

# Adaptively Selecting a Printer Color Workflow

Kristyn Falkenstern<sup>1,2,4</sup>, Nicolas Bonnier<sup>1</sup>, Hans Brette<sup>2</sup>, Mehdi Felhi<sup>1,3</sup>, and Françoise Viénot<sup>4</sup>.

<sup>1</sup>Océ Print Logic Technologies S.A. (France). <sup>2</sup>Institut TELECOM, TELECOM ParisTech, LTCI CNRS (France). <sup>3</sup>LORIA, UMR 7503, Université Nancy 2, (France). <sup>4</sup> Muséum National d'Histoire Naturelle (France), Centre de Recherche sur la Conservation des Collections (France).

## Abstract

We present a novel approach to adaptively selecting a color workflow solution per document. The success of a color management solution is often dependent on the document's content. In most workflows, either a compromise between options is chosen for all documents or each document is manually processed. Increased interest in color management has led to more options that a user may choose between, with a variety of choices that impact the perceived quality of the reproductions.

Our proposed method automatically selects a color workflow (output profile and rendering intent) for each input document, dictated by the document's characteristics and a set of color workflow performance tests. The choice of performance tests is specific to each of the predefined quality attributes. A selection engine uses these results, weighs them and makes a recommendation on which workflow to apply. The experimental results indicate that the selection engine can determine which rendering intent to apply, but more work is needed in selecting the exact color workflow when the profiles are of similar quality.

## Introduction

Color management is an essential component in the printing industry to attain precise and repeatable color workflows. The International Color Consortium (ICC) color profiles are used to store the necessary information to transform color data between a device color space and an independent color space [1]. An ICC output profile includes four different rendering intents (referred to as B2An) which are used to address the varying reproduction goals a user may have [2]. Our work includes two of these rendering intents, the colorimetric media-relative (B2A1) and the perceptual (B2A0). The colorimetric intent, aims for as close of a colorimetric match, with the specified viewing conditions and device constraints, as possible. The perceptual rendering intent aims to reproduce pleasing images, a concept that is not clearly defined[3], but allows the profile creator to be more flexible on how the colors are mapped between color spaces. Determining which profile and rendering intent to use is a challenging task with the vast number of color workflows (profile and rendering intent) a user must choose between.

Much effort has been made towards creating an adaptive processing workflow where the final processing is driven by the input document's content [4–6]. Creating a document driven adaptive selection engine is necessary for our model. Much of the previous work has used a training set of documents to create rules that require two inputs: document features and observer preference ratings. Features have been used in the literature to summarize differences between documents or to group documents into categories [7]. Sun [8] described features as a summary of the image

properties that may represent any attribute of an image, for example: image gamut, type, histograms, texture or layout. Other works used the terms statistics, properties, factors, image characteristics and descriptors to describe a document's properties [4–8].

The motivation of our model is to be able to add and remove color workflows easily, which excludes the use of time-consuming observer data as the preference input. Our model replaces the user input with a set of performance results derived from metric tests. Each metric compares the color workflow performance of a specific perceptible Quality Attribute (QA). In our previous work[9], we summarized a set of key QAs that were used to summarize the potential performance differences between color workflows. Our current list includes: *colorimetric accuracy*, *colorfulness*, *gamut boundary*, *smoothness*, *details*, *shadow details*, *highlight details* and *neutrals*. The eight QAs are referred to as  $QA_i$ , where  $i = 1, \dots, 8$  the index of each QA.

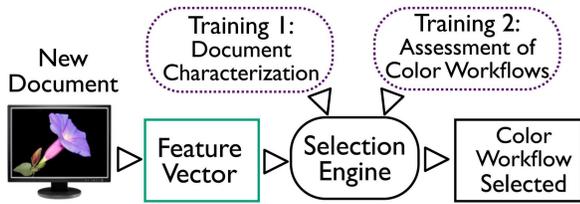
Our hypothesis is that an observer's preference between color workflow options is dependent on two variables: the perceptible differences between the reproductions and the content of the input document. If we can make a connection between the content of the document (features) and the differences between color workflows then we can predict which color workflow to apply for each new input document.

The rest of the paper is organized as follows. We start by giving an overview of the proposed method and describe the selection rules in detail. The rules are then used to make a selection in the next section. Once the selection is made, an observer evaluation is used to verify the performance of the selection engine. We conclude with a summary of results and future work.

