Interval HSV: Extracting Ink Annotations

John C. Femiani Anshuman Razdan
Arizona State University
Polytechnic Campus, 7001 E. Williams Field Rd.
Bldg. 140, Mesa, AZ 85212, USA
john.femiani@asu.edu

Abstract

The HSV color space is an intuitive way to reason about color, but the nonlinear relationship to RGB coordinates complicates histogram analysis of colors in HSV. We present novel Interval-HSV formulas to identify a range in HSV for each RGB interval. We show the usefulness by introducing a parameter-free and completely automatic technique to extract both colored and black ink annotations from faded backgrounds such as digitized aerial photographs, maps, or printed-text documents. We discuss the characteristics of ink mixing in the HSV color space and discover a single feature, the upper limit of the saturation-interval, to extract ink even when it is achromatic. We form robust Interval-HV histograms in order to identify the number and colors of inks in the image.

1. Introduction

Paper documents may be photocopied, faxed, reviewed, folded, or marked with hand drawn annotations. For example reviewers may mark corrections onto a physical copy of a journal paper, or planners may circle important regions of an aerial photograph. Unfortunately, annotations make it very difficult to run image processing routines on the images when they are later digitized. We believe that existing document processing methods have the potential to improve their accuracy at recognizing text in images if annotations drawn in ink are first extracted to a separate image so that they can be analyzed individually.

This paper is motivated by issues we faced while using color histograms in hue, saturation, and value (HSV) to identify common types of ink used for annotation. We considered alternative color spaces such as CIELAB or YUV, but chose to analyze colors using HSV because we interpret our input colors as RGB *reflectances* and HSV has special invariance properties with respect to our ink mixing model (section 3). Furthermore, the input RGB indices fill a color-cube which the HSV hexcone was designed to fit perfectly. The fundamental problem is that uniform RGB digitization errors are not uniform after mapping to HSV coordinates. This can cause gaps in histograms as illistrated in Fig. 1. To address this we present a new *interval* approach to using HSV that not only leads to a simple, intuitive, robust, and fast density estimation approach using simple histograms, but it also reveals a powerful new attribute, the *saturation upper limit*, that can discriminate both chromatic and achromatic inks from background colors in many images. The utility of interval HSV is demonstrated with a simple histogram based image segmentation that is more appropriate than state of the art techniques for color image segmentation such as mean shift [1] and other similar methods when applied to the problem of ink annotation segmentation. This low level color based approach provides a *language agnostic* and a fast method to separate handwriting, stamps, decorations, and other artifacts from faded document images, so optical character recognition (OCR) techniques can be applied to a single ink-layer (presumably, of machine-print).

Intervals in HSV are used to overcome numerical issues which have complicated histogram analysis [2, 3, 4, 5]. We make the following contributions:

1. We discover a new feature, the saturation-interval upper limit, that illustrates the merit in treating discrete colors as intervals by reducing the complexity of extracting both chromatic and achromatic inks.

2. We present a robust way to compute joint HSV or just HV histograms. This allows individual

978-1-4244-3991-1/09/$25.00 ©2009 IEEE

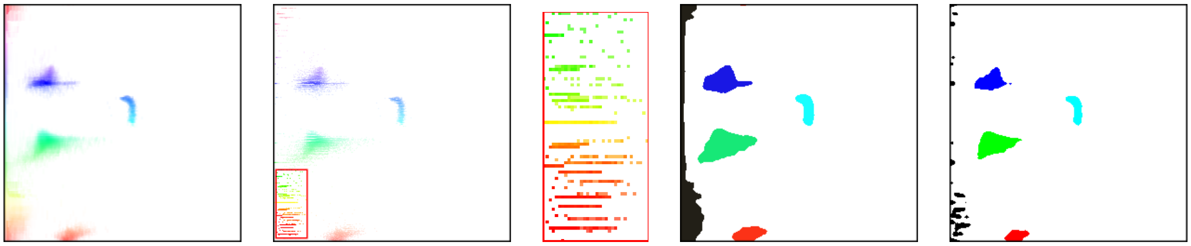

(a) (b) (c) (d) (e)

Figure 1. Histograms of an image with five pens, including one black, are shown. In each the vertical axis is hue and the horizontal is value. (a) was computed with interval arithmetic, (b) was computed without. Both were low-pass filtered and thresholded to form (d) and (e). (c) is a zoomed view of the portion of (b) within the red rectangle. Observe that (e) has split the black pen into several disconnected regions.

inks to be identified using connected decision regions of the histogram.

