https://doi.org/10.2352/ISSN.2470-1173.2021.16.COLOR-245

Banding defect detection and image quality classification [1]

Yi Yang [a], Eric Maggard [b], Yousun Bang [c], Minki Cho [c], Jan P. Allebach [a]
[a]. Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA.
[b]. HP Inc., Boise, ID 83714, USA.
[c]. HP Inc., Suwon City, South Korea.

Abstract

Banding has been regarded as one of the most severe defects affecting the overall image quality in the printing industry. There has been a lot of research on it, but most of them focused on uniform pages or specific test images. Aiming at detecting banding on customer's content pages, this paper proposes a banding processing pipeline that can automatically detect banding, identify periodic and isolated banding, and estimate the periodic interval. In addition, based on the detected banding characteristics, the pipeline predicts the overall quality of printed customer's content pages and obtains predictions similar to human perceptual assessment.

1. Introduction

Banding is one of the most difficult image defects which is a one dimensional, isolated or periodic, luminance and/or chromatic variation induced by the vibration of different printer components [1]. This defect was categorized under the macro uniformity image quality attribute in the paper [2], which evaluates the overall image quality of the printed image and considers banding as one of the most severe defects that affect the overall perceived image quality.

Much previous work has been done to study various aspects of banding defects, including defect stimulation, visual analysis, Fourier domain analysis, and banding reduction. A common approach to assess the visibility of banding is to conduct psychophysical experiments and collect the subject's perceptual evaluation [3, 4, 5]. Some methods analyzed banding in the Fourier domain, implemented in one dimension [6] or two dimensions [7]. Several works have also been done at the printer mechanism level to reduce banding, and have produced encouraging results for banding reduction [7, 8].

We draw inspiration from the banding detection work of Zhang et al. [9, 10, 11]. The algorithm that they developed can detect banding on printed uniform color pages, and can classify periodic and aperiodic banding. However, this work cannot detect banding defects on customer's content pages. We remedy this with a new design that allows automatic detection of banding defects on customer's content pages. Furthermore, we design a scheme for predicting the overall quality based on banding features and perceptual assessment, which yields promising predictions.

2. Methodology

The overall goal of our work is to automatically detect aperiodic and periodic banding on the printed customer's content pages and automatically analyze the print quality according to the banding severity.

As shown in Figure 1, the input is the master - test image pair. The master image is the digital original, and the test image is the scanned printed customer's content pages.

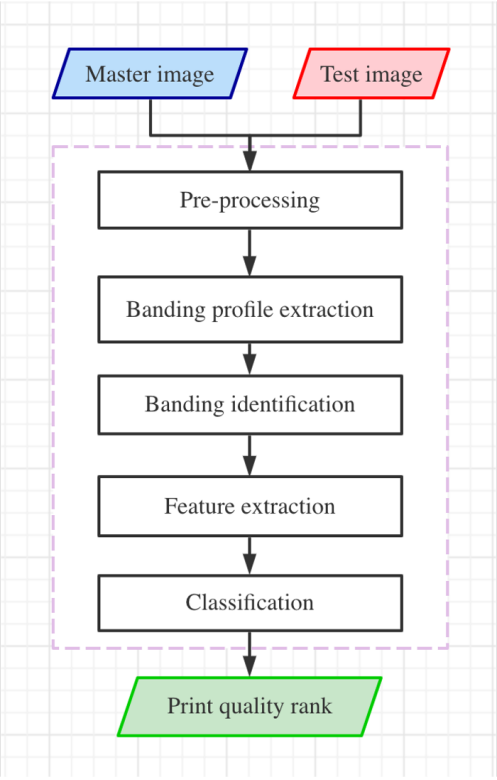

Figure 1. *Overall pipeline of banding detection and analysis.*

[1]Research supported by HP Inc., Boise, ID, USA

We first pass the master-test image pair into the pre-processing pipeline, which includes image registration, region of interest (ROI) extraction, and ROI alignment. Next, we apply banding profile extraction and banding identification to localize and identify banding defects. We extract meaningful features of banding defects, which can indicate the severity of the defects. Then we use machine learning models to classify these features into four categories from A to D (A symbolizes almost perfect print quality and D symbolizes the worst print quality. The ground truth is labeled by human subjects).

2.1 Pre-processing

Image registration is the first step of pre-processing, which aligns the test image with its master image, to eliminate the misalignment caused by the scanning process.