## Proposed Method

There are two trainings used in our model, trainings 1 and 2, which generate the rules that are used by the selection engine to choose a color workflow. A training documents set  $D_T$  is used with both. Training 1 creates weights that are used to determine which  $QA_i$  are most important for the input documents, this step only considers the document's features and not the color workflows. Training 2 ranks the performances of the color workflows by finding the differences between a set of target documents  $D_{TG}$  and the output of the processed documents using difference metrics.

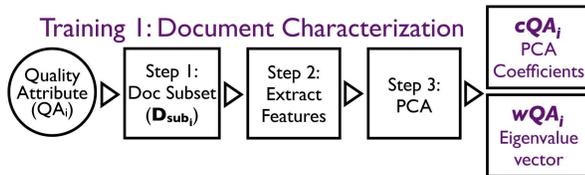
Figure 1 illustrates the image pipeline for a new document that needs a color workflow chosen. The document is first converted to a device independent space. Next, a feature vector is extracted from the document. The extracted features describe the key characteristics of the new input document. The selection engine uses the rules generated in the trainings to make a decision on which color workflow to apply.



**Figure 1.** The document is transformed to a device independent colorspace, CIECAM JCh. Next, the document features are extracted and sent to the selection engine. The selection engine uses the rules from Trainings 1 and 2 and the feature vector to make a decision on which color workflow to apply.

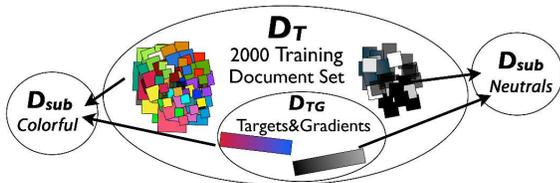
### Training 1: Document Characterization

Before starting the steps in training 1, we need to determine both the document training set  $D_T$  and the feature list. Figure 2 illustrates the three steps in training 1. First, a document subset  $D_{sub_i}$  is selected for each  $QA_i$ . Next, the feature vector is extracted from the  $D_{sub_i}$ . Then Principal Component Analysis (PCA) is used to project the training features in a new component space, where the eigenvalue vectors  $wQA_i$ , the PCA coefficients  $cQA_i$ , and the coordinates of the original features in the new coordinate system  $FQA_{sub_i}$ , are stored and used by the selection engine. These steps are repeated for each  $QA_i$ .



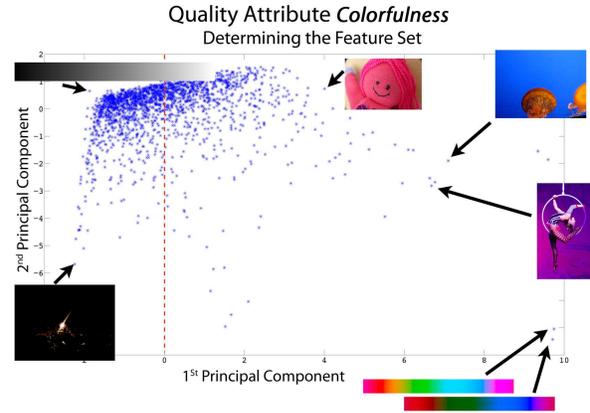
**Figure 2.** For each  $QA_i$  we manually found a subset of documents  $D_{sub_i}$  that have a specific  $QA_i$ . The feature vectors are extracted from the  $D_{sub_i}$ . PCA is used to project the extracted features into a new component space for the given  $QA_i$ . The  $cQA_i$  and the  $wQA_i$  are stored and used by the selection engine to determine the importance of the  $QA_i$  for a new document.

We have a training set of more than 2000 documents  $D_T$  that is used to generate the training rules. Most of the  $D_T$  were downloaded from the MIRFLICKR Database<sup>1</sup>, with the creative commons license from flickr. Additionally, the  $D_{TG}$ , used in training 2 are included in the  $D_T$  document set. From the  $D_T$  we manually selected a subset of documents that showed a dominance of the given  $QA_i$ , i.e. if the attribute was *shadow details* we found 100 documents that had a significant amount of dark pixels and shadow details, see Figure 3. In total, 8 document subsets  $D_{sub_i}$  were created with 100 documents in each.



**Figure 3.** 8 document subsets  $D_{sub_i}$  were manually selected for each of the  $QA_i$ , from the  $D_T$ , which included both complex documents and targets.

