A Simplified iCAM-like Image Appearance Pipeline Implementation, Results and Analysis

Shankhya Debnath

August 15, 2025

Abstract

This document describes the design, mathematical formulation, implementation, and experimental evaluation of a simplified iCAM-like image appearance pipeline. The pipeline implements the core appearance-preserving ideas from iCAM family models: (1) local adaptation estimation, (2) local tone mapping, and (3) chromatic preservation during luminance remapping. Results and quantitative summaries for example images are presented, followed by discussion on limitations and next steps for a closer reproduction of iCAM.

1 Introduction

The iCAM model family (image Color Appearance Model) was developed to predict how images appear under different viewing conditions by combining color appearance modeling with multi-scale spatial operations. Full iCAM implementations include opponent-space nonlinearities, band-pass spatial decomposition, surround induction, and more. In many practical applications a simplified pipeline that captures the key perceptual effects (local adaptation and appearance-preserving tone mapping) is very useful.

This report documents a simplified but practical implementation that:

- linearizes sRGB to linear RGB and converts to CIEXYZ,
- estimates a local luminance (adapting) map using an edge-preserving filter (bilateral) or Gaussian blur,
- performs a local tone mapping (heuristic inspired by Reinhard/local operators),
- reconstructs XYZ with chromatic preservation (scale X and Z by the luminance ratio),
- converts back to sRGB for display and comparison.

2 Background and Theory

This section provides the mathematical details of each processing step used in the pipeline.

2.1 sRGB ↔ linear RGB

sRGB images are encoded with a gamma-like nonlinearity. The linearization from sRGB value $c_{srgb} \in [0,1]$ to linear RGB c_{lin} is:

$$c_{lin} = \begin{cases} \frac{c_{srgb}}{12.92}, & c_{srgb} \le 0.04045 \\ \left(\frac{c_{srgb}+0.055}{1.055}\right)^{2.4}, & c_{srgb} > 0.04045 \end{cases}$$

The inverse (linear-to-sRGB) is:

$$c_{srgb} = \begin{cases} 12.92\, c_{lin}, & c_{lin} \le 0.0031308 \\ 1.055\, c_{lin}^{1/2.4} - 0.055, & c_{lin} > 0.0031308 \end{cases}$$

2.2 Linear RGB ↔ CIEXYZ

The linear sRGB to CIEXYZ conversion uses the standard D65-based matrix:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = M_{sRGB\to XYZ} \begin{bmatrix} R_{lin} \\ G_{lin} \\ B_{lin} \end{bmatrix}, \qquad M_{sRGB\to XYZ} = \begin{bmatrix} 0.4124564 & 0.3575761 & 0.1804375 \\ 0.2126729 & 0.7151522 & 0.0721750 \\ 0.0193339 & 0.1191920 & 0.9503041 \end{bmatrix}.$$

The inverse $M_{XYZ\to sRGB} = M_{sRGB\to XYZ}^{-1}$ converts XYZ back to linear sRGB.

2.3 Bradford Chromatic Adaptation (optional)

To simulate illuminant change (not used by default), the Bradford CAT is performed by mapping XYZ to an approximate cone response domain (LMS) via M_{br}, scaling LMS to match destination whitepoint, and mapping back:

$$\mathrm{LMS} = M_{br}\,\mathrm{XYZ}, \qquad \mathrm{LMS}' = D \cdot \mathrm{LMS}, \qquad \mathrm{XYZ}' = M_{br}^{-1}\,\mathrm{LMS}',$$

where $D = \mathrm{diag}\left(\frac{\mathrm{LMS}_{dst}}{\mathrm{LMS}_{src}}\right)$. Partial adaptation can be implemented by interpolating D with identity (adaptation strength parameter).

2.4 Local adaptation estimation

Let $\mathrm{XYZ}(x, y)$ denote the per-pixel tristimulus values and $Y(x, y)$ the luminance channel. The model estimates a locally-adapting luminance map $Y_{loc}(x, y)$ using either:

- an edge-preserving bilateral filter applied to $Y(x, y)$ (preferred), or
- a simple Gaussian blur with standard deviation σ (fallback).

The bilateral filter better preserves luminance edges (helpful to avoid halo artifacts) while still providing local averages that approximate visual adaptation.

