

Nail Polish Try-On: Realtime Semantic Segmentation of Small Objects for Native and Browser Smartphone AR Applications

Brendan Duke Abdalla Ahmed Edmund Phung Irina Kezele Parham Aarabi

ModiFace Inc.

{brendan,abdalla,edmund,irina,parham}@modiface.com

Abstract

We provide a system for semantic segmentation of small objects that enables nail polish try-on AR applications to run client-side in realtime in native and web mobile applications. By adjusting input resolution and neural network depth, our model design enables a smooth trade-off of performance and runtime, with the highest performance setting achieving 94.5 mIoU at 29.8ms runtime in native applications on an iPad Pro. We also provide a postprocessing and rendering algorithm for nail polish try-on, which integrates with our semantic segmentation and fingernail base-tip direction predictions. Project page: <https://research.modiface.com/nail-polish-try-on-cvprw2019>

1. Introduction

We present our end-to-end solution for simultaneous realtime tracking of fingernails and rendering of nail polish. Our system locates and identifies fingernails from a videostream in realtime with pixel accuracy, and provides enough landmark information, e.g., directionality, to support rendering techniques.

We collected an entirely new dataset with semantic segmentation and landmark labels, developed a high-resolution neural network model for mobile devices, and combined these data and this model with postprocessing and rendering algorithms for nail polish try-on.

We deployed our trained models on two hardware platforms: iOS via CoreML, and web browsers via TensorFlow.js [8]. Our model and algorithm design is flexible enough to support both the higher computation native iOS platform, as well as the more resource constrained web platform, by making only minor tweaks to our model architecture, and without any major negative impact on performance.

1.1. Contributions

Below we enumerate our work’s novel contributions.

- We created a dataset of 1438 images sourced from photos and/or videos of the hands of over 300 subjects, and annotated with foreground-background, per-finger class, and base-tip direction field labels.

- We developed a novel neural network architecture for semantic segmentation designed for both running on mobile devices and precisely segmenting small objects.
- We developed a postprocessing algorithm that uses multiple outputs from our fingernail tracking model to both segment fingernails and localize individual fingernails, as well as to find their 2D orientation.

2. Related Work

Our work builds on MobileNetV2 [7] by using it as an encoder-backbone in our cascaded semantic segmentation model architecture. However, our system is agnostic to the specific encoder model used, so any existing efficient model from the literature [5, 9, 10, 12] could be used as a drop-in replacement for our encoder.

Our Loss Max-Pooling (LMP) loss is based on [1], where we fix their p -norm parameter to $p = 1$ for simplicity. Our experiments further support the effectiveness of LMP in the intrinsically class imbalanced fingernail segmentation task.

Our cascaded architecture is related to ICNet [11] in that our neural network model combines shallow/high-resolution and deep/low-resolution branches. Unlike ICNet, our model is designed to run on mobile devices, and therefore we completely redesigned the encoder and decoder based on this requirement.

3. Methods

3.1. Dataset

Due to a lack of prior work specifically on fingernail tracking, we created an entirely new dataset for this task. We collected egocentric data from participants, who we asked to take either photos or videos of their hands as if they were showing off their fingernails for a post on social media.

Our annotations included a combination of three label types: foreground/background polygons, individual fingernail class polygons, and base/tip landmarks to define per-polygon orientation. We used the fingernail base/tip landmarks to generate a dense direction field, where each pixel




Figure 1: A top level overview of our nail segmentation model. In this example, the input resolution is $H \times W$. See Figure 2 for fusion and output block architectures.




Figure 2: The fusion module (left) and output branch (right) of our model. We fuse upsampled low-resolution, deep features F_1 with high-resolution, shallow features F_2 to produce high-resolution fused features F_2' in the decoder. The output branch is used in each fusion block, and at the full-resolution output layer of the model.

contains a base-to-tip unit vector for the fingernail that that pixel belongs to.

Our annotated dataset consists of 1438 annotated images in total, which we split into train, validation and test sets based on the participant who contributed the images. The split dataset contains 941, 254 and 243 images each in train, val and test, respectively.

3.2. Model

The core of our nail tracking system is an encoder-decoder convolutional neural network (CNN) architecture trained to output foreground/background and fingernail class segmentations, as well as base-tip direction fields. Our model architecture draws inspiration from ICNet [11], which we improved to meet the runtime and size constraints

of mobile devices, and to produce our multi-task outputs. A top-level view of our model architecture is in Figure 1.