<sup>1</sup><http://press.liacs.nl/mirflickr/>



**Figure 4.** The  $D_T$  documents are plotted in the 1<sup>st</sup> and 2<sup>nd</sup> component space. The features in this test were: percentage of pixels out of gamut, mean chroma, busyness of the chroma channel, and percentage of pixels not in the 3 neutral 3D color bins (bins 1:3).

As illustrated in Figure 2, the next step in training 1 is to extract the  $D_{sub_i}$  feature vectors, but first we must determine the feature list. Our starting set of relevant features included:

- 1D histogram statistics of lightness and chroma: mean, median, standard deviation, skewness, kurtosis [4],
- 3D color histogram bins [4]: B&W, saturated, dark, light,
- Ratio of out of gamut pixels,
- Hasler[10] and Cui's[11] colorfulness,
- Spatial Frequency: entropy [12], average local range (MATLAB rangefilt), and busyness[13].

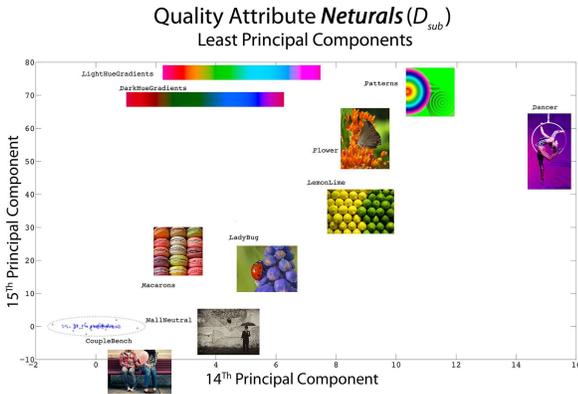
Given the large number of features to consider, reducing the feature vector dimensions is necessary, a method known as feature reduction [14] is applied to reduce the feature vector dimensions. The  $D_T$  was used to determine the final feature set. We wanted to separate each  $D_{sub_i}$  from the non subset  $D_{non_i} = D_T - D_{sub_i}$ . For each  $QA_i$ , the features that were able to separate the  $D_{sub_i}$  from the  $D_{non_i}$  were included in the final set of features, as illustrated in Figure 4. Our strategy was an iterative learning, which included the following steps:

1. For each  $QA_i$ , include all features that are likely to separate the  $D_{sub_i}$  documents from the  $D_{non_i}$ ,
2. PCA is used to determine the contribution of each of the features to the primary components,
3. The  $D_T$  features are plotted in the new space, if the  $D_{sub_i}$  are extreme points then this feature is successful, see Figure 4,
4. Features that did not separate the  $D_{sub_i}$  documents or did not contribute to the principle components were removed[4].

Figure 4 illustrates the features projected in the 1<sup>st</sup> and 2<sup>nd</sup> components of the  $QA$  colorfulness space. Features were added and removed until the *colorfulness*  $D_{sub}$  were the most extreme points or had the largest distance from the center of the data set. A final feature list was determined which included the following 15 features:

- 1D histogram lightness statistics: mean, standard deviation, median and skewness,

- 1D histogram chroma statistics: mean and standard deviation,
- 3D histogram color bins: 1 (lightness $\leq$ 40, chroma $\leq$ 20 and all hues), 2 (20<lightness $\leq$ 70, chroma $\leq$  20 and all hues), 3 (lightness $\geq$ 70, chroma $\leq$ 20 and all hues), sum of the very saturated bins (20-27: all lightness, chroma $\geq$ 50 and all hues) and the medium saturation bins(4-19: all lightness, 20 $\leq$ chroma<50 and all hues),
- Busyness for lightness and chroma,
- Entropy and local range,
- % of pixels out of gamut.



**Figure 5.** The 14<sup>th</sup> and 15<sup>th</sup> principal components of the neutral  $D_{sub}$  set and some of the  $D_{non}$  documents are plotted. The neutral  $D_{sub}$  features are circled, close to the origin. The  $D_{non}$  documents have a varying amount of neutrality.

