Bayesian Helmholtz Stereopsis with Integrability Prior (supplementary material)

Nadejda Roubtsova and Jean-Yves Guillemaut

SECTION A. CORRELATION-BASED DNPRIOR AND INTEGRABILITY

The normalised correlation angle ϕ_{ph-g} is:

$$\phi_{ph-g} = \pi^{-1}\left|arcsin(\mathbf{n}_{prj,x} \cdot \mathbf{t}_x)\right|,$$

where $\mathbf{t}_x = ((\delta x)^2 + (z_2 - z_1)^2)^{-\frac{1}{2}}[\delta x, 0, (z_2 - z_1)]^\top$ and $\mathbf{n}_{prj,x} = (n_x^2 + n_z^2)^{-\frac{1}{2}}[n_x, 0, n_z]^\top$. The absolute value operator is introduced for consistency with the original cosine-based formulation of the correlation angle in Equation (8) of the paper in the ϕ_{ph-g} range of $[0, \frac{\pi}{2}]$. Normals outside of the range point into the surface and can be eliminated as inconsistent prior to optimisation. Hence, we can write:

$$\begin{aligned}
(\mathbf{n}_{prj,1,x} \cdot \mathbf{t}_x)^2 &= \frac{(n_x\delta x + n_z(z_2 - z_1))^2}{(n_x^2 + n_z^2)((\delta x)^2 + (z_2 - z_1)^2)} \\
&= \frac{\left(\frac{(z_2 - z_1)}{\delta x} + \frac{n_x}{n_z}\right)^2 n_z^2 \delta x^2}{(n_x^2 + n_z^2)((\delta x)^2 + (z_2 - z_1)^2)} \\
&= \frac{\delta x^2}{\delta x^2 + (z_2 - z_1)^2}\left(\frac{(z_2 - z_1)}{\delta x} + \frac{n_x}{n_z}\right)^2 \frac{n_z^2}{(n_x^2 + n_z^2)} \\
&= \frac{1}{1 + \frac{(z_2 - z_1)^2}{\delta x^2}}\left(\frac{(z_2 - z_1)}{\delta x} - g_{x,1}\right)^2 \frac{1}{1 + g_{x,1}^2} \\
&= \frac{(\frac{z_2 - z_1}{\delta x} - g_{x,1})^2}{(\frac{z_2 - z_1}{\delta x} - g_{x,1})^2 + (\frac{z_2 - z_1}{\delta x} g_{x,1} + 1)^2} \\
&= \frac{e_{1,dn}^{Horn}}{e_{1,dn}^{Horn} + (\frac{z_2 - z_1}{\delta x} g_{x,1} + 1)^2}.
\end{aligned}$$

The resulting relationship of the correlation-based DNprior E_{dn}^{corr} to E_{dn}^{Horn} is complex:

$$E_{dn}^{corr} = \frac{1}{2\pi^2}\left|arcsin\left(\sqrt{\frac{e_{1,dn}^{Horn}}{e_{1,dn}^{Horn} + (\frac{z_2 - z_1}{\delta x} g_{x,1} + 1)^2}}\right)\right|^2 + \frac{1}{2\pi^2}\left|arcsin\left(\sqrt{\frac{e_{2,dn}^{Horn}}{e_{2,dn}^{Horn} + (\frac{z_1 - z_2}{\delta x} g_{x,2} + 1)^2}}\right)\right|^2.$$

SECTION B. MRF PARAMETER α OPTIMISATION

The graphs show the Middlebury depth accuracy at threshold of 90% as a function of relative data/smoothness weighting parameter α for different priors and datasets. Middlebury accuracy of $x\,mm$ at threshold $th = N\%$ indicates that $N\%$ of reconstruction vertices lie within $x\,mm$ of the ground truth:
Seitz, S., Curless, B., Diebel, J, Scharstein, D., Szeliski, R. "A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms", CVPR, pp.519-528, 2006.
The α parameter optimisation sweep is performed on data significantly corrupted by Gaussian noise (normalised noise variance of 0.001 i.e. ± 2072 intensity levels).

