Luminance-Preserving and Temporally Stable Daltonization

Supplemental material

P. Ebelin[1], C. Crassin[1], G. Denes[2,3], M. Oskarsson[4], K. Åström[4], and T. Akenine-Möller[1]

[1]NVIDIA [2]The Perse School Cambridge [3]University of Cambridge [4]Centre for Mathematical Sciences, Lund University

1. Color Vision Deficiency Simulation

For the sake of reproducibility, we outline exactly how CVD simulation is done in our system. There are three commonly used methods for color vision deficiency simulation, namely the seminal works by Brettel, Viénot, and Mollon [BVM97, VBM99] and the approach suggested by Machado et al. [MOF09]. The latter work focuses on less severe cases of CVD and uses the former as a reference for dichromacy. Our work uses Viénot et al.'s simulation method [VBM99]. Viénot et al. simplify the method suggested in their earlier work and have it assume the same chromaticities as are used in the sRGB standard [Poy03]. The simplification approximates Brettel et al.'s projection onto two halfplanes by a projection on a single plane. This approximation, together with their linear approximation of the Judd-Vos modification in XYZ, implies that we can express Viénot et al.'s method as a single matrix multiplication given linear sRGB colors. In particular, we can construct the simulation matrix, $\mathbf{M}_S$, by

$$\mathbf{M}_S = \mathbf{M}_{\text{LMS}\rightarrow\text{RGB}}\mathbf{M}_{\text{Projection}}\mathbf{M}_{\text{RGB}\rightarrow\text{LMS}}, \tag{1}$$

where

$$\begin{aligned}
\mathbf{M}_{\text{RGB}\rightarrow\text{LMS}} &= \begin{pmatrix} 17.8824 & 43.5161 & 4.11935 \\ 3.45565 & 27.1554 & 3.86714 \\ 0.0299566 & 0.184309 & 1.46709 \end{pmatrix}, \\
\mathbf{M}_{\text{LMS}\rightarrow\text{RGB}} &= (\mathbf{M}_{\text{RGB}\rightarrow\text{LMS}})^{-1}, \text{ and} \\
\mathbf{M}_{\text{Projection}} &= \begin{pmatrix} 0 & 2.02344 & -2.52581 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \text{ or} \\
\mathbf{M}_{\text{Projection}} &= \begin{pmatrix} 1 & 0 & 0 \\ 0.494207 & 0 & 1.24827 \\ 0 & 0 & 1 \end{pmatrix},
\end{aligned} \tag{2}$$

respectively, depending on whether we simulate protanopia or deuteranopia.

Thus, to simulate dichromacy for a given image, $\mathbf{I}'$, stored in sRGB, we first linearize the pixel values, yielding $\mathbf{I}$ after which we multiply each pixel's RGB colors by $\mathbf{M}_S$ (with $\mathbf{M}_{\text{Projection}}$ determined by the assumed dichromacy type). This yields $\mathbf{I}_S$, which we apply the sRGB transform to and obtain the dichromacy simulation result, $\mathbf{I}'_S$.

2. CBFM simulation

In this section, we discuss our simulation method for the Computer-Based Farnsworth-Munsell (CBFM) 100-hue test [Far43, GPD*14]. We first generate the undistorted color of each cap based on Yang et al. [YRWL08]. To compute the *visible color* of each cap, we explore the following stimuli options:

- Original: colors are left unchanged
- Original+CVD: each color is transformed using CVD simulation
- Our+CVD: our proposed daltonization is applied before colors are transformed using CVD simulation
- Farup+CVD: Farup's daltonization [Far20] is applied before colors are transformed using CVD simulation.

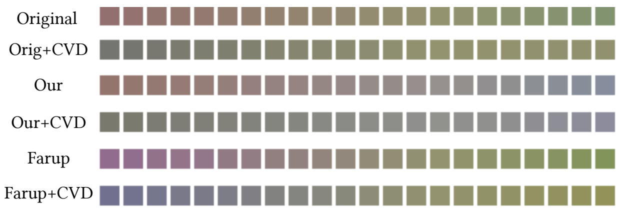

Figure 1: *First box of colours in the CBFM simulation, as ordered by our simulation. Ordering is substantially harder with CVD, as several color caps can map to indistinguishable colors.*

