ColorNet: Use Case of Artificial Intelligence In Sports Broadcasting

Michelle Mayer, Mason Fox, Jake Liguori, Anthony Dempsey,
Jeremy Spooner, Erica Walker, and Hudson Smith

Keywords: automation, artificial intelligence, color management, screen, video

Abstract

ColorNet is a machine learning model developed at Clemson University that color corrects targeted portions of live video feeds in real-time. The original ColorNet 1.0 was developed to solve the problem of Clemson orange appearing incorrectly during football broadcasts. The machine learning algorithm can detect which pixels on the screen need color adjustment and shift those pixels to the correct color without noticeable latency in the broadcast and without negatively impacting surrounding colors. Production beta tests proved successful at correcting Clemson orange during live sporting events. ColorNet 2.0 was proposed as an improved solution that could correct any specified team color using image segmentation. The goal of ColorNet 2.0 is to use a reference color as an input and segment out portions of the screen containing that color. Then, the technician can adjust those areas of the screen according to the desired color specifications. This new approach would allow the model to be more universally applicable because it does not require new training data for each color the user wants to correct. It would also allow technicians more control over the final appearance of the brand colors. The current work presents a novel data augmentation strategy that synthetically expands the available training data to include all ACC colors. Future work will focus on the development of the neural network architectures needed to accurately automate the segmentation of the targeted brand colors.

Introduction

Sports broadcasts are some of the most-watched events on national television (2014 U.S. Online Video Rankings, 2014; World Heritage Encyclopedia, n.d.). Companies want to ensure viewers see a proper representation of their brands (Jin et al., 2019).

Clemson University

Through the use of logos, slogans, and team colors, sports brands build a strong relationship between fans and their teams. One way to do this is to have brand color consistency (Chang & Lin, 2010; Moser, 2003). The artificial intelligence-based software, ColorNet, aims to provide brand color consistency throughout live broadcasts without negatively impacting surrounding elements on the screen.

Various factors impact how the viewer sees their favorite sports team's colors on screen including weather, lighting, and camera settings. The ColorNet research team, composed of students and faculty from Clemson University, noticed this was a problem when watching the University's football team in what seemed to be uniforms from the opposing team (Mayes, E. et al., 2020.; Walker, E.B. et al., 2020). Fans across social media noticed this discrepancy as well and complained on social media (Figure 1). One of the most distinct cases occurred during the 2018 "Purple Out" game. Despite the players' uniforms being all purple, they appeared to be wearing blue to fans watching the live broadcast (Figure 2).

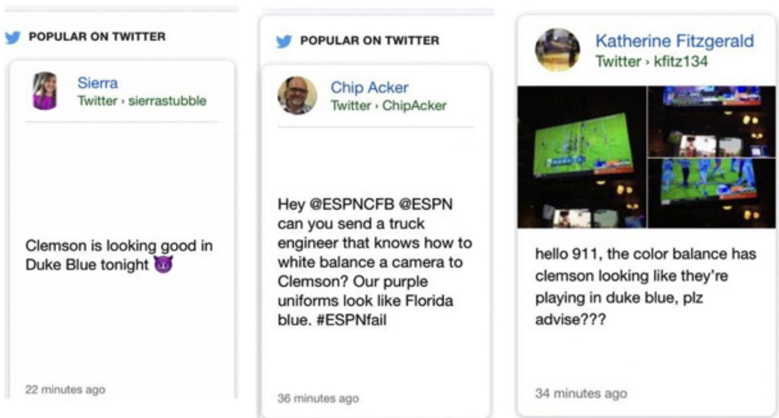

Figure 1: *Different fan reactions to the improper brand colors*

Figure 2: *The left image shows the live broadcasts footage. The right image shows corrected color*

In the current broadcasting system, changes can be made to color correct the picture by adjusting red, green, and blue levels in the image. However, these changes apply to the entire frame, rather than selected regions such as the player's uniforms or other areas that contain the brand color. Through the patented ColorNet system, changes are applied to the brand colors across certain pixels without affecting other areas.

