GONE Offers Nice Editing

Jing Lim*, Travis McGuire**, Alexis Goh*

*Department of Computer Science Stanford University, **Department of Electrical Engineering Stanford University

{jinglim2, tmcgurie, gweiying}@stanford.edu

Introduction

End-to-end **detection** and **deletion** of background **people** from images using **Mask R-CNN**[1] and **Generative Inpainting**[2] on **MS COCO 2014**[3]

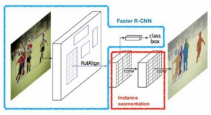

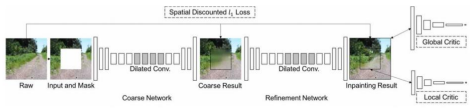

Experiments & Hyperparameter Tuning

Database Correction

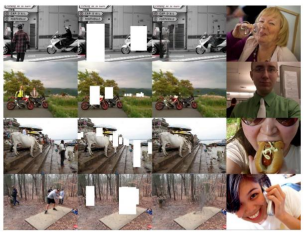

Left to right: Model Baseline, Large Crop Examples

	With large crop	Without large crop
Good	13	19
Okay	12	11
Striped	23	41
Solid	52	29
Total	100	100

Table 1: Number of classifications by quality and crop size

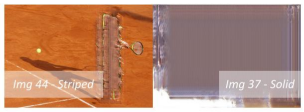

Inpainted with pretrained weights - Places256

Box vs Mask vs Dilated Mask

Left to Right: Box, Mask, Dilated Mask Crops

Left to right: Inpainting on Box, Mask, Dilated Mask

Learning Rate

Learning rate: 1e-05, 1e-04, 1e-03

Other **hyperparameter tunings** with *less conclusive* findings:

- Training with **'no people'** dataset
- **ℓ1-pixelwise**[4] reconstruction loss

Results & Discussion

Original, Results from epochs 4,8,12 of trained GAN on MS COCO

Initial results were *poor*, but **improved** through **error analysis**. Future experiments to train on **'no people'** data and **vary ℓ1-pixelwise** loss values seem **promising** for further improvement.