Consistent Display of Clemson Brand Colors Using Artificial Intelligence

Erica Walker, Hudson Smith, Emma Mayes,
John Paul Lineberger, Michelle Mayer, and Andrew Sanborn

Keywords: Artificial Intelligence, Color Management, Color Measurement

Abstract

When watching a broadcast of an athletic event on a digital display, fans notice when there are inconsistencies in brand colors. Brand colors are a valuable asset built through consistency and repetition. The accurate display of color is impacted by factors including the capture source(s), display screen technology, and ambient lighting which can change drastically during a live event due to changes in time of day and weather. Environmental shifts require adjustments to the video feed to maintain visual consistency. In the current workflow, this is done manually by a technician who color-corrects up to two dozen incoming camera feeds in real-time. The primary question of this research is: Can ColorNet, a neural network-based algorithm, automatically color correct for accurate and consistent display of brand colors in real-time video without impacting non-targeted colors? ColorNet is a patent-pending artificial intelligence (AI) technology that applies a machine learning model to adjust each video frame pixel-by-pixel to produce a color-corrected video output. For this study, ColorNet is demonstrated using Clemson University's brand color, Pantone 165 (orange). The model was trained using a collection of corresponding original and color-corrected frames from Clemson athletic events. Manual color correction was completed using Adobe Premiere Pro to produce this dataset. The current model is able to adjust only the targeted brand colors without shifting surrounding colors in the frame, generating localized corrections while adjusting automatically to changes in lighting; this is validated through analysis of the impact of ColorNet adjustments on the full color spectrum. The production viability of ColorNet was demonstrated through alpha and beta tests, and a Delta E (dE) color distance analysis was used to determine if the adjustment for the targeted colors was within a reasonable tolerance of the brand specification. Further progress is being made to expand the model to additional brand colors and explore applications beyond sports broadcasting.

Clemson University

Introduction

Consistent use of brand colors and assets such as a logo or typeface is an important aspect of a company's marketing plan. For example, McDonald's red and yellow colors are instantly recognizable. These colors are quickly recognizable because of their consistency across all of the consumer touchpoints from an ad on TV to the packaging that your food comes in. Consistent brand colors across all visual mediums help to build brand recognition and consumer loyalty (Budelmann et al., 2010; Chang and Lin, 2010). Therefore, brand color specifications are provided in multiple formats, including Pantone numbers, RGB values, hex code, and CMYK values. Like any other corporate brand, sports teams also need to control their brand colors. Fans notice when the broadcast of a sporting event fails to consistently display the team colors on screen (Conti and Walker, 2019) and they express their displeasure on websites such as Twitter, figure 1.

Figure 1: *Fans expressed displeasure on Twitter regarding how uniform colors looked on-screen during a football game*

The display of brand colors on screen is impacted by many factors including the screen technology and gamut limitations, the video input source, and the ambient lighting. Most of these can be adjusted prior to the start of a game, but during an outdoor event, ambient lighting changes continuously as the daylight and weather shift. To maintain visual continuity, the color output needs to be manually adjusted throughout the event. Human perception also affects what appears as accurate to the viewer. For instance, if part of a jersey is in a shadow, the brand color may need to shift to accommodate a change in lighting. If it still matches the specified RGB values exactly, it will appear over-manipulated because a perfect brand color match does not account for the visual impact of the shadow on the jersey.

The current standard for managing color and lighting adjustments in television broadcast of sports programming involves a technician who adjusts footage from all of the video feeds in real-time. In a typical setup, the technician views all of the incoming feeds as a grid on a set of screens and a smaller, color-accurate monitor to make any necessary adjustments to an individual camera feed. The technician wants to match the brand colors during this process, but also considers overall consistency of color and brightness across all the camera feeds. The viewer is more likely to notice inconsistencies between each camera feed rather than be distracted by consistent, but inaccurate brand colors. Before the game starts the technician sets a white balance for each of the cameras. Although setting the white balance could be repeated during the game as conditions change, it would be difficult to repeat the process in the middle of an event. Instead, the technician adjusts the aperture and the red, green, and blue values to control the overall color and brightness. These changes modify the color across the entire frame and not just the pixels that should display the correct brand color.

