own true values.

So we want to incorporate different target types, their unique values, standardized QA reporting, custom project demands and guidelines, and provide insight for the customer. This means that the way forward for QA is using databases, accessible from anywhere. We propose this novel method in the form of the website Delta (Fig. 1), a QA platform for anyone from anywhere.

Figure 1. *The logo of the platform.*

Background

With the ever increasing digitization volumes, Picturae (formerly Pictura Imaginis) has had to evolve from its creatively organized workstations in 2009, to the industrialized environment it is today, in order to maintain efficiency and quality while keeping up with the pace. For example, transparency digitization takes place in pressurized rooms with operators wearing antistatic clothing. Custom machines have been made to clean originals (items to be digitized) with ionized air, and webtools have been developed for datamanagement, 1:1 completeness control and cropping. All of the processes involved in digitization have been internally standardized [1] to facilitate an efficient and consistent workflow that can accommodate millions of scans of various types each year.

After an image had been scanned manually, it will follow an automated trajectory of post-processing on the servers. Everything is clearly structured and managed digitally, for example, the management can see how many scans are being made in real-time on each system and for which customer. One of the bottlenecks of the industrialization and standardization proved to be the image quality assessments.

I will sketch the workflow of QA using the static, old approach. First, an operator checks what customer he is going to work for on his workstation. Then he chooses the appropriate test target(s) and checks what settings his workstation should be on. The operator digitizes the target and takes the digital image 'upstairs' to the datamanagement which have the QA software installed. The result is displayed, and has to be interpreted (preferably by a qualitymanager) and compared to the specific demands of the project; and some projects demand as much as five different test targets, to be checked with different QA software. Then the results can be exported to a folder that includes the results, often even in the form of a screenshot of the GUI with the results. Now imagine that 20 operators come in at 08:15 and all want their quality analyzed, and again at 13:00 and 18:00. Then imagine that upper-management or the customer wants to compare all the QA results of a different time, project, operator or system with each other. This lead our R&D group to come up a new dynamic, yet standardized approach to QA that fits in our workflow.

Different Targets

Roughly speaking, our main focus in QA lies in obtaining a consistent effective resolution, and a consistent color accuracy over the image. Therefore we need a target or combination of targets to achieve insight in the previous parameters. Which specific target(s) you use to obtain those results should not be of influence on the results themselves, and is therefore open to personal preference. We see this in the existence of various QA targets like notably the Golden Thread (GT) and the Universal Test Target (UTT) which are, objectively, both as valid. Therefore, they are among the targets which delta can automatically detect and analyze.

Our own personal preference in terms of a one-shot test target is a black background with a calibrated X-Rite Digital Colorchecker SG (SG) in the middle, with in each corner of the image a QA-62 SFR target, which provides us with all the essential data. We use this method whenever a project does not have a test target preference. The reason we use this setup is because we consider these as widely recognized and proven test targets.

Other test targets such as the Kodak Gray Scale (Q-13), Kodak Color Control Patches (Q-14) and the Colorchecker Mini or Passport (CC) are also widely accepted and used, and are also supported in Delta.

The Workflow

When the target has been digitized, the browser is opened and the operators first choose which project they are going to work on (to get the proper guidelines and thresholds). Then they just drag and drop the image on the upload window, and after around 15 seconds the result is displayed. For operators, this is as easy as 'green means go' (Fig 2). This is even literally true, as on every cubicle there is a small monitor displaying workstation statistics, connected to a red/green fluorescent tube above the set. This monitor is wirelessly connected to our servers, and lights up the green tube when the workstation is in checked and in operation. The monitors displays things as the project, dpi, ΔE, daily hours of operation and production.

When in-depth knowledge of the quality is of interest, elements in the browser are clickable and expandable, like for instance an MTF curve as displayed in Fig 3. Most of the analyses are valid, and therefore it is for example not necessary for a scan-operator to be able to analyze this MTF curve when its characteristics have been objectively verified. Whenever a QA is invalid, the qualitymanager is usually one click away from identifying the problem. Usually the operator can identify the problem himself, as Delta provides quick hints for improvement, like refocusing, decreasing sharpening, ISO or lighting uniformity. And because all the QA's are stored in the database, the quality manager has his own display where he sees the quality on all the sets. He can identify degrading parameters immediately, and improve them when the workstation is unoccupied.