3. We employ decisions which are based predominantly on color at each pixel rather than structure [6, 7] and we do not require complex or iterative feature space clustering techniques [8, 9, 10].

1.1. Problem

Given a digitized mixed image $\mathcal{M}$ such that each pixel is a subtractive mixture of an underlying background image U and a partly transmissive ink $P_1, P_2, ..., P_N$, or no ink P_0, our aim is to deduce the number of inks (N) and decide which if any of the N pens influence each pixel. Output is a set of labels $\mathcal{L}$ so that according to the mixing model discussed in section 3.3 the mixed color is $M_i = P_{\ell_i} U$. Fig. 2 provides examples of valid input images along with pseudo-color output images.

1.2. Scope and Assumptions

We use an ink mixture model in section 3.3 that assumes a uniform-thickness for the layer of ink and does not allow inks to scatter or reflect light. Inks which have significant reflectance, such as grease-pens or chalk, are outside the scope of this paper and may or may not be identified correctly. Furthermore, we presume an achromatic background layer, but we extend this to other backgrounds with simple preprocessing. In practice RGB values are sometimes quantized to $\widehat{R}\widehat{G}\widehat{B}$ indices and stored so that they are *device dependent*, but since accurate color profiles for the scanner are generally not available, we treat the RGB values as estimates of reflectance. Backgrounds with linear structures and high saturation are outside the scope of this paper, and will cause the proposed algorithm to under-segment certain inks.

2. Prior Art

Color segmentation approaches include clustering or thresholding colors directly in RGB or YUV [7, 11], treating colors in unstable regions of HSV differently [3, 12], or using error propagation techniques with variable kernel estimation [13, 2]. The last approach is most similar to our approach, but interval arithmetic has advantages because we can identify ranges in HSV even though the transform is not differentiable.

Various authors have used clustering, quantization, and thresholding or partitioning techniques. The *serialized k-means* algorithm [9] uses a sliding window of pixels and applies k-means algorithm to feature-vectors in a combined R, G, B, H, S, L color-space. Each time the window is advanced, the k-means algorithm is used to search for new clusters; starting with the centers of the previous clusters as seeds.

Comaniciu et. al. cluster colors using the *mean shift* algorithm [1]. Mean shift is a hill-climbing approach to locate the center of the nearest region of maximum density, or *mode*. This approach uses a combined feature space of L*, u*, v* and spatial x, y coordinates so that spatial proximity is taken into account while clustering. Unlike the proposed interval HSV approach, the mean shift algorithm requires an iterative gradient descent operation at each point, and it requires a costly nearest-neighbor query at each iteration. Commaniciu does several things to avoid this, including an additional vector quantization step. Unlike mean shift the proposed approach does not require nearest-neighbor queries or multiple iterations per pixel, so it scales well and is simple to implement.

Gevers and Stokman [13] use error propagation to recognize objects based on robust histograms of color invariants. For each measure RGB a normal distribution is identified in the desired color space. Burns and Berns [5] provide detail on this kind of error propaga-

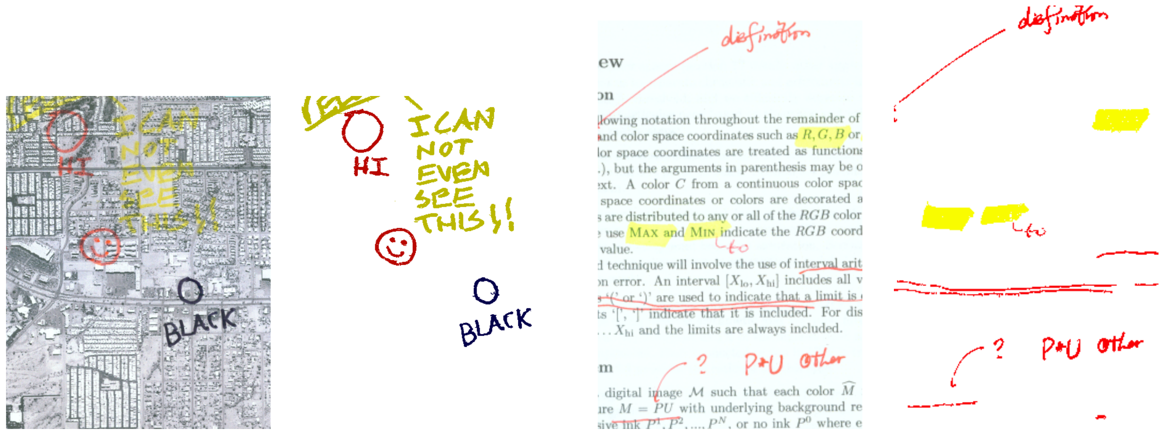

(a) (b) (c) (d)

