Informed MCMC with Bayesian Neural Networks for Facial Image Analysis

Adam Kortylewski*, Mario Wieser*, Andreas Morel-Forster*, Aleksander Wieczorek, Sonali Parbhoo, Volker Roth, Thomas Vetter

** equal contribution*

Department of Mathematics and Computer Science, University of Basel, Switzerland

University of Basel

Motivation

Computer vision tasks are often **difficult** because of the large **variability** in the data that is induced by changes in light, background, partial occlusion as well as the pose, texture and shape of objects.

Generative approaches to computer vision such as **3D Morphable Models** (3DMM) [1] allow us to overcome this difficulty by **explicitly modeling** the physical **image formation process**.

Facial Image Analysis with 3DMM

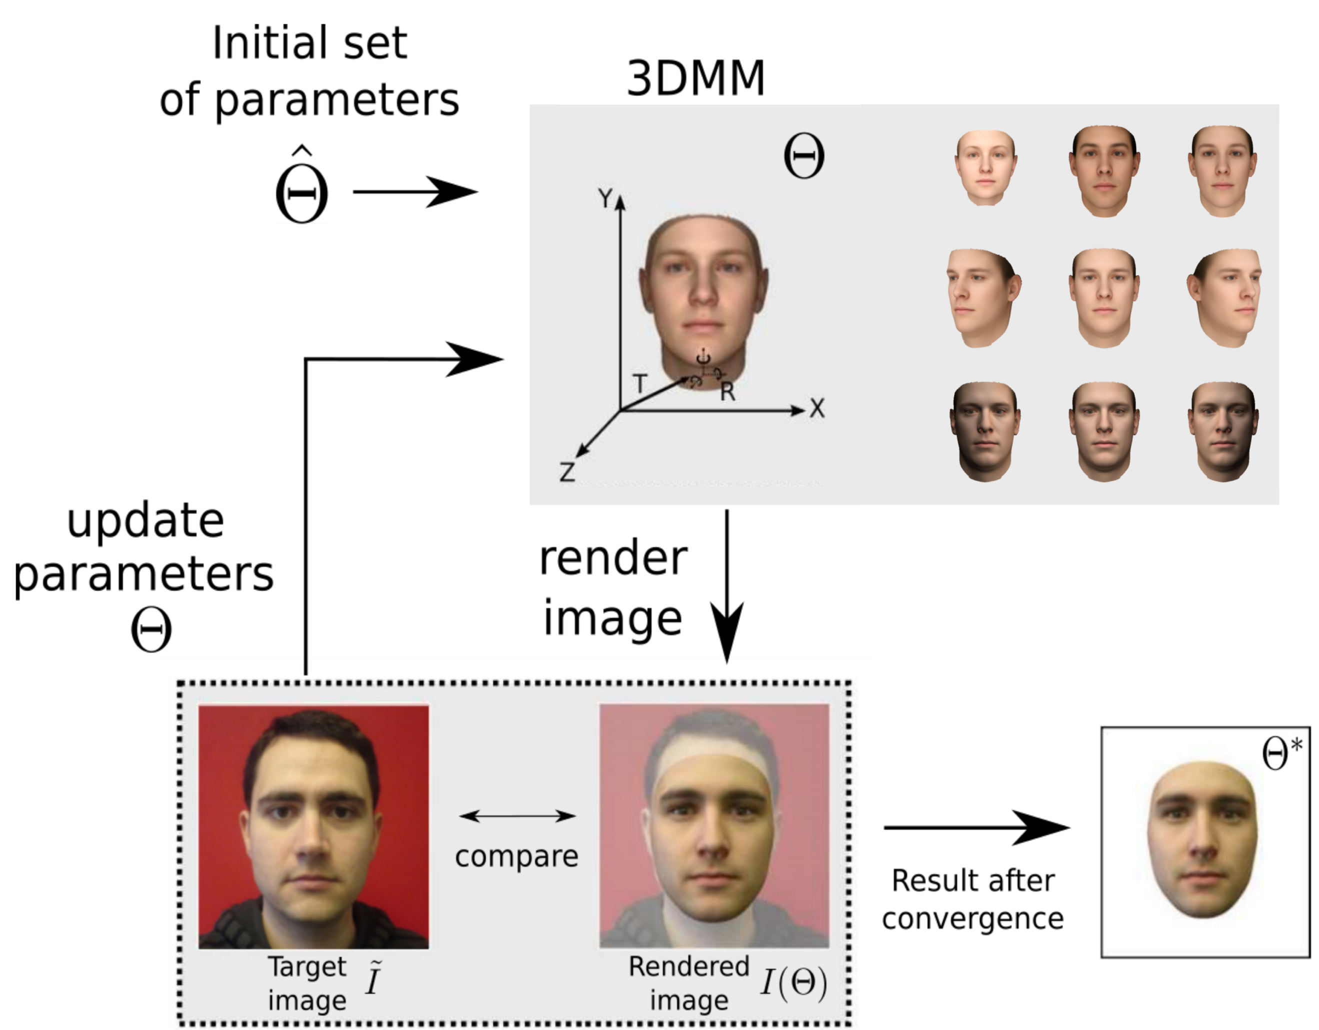

Posterior Estimation: We sample $p(\Theta|\tilde{I})$ with the Metropolis-Hastings (MCMC) in a **two-step** process:

- Generate a new point from the proposal distribution: $\Theta_{t+1} \sim Q(\cdot|\Theta_t)$.
- Accept with acceptance probability: $A(\Theta_{t+1}, \Theta_t) = \min\left(1, \frac{p(\Theta_{t+1})Q(\Theta_t|\Theta_{t+1})}{p(\Theta_t)Q(\Theta_{t+1}|\Theta_t)}\right)$.

