https://doi.org/10.2352/EI.2026.38.10.HVEI-213

Closed-Loop Color Refinement in Camera ISP via Post-ISP Color Feedback

Huanzhao Zeng, Google Inc. San Diego, CA, USA

Abstract

A closed-loop color feedback algorithm that leverages post-ISP statistics to improve camera color quality is presented. Unlike traditional approaches, which evaluate white balance and color early in the pipeline and tune individual modules in isolation, the proposed method assesses color near the end of the ISP pipeline, compares it against target perceptual colors, and feeds the resulting deviations back to upstream processing blocks. This enables dynamic adjustment of AWB and color-related parameters to achieve desired perceptual color outcomes. The framework addresses key limitations of conventional color processing, including (1) evaluating AWB in the raw domain where perceived color cannot be reliably assessed, (2) the inability of fixed color-tuning parameters to compensate for deviations introduced by other ISP blocks, and (3) the lack of coordinated color evaluation across modules.

We further demonstrate an application of this framework for skin-tone improvement. The system takes face regions, filters non-skin pixels, computes representative skin color statistics, compares them with target skin colors, and derives adjustment parameters that update color tunings for the current or subsequent frame. This example illustrates the flexibility and effectiveness of the proposed closed-loop approach for perceptually guided color enhancement or accurate color reproduction.

Introduction

In mobile-camera ISP pipelines, AWB (auto white balance), CCM (color correction matrix), and CLUT (color lookup table) are the primary modules responsible for white-balance estimation and color transformation (Fig. 1) [1]. AWB is typically computed from down-sampled sensor-raw images or raw-domain statistics, meaning the estimation occurs entirely in the sensor color space. Because this evaluation does not account for subsequent ISP color processing, it cannot reliably reflect the final perceptual color tint.

Color tuning for CCM and CLUT is similarly influenced by multiple upstream ISP operations, including black-level correction, lens-shading correction, tone mapping, and noise reduction. Inconsistent behavior among these blocks can lead to suboptimal tuning. Moreover, many of these processes are scene- or content-dependent, making it difficult for pre-tuned AWB, CCM, and CLUT parameters to capture all relevant variations. As a result, white-balance and color shifts may occur under conditions not fully represented during tuning. For example, face-based AWB computed from raw-domain face statistics may not account for color shifts introduced later in the pipeline. Likewise, CCM and CLUT parameters tuned for specific scenarios may fail to deliver the intended output when other ISP blocks alter the color characteristics. Even with advanced adaptive algorithms for AWB, CCM, and CLUT, the combined effect of all ISP stages may still deviate from the desired perceptual color due to incomplete modeling of cross-module color interactions.

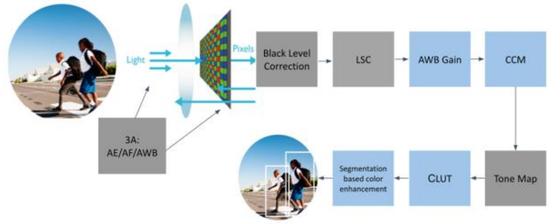

Figure 1. *Main color processing blocks in the camera ISP pipeline*

A further fundamental challenge arises from the metameric mismatch between CMOS sensors and the human visual system. Although both rely on a tri-chromatic representation of color, their spectral sensitivity functions differ significantly (Fig. 2). As a result, the RGB responses produced by integrating a light source's spectral power distribution are not identical for the two systems. A camera calibrated to match human perception under one illuminant may therefore behave differently under another. Consequently, even when CCMs and CLUTs are optimally characterized for a specific lighting condition, the resulting color transformation may not remain accurate across varying illuminations.

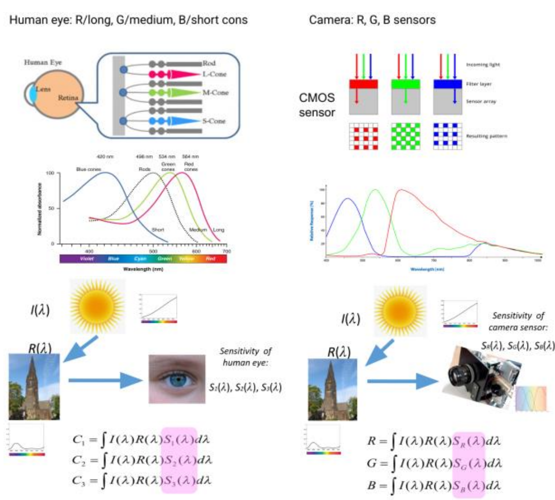

Figure 2. *Color signal response by human eye v.s. by camera sensor*

To address these limitations, we introduce a closed-loop feedback mechanism within the ISP pipeline that monitors the camera's output colors, analyzes deviations from desired targets, and feeds this information back to upstream processing blocks to adjust their parameters accordingly. A post-ISP statistics module can be integrated into the hardware to support real-time operations. The next section describes the proposed framework in detail, followed by an example application in Section 3. The final section provides concluding remarks.