Quality

Since Delta has been in use, we are able to maintain an average $\Delta E'76$ as low as 1. Operators come in at 08:15 and begin actual production around 08:20. The very same system is in use for both document scanning (books and pages) as well as large format (A0+) and transparencies. They might use targets of different make or model, but as long as the desired parameters can be assessed we can use them for verification.

Targets are Unique

As previously stated, one source of error that is often unaccounted for occurs when you compare your test target to the reference values. Are these the reference values of your own target, or those of the average target of the same type, and what is this average derived from? Is this methodologically correct? For example, the values of color patches of the same target can differ due to manufacturing variability, and degrade due to the amount of use, storage and light exposure. The manufacturer might even slightly change certain values in the future. This why you have to identify the reference values for your own targets, and attach a serial number to the target so that the target on your image can be matched with its proper reference. This is why Delta supports the readout of QR-codes. QR-codes are more robust than bar codes (or handwriting) due to Forward Error Correction and they can store a relatively large amount of data. They should be placed on or near the target, so that Delta can compare the analyzed values from your image to the values from the reference database. It not only promotes accuracy, but also consistency - especially when you have multiple targets of the same type in use.

Currently Supported Types

Currently, the input can be of any common type; JPEG, PNG, TIFF, 8-bit, 16-bit, etc. It can handle all the common colorspaces, reference white's or custom gamma's. Among the results are $\Delta E'76$, $\Delta E'94$ *(T&G)*, $\Delta E'2000$, $\Delta E'CMC$, *MTF50*, *MTF10*, sampling efficiency, over-sharpening, uniformity, noise, dpi.

We have tested it with images as large as 80MP. It supports test targets like the Golden Thread Object and Device Level target, the UTT, Scan Reference Chart (SRC), SG, CC, Q-13, Q-14 and QA-62. Delta does not need to know what you are uploading a-priori, it detects this automatically and will process the separate targets accordingly. Multiple targets per image are not a problem (even of the same type), and it is also possible to upload multiple images at a time. The target detection is scale invariant though we have restricted rotation invariance in order to avoid complications when analyzing slanted edges.

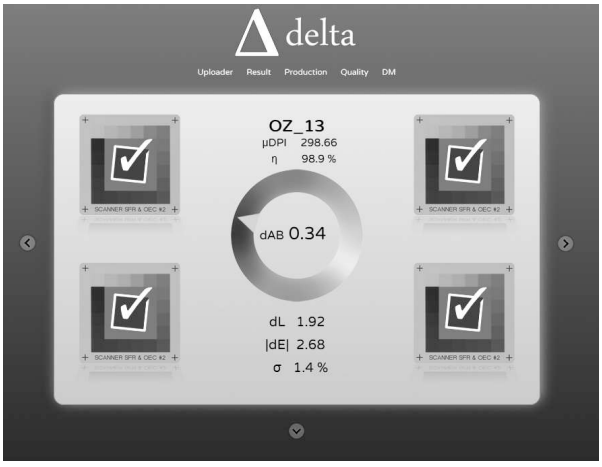

Figure 2. *Example of the current basic output returned to an operator when one SG and four QA-62's have been detected.*

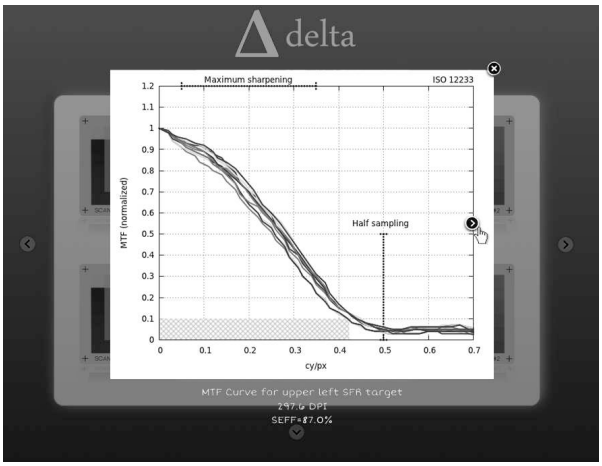

Figure 3. *In-depth info (like an MTF curve) is just a click away.*

Framework

The framework is displayed in Fig. 4, were you can see what happens to an image when it is uploaded. Uploading an image and waiting for it to be processed takes around 15 seconds. Then the results are displayed, and you could download some information or cycle through the data. That is what happens on the client-side.