The training 1 weights were determined with a PCA classification approach. For each  $QA_i$ , we calculated a weight of belonging  $wQA_i$  from the eigenvalue vectors for each  $D_{sub_i}$ . The  $D_{sub_i}$  features are projected using Matlab's princomp function. The function outputs the coefficients  $cQA_i$  used to transform the feature vector into the principal component space, the new coordinates of the training features  $FQA_{sub_i}$ , and the variance value that corresponds to each of the 15 (number of features) principal components  $\lambda_k$ , where  $k = 1, \dots, 15$  the index of each principal component. As  $k$  increases, the  $\lambda_k$  values decrease. For all of our  $D_{sub_i}$ , the final component variance was very small, illustrated in Figure 5. Since we are interested in finding the least amount of variance between the  $D_{sub_i}$  features, our training weighting vectors  $wQA_i$  are found with:

$$\vec{wQA_i} = 1 - \frac{|\lambda_k|}{|\lambda_{max}|}, \quad (1)$$

where  $\lambda$  is the vector variances of the projected features for the  $D_{sub_i}$ .  $\lambda_{max}$  is the maximum variance or the 1<sup>st</sup> principal component variance and  $k$  is the component index.

With this approach, we find the features that cause the least amount of variance for each  $D_{sub_i}$ . A new document is part of a group if, when transformed by the  $cQA_i$  coefficients, we find a very small value in the last component. If it has a very large value for the last component then this  $QA_i$  is not important. Figure 5

illustrates the clustering of the neutral (low saturation)  $D_{sub}$  compared to 10  $D_{non}$ , which have a varying amount of belonging to the neutral group (distance from the center of the data set).

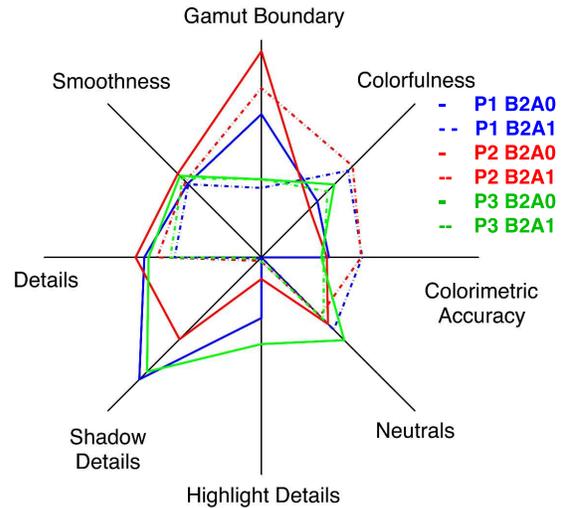
## Training 2: Color Workflow Assessment

**Table 1: There are three different profiles (P1, P2, and P3) and two rendering intents (perceptual B2A0 and media-relative B2A1) used in our work, six color workflows in total, for more details on the profiles see Color Workflow Details 10.**

Workflow Name	Profile	Rendering Intent	PCS
P1 B2A0	Océ P1	Perceptual B2A0	CIELAB
P1 B2A1	Océ P1	Media-Relative B2A1	CIELAB
P2 B2A0	Custom P2	Perceptual B2A0	CIE XYZ
P2 B2A1	Custom P2	Media-Relative B2A1	CIE XYZ
P3 B2A0	Onyx P3	Perceptual B2A0	CIELAB
P3 B2A1	Onyx P3	Media-Relative B2A1	CIELAB

In training 2 we compare an original set of documents to the output of the processed (color workflows applied) reproductions. For each original document we have  $j$  reproductions, where  $j = 1, \dots, 6$ , for each of the six color workflows. Training 2 uses the same  $QA_i$  to address the different ways to evaluate a color workflow's performance: *colorimetric accuracy*, *colorfulness*, *gamut boundary*, *smoothness*, *details*, *shadow details*, *highlight details*, and *neutrals*. The metric differences between the original and each reproduction are used to compare the performances of the color workflows, (see Table 1 and Figure 10 for details on the six color workflows). For a full description of the metric tests and target generation see [9].

## Quality Attribute: difference test results

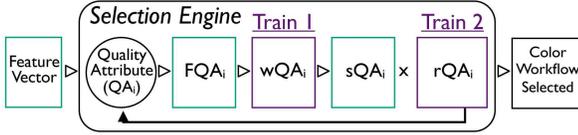


**Figure 6.** The results from training 2: the further from the origin the workflow performed. All color workflows were ranked first at least once, except P3 B2A1. The P2 B2A0 workflow was the most successful with the most number of attributes: gamut boundary, smoothness and details.

