InverseFaceNet: Deep Monocular Inverse Face Rendering

Hyeongwoo Kim [1] Michael Zollhöfer [1,2] Ayush Tewari [1] Justus Thies [3] Christian Richardt [4] Christian Theobalt [1]

[1] MPI Informatics [2] Stanford University [3] Technical University of Munich [4] University of Bath

CVPR 2018 SALT LAKE CITY • JUNE 18-22

Motivation

Face reconstruction from a single image

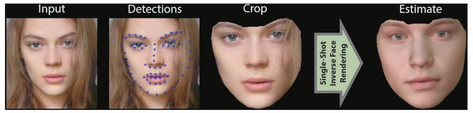

- High dimensional problem
- Runtime constraint
- Face identification and editing

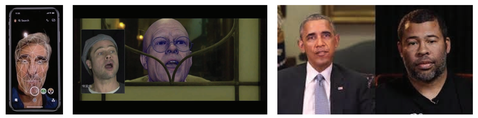

Contributions

Deep convolutional inverse rendering framework for faces

Training Corpus → Training Corpus with Domain Shift → Real World Images without Ground Truth

- Joint estimation of all face rendering parameters
- Single input image
- Single shot inference
- Real-time
- Domain shift from synthetic to real-world imagery

Method

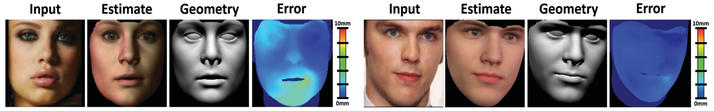

Parameterization

Rendering face images with 350 parameters

$$\boldsymbol{\theta} = \left(\mathbf{R}, \boldsymbol{\theta}^{[\mathrm{s}]}, \boldsymbol{\theta}^{[\mathrm{e}]}, \boldsymbol{\theta}^{[\mathrm{r}]}, \boldsymbol{\theta}^{[\mathrm{i}]}\right) \in \mathbb{R}^m$$

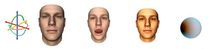

$$P = (\mathrm{R}, \theta^{[s]}, \theta^{[e]}, \theta^{[r]}, \theta^{[i]})$$

pose shape expression reflectance illumination

Sampling

Synthetic images to train

Distribution provided by a PCA face model

$$P(\boldsymbol{\theta}) = P(\mathbf{R})\, P(\boldsymbol{\theta}^{[\mathrm{s}]})\, P(\boldsymbol{\theta}^{[\mathrm{e}]})\, P(\boldsymbol{\theta}^{[\mathrm{r}]})\, P(\boldsymbol{\theta}^{[\mathrm{i}]})$$

$$\begin{aligned}
\alpha, \beta &\sim \mathcal{U}(-40^\circ, 40^\circ) \\
\gamma &\sim \mathcal{U}(-15^\circ, 15^\circ) \\
\boldsymbol{\theta}^{[\mathrm{s}]}, \boldsymbol{\theta}^{[\mathrm{r}]} &\sim \mathcal{N}(0, 1) \\
\boldsymbol{\theta}^{[\mathrm{e}]} &\sim \mathcal{U}(-12, 12) \\
\boldsymbol{\theta}^{[\mathrm{i}]} &\sim \mathcal{U}(-0.2, 0.2)
\end{aligned}$$

Loss

Parameter-space loss with statistical normalization

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\theta}_g) &= \left\|\boldsymbol{\theta} - \boldsymbol{\theta}_g\right\|_{\mathbf{A}}^2 \\
&= (\boldsymbol{\theta} - \boldsymbol{\theta}_g)^\top \underbrace{\mathbf{A}}_{\boldsymbol{\Sigma}^\top \boldsymbol{\Sigma}} (\boldsymbol{\theta} - \boldsymbol{\theta}_g)
\end{aligned}$$

$$\mathrm{diag}(\omega_{\mathbf{R}} \mathbf{1}_3, \omega_{\mathrm{s}} \boldsymbol{\sigma}^{[\mathrm{s}]}, \omega_{\mathrm{e}} \boldsymbol{\sigma}^{[\mathrm{e}]}, \omega_{\mathrm{r}} \boldsymbol{\sigma}^{[\mathrm{r}]}, \omega_{\mathrm{i}} \mathbf{1}_{27}) \in \mathbb{R}^{m \times m}$$