Figure 2. Examples of the input (a,c) and pseudocolored output (b,d) that one can expect from the proposed method. The colors in (b,d) are the mean colors of the pens in (a,b) with contrast adjusted to increase visibility. Image (a) is an annotated aerial photo of Phoenix, Arizona (printed onto transparency, obtained from www.usgs.gov), and (b) is an annotated revision of this paper.

tion analysis; but it requires differentiable transform functions, and the mapping to HSV is defined piecewise with the derivatives that are not continuous. The proposed approach overcomes this limitation using interval arithmetic, which allows us to precisely track the range in terms of hard limits so we can threshold on S_{hi} (section 3.3).

3. Method

Our approach has the following major steps:

1. Preprocess the image in order to reduce the saturation of large, homogeneous areas of color that are not likely to contain annotations while keeping the inked areas saturated.
2. Form a histogram of saturation interval upper-limits.
3. Form a 2D value × hue (HV) histogram of annotated pixels using interval arithmetic to cope with singularities, and threshold the histogram at 1% of the highest peak. Each connected region in HV is treated as a distinct ink.
4. Label regions of the image based on pixel HSV representations. Morphological operations recover pixels otherwise missed in the HV histogram.

The proposed method is simple because it can be implemented using arrays for histograms, and does not need complex spatial indexing techniques or nearest-neighbor algorithms in order to estimate a probability density function or find its modes. The approach is almost parameter free; the only parameters controlling the results are the 1% threshold used in step 3, and possibly the filter-size used for preprocessing in step 1 (section 3.1).

3.1. Preprocessing

Documents are subject to various processes which may desaturate and brighten, the images. The proposed segmentation technique works best when ink marks are darker or more saturated than the background image, but this is not always the case. To allow our method to work on these images, we start by adjusting the brightness before processing by looking for the darkest pixel in the image and subtracting its value from all other pixels.

Sometimes the fading effect is not isotropic in the RGB coordinates, causing the image to have a tinted effect. Some images, while faded, still have low frequency variation in chroma. In our implementation we pre-process these images by rescaling RGB coordinates so that below a certain frequency the image is unsaturated. Let F be a low-pass filter applied to the original image $\mathcal{M}$. Then the preprocessed image $\mathcal{M}'$ is the result of

$$\mathcal{M}' = \mathcal{M} \times \frac{\mathtt{Value}(\mathcal{M} * F)}{\mathcal{M} * F}$$

where the multiplication operator $(\times)$ is done element by element, and the $\mathtt{Value}(\cdot)$ function returns the arithmetic mean of R, G, and B. In our experiments we use a box filter with a square window of $4\times\sigma+1$, where σ is a rough estimate of the expected pen thickness at the image's resolution.

3.2. Computing HSV from RGB

In this section we provide the formulas to transform discretized $\widehat{R}\widehat{G}\widehat{B}$ indices into an interval of HSV indices. The well known transformation from RGB to HSV [14]works for real values, but in practice our colors are not real numbers. Instead they are quantized to indices that are positive scaled integers with a scale of L, where rounding truncates so that $\widehat{R} = \min(L-1, \lfloor RL \rfloor)$ and L is 256 for images with 8 bit channels. This mapping transforms half-open intervals to quantization indices. The error added by quantization causes significant problems when mapping these colors to HSV because division by small numbers is involved. The problem is particularly evident with *unsaturated colors* and *dark colors* such as those produced by a black or brown pen.

In order to base our decisions on histograms using HSV coordinates in sections 3.3 and 3.4 we must solve this issue with discretization. Fig. 1 shows the kind of problem we can expect if errors are not addressed.

