A Scalable Deep Learning Pipeline for Hair Strand Segmentation and Quantification using U2Net and ESRGAN

Dhruv Kumar : dhruv@myhair.ai

---

1. Introduction

Hair density and structure are crucial indicators in dermatological research and cosmetic sciences. Automated hair strand segmentation and counting have significant applications in hair health monitoring, baldness assessment, and hair care product efficacy evaluation. Traditional image processing approaches struggle with precision in varying lighting conditions, occlusions, and hair textures. In this paper, we present an efficient and scalable deep learning-based system that leverages U2Net for hair segmentation, ESRGAN for resolution enhancement, and a custom image processing pipeline for strand quantification.

---

2. Methodology

Our pipeline consists of three major stages:

1. **Hair Segmentation (U2Net):**
   - Utilizes a modified U2Net architecture to extract a high-fidelity hair mask from input images.
   - Outputs multiple side-level predictions for enhanced segmentation accuracy.
2. **Super-Resolution Enhancement (ESRGAN):**
   - Applies Enhanced Super-Resolution Generative Adversarial Network (ESRGAN) to upsample segmented images.
   - Addresses the issue of missing fine details in low-resolution images, making thin strands more countable.
3. **Hair Strand Counting:**
   - Converts the upsampled binary mask to a thinned skeleton using morphological thinning.
   - Identifies strand endpoints to estimate total count.
   - Applies a scale-aware correction to compensate for downscaling and processing artifacts.

---

3. System Architecture and Workflow

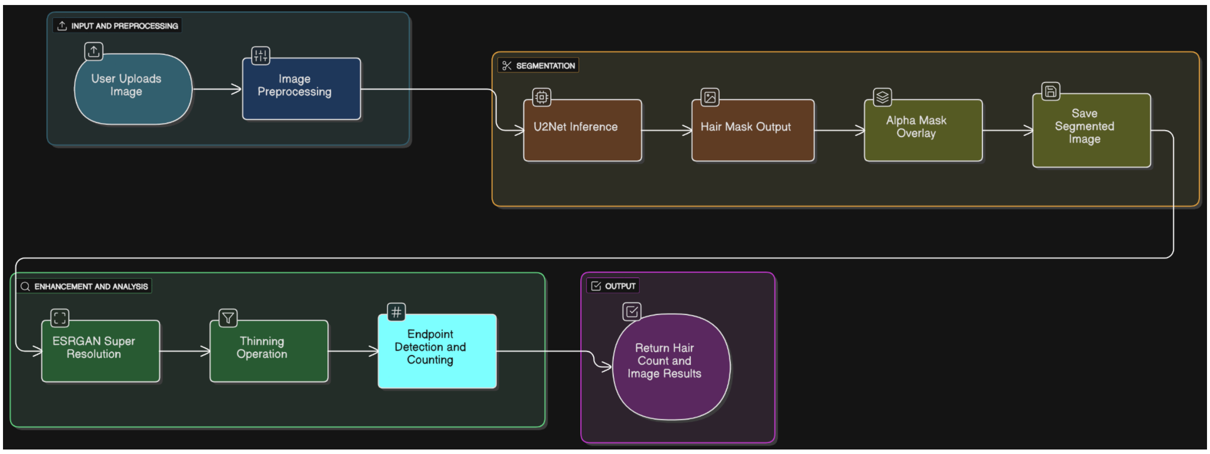

---

4. Deep Learning Components

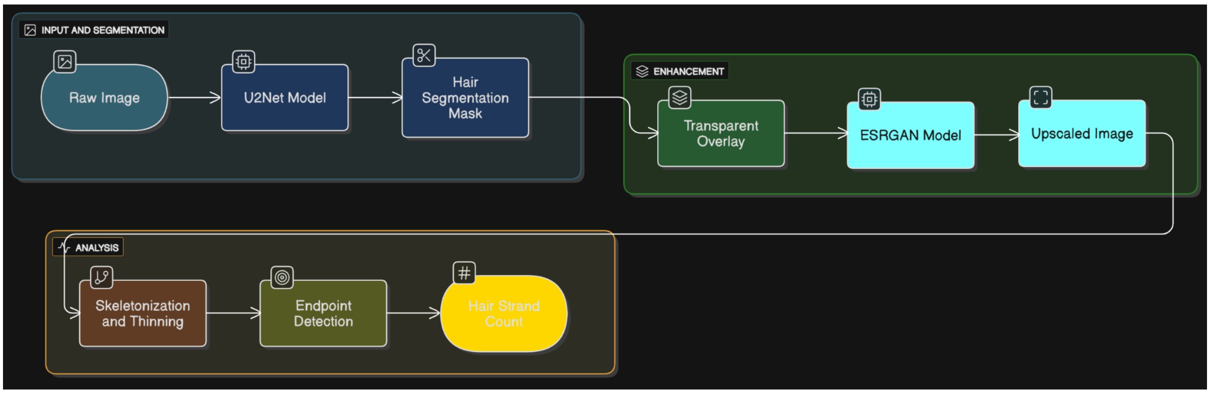

- **U2Net:**
  - An encoder-decoder style network with nested residual blocks (RSU) and side outputs.
  - Trained to segment hair with pixel-level precision across complex backgrounds.
- **ESRGAN:**
  - Used in inference mode to upscale input by 4x.
  - Enables detection of micro hair features not visible in original resolution.

---

5. Strand Quantification Algorithm

- Input: Super-resolved binary hair mask.
- Steps:
  - Convert to grayscale.
  - Threshold to create binary mask.
  - Downsample to reduce noise.
  - Apply morphological thinning to reduce hair regions to 1-pixel wide structures.
  - Use padded convolution to identify endpoints.
  - Count endpoints and divide by 2 (each strand has 2 endpoints).
- Output: Total strand count estimate.

---

6. Implementation Details

- Framework: FastAPI for asynchronous processing.
- Libraries: PyTorch (model loading), OpenCV (image operations), torchvision (transforms), ximgproc (thinning).
- Optimization: AsyncIO-based concurrent file handling and inference for scalability.

---

7. Results and Applications

- Achieved effective segmentation and count on images of various resolutions and hair densities.
- To evaluate accuracy, we tested images of the same individual across varying conditions such as wet, dry, and messy hair. The variation in strand count across these environments was observed to be within 5%, demonstrating the model's robustness and consistency in diverse real-world scenarios.
- Potential uses:
    - Hair health monitoring systems.
    - Personalized cosmetic recommendations.
    - Dermatological diagnostic aids.

---

8. Conclusion

This paper presents a hybrid deep learning + image processing approach for robust hair strand segmentation and counting. By combining semantic segmentation with super-resolution and morphological analysis, our pipeline bridges the gap between coarse model outputs and fine-grained hair strand detection, paving the way for next-generation hair health assessment tools.