For our test images, we assume that the geometric transformation involves only a small skew angle rotation and a small translation along the x-axis and the y-axis, so we would like to find several matched key points to compute the best transformation. We convert images from RGB to grayscale to reduce the image dimension, and then resample the image using a 1/3 downsampling rate to save computation. Next, we apply histogram matching to the test image based on the master image to match the image gray values, and then we use Harris Corner Detection [12] to extract key points on the master image and test image, respectively. The feature descriptor is constructed using the 31 $\times$ 31 local areas surrounding each key point. By minimizing the sum of squared differences (SSD) and computing the ratio SSD, we find the best-matched key point pair. We then use the Maximum Likelihood Estimation Sample Consensus (MLESAC) [13] to take advantage of multiple matched point pairs. The transformation estimated by MLESAC is more robust than by random sample consensus (RANSAC) [14].

After obtaining the aligned test image, we would like to generate the object map, and divide the entire image into different ROIs according to the object map. The details of this part were explained in our previous work [15].

In this paper, all the following processes are implemented on the ROI, but they can also be implemented on the entire image.

2.2 Banding defect detection

In this section, we extract the banding profile, and then localize and identify banding based on the profile. Next, we refine and link the banding result, measuring the amount per unit area. For each band, we get its height, width, and distance, to describe the severity of the banding. In addition, periodic banding is identified and its repetitive interval is estimated.

2.2.1 Banding profile extraction

Figure 2 shows the steps of the banding profile extraction. First, we de-screen the input image or ROI to remove the halftone

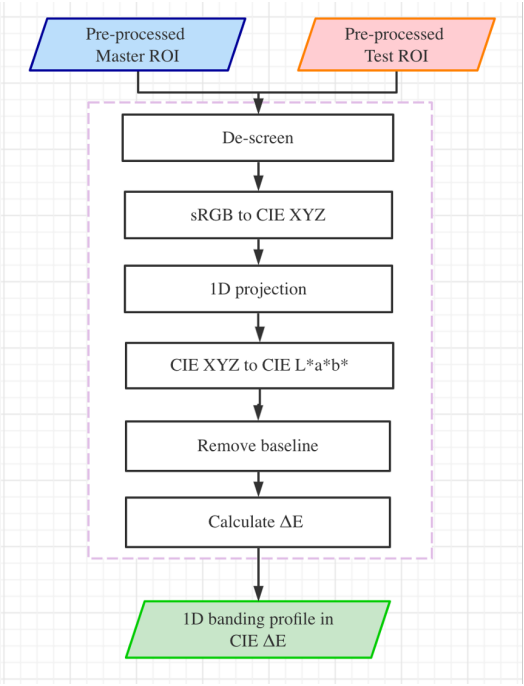

Figure 2. *Pipeline of banding profile extraction.*

patterns. A previous work [11] used a median filter with a window size of 0.05 inches (15 $\times$ 15 pixels for an image scanned at 300 dpi) to de-screen for uniform color pages and achieved good performance. But for our customer's content pages, after several tests, we found that a Gaussian filter with a size of 0.02 inches (7 $\times$ 7 pixels for our pages scanned at 300 dpi) is the best approach to do de-screening.

Then, we convert the image or ROI from the sRGB color space to the CIE1931 XYZ color space (hereinafter abbreviated as CIE XYZ). Next, in the CIE XYZ color space, the one dimensional projection of the image along the scan direction is computed by calculating the mean value of each line in this direction.

Since the banding may fade along the scanning direction, the 1D projection of the banding may become inconspicuous if we compute it across the entire image. Accordingly, we divide the image into three parts along the scanning direction and perform three separate projections, respectively. Independent analysis of each projection can yield more accurate results. For each part, we compute 1-D projections of the X, Y, and Z channels using the following formula:

$$projection_i[m] = \frac{1}{N}\sum_{n=1}^{N} image_i[m,n] \tag{1}$$

where M is the number of pixels in the process direction and N is the number of pixels in the scan direction, which is perpendicular to the process direction; and $[m,n]$ represents the coordinates along the process direction and the perpendicular direction, respectively.

Then, the 1-D projections of the X, Y, and Z channels are converted to the CIE 1976 $L^*a^*b^*$ color space (with a 2° observer

and D65 illumination). Following this, to make the banding more distinct from the image texture, we deduct its baseline from each 1-D projection signal. The baselines are obtained by applying a 1-D median filter with a size of 0.02 inches (7 pixels at 300 dpi) to each 1-D projection, respectively. Then, we compute the *CIE* ΔE and signed *CIE* ΔE using the three baseline-removed projections in CIE $L^*a^*b^*$ color space.