The color caps are split into four 'boxes', where in each box the first and last color cap is fixed. We compute ΔE_{00} distances between each pair of caps within each box. As a ΔE_{00} of '1' is often regarded as the detection threshold, we can consider ΔE_{00} differences the output of a pairwise comparison experiment, and can compute the probability of detecting the difference between two color caps using a cumulative normal distribution function ($\sigma = 1.4826$) [POMZ*19]. This can then in turn be transformed into a probability of *indistinguishability* as

$$P(C_a = C_b) = 2(1 - \text{cndf}_\sigma(\Delta E_{00}(C_a, C_b)), \tag{3}$$

where cndf is the cumulative normal distribution function with $\sigma = 1.4826$.

To simulate a possible strategy of a human participant, we split our simulation into two stages:

1. initial crude ordering,
2. sequence clean-up.

During the initial crude ordering, we start with a *sorted list*, which contains only the first fixed cap. We then repeatedly look for a color cap that is close to the last cap of the sorted list and append it to the sorted list. As this phase is only a crude ordering, we pick randomly from color caps within two units of ΔE_{00}.

Proceedings published by Eurographics - The European Association for Computer Graphics.

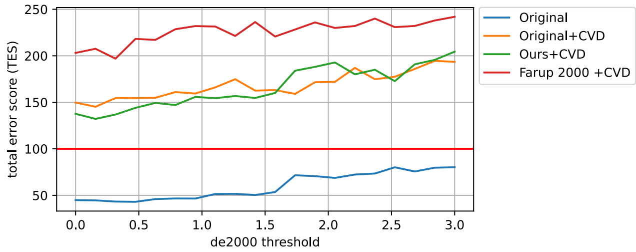

Figure 2: *The impact of the ΔE_{00} threshold in the initial crude sorting of color caps. Higher ΔE_{00} thresholds result in higher TES scores (with 50 clean-up steps) averaged over 30 simulations. However, the relative performance of each technique is not affected.*

During the second phase, we simulate the human participant cleaning up a semi-sorted box. Specifically, we randomly pick a color cap C_x and attempt to move it to a location where it better fits in the sequence of colors. More formally, we can find the value of a sequence of color caps C_0 to C_n as

$$P(C_0...C_n) = \prod_{i=1}^{n} P(C_i = C_{i-1}), \tag{4}$$

where $P(C_i = C_{i-1})$ indicates that probability that color C_i is indistinguishable from color C_{i-1}. A human participant will attempt to maximize the value of this expression, but as this is achieved via a number of visual comparisons, the process is not expected to be error free. To account for this, we compute the value of each possible location (i.e., we compute the value of each sequence assuming that the color cap has been placed at a given location), then select from this randomly with the probability of picking a location proportionate to the relative value of a resulting sequence. We then repeat this clean-up step N times.

We repeat the simulation multiple times to compute an average TES score for each stimulus. Note that there are two parameters in our model: (1) the ΔE_{00} threshold at which we cut off in the initial stage, and (2) N, which is the number of clean-up iterations we perform. Figure 2 shows that the choice of the initial ΔE_{00} threshold has some impact on the final TES scores, but it does not seem to affect the relative performance of the daltonization techniques. We found that ΔE_{00}=2 and N=50 with 20 repetitions provides stable and plausible results.

3. Examples of Temporal Instability

Most previous approaches do not provide temporal stability. In Figure 3, we see an example where the daltonization algorithm by Machado and Oliveira [MO10] exhibits global color variations during camera viewpoint changes. Huang et al.'s method [HZC*22] also generates inconsistent results with small changes in camera viewpoint, as illustrated in Figure 4. Note that, because Huang et al.'s daltonization [HZC*22] assumes Machado et al.'s CVD simulation algorithm [MOF09], said algorithm is used for the results in Figure 4. The temporal stability of our approach is shown in the accompanying video and compared with that of Machado and Oliveira's [MO10] and Farup's [Far20] algorithms.

Figure 3: *Color variation observed with Machado and Oliveira's algorithm [MO10] during a camera movement after CVD simulation. Note that our method (bottom) provides temporal consistency. Simulated CVD: protanopia.*

Figure 4: *The method by Huang et al. [HZC*22] can generate substantially different results depending on image content. This gives rise to occasional flicker during animation.*

4. Additional Image Results

In some cases, our daltonization algorithm allows making color details that would become invisible to a person with dichromacy visible again. An example of such a result for a real-time rendered scene is shown in Figure 5.

