Predicting Range of Acceptable Photographic Tonal Adjustments

Ronnachai Jaroensri, MIT CSAIL
Sylvain Paris, Adobe Research
Aaron Hertzmann, Adobe Research
Vladimir Bychkovsky, Facebook, Inc.
Frédo Durand, MIT CSAIL
http://groups.csail.mit.edu/graphics/acceptable_adjust/

Abstract

There is often more than one way to select tonal adjustment for a photograph, and different individuals may prefer different adjustments. However, selecting good adjustments is challenging. This paper describes a method to predict whether a given tonal rendition is acceptable for a photograph, which we use to characterize its range of acceptable adjustments. We gathered a dataset of image "acceptability" over brightness and contrast adjustments. We find that unacceptable renditions can be explained in terms of over-exposure, under-exposure, and low contrast. Based on this observation, we propose a machine-learning algorithm to assess whether an adjusted photograph looks acceptable. We show that our algorithm can differentiate unsightly renditions from reasonable ones. Finally, we describe proof-of-concept applications that use our algorithm to guide the exploration of the possible tonal renditions of a photograph.

1. Introduction

Photographic adjustment is a subjective matter: a photograph can be modified with different exposure, color and other adjustments in various ways, depending on the user's intent and preference. However, performing these adjustments is a difficult task because there are many parameters to adjust. Presets and other completely automatic methods have been proposed [1, 2, 11, 12, 20]. However, when more controls are desired, users must fall back on tedious parameter manipulation. As the space of parameters grows, the exploration naturally becomes more difficult [14].

This paper is inspired by the observation that, while many different adjustments may be valid for a given photograph, there are many, more-extreme adjustments that would also never be chosen for normal photo touch-up. We refer to this space of normal adjustments as *acceptable*. For example, Figure 1 shows a photograph where a user may prefer an accented highlight on the hair (Fig. 1(a)), or a good exposure on the entire face (Fig. 1(b)). However, extreme adjustments lead to photographs that are too degraded to be normally usable (Figs. 1(c) and 1(d)). We call these photographs *unacceptable*. Of course, precisely defining the notion of acceptable is in itself a challenging task.

Our goal is not to estimate the absolute aesthetics of an adjusted photograph. For this work, acceptability has more to do with recognizable content and proper exposure than it does with composition and color harmony of the subject. We seek to study how the quality of a natural photograph changes as we explore its renditions across the adjustment space. This contrasts our work with techniques in the aesthetics assessment literature where acceptable rendition is essentially a minimum requirement [5, 6, 13, 15, 16].

This paper describes a method for predicting the acceptable set of tonal adjustments for a given photograph. There are four main contributions of this work.

1. We gather a crowd-sourced dataset of labeled photographs that allow us to analyze the notion of acceptability. This dataset is available online.

2. We find, from this dataset, that in a two-parameter adjustment space, photograph corruption can be categorized into three forms: over-exposure, under-exposure, and low contrast.

3. We introduce a learned model that predicts whether a given rendition is acceptable. We use this predictor to define the range of acceptable adjustments by sampling the parameter space and analyzing each sample. We show that this model makes accurate prediction of acceptable renditions, even for the ones adjusted by a more complex model.

4. We demonstrate the benefits of our algorithm with proof-of-concept editing applications that adapt to the predicted ranges.

1.1. Related Work

There exist several approaches to automatic photo adjustment. One is to use machine-learning techniques based on manually built training sets [1, 2, 11]. Another option is to directly translate standard photographic guidelines into an automatic algorithm [12, 20]. The common point of these

