

Dating Historical Color Images

Frank Palermo¹, James Hays², and Alexei A. Efros¹

¹ Carnegie Mellon University, Pittsburgh PA 15213, USA

² Brown University, Providence RI 02912, USA

{fpalermo, efros}@cs.cmu.edu, hays@cs.brown.edu

Abstract. We introduce the task of automatically estimating the age of historical color photographs. We suggest features which attempt to capture temporally discriminative information based on the evolution of color imaging processes over time and evaluate the performance of both these novel features and existing features commonly utilized in other problem domains on a novel historical image data set. For the challenging classification task of sorting historical color images into the decade during which they were photographed, we demonstrate significantly greater accuracy than that shown by untrained humans on the same data set. Additionally, we apply the concept of data-driven camera response function estimation to historical color imagery, demonstrating its relevance to both the age estimation task and the popular application of imitating the appearance of vintage color photography.

1 Introduction

The problem of determining the era in which a photograph was taken is important to various areas of study. Historians, genealogists, and conservationists tasked with the preservation of photographic artifacts may all find knowledge of a photograph's age to be useful. For this reason, there is a rich body of literature suggesting various *manual* approaches to the dating of photographs. Works targeted primarily at genealogists ([21, 22, 26]) typically advocate the study of visual clues within the photograph, particularly the fashion and hair styles worn by human subjects. Domain-specific knowledge can be quite useful here, e.g. Storey [23] focuses entirely on the dating of military photographs. In contrast, works targeted primarily at conservationists tend to advocate study of the photographic artifact itself as opposed to image content. For example, Messier [14] notes that exposing a photographic print to ultraviolet light may yield clues about the print's age by revealing the presence of optical brightening agents in the paper. We note in particular a reference volume by Wilhelm and Brower [28] which provides a comprehensive overview of the aging of color photographs and motion pictures as well as guidelines for their long-term preservation.

Although the emergence of Internet imagery and metadata has enabled progress on novel computer vision tasks, historical photographs were poorly represented online until fairly recently. In January 2008, Flickr and The Library of Congress began a pilot project to make publicly-held historical photography collections



Fig. 1. Color photography has been in widespread use for over 75 years, and various factors have caused its appearance to continuously evolve during that time. Observe the striking differences in the “looks” of these four photographs taken in (from left): 1940, 1953, 1966, and 1977.

accessible online. Known as “The Commons” [30], the collection has greatly expanded since that time: as of this writing, there are 56 cultural heritage institutions participating and the Library alone shared more than 16,000 historical photographs. In spite of this, there has to date been little computational effort to analyze this increasing quantity of data. Kim et al. [12] modeled the temporal evolution of topics in collections of Flickr images, but only over a comparatively short time span of several recent years. The only prior work regarding our task of historical photograph dating on a *decades-long* time span is by Schindler et al. [19, 20] who describe methods for temporally ordering a large collection of historical city-scape images by reconstructing the 3D world using structure-from-motion techniques which require many overlapping images of the same scene. In contrast, we seek to develop methods which are applicable to diverse scene types and require only a single input image. We are unaware of any existing publications which have pursued this goal, and we hope to encourage research along these lines by demonstrating the feasibility of the approach.

When considering methods from the existing literature which may be applied to dating historical photographs via digital images, we are obviously limited to those based solely on image content rather than inspection of the photograph as a physical artifact. Methods which analyze fashion trends could be useful, but are only applicable to photographs containing at least one human subject. We have found that in a collection of approximately 230,000 Flickr images taken prior to 1980 (see Section 3.1), a state-of-the-art face detection algorithm was able to locate faces in fewer than one-third of the images tested. It is therefore likely that the majority of historical images currently available online do not contain human subjects of sufficiently high resolution to be useful. Man-made objects that are well correlated to specific eras (e.g. cars, trains, etc.) are another intriguing source of temporal information. However, object recognition research is currently not sufficiently advanced to be reliable for this task. We are therefore left with more basic methods which analyze low-level image properties in search of significant statistical correlations. Among such properties, we have found that color is one of the most informative, as color reproduction processes have varied significantly over time (see Figure 1).

