May 2022

CONTENT-BASED DISPLAY COLOR ENHANCEMENT USING DYNAMIC MEMORY COLOR PROTECTION

Farnaz Agahian

Daniel Solomon

CONTENT-BASED DISPLAY COLOR ENHANCEMENT USING DYNAMIC MEMORY COLOR PROTECTION

ABSTRACT

A computing device (e.g., a smartphone, a laptop computer, a tablet computer, a smartwatch, etc.) may apply color enhancement (e.g., by modifying digital values of red, green, and blue additive primaries to increase color saturation of colors originally in the sRGB color space) to content (e.g., an image, a video, etc.) without causing objects within the content to appear unrealistic to a user. For example, the computing device may apply a machine learning model to label objects (e.g., based on image features) within the content. Based on the labels, the computing device may determine whether color enhancement should be applied to the corresponding objects or not (e.g., by looking up whether the label is on a list of objects to which color enhancement should not be applied). Based on this determination, the computing device may selectively apply color enhancement to some objects within the content. In this way, the computing device may increase color saturation of a part of the content (which a user may desire) while maintaining (from the perspective of a user) a substantially realistic appearance of the content overall, potentially improving the user experience.

DESCRIPTION

FIG. 1 below is a conceptual diagram illustrating a system 100 that performs content-based color enhancement in accordance with techniques of this disclosure. As shown in FIG. 1, computing device 100 includes a presence-sensitive display 102 (“display 102”), one or more processors 104, and one or more storage devices 106. As further shown in FIG. 1, storage devices 106 stores a color protection module 108, a color enhancement module 110, and a content repository 112.