The quantization operation $RGB \mapsto \widehat{R}\widehat{G}\widehat{B}$ is not invertible, but its preimage[1] is a set of half-open intervals $[R], [G], [B]$ of the form $[R] = [R_{lo}, R_{hi}) = \left[\frac{\widehat{R}}{L}, \frac{\widehat{R}+1}{L}\right)$. We find intervals $[H], [S], [V]$ using interval arithmetic to adapt the standard formulas [14], and we deal explicitly with regions where the mapping $RGB \mapsto HSV$ is undefined by setting them to the entire range of $[0, 1]$. Intervals $[\textsc{Max}]$ and $[\textsc{Min}]$ are set to one of $[R], [G], [B]$ depending on which of $\widehat{R}, \widehat{G}, \widehat{B}$ is the maximum or minimum.

In interval HSV, the value interval is

$$[V] = [\textsc{Max}] = \left[\frac{\widehat{\textsc{Max}}}{L}, \frac{\widehat{\textsc{Max}}+1}{L}\right) \qquad (1)$$

We divide by $[\textsc{Max}]$ to find $[S]$, but division by an interval which crosses 0 is undefined so we replace $[S]$ by $[0, 1]$ whenever $[\textsc{Max}]$ includes 0.

$$[S] = \begin{cases} [0,1] & \widehat{\textsc{Max}} = 0 \\ \left(\frac{\widehat{\textsc{Max}}-\widehat{\textsc{Min}}-1}{\widehat{\textsc{Max}}+1}, \frac{\widehat{\textsc{Max}}-\widehat{\textsc{Min}}+1}{\widehat{\textsc{Max}}}\right) & \text{else.} \end{cases} \qquad (2)$$

[1] In this context we are talking about the preimage or image of a *function*, not a digital image or *picture*.

The hue interval $[H]$ involves another interval called $[D]$ where

$$[D] = \frac{1}{L}\left(\widehat{\textsc{Max}} - \widehat{\textsc{Min}} - 1, \widehat{\textsc{Max}} - \widehat{\textsc{Min}} + 1\right)$$

When we divide by $[D]$, if $[D]$ crosses 0 the result is forced to $[0, 1]$

$$[H] = \begin{cases} [0,1] & \text{if } D_{\text{lo}} \le 0 \\ \left(\frac{\widehat{G}-\widehat{B}-1}{6LD_{\text{hi}}}, \frac{\widehat{G}-\widehat{B}+1}{6LD_{\text{lo}}}\right) & \text{else if } \widehat{\textsc{Max}} = \widehat{R} \\ \left(\frac{\widehat{B}-\widehat{R}-1}{6LD_{\text{hi}}}, \frac{\widehat{B}-\widehat{R}+1}{6LD_{\text{lo}}}\right) + \frac{1}{3} & \text{else if } \widehat{\textsc{Max}} = \widehat{G} \\ \left(\frac{\widehat{R}-\widehat{G}-1}{6LD_{\text{hi}}}, \frac{\widehat{R}-\widehat{G}+1}{6LD_{\text{lo}}}\right) + \frac{2}{3} & \text{else.} \end{cases}$$

3.3. Thresholding the Saturation

According to the Kubelka Munk theory of reflectance an ink with no reflectance and uniform thickness alters a background color according to the familiar formula for subtractive color mixing $M = PU$ where M is the reflectance in the mixed image before quantization, P is reflectance of a pen's ink on a white background, and U is the reflectance of the underlying material. In this section we introduce a theorem that justifies the use of S_{hi} to separate ink from a grayscale background.

Lemma 3.1. *The value of a color mixed with a subtractive ink is always less than or equal to the value of the achromatic background* ($\texttt{Value}(M) \le \texttt{Value}(U)$).

Proof. Since P is a reflectance it is constrained to the interval $[0, 1]$ and $\textsc{Max}(PU) \le \textsc{Max}(U)$.

□

Lemma 3.2. *If an achromatic background $U > 0$ then the saturation of a mixed color is equal to the saturation of the pen used to mark it, $S(M) = S(P)$, and is invariant with respect to the intensity of the achromatic background material.*

Proof. Since U is a nonzero reflectance it is constrained to the interval $(0, 1]$, and U has equal R, G, and B because it is achromatic. Therefore

$$S(M) = \frac{\textsc{Max}(M) - \textsc{Min}(M)}{\textsc{Max}(M)} = \frac{\textsc{Max}(P)U - \textsc{Min}(P)U}{\textsc{Max}(P)U} = S(P)$$

□

Lemma 3.3. *If an achromatic background $U > 0$ then the hue of a mixed color is equal to the hue of the pen used to mark it, $H(M) = H(P)$, and is invariant with respect to the intensity of the achromatic background material.*

Proof. Although H is piecewise defined, for each case the hue H is always a constant plus a fraction with a linear combination of R, G, B in the numerator and denominator. The ratios are invariant to U because it is a common factor and the constant part of equation for hue depends only on which of R, G, B is maximal and is therefore invariant to multiplication by a positive constant U.