Table 2 reports the results from training 2. Figure 6 illustrates the workflows ranked each other. The scales between tests are not the same, differences should not be compared

**Table 2: The evaluation results of each  $QA_i$ . The bold purple values are the workflows that were most successful.**

Quality Attribute	P1 B2A0	P1 B2A1	P2 B2A0	P2 B2A1	P3 B2A0	P3 B2A1
Colorimetric Accuracy: CIE $\Delta E^*_{94}$ (Color Target)	8.26	<b>6.42</b>	8.38	6.48	8.68	8.55
Colorfulness: $\Delta Cui^*$ 's Colorfulness [11] (Color Target)	1.27	0.87	1.37	<b>0.81</b>	1.05	1.14
Gamut Boundary: $\Delta LCh^*_{STDV}$ [15] (Gradients)	0.66	0.32	<b>0.95</b>	0.78	0.37	0.36
Smoothness: $2^{nd}$ Derivative[16] (Gradients)	3.12	3.14	<b>2.73</b>	3.03	2.81	2.91
Details: VIF [17] (Complex Documents)	0.77	0.70	<b>0.79</b>	0.74	0.76	0.71
Shadow Details: $\Delta L^*_{STDV}$ [9] (Target $t \leq 20$ )	<b>3.98</b>	0.03	2.67	0.11	3.73	0.04
Highlight Details: $\Delta L^*_{STDV}$ (Target $\geq 80$ )	0.14	0.00	0.05	0.01	<b>0.20</b>	0.01
Neutrals: CIE $C^* \times \Delta h^*$ [9] (Target CIE $\Delta L^*$ )	5.00	2.64	2.83	3.09	<b>2.30</b>	2.97



**Figure 7.** The feature vector of the input document is transformed  $FQA_i$  and multiplied by the weights  $wQA_i$  from training 1, which results in the  $sQA_i$ . The  $sQA_i$  are then applied to the ranked scores from training 2,  $rQA_i$ , which gives the color workflow recommendation.

between axes. The profile with the largest area, has been the most successful with the largest number of attributes.

The results from training 2 illustrate that the color workflows are competitive and similar in quality. They were all ranked first for at least one assessment, except **P3 B2A1**. The largest difference between color workflows is the rendering intents. For most of the  $QA_i$  assessments, either the perceptual (B2A0) workflows were successful or the media-relative (B2A1) workflows. The B2A0 workflows performed better with details and the B2A1 workflows performed better with maintaining the colors. For the *colorimetric accuracy*  $QA$ , **P1 B2A1** was ranked first and **P2 B2A1** was a very close second. The **P2 B2A1** workflow was ranked first for the *colorfulness* attribute. **P1 B2A0** was most successful with the *shadow details*, while **P3 B2A0** was successful with the *neutrals* and *highlight details*. **P2 B2A0** did well with maintaining the *gamut boundary* differences, *smoothness* and the *detail*.

## The Selection Engine

Figure 7 illustrates the processing pipeline for a new document. First the feature vector of the new document is extracted and inputted to the selection engine. It is transformed into the component space by the  $cQA_i$  coefficients calculated in training 1. The new vector  $FQA_i$  is then element-wise multiplied by the eigenvalue vector  $wQA_i$ , which results in the score  $sQA_i$ :

$$sQA_i = \sum_{k=1}^{15} (|FQA_{i,k}| \odot wQA_{i,k}), \quad (2)$$

where  $sQA_i$  is the significance score of the  $i^{th}$   $QA$  for the new document. Equation 2 is also used to find the training scores  $sQA_{sub_i}$ , where the projected training features  $FQA_{sub_i}$  from training 1 are used.

Figure 8 illustrates the  $sQA_i$  percentage for a new set of documents  $D_V$ , using the  $wQA_i$  from training 1. Documents with  $sQA_i$  values larger than  $\mu$  were removed, given

$$\mu = \overline{sQA_{sub_i}} + (2.58 \times \sigma), \quad (3)$$

where  $\overline{sQA_{sub_i}}$  is the average  $sQA_{sub_i}$  and  $\sigma$  is the standard deviation of  $sQA_{sub_i}$ . This eliminates  $QA_i$  that are not present or have an insignificant presence from impacting the final selection.

