ENHANCING PRODUCT IMAGES FOR CLICK-THROUGH RATE IMPROVEMENT

Rakuten Institute of Technology http://rit.rakuten.co.jp facebook.com/rakutenRIT twitter.com/RakutenRIT

Yeongnam Chae
Mitsuru Nakazawa
Björn Stenger

Abstract

This paper proposes a statistical method to enhance image quality in order to increase the click-through rate (CTR) of product images. We build a joint probability model of global image features for photos of different product categories. The images are modified in terms of brightness, contrast, and sharpness in order to increase the expected CTR. The effectiveness of the method is evaluated using a perceptual user study, comparing it to histogram equalization methods, and by conducting an A/B test over one week on the e-commerce site Rakuten Ichiba.

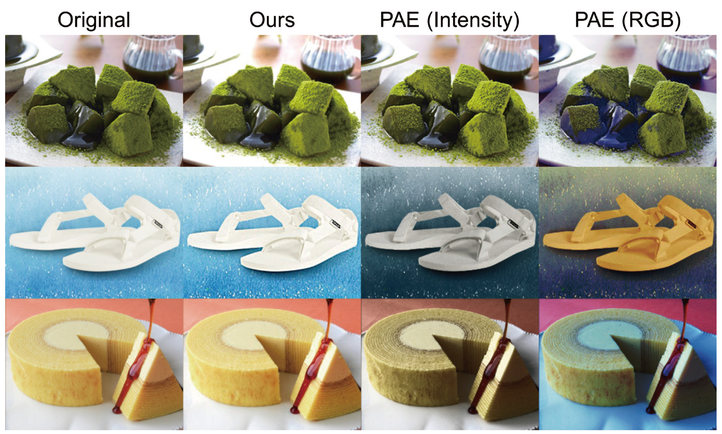

Image Enhancement Examples

Statistical Image Enhancement

The CTR of an image is affected by various factors. Here we are interested in how the image quality affects the CTR, and consider all other factors as nuisance variables. We formulate the task as finding a suitable function f_θ, with parameters θ which, when applied to an input image I leads to a high likelihood value of the CTR value, r:

$$\theta^* = \text{argmax}_\theta \; p(r|f_\theta(I)). \tag{1}$$

$$f_\theta(I_L) = (\alpha I_L + \beta) - \mathcal{G}_\sigma * I_L, \tag{2}$$

where I_L is the luminance channel of the image in Lab representation, and $\theta = (\alpha, \beta, \sigma)$ contains the scaling parameter α, the offset β and the σ to determine the variance in the Gaussian kernel $\mathcal{G}\sigma$. We take a learning approach to determine the parameters of the enhancement function. In order to estimate the perceptual effect of the image enhancement, we first extract global image features $\phi(I)$ and learn a mapping from these features to the expected CTR value

$$E(r|\phi(I)), \tag{3}$$

$$\phi(I) = (\varphi_b, \varphi_c, \varphi_s, \varphi_{sal})^T, \tag{4}$$

where φ_b is the mean brightness, φ_c is the image contrast, φ_s is the sharpness and φ_{sal} is the saliency area of image.

Local Neighborhood Weighting

We use weighting of the expected CTR value function in a local neighborhood of the input image features $\phi(I)$ to obtain the following objective function for an image I_j:

$$\mathcal{S}(\phi(I_j)) = w(\phi(I), \phi(I_j))\, E(r(\phi(I_j))), \tag{7}$$

by applying Gaussian weighting with zero mean and standard deviation σ = 0.01

$$w(\phi(I), \phi(I_j)) = \mathcal{G}_{\mu,\sigma}(||\phi(I) - \phi(I_j)||))/\zeta_w, \tag{8}$$

where ζ_w is a scaling factor

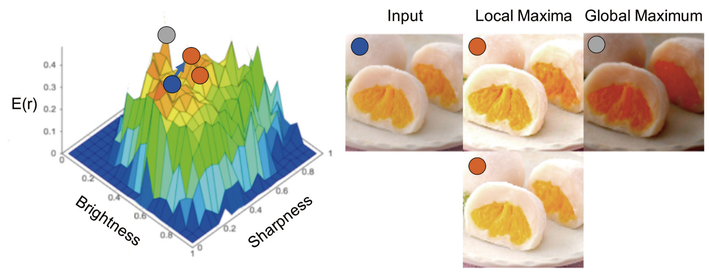

Nearby local maxima vs. global maximum

Preference survey and A/B Test

We collected 127,121 images of 26 different product categories, collected from an e-commerce site. We carried out two experiments to assess the performance of the method, a user preference study, and an A/B test. For the two experiments, we used 48 images of 270 × 200 resolution, selected from four product categories: fashion, food, interior products and sports goods.

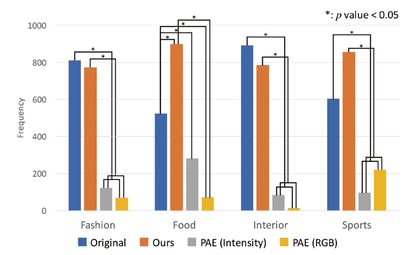

User preference study *A The proposed method performs well on the food and sports goods categories, while performing slightly worse on fashion and interior categories.*

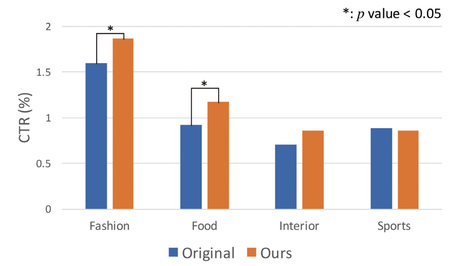

Result of A/B test *Improvements are observed for fashion, food, and interior product photos. Only results on fashion and food are statistically significant.*

Preference survey :

For each image we computed the enhanced image using the proposed method, and two enhancement methods based on HE. These images were displayed in randomized order to 148 users together with the question:
Which is the image that makes you want to purchase the product?

A/B Test :

We carried out an A/B test on an e-commerce website. The original images and enhanced images were shown at the same location of the page, and all other content of the page being identical. The number of impressions and the number of click-throughs were measured over the period of one week, yielding a total of 100,720 impressions and 1,119 clicks.

Conclusion

In this paper we have presented a statistical image enhancement method, adjusting image brightness, contrast, and sharpness, with the objective to increase the click-through rate of product images. We have shown in a user preference study that our method outperforms standard enhancement methods using histogram equalization. In an A/B test over one week we have measured the change of CTR over the period of a week and found an average increase of 15.7%, demonstrating the effectiveness of the method.