Currently, the system only addresses Clemson brand colors because Clemson game footage was used for the machine learning data set; since the current build of the software, ColorNet 1.0, needs large amounts of manually created data from each brand color, this model is not ideal for use by other teams. Therefore, building on the strengths of the original software, ColorNet 2.0 will implement segmentation to select different colored regions within each frame and make the appropriate color adjustments based on whichever teams are playing. While this paper will not present the results of this new technique, the design philosophy and decisions used to pursue these advancements will be discussed in depth.

Methods

ColorNet 1.0 was the first successful, but limited, approach to AI-based brand color correction. The ColorNet 1.0 model uses regression to predict the correct RGB values of each pixel. It begins with a Convolution Neural Network (CNN), which is trained using paired frames from original footage: an uncorrected and the manually corrected version of each frame. The network learns to predict the difference in RGB values from the original footage to the corrected. Consequently, the network determines where differences occur, such as the orange in the team jerseys, and where no corrections are needed like the green of the turf field. This allows for targeted correction of the brand color without affecting other colors in the image (Mayes, E. et al., 2020.; Walker, E.B. et al., 2020).

By generating predictions for the sequence of original video frames, ColorNet 1.0 can be used to correct an entire video. Likewise, using the model in real-time allows for a corrected live broadcast. In order to seamlessly fit into the broadcasting workflow, the ColorNet model must be capable of applying its corrections almost instantaneously.

While ColorNet 1.0 was successful in adjusting Clemson's brand orange in live broadcasting, it has several limitations. First, ColorNet 1.0 can only adjust one brand color. If the opponent's brand color needed correction, this would require a separate model trained only on data from that team's colors. This would be a tedious and computationally expensive practice to apply similar models to several different teams. Second, ColorNet 1.0 does not allow for manual adjustment by the production team. Any machine learning model, especially regression, will be inherently imperfect. While the underlying Neural Network has been designed to produce the most accurate adjustments possible, the possibility of mistakes always exists. Consequently, a producer running a broadcast with ColorNet 1.0 might notice the corrections are not as accurate as they need to be, but have no ability to adjust them manually.

ColorNet 2.0 was designed to address these limitations. Like ColorNet 1.0, ColorNet 2.0 is a machine learning model using a Convolutional Neural Network (CNN). Rather than predicting the difference in RGB values between original and corrected frames, ColorNet 2.0 will find all pixels that are similar to a target color which is passed as an additional input into the model. Essentially, it finds segments of the input image near the specified input color in the color spectrum.

Segmentation is widely used for object detection in other fields where it identifies and labels different parts of an image (Minaee, S. et al., 2021). Here, we have no specific object to segment, but instead a chosen part of the color spectrum. It is important to consider that lighting and positioning will result in several different shades of the same brand color appearing in the image (e.g. the jersey of a player on a shaded part of the field will show as a darker shade of orange than the jersey of a player standing in the sun). Similar to the approach in ColorNet 1.0, this is addressed through a series of convolutions within the CNN. As the model analyzes each frame in the live footage, it uses these convolutions to learn underlying features and then captures the appropriate range of the target color in the segmentation.

In ColorNet 1.0 the network was trained using manually corrected frames. Since ColorNet 2.0 is not concerned with the differences in RGB values between frames, but instead needs to segment the target color, the network is trained using black and white masks. If the target color is Clemson orange, a new image is created in which every pixel that was the target orange is now white while every other pixel becomes black. This transforms our original image to a mask where each pixel is no longer a RGB value, but rather a binary value: a 1 for the target color regions and 0 all other areas of the image (Figure 3).

As mentioned previously, a limitation to ColorNet 1.0 is that each brand color that needs to be corrected requires a new model to be trained using another manually created dataset with that team's colors. This was also a constraint in the development of ColorNet 2.0 because the dataset used to train the model is exclusively Clemson footage that lacks a diverse set of colors that the model should be capable of segmenting. These shortcomings are addressed through an augmentation technique in which an existing brand color within a frame is adjusted to a randomly chosen Atlantic Coast Conference (ACC) color. This expands the dataset to include all ACC colors, including those that are not naturally present within the given footage.