Interest in vintage color photography (and particularly in imitating the appearance thereof) has recently skyrocketed with the release of popular mobile ap-

plications such as CameraBag [16], Hipstamatic [25], and Instagram [11]. While the fundamental technique of manipulating color channel response curves is well known (e.g. [24]), we are unaware of any existing publication which has attempted to learn such models using a data-driven approach.

Because it is closely related to both a novel computer vision task and a popular application, we consider the analysis of historical color imagery an excellent area for investigation. This paper, therefore, will focus primarily on the role of color in dating historical photographs.

2 Overview

One can define the problem of estimating a historical photograph’s age in a number of ways. Of particular concern, however, is the relative scarcity and decidedly non-uniform distribution of available historical image data. Overall data availability is highly skewed toward more recent years and the rarity of *color* images from distant years only intensifies this effect. One practical approach is to split the available historical images into discrete time interval classes (over which uniformity may be enforced) and define the problem as one of classification. While many definitions of these time intervals are possible, we note that decades (e.g. “the 1950s”) provide an intuitive grouping which both corresponds well to cultural trends and yields a reasonable solution to the problem of data non-uniformity. We therefore define the classification problem of determining the decade during which a historical color photograph was taken.

Discriminative models which partition by time have the advantage of considering the “complete” appearance of a historical color photograph as it is seen today. As detailed by Wilhelm and Brower [28], certain types of color photographs experience fading and other age-related changes which significantly alter their appearance fairly rapidly. Since these effects are highly dependent on the time elapsed since capture, partitioning by time is vital. However, such models constructed only from low-level image statistics may not give particular insight into the equally important behavior of the specific color imaging process at the time of capture (e.g. for the aforementioned synthesis of novel images which possess the appearance of a specific vintage color process). We may therefore choose to partition historical image data by imaging process, noting that process-specific models are not only relevant to that application, but can also assist us with age estimation: if we are able to match a photograph to the vintage color imaging process which created it, then changes in the relative popularity of these processes over time can implicitly yield probability distributions for the age of the photograph. (Of course, such age estimates could be very imprecise – for example, Kodak’s famed Kodachrome film remained in production with only comparatively minor changes for an astounding 74 years from 1935–2009).

In Section 3, we will discuss the construction of novel historical image datasets which allow us to analyze both the aforementioned decade classification task (Section 3.1) and the modeling of specific vintage color imaging processes (Section 3.2). In Section 4, we will present our approach for data-driven modeling of

vintage color processes (which makes use of a technique originally intended for the radiometric calibration of digital cameras [13]). In Section 5, we will then use the process models of Section 4 along with a number of other features to perform the decade classification task. In spite of the inherent difficulty of the problem, we show that our features and standard machine learning techniques are able to achieve promising results, and we present user study data which indicates that our system’s performance is much better than that of untrained humans at the same task. Finally, in Section 6, we will briefly discuss a technique which applies the process models from Section 4 to the popular application of imitating the overall appearance of vintage color photographs.

3 Databases

Our experiments are enabled by two novel historical image datasets which we have made available online. We will briefly discuss each dataset here before turning our attention to the experiments themselves.

3.1 Historical Color Image Dataset (for Classification by Decade)

As discussed in Section 2, we wish to explore the classification problem of determining the decade during which a historical color photograph was taken, ideally extending back to the introduction of commercial color negative film in the mid-1930s. Our desire for a chronologically uniform distribution of data is complicated by the highly-skewed availability of historical image data on Flickr. Achieving a uniform distribution over the interval 1930–1980 requires discarding over 75% of the available data, and we ultimately select this five-decade span as a compromise between lengthening the interval to the point where data availability is even more skewed, or shortening it to the point where the problem is no longer interesting.

Beginning with a collection of approximately 230,000 Flickr images taken prior to 1980, we perform automated removal of monochromatic images. The remaining images are manually inspected to remove non-photographic content (e.g. scans of vintage artwork) and any remaining monochromatic images. Finally, a random subsampling and decimation is performed to create a dataset containing an equal number of historical color images for each decade (1,375 images total). This dataset is used for the experiments of Section 5; example images are shown in Figure 3 of that section.