SPHERE

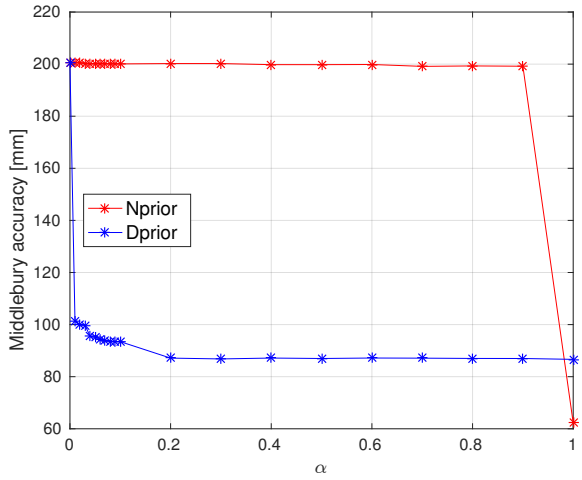

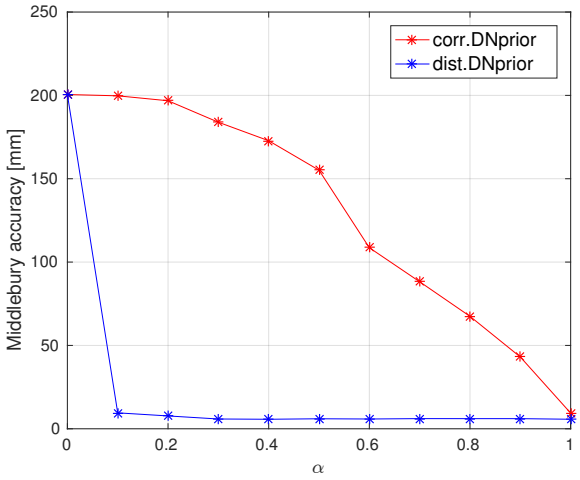

PEAR

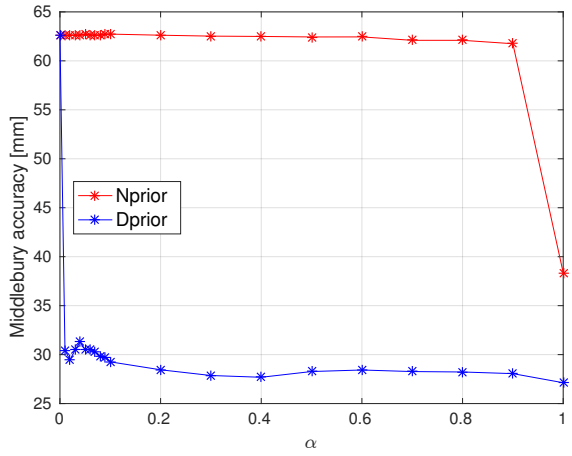

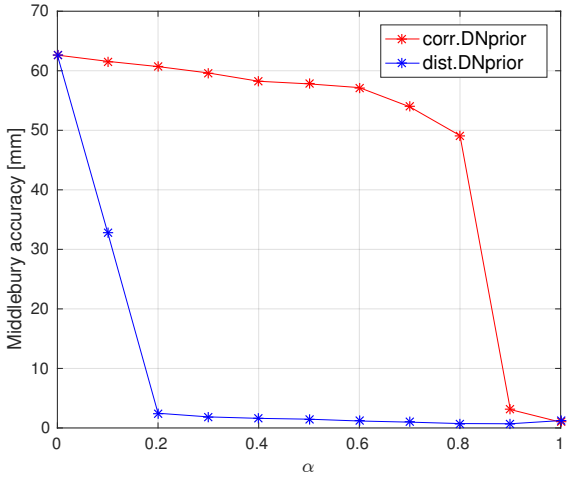

BUNNY

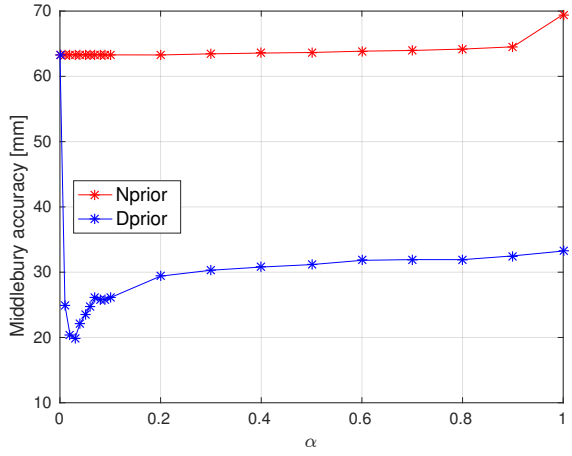

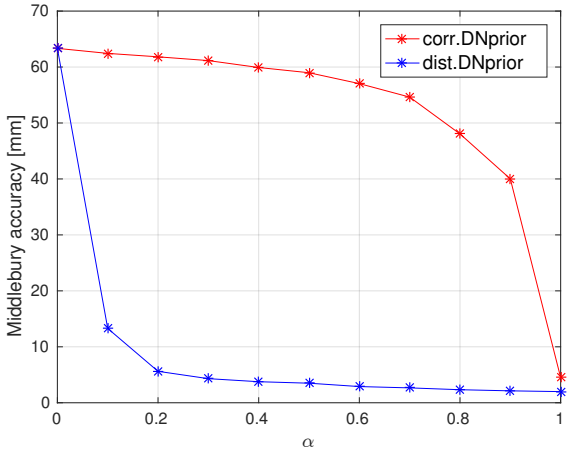

Example intensity images for the synthetic objects from the main paper are presented below to give an indication of the geometry, reflectance