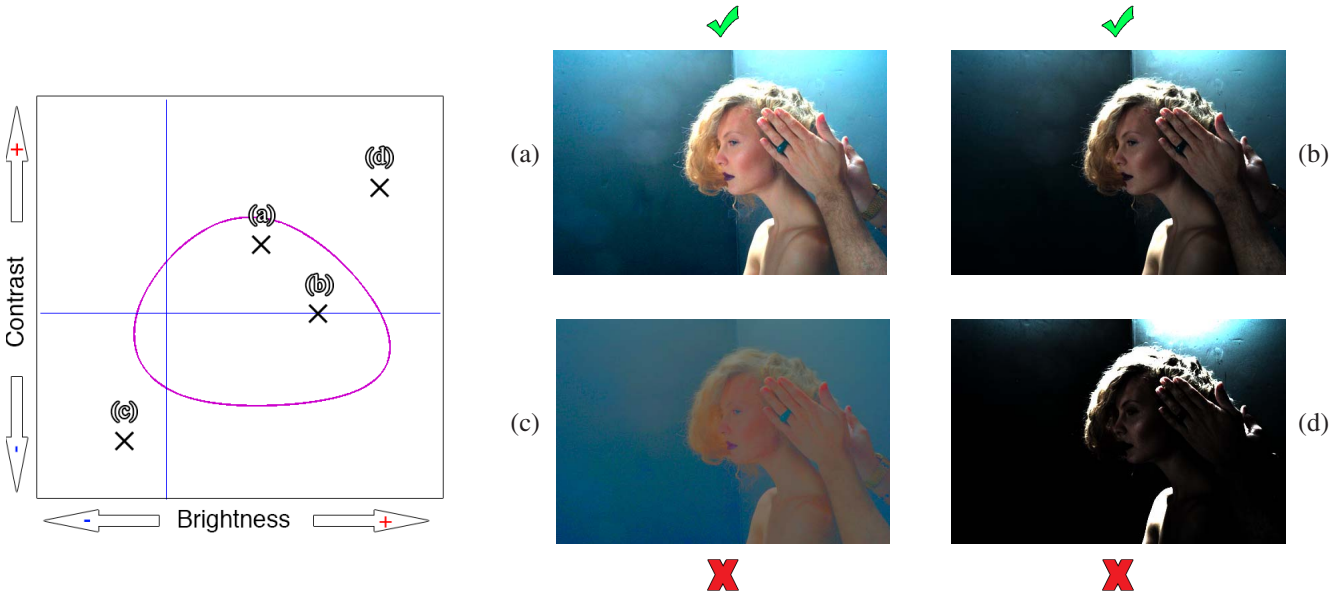

Figure 1: *Adjustment Space Exploration.* There are many valid ways to adjust an image. Our two-dimensional adjustment space is shown on the left, with the origin (no adjustment) at the crosshairs. Four adjustments are shown in the space with Xs: two acceptable, two unacceptable. The magenta curve shows the boundary between acceptable and unacceptable adjustments, obtained by crowdsourcing: outside the boundary, the images no longer look acceptable. We present a machine learning algorithm to estimate this boundary for new images.

approaches is they predict a single valid adjustment. In contrast, we aim for the set of all such adjustments.

Our work is also related to automatic aesthetic assessment, e.g., [5, 6, 13, 15, 16]. These methods assign a score to images depending on their visual quality but do not modify them. They include subject content into their consideration, and a strong emphasis is put on image composition. By contrast, we focus on the effect of adjustments on the image renditions. That being said, our approach shares the same strategy of using machine learning to exploit training data, and some aspects of our image descriptor are inspired by those used in these techniques.

Recently, Koyama *et al*. [14] proposed a user interface for exploring the range of valid adjustments, showing that adapting controls to valid ranges improves user experience and speeds up their exploration process. Their work required crowd-sourcing for each image to define the quality of different adjustments. Our work is complementary, as we show how to define the range of acceptable adjustments for new images without requiring crowd-sourcing.

2. A Crowd-sourced Dataset

Our dataset consists of photographs adjusted to different renditions, and the corresponding binary labels {acceptable, unacceptable} specified by human subjects. We used the MIT–Adobe FiveK dataset because it provides a variety of high quality photographs in both edited and original renditions [1]. We started from 500 randomly selected photographs and manually adjusted the selection to ensure a variety of natural scenes and conditions, e.g., day vs. night, indoor vs. outdoor.

"Acceptable" is an ambiguous term. In our case, we study the effect an adjustment has for a given input photographs, rather than comparing the aesthetics of different photographs. Properly adjusted, a beautiful rose should be equally "acceptable" as a pile of dead leaves. In our experiment, we specifically instructed our participants to ignore the subject content of photographs, and provided them contrasting examples of professionally adjusted photographs with good and bad compositions in our instruction.

Furthermore, we found that some extreme adjustments can create non-photorealistic renditions that our experiment participants might find interesting (Figure 3). We collected a small set of data that allowed such effects to be accepted, but there were not enough agreement among the participants to create a meaningful signal. Thus, in the final version of the survey, we excluded these effects, and instructed our participants to reject such renditions.

We crowd-sourced our data through Amazon Mechanical Turk (MTurk). While this has several drawbacks such as the lack of control for viewing conditions, our MTurk experiment was representative of how people typically view photographs on their computers. Furthermore, as pointed

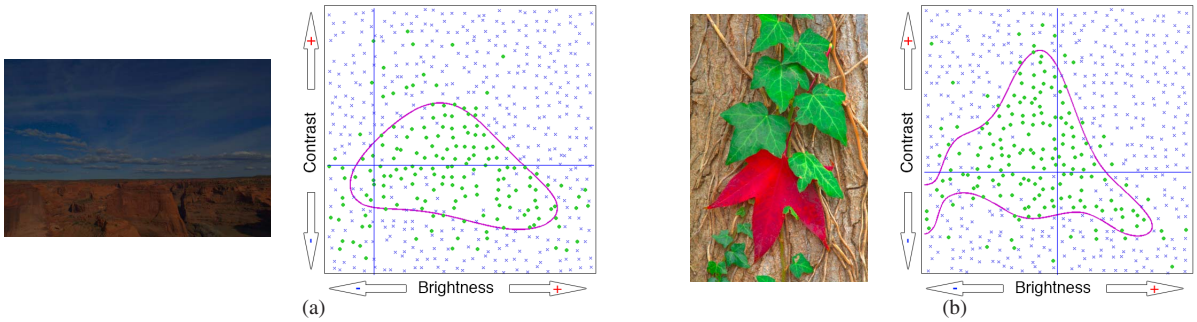

Figure 2: Sample unedited photographs with their corresponding crowdworker data. The green dots represent *acceptable* adjustments, and the blue x's represent *unacceptable* adjustments. The blue crosshair is zero adjustment. Shown boundaries (magenta) are calculated using SVM with radial basis function kernel on the adjustment values.