The formula for calculating *CIE* ΔE is:

$$CIE\ \Delta E = \sqrt{\sum_{c=L^*,a^*,b^*} (Original_c - Baseline_c)^2} \qquad (2)$$

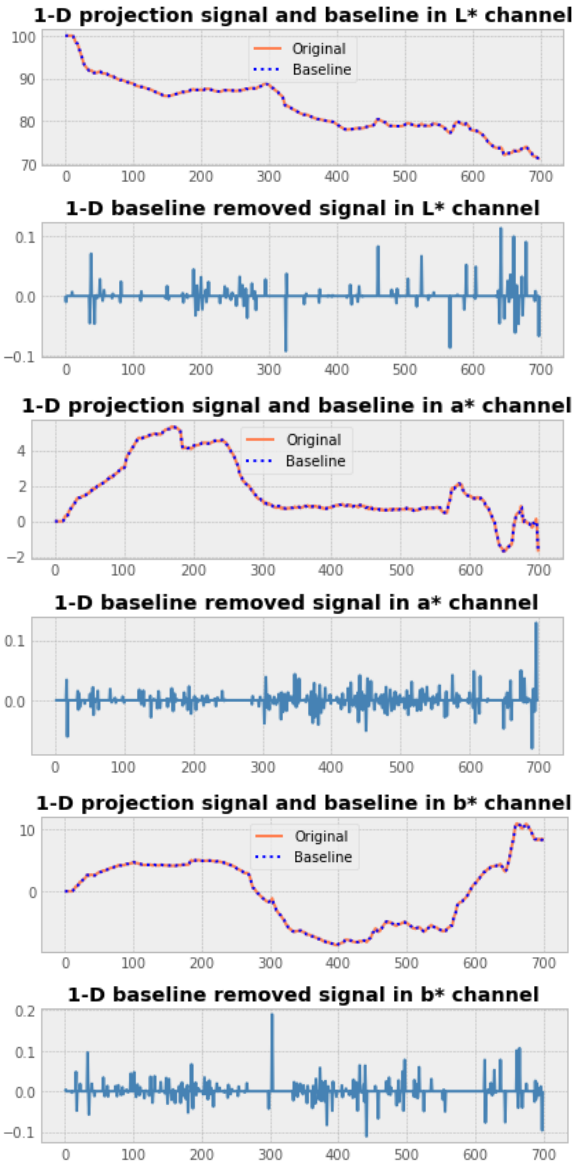

Figure 3. *1-D projection signals and baseline-removed signals.*

Since the L^* channel indicates the lightness information, for the signal of L^*, if it is higher than the baseline, the banding in this position is a light banding; otherwise, it is a dark banding.

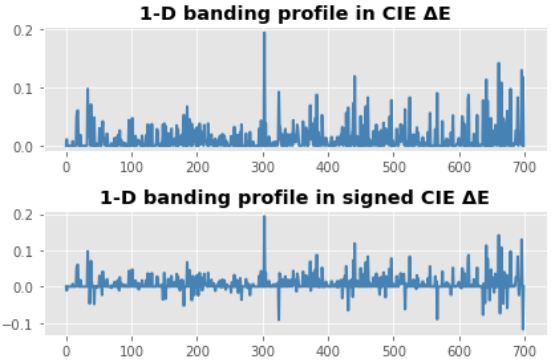

Figure 4. *1-D banding profile in CIE* ΔE *and signed CIE* ΔE *units.*

CIE ΔE can indicate the color difference in the CIE $L^*a^*b^*$ color space, but because it is non-negative, the light/dark information about the banding is lost. We calculate the signed *CIE* ΔE by multiplying *CIE* ΔE by the sign of the L^* channel baseline-removed projection, which makes the color difference also include lightness information. Figure 3 shows the 1-D projection signals and baseline-removed signals in the L^*, a^*, and b^* channels, respectively. Figure 4 shows the 1-D banding profile in *CIE* ΔE and signed *CIE* ΔE units.

Once we get the *CIE* ΔE, we need to find local maxima in *CIE* ΔE. A local maximum is found by comparing adjacent values. Once the local maximum is determined, its center point, height, and width will be calculated, as shown in Figure 5. The blue curved lines are signed ΔE. The red points signify the center position of each detected banding, which also is the local maximum value of the signed ΔE. The green vertical line is the height for each peak, which is obtained by computing a vertical distance between the peak and its lowest profile line. The two yellow lines represent the total width and the width at half the height of each peak.

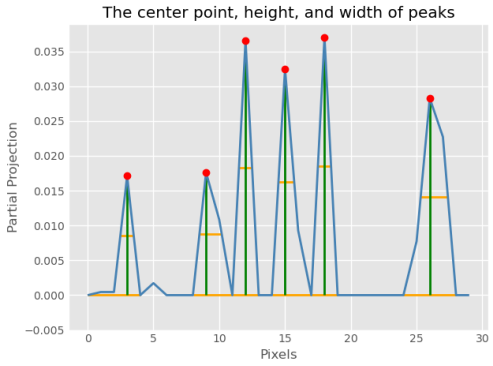

Figure 5. *The center point, height, and width of peaks.*

Then, we customized some thresholds for height, width, and distance between adjacent bandings according to the properties of the banding defect, and finally selected local maxima that met the threshold as bands. We set the threshold for local maxima according to a previous work [11] and our several experiments.