Problem

- Time to **convergence** of the posterior inference strongly depends on a **careful design** of the **proposal distribution** $Q(\cdot|\Theta_t)$.

Informing MCMC with BNN

Informed sampling:

- decompose the proposal distribution into **local** Q_L and **global** Q_I as in [3],

$$\Theta_{t+1} \sim \alpha Q_L(\cdot|\Theta_t) + (1-\alpha)Q_I(\cdot|x).$$

- make global Q_I depend on data (estimate conditional density $\Theta \mid x$).

Bayesian Neural Networks:

- for model uncertainty, use a prior distribution on network's weights:

$$p(W \mid X, \Theta) = \frac{p(X, \Theta \mid W)p(W)}{\int p(\Theta \mid X, W)p(W)\mathrm{d}W},$$

- for data uncertainty, define a Gaussian likelihood $\mathcal{N}(f^W(x), \sigma^2)$ on 3DMM parameters,
- combine model and data uncertainties for $Q_I(\cdot|x)$ as in [2].

Our approach:

- adopt the general informed sampling approach,
- estimate global distribution $Q_I(\cdot|x)$ with a BNN.

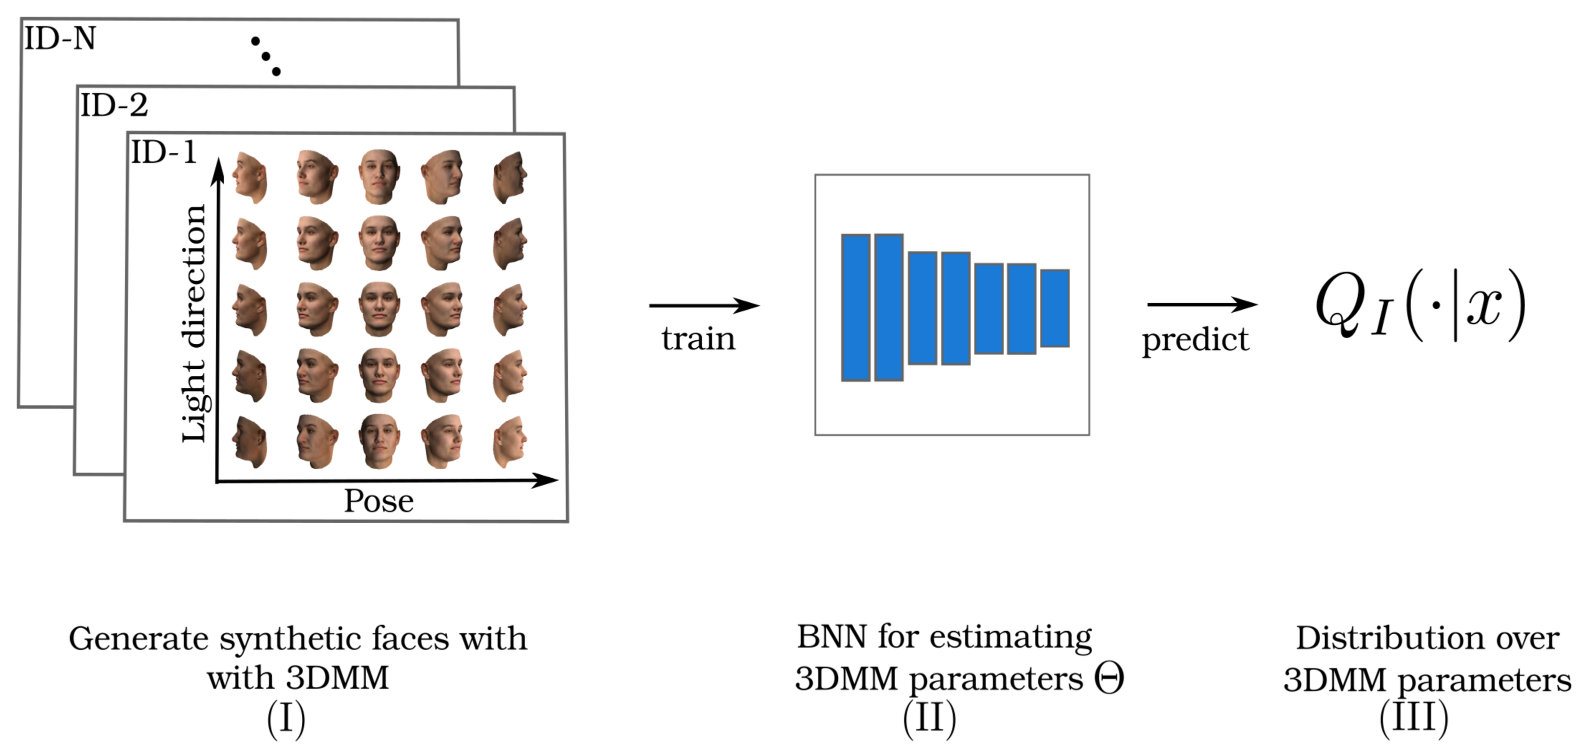

Contributions

- We estimate a **global, image-dependent proposal** distribution.
- BNN-Informed MCMC **significantly improves exploration** of maximal posterior regions.