out by Harris *et al*., the quality of MTurk data approaches that of controlled experiment when there are enough data points [8]. While Harris *et al*. did not provide a threshold for our specific experiment, we ensured that we sampled our adjustment space densely enough to see consistent trends within our data. Additionally, we allowed each participants to see only one rendition per photograph. This ensures that a photograph is seen by multiple persons, and the data represents the opinion of the crowd. We did not show a reference photograph to the participants, and they were forced to choose acceptable/unacceptable for each photograph as they saw it.

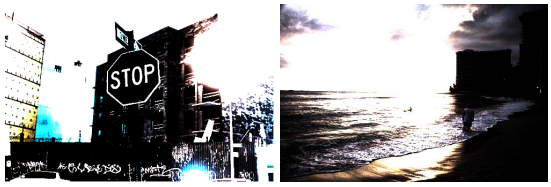

Figure 3: Extreme adjustment that might be seen as interesting by some users. Because answers were not consistent on those, we guided crowdworkers to reject them.

In this work, we focus on tonal adjustments and seek to minimize color-related effects. In particular, we found that using the original white balance and saturation can sometimes make a photo look bad independently of the brightness and contrast settings. To avoid this issue, we applied the white balance and saturation of one of the expert retouchers of the MIT–FiveK dataset (we selected the same expert as Bychkovsky et al. in their study) and kept the other parameters as is. This color-adjusted photo is what we used as the zero adjustment on top of which we applied tonal adjustments.

We performed our brightness and contrast adjustments in Chong et al.'s color space because it decouples the chrominance and the luminance channel [4]. Our brightness adjustment roughly corresponds to a multiplication of the linear RGB channel, while the contrast adjustment corresponds to a gamma curve. More detail of the experiment protocol and adjustment model is given in the supplementary material.

The number of samples we need grows exponentially with the adjustment space dimensionality. Studying a 2D adjustment space allows us to sample densely while keeping the size of our data manageable. Despite its simplicity, brightness and contrast represent the most important tonal adjustments [1]. As we shall see in Section 5, the learning from this space applies in higher dimensions as our machine-learning technique also makes good prediction in a more complex adjustment space.

3. Data Analysis

To gain some insight and guide the design of our algorithm, we analyze the distribution of our data in the brightness $\times$ contrast space. We observe a dense region with accepted adjustments in the middle of the sampling window. The density of the accepted points decreases away from the center. Figure 2 shows a few photographs from our dataset and their corresponding acceptable regions. Additionally, we observe significant image-to-image variations of the acceptable region with the exposure of the original photograph playing a major role. For instance in Figure 2(a), the region is shifted toward the high-brightness values because the input is under-exposed, while the region in Figure 2(b) is centered because the corresponding original photograph is well exposed. Furthermore, we also observe significant variations in the size of accepted regions.

Inspecting the acceptable regions for many images, we observe that the acceptable region frequently has a triangle-like boundary with curved edges (Fig. 2). This trend is further confirmed with a kernel density estimate of the ac-

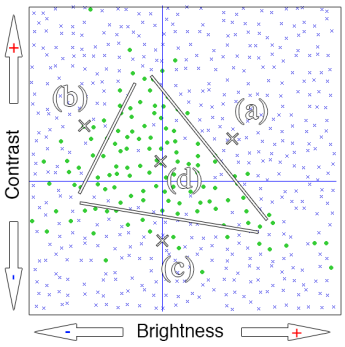

(a) crowd-sourced data

(a) over-exposed

(b) under-exposed

(c) low contrast

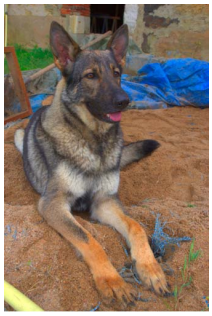

(d) original

Figure 4: Failure cases in tonal adjustments. Each edge of the triangle represents different failure cases: over-exposure (a), under-exposure (b), and low contrast (c). The original photograph at zero adjustment is shown in (d).

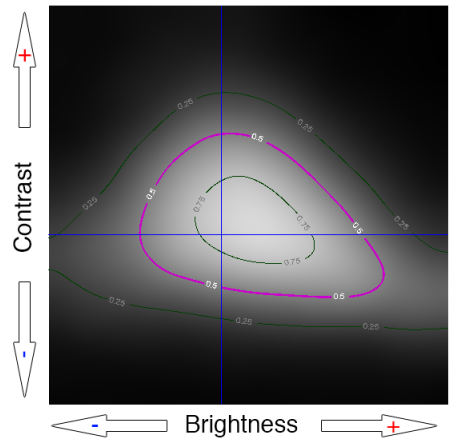

Figure 5: Kernel density estimate of accepted data points. An arbitrary contour is drawn in magenta to illustrate the triangular shape, which indicates three cases of image degradation: overexposure, underexposure, and low contrast.

cepted adjustments (Fig. 5). We found that the triangular shape indicates three issues (Fig. 4): over-exposure, under-exposure, and low contrast.

