ABSTRACT

The prime goal of digital imaging techniques is to reproduce the real-world appearance of the scene faithfully. Traditional camera sensors often struggle to handle real-world lighting contrasts such as bright lamps at night and shadows under the sun. The ability to capture and display both highlights and shadows in the same scene is characterized by the dynamic range of the device. Low Dynamic Range (LDR) cameras are incapable of representing the wide dynamic range of the actual scene. The captured image regions often turn out to be either too dark (underexposed) or too bright (overexposed). The recent advancements in High Dynamic Range (HDR) imaging have proved to overcome these low dynamic range issues and can plausibly reproduce human visual system characteristics. However, they require high-end HDR cameras, which currently seem out of reach of regular consumers.

In this thesis, we propose a deep neural network pipeline that attempts to reconstruct an HDR image from a single LDR image. We formulate this problem as an image-to-image (I2I) translation task. To this end, we present a novel conditional GAN (cGAN) based framework trained in an end-to-end fashion over the HDR-REAL and HDR-SYNTH datasets. Since, saturation in overexposed regions makes this task challenging, our network uses an overexposed mask obtained from a pre-trained segmentation model to facilitate the hallucination task over saturated regions. We demonstrate the effectiveness of the proposed method by performing an extensive quantitative and qualitative comparison with state-of-the-art approaches for the single-image based HDR reconstruction.