□

Definition 3.4. *A pair of inks P_i and P_j are* ***separable*** *based on color space coordinate X by a threshold τ if the domain of X can be partitioned into two disjoint intervals $[X_i]$ and $[X_j]$ with a common bound at τ and $M = P_i U$ if and only if $X(M) \in [X_i]$ and $M = P_j U$ if and only if $X(M) \in [X_j]$.*

We offer the following theorem for achromatic inks.

Theorem 3.5. *If achromatic inks P_i and P_j are separable based on V_{lo}, then they are separable based on S_{hi}.*

Proof. Any achromatic ink satisfies $\text{MAX} = \text{MIN}$ by definition, so $S_{hi} = \frac{1}{V_{lo}}$ according to equations (2) and (1) . Therefore for any τ, if $V_{lo} < \frac{1}{\tau}$ then it follows that $S_{hi} > \tau$.

□

The most important pen to identify is the 100% transparent background-pen P_0 because it occupies the majority of the pixels. Theorem 3.5 suggests that we identify a threshold τ so we can decide that pen P_0 influences all colors where $S_{hi} < \tau$ is true. We choose Otsu's threshold selection technique to identify a threshold $\widehat{\tau}$ using a histogram of discretized $\widehat{S}_{hi}$. Fig. 3(b) shows results on images with chromatic and achromatic inks. Note that it is often successful even for the black pen. We improve further on this result by region growing in section 4.

3.4. Grouping by Hue and Value

Section 3.3 identified colors that were unmixed (i.e. mixed with P_0). In this section we use interval HSV to form a $2D$ joint value and hue histogram and distinguish between the remaining pens. If inks are designed so that a human can distinguish between them, then it is very unlikely that two different pens will have similar hue and value. Eliminating the colors that are rare in any pen (e.g. less than 1% of the most frequent color) leaves remaining colors in connected regions of the $\widehat{H}\widehat{V}$ plane that can be attributed to a single pen.

Consider two joint HV-histograms $\mathbf{g}$ and $\mathbf{h}$. Let $\mathbf{g}_{i,j}$ be the measured number of colors that are *not* in the background, with $i = \widehat{H}$ and $j = \widehat{V}$. Let $\mathbf{h}_{i,j}$ count the expected number if we assume colors are uniformly distributed in the discrete interval $[\widehat{H}] = \lfloor H_{lo} \rfloor \ldots \lceil H_{hi} \rceil$, which includes i. In the $\mathbf{h}$ formulation each color $\widehat{M}$ contributes equally to each histogram bin in the intervals $[\widehat{H}]$ and $[\widehat{V}]$. The difference between these two histograms is evident in Fig.1. Notice that $\mathbf{h}$ should have only a few well defined peaks corresponding to the colors of the inks used in the image. Without interval arithmetic these peaks are divided into many disconnected pieces. This is undesirable because connected component labeling can not be applied to the histogram in order to count the number of pens.

We expect it to be rare but not impossible that pens can produce the same colors, so we threshold $\mathbf{h}$ by $0.01 \times \max(\mathbf{h}_{i,j})$and then label connected components to identify N decision regions $\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_N$ for N pens so that we decide $M = P_k U$ if the projection of $\widehat{M}$ onto the HV plane intersects $\mathcal{R}_k$. When we find connected components we take care to recognize that the H axis is cyclic. Fig. 1(c) shows the identified components colored according to the average color (with the average computed in RGB).

4. Generating Output

Output is a corresponding raster of labels $\mathcal{L}$ such that each $M_{x,y} = P_{\ell_{x,y}} U$. In this section we describe how a simple iterative region growing technique is used to set the labels. Our region growing process starts with a seed labeling $\mathcal{L}^0$ and iteratively changes labels to produce a sequence $\mathcal{L}^1, \mathcal{L}^2, \ldots \mathcal{L}^\infty = \mathcal{L}$. Notationally we use $\mathcal{L}^k$ for the label set after k iterations.