These three cases will inform the design of the machine-learning features that we use to predict whether an adjustment is acceptable. In Section 5.2, we shall see that these observations also generalize to a higher-dimensional tonal adjustment space.

4. Predicting Acceptability

In this section, we build upon the previous observations to design our predictor. Formally, given an input tuple (I, x), we seek to predict binary label $g \in \{acceptable, unacceptable\}$, where I is the input image, and x is the adjustment value. We first describe two toy predictors that make their predictions solely on adjustment values. Then, we explain our technique, which makes its prediction based on adjusted photographs and is completely agnostic of the adjustment model. The toy predictors will help characterize our data and will provide useful performance comparisons to our technique.

4.1. Toy Predictors

In order to characterize the variations in our data, we create two toy predictors that only use the adjustment value as their feature: the *average-adjustment* predictor, and the *per-image* predictor. These toy predictors define performance bounds for the machine-learning algorithm that we will describe in Section 4.2.

The average-adjustment predictor is trained purely on adjustment value and ignored image dependency completely. This corresponds to a naive form of machine learning where the prediction is always the same and equal to the average of the training set. While a random predictor would be worse, we nonetheless expect any sophisticated predictor to perform better than this simple technique. The per-image predictor, on the other hand, is a collection of classifiers, each of which was trained on one of the images from the dataset. These predictors generalize poorly to unseen photos due to their minimal training set, but perform ideally on their corresponding training image. Because of this property, they provide an upper bound for the subsequent evaluation, i.e., we do not expect any generic predictor to perform better on a given photo than a predictor trained specifically on that photo. These predictors also help us quantify the irreducible classification error due to the progressive transition between acceptable and unacceptable classes.

We use Support Vector Machines with Radial Basis

Function kernel as the algorithm for these toy predictors. Our implementation is based on LibSVM for the per-image predictor, and on BudgetedSVM with LLSVM approximation scheme for the average-adjustment predictor [3, 7]. We did 8-fold cross-validation to choose the slack cost.

4.2. Our Image-Based Predictor

Instead of using adjustment values in the prediction, our machine-learning algorithm learns directly from the adjusted photos, and it is completely agnostic to the adjustment model. This will allow us to apply its learning towards more general tonal adjustment models such as those that includes highlight and shadow adjustments. We design our features to capture the three observed failure cases: over-exposure, under-exposure, and low contrast. We experimented with many different sets of features, the best option is as follows:

1. *Fraction of Highlight and Shadow Clipping:* Over-exposure and under-exposure appear when pixels are clipped. We calculated the fraction of such clipped pixels in both directions as over-exposure and under-exposure features. We computed these quantities globally, weighted by image saliency [10], and by manual sky annotation [19].

2. *Luminance Histogram:* A standard practice among photographers is to observe the luminance histogram to expose a picture. Inspired by this, we added a 10-bin histogram to our feature vector. We found that weighing by the saliency values performs better.

3. *RMS Luminance Contrast:* A common way to measure contrast in images is the rms contrast [18]. In addition to using saliency weighing scheme as before, we also calculated this feature on each scale of a Gaussian pyramid and local sliding windows. This helps capture the effect of contrast in different scales, and locally.

4. *Loss of Gradient:* Varying contrast settings changes the strength of edges in the image. An image region loses its saliency when its edges become weak. This feature counts the fraction of edge pixels whose gradients fall below a constant threshold that is predetermined experimentally (see supplemental material).

5. *Local Band-Limited Contrast:* Peli suggests that contrast in each spatial frequency band may effect our perception differently [18]. We divided the images into 9 bands and computed their contrast as described by Peli. Haun and Peli suggest a winner-takes-all pooling approach which we follow with a power mean using exponent $p = 4$ to average the per-pixel contrast values into a single number [9].

The total dimensionality of our feature space is 48. The features are normalized such that each dimension is zero-mean with a standard deviation of 1. We describe the details of our feature in supplementary material. We use MATLAB© implementation of Logitboost and cross-validate to choose the number of ensembles [17]. We also experimented with linear and RBF SVM, but the Logitboost outperformed both SVMs considerably.

5. Evaluation

As described in the previous section, we designed our algorithm to make predictions directly from the adjusted photographs, independently of the adjustment model. In this section, we show that our algorithm generalizes to unseen photographs and that it performs well on photographs adjusted by a more sophisticated model.