Since this process of adjusting the colors on the fly alters the entire frame, it is difficult to target the team's brand colors without negatively impacting the other colors in the frame. In response to this challenge, the research team developed ColorNet to apply on-the-fly color correction that only targets appropriate pixels within each frame. The primary question of this research is: Can ColorNet, a neural network-based algorithm, automatically color correct for the accurate and consistent display of brand colors in real-time video without impacting non-brand other colors?

Methods

ColorNet is a patent-pending solution based on computer vision technology developed to address this color management challenge. Frames with incorrect color are run through the system and color-correct frames are output. The corrected frames have only the targeted colors adjusted and the rest of the frame is left unchanged. ColorNet uses a supervised machine learning (ML) model which learns to map from the input to output frames based on a dataset of examples.

ColorNet was trained on how to color correct video frames by giving it a dataset of paired raw and brand-color-corrected frames. The current version of ColorNet was trained on thirteen two-minute clips from a single Clemson University Football game. These were manually color corrected in Adobe Premiere Pro using the Lumetri Color panel and a brand-accurate color matte for reference. Objects in the frame that should be Clemson Orange such as the jerseys and helmets were adjusted to Pantone 165 or a perceptually natural approximation based on natural highlights and shadows. Using the color picker tool a specific color range was selected and adjusted to match Clemson's brand colors. The color picker selection targeted the brand-color changes and avoided unnecessary adjustments to areas of the frame that should not be altered like grass, skin tones, and the opposing team's jerseys. Based on the color correct frames, a set of corresponding "correction masks" were created that specify the numerical pixel difference between the color correct frames and the original frame.

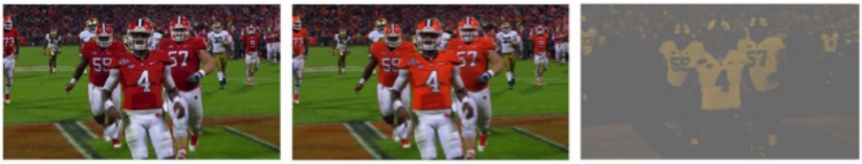

Figure 2: *An example of a color-incorrect image, its corresponding color correct image, and the "correction mask" created based on their difference*

In order to minimize the similarity between samples while still having a sufficiently large training dataset, every 14th frame was sampled from the thirteen two-minute clips resulting in a total dataset size of 4,291 pairs of original and corrected frames. From these, a customized, simple neural network model architecture named ColorNet was developed. This model was trained using the machine learning library Pytorch on 3,432 frames leaving the remaining 859 frames to validate the resulting model. The model was trained using a single GPU compute node on the Clemson University Palmetto Cluster. After multiple iterations, the current ColorNet architecture effectively color corrects Pantone 165, Clemson Orange, with only 62 parameters. This small size of this model results in very fast inference rates as necessary to color correct video feeds running at 60 frames per second. We evaluate model performance both qualitatively and quantitatively on the remaining 859 images.

Currently, a similar process is being used to expand and refine ColorNet to include a secondary brand color and address unfavorable color shifts in the initial prototype. Since the original dataset consisted of clips pulled from a single game, it is important to expand the dataset to include a wider variety of situations that the algorithm might encounter such as differences in weather, ambient lighting, and opposing team colors. In order for ColorNet to adjust for a secondary brand color, a dataset including the second color with paired original and corrected frames was developed and added to the training process. Throughout the process of expanding and refining the model, the effectiveness of the algorithm was evaluated using Delta E (dE) CIE 2000 to determine which colors were being modified and confirm that only brand-specific colors were being targeted.