We start by defining the seed labels $\mathcal{L}^0$ so that $\ell^0_{x,y} = i$ if the HVprojection of $\widehat{M_{x,y}}$ is in $\mathcal{R}_i$. For each pen P_i, we define the expected mixed color as the average color of the mixed pixels in region $\mathcal{R}_i$. In each subsequent iteration $\mathcal{L}_{k+1}$, unmarked pixels are given the same label as their most similar adjacent pixel in the image according to $\mathcal{L}_k$, with similarity measured using the

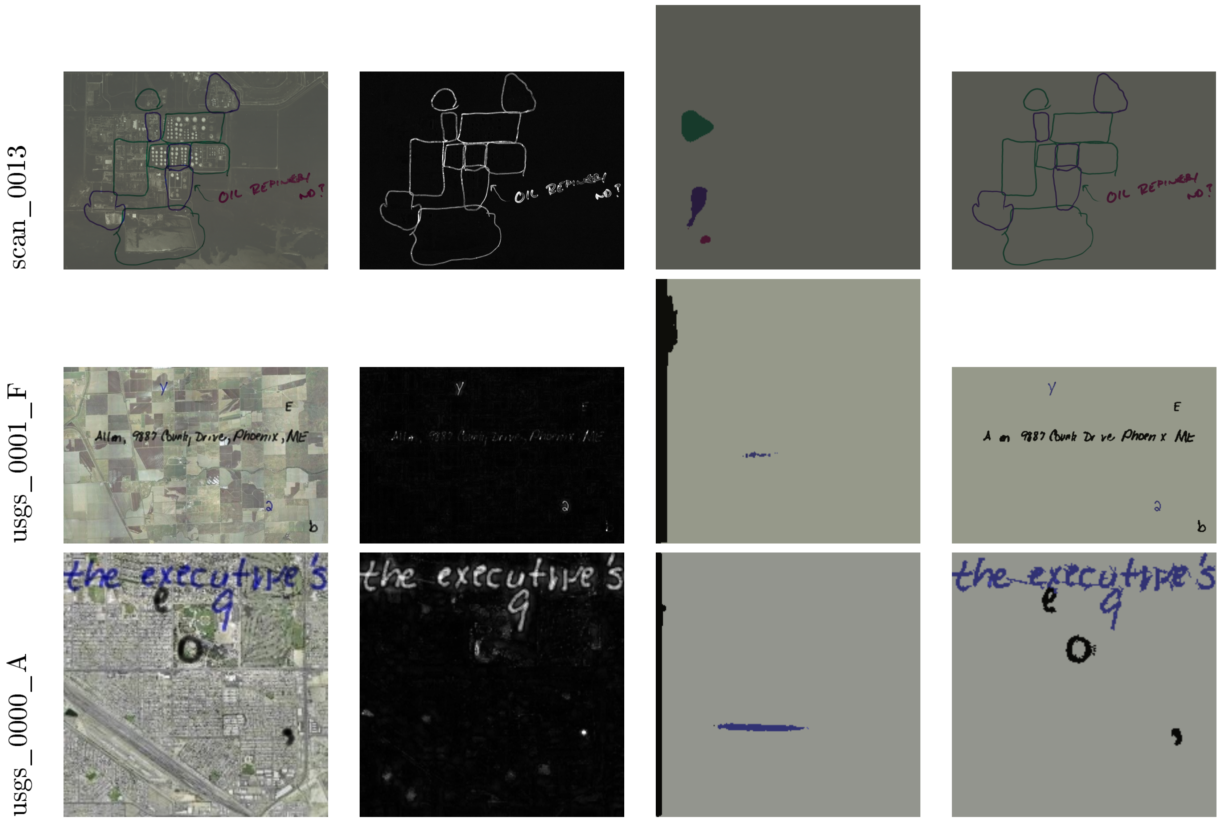

(a) (b) (c) (d)

Figure 3. Before and after examples of the proposed ink-extraction technique. Each row shows (a) the original image, (b) the saturation upper limits, (c) the value/hue histogram, and (d) the final segmented image.

Euclidean norm in RGB space for simplicity[2]. The process converges when $\mathcal{L}^k = \mathcal{L}^{k+1}$.

5. Evaluation

We used two sets of images for testing, the first was generated from images from the on-line Teraserver managed by the United States Geological Survey (USGS), and the second is a dataset provided by the National Geospatial Agency. The USGS dataset is a synthetic dataset generated using color background images acquired from the USGS terraserver, with handwriting annotations from the Unipen [15] dataset digitally composited using the Kubelka Munk mixing equations and various pigment settings to simulate blue and black soft-tipped markers.