Figure 6 shows a close-up rendering of a tiger. We note that both our algorithm and Farup's improve the contrast in the image. Farup makes the orange in the tiger more purple, which becomes blue after CVD simulation, while our algorithmn makes the greens of the trees more blue. Here, our version provides a more natural experience overall. The same can be said about the results for the photograph of the lady (Figure 7), where our choice of colors again results in a more natural experience, both for the person with dichromacy and the one without it, compared to Farup's output. Note, however, that Farup better preserves the colors of the trees.

Figures 8 and 9 show images of nebulae generated by the NASA/ESA Hubble space telescope. In the first figure, our method provides a clearer separation between the red/brownish colors and the blue/turquoise colors, compared to Farup's algorithm. In Figure 9, we can see that the CVD simulation of the original image is mostly yellow and has lost a lot of contrast. Both our method and Farup's method increase the contrast mainly by adding blue colors. We leave it up to the reader to determine which one is preferable, and for future work, it would be interesting to have a user study where people with CVD could evaluate these images.

In Figure 10, our method and Farup's have been applied to a

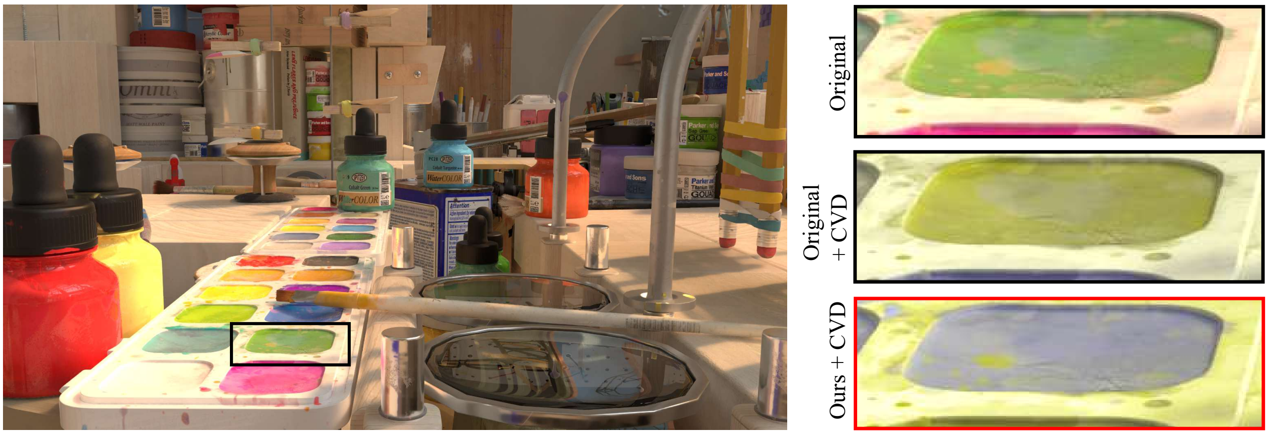

Figure 5: *Without special color treatment, some color details like the painting spots visible in the insets on the right can become invisible to a person with dichromacy (protanopia; center inset). Our daltonization algorithm makes these color details visible again (bottom inset).*

Table 1: *Average absolute luminance difference compared to* **R**.

CVD type	No daltonization	Farup [Far20]	Our
Protanopia	0.035	0.039	**0.001**
Deuteranopia	0.019	0.016	**0.002**

colorful image of pencils. This is an example of a difficult image, where it seems that neither our method nor Farup's provide improvements over using no daltonization (middle row). As stated in the introduction in the main paper, there are only about 0.4% of the colors left to use after CVD simulation for dichromats, which makes it challenging to provide a satisfactory daltonized result in all cases. While our method and Farup's often provide improvements, this is a case where they do not. We see examples of this in Figure 12 as well. There, we show a collection of Ishihara plates [Ish18], which are a common part of color vision deficiency diagnostics. Our recoloring makes it possible for the dichromat to see the information in many of the plates, but it also makes it harder in a few of the cases. The same is true with Farup's algorithm.