Discussion

Model Testing Results

To quantify ColorNet's performance, we calculate the Structural Similarity (SSIM) and Mean Absolute Error (MAE) metrics of the model on the 859 withheld images. Both metrics quantify the similarity between the predicted correction mask and the correction mask produced by a human annotator. The SSIM metric ranges from -1 to 1 with 1 corresponding to perfect similarity. The current version of ColorNet achieves an SSIM of 0.905. In the case of images with pixel values coded on a scale from 0 to 1, the MAE metric ranges from 0 to 1. ColorNet achieves an MAE of 0.133. When evaluating these results, it's important to keep in mind that in the color correction task, the majority of input pixels should receive no corrections. Thus, a model which simply predicts no corrections can achieve a good score according to SSIM and MAE. We are currently working to address this issue through the development of customized metrics for image correction. For now, we point to the qualitative results below as confirmation that the model is performing better than this no-corrections baseline.

In order to quantify and refine ColorNet, we developed a customized measuring tool to describe how selective the model is when adjusting colors. Ideally, we want the model to only adjust values that are in the tonal range close to the specified brand color and leave all other colors on the screen alone. It is especially important from a branding standpoint to ensure that ColorNet is not negatively impacting the brand color of the opposing team.

To see what colors are being modified by ColorNet and by how much, the python-colormath module was used to import a tool that calculates the difference between two RGB values using the formula for Delta E (dE) CIE 2000 (Sharma et al., 2005). A color bar was built by incrementally changing single RGB values so that each horizontal pixel in the color bar has a unique RGB value. Creating the color bar using this process allows us to assess the impact of ColorNet on a wide range of individual colors in the RGB space.

After the color bar was created, each color was pulled from the bar and used to fill a 10 pixel x 10 pixel, single-color square which was then passed through the model. Then the dE was calculated between the original color matte and the resulting color matte output by the model. Results for several specific colors are shown in , figure 3. A grey correction mask indicates that there was no shift in the RGB values between the original and corrected color square; while a tan mask indicates the degree of the resulting color adjustment. As expected, the model noticeably shifts color for input colors in the orange spectrum. Undesirably, the model also shifts colors for the magenta and pink inputs.

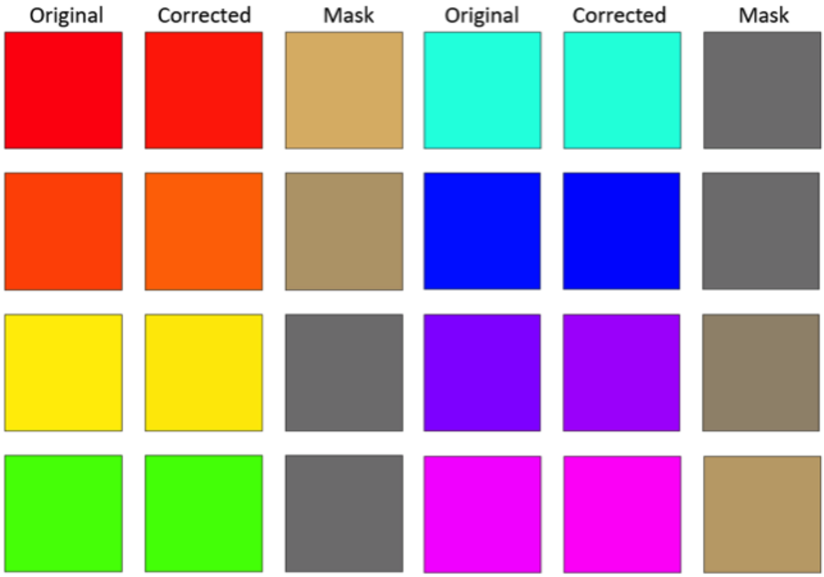

Figure 3: *Original 10x10 pixel input swatches selected from the RGB color bar, Corrected swatches as output by the model, and the corresponding correction masks.*

To get a more comprehensive look at the effect on color, the original color bar was then displayed on top of a new color bar corresponding to the outputs generated by ColorNet, figure 4. Above the color bars is a plot of the dE color shift for the colors shown below the horizontal axis. The peaks in dE correspond to the largest differences between the original and model-generated colors, which happen in the magenta and orange segments of the color spectrum. The irregularity of the curve is characteristic of machine learning in that the correlation between the input and output is not always intuitive but rather based on subtle intricacies found in the dataset used to train the model.