Proposed Solution

Figure 3 illustrates the overall flow of color processing from the camera's raw image to the final output. For clarity, several ISP modules—such as black-level correction, lens-shading correction, demosaicing, tone mapping, and noise reduction—are omitted because they do not directly influence color transformation. In typical mobile-camera hardware pipelines, low-resolution images down-sampled from the sensor-raw domain, or hardware-generated statistical summaries, are used for AWB, AE, and AF analysis. We refer to these low-resolution images or preprocessed data collectively as *stats*.

Traditionally, AWB stats are computed in the sensor-dependent raw domain, which is not aligned with any standard perceptual color space. AWB estimation therefore relies heavily on these raw-domain statistics, supplemented by heuristic cues such as brightness value (BV), ambient-light sensor readings, and face information. Because this analysis occurs before ISP color processing, it cannot accurately predict the final perceptual color tint, often leading to unintended color biases in the output image. A more reliable strategy is to evaluate AWB and color later in the pipeline—after key color transformations—so that color can be assessed in a device-independent and perceptually meaningful color space.

A similar limitation exists for CCM and CLUT tuning. These parameters are typically optimized for predefined lighting conditions and scene categories. During image processing, the ISP classifies the scene using heuristic rules and applies to the corresponding tuning set. If the classification is inaccurate or the scene does not match any predefined category, the resulting color output may deviate from expectations. These challenges highlight the need for a more adaptive and perceptually grounded color-processing strategy.

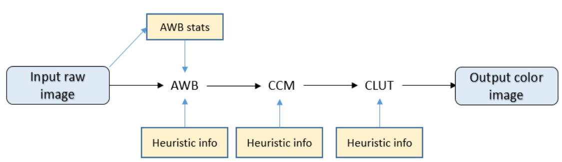

Figure 3. *Major color adjustment blocks in a typical mobile phone camera pipeline*

A key observation is that without evaluating and feeding color information from the end of the ISP pipeline back to earlier stages, many color errors and shifts cannot be reliably detected or corrected. Because the colors produced at the end of the pipeline are what users ultimately perceive, evaluating color at or near this stage provides a far more accurate basis for white-balance decisions and color correction. By analyzing output colors after the major ISP color transformations, the system can capture errors introduced by ISP processing as well as inaccuracies in AWB estimation. These deviations can then be fed back to upstream modules to adjust AWB and color-processing parameters for subsequent frames, thereby reducing AWB errors and mitigating color shifts.

To implement this concept, we introduce a post-ISP statistics block that generates color metrics at or near the end of the ISP color pipeline and provides additional analysis to AWB and color modules (Fig. 4). These statistics enable evaluation of output color—such as skin tone or white balance—in a device-independent perceptual color space. The resulting feedback is used to correct color issues by adjusting AWB gains or, when necessary, updating CCM or CLUT parameters. This establishes a closed-loop color-feedback mechanism that continuously refines color accuracy.

The same closed-loop framework can be extended beyond color to analyze brightness and contrast, enabling the ISP to refine global or local luminance characteristics, including face brightness and contrast. This flexibility makes the approach broadly applicable to a range of perceptual image-quality improvements.

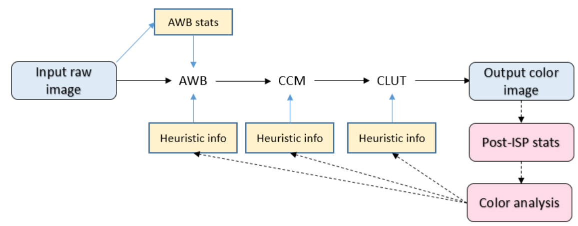

Figure 4. *A closed-loop color feedback process in the camera pipeline for AWB and color adjustments*

A Use Case Example: Face Color Improvement

The closed-loop color framework can be applied to improve image quality across multiple dimensions, including white balance, sky color, skin tone, and overall brightness and contrast. In this section, we focus specifically on its application to face-color adjustment. Although the same mechanism can be used to refine AWB by updating white-balance gains [2], our discussion here emphasizes face-color improvement through CCM adjustment.