5. Preserving Luminance

For this comparison, we use a 4k×4k image, **R**, containing every 24-bit RGB color. In Table 1, we show the average absolute luminance difference between **R** and dichromacy-simulated versions of it, before and after daltonization. The results show that our method preserves luminance substantially better compared to Farup's algorithm [Far20], and we retain the luminance that the dichromat loses when observing the original image.

6. Back-projection and Lower Severity CVD

Even though our approach is designed to improve color separation and contrast preservation in case of dichromacy, our back-projection minimizes the color change between our output and the original image. While our approach may be sub-optimal for improving images for people with lower severity CVD, for whom some of the chrominance is retained, they too may benefit from our transform to some degree because of our back-projection. This visual proximity to the original image can be observed in the results presented in Figures 6, 7, 8, 9, and 11.

7. Performance

We implemented our daltonization algorithm in Falcor [KCK*21]. Our fastest implementation of the algorithm simply precomputes the entire transform into a three-dimensional texture and performs a single lookup in that texture per pixel. We applied our daltonization after path tracing, denoising, tone mapping, and accumulation and measured the time of our render pass on an NVIDIA RTX 4090 GPU with driver version 527.56 at resolution 3840 × 2160. The implementation with a three-dimensional lookup table runs in approximately 0.2 ms, while our slower variant, which uses two-dimensional textures, more compute, and less memory, requires approximately 0.4 ms. One may trade-off accuracy for speed through the number of texels in the two- and three-dimensional textures.

Acknowledgments We are grateful to the people sharing their scenes: ZERODAY ©beeple [Win19] and BISTRO ©Amazon Lumberyard [Lum17]. We thank Robert Toth, Pierre Morreau, Jacob Munkberg, Jon Hasselgren, Ruth Rosenholtz, Andrew Russell, Rachel Brown, Anjul Patney, Ekta Prashnani, and Amrita Mazumdar for helpful discussions and feedback. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

[BVM97] BRETTEL H., VIÉNOT F., MOLLON J. D.: Computerized Simulation of Color Appearance for Dichromats. *Journal of the Optical Society of America A 14*, 10 (Oct 1997), 2647–2655. 1

[Far43] FARNSWORTH D.: The Farnsworth-Munsell 100-Hue and Dichotomous Tests for Color Vision. *Journal of the Optical Society of America 33*, 10 (Oct 1943), 568–578. 1

[Far20] FARUP I.: Individualised Halo-Free Gradient-Domain Colour Image Daltonisation. *Journal of Imaging 6*, 11 (2020). 1, 2, 3, 4, 5, 6

[GPD*14] GHOSE S., PARMAR T., DADA T., VANATHI M., SHARMA S.: A New Computer-Based Farnsworth Munsell 100-Hue Test for Evaluation of Color Vision. *International Ophthalmology 34*, 4 (2014), 747–751. 1

Figure 6: *Image of a rendered tiger. Note the bluish colors of the tiger in the image by Farup and the increased contrast in the background in our image. Simulated CVD: deuteranopia.*

[HZC*22] Huang W., Zhu Z., Chen L., Go K., Chen X., Mao X.: Image Recoloring for Red-Green Dichromats with Compensation Range-based Naturalness Preservation and Refined Dichromacy Gamut. *The Visual Computer 38* (2022), 3405–3418. 2

[Ish18] Ishihara S.: Tests for Color Blindness. *American Journal of Ophthalmology 1*, 5 (1918), 376. 3

[KCK*21] Kallweit S., Clarberg P., Kolb C., Yao K.-H., Foley T., Wu L., Chen L., Akenine-Möller T., Wyman C., Crassin C., Benty N.: The Falcor Rendering Framework, August 2021. URL: https://github.com/NVIDIAGameWorks/Falcor. 3

[Lum17] Lumberyard A.: Amazon Lumberyard Bistro, Open Research Content Archive (ORCA), July 2017. URL: http://developer.nvidia.com/orca/amazon-lumberyard-bistro. 3

[MO10] Machado G. M., Oliveira M. M.: Real-Time Temporal-Coherent Color Contrast Enhancement for Dichromats. *Computer Graphics Forum 29*, 3 (2010), 933–942. 2

[MOF09] Machado G. M., Oliveira M. M., Fernandes L. A. F.: A Physiologically-based Model for Simulation of Color Vision Deficiency. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009), 1291–1298. 1, 2