The thresholds for height, width and distance are shown in the table below and Figure 6 shows the peaks refined by the different thresholds.

Table1: Threshold for local maxima.

Height	$Mean_{height} + \frac{1}{2} Stddev_{height}$
Width	2 pixels
Inter-band Distance	5 mm (59 pixels at 300 dpi)

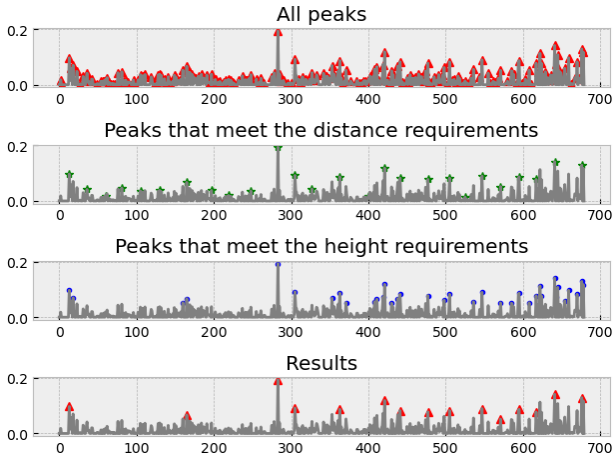

Figure 6. *Peaks refined by different thresholds.*

Image features may also lead to local maxima in the 1D projection. To distinguish the local maxima caused by banding defects from those caused by image features, we use the master image as a reference. Although we have performed a global image registration between the master image and the test image, we found that for a small ROI, there will still be a slight translation misalignment. Therefore, we use cross-correlation to locally register the test ROI and the master ROI pair, and calculate the signed offset of the test ROI relative to the master ROI in the process direction, which is denoted by o. Then, we extract banding profiles to obtain two sets of local maxima (master set and test set). Obviously, the master sets include all the local maxima caused by the image features, so we only select the local maxima in the test set but not in the master set. For each local maximum position $p_{test}[i]$ in the test ROI, if it is caused by image features, it must have a local maximum on the position $p_{master}[j]$ of the master ROI. The relationship between the two positions is

$$position_{test}[i] = position_{master}[j] + o \tag{3}$$

where i indicate i-th peak in the test set and j indicate j-th peak in the master set

Through this step, we can obtain the local maximum caused by banding defects, which will not be affected by the image features or content.

2.2.2 Repetitive interval calculation

Periodic bandings are an important print defect, which can help us figure out problems that occurred in an internal printer rotating component. To define periodic bands and estimate their repetitive interval, we use a search strategy to select periodic peaks from the isolated peaks detected in the previous section, estimate an approximate interval, and then optimize it until an accurate result is obtained.

We first compute all intervals between neighboring bands,

$$interval_i = Band\ Position_{i+1} - Band\ Position_i$$

and sort them by length. So we obtain a sorted array of lengths $[interval_1, interval_2, ..., interval_i]$. If there are periodic bandings in the test image, at least three isolated bands are identified as periodic bands. Thus the length of at least two intervals is the same or similar.

We set an upper bound ($interval \times 1.1$) and a lower bound ($interval \times 0.9$) for each interval of a different length, then repeatedly scan all intervals to find the bounded intervals, group them, and update their mean value as the new rough repetitive interval. This step is repeated until there is no change between the new rough interval and the old rough interval. After that, we use a new upper bound ($interval \times 1.05$) and a new lower bound ($interval \times 0.95$) of the updated interval, and re-scan all intervals to find all bounded intervals and update the mean value as the new interval. We compute the banding occurrence ratio (BOR) as expressed by the following formula (4) and select the interval with the largest BOR as the periodic interval.

$$BOR = \frac{n \times I}{h} \tag{4}$$

where n is number of intervals that fall in this bin, h is the height of the ROI and I is average value of the intervals in this bin.