We have applied the  $sQA_i$  to 18 verification documents  $D_V$ , displayed in Figure 8.  $D_V$  is also used in the observer evaluation test. The *face* document is the only one with all eight  $QA_i$ . The very dark document, *bridge*, only has one attribute, *shadow details*. Many documents have a large amount of the *details* attribute, shown in blue. Most of  $D_V$  have between five and six  $QA_i$ . The final color workflow is selected by:

$$S = \operatorname{argmin}_j \left( \sum_{i=1}^8 \left( \frac{1}{sQA_i} \otimes rQA_i^T \right)_{i,j} \right), \quad (4)$$

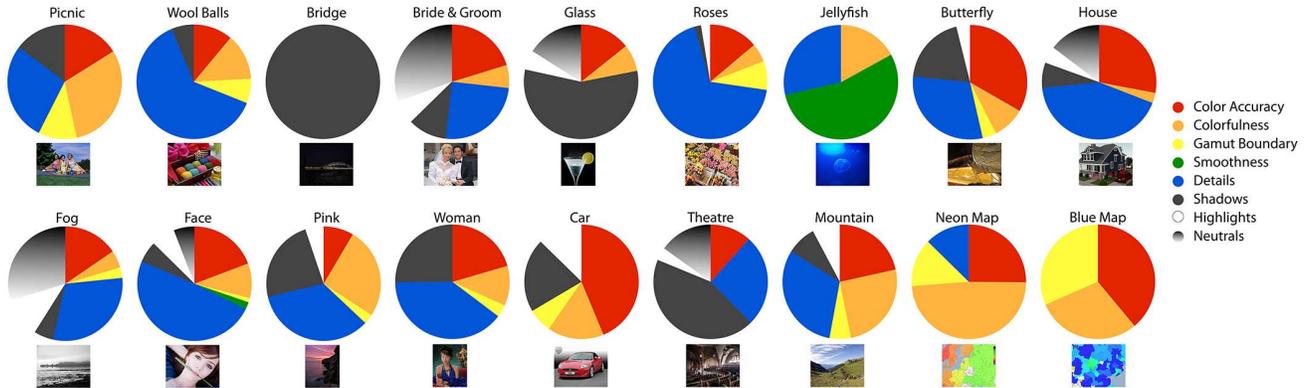
where  $S$  is the index of the selected color workflow and  $rQA$  is the ranking from the metric assessment in training 2. As an example, if it is determined from the rules in training 1 that a new document only has 2  $sQA_i$ : *shadow details* (60%) and *smoothness* (40%). The engine would choose **P3 B2A0**, which was ranked second for both and had the minimum sum of 200, as compared to 260 and 220 for **P1 B2A0** and **P3 B2A0** respectively. The selection engine results for the  $D_V$  are compared to the observer preferences in Table 3.

## Evaluating the Selection Engine

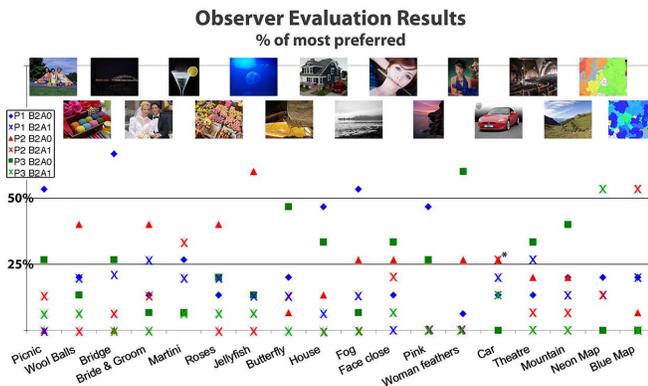
The selection engine was tested against an observer evaluation. All reproductions were printed on the Océ ColorWave 600 wide format printer with the LFM090 uncoated Océ Top Color Paper. For details on the six color workflows, (three profiles, **P1**, **P2**, and **P3**) and two rendering intents (perceptual B2A0 and media-relative B2A1), see Table 1. For each of the 18  $D_V$ , the six color reproductions were viewed simultaneously, at a distance of approximately 24 inches. The  $D_V$  is entirely different from the  $D_T$ . The observers were asked to choose the reproduction that they most preferred based on overall image quality. The color temperature of the room was approximately 5100K and the illumination was 560 lux. The 15 observers ranged in age, experience, ethnicity and gender. The results of the observer study are shown in Figure 9.

During the evaluation, the observers informally discussed the motivation behind their final choice. They were easily able to distinguish between two groups of documents that corresponded to the two rendering intents. Once they identified the two groups they chose one group and discarded the other. From the three remaining reproductions, the observers commented on having difficulty choosing just one reproduction and that they did not have a strong preference.

Table 3 compares the selection engine's choice to the ob-



**Figure 8.** The results of the training 1 weights applied to the  $D_V$  are illustrated. The color related attributes are in the warmer colors (red, orange and yellow), smoothness is in green, details are in blue, shadow details in black, highlight details in white and the neutrals are a gradient. Only the face document has all of the  $QA_i$  included, most have between five and six attributes. The bridge, jellyfish and blue map have the fewest number of  $QA_i$ , less than three.