The post-ISP statistics block computes face-color metrics from the ISP output and evaluates whether the measured face color falls within a desired target range. When deviations are detected, the resulting information is fed back to upstream ISP modules to adjust color-processing parameters accordingly.

The post-ISP statistics block can be configured to accept one or more regions of interest (ROIs) and analyze the color characteristics within each. When an ROI corresponds to a detected face, the block extracts color statistics specifically for that region. As illustrated in Fig. 5, the initial face ROI (red rectangle) produced by the face-detection algorithm often includes background pixels. To improve accuracy, this ROI is refined to a smaller, tighter region (green rectangle) before being passed to the post-ISP statistics block. Non-skin pixels—such as very dark or bright pixels, or pixels corresponding to sunglasses or face masks—are removed using a screening algorithm. The remaining skin pixels are then used to compute reliable face-color statistics.

A preferred skin color set is defined for face color adjustment. We started with the Monk Skin Tone (MST) scale (see Fig. 6 and Table 1) [3-4]. The MST was originally defined in RGB. They were converted to CIELAB for evaluation [5]. In this MST data set, lightness (L*) of the three lightest tones does not monotonically decrease. Chroma (C*) and hue angle (H* in degrees) are not smooth (see Fig. 7). Furthermore, light skin tones are very greenish. After investigating and testing the skin color adjustment, it was determined that they could not be used without significant modifications.

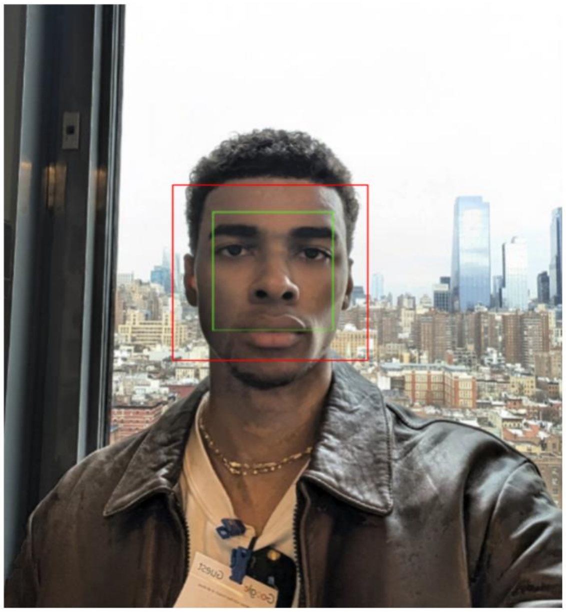

Figure 5. Face ROI by face detection (red rectangle), and modified face ROI passed to the post-ISP for face color evaluation

Figure 6. Monk Skin Tone scale

Table 1 RGB and CIE LCH values of Monk Skin Tones (MST)

MST	Hex (rgb)	L*	C*	H*
1	#f6ede4	94.42	5.62	74.49
2	#f3e7db	92.27	7.56	74.19
3	#f7ead0	93.09	14.2	89.13
4	#eadaba	87.57	17.75	88.52
5	#d7bd96	77.9	23.39	81.47
6	#a07e56	55.14	27.85	73.77
7	#825c43	42.47	23.94	59.02
8	#604134	30.68	17.72	48.81
9	#3a312a	21.07	6.54	65.71
10	#292420	14.61	3.82	67.19

