October 12, 2018

Improved White Balance Based on Multiple Photos of Same Subjects

Thomas Deselaers

Victor Carbune

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Improved white balance based on multiple photos of same subjects

ABSTRACT

This disclosure describes techniques to perform improved automatic white balance correction to digital images in a user's collection. With user permission, multiple photos that depict the same subjects are found in the user's collection in photos depicting a variety of different lighting and environments, and an automatic white balance correction is determined based on the multiple photos. The white balance correction can be automatically applied to new photos and/or older photos in a user's collection. Such features enable automatic white balance to more accurately applied and produce higher quality photos with accurate colors.

KEYWORDS

- white balance
- color correction
- photo library
- image enhancement
- photo subjects

BACKGROUND

Many cameras include an automatic white balance feature, in which the camera automatically adjusts the colors of the pixels of a captured image to produce more accurate colors. The white balancing attempts to compensate for a particular "temperature" produced by source lighting that may illuminate the scene of the image, e.g., reduce a yellow tone produced by yellow light, a blue tone produced by fluorescent light, etc. In some cases, the automatic white balance feature of a camera is inaccurate, producing a distorted or off-key color range in images taken with the camera.

<u>DESCRIPTION</u>

This disclosure describes techniques that adjust white balance (color correction) in digital images based on one or more other images in a user's collection. Described techniques determine a more accurate white balance by using multiple images that depict the same or similar subjects under different lighting conditions and camera settings. Metadata related to the images (time of capture, location of capture, settings of capture device, etc.) can also be used in white balance determination. For example, multiple such images are searched for and located in a user's collection. These images and their characteristics are used as inputs to a model to determine an accurate white balance for images depicting the subject. The techniques are implemented upon specific user permission to access the one or more other images and to perform color correction. Only the images and image metadata for which the user grants permission are utilized to determine white balance adjustments.

Many users capture digital images of the same subjects repeatedly, such as people known to the user and objects, e.g., in familiar locations. A user's image collection typically includes multiple images of the same subjects, where many of these images were captured in different locations or environments and under different lighting conditions, such as indoors at a party, outdoors in sunlight, outdoors in shade cast by any of various objects, etc. With user permission, the described techniques make use of such multiple images that show the same subject to determine white balance adjustments.

The described techniques include searching other images to identify the same subjects in the other images, when user provides consent, and determining one or more white balance settings for one or more images based on the other images. The white balance setting modifies

image colors to ensure that the same subject is shown with the same colors across all the images of that subject.

Obtaining Multiple Images Depicting a Subject

The techniques can be utilized, e.g., when a new image is captured and white balance adjustment is to be determined, or to determine white balance adjustments for a group of stored images of the user's collection. With user permission, when a new image is captured, other images in the user's collection that depict the same subject as the new image are identified. To adjust white balance for a stored group of images, the same subjects are identified among the stored group of images. The subject of an image can be one or more particular objects (e.g., person, animal, physical object, etc.), a landscape or other depicted setting, etc.

One or more machine learning models can run on a camera (or other capture device) to locate images that have the same (or similar) subjects. For example, with user permission, a face detector can be used for isolating regions in images that include faces and detecting faces that belong to the same person. An object classifier can be run to identify the classes of different objects in images. Segmentation models can be used to segment objects in images. In some cases, the machine learning models can provide information about the colors of detected objects in the images that can be used in a determination of white balance.

With user permission, the machine learning models for detecting types or classifications of subjects can be trained using all of the images in a user's collection of images, including the metadata of the images. In some examples, a model can be trained using both positive and negative examples of different classifications or types. For example, image metadata can be used as a factor in identifying that a detected subject is the same subject as in other images. A machine learning model can be provided metadata such as capture location, capture time, and

camera settings of the capture device, which can indicate similar subjects in particular images, e.g., if those images were captured at the same location or within a narrow time range.

If the same subjects are detected in images, the images depicting the same subjects can be clustered into different classifications such that each different subject is represented by an associated classification or identification. In an example of detecting faces, the identities of users of the faces need not be determined, since the system need only identify that depicted faces belong to a particular person, e.g., assigned an internal person classification or identification.

The described machine learning models can be implemented on a capture device, e.g., a camera or other mobile device, and/or can be fully or partially implemented on a server or other remote device in communication with the capture device.

In some cases, heuristics and rules can be used instead of or in addition to machine learning models. Particular metadata of the images can be associated with particular heuristic scores indicating a likelihood of image having same subject. For example, metadata such as geographic location of capture can indicate if two images were captured in the same location, indicating a greater likelihood that the same object is depicted in those images. Similarly, timestamp metadata can indicate time of capture of images, such that if the time of capture is within a threshold time range of each other, the images are more likely to depict the same subject (such as in a burst of images captured in a short range of time).

Determining and applying white balance correction to images

After multiple images of the user's collection which include the same subject have been identified, a white balance correction is determined based on those multiple images.