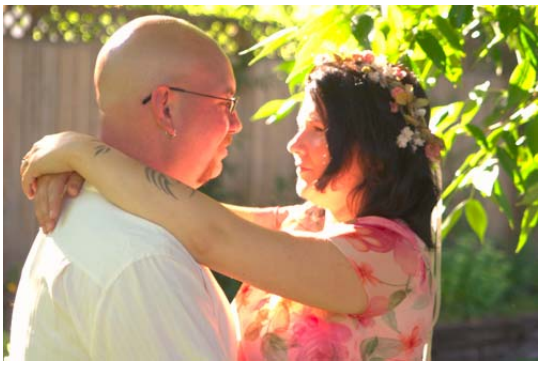

(a) input photo

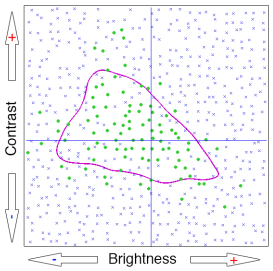

(b) our algorithm

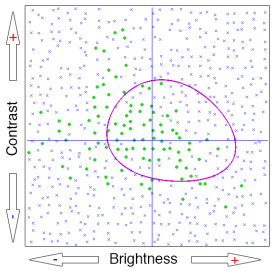

(c) average-adjustment predictor

Figure 6: Our algorithm adapts to the exposure of the input photo (a). It allows it to better represent the space of acceptable adjustments (b) whereas the average-adjustment predictor includes many unacceptable samples in its acceptable range and also misses many acceptable points (c).

5.1. Prediction Result

Table 1 summarizes our prediction result. Here, we exclude points that are likely to be noise such as those that are on the wrong side of boundary and outside the margin

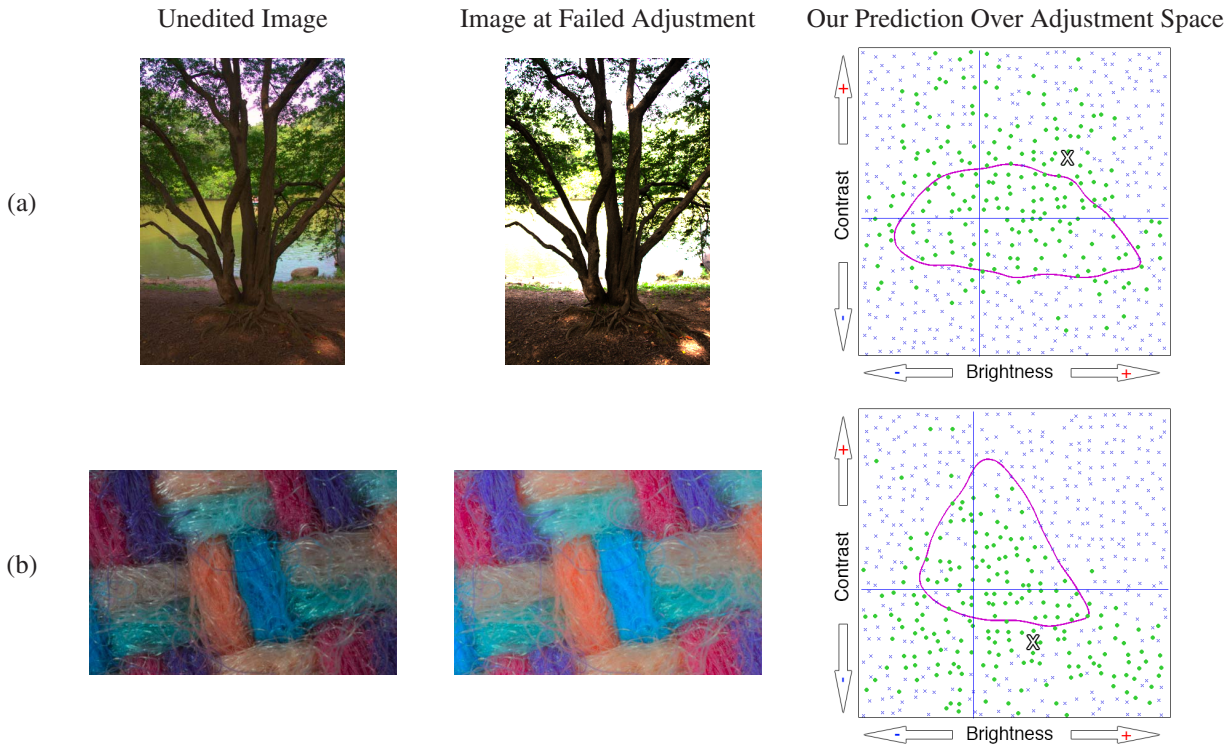

Figure 7: *Our Algorithm's Failure Cases* The adjustment values are marked with X's in the third column. High dynamic range (a) appeared overexposed, while textured scene (b) were allowed contrast in unusual range, causing misclassifications.

Predictor	Error (%)	Precision	Recall	F-Score
Average Adjustment	15.6	0.68	0.55	0.61
Per-Image	10.6	0.80	0.68	0.73
Ours	14.1	0.71	0.62	0.66

Table 1: *Performance Comparison.* As expected, our predictor performed better than the naive average adjustment but not as well as the per-image fit (remember that this latter option does not generalize to unseen photographs, it is only provided as a reference).

of the per-image SVM. All metrics are calculated on the remaining data points. After removing noise, we have approximately 23% positive points left.