We initialized the encoder of our model with MobileNetV2 [7] model weights pretrained on ImageNet [2]. We used a cascade of two $\alpha = 1.0$ MobileNetV2 encoder backbones, both pretrained on 224×224 ImageNet images. The encoder cascade consists of one shallow network with high-resolution inputs (stage_high1..4), and one deep network with low-resolution inputs (stage_low1..8), both of which are prefixes of the full MobileNetV2. To maintain a higher spatial resolution of feature maps output by the low-resolution encoder we changed stage 6 from stride 2 to stride 1, and used dilated $2 \times$ convolutions in stages 7 and 8 to compensate, reducing its overall stride from $32 \times$ to $16 \times$.

Our model’s decoder is shown in the middle and bottom right of Figure 1, and a detailed view of the decoder fusion block is shown in Figure 2. For an original input of size $H \times W$, our decoder fuses the $\frac{H}{16} \times \frac{W}{16}$ features from stage_low4 with the upsampled features from stage_low8, then upsamples and fuses the resulting features with the $\frac{H}{8} \times \frac{W}{8}$ features from stage_high4. The convolution on F_1 is dilated to match F_1 ’s $\times 2$ upsampling rate, and the elementwise sum (to form F_2') allows the projected F_2 to directly refine the coarser F_1 feature maps.

As shown in Figure 2, a 1×1 convolutional classifier is applied to the upsampled F_1 features, which are used to predict downsampled labels. As in [3], this “Laplacian pyramid” of outputs optimizes the higher-resolution, smaller receptive field feature maps to focus on refining the predictions from low-resolution, larger receptive field feature maps.

Also shown in Figure 2, our system uses multiple output decoder branches to provide directionality information (i.e., vectors from base to tip) needed to render over fingernail tips, and fingernail class predictions needed to find fingernail instances using connected components. We trained these additional decoders to produce dense predictions penalized only in the annotated fingernail area of the image.

3.3. Criticism

To deal with the class imbalance between background (overrepresented class) and fingernail (underrepresented class), in our objective function we used LMP [1] over all pixels in a minibatch by sorting by the loss magnitude of each pixel, and taking the mean over the top 10% of pixels as the minibatch loss. Compared with a baseline that weighted the fingernail class’s loss by $20 \times$ more than the background, LMP yielded a gain of $\approx 2\%$ mIoU, reflected in sharper nail edges where the baseline consistently over-segmented.

We used three loss functions corresponding to the three outputs of our model shown in Figure 2. The fingernail class

and foreground/background predictions both minimize the negative log-likelihood of a multinomial distribution given in Equation 1, where c is the ground truth class and x_c^{ij} is the pre-softmax prediction of the c th class by the model. Each of the following equations refers to indices i and j corresponding to pixels at $(x, y) = (i, j)$.

$$\mathcal{L}_{class,fgbg}^{ij} = -\log\left(\frac{\exp x_c^{ij}}{\sum_{c'} \exp x_{c'}^{ij}}\right) \quad (1)$$

In the case of class predictions, $c \in \{1, 2, \dots, 10\}$, while for foreground/background predictions $c \in \{1, 2\}$. We used LMP for foreground/background predictions only; since fingernail class predictions are only valid in the fingernail region, those classes are balanced and do not require LMP.

$$\ell_{fgbg} = \frac{1}{N_{thresh}} \sum_{ij} [\mathcal{L}_{fgbg}^{ij} > \tau] \cdot \mathcal{L}_{fgbg}^{ij} \quad (2)$$

In Equation 2, $N_{thresh} = \sum_{ij} [\mathcal{L}_{fgbg}^{ij} > \tau]$, and threshold τ is the loss value of the $[0.1 \times H \times W]$ th highest loss pixel. The $[\cdot]$ operator is the indicator function.

For the direction field output, we apply an L_2 loss (other loss functions such as Huber and L_1 could be used in its place) on the normalized base to tip direction of the nail for each pixel inside the ground truth nail.

$$\ell_{field} = \frac{1}{H \times W} \sum_{ij} \|\hat{\mathbf{u}}_{field}^{ij} - \mathbf{u}_{field}^{ij}\|_2^2 \quad (3)$$

The field direction labels are normalized so that $\|\mathbf{u}_{field}^{ij}\|_2 = 1$. Since the direction field and the fingernail class losses have no class imbalance problem, we simply set these two losses to the means of their respective individual losses (i.e., $\ell_{class} = \frac{1}{H \times W} \sum_{ij} \mathcal{L}_{class}^{ij}$). The overall loss is $\ell = \ell_{fgbg} + \ell_{class} + \ell_{field}$.