The results on the USGS data set were evaluated manually by counting the number of symbols in increments of $\frac{1}{2}$, to allow for partially segmented symbols, that were correctly extracted from the image. We show some problem images with explanations in Fig. 4. The results of the proposed algorithm on the USGS data-set are summarized in Table 1.

Table 1. IHSV on the USGS dataset.

Black	Blue	All Pens
$\frac{91}{103}$	$\frac{80}{81}$	$\frac{171}{184} = 93\%$

The NGA data set is a group of 44 aerial photographs with hand drawn annotations made by an analyst. The images are RGB scans of originally grayscale films with colored and black ink annotations. In addition to four basic colors of red, green, blue, and black ink, these images also include a tiny amount of purple ink, yellow ink, and an additive mark that appears to have been made by a pink grease pen. The images are much larger than the USGS dataset, and because these images are authentic rather than digitally composited they exhibit a number of more varied artifacts. The results of our algorithm on the NGA dataset are summarized in Table 2.

[2] We used the RGB space to measure distances rather than attempting to use a perceptual distance measure in order to make the algorithm simpler, and also because inks most directly change the reflectance of the image rather than its perceived color.

Table 2. IHSV on NGA Dataset.

Red	Green	Blue	Black	Other	All
$\frac{407.5}{436}$	$\frac{239}{240}$	$\frac{234}{247}$	$\frac{133}{166}$	$\frac{11}{16}$	$\frac{1024.5}{1105} = 93\%$

From these experiments we conclude that the proposed approach has a recall of about 93% for extracting ink from aerial photographs, and that the recall is close to 80-88% when ink is black, although in practice it will probably be closer to 80% than 88%. If black ink is excluded, the recall is 95-99%. Fig. 3 shows typical input and output images from both datasets. The numbers in Tables 1 and 2 are based on the ability to identify the inked portions of the image *and* group similar colored inks properly. When black ink was misclassified as blue or red, as in Fig. 4(a,b,c), we did not include it in the table even though we were able to extract the marks.

The proposed approach is difficult to compare with popular color segmentation algorithms such as mean shift or serialized k-means because those algorithms focus on clustering the colors effectively. Our proposed interval IHSV representation of colors tackles an orthogonal issue, i.e. which features to base the decision on. Since the mean shift algorithm [1] tends to treat the $L^*u^*v^*$ color space without special attention to chroma we find that it can fail to distinguish background variation from foreground variation as we show in Fig. 5.

6. Conclusions

We have presented a solution to the handwritten annotation extraction problem over monochromatic or achromatic backgrounds. With preprocessing, we have extended this algorithm to work on background that are color aerial photographs. Our solution produces accurate results when subtractive pigments are used, and when achromatic inks are separable from the background by some threshold. Our main contributions are 1) interval arithmetic to make analysis of HSV distributions tenable, and 2) a single feature, S_{hi}, to separate both achromatic and chromatic pigments from the background. We are also the first that we know of to 3) address the segmentation of transmissive colored handwriting over photographic images.

References

[1] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *TPAMI*, vol. 24, no. 5, pp. 603–619, 2002.

[2] P. S. S. Romaní and E. Montseny, "On the reliability degree of hue and saturation values of a pixel for color image classification," in *FUZZ*, May 2005, pp. 306–311.

[3] S. Sural, Q. Gang, and S. Pramanik, "Segmentation and histogram generation using the hsv color space for image retrieval," in *ICIP*, vol. 2, 2002, pp. 589–592.

[4] T. Gevers, "Robust histogram construction from color invariants," in *ICCV*, vol. 1, 2001, pp. 615–620.

[5] P. D. Burns and R. S. Berns, "Error propagation analysis in color measurement and imaging," *CR&A*, vol. 22, no. 4, pp. 280–289, 1997.

[6] Y. Liu, S. Goto, and T. Ikenaga, "A robust algorithm for text detection in color images," in *ICDAR*, 2005, pp. 399–405.

[7] C. Rother, V. Kolmogorov, and A. Blake, ""grabcut": interactive foreground extraction using iterated graph cuts," *SIGGRAPH*, vol. 23, no. 3, pp. 309–314, 2004.

[8] U. Garain, T. Paquet, and L. Heutte, "On foreground-background separation in low quality color document images," *IJDAR*, vol. 8, pp. 47–63, 2006.

[9] Y. Leydier, F. L. Bourgeois, and H. Emptoz, "Serialized k-means for adaptative color image segmentation," in *DAS VI*, 2004, pp. 252–263.