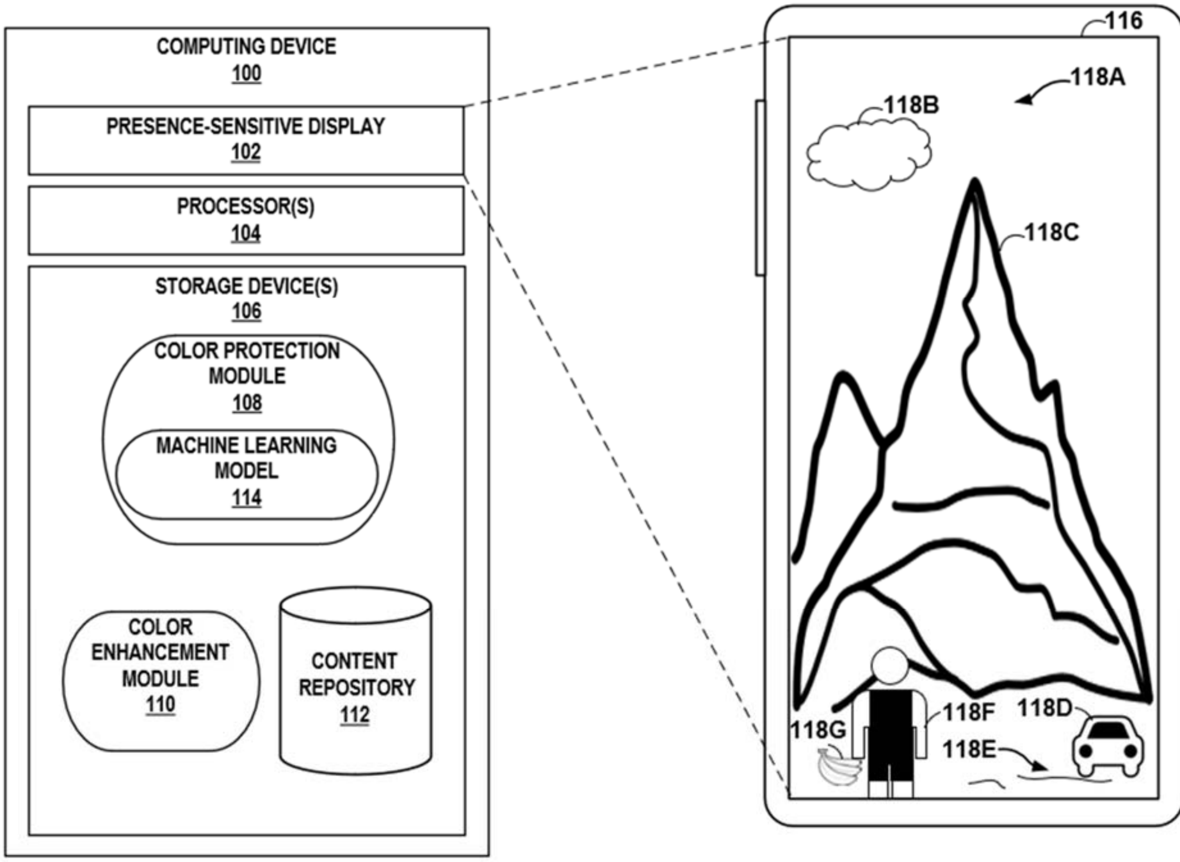

FIG. 1

Display 102 may function as an input device for computing device 100 using a touchscreen, pressure sensitive screen, an acoustic pulse recognition touchscreen or another presence-sensitive screen technology. Display 102 may function as an output device for computing device 100 using any one or more of display technologies (e.g., liquid crystal display, dot matrix display, light emitting diode (LED) display, mini-LED display, organic light emitting diode (OLED) display, electronic ink display, or similar display technology).

Display 102 may be configured to display a variety of colors. As used here, the color gamut of display 102 is defined as the range of colors display 102 can reproduce. In general, the color gamut of a conventional presence-sensitive display may be based on a Red-Green-Blue (RGB) color model (e.g., an additive color model in which the red, green, and blue primary colors of light are combined to reproduce an array of colors). In this disclosure, a color gamut based on the standard RGB model may be referred to as a sRGB color space. A sRGB color space may be defined by chromaticity coordinates of the red, green, and blue additive primaries.

Processors 104 may implement functionality and/or execute instructions associated with computing device 100. Examples of processors 104 include one or more of an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), an application processor, a display controller, an auxiliary processor, a central processing unit (CPU), a graphics processing unit (GPU), one or more sensor hubs, and any other hardware configure to function as a processor, a processing unit, or a processing device. In some examples, processors 104 may represent a system on a chip (SoC) that includes an integrated circuit for implementing one or more of the above referenced examples of processors 104, along with supporting memory and/or storage, and possibly various interfaces, modems, etc. as a single package.

Storage devices 106 may include one or more computer-readable storage media. For example, storage devices 106 may be configured for long-term, as well as short-term storage of information, such as instructions, data, or other information used by computing device 100. In some examples, storage devices 106 may include non-volatile storage elements. Examples of such non-volatile storage elements include magnetic hard disks, optical discs, solid state discs, and/or the like. In other examples, in place of, or in addition to the non-volatile storage elements, storage devices 106 may include one or more so-called “temporary” memory devices, meaning that a primary purpose of these devices may not be long-term data storage. For example, the devices may comprise volatile memory devices, meaning that the devices may not maintain stored contents when the devices are not receiving power. Examples of volatile memory devices include random-access memories (RAM), dynamic random-access memories (DRAM), static random-access memories (SRAM), etc.

Due to improvements in display technology, display 102 may have a color gamut larger than the sRGB color space. In these examples, display 102 may be referred to as a wide color gamut (WCG) display. Examples of WCG displays may include OLED displays. Although display 102 may be able to produce more saturated colors (e.g., vibrant tones) due to the WCG capabilities of display 102, media content (e.g., videos, images, etc.) may be processed for display on legacy sRGB displays. Thus, when such content is displayed on display 102 in these examples, the WCG capabilities of display 102 may not be utilized (e.g., because colors outside of the sRGB gamut are clipped or mapped to similar colors within a calibrated sRGB color gamut).

To address the above issue, computing device 100 may execute color enhancement module 110. Color enhancement module 110 may modify (e.g., for each pixel of display 102) digital values of red, green, and blue additive primaries to increase color saturation of colors originally in the sRGB color space. While a user of computing device 100 may generally prefer the color enhancement to a variety of content, color enhancement module 110 may modify the colors of certain objects (e.g., plants, food, dirt, human skin, animals, the sky, etc.) so much that the objects no longer appear realistic. For example, color enhancement module 110 may enhance the colors of an image of a blue sky, and as a result the color of the sky may change from blue to purple. Such “uncanny” color enhancements may be displeasing to a user, particularly when the user already has a memory or expectation of what color an object is supposed to be.

