Friction Prediction via Illumination Ratios

Shayan Nasiri
Department of Computer Science, University of Toronto

Motivation

What? With advancements in the robotics industry, robot stabilization is an important consideration to prevent falls and damages

Why? Existing methods often require contact with surfaces to predict slipperiness, or are computationally too expensive to run online

How? Apparent reflectance and micro-texture correlate with slipperiness. Cross-illumination ratios suppress shading and expose better reflectance cues, allowing better friction prediction. This project aims to derive features from multi-illumination ratios and train a linear regression model for predicting coefficients of friction (CoF)

ROI (cropped)

Related Work

- **Brandão et al.** built the GTF dataset. Baselines use image cues from a Retinex shading layer (ShadMax, ShadStd) and gradients (GradMu). Best results come from utilizing material labels and word embeddings (WordMM/MS), which rely on knowing or guessing the material, not great for online use [1].
- **Zhang et al.** use a special optics rig (camera + beam splitter + parabolic mirror) to make reflectance disks, then a CNN to predict CoF. Works well but hardware is heavy and not robot-mountable [2].
- **Finlayson et al.** show a simple channel-wise diagonal (von Kries) transform models illumination changes. We use this to simulate colored LEDs and then take cross-illumination luminance ratios for feature extraction [3].

References

[1] Brandão, Hashimoto, Takanishi, *Friction from vision: A study of algorithmic and human performance with consequences for robot perception and teleoperation*, IEEE-RAS International Conference on Humanoid Robots, 2016

[2] . Zhang, K. Dana, and K. Nishino, "Friction from reflectance: Deep reflectance codes for predicting physical surface properties from one-shot in-field reflectance," in Lecture Notes in Computer Science,2016

[3] Finlayson, Drew, Funt, *Diagonal transforms suffice for color constancy*, International Conference on Computer Vision, 1994

New Technique

Simulate controlled color illuminations via a von Kries diagonal transform, convert each to luminance, take log-ratios against the white-light luminance to suppress shading, and summarize the ratios (mean or std) to predict CoF with a simple linear regression (LR) model