[POMZ*19] Perez-Ortiz M., Mikhailiuk A., Zerman E., Hulusic V., Valenzise G., Mantiuk R. K.: From pairwise comparisons and rating to a unified quality scale. *IEEE Transactions on Image Processing 29* (2019), 1139–1151. 1

[Poy03] Poynton C.: *Digital Video and HDTV – Algorithms and Interfaces*. Morgan Kaufmann, 2003. 1

[VBM99] Viénot F., Brettel H., Mollon J. D.: Digital Video Colourmaps for Checking the Legibility of Displays by Dichromats. *Color Research & Application 24*, 4 (1999), 243–252. 1

[WGY*18] Wolski K., Giunchi D., Ye N., Didyk P., Myszkowski K., Mantiuk R., Seidel H.-P., Steed A., Mantiuk R. K.: Dataset and Metrics for Predicting Local Visible Differences. *ACM Transactions on Graphics 37*, 5 (2018). 4

[Win19] Winkelmann M.: Zero-Day, Open Research Content Archive (ORCA), November 2019. URL: https://developer.nvidia.com/orca/beeple-zero-day. 3

[YRWL08] Yang S., Ro Y., Wong E., Lee J.-H.: Quantification and standardized description of color vision deficiency caused by anomalous trichromats—part i: simulation and measurement. *EURASIP Journal on Image and Video Processing 2008* (2008), 1–9. 1

Figure 7: *Results for a photograph of a lady leaning against a stone. Without daltonization, a lot of contrast is lost to the dichromat. When daltonization is added using either daltonization algorithm, contrast is retained, but the images look less natural as the skin tone or leaf colors changes significantly. Simulated CVD: deuteranopia. The image was retrieved from the LocVis database [WGY*18].*

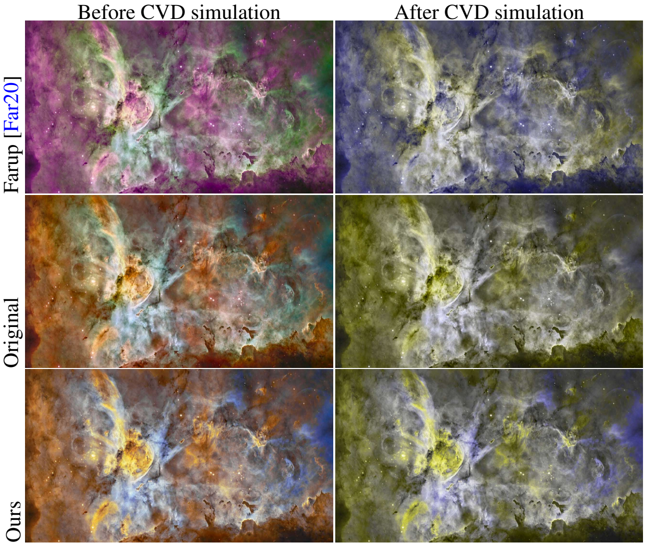

Figure 8: *Results for a capture of the Carina Nebula. Comparing our output to Farup's, we observe that ours look more similar to the original (middle-left), even to the dichromat, as can be seen by comparing the images to the right. Simulated CVD: protanopia. Original image credit: ESA/Hubble, "star birth in the exteme."*

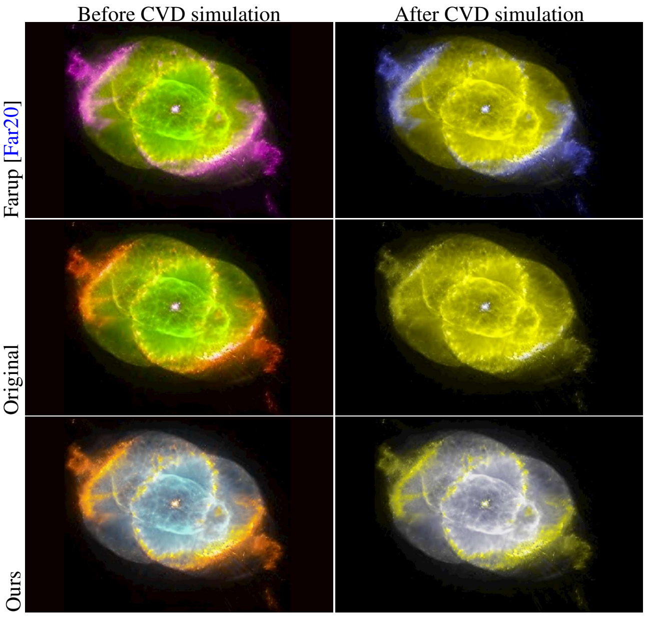

Figure 9: *Results for a capture of the Cat's Eye Nebula. When the image contains only reds, yellows, and greens, a dichromat will have significant issues distinguishing between the colors, as can be seen in the middle-right image. Both daltonization algorithms improve on this, but ours also preserves luminance. Simulated CVD: protanopia. Original image credit: ESA/Hubble.*