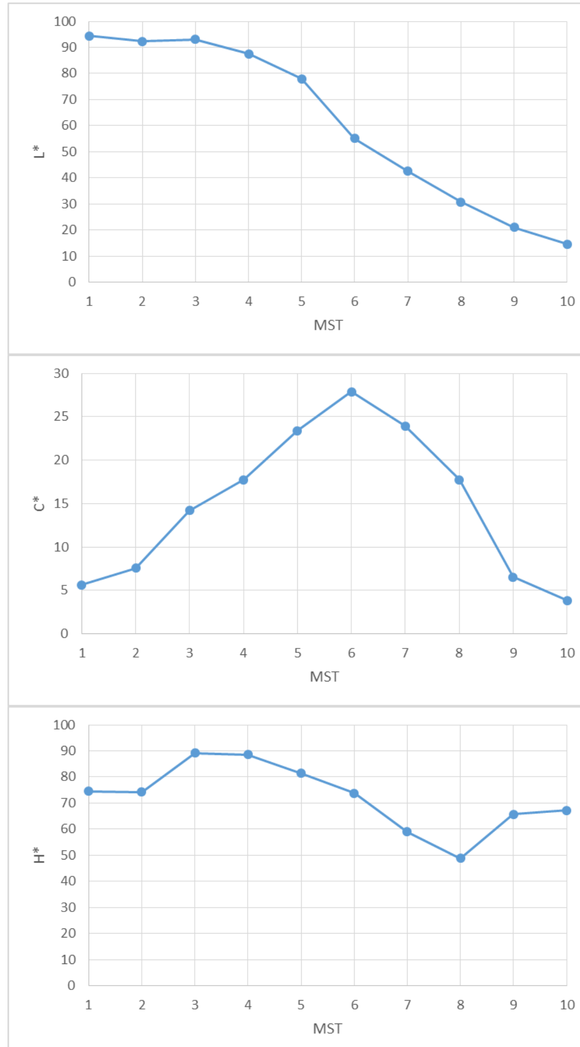

*Figure 7. L*C*H* of Mon Skin Tone Scales*

Our first modification targeted hue. We found that a hue angle of approximately 46° in CIELAB under D65 illumination produced visually reasonable results for initial experiments, and we adopted this value as the target hue. Through further testing across diverse photo data sets, we also refined chroma values. Figure 8 compares the original MST L–C distribution with our modified data set. The revised set smooths chroma across all tone levels and extends the range by adding black and white reference points. Limitations of using a single modified set for all skin-tone types will be discussed later.

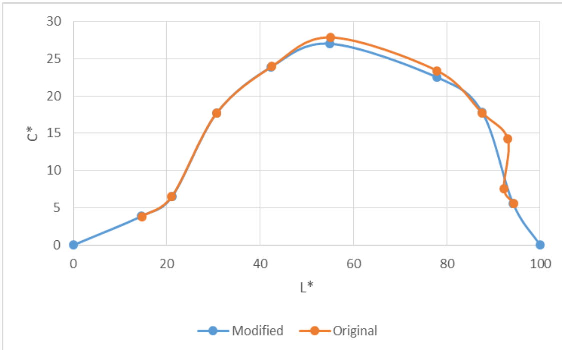

Figure 8. *The original MST chroma and our modifications*

Figure 9. *Skin color saturation is improved in this backlit image.*

Once the face color is computed from the refined ROI, a corresponding target color is computed from the preferred skin-color data set. A chroma-scaling factor and a hue-shift value are then derived to move the measured face color toward the target. If the measured color falls within predefined tolerance thresholds, no adjustment is applied; otherwise, the color is shifted toward the target region. The resulting adjustment parameters are used to compute the final corrected skin color.

Depending on implementation or hardware capabilities, these adjustments may be applied through CCM or CLUT modification. Because CCM is a global transform, applying a skin-specific correction through a single CCM may degrade non-skin colors. If the hardware supports segmented CCMs, a dedicated skin-region CCM can be modified for localized correction. When using a CLUT-based approach, a skin-color region within the CLUT is defined, and adjustments are applied only within that color region, with smooth transitions to surrounding non-skin colors.

Using a single target color set for all skin-tone types may not yield optimal results for every individual. To achieve robust performance across diverse skin tones, target colors should be defined per skin-tone category. A final target skin color is then determined based on both the measured skin lightness and the identified skin-tone type.

Experimental Verification

In the following samples, the left is the original, and the right is the result after applying a skin CCM to modify skin colors. The CCM is applied globally to the whole image in this demonstration, so non-skin regions are also modified. Because the segmentation CCM feature can be enabled to apply a skin CCM to the human body region only in the real system, we can ignore non face/body regions in the evaluation.

We tested the algorithm on various image data sets. Some images were not modified because the algorithm detected that the original skin colors were already within the target color preference. For those images that were modified, we selected a few for discussion.

In backlit scenes, after the brightness and contrast enhancement, skin color saturation may not be ideal. In this backlit image in Fig. 9, the lower than ideal face color saturation is increased after applying the face color adjustment method.

In the photo in Fig. 10, the algorithm detects that the original skin color saturation a lower than ideal and it is increased after the processing.