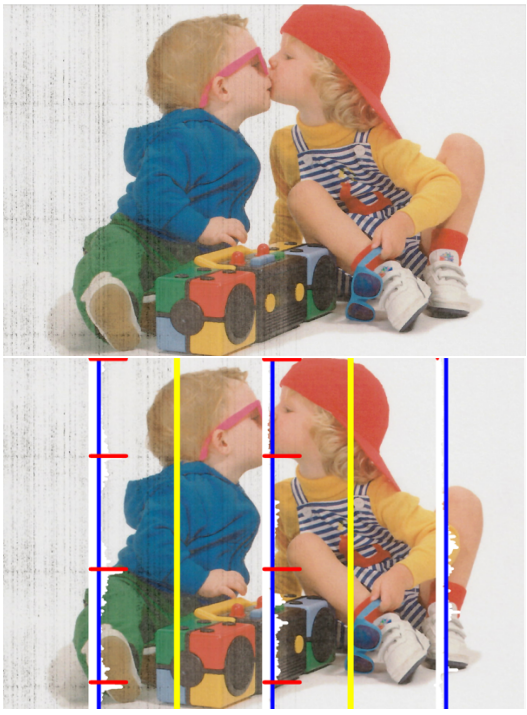

Figure 7. *A customers' content page and its banding detection results.*

Figure 7 shows the test image and its banding detection results. The yellow lines indicate the boundary between the three areas; the blue lines indicate the center lines of each area; the white projection around the blue line is the 1D projection of this area; the red short lines perpendicular to the blue line represent periodic bandings. We also calculated that the repetitive interval of the periodic bands of this image is 315.33 pixels (26.67 mm at 300 dpi).

2.3 Banding feature extraction and quality ranking classification

In the previous section, we detected banding in the image or ROI. In this section, we describe the extraction of banding features, which represent the severity of banding. Then, we develop classification models that can predict the print quality based on the severity of banding.

2.3.1 Data preparation and banding feature selection

Our test images are provided by HP Inc. All test images are scanned pages of printed images, with various print defects of varying severity. We select 800 images (each image has a size of 3100 $\times$ 2400 pixels and a resolution of 300 dpi), and after the aforementioned pre-processing, they are aligned with their master image and cropped into ROIs.

We set four levels from A to D to represent the print quality. A symbolizes almost perfect print quality, and the subsequent grades B, C, and D indicate that the banding defects are becoming more and more obvious and the print quality is getting worse and worse. Since we would like to use machine learning models to predict print quality based on perceptual assessment, three of our laboratory members labeled the ground truth manually. A set of ground truth samples is shown in Figure 8.

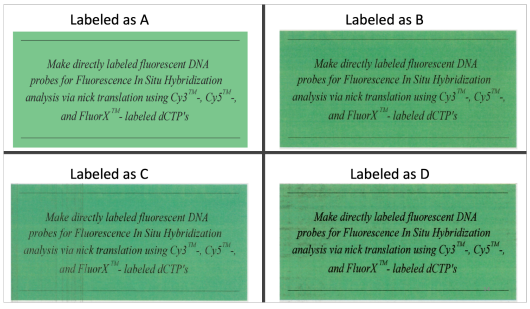

Figure 8. *A set of ground truth samples. The reader is advised to zoom in to see the banding defects.*

For the banding defects, our pipeline can automatically detect the total number of bands in each ROI, and the position coordinate, length, width, and prominence of each band. Among them, the prominence shows the "local protrusion" of each band. It is defined as follows: let B denote the current peak, A denotes the peak immediately to the left of B, and C denotes the peak immediately to the right of B. Let h_i, $i = A, B, C$ denote the height of each of these peaks, then the prominence for peak B is defined as

$$prominence_B = h_B - \max(h_A, h_B, h_C) \tag{5}$$

The following table shows the components of the detectable banding feature vector.

Table 2: Components of the detectable banding feature vector for an ROI.

1.	The total number of bands
2.	The absolute mean value of signed delta E
3.	The standard deviation of signed delta E
4.	ROI area
5.	The average height
6.	The maximum height
7.	The average width
8.	The maximum width
9.	The average prominence

As detailed above, we have calculated various features, which are sufficient to define the shape and state of the banding. However, for machine learning models, too many input features often lead to model overfitting.

The ANOVA F-statistic [16] is a popular feature selection technique. In [17], researchers found that the traditional ANOVA F-statistic is proper for selecting features for classification modeling problems where the inputs are numeric and the predictor outputs are categorical. For our model, all banding features are numeric, and we would like to obtain four levels (from A to D) of predictors to show the print quality. Therefore, our classification modeling problem contains the proper inputs and outputs to use the ANOVA F-statistic for feature selection.

We use the Python scikit-learn library function f-classif to implement feature selection. This function computes the ANOVA F-statistic of the given inputs and returns the feature importance score based on the F-statistic score and p-value of each feature. First, we randomly split our images into 60% training images, 20% validation images, and 20% test images, and pass them into our banding detection pipeline to compute feature vectors. The validation image set is used for feature selection and model selection. The test data is used to analyze the performance of the final model.

