


End-to-end Piecewise Unwarping of Documents

Sagnik Das^{2,1}, Kunwar Yashraj Singh¹, Jon Wu¹, Erhan Bas¹, Vijay Mahadevan¹, Rahul Bhotika¹, Dimitris Samaras²



Motivation

- Prior patch-based local unwarping approaches are:
 - **Not end-to-end trainable.**
 - **Image level patch stitching strategies leave artifacts.**
- Prior **end-to-end** approaches predict a global unwarping grid, often resulting:
 - **Less robust local unwarping.**
 - **Undesired warped regions.**



(a) Proposed, (b) DewartNet

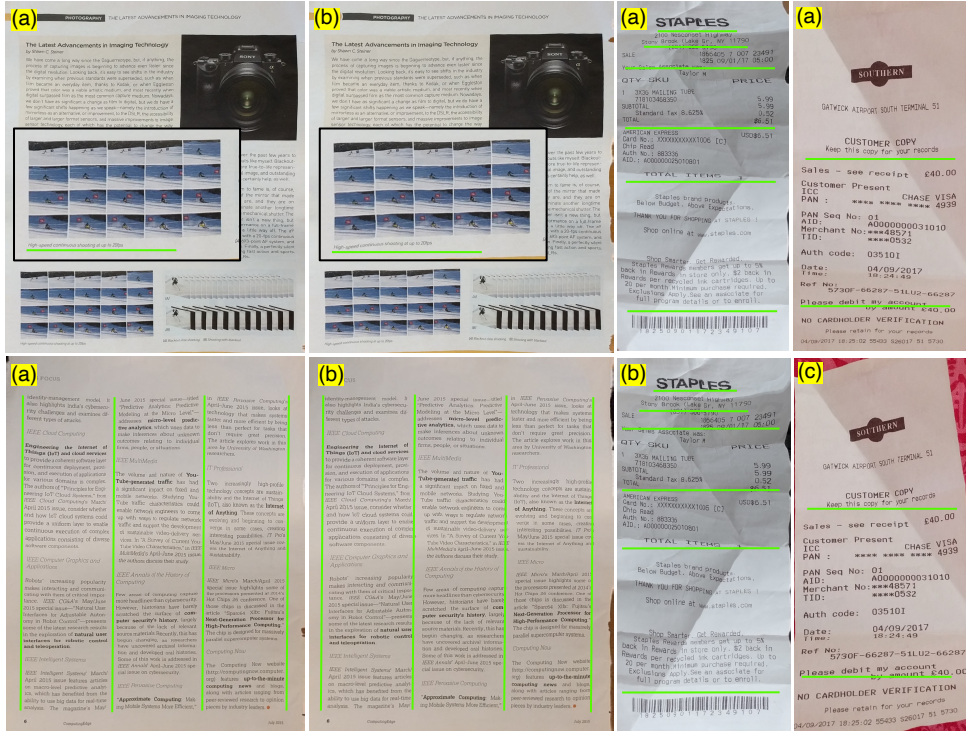
(a) Proposed, (b) DocProj

Contributions

- The first **end-to-end trainable piece-wise** framework which unwarpes local deformations.
- The first **fully differentiable stitching network**: takes per-patch unwarping maps as input and produces a global unwarping map.
- We show significant **improvement in local unwarping quality**, which results in:
 - Better (+5%) image similarity metric (MS-SSIM).
 - Better and more stable (+3%) OCR Performance.

Acknowledgements:
This work was partially supported by the Partner University Fund, the SUNY2020 ITSC, and a grant from Amazon AI.

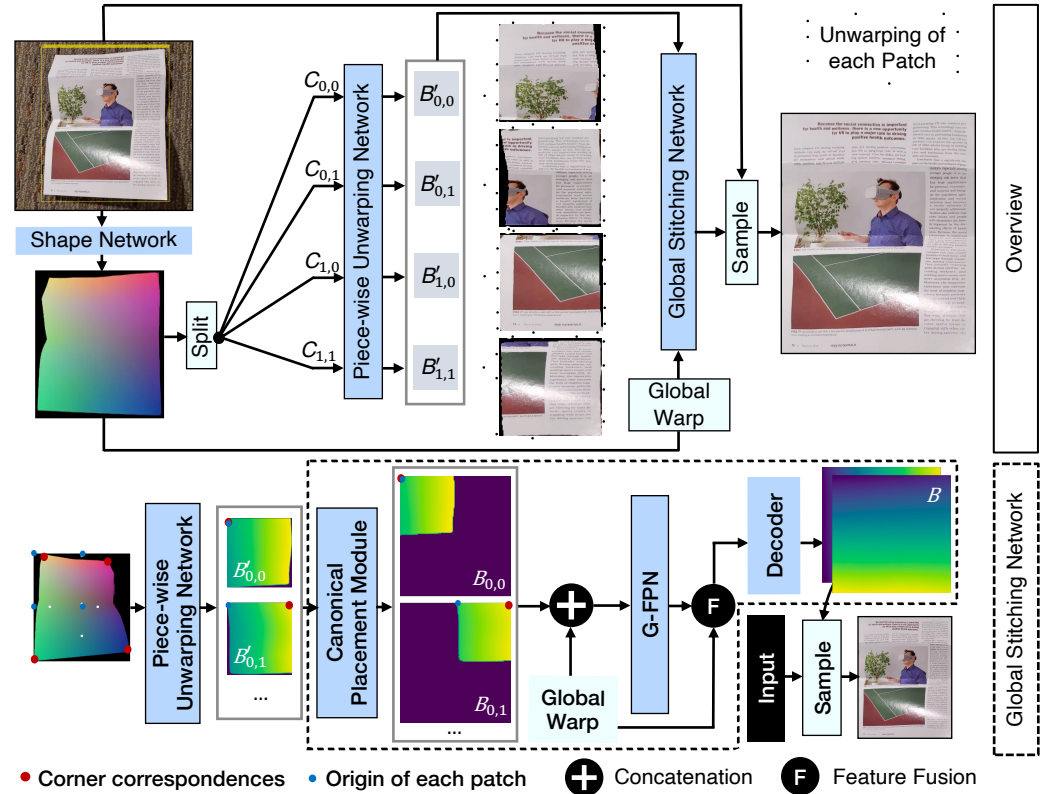
Experimental Results



Follow green cue lines for better visualization.

(a) Proposed, (b) DewartNet (c) DocProj

Method: Piecewise Unwarping



• Corner correspondences • Origin of each patch + Concatenation F Feature Fusion

References:
 1. Ke Ma, Zhixin Shu, Xue Bai, Jue Wang, and Dimitris Samaras. DocU-Net: Document Image Unwarping via A Stacked U-Net. In Proc. CVPR. IEEE, 2018.
 2. Sagnik Das, Ke Ma, Zhixin Shu, Dimitris Samaras, and Roy Shikhot. DewartNet: Single-image document unwarping with stacked 3D and 2D regression networks. ICCV, 2019.
 3. Xiaoyu Li, Bo Zhang, Jing Liao, and Pedro V. Sander. Document Rectification and Illumination Correction using a Patch-based CNN. ACM Transactions on Graphics (TOG), 2019