Not surprisingly, the per-image predictor does best in all performance metrics and the average-adjustment predictor does significantly worse. For the latter, its relatively high precision, and relatively low recall suggests that its predictions are generally conservative. Nonetheless, our algorithm performs better than the average predictor. Using the F-score that balances precision and recall, our algorithm is about half way between the baseline and the ideal case.

The performance of the average-adjustment predictor heavily relies on the fact that most input photographs are well exposed. The prediction fails on an under- or over-exposed image (Fig. 6). Our algorithm, on the other hand, is based on directly calculating the image statistics, and it finds the correct location of the acceptable region.

The failure cases of our predictor are often the result of errors in scene understanding. For instance, in Figure 7(a), human observers tolerate saturating the river as long as the tree that is the main element of the picture looks good. Our algorithm misses the layering of the scene and predicts such saturation as unacceptable. Figure 7(b) shows a different scenario where the photograph is a pure texture without recognizable content. In this case, human subjects allow low contrast adjustments below what our algorithm predicts.

5.2. Going beyond Brightness and Contrast

Because our algorithm is agnostic to the adjustment model, it generalizes beyond the brightness–contrast adjustment space. To verify this, we collected a small dataset with highlight and shadow adjustments added to the adjustment model. The adjustment range was controlled such that half of the dataset was acceptable. Variations in brightness and contrast adjustments were limited to ensure that we observe mainly the effect of the added dimensions. After dis-

carding low-quality data as done previously, we had 4,500 data points from 50 unseen photographs, with 58% accepted points.

Predictor	Error (%)	Precision	Recall
Ours	**24.8**	**0.75**	0.87
Average Adjustment	39.2	0.61	**0.96**
Per-Image	35.8	0.64	0.90
Chance	41.8	0.58	1.0

Table 2: *Performance in High Dimensional Adjustment Space.* Our algorithm's performance remained the same. Performance of both toy predictors were near chance, which indicates that the effect of added adjustment dimensions dominates human judgment. Chance predictor is provided for reference, and, in this case, it does best if it always gives positive prediction.

Table 2 summarizes the performance of our algorithm on this dataset. Our algorithm does significantly better than the average-adjustment and per-image predictors (both toy predictors ignore additional adjustments and predict based on brightness and contrast only). In fact, the toy predictors' misclassification rates are about the same as chance, and the high recall indicates that they mostly yield positive predictions on this dataset. This shows that the added dimensions have a significant effect over human's judgement of the image quality. We cannot directly compare this experiment with the brightness–contrast only experiment because the sampling rates differ (sampling the 4D space as finely as the brightness–contrast space and crowd-sourcing votes is impractical due to the sheer number of samples that would be required); nonetheless, the precision rating suggests that our algorithm does equally well on this 4D dataset as on the brightness–contrast data.

6. Applications

In this section, we demonstrate a few applications that benefit from our predictor. More examples and demos can be found in the supplementary material.

6.1. Most Favorable Adjustment Prediction

There are many possible schemes to deduce a single most favorable point from range of acceptable adjustments, e.g., using the adjustment that is the farthest from the boundary. In our work, the predictor also provides us with a score, and we choose the adjustment with highest score as our most favorable point. The average adjustment predictor is independent of the image content, so it always predicts a constant value as its most favorable adjustment. Using its prediction is nearly equivalent to no operation, and it will fail when the input is too different from the average (Fig. 8). Our algorithm respects the image content and yields similar renditions to that of the ideal per-image predictor. Note that our algorithm is trained on crowd-sourced data and as such reflects the consensus of the crowd. An interesting avenue for future work would be to tailor our predictions to a specific user akin to previous work like [1, 11].

6.2. Smart Preset

Presets are a simple method for achieving consistent looks on an image collection. However, blindly copying and pasting adjustment values to all images is bound to fail because each image has different dynamic ranges, and camera metering is not perfect. Our smart preset clamps the adjustment value to the predicted boundary by shrinking it towards the most favorable adjustment found earlier.

Figure 9 illustrates a case where the smart preset is effective. The sky in the original image is properly exposed, but the photographer wanted to reveal the building in the bottom of the picture. Since the building is dark, direct application of the adjustment value to other images causes overexposure in almost all cases. Our smart preset limits the extent to which the adjustment is allowed to alter an image, which generates better visual results.

6.3. Adjustment Space Exploration Assistance

There are many ways we can use the prediction to assist users in exploring the adjustment space. Koyama et al. describe an adjustment interface where each sliders is color-coded to indicate the direction that would lead to a more desirable result [14]. The crowd-sourced underlying "goodness" function also allows them to make a better adjustment randomization by biasing towards adjustment with higher scores. However, for each unseen image, the "goodness" function needs to be crowd-sourced, which takes time and money. Our work complements theirs by providing the required "goodness" score on unseen images.

7. Conclusion

We collected and studied a dataset of acceptable adjustment over brightness and contrast adjustment. To characterize our dataset, we described two adjustment value-based toy predictors: the average-adjustment and per-image predictors. We found significant performance gap between these two toy predictors, which suggests that image-dependent analysis is needed.