**Figure 9.** The observer evaluation results are plotted. All color workflows were most preferred for at least one document, except **P1 B2A1**. With seven documents, the observers agreed on their most preferred reproduction. In general, the observers chose **P2** with the more colorful originals: wool balls, roses, jellyfish, car and blue map. **P1** was preferred with documents that had more shadow details and **P3** was preferred with documents that had more yellows, greens and neutrals.

servers' choice for the  $D_V$ , with consideration to each of our evaluation aims. An ideal system would always choose the same document as the majority of the observers, however there are other ways to evaluate the selection engine's performance. We aim to have our selection engine choose:

1. the same color workflow as the observers,
2. the same rendering intent as the observers,
3. the same color workflow, when the majority of observer's agree (> 50%),
4. one of the two most preferred workflows (observers first or second choice),
5. one of the preferred workflows (> 20%), when the observers do not have a strong preference (weak obs pref),
6. never select a 'bad' workflow (observer's never chose or least preferred).

The selection engine chose the same rendering intent as the observers for 89% of the  $D_V$ , selecting the exact color workflow

**Table 3:** the results of our evaluation aims show that the Selection Engine (SE) was successful in correctly choosing the rendering intent with 89% of the  $D_V$ . The engine chose the same color workflow as the observers for 33% of the documents. The selection engine chose either the observer's first choice or their second choice with 61% of the documents.

Evaluation	SE	# Doc	Percentage
1. Same color workflow	5	18	33%
2. Same rendering intent	16	18	89%
3. Same when a majority	4	7	61%
4. One of the two most preferred	10	18	56%
5. A preferred (weak obs pref)	7	11	64%
6. Chose bad workflow	0	18	0%

was more challenging. The selection engine chose the same color workflow as the observers for six documents, 33%. If the observer's second choice is included, the engine's success increased to 61%. When investigating the instances where the observers showed a strong consensus on their preference (more than 50%), the selection engine was successful with four out of seven documents, 57%. There were no occurrences where the engine chose a workflow that nobody chose or that was least preferred by the observers. The color workflow performances were competitive, five of them were most preferred for at least one of the  $QA_i$  in training 1 and most preferred by the observers, which makes automatically selecting one more challenging.

The instances where the selection engine did not choose the same color workflow as the observers, **P3 B2A0** was usually involved. The two documents that the selection engine chose **P3 B2A0** and the observers did not were *fog* and *bride & groom*. The selection engine chose **P3 B2A0** because of *details*, *shadow details*, and largely from the *neutrals QA*. With the *fog* document, the observers chose **P1 B2A0** which had a slight red tone. Although, the original was not available to the observers, it was a grayscale image. The results from training 2 ranked **P1 B2A0** the lowest in the *neutrals QA*, but the observers preferred the warmer reproduction.

When the observers chose **P3 B2A0** and the selection engine did not, the documents often had large regions of in-gamut, low to medium saturation of pink, yellow and green hues (*butterfly*, *face*, *woman* and *theatre*). When we compared the color differences

3 ICC v2 CMYK output profiles were chosen for this work, see Table 1. The profiles in the set were chosen because they were comparable in quality and a user of this printer would have access to them. Profile, **P1**, is available for download from Océ's media guide. **P2** was created with an Océ internal tool from the same measurement data that was used to make **P1**. The reproductions using Profiles **P1** and **P2** were all processed through the Océ Power Controller M+. **P3** is commercially available through the ONYX Driver and Profile DownloadManager and developed to be used with Onyx ProductionHouse.

Figure 10. Color Workflow Details

between the original and the reproductions, using the perceptual intent, **P3** was the only profile that changed these in-gamut colors. The dominant colors in these documents had a lower CIE  $L^*$  value and their hues shifted counter-clockwise, in the positive CIE  $a^*$  direction. As of yet our selection engine does not look at intentional hue shifts or other types of image enhancements. Our metrics compare the reproduction to the original. In the future we would like to consider using color dependent preference metrics, where there are intentional color shifts. For example, the *woman* document is dominated by skin-tones, the observers preferred the reproduction which darkened her skin and shifted it towards red.