3.2 Color Photographic Film Process Dataset

We also construct a dataset to facilitate modeling of several vintage color imaging processes as well as modern digital imagery. In all cases, our selections are motivated by the availability of significant amounts of training data as well as the relative popularity of each process with photographers or as a target of digital simulation.

Modern Digital Imagery: Modern digital cameras are represented by the Nikon D90, a digital SLR camera which as of this writing was one of the most popular cameras among Flickr users.

Technicolor Process 4: “Glorious Technicolor” was the dominant process of color motion picture photography from 1934 until the early 1950s [8]. Although Technicolor itself was not used for still photography, it was conceptually similar to techniques such as the tri-color carbro process which were used by professional photographers at roughly the same time. Attempts to digitally recreate Technicolor may be found in various software packages intended for still image processing (e.g. [2]), and were used to notable effect in the 2004 motion picture *The Aviator* [1]. Some motion pictures photographed in Technicolor are no longer afforded copyright protection in the United States; given our modest data requirements, frames captured from such films are sufficient for our needs. Following the application of a video shot detection algorithm [7], we randomly select a single frame from each shot. This collection of frames is then decimated by using the Gist descriptor [17] to identify and discard duplicate scenes, and then by drawing a random selection from the unique scenes discovered in all films processed.

Kodachrome: Immortalized in song by Paul Simon, Kodachrome’s unique appearance and unrivaled longevity in the marketplace (1935–2009) earned a loyal following among photographers. Like Technicolor, it is a staple of software packages which emulate vintage color processes [2]. Flickr contains several “groups” dedicated to Kodachrome imagery which contain thousands of example images.

Ektachrome: Developed by Kodak in the early 1940s as an easier-to-develop alternative to Kodachrome, Ektachrome also became quite popular. Unfortunately, unlike the fade-resistant Kodachrome, early Ektachrome images have long-term stability described by Wilhelm and Brower as “*extremely poor*” [28].

Velvia, Sensia, and Provia: Introduced by Fuji in 1990 [28], Velvia is a “modern” color film type which (along with Sensia and Provia) remains in production today. As such, these films effectively bridge the chronological gap between discontinued historical color film types and digital imagery of the present day.

Our general approach is the same for each of the seven photographic processes: beginning with a random selection of images covering a diverse array of scene types photographed by each process, we manually remove images which have been extensively processed for artistic effect. A further random selection is then performed leaving a total of 300 images per process. Models of each color process will first be used indirectly (by enabling the construction of features) in the classification experiments of Section 5. Later, in Section 6, the models will be used directly to enable experiments on imitating the appearance of vintage color imagery. We now turn our discussion to our proposed method for creating these models.

4 Modeling Historical Color Film Processes

Having in the previous section constructed a dataset of several historical color imaging processes, we now propose a data-driven technique for modeling the behavior of historical color processes which relies only on such publicly available historical image data. Such modeling is not only relevant to the problem of dating images, but also crucial to the popular application of digitally imitating the appearance of historical color imagery.

Our proposal is to apply a technique for estimating camera response functions to the domain of historical color images. Kuthirummal et al. [13] recently proposed a clever data-driven technique to recover the non-linear inverse response function of a given digital camera model from an online collection of real-world photographs taken by that camera. Two key attributes of their approach are well-suited for use with digital scans of historical color imagery. First, their method requires neither physical access to the camera nor the availability of specialized image content (such as calibration targets, or multiple views of the same scene). Such requirements could be impossible to meet for historical color imaging processes, effectively eliminating from consideration methods such as that of Chakrabarti et al. [5]. Secondly, their method requires surprisingly few images to be effective: the authors state that in some cases fewer than 50 images are required to recover the inverse response function to within 2% RMS fitting error of ground truth. This too is an important benefit, given the sparsity of color images from distant years. The authors describe only the application of their method to images originally captured by a digital camera, and we are unaware of any work which has proposed extending the technique to digital scans of images originally captured on film.