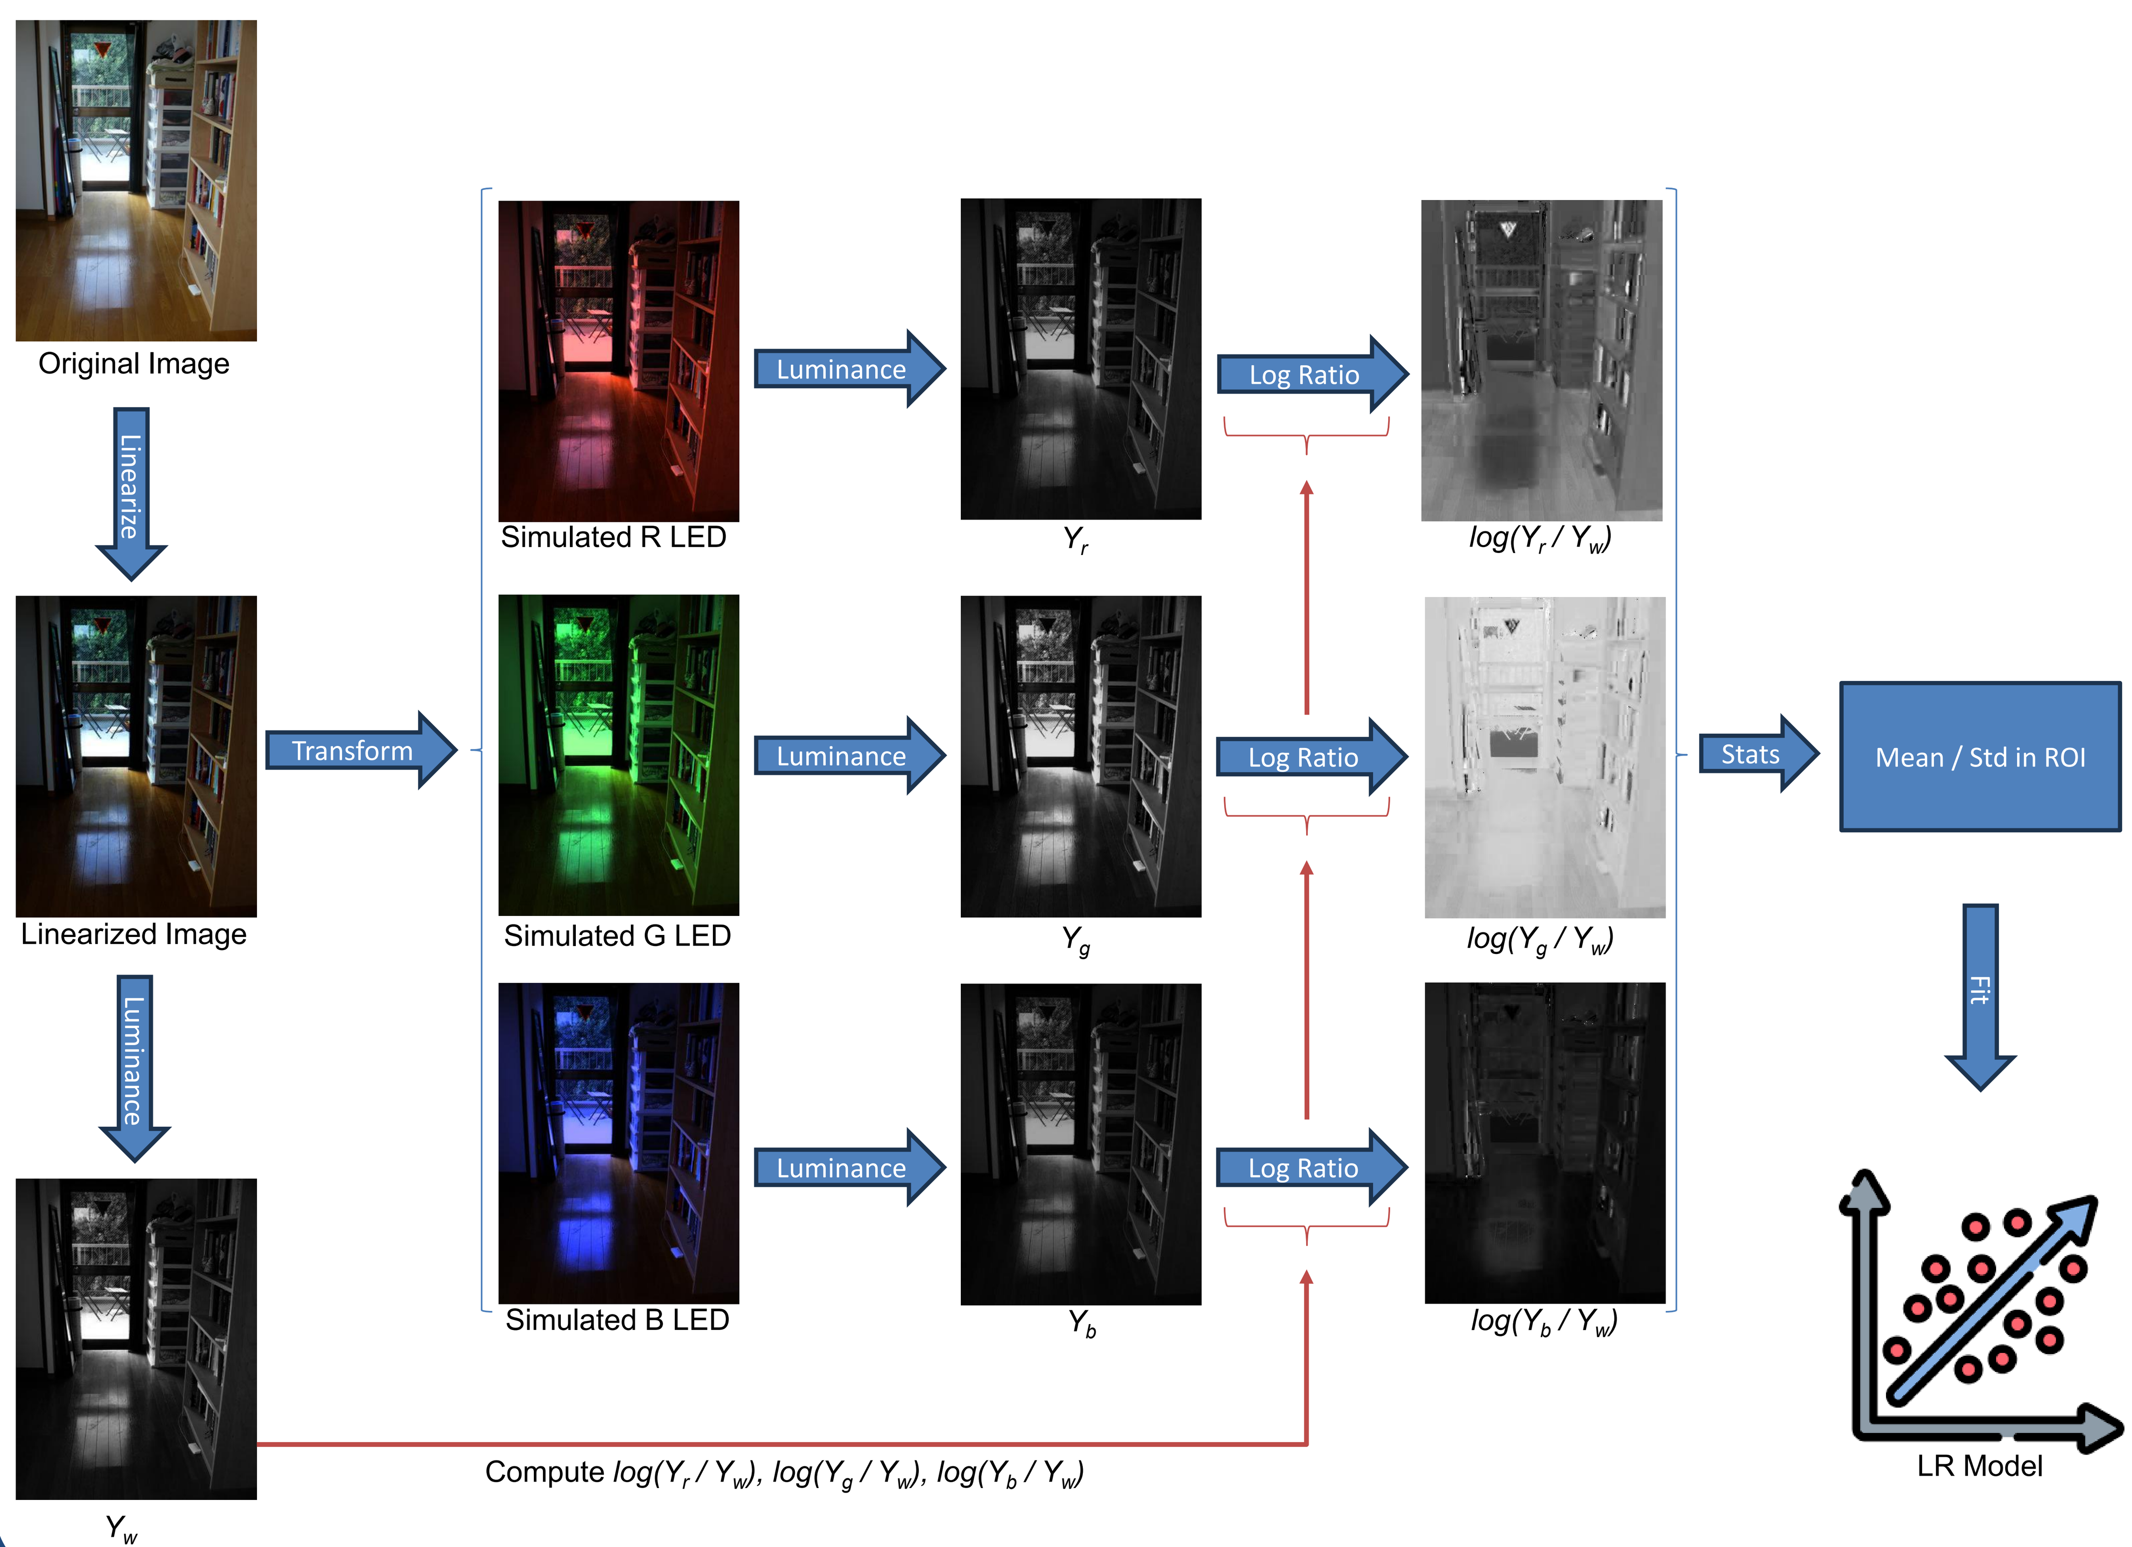

Experimental Results

We evaluate the method using **root mean squared error (RMSE)** under 10-fold cross validation, and compare performance against **DINOv2-based learned features**

- Image Only Baseline

Feature	10-CV RMSE
ShadMax	0.173
ShadStd	0.177
GradMu	0.172

- Cross Illumination Ratio Features

Ratios Used	10-CV RMSE (mean)	10-CV RMSE (std)
$log(Y_r / Y_w)$, $log(Y_g / Y_w)$, $log(Y_b / Y_w)$	0.168	0.195
$log(Y_r / Y_w)$, $log(Y_g / Y_w)$	0.165	0.188
$log(Y_r / Y_w)$, $log(Y_b / Y_w)$	0.167	0.179
$log(Y_g / Y_w)$, $log(Y_b / Y_w)$	0.170	0.181
$log(Y_r / Y_w)$	**0.163**	0.177
$log(Y_g / Y_w)$	0.166	0.175
$log(Y_b / Y_w)$	0.169	0.172

- DINOv2 Learned Features

Features	10-CV RMSE
DINOv2	**0.165**
DINOv2 + $log(Y_r / Y_w)$	**0.165**

Conclusions

- Cross-illumination ratios with linear regression **improve** CoF prediction over Brandão et al.'s image-only baselines (best RMSE **0.163 vs. 0.172–0.177**) while **avoiding complex pipelines**, making the method lightweight and suitable for real-time use.
- The approach marginally **outperforms** a modern learned-feature baseline (DINOv2, **0.165 RMSE**).
- Adding $log(Y_r / Y_w)$ to DINOv2 features brings no further gain (still 0.165), suggesting DINOv2 already encodes similar features.