Figure 9 shows the feature importance score of each input feature. This clearly shows that Feature 1 (number of bands), Feature 5 (the average height), Feature 6 (the maximum height) and Feature 8 (maximum width) are the most important. In addition, Feature 2, Feature 4, and Feature 9 also have slightly higher scores. In order to decide the number of features most suitable for the model, we use logistic regression with an L1 penalty to iteratively estimate the accuracy of choosing 3 to 9 features. For the same validation data, the accuracy of selecting 3 features is 71.43%, the accuracy of selecting 4 features is 76.19%, and the accuracy of selecting 5 features is as high as 80.95%. The accuracy goes back to about 70% if 6 and subsequent higher numbers of features are selected. Therefore, we chose the top 5 fea-

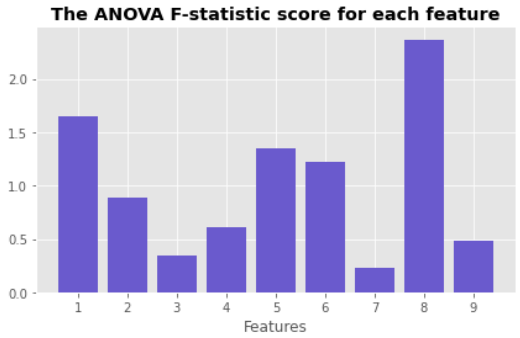

Figure 9. *The ANOVA F-statistic score for each feature.*

tures ranked according to the importance score, which are the total number of bands in an ROI, the absolute mean value of signed delta E, the average banding height, the maximum banding height, and the maximum banding width within an ROI.

2.3.2. Classification models

We use four popular machine learning models to build predictors: logistic regression, support vector machine, KNN and random forest. There are two heuristics to use binary classification for multi-class classification: One-vs-Rest (aka OvR, One-vs-All or OvA), and One-vs-One (OvO) [18].

OvR means dividing a multi-class classification into multiple binary classification problems. Then, we train a binary classifier for each binary classification problem, and for each data sample, use the most confident model to make the prediction. For our case, the multi-class classification problem can be divided into four binary classifications, as shown below:

Binary classification problem 1: A vs [B,C,D]

Binary classification problem 2: B vs [A,C,D]

Binary classification problem 3: C vs [A,B,D]

Binary classification problem 4: D vs [A,B,C]

OvO also splits the multi-class classification into binary classification problems, but uses a one-to-one strategy. Compared with the OvR, OvO requires more binary models. The formula used to calculate the number of binary models is as follows:

$$\#\ of\ binary\ models = \frac{NumClasses(NumClasses–1)}{2} \tag{6}$$

So, our problem with four types: A, B, C, and D, can be divided into several binary classifications as follows:

Binary classification problem 1: A vs B

Binary classification problem 2: A vs C

Binary classification problem 3: A vs D

Binary classification problem 4: B vs C

Binary classification problem 5: B vs D

Binary classification problem 6: C vs D

Each binary classifier predicts a class label. When we input the test data into the classifier, we will get the final result based on the majority count of the outputs of all the binary classifiers.

2.3.3 Results

We mentioned that we use a random splitting tool to randomly split the data sets into 60% training data, 20% validation data, and 20% test data.

For skewed classes, accuracy cannot sufficiently indicate the performance of the model. Two supplementary evaluation scores, balanced accuracy [19], and F1 score [20] are computed based on the test data, to avoid over-performance estimation of unbalanced data sets.

$$Accuracy = \frac{TP+TN}{Total\ samples} \tag{7}$$

$$Balanced\ Accuracy = \frac{1}{2}\left(\frac{TP}{TP+FN} + \frac{TN}{TN+FP}\right) \tag{8}$$

$$F_1 = \frac{TP}{TP+\frac{1}{2}(FP+FN)} = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{9}$$

$$\begin{cases} Precision & = \frac{TP}{TP+FP} \\ Recall & = \frac{TP}{TP+FN} \end{cases} \tag{10}$$

where, TP = number of true positives;
TN = number of true negatives;
FP = number of false positives;
FN = number of false negatives.

The following table shows the models we selected and the corresponding accuracy, balanced accuracy, and F1 results. The best decision is made by the random forest. We conclude that random forest is based on ensemble learning and introduces randomness, so that the model can handle outliers and solve the problem of data imbalance, thus it better adapts to our data. In addition, all the prediction results on the test dataset have notable performance, which can help us quickly inference to predict the print quality.

Table 3: Score results for various classifiers.