A second consideration on why the selection engine did not choose **P3 B2A0** when the observers did, involves the weighing of the  $QA_i$ . **P3 B2A0** was often the most *colorful* of the three B2A0 workflows, but it was last in maintaining the *details*. When the observers were evaluating the workflows, once they chose the rendering intent, they then looked at the complimentary  $QA_i$ . If they chose B2A0 and their decision was primarily based on the *details*, they would then look at the *colorfulness* of the document or a  $QA_i$  that the B2A1 performed well with, instead of choosing the B2A0 workflow with the most *details* of the three. Following the observer's behavior and eliminating workflows in two steps, first the rendering intent and then make a final decision on which profile, may improve the selection engine's performance.

## Conclusion and Future Work

We have proposed a novel method of automatically selecting a color workflow based on the statistical content of the input document and a set of rules. Additionally, the decision rules do not rely on observer data, but rather on a set of metric performance tests. The workflow assessment and the document characteristic rules are both created with the same set of quality attributes. The performance of one workflow may be strong in one area but weak in another, this allows the document's properties to dictate which workflow should be applied.

For a first training, our results are promising. We have successfully determined which rendering intent to select. The final selection between profiles needs to be refined with consideration to the evaluation aims.

## References

- [1] International Color Consortium. ICC White Paper 7: The role of ICC profiles in a colour reproduction system. 23/03/11: www.color.org, Dec 2004.
- [2] International Color Consortium. ICC White Paper 9: Common Color Management Workflows and Rendering Intent Usage. 23/03/10: www.color.org, Mar 2005.
- [3] International Color Consortium. ICC White Paper 2: Perceptual Rendering Intent Use Case Issues. 23/03/10: www.color.org, Jan 2005.
- [4] Pei-Li Sun and Zhong-Wei Zheng. Selecting appropriate gamut mapping algorithms based on a combination of image statistics. In *Color Imaging X: Processing, Hardcopy, and Applications*, volume 5667, pages 211–219, San Jose, CA, January 2005. Proceedings SPIE IS&T.
- [5] Todd D. Newman, Timothy L. Kohler, and John S. Haikin. Dynamic gamut mapping selection, September 2005.
- [6] Asaf Golan and Hagit Hel-or. Novel workflow for image-guided gamut mapping. *Journal of Electronic Imaging*, 17(3):033004, July-Sep 2008.
- [7] Timothée Royer. Influence of image characteristics on image quality. Master's thesis, Ecole Nationale Des Sciences Geographiques and Gjøvik University College, France, Nov 2010.
- [8] Pei-Li Sun. *The Influence of Image Characteristics on Colour Gamut Mapping for Accurate Reproduction*. PhD thesis, University of Derby, Derby, UK, July 2002.
- [9] Kristyn Falkenstern, Nicolas Bonnier, Hans Brettel, Marius Pedersen, and Françoise Viénot. Using Metrics to Assess the ICC Perceptual Rendering Intent. In *Image Quality and System Performance*, Proceedings of SPIE/IS&T Electronic Imaging, San Francisco, CA, Jan 2011. SPIE.
- [10] S. Hasler and S. Susstrunk. Measuring colorfulness in real images. volume 5007, pages 87–95, 2003.
- [11] Chengwu Cui and Steven F. Weed. Colorfulness? in search of a better metric for gamut limit colors. In *PICS*, volume 3, pages 183–187, Portland, OR, March 2000.
- [12] Rafael C. Gonzalez, Richard E. Woods, and Steven L. Eddins. *Digital Image Processing Using MATLAB*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2003.
- [13] M. Orfanidou, S. Triantaphillidou, and E. Allen. Predicting image quality using a modular image difference model. In Susan P. Farnand and Frans Gaykema, editors, *Image Quality and System Performance V*, volume 6808, page 68080F, San Jose, CA, USA, Jan 2008. SPIE.
- [14] Amir Navot, Lavi Shpigelman, Naftali Tishby, and Eilon Vaadia. Nearest neighbor based feature selection for regression and its application to neural activity. In *Advances in Neural Information Processing Systems 18*, pages 995–1002. MIT Press, 2006.
- [15] Kristyn Falkenstern, Nicolas Bonnier, Hans Brettel, Marius Pedersen, and Françoise Viénot. Using Image Quality Metrics to Evaluate an ICC Printer Profile. In *18th Color Imaging Conference*, San Antonio, TX, Nov 2010. IS&T.
- [16] P. Green. *Color Management: Understanding and Using ICC Profiles*. Wiley-IS&T Series in Imaging Science & Technology, Feb 2010.
- [17] Hamid R. Sheikh and Alan C. Bovik. Image information and visual quality. 15:430:444, 2006.