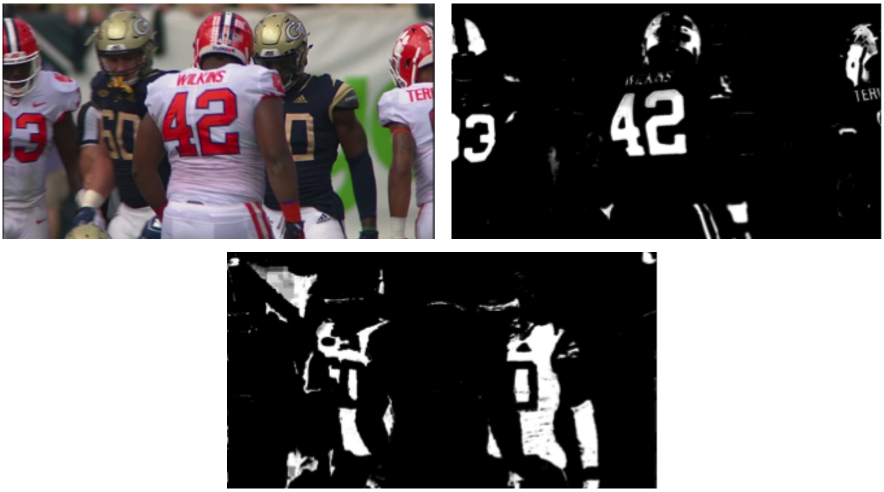

Figure 3: Segmentation masks for selecting Clemson orange (bottom left) and Georgia Tech navy blue (bottom right)

This augmentation process occurs exclusively during model training to expand the diversity of the training dataset. After being randomly selected for augmentation, the original frame is used alongside the true target mask to segment out which regions will be modified. All pixels in the original frame with an associated value of 1 in the mask (white) are adjusted to a predetermined, random ACC color. All other pixels remain unchanged. This results in a final frame in which all of the original brand color has been replaced with the brand color of a different, existing team. Upon completion of training, the model will have been exposed to a variety of unique colors outside of the original dataset.

Figure 4: Augmentation results from changing Clemson orange to blue

Through the training process, the network learns where the target color is likely to be. Then, the model can receive an uncorrected image and assign each pixel a probability of being the brand color. While the threshold can be adjusted, we assume a probability above .5 indicates the pixel should be labeled as the brand color (1, or white in the mask). The network outputs a pixel mask that labels which pixels in the input frame correspond to the target brand color enabling color correction in those specified areas.

One frame can be corrected several times. For example, first for Clemson's orange, then purple, then the other team's branded-shade of blue. Each prediction takes in the color value and finds the correlated pixels to adjust. The primary advantage of this method over ColorNet 1.0 is that it is no longer necessary to create separate models for each color. Rather, multiple segmentations can be applied by the same model using different target colors. In addition, with the color automatically segmented, the producer will have the ability to fine-tune the adjustment using a custom software or hardware solution rather than the model making the correction without manual oversight.

Discussion

When analyzing the performance of ColorNet 1.0, the python-colormath module was used to calculate the difference in color of RGB values using a formula for Delta E CIE 2000. A continuous colorbar was generated where individual RGB values were passed through the model which allowed the resulting adjustments to be more easily evaluated and analyzed visually. The results can be seen in Figure 5 where the uncorrected and corrected colorbars are compared. A higher Delta E value represents a more significant difference in the original and the adjusted color value after running through the model. One unexpected outcome of this test was the spike in the magenta area of the colorbar. This can be attributed to the inherent nature of machine learning. If these colors are rarely present within the trading data, there can be irregularities in the resulting effect on these "unknown" colors.

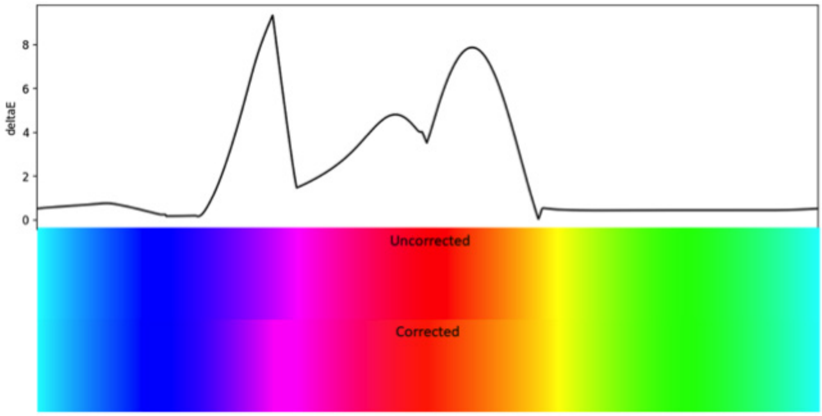

Figure 5: *Delta E test showing an unexpectedly high Delta E for the magenta range of the colorbar*