Consider a scenario where a new image has been captured and previous images of the user's collection that depict the same subject as the new image have been located. If a sufficient number of previous images depicting the same subject have been located (e.g., above a threshold number of images), white balance adjustment of the new image is triggered. Furthermore, white balance adjustments of the previous images is performed in some cases as described below, e.g., if such adjustment has not already been applied to the previous images, or to update/ improve white balance of the previous images.

If the previous images have been white balanced, then the colors of the subject is accurate in these images. The colors of the subject in the new image are compared to one or more colors in the previous images. For example, an average of the colors (or other parameter based on the colors) of the subject from the previous images is determined (e.g., an average of one of the colors of the subject). The corresponding colors of the new image subject are compared to this average. If the colors of the new image subject are different than the average colors, then the colors of the new image are adjusted using a white balance value such that the subject in the new image has the same colors as the average color from the previous images.

In some cases, only particular other images are used to determine the average color, and/or particular other images can be weighted higher in the determination of the average color. For example, the particular images can include images considered more likely to be closer in color to the new image, based on lighting, environment, and/or other conditions of capture. Metadata of the images is examined, and the other images that have the same or similar time of capture, location of capture, camera settings, and/or other metadata as the new image are used in the determination of the average color (while some of the other images may be excluded

from the average). Alternatively or additionally, other images that have similar metadata as the new image can be assigned a higher weight in the determination of the average color.

In another example, particular other images are weighted higher if it is known or estimated that these images have accurate white balance, or lower if white balance in these images is known to be inaccurate or unreliable. For example, colors of a previous image may have been manually corrected by a user, e.g., using a manual white balance operation. In some cases, metadata of an image can include a manually white-balanced value to indicate this condition. A manual correction is assigned a higher weight than automatic (device) color correction since the user's color preferences are exhibited directly by such a correction.

Other image metadata can be examined to indicate whether colors for a particular other image are likely to be accurate, and can indicate whether to weight that image higher or lower. For example, timestamp metadata indicating an image was captured at sunset can indicate that, if particular colors of the image are consistent with that time of day, that image can be weighted higher. Metadata indicating capture location and camera settings can be similarly used to weight particular other images. In another example, particular other images that have been marked or designated as having accurate colors by a user can be assigned higher weight.

Once the average color of the subject is determined, the new image is white balanced (e.g., normalized) to that color. For example, the pixel colors of the new image are adjusted so that the colors of the subject in the new image match the average color of the subject in the other images.

The white balance correction may have the most accurate result when applied to an image in a raw format. For example, if a raw version (e.g., uncompressed and unprocessed) of the new image is available, then the white balance correction can be applied to the raw version.

The above-described techniques can also be used to white-balance images that were captured at a previous time, e.g., the other images used in the white balance correction of the new image. For example, such images may have been previously white-balanced using a different technique that may be less accurate, or not white-balanced. This may occur, for example, if a person starts capturing images of a new subject (e.g., a new friend, etc.), and previous images may not yet have been white balanced according to described techniques, e.g., due to having an insufficient number of other images of the same subject. This technique allows uncorrected or inaccurately-corrected images to be improved in white balance over time, e.g., after several images of a subject have been captured.

White balancing the other images can be similar to white balancing the new image. For example, the color(s) of the subject in each of the images can be compared to the average color, and the colors of the image adjusted if needed.

In further examples, a group of images can be white balanced according to these techniques. Images having the same subject can be identified within a stored group of images. If a sufficient number of particular images depicting the same subject are identified within a group of images, white balancing for the particular images is triggered. The average color of a particular subject (e.g., weighted) is determined from the particular images as described above. Each of the images having the same subject is adjusted by white balancing the colors of the pixels of each particular image such that the subject in that image attains the average color.

Furthermore, stored images can be updated with more accurate white balance correction over time. For example, new images of a subject may be captured, allowing a more accurate white balance correction of other images depicting that subject based on those new images. The more accurate correction can be applied to the other (e.g., previously-stored) images of a

subject, e.g., after certain number of new images of that subject are captured or obtained by a device (e.g., loaded onto the device).

Some implementations can utilize other images captured by different cameras and capture devices (and/or different users) to determine white balance as described herein, if the owning user's consent has been obtained for the images accessed. For example, other images depicting the same subject that were captured by a first camera may have been white balanced. A transformation can be determined that converts a color (or average color) of the subject as captured by the first camera to a color appropriate for the second camera, taking into account differences in capture characteristics and different camera settings of the first camera and second camera. For example, an accurately white balanced subject captured by the first camera can be compared to an accurately white balanced subject captured by the second camera to obtain the transformation. In some examples, the transformation can convert an average color determined from the colors of the subject depicted in the first camera's images to an average color for use with images of the subject captured by the second camera.