Model	Accuracy	Balanced accuracy	F1
LR(OVR)	0.81	0.75	0.77
SVM (OVO) (RBF)	0.82	0.78	0.80
SVM (OVR) (RBF)	0.82	0.78	0.80
KNN	0.86	0.83	0.85
Random Forest	0.91	0.85	0.90

3. Conclusion

In this paper, we developed a banding processing pipeline and print quality classifier to diagnose printer defects and evaluate print quality. Our pipeline includes banding profile extraction, periodic banding interval estimation, and print quality classification. Our results show that banding defects can be automatically detected, the periodic interval of periodic banding can be estimated, and higher accuracy, balanced accuracy, and F1 score can be obtained based on the banding feature vectors.

References

[1] Robert Chung and Matthew Rees, "A survey of digital and offset print quality issues", Rochester: Printing Industry Center at RIT, vol. PICRM-2006-04, 2007.

[2] D. René Rasmussen, Edul N. Dalal and Kristen M. Hoffman, "Measurement of macrouniformity: Streaks, bands, mottle and chromatic variations", *IS&T* PICS Conference Proceedings, 2001.

[3] Peter D. Burns, Jonathan B. Phillips and Don Williams, "Adapting the ISO 20462 softcopy ruler method for online image quality studies" *IS&T* Electronic Imaging, 2013.

[4] Elena A. Fedorovskaya and Huib De Ridder, "Subjective matters: from image quality to image psychology" *IS&T* Electronic Imaging, 2013.

[5] Yi Yang, Utpal Sarkar, Isabel Borrell and Jan P. Allebach, 2020. "Inkjet quality ruler experiments and print uniformity predictor" *IS&T* Electronic Imaging, 2020.

[6] Pei-Ju Chiang, Nitin Khanna, Aravind K. Mikkilineni, Maria V. Ortiz Segovia, Sungjoo Suh, Jan P. Allebach, George T.-C. Chiu and Edward J. Delp, "Printer and scanner forensics," Signal Processing Magazine, IEEE, vol. 26, no. 2, pp. 72–83, 2009.

[7] Ahmed H. Eid; Mohamed N. Ahmed, Brian E. Cooper and Edward E. Rippetoe, "Characterization of electrophotographic print artifacts: Banding, jitter, and ghosting," Image Processing, IEEE, vol. 20, pp. 1313–1326, 2011.

[8] Cheng-Lun Chen, George T.-C. Chiu, Jan P. Allebach, "Banding reduction in electrophotographic process using human contrast sensitivity function shaped photoreceptor velocity control," Journal of Imaging Science and Technology, vol. 47, pp. 209–223, 2003.

[9] Jia Zhang, Stephen Astling, Renee Jessome, Eric Maggard , Terry Nelson, Mark Shaw and Jan P. Allebach. "Assessment of presence of isolated periodic and aperiodic bands in laser electrophotographic printer output" *IS&T* Electronic Imaging, 2013.

[10] Jia Zhang and Jan P. Allebach. "Estimation of repetitive interval of periodic bands in laser electrophotographic printer output" *IS&T* Electronic Imaging, 2015.

[11] Jia Zhang. An Investigation of Print Quality Defects: Psychophysical Evaluation of Content Masking, Development of Web-based Troubleshooting Tools, and Analysis of Sharp Roller Bands, Ph.D. Dissertation, Purdue University, West Lafayette, IN, 2016.

[12] Jie Chen J, Li-hui Zou, Juan Zhang and Li-hua Dou. "The Comparison and Application of Corner Detection Algorithms". Journal of Multimedia, vol. 4, pp. 435-441, 2009.

[13] Philip HS. Torr and Andrew Zisserman. "MLESAC: A new robust estimator with application to estimating image geometry." Computer Vision and Image Understanding, vol. 78, pp. 138-156, 2000.

[14] Konstantinos G. Derpanis, "Overview of the RANSAC Algorithm," Image Rochester NY, 4(1), pp. 2-3, 2010.

[15] Runzhe Zhang, Yi Yang, Eric Maggard, Yousun Bang, Minki Cho and Jan P. Allebach, "A comprehensive system for analyzing the presence of print quality defects", *IS&T* Electronic Imaging, 2020.

[16] Henry Scheffe, The Analysis of Variance, Wiley Classics Library, Wiley-Interscience Publication, New York, NY, 1999.

[17] Max Kuhn, and Kjell Johnson, Feature Engineering and Selection: A Practical Approach for Predictive Models. Chapman Hall/CRC Data Science Series, Boca Raton, FL, 2020.

[18] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar, Foundations of Machine Learning. MIT Press, Cambridge, MA, 2018.