These results were a key consideration when transitioning into ColorNet 2.0 because the new goal was to be capable of targeting any specified brand color within a frame, rather than just Clemson orange. The expansion in flexibility of this model revolves around the target color as a new, additional input. As opposed to the previous model, ColorNet 2.0 will no longer act as a "black box" that automatically detects and makes adjustments on its own. It will instead have an element of user interaction in which producers are capable of using their expertise to determine the brand colors present in the game and fine-tune the adjustments that they see fit.

To accomplish this, a new interface will be designed that will allow the user, such as a producer of a broadcast, to manually select multiple brand colors they would like to segment out from the frame. Based on this input, ColorNet 2.0 will generate segmentation masks for each frame within the live footage. For each brand color segmented, the producer will also select, and manually adjust, a desired target color. Based on the segmentation masks and inputs selected through the interface, a color shift can be applied to produce a final image that addresses both team's brand colors. This interface and the ColorNet 2.0 model are still in development.

Conclusion

The success of ColorNet 1.0 in correcting Clemson orange during a live event piqued interest in expanding the machine learning algorithm for a wider array of broadcast applications. One of the most significant limitations of the initial model was its inability to correct for other brand colors without procuring data and re-training the model to recognize each additional color. An image segmentation model can circumvent this problem by segmenting specific regions of the image that need color adjustment. Then the technician could adjust the selected pixel regions without negatively impacting color on other areas of the screen. We have presented a new data-augmentation approach that generates synthetic data with substituted brand colors to expand the variability of our training dataset to all ACC colors. The team is now developing the neural architectures needed to accurately automate the segmentation of the targeted brand colors.

Acknowledgements

This project was supported by a Clemson University Research Foundation Technology Maturation Grant.

References

2014 U.S. Online Video Rankings. (2014, February 21). comScore, Inc. http://www.comscore.com/Insights/Press-Releases/2014/2/comScore-Releases-January-2014-US-Online-Video-Rankings

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

Jin, C., Yoon, M., & Lee, J. (2019). The influence of brand color identity on brand association and loyalty. *Journal of Product & Brand Management, 39, 88.* https://doi.org/10.1108/JPBM-09-2017-1587

Mayes, E., Lineberger, J.P., Mayer, M., Sanborne, A., Smith, H.D., & Walker, E.B. (2020). Automated Brand Color Accuracy for Real Time Video. *Proceedings of the 2020 NAB Broadcast Engineering and Information Technology Conference*. NAB Broadcast Engineering and Information Technology (BEIT) Conference, Las Vegas, NV. Retrieved May 15, 2020, from https://nabpilot.org/beitc-proceedings/2020/automated-brand-color-accuracy-for-real-time-video/

Minaee, S., Boykov, Y. Y., Porikli, F., Plaza, A. J., Kehtarnavaz, N., & Terzopoulos, D. (2021). Image segmentation using deep learning: A survey. IEEE *Transactions on Pattern Analysis and Machine Intelligence*.

Moser, M. (2003). United We Brand: How to Create a Cohesive Brand That's Seen, Heard, and Remembered. *Harvard Business School Press*. https://play.google.com/store/books/details?id=k17pnQEACAAJ

Walker, E.B., Smith, D.H., Lineberger, J.P., Mayer, M., Mayes, E., & Sanborne, A. (2020). ColorNet: A Neural Network-Based System for Consistent Display of Brand Colors for Video. SID Display Week 2020. *Display Week*.

World Heritage Encyclopedia. (n.d.). LIST OF MOST WATCHED TELEVISION BROADCASTS. *WorldHeritage.org*. Retrieved May 31, 2021, from http://cdn-cache.worldheritage.org/articles/List_of_most_watched_television_broadcasts