Machine learning models that can be used in the system are trained and implemented upon user permission to access user data that serves as input to the models. Users are provided with options to indicate their permission or denial of permission for access to various data, e.g., images, image metadata, video, and other content in the user's image library, contextual factors such as time, location, application in use, etc. In using the models, use is made only of user-permitted data. One or more of the models are not implemented, if users deny permission. Model training is performed based on generalized data that is not attributable to individual users, and/or training is performed only locally on a user device with user data, e.g., using a federated learning approach.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user's activities, social network, or social actions, profession, a user's preferences, or a user's current location), and if a user device is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.
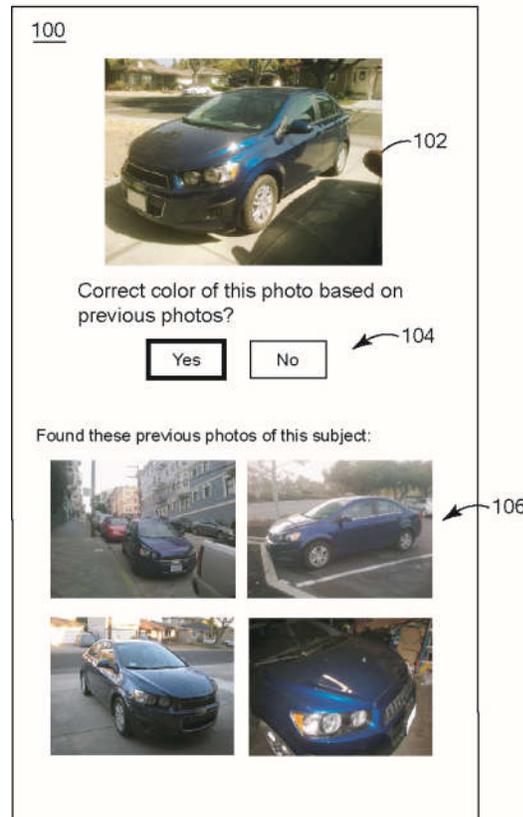
User Interface Examples

Fig. 1

Fig. 1 shows one example of a user interface (100) that can be displayed on a device, such as a mobile device, of an application that enables automatic white balance adjustment of images. For example, the user interface can be provided by an application that executes on a mobile device that includes a camera. In this example, a new image (102) was captured by the device and is displayed with a prompt (104) asking whether the user wishes to white balance the new image.

The user interface can also displays several other images (106) of the user's collection, which were identified by the system (e.g., by the mobile device or a connected server) as depicting the same subject as the new image and which may be used for white balance

adjustments for the new image. In this example, a car subject is detected in the new image, and the same car subject is detected in the other images. Furthermore, the other images depict the subject in different lighting conditions, allowing a more accurate white balance correction to be performed.

Upon receiving confirmation from the user, a white balance correction is applied to the new image based on the colors of the car subject in the other images, as described above. The colors of the previous images are used to adjust the colors of the new image. In some instances, the white balance correction can be performed automatically, without displaying the prompt or receiving user confirmation.
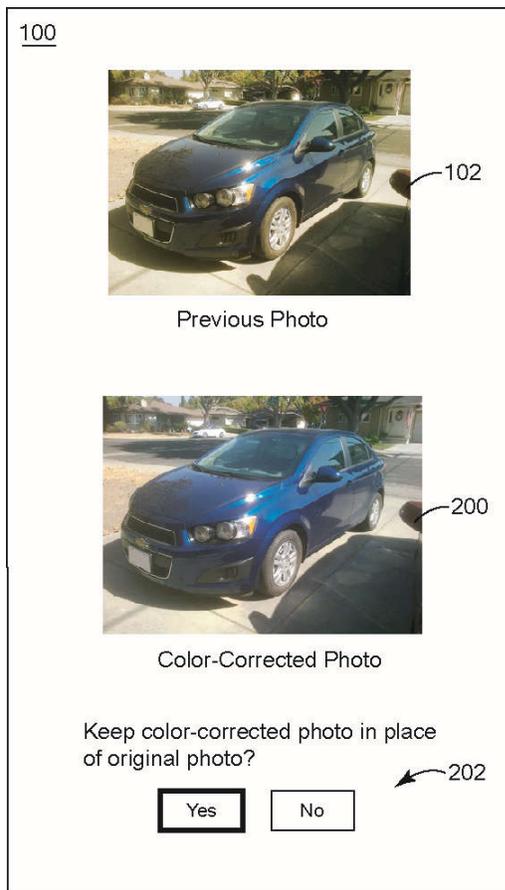

Fig. 2

Fig. 2 shows the user interface after determining the white balance correction for the new image. The original image is displayed, followed by a display of an adjusted image (200) resulting from the white balance operation. A prompt (202) is displayed asking whether the user wishes to replace the original image with the color corrected image, or keep the original image as a backup, etc.

CONCLUSION

This disclosure describes techniques for automatically determining and providing white balance correction of images based on other images in a user's collection. Colors of a particular subject found in the other images can be used to determine white balance values for newly-captured images that depict that same subject, enabling white balancing that is more accurate than existing automatic white balance techniques due to being based on a wider range of images showing the subject, including in different lighting conditions and environments. Further, more accurate white balance correction can be applied to stored images and updated based on new images that are later captured, such that image colors in stored images can be made more accurate over time.