3.4. Postprocessing and Rendering

Our postprocessing and rendering algorithm uses the output of our tracking predictions to draw realistic nail polish on a user’s fingernails. Our algorithm uses the individual fingernail location and direction information predicted by our fingernail tracking module in order to render gradients, and to hide the light-coloured distal edge of natural nails.

We first use connected components [4] on the class predictions to extract a connected blob for each nail. We then estimate each nail’s angle by using the foreground (nail) softmax scores to take a weighted average of the direction predictions in each nail region. We use the fingernails’ connected components and orientations to implement rendering effects, including rendering gradients to approximate specular reflections, and stretching the nail masks towards the fingernail tips in order to hide the fingernails’ light-coloured distal edges.

| Model | mIoU |
|-----------|------|
| Baseline | 91.3 |
| + LMP | 93.3 |
| + Cascade | 94.5 |




Figure 3: Runtime/mIoU comparison for model settings on an iPad Pro device (right) and ablation study of LMP and cascade model architecture features (left).

3.5. Performance and Runtime Results

We trained our neural network models using PyTorch [6]. We deployed the trained models to iOS using CoreML, and to web browsers using TensorFlow.js [8].

Given current software implementations, namely CoreML and TensorFlow.js, and current mobile device hardware, our system could run in realtime (i.e., ≥ 10 FPS) at all resolutions up to 640×480 (native mobile) and 480×360 (web mobile), for which we trained on randomly-cropped input resolutions of 448×448 and 336×336 , respectively.

In Figure 3 we evaluate performance in binary mIoU on our validation set against runtime on an iPad Pro device on CoreML (native iOS) and TensorFlow.js (Safari web browser) software platforms using two variants of our model (55-layer and 43-layer encoders) at three resolutions (288px, 352px, and 480px), and compare the effect of introducing LMP and the low/high-resolution cascade into our model design.

4. Workshop Demo

For our submission we propose both a web browser-based and iOS demo of our virtual nail polish try-on technology. We will provide an iOS tablet device for users to try the virtual nail try-on application in person, as well as a web link at <https://ola.modiface.com/nailsweb/cvpr2019demo> (best viewed on a mobile device using Safari on iOS, Chrome or Firefox on Android) that users can access with their own devices and use the technology by enabling camera access in the browser (“https” prefix required).

5. Conclusion

We presented a fingernail tracking system that runs in realtime on both iOS and web platforms, and enables pixel-accurate fingernail predictions at up to 640×480 resolution. We also proposed a postprocessing and rendering algorithm that integrates with our model to produce realistic nail polish rendering effects.

References

- [1] Samuel Rota Bulò, Gerhard Neuhold, and Peter Kontschieder. Loss max-pooling for semantic image segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 2
- [3] Golnaz Ghiasi and Charless C. Fowlkes. Laplacian reconstruction and refinement for semantic segmentation. In *ECCV*, 2016. 2
- [4] C. Grana, D. Borghesani, and R. Cucchiara. Optimized block-based connected components labeling with decision trees. *IEEE Transactions on Image Processing*, 2010. 3
- [5] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size. *arXiv:1602.07360*, 2016. 1
- [6] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. 3
- [7] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1, 2
- [8] Daniel Smilkov, Nikhil Thorat, Yannick Assogba, Ann Yuan, Nick Kreeger, Ping Yu, Kangyi Zhang, Shanqing Cai, Eric Nielsen, David Soergel, Stan Bileschi, Michael Terry, Charles Nicholson, Sandeep N. Gupta, Sarah Sirajuddin, D. Sculley, Rajat Monga, Greg Corrado, Fernanda B. Viégas, and Martin Wattenberg. Tensorflow.js: Machine learning for the web and beyond. *arXiv preprint arXiv:1901.05350*, 2019. 1, 3
- [9] Robert J Wang, Xiang Li, and Charles X Ling. Pelee: A real-time object detection system on mobile devices. In *Advances in Neural Information Processing Systems 31*, 2018. 1
- [10] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [11] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnets for real-time semantic segmentation on high-resolution images. In *ECCV*, 2018. 1, 2
- [12] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1