By evaluating the dE, it was discovered that the model currently shifts magenta and purple. This is an inherent limitation of machine learning as the model had not previously been given training data that contained these color values, so it produced unexpected results. The training dataset is now being expanded to include footage from a “Pink Out” as well as a Purple Military Appreciation game so that ColorNet will be better prepared to handle color-correction in this tonal range. Expanding the type of footage in the training dataset should make ColorNet more selective and precise when shifting colors.

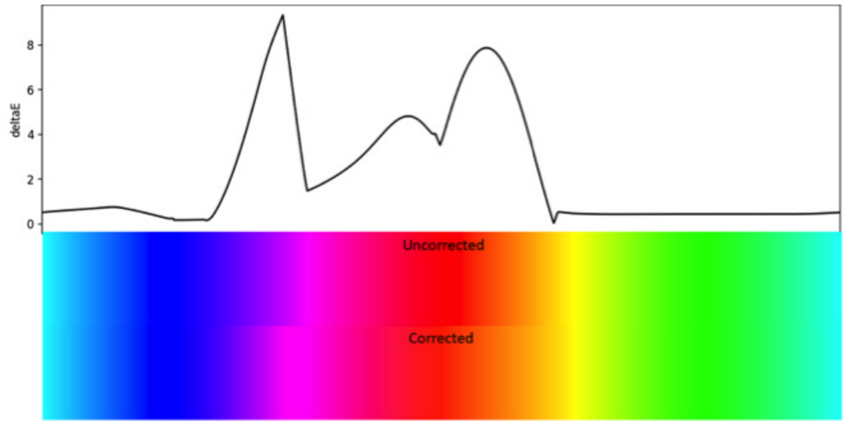

Figure 4: *The corrected color bar on the bottom, original on top, and corresponding Delta E above*

Beta Test Results

Before implementation of ColorNet in production environment, tests must take place—not only of the software for the color correction algorithm but also of the hardware that the system will be running on. The current production-level prototype uses off the shelf, consumer hardware which is able to perform 1080p color correction at 60+ FPS.

An important step in ensuring that the model can perform in a broadcast setting was testing the hardware for performance limitations. To accomplish this, a beta test was performed where ColorNet was implemented into the current video board system and displayed on the jumbotron in Littlejohn Coliseum during a live basketball game. This allowed testing of the software and hardware capabilities in a lower-risk and more controlled, indoor environment. This athletic event was chosen for the beta test because a football game has a large live audience and additional television broadcast considerations.

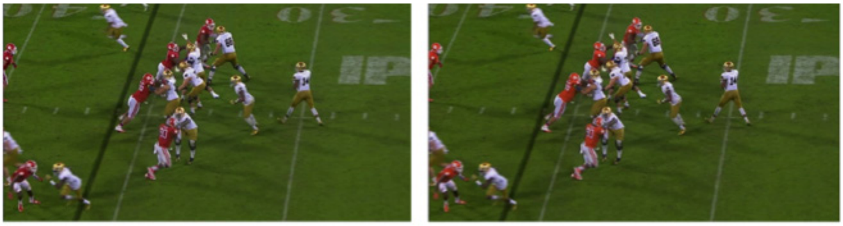

Figure 5: *Uncorrected frame displayed on left and ColorNet corrected image displayed on the right*

The success of the beta test was analyzed based on two metrics: acceptable color correction and hardware performance. Acceptable color correction was measured visually on the displays. Although we had demonstrated that ColorNet performs visually consistent and accurate color correction when given images and videos from previous Clemson football games (Figure 5), the model had not been trained or tested in the basketball, indoor environment.