The underlying feature used in [13] is a joint histogram of pixel intensities occurring in each color channel near the center of the image. The camera response function is estimated by an optimization approach which attempts to minimize the symmetric Kullback-Leibler divergence between the transformation of the intensity histograms into irradiance histograms and a non-parametric, camera-independent prior on irradiances. For the purpose of imitating the appearance of historical color imagery, estimating the camera response function of a historical color imaging process is desirable (and will be demonstrated in Section 6). However, for the purpose of discriminating between color imaging processes (for the task of dating historical color images in Section 5), it should not be necessary to explicitly estimate the camera response function. Simply examining the differences between the distributions in pixel intensity space should suffice because in this context, the camera response functions are merely the particular mappings required to transform each of the different intensity histograms into the single non-parametric, camera-independent irradiance prior. We will exploit this observation in the design of our proposed features for dating historical color images, and now turn our discussion to that topic.

5 Classifying Images by Decade

For each image in the five decade database of Section 3.1, we compute a number of novel image features (some using the models described in Section 4), as well as several popular features for conceptually similar tasks such as scene categorization. We then present performance results for linear support vector machine (SVM) classifiers trained using these features, and compare those results with the results of a user study measuring the performance of untrained humans at the same task on the same dataset.

5.1 Features

We evaluate seven image features. The first four features are novel and fairly domain specific while the last three are standard features for image retrieval or scene categorization.

“Process Similarity Feature”: As discussed in Section 4, Kuthirummal et al. explored the use of joint histograms encoding the co-occurrence of RGB pixel intensities for the purpose of radiometric calibration of digital cameras [13]. In that work, the feature was defined by examining four-connected neighborhoods in a 32 by 32 pixel block near the center of the image (i.e. away from vignetting effects), and creating 256 by 256-dimensional histograms of intensity *co-occurrences* for each color channel (i.e. 196,608 dimensions overall) which were then summed over many images.

Our first application of these pixel intensity joint histograms is as follows. First, we sum the histograms as defined in [13] over the images for each photographic process from our color photographic process data set, thereby creating seven cumulative process histograms (see Sections 3.2 and 4). We then treat the histogram from the single decade classification test image as an estimate of the distribution of the photographic process by which the test image was created. As in the optimization procedure of [13], we calculate the symmetric Kullback-Leibler divergence between the test image’s histogram and each of the seven process histograms, also ignoring under- and over-exposed pixels and underpopulated histogram bins as in [13]. We treat the resulting distances as a measure of the similarity of the test image to images photographed by a given color process, which should yield implicit information about the age of the image (as discussed in Section 2). Taken together, the distances form a seven-dimensional feature vector which we refer to as the “process similarity feature”.

Color Co-occurrence Histogram: We additionally calculate the histograms described above at a reduced dimensionality of 32 by 32 per color channel and evaluate their performance as a feature directly.

Conditional Probability of Saturation Given Hue: Haines [8] has argued that the unique appearance of images produced by historical color imaging processes is at least partly due to differences in their reproduction of certain hues, especially with regard to saturation. Based on this intuition, we propose a feature that encodes the correlates of hue and saturation in the CIELAB space.³ We utilize thresholds on chroma to discard pixels which are too close to grey to give reliable hue measurements, then estimate the conditional distribution of pixel saturation given hue by calculating finely-binned two-dimensional histograms over the hue and saturation correlates. The final descriptor encodes the mean and standard deviation of the saturations observed for each of 256 distinct hues for a total of 512 feature dimensions.

Hue Histogram: We examine a one-dimensional histogram which assigns pixels into 128 bins based solely on the hue correlate described above. This histogram may potentially capture changes in the “palette” of colors (e.g. of man-made products and structures) observed during a given era.

Gist Descriptor: The gist descriptor of Oliva and Torralba has been shown to perform well at the tasks of scene categorization [17, 18] and the retrieval of semantically and structurally similar scenes [9]. This feature could potentially capture changing tendencies in scene types or photographic composition over time (if indeed such trends exist), and might also capture vignetting or differences in high frequency image texture which are more related to the photographic process than to the scene. We create gist descriptors for each image with 5 by 5 spatial resolution, where each bin contains that image region’s average response to steerable filters at 6 orientations and 4 scales. Unlike the remaining features, the gist descriptor does not encode color information.

Tiny Images: Torralba et al. examined the use of drastic reductions of image resolution for scene classification [27]. We evaluate the performance of these “tiny images” at a very low spatial resolution (32 by 32) which makes this feature effectively blind to image texture, preserving only the overall layout and coarse color of the scene.