2.5 Local tone mapping

We use a simple heuristic local tone mapping inspired by local Reinhard operators. Define the global mean luminance $\overline{Y} = \langle Y(x, y)\rangle$ and a "key" parameter k controlling overall exposure. Compute a scaled luminance:

$$L_s(x, y) = k \cdot \frac{Y(x, y)}{\overline{Y}}.$$

We then compress L_s using the local adapting luminance Y_{loc}:

$$\mathrm{denom}(x, y) = 1 + \frac{L_s(x, y)}{1 + \frac{Y_{loc}(x,y)}{\overline{Y}}}, \qquad Y_{out}(x, y) = \frac{L_s(x, y)}{\mathrm{denom}(x, y)}.$$

This form compresses high local luminance values more strongly while allowing darker regions to remain visually distinct.

2.6 Chromatic preservation / reconstruction

To preserve perceived chromaticity and hue while changing luminance, the pipeline scales the chromatic channels proportionally to the change in luminance:

$$X'(x,y) = X(x,y) \cdot \frac{Y_{out}(x,y)}{Y(x,y)+\epsilon}, \quad Y'(x,y) = Y_{out}(x,y), \quad Z'(x,y) = Z(x,y) \cdot \frac{Y_{out}(x,y)}{Y(x,y)+\epsilon}.$$

This keeps x, y chromaticity approximately constant and produces an XYZ image consistent with the mapped luminance. The small ϵ avoids division by zero.

$$\begin{aligned} \text{sRGB (nonlinear)} &\xrightarrow{\text{linearize}} \text{linear RGB} \\ &\xrightarrow{M_{\text{sRGB}\to XYZ}} \text{XYZ} \\ &\xrightarrow{\text{local } Y_{\text{loc}}} \\ &\xrightarrow{\text{tone map}} Y_{\text{out}} \\ &\xrightarrow{\text{chromatic preserve}} \text{XYZ}' \\ &\xrightarrow{M_{XYZ\to\text{sRGB}}} \text{linear RGB}' \\ &\xrightarrow{\text{gamma}} \text{sRGB}'. \end{aligned}$$

3 Experimental Setup

All experiments reported here used the demo script with the default parameters unless stated otherwise:

- key $k = 0.18$ (typical photographic key),
- Gaussian fallback $\sigma = 16$ pixels,
- bilateral parameters: $\sigma_{color} = 0.08$ (on normalized Y), $\sigma_{spatial} = 7$ px,
- both bilateral and gaussian modes run for direct visual comparison.

4 Results

From visual inspection and the basic luminance statistics:

- The bilateral-based local adaptation tends to preserve edges and reduce haloing near high-contrast boundaries relative to the gaussian variant; texture detail near edges is better preserved.
- The local tone mapping compresses highlights and reduces overall luminance contrast while maintaining local visibility of midtones and shadows.
- Chromaticity (hue) is largely preserved by the chromatic-preservation step; saturation may appear to shift slightly because absolute luminance scaling changes perceived saturation.

5 Discussion and Inference

5.1 How the pipeline maps to perceptual ideas

The pipeline captures two central perceptual mechanisms:

1. **Adaptation:** Human visual sensitivity is governed by local adaptation levels; bright local surrounds desensitize perception of luminance increments. Estimating Y_{loc} approximates this effect.

2. **Color constancy / chromatic preservation:** Maintaining chromaticity while changing luminance preserves perceived hue and avoids color casts; our simple proportional scaling is an effective practical approximation.

5.2 Limitations

- The current model does not include a multi-scale band-pass decomposition or explicit CSF weighting; such components in full iCAM produce more accurate spatial dependency behavior.

- Chromatic induction (surround-induced chromatic shifts) is not modeled explicitly.

- The tone mapping function is heuristic; for perceptual fidelity across many images, a more principled mapping (using appearance-space nonlinearity) is better.

- Metrics: we provided only simple luminance-based statistics; perceptual evaluation requires psychophysical measurements or perceptual quality metrics (MS-SSIM, HDR-VDP, etc.).

6 Conclusions

The implemented iCAM-like pipeline provides a compact and practical approach to include local adaptation and appearance-preserving tone mapping in image processing. It is suitable as a building block for image editing, display-prep, and perceptual preprocessing.

This can be taken further through the following ways:

- Implement multi-scale (Laplacian pyramid) decomposition with band-pass filters weighted by a CSF (closer to iCAM).

- Add explicit opponent-space nonlinearities and surround induction steps from iCAM.

- Incorporate an objective perceptual metric for quantitative evaluation; run subjective tests for validation.

- Provide an interactive UI (Streamlit) to tweak parameters (key, bilateral sigmas, scale) and preview effects live.

References

1. Fairchild, M. D. (2013). *Color Appearance Models*, 3rd ed. Wiley. (for color appearance theory and models)

2. Local tone mapping literature (Reinhard et al., 2002) and iCAM original papers for full model details.