[10] T. Perroud, K. Sobottka, H. Bunke, and L. Hall, "Text extraction from color documents - clustering approaches in three and four dimensions," in *DAR*, 2001, p. 937.

[11] Y. S. J. W. H.D. Cheng, X.H. Jiang, "Color image segmentation: Advances and prospects," *PR*, vol. 34, no. 12, pp. 2259–2281, 2001.

[12] C. Zhang and P. Wang, "A new method of color image segmentation based on intensity and hue clustering," *ICPR*, vol. 3, p. 3617, 2000.

[13] T. Gevers and H. Stokman, "Robust histogram construction from color invariants for object recognition," *TPAMI*, vol. 26, no. 1, pp. 113–117, 2004.

[14] A. R. Smith, "Color gamut transform pairs," in *SIGGRAPH*. ACM Press, 1978, pp. 12–19.

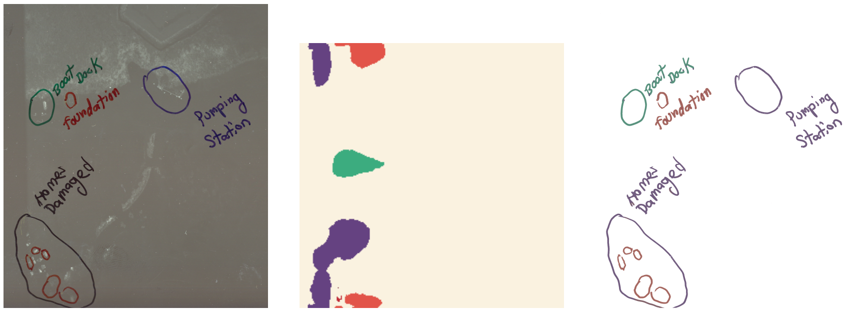

(a) (b) (c)

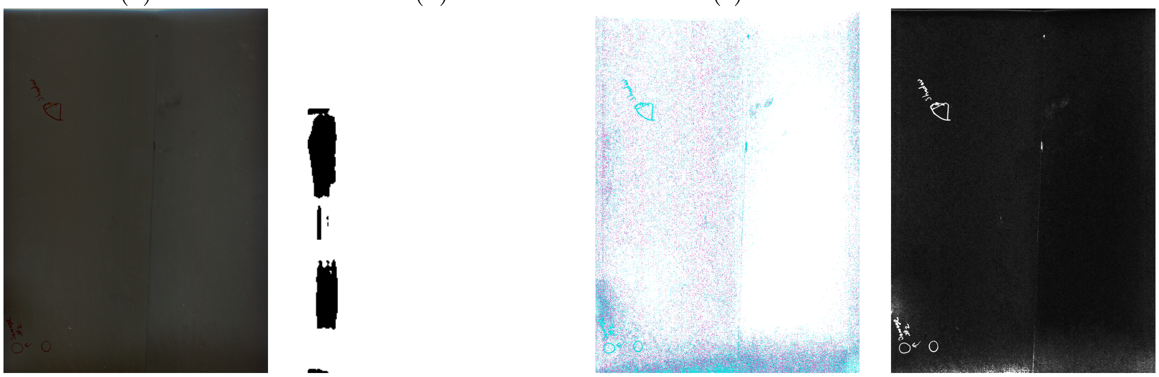

(d) (e) (f) (g)

Figure 4. Examples of poor performance. (a) from the NGA dataset (75% recall) with fairly well defined inks, (b) Components from the value/hue histogram, black and purple inks merged, (c) the resulting ink-colors image, with 0% recall for black ink only. (d) The worst case image (0% recall). This is a dark and noisy image with little text. (e) the components identified from the histogram, the ink cluster merges with the background blobs (f) the pen-labels rendered in random colors, and (g) the saturation upper limit for each pixel. The background noise is scattered but frequent enough to skew the histogram.

(a) (b) (c)

Figure 5. A comparison of our approach with accelerated the version of mean shift from [1], (a) is the original image, (b) is the result of mean shift with spatial bandwidth = 7, color bandwidth = 6.5, and minimum region = 20, and (c) is the result of our method.

[15] I. Guyon, L. Schomaker, R. Plamondon, M. Liberman, and S. Janet, "Unipen project of on-line data exchange and recognizer benchmarks," in *ICPR*, vol. 2, 1994.