Our dataset showed that unacceptable renditions can be explained by over-exposure, under-exposure, and low contrast. We designed a machine-learning feature based on this observation, and showed that it performs significantly better than the average toy predictor and that it generalizes to higher-dimensional tonal adjustment models.

Through the range of acceptable adjustment predictions, we demonstrated three proof-of-concept applications that facilitate adjustment space exploration. We hope that these applications would create better user experience with complex photographic adjustments, allowing users to work faster, while keeping the quality of their work high. We will share our dataset to help reproduce these results.

References

[1] V. Bychkovsky, S. Paris, E. Chan, and F. Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 97–104. IEEE, 2011. 1, 2, 3, 7

[2] J. C. Caicedo, A. Kapoor, and S. B. Kang. Collaborative personalization of image enhancement. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 249–256. IEEE, 2011. 1

[3] C.-C. Chang and C.-J. Lin. 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. 5

[4] H. Y. Chong, S. J. Gortler, and T. Zickler. A perception-based color space for illumination-invariant image processing. In *ACM Transactions on Graphics (TOG)*, volume 27, page 61. ACM, 2008. 3

[5] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Studying aesthetics in photographic images using a computational approach. In *Computer Vision–ECCV 2006*, pages 288–301. Springer, 2006. 1, 2

[6] R. Datta, J. Li, and J. Z. Wang. Algorithmic inferencing of aesthetics and emotion in natural images: An exposition. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 105–108. IEEE, 2008. 1, 2

[7] N. Djuric, L. Lan, S. Vucetic, and Z. Wang. Budgetedsvm: A toolbox for scalable svm approximations. *Journal of Machine Learning Research*, 14:3813–3817, 2013. 5

[8] M. D. Harris, G. D. Finlayson, J. Tauber, and S. Farnand. Web-based image preference. *Journal of Imaging Science and Technology*, 57(2):20502–1, 2013. 3

[9] A. M. Haun and E. Peli. Perceived contrast in complex images. *Journal of vision*, 13(13):3, 2013. 5

[10] T. Judd, K. Ehinger, F. Durand, and A. Torralba. Learning to predict where humans look. In *Computer Vision, 2009 IEEE 12th international conference on*, pages 2106–2113. IEEE, 2009. 5

[11] S. B. Kang, A. Kapoor, and D. Lischinski. Personalization of image enhancement. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1799–1806. IEEE, 2010. 1, 7

[12] L. Kaufman, D. Lischinski, and M. Werman. Content-aware automatic photo enhancement. In *Computer Graphics Forum*, volume 31, pages 2528–2540. Wiley Online Library, 2012. 1, 2

[13] Y. Ke, X. Tang, and F. Jing. The design of high-level features for photo quality assessment. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 419–426. IEEE, 2006. 1, 2

[14] Y. Koyama, D. Sakamoto, and T. Igarashi. Crowd-powered parameter analysis for visual design exploration. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*, pages 65–74. ACM, 2014. 1, 2, 7

[15] W. Luo, X. Wang, and X. Tang. Content-based photo quality assessment. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2206–2213. IEEE, 2011. 1, 2

[16] Y. Luo and X. Tang. Photo and video quality evaluation: Focusing on the subject. In *Computer Vision–ECCV 2008*, pages 386–399. Springer, 2008. 1, 2

[17] MATLAB. *Statistics Toolbox Release 2012a*. The MathWorks Inc., Natick, Massachusetts, 2012. 5

[18] E. Peli. Contrast in complex images. *JOSA A*, 7(10):2032–2040, 1990. 5

[19] K. Subr, S. Paris, C. Soler, and J. Kautz. Accurate binary image selection from inaccurate user input. In *Computer Graphics Forum*, volume 32, pages 41–50. Wiley Online Library, 2013. 5

[20] L. Yuan and J. Sun. Automatic exposure correction of consumer photographs. In *Computer Vision–ECCV 2012*, pages 771–785. Springer, 2012. 1, 2