L*a*b* Color Histogram: A fairly standard color histogram feature which has previously been investigated in the context of scene recognition [29] and image geolocation [10]: a joint histogram in CIELAB color space having 4, 14, and 14 bins in L*, a*, and b* respectively for a total of 784 feature dimensions.

³ Hue is formally defined in the CIELAB space and the related CIE L*C*h space. Saturation is *not* formally defined in CIELAB, but a formula analogous to that used in the CIELUV space is an oft-used approximation which is adequate for our purposes.

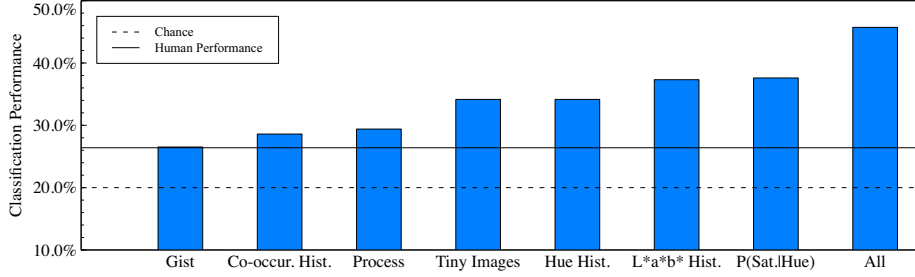


Fig. 2. With a combination of carefully designed color features, standard linear support vector machine learning is able to classify historical color images by decade with accuracy significantly exceeding that of untrained humans (45.7% vs. 26.0%).

5.2 Decade Classification: Support Vector Machine Performance

We evaluate the performance of these features by training a set of linear support vector machines (SVMs) in a one-vs-one configuration ([4],[6]) to solve the five-way decade classification problem described above. In all of our experiments, 50 of the available images for each decade were randomly selected for testing, while the remaining 225 images of each decade were used for training the SVM classifiers using ten-fold cross-validation. We report the accuracy of the classifiers on the test sets, averaged over at least ten random training/testing splits for each feature. Our results are summarized by Figure 2.

The gist descriptor yields the lowest performance (26.5%). The comparatively poor performance of this feature may indicate that the information about scene structure which it captures is not as powerful of a cue for this task as the color information which it ignores. Tiny images (which are coarser than gist, but do capture color information) perform somewhat better (34.2%), as do hue histograms (also 34.2%), which may indicate that changes in the color “palette” over time are a useful cue for this task.

Performances of the co-occurrence histograms in RGB space (at 28.6%) and the “process similarity feature” (at 29.4%) are modest when evaluated individually. The most powerful individual features are the L*a*b* histogram (37.3%), and the conditional probability of saturation given hue (37.6%). Performance of the L*a*b* histogram has previously been analyzed in the literature ([10],[29]), and we find it applicable to this task as well. The best individual performance, however, was achieved by exploiting the observation (in e.g. [8]) that certain vintage color imaging processes reproduced select hues with uniquely intense saturation.

Classification errors between the various features are not entirely correlated, and a noticeable performance improvement (to 45.7%) is obtained by fusing all of the aforementioned features together. Unfortunately, our experimentation with kernel SVMs showed that while they did provide modest performance improvements for certain features, they were prone to overfitting on others; we therefore do not report performance results for this approach.

Figure 3 depicts the five images our classifier most confidently assigned to each of the five decades in our data set. Only one of these twenty-five most confident classifications is incorrect.