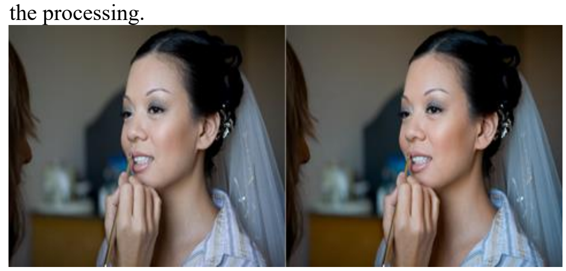

Figure 10. *Skin color saturation increases.*

In Fig. 11, the original skin color is greenish and is fixed after the closed-loop color refinement.

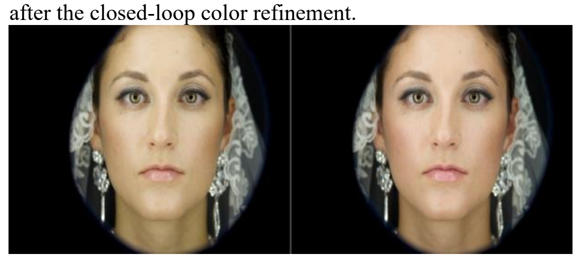

Figure 11. *The greenish skin tone is fixed.*

In Fig. 12, the original skin color is pale and greenish. Both issues are fixed after the closed-loop color adjustment.

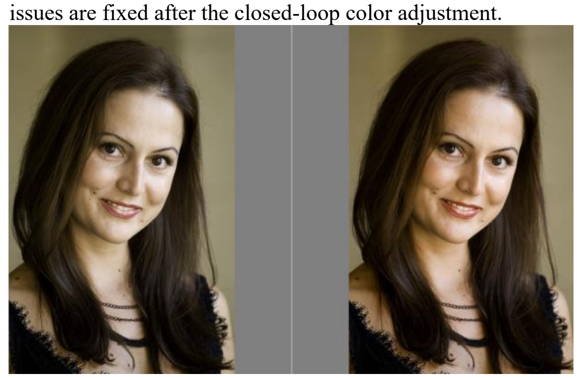

Figure 12. *The pale and greenish skin tone is fixed.*

In Fig. 13, the original skin color looks close to ideal. After the processing, the skin color hue is slightly reduced to fix a very slight greenish tint, and the color saturation is slightly increased. This result is more preferred by the author's preference.

Figure 13. *The skin color is improved by slightly shifting skin hue to the reddish direction and very slightly increasing color saturation.*

In Fig. 14, the skin color saturation is increased slightly in our default parameter setup. Depending on the color preference setting, we may adjust color preference parameters to not increase the color saturation in this case.

Figure 14. *The skin hue is not adjusted, but the saturation is increased slightly.*

In Fig. 15, the purple skin tone is reduced. Because the hue adjustment to maximize the skin color optimization is larger than the maximum hue adjustment limit allowed in our setting parameters, the purple tone is not completely removed. Because of that, a little bit of the original purple tone is kept.

Figure 15. *The purple skin tone is reduced.*

In Figure 16, the orange skin tone is fixed, but the skin color might still not be ideal. In Figure 17, the adjusted face color is less red and less saturated. However, it becomes more orange, and this might not be preferred. The skin tone type was not applied in this experiment. From the results in these two images, we observed the limitation of applying a single target reference for all skin tone types. To further improve these skin tones, face detection should include skin tone type identification. Preferred skin colors of different skin tones and different culture backgrounds are different [5, 6]. By setting different target preferences for different tone types and enabling skin tone type detection, output skin colors in these two images should be improved.

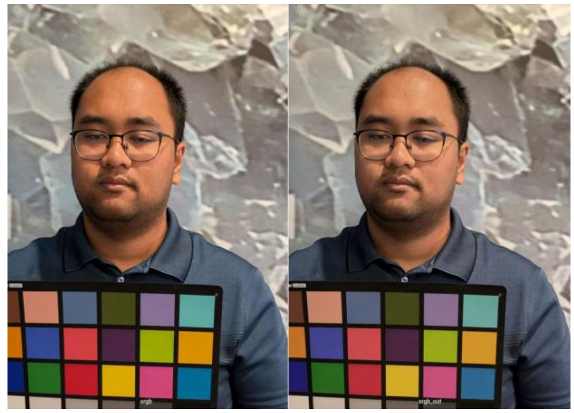

Figure 16. *The orange skin tone on the face is fixed.*

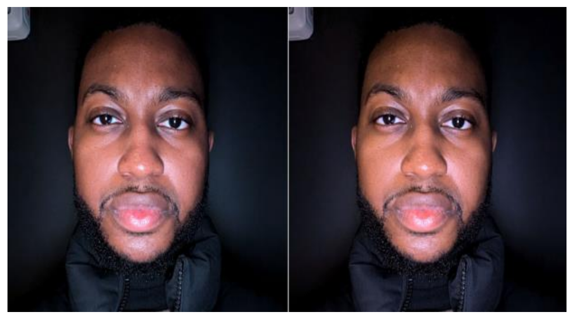

Figure 17. *The skin hue and saturation are adjusted, but skin color is still not ideal.*