The client sends uploads the file to the server, where it is picked up by the actual computation server. A C++ program is invoked using various open-source libraries to handle the image and perform the detection. Because speed is important, we use Graphical Processing Units (GPU's) to do computations around 20x faster than the CPU. For the target detection we use a custom Fourier matching algorithm, which has proven to be the most robust for our application. This is due to the repetitive nature of test targets. We had initially implemented other matching algorithms among which SIFT [2] and SURF [3], that were efficient but not robust enough for most targets, mostly due to their repetitive nature.

After a target has been found and read, it searches the images as long as it can not find any more targets in the image. Then it outputs a XML structure, which is saved and read in by the database. It is then compared with its own reference, or when a serial number is inexistent, with an average reference. Then, in order to display results, it will be compared to the guidelines of preference, be it FADGI, Metamorfoze, or your own.

Application Programming Interface

The server side computation is invoked by the client side website, and reports its results back into the system. But since this computation system is independent, it can operate independently from the website. This means that for instance the 20.000 scans made each day can also be automatically processed, if a supported test target is visible in the image. That way we can easily check for consistency, or even include metadata inside each image that contain quality information. We have used the API recently when a customer wanted custom quality reporting of 300 target images reported in an Excel file in order to check quality consistency themselves. Providing customers the insight into the quality of their digital reproduction should be encouraged, which requires no effort when this data is already in the cloud.

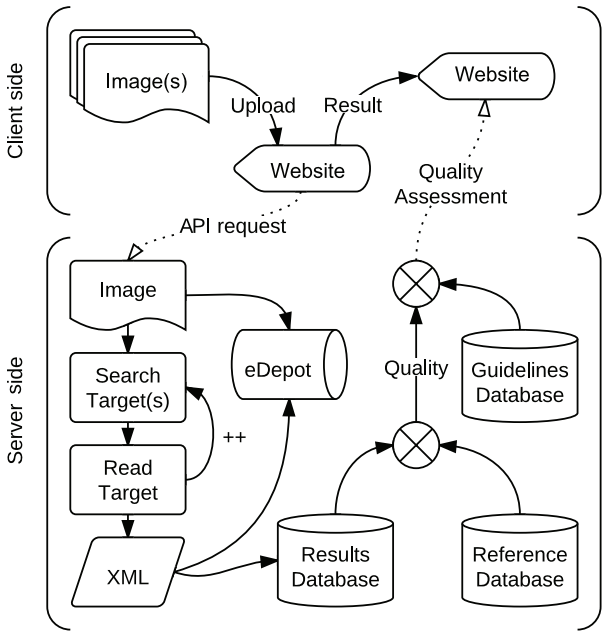

Figure 4. *The flowdiagram when uploading a test target.*

Business Model

Picturae's business focuses on digitization, websites and storage. It has no intentions of commercially exploiting Delta, not does it hold any monetary interest or stakes in any related target or software. The use of Delta is free of charge for every customer doing business with Picturae. We also offer access to the platform free of charge to low-volume users. Commercial parties or instances that would like to incorporate it into their own workflow will be offered a SAAS solution with a small fee to cover operational expenses and R&D costs.

Conclusion

This novel way of quality analysis is the way forward in standardization of digital image quality quantification. Moving to the cloud means unrestricted accessibility and keeping data related and organized in a way that can be adapted to anyone's preference.

Acknowledegments

Gratitude to Don Williams of Image Science Associates and Peter Burns of Burns Digital Imaging for advice and their excellent work on quality analysis through publications and education.

References

[1] O. Slijkhuis, Digitize More, Pay Less: An Industrial Evolution in the Digitization of Cultural Heritage, Proc. Archiving, pg. 227. (2011).

[2] D. G. Lowe, Object Recognition from Local Scale-invariant Features, Proc. Computer Vision, pg. 1150-1157. (1999).

[3] H. Bay, T. Tuytelaars, L.V. van Gool, SURF: Speeded Up Robust Features, Proc. Computer Vision, pg. 404-417. (2006).

Author Biography

Tim Zaman is a R&D employee at Picturae, and a MSc Biorobotics student at Delft University of Technology in The Netherlands.