In accordance with techniques of this disclosure, computing device 100 may execute color protection module 108. Color protection module 108 may identify objects that have expected colors (sometimes referred to here as "protected objects") within content (e.g., an image, a video, etc.). Color enhancement module 110 may apply color enhancement to regions that do not contain protected objects. Additionally, color enhancement module 110 may not apply color enhancement to regions that do contain protected objects. In this way, color protection module 108 may protect objects that have expected colors from color enhancement that may otherwise cause the protected objects to look unrealistic.

As described above, color protection module 108 may identify objects that have expected colors within content for display by display 102. Color protection module 108 may include a machine learning model 114 trained to label objects within the content. A training module (e.g., executing at a computing system, such as a remote server) may train machine learning model 114 on a training dataset that includes training examples labeled as belonging (or not belonging) to one or more classes. In some examples, the training module may use supervised learning techniques to train machine learning model 114 to perform classification.

Machine learning model 114 may perform one or more computer vision techniques (e.g., semantic segmentation, instance segmentation, etc.,) to label objects within content. For instance, machine learning model 114 may label pixels of an image with a corresponding class (e.g., grass, tree, sky, cat, building, car, bike, person, etc.) of what machine learning model 114 predicts is being represented. Machine learning model 114 may label objects based on image features (e.g., a piece of information about the content of an image). Image features may include properties such as corners, edges, regions of interest points, ridges, patterns, etc. In some cases, the objects labeled by machine learning model 114 may be image features. For instance, machine learning model 114 may process image 116 to extract image features (e.g., through a form of content segmentation) and apply labels to each of the image features.

A contiguous set of pixels having the same label may form a segment (i.e., a region of pixels of content) representing an object identified by the label of the contiguous set of pixels. As an example, machine learning model 114 may label pixels of an image 116 stored in content repository 112. In this example, machine learning model 114 may label a segment 118A as "sky," a segment 118B as "cloud," a segment 118C as "mountain," a segment 118D as "car," a segment 118E as "dirt," a segment 118F as "human skin," a segment 118G as "bananas," (collectively, "segments 118"), etc.

Based on the labels provided by machine learning model 114, color protection module 108 may determine whether segments 118 are protected objects or not. For example, color protection module 108 may determine whether any of segments 118 has a label on a protected object list. Examples of objects on the protected object list include, but are not limited to, human skin, fruits, vegetables, plants, sky, clouds, and grass. Accordingly, color protection module 108 may determine that segments 118A-118B and 118F-118G in image 116 are protected objects and that the remaining portions of image 116 (e.g., segments 118C-118E, etc.) are not protected objects. Based on this determination by color protection module 108, color enhancement module 110 may apply color enhancement to the regions of image 116 containing segments 118C-118E but not apply color enhancements to the regions of image 116 containing segments 118A-118B and 118F-118G.

Computing device 100 may locally execute a version of machine learning model 114 while a version of machine learning model 114 is trained on a computing system. In some examples, computing device 100 may update the version of machine learning model 114 installed at computing device 100 by downloading a more recent version of machine learning model 114 (e.g., that has been further trained on the computing system). In this way, the accuracy of machine learning model 114 at labeling objects in content may improve. Similarly, color protection module 108 may update the protected object list by downloading a more recent version of the protected object list. In this way, the ability of color protection module 108 to maintain a substantially realistic appearance of content may improve.

A brief overview of example machine-learned models and associated techniques has been provided by this disclosure. For additional details, readers should review the following references: Machine Learning A Probabilistic Perspective (Murphy); Rules of Machine Learning: Best Practices for ML Engineering (Zinkevich); Deep Learning (Goodfellow); Reinforcement Learning: An Introduction (Sutton); and Artificial Intelligence: A Modern Approach (Norvig).

It is noted that the techniques of this disclosure may be combined with any other suitable technique or combination of techniques. As one example, the techniques of this disclosure may be combined with the techniques described in Chinese Patent Application Publication No. CN112887582A. In another example, the techniques of this disclosure may be combined with the techniques described in Arjan Gijsenij et al., "Computational color constancy: survey and Experiments", IEEE Transactions on Image Processing, September 2011. In yet another example, the techniques of this disclosure may be combined with the techniques described in PCT Patent Application Publication No. WO2019012112A1. In yet another example, the techniques of this disclosure may be combined with the techniques described in Simone Bianco and Raimondo Schettini, "Color Constancy Using Faces", IEEE Conference on Computer Vision and Pattern Recognition, June 2012.