Skin color preference is highly dependent on cultural backgrounds. The target color preferences may be set slightly differently for different cultural backgrounds. Personal color preference is one more factor. The target preferences set for this test will not meet everyone's preferences. In some of the above samples, some people may argue that they prefer originals. Personal color preferences may be applied to adjust target skin color preferences for the processing. To implement the personal color preferences on a camera, we may let users set their target color preferences and then pass personal color preference parameters into the algorithm.

Conclusions

We presented a closed-loop color-feedback mechanism integrated into the ISP pipeline to improve camera color quality. By evaluating color at the perceptual end of the pipeline and feeding this information back to upstream modules, the proposed framework addresses key limitations of traditional ISP designs, where AWB, CCM, and CLUT are tuned independently and without end-to-end color awareness. This closed-loop strategy enables coordinated, perceptually guided adjustments that reduce white-balance errors, mitigate color shifts, and improve overall color accuracy and color preference. Experimental results demonstrate the effectiveness of the approach. A dedicated post-ISP hardware statistics module can further accelerate processing and support real-time deployment, particularly for video applications.

References

[1] Phil Green, Fundamentals and Applications of Colour Engineering (2024 John Wiley & Sons, Ltd).

[2] Liqing Wang, Yuechen Zhu, Xiaoxuan Liu, Ming Ronnier Luo, “Improve image white balance by facial skin color”, https://doi.org/10.2352/CIC.2023.31.1.9, 2023

[3] Monk Skin Tone Scale, https://en.wikipedia.org/wiki/Monk_Skin_Tone_Scale.

[4] Start using the Monk Skin Tone Scale, https://skintone.google/get-started.

[5] Huanzhao Zeng and Ronnier Luo, Colour and Tolerance of Preferred Skin Colours on Digital Photographic Images (Color Res Appl 2013; 38:30–45).

[6] Zeng, Huanzhao (2011) Preferred skin colour reproduction. PhD thesis, University of Leeds.

Author Biography

Huan Zeng is a color imaging scientist at Google, specializing in the development of color imaging architectures and algorithms for mobile devices. He previously held similar roles in Meta, Qualcomm, and Hewlett-Packard, contributing to a broad range of color imaging technologies and color management systems. He received a PhD in Coulor Science from the University of Leeds and an MS in Imaging Science from Rochester Institute of Technology.

JOIN US AT THE NEXT EI!

Imaging across applications . . . Where industry and academia meet!

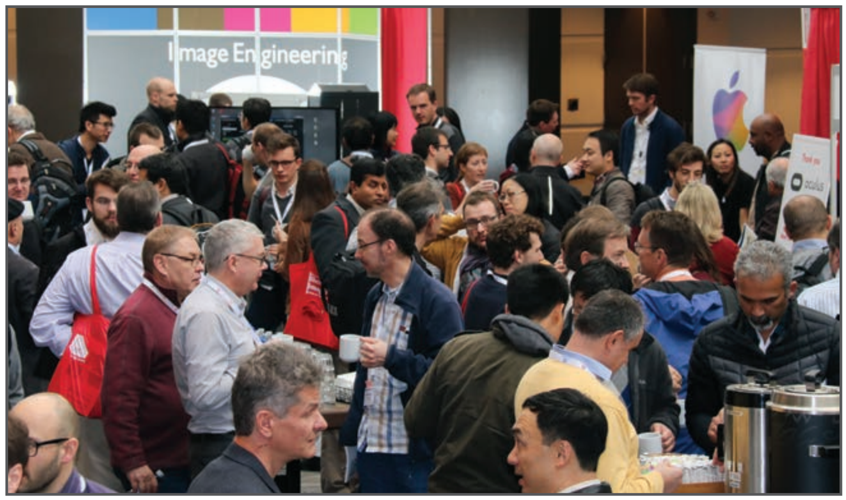

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org