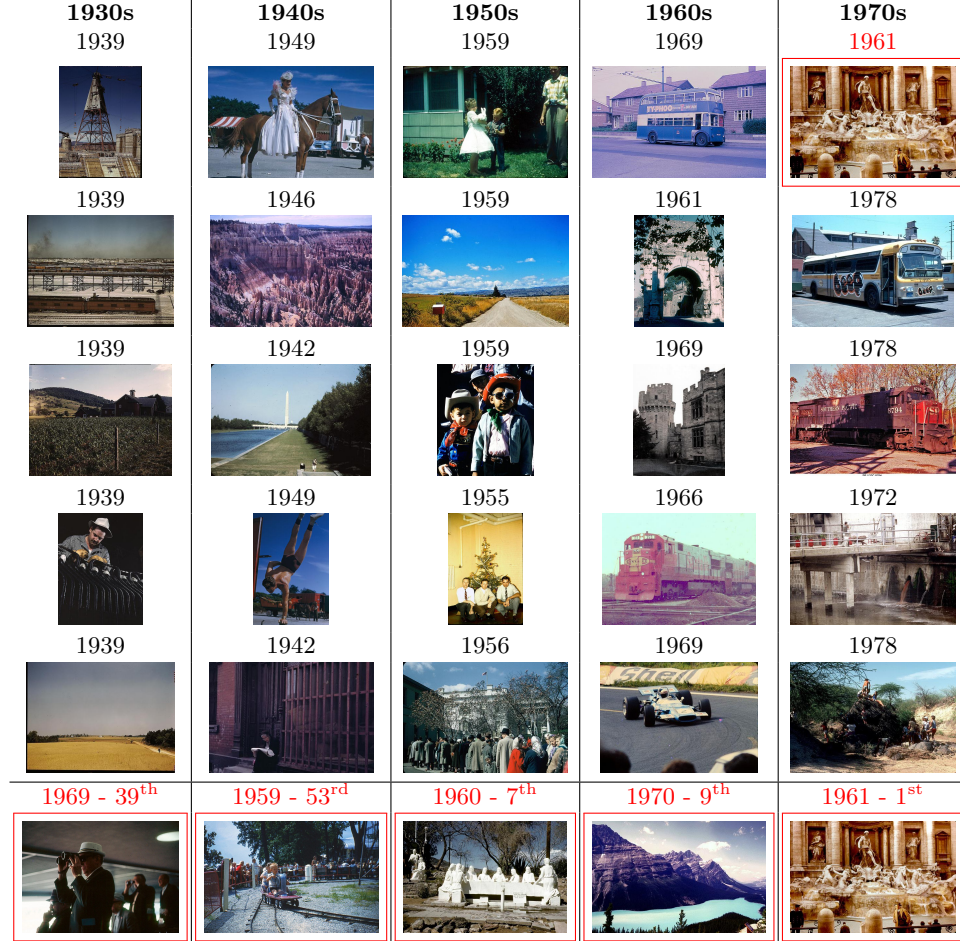


Fig. 3. The five photographs which our classifier most confidently assigned to each of the decades in our data set (arranged by descending confidence) and the most confident *incorrect* classifications from each decade. Each image is labeled with the actual year during which it was taken, and the misclassifications are also labeled with the ordinal rank of their confidence.

The figure additionally shows the most confident *incorrect* classifications from each decade. Images misclassified as 1950s, 1960s or 1970s images tend to be confident classifications (e.g. 7th most confident overall) made near the boundaries between decades. This behavior may be related to our decision to approach the

problem as one of classification rather than regression (as discussed in Section 2). In contrast, 1930s and 1940s misclassifications are made far less confidently (e.g. 53rd most confident overall) and are somewhat less predictable.

5.3 Decade Classification: Human Performance

In order to ascertain the performance of untrained humans for the same classification task, we conducted a user study via the Amazon Mechanical Turk [3] platform. Workers on the service were shown an image from our dataset and told that it was taken sometime between 1930 and 1980. They were then asked to indicate via multiple choice the decade (i.e. “1930s”, “1940s”, “1950s”, “1960s”, or “1970s”) during which they believed the photograph was taken; there was no enforced time limit for the decision to be made. Each of the 1,375 images in our dataset was evaluated by six unique workers, yielding a total of 8,250 human classifications.

Analysis of the responses obtained reveals an overall performance of only 26.0% at correctly answering the five-way decade classification problem, a small gain over random chance at 20.0%. We note that even our least powerful features individually show better overall accuracy for this task, and that the overall performance of the combined feature classifiers is nearly twice this level.

6 Imitating the Appearance of Historical Color Processes

Modification of color channel response curves is a commonly known technique for imitating the appearance of historical color imagery (see e.g. [24]). To date, the proper curves for each imaging process have commonly been found by “trial-and-error” experimentation; while developers of proprietary software packages might have employed a more formal approach, we are not aware of any publication which reveals such a method to the literature. While we did not need to explicitly estimate the camera response function of a vintage imaging process to obtain temporally discriminative information in Section 5, here we do attempt to recover the color channel curves used by e.g. [24] to imitate vintage color imagery.