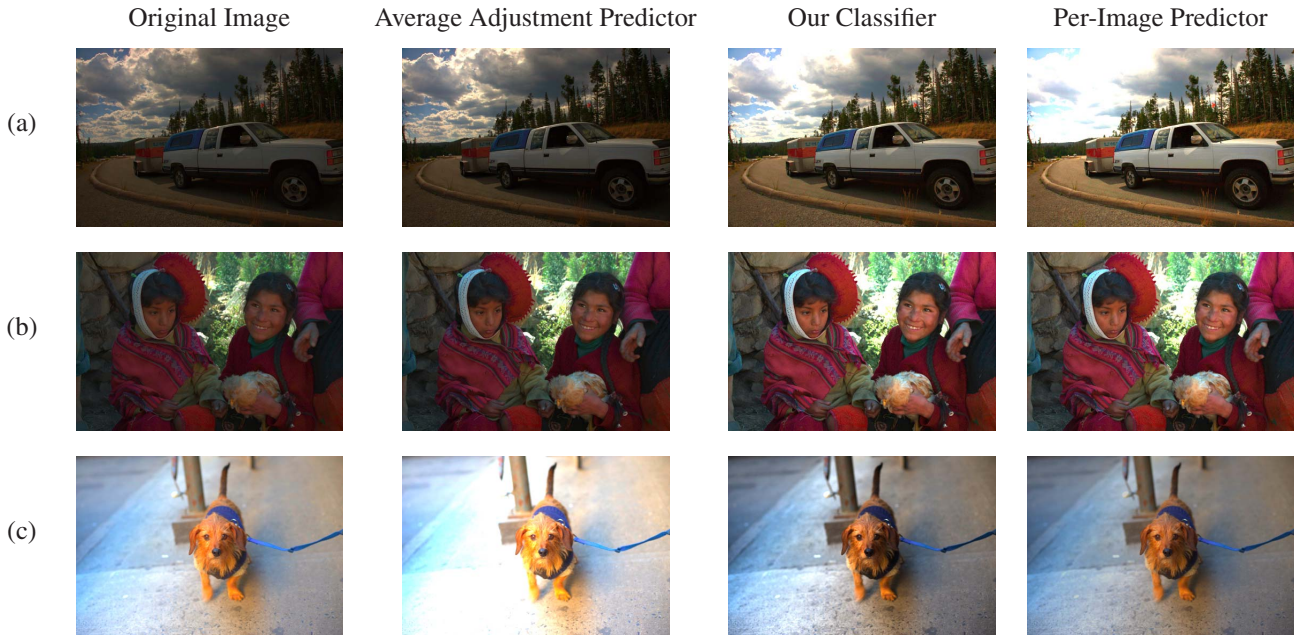

Figure 8: *Most Favorable Point Prediction.* Our prediction adapts well to the image content, while the average toy predictor returns always the same value, which fails when the image is not already properly exposed.

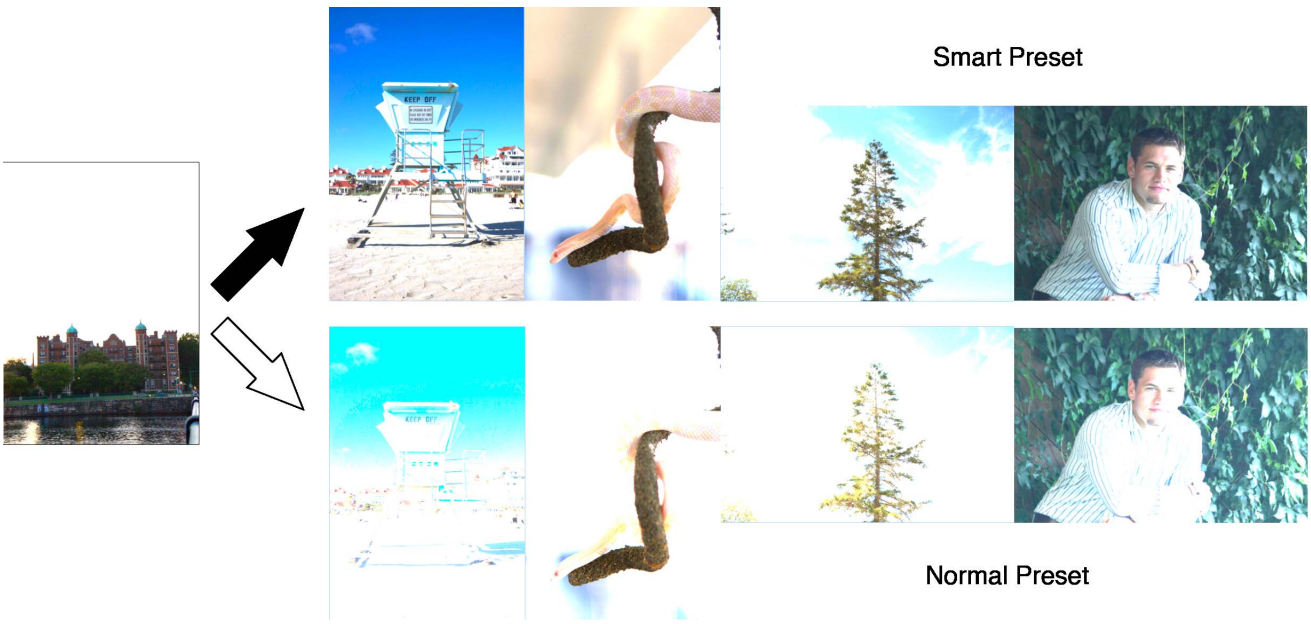

Figure 9: *Smart Preset vs Normal Preset.* The top row shows the result of our smart preset, and the bottom row the standard preset, which directly copies the adjustment to the destination image. Our algorithm predicts acceptable adjustments, which the smart preset uses to control the amount of adjustment done to an image. The resulting images exhibit a high brightness similar to that of the model while remaining aesthetically pleasing, even though dynamic ranges of the source and destination images are different.