Resampling

Domain shift from synthetic to real images

New proposals from a Gaussian distribution

$$P(\mathbf{I}_r(\boldsymbol{\theta}) \mid \mathbf{I}_{\mathrm{real}}) \sim \boldsymbol{\theta}(\mathbf{I}_{\mathrm{real}}) + \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}^2)$$

$$\alpha, \beta, \gamma : \mathcal{U}(-5^\circ, 5^\circ) \quad \boldsymbol{\theta}^{[\mathrm{s}]} : \mathcal{N}(0, 0.05)$$

$$\boldsymbol{\theta}^{[\mathrm{r}]} : \mathcal{N}(0, 0.2) \qquad \boldsymbol{\theta}^{[\mathrm{e}]} : \mathcal{N}(0, 0.1) \qquad \boldsymbol{\theta}^{[\mathrm{i}]} : \mathcal{N}(0, 0.02)$$

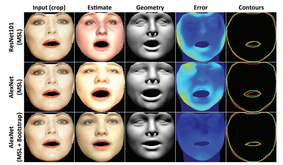

Approach	Training iterations	Training time [h]	Test time [ms / image]	Photometric error [8 bits]	Geometric error [mm]	Intersection over union [%]
AlexNet [31]	75K	4.14	**3.9**	46.26 ± 12.42	2.91 ± 0.99	90.44 ± 3.81
+ model-space loss	75K	4.36	**3.9**	39.71 ± 9.86	2.77 ± 1.00	92.51 ± 2.59
+ bootstrap (= **InverseFaceNet**)	75K*	29.40	**3.9**	34.03 ± 7.56	**2.11** ± **0.84**	**93.96** ± **2.08**
ResNet-101 [21] + model-space loss	150K	40.99	21.0	41.23 ± 10.58	2.54 ± 0.87	92.07 ± 2.87
MoFA [58]	—	—	**3.9**	**17.23** ± **4.42**	3.94 ± 1.34	84.20 ± 4.23

Results

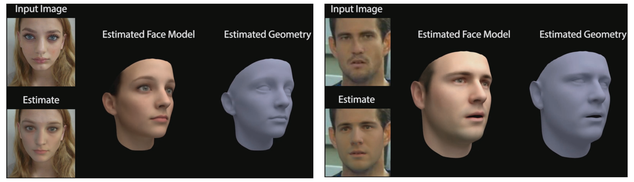

Quantitative comparison

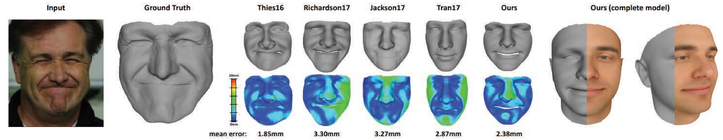

Comparison to optimization-based approaches

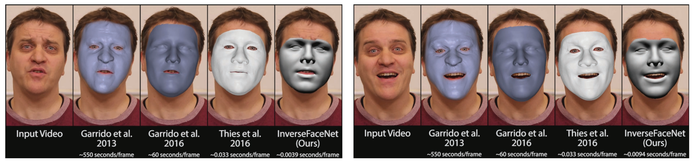

Comparison to learning-based approaches

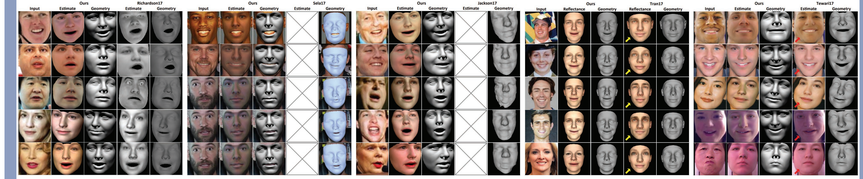

ERC Starting Grant project CapReal (335545)

Max Planck Center for Visual Computing and Communications (MPC-VCC)

RCUK grant CAMERA (EP/M023281/1)