From the data set described in Section 3.2, we applied the technique of [13] (detailed in Section 4) “as is” to estimate response functions for modern digital imagery and two variants of Technicolor Process 4. The estimated response functions (depicted by the second row of Figure 4) reveal the significant differences which give rise to the unique appearances of images captured by these color imaging processes (examples of which are shown in the first row of Figure 4).

Given the camera response function estimates from Section 4, we transform an image from one color process to another by applying first the inverse response function of the source color process (thereby converting pixel intensities to estimates of scene irradiances), then the response function of the target process (thereby producing estimates of pixel intensities for the target system). A complication arises from the fact that the response functions (and hence the irradiance estimates) are known only up to a scale factor: naive application of

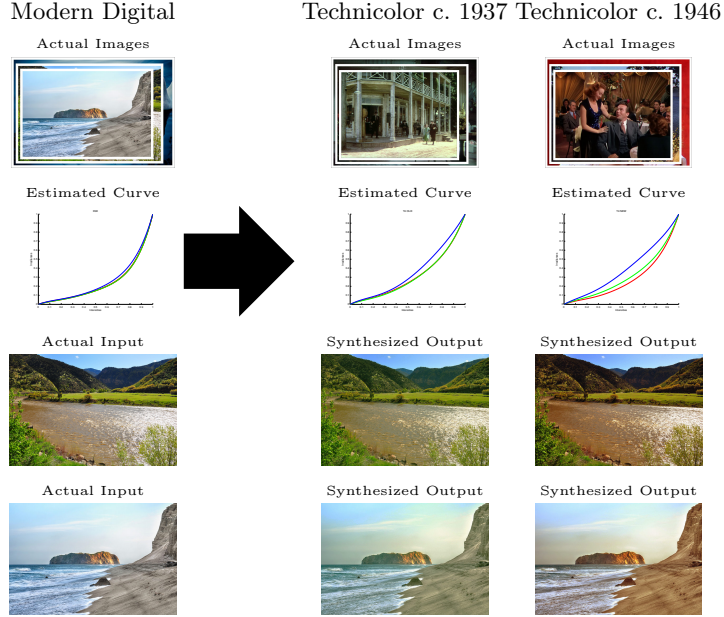


Fig. 4. Collections of vintage color images (Row 1) yield estimated color response curves (Row 2) which may be used to imitate vintage color imagery (Rows 3 & 4).

the above approach can result in drastic changes of brightness which overwhelm the comparatively subtle changes of image color. We address this issue by introducing a constant scaling factor between the application of the source and target response functions and search (via the Nelder-Mead method [15]) for the scaling which minimizes the difference between the input and output image brightness (measured as the mean value of the L^* channel in CIELAB space).

Example synthesis results are shown in Figure 4, which depicts the transformation of modern digital landscape images to imitate the highly-saturated appearance of the later Technicolor process and the more subdued appearance of its predecessor.

7 Discussion

Inspired by the recent online availability of large quantities of historical color image data, we introduced the novel task of automatically classifying a historical color photograph into the decade during which it was taken. We showed that given carefully designed features, color information provides a sufficiently strong cue for relatively simple machine learning techniques (i.e. linear SVMs) to achieve promising results significantly better than untrained humans for this difficult task. Finally, we introduced the novel application of data-driven camera response function estimation to historical color imagery, demonstrating both the

creation of discriminative features for the aforementioned classification task, and an approach to the popular application of imitating the appearance of historical color photographs more principled than that currently granted to the literature.

While online availability of historical images will no doubt continue to improve, it remains a fundamental truth that far fewer such images were ever captured (or survive to the present day) than are now captured through the widespread use of digital cameras. An interesting challenge of this research domain will continue to be the need to intelligently use the sparse historical training examples from distant years. Future research might also include efforts to learn the historical color processes which are present in unlabeled training data, or higher-level reasoning about image content which could extend this process to the significant fraction of historical images which are *not* color, but rather monochromatic.