During the beta test, the model made the team's basketball jerseys and other pixels containing Pantone 165 within the frame appear brand-accurate on the jumbotron without altering other colors within the frame, Figure 4. The graphic elements along the left and bottom of the display did not run through the model and were brand color-correct. The color correction on clothing was satisfactory, but the paw painted on the court appeared to be over-corrected. In order to address this, the dataset used to train ColorNet needs to be expanded to include footage from basketball games to better correct for the reflective surface of the floor and the indoor lighting present in this facility. Introducing basketball clips into the dataset should rectify this issue.

Figure 6: *ColorNet corrected footage is displayed on the jumbotron during the beta test at a Clemson women's basketball game*

The second metric for the beta test was hardware performance. During the beta test, the production model processed a continuous 60 fps feed with approximately 1.5-second latency, which is an acceptable delay during live events. However, the hardware system was hitting memory capacities every 10 to 15 minutes, so it had to be restarted whenever the video feed paused on the jumbotron for the display of pre-made static graphics. After the beta test, this issue was addressed by adjusting the video stream buffer and now the system can run indefinitely with no memory issues or reboots necessary.

Conclusion

The ColorNet model automates a previously subjective manual process and is scalable for real-time color correction. Although ColorNet will benefit from further training and refinement to make the color correction process increasingly precise, the beta tests and dE evaluation process shows proof of concept that ColorNet is a novel way to address brand color consistency for display on screen. Though the current ColorNet model specializes in correcting for Clemson Orange (Pantone 165), in principle, the technology can be trained to address any brand color. Training the model to correct for additional colors only requires developing a corresponding dataset for training. Currently, the ColorNet model is being expanded to also color correct for Clemson Purple, or Pantone 268 C, the University's other primary brand color.

The nature of the model allows for expansion into other applications beyond athletics where brand accurate color on screen is important. ColorNet has the capability to correct color in any digital display context where a brand would want to closely regulate the consistent display of their color(s). It should be noted that ColorNet does not override display-based color settings, but only provides a color-corrected feed. The screen settings could still negatively impact color for the viewer. Additional markets may include medical, film and television, display technology, professional color correction software, natural resource drilling, and others.

Acknowledgments

This research was supported by a Clemson University Technology Maturation Grant.

References

Budelmann K, Kim Y and Wozniak C (2010) Brand Identity Essentials: 100 Principles for Designing Logos and Building Brands. Rockport Publishers. Available at: https://play.google.com/store/books/details?id=yTI4xVI1MhQC.

Chang W-L and Lin H-L (2010) The impact of color traits on corporate branding. African Journal of Business Management 4(15). Academic Journals: 3344. Available at: https://pdfs.semanticscholar.org/6a4c/5c9c5744648165e0a10a0c4a61c0b066a8df.pdf.

Conti B and Walker EB (2019) Are You for the Orange Team or the Red Team? An exploratory study on displaying brand colors on large format LED screens at live sporting events. Visual Communications Journal 55(1).

Sharma G, Wu W and Dalal EN (2005) The CIEDE2000 Color-Difference Formula: Implementation Notes, Supplementary Test Data, and Mathematical Observations. Color Research & Application 30: 21–30. Available at: http://www2.ece.rochester.edu/~gsharma/ciede2000/ciede2000noteCRNA.pdf.

Chang, W.L. & Lin, H.L., 2010. The impact of color traits on corporate branding. African Journal of Business Management, 4(15). Available at: https://pdfs.semanticscholar.org/6a4c/5c9c5744648165e0a10a0c4a61c0b066a8df.pdf.

Krizhevsky, A., Sutskever, I., & Hinton, G. E., 2012. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems 25 (pp. 1097–1105). Available at: http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-Convolutional-neural-networks.pdf.

Rawat W, Wang Z, 2017. Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review. Neural Comput. Available at: http://dx.doi.org/10.1162/NECO_a_00990.

Wang Z, Bovik AC, Sheikh HR, Simoncelli EP, 2004. Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process [Internet]. Available at: https://www.ncbi.nlm.nih.gov/pubmed/15376593.