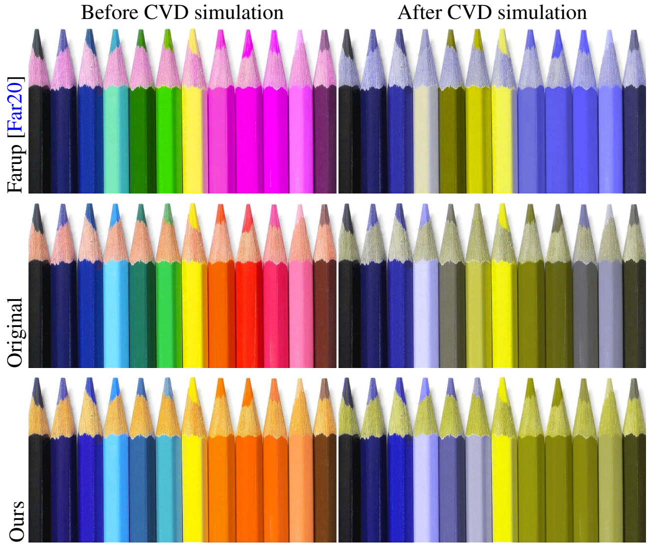

Figure 10: *Results for a photograph of pencils. This photograph contains several vivid and distinct colors and is a case where daltonization does not help. Simulated CVD: protanopia. Original image credit: Brokenarts on freeimages.com*

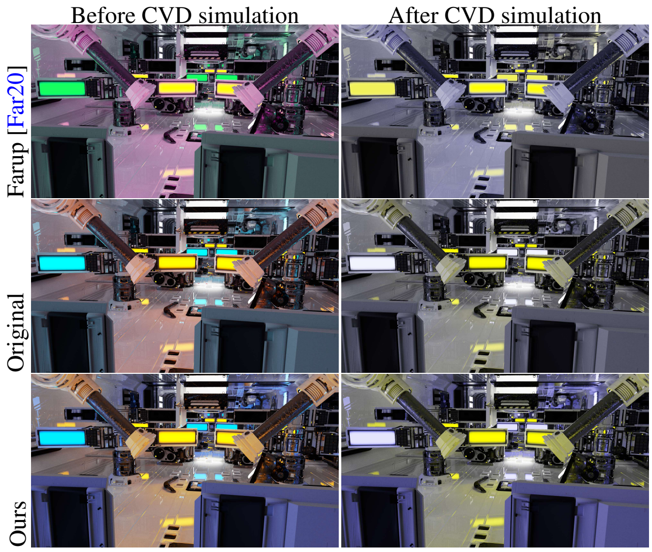

Figure 11: *Results of a set of recoloring algorithms for a rendered image of the ZeroDay scene. Note that Farup's method generates an image after CVD simulation (top-right) where all lamps, which were originally either yellow or turquoise, are all yellow. Our method generates an image with clearly different colors for these lamps. This is also a case where our image before CVD simulation is clearly more similar to the original than the image generated by Farup's method. Simulated CVD: protanopia.*

Proceedings published by Eurographics - The European Association for Computer Graphics.

Figure 12: *Results for a collection of Ishihara plates. The first two rows show benefits from daltonization, i.e., we can make out more numbers after simulation when daltonization is active. The middle plate in the middle row is an example where our algorithm is the only one successfully showing a number. We show worse results in the last plate on the next row, where no daltonization or Farup makes it easier to distinguish the number. On the last row, Farup's method is the only one that clearly shows a 5 in the second plate, however, Farup also seems to indicate that the middle plate (bottom row) says 45, which the original image does not easily reveal.*
Simulated CVD: deuteranopia. Original image credit: Kaludov on VectorStock.com.