Experiments

Dataset: We use a sample of 80 face images from the CMU-Multipie face dataset, sampled from Session-01 using the frontal cameras.

Experiment 1: Probabilistic Estimation of 3DMM Parameters

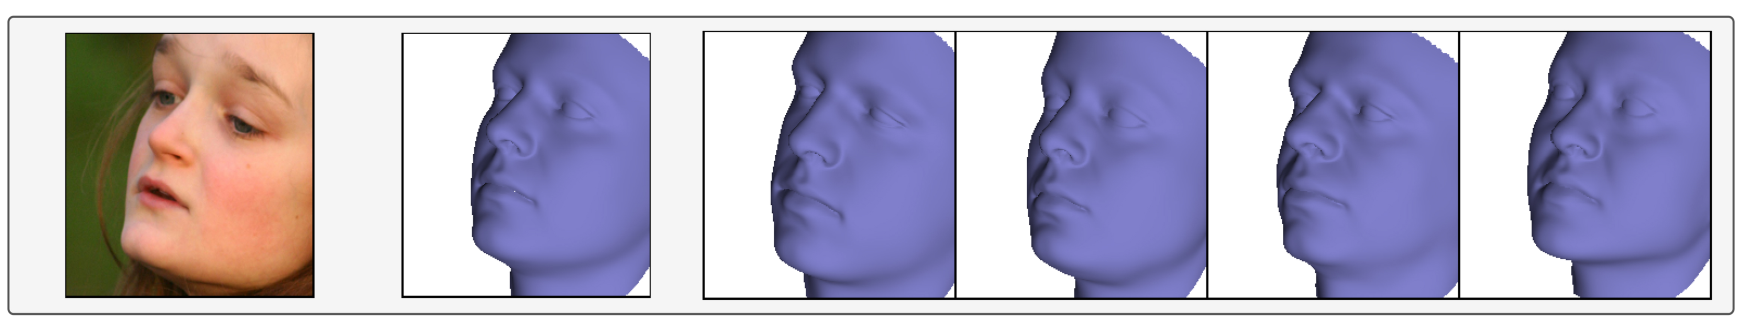

Experiment 2: Posterior Estimation with BNN-informed MCMC.

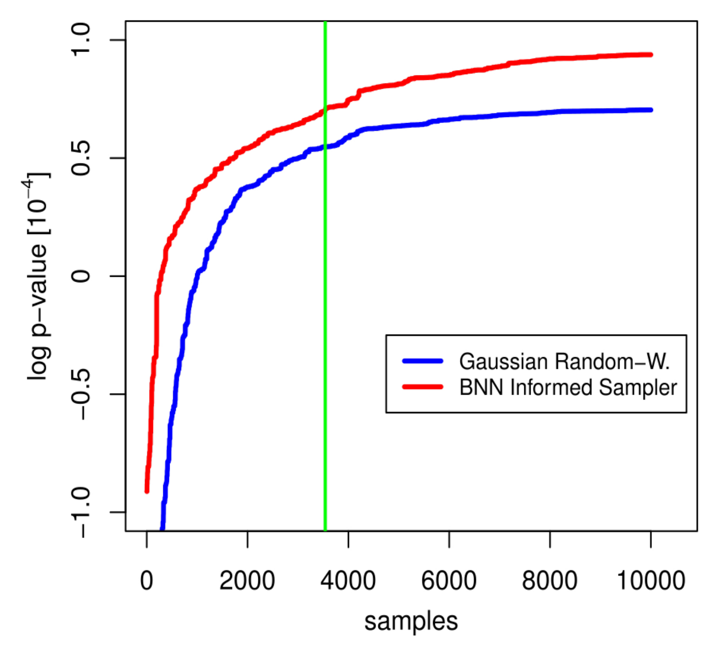

Ours

Uninformed

References

[1] Blanz, Volker and Vetter, Thomas; A morphable model for the synthesis of 3D faces; 1999
[2] Kendall, Alex and Gal, Yarin; What uncertainties do we need in bayesian deep learning for computer vision?; 2017
[3] Jampani, Varun, el al.; The informed sampler: A discriminative approach to bayesian inference in generative computer vision models; 2015

Acknowledgement: A.K. is supported by a Novartis University of Basel Excellence Scholarship for Life Sciences. M.W., A.W. and S.P are partially supported by the Swiss National Science Foundation (SNF), SystemsX.ch and the National Center of Competence in Research MARVEL. We gratefully acknowledge the support of NVIDIA with the donation of a Titan Xp.