[19] Digna R. Velez, et al, "A balanced accuracy function for epistasis modeling in imbalanced datasets using multifactor dimensionality reduction." Genetic Epidemiology: the Official Publication of the International Genetic Epidemiology Society, 2013.

[20] David M.W. Powers, "Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation." arXiv preprint arXiv:2010.16061, 2020.

Author Biography

Yi Yang received her B.S in Geomatics Engineering from Wuhan University in 2013 and M.S in Geomatics Engineering from Chinese Academy of Sciences in 2016. She is currently working on a Ph.D. in Electrical and Computer Engineering at Purdue University. Her primary area of research has been image processing, computer vision and machine learning.

Eric Maggard received his B.S. degree in Physics from Northwest Nazarene University, Nampa, Idaho in 1991 and the M.S. degree in Computer Science specializing in image analysis and processing from Walden University in 2006. He is an Expert Imaging Scientist in the LaserJet Hardware Division and has developed programs and image quality algorithms for the last 15 years that are used in the testing of LaserJet print and scan image quality.

Minki Cho is an engineer with HP Printing Korea. He received B.S.(1997) and M.S.(1999) in electrical engineering from the Inha University, Korea. From 2003 to 2017, he worked for Samsung Electronics and Samsung Advanced Institute of Technology. His research areas are print image processing, print image quality diagnosis and calibration.

Yousun Bang is a manager of Image Quality Part in the Imaging Lab at HP Printing Korea Co. Ltd. She received her BS and MS in mathematics from Ewha Womans University, Seoul, Korea in 1994 and 1996, and her Ph.D. in the School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana in 2005. She worked for Samsung Advanced Institute of Technology and Samsung Electronics Company from 2004 to 2017.

Jan P. Allebach is Hewlett-Packard Distinguished Professor of Electrical and Computer Engineering at Purdue University. Allebach is a Fellow of the IEEE, the National Academy of In-

ventors, the Society for Imaging Science and Technology (IS&T), and SPIE. He was named Electronic Imaging Scientist of the Year by IS&T and SPIE, and was named Honorary Member of IS&T, the highest award that IS&T bestows. He has received the IEEE Daniel E. Noble Award, the IS&T/OSA Edwin Land Medal, the IS&T Johann Gutenberg Prize, and is a member of the National Academy of Engineering.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

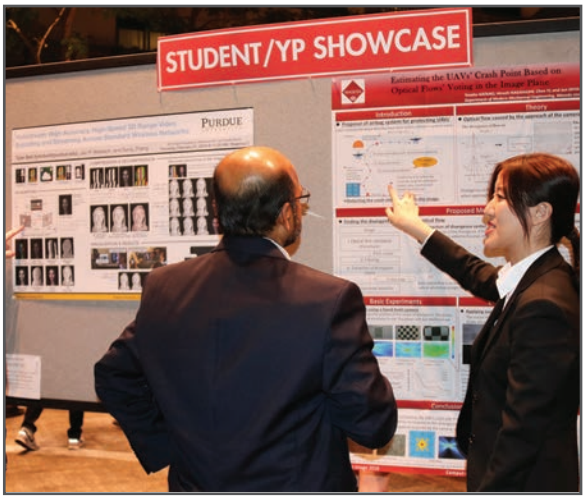

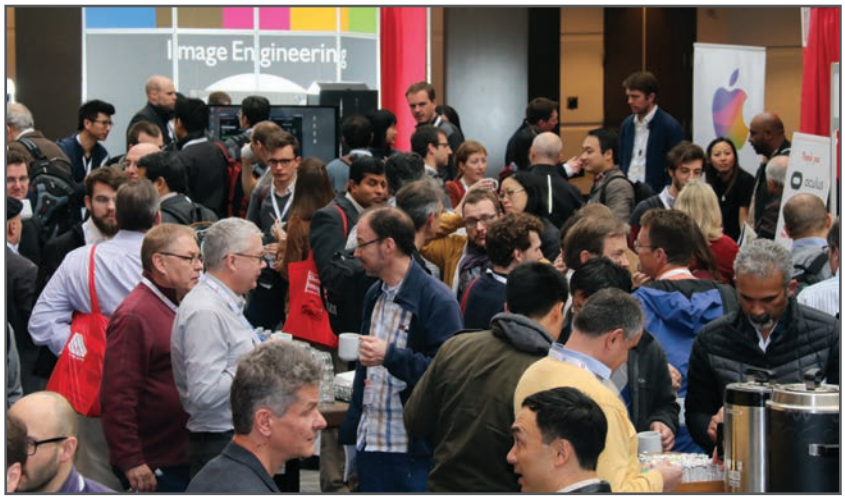

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org