Acknowledgments: This research was supported by a Google Research Award to James Hays and Alexei A. Efros and NSF CAREER Award 1149853 to James Hays.

References

1. Aviator visual effects - behind the scenes. <http://www.aviatorvfx.com/>
2. Alien Skin Software, LLC.: Alien skin software: Exposure 4 examples. http://www.alienskin.com/exposure/exposure_examples.aspx
3. Amazon.com, Inc.: Amazon mechanical turk. <http://www.mturk.com/>
4. Athanasopoulos, A., Dimou, A., Mezaris, V., Kompatsiaris, I.: Gpu acceleration for support vector machines. In: Proceedings of the 12th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS 2011) (2011)
5. Chakrabarti, A., Scharstein, D., Zickler, T.: An empirical camera model for internet color vision. In: Proceedings of the British Machine Vision Conference (BMVC) (2009)
6. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology 2, 27:1–27:27 (2011), software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
7. Dugad, R., Ratakonda, K., Ahuja, N.: Robust video shot change detection. In: IEEE Workshop on Multimedia Signal Processing (1998)
8. Haines, R.W.: Technicolor Movies: The History Of Dye Transfer Printing. McFarland & Company, Inc. (1993)
9. Hays, J., Efros, A.A.: Scene completion using millions of photographs. ACM Transactions on Graphics (SIGGRAPH 2007) 26(3) (2007)
10. Hays, J., Efros, A.A.: im2gps: estimating geographic information from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2008)
11. Instagram, Inc.: Instagram. <http://instagram.com>
12. Kim, G., Xing, E.P., Torralba, A.: Modeling and analysis of dynamic behaviors of web image collections. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV (5). Lecture Notes in Computer Science, vol. 6315, pp. 85–98. Springer (2010)

13. Kuthirummal, S., Agarwala, A., Goldman, D.B., Nayar, S.K.: Priors for large photo collections and what they reveal about cameras. In: Forsyth, D.A., Torr, P.H.S., Zisserman, A. (eds.) ECCV (4). Lecture Notes in Computer Science, vol. 5305, pp. 74–87. Springer (2008)
14. Messier, P.: Notes on dating photographic paper. *Topics in Photographic Preservation* 11, 123–130 (2005)
15. Nelder, J., Mead, R.: A simplex method for function minimization. *Computer Journal* (7), 308–313 (1965)
16. Nevercenter Ltd. Co.: Camerabag. <http://nevercenter.com/camerabag/mobile/>
17. Oliva, A., Torralba, A.: Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision* 42(3), 145–175 (2001)
18. Oliva, A., Torralba, A.: Building the gist of a scene: the role of global image features in recognition. *Progress in Brain Research* 155 (2006)
19. Schindler, G., Dellaert, F.: Probabilistic temporal inference on reconstructed 3d scenes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010)
20. Schindler, G., Dellaert, F., Kang, S.B.: Inferring temporal order of images from 3d structure. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007)
21. Severa, J.: *Dressed for the Photographer: Ordinary Americans & Fashion, 1840–1900*. The Kent State University Press (1997)
22. Shrimpton, J.: *Family Photographs & How to Date Them*. Countryside Books (2008)
23. Storey, N.: *Military Photographs & How to Date Them*. Countryside Books (2009)
24. Sulonen, P.: Simulating film effects with curves. http://www.prime-junta.net/pont/How_to/100_Curves_and_Films/_Curves_and_films.html
25. Synthetic, LLC: Hipstamatic. <http://hipstamatic.com/>
26. Taylor, M.A.: *Uncovering Your Ancestry through Family Photographs*. Family Tree Books, second edn. (2005)
27. Torralba, A., Fergus, R., Freeman, W.T.: 80 million tiny images: a large dataset for non-parametric object and scene recognition. *IEEE Transactions On Pattern Analysis And Machine Intelligence* 30(11), 1958–1970 (2008)
28. Wilhelm, H., Brower, C.: *The Permanence and Care of Color Photographs: Traditional and Digital Color Prints, Color Negatives, Slides, and Motion Pictures*. Preservation Publishing Company (1993)
29. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010)
30. Yahoo!, Inc.: Flickr: The commons. <http://www.flickr.com/commons/>