and the degree of corruption by Gaussian noise of these datasets.

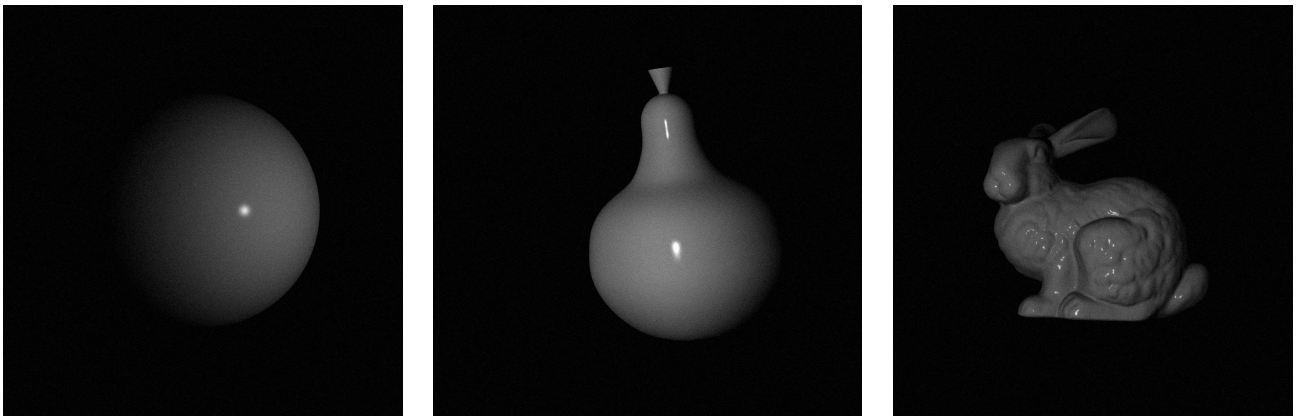

sphere pear bunny

Further, we show the following mesh reconstructions (visualised with flat shading) for the noise corrupted synthetic data corresponding to the depth/normal error analysis presented in Figure 3 and Table 1 of the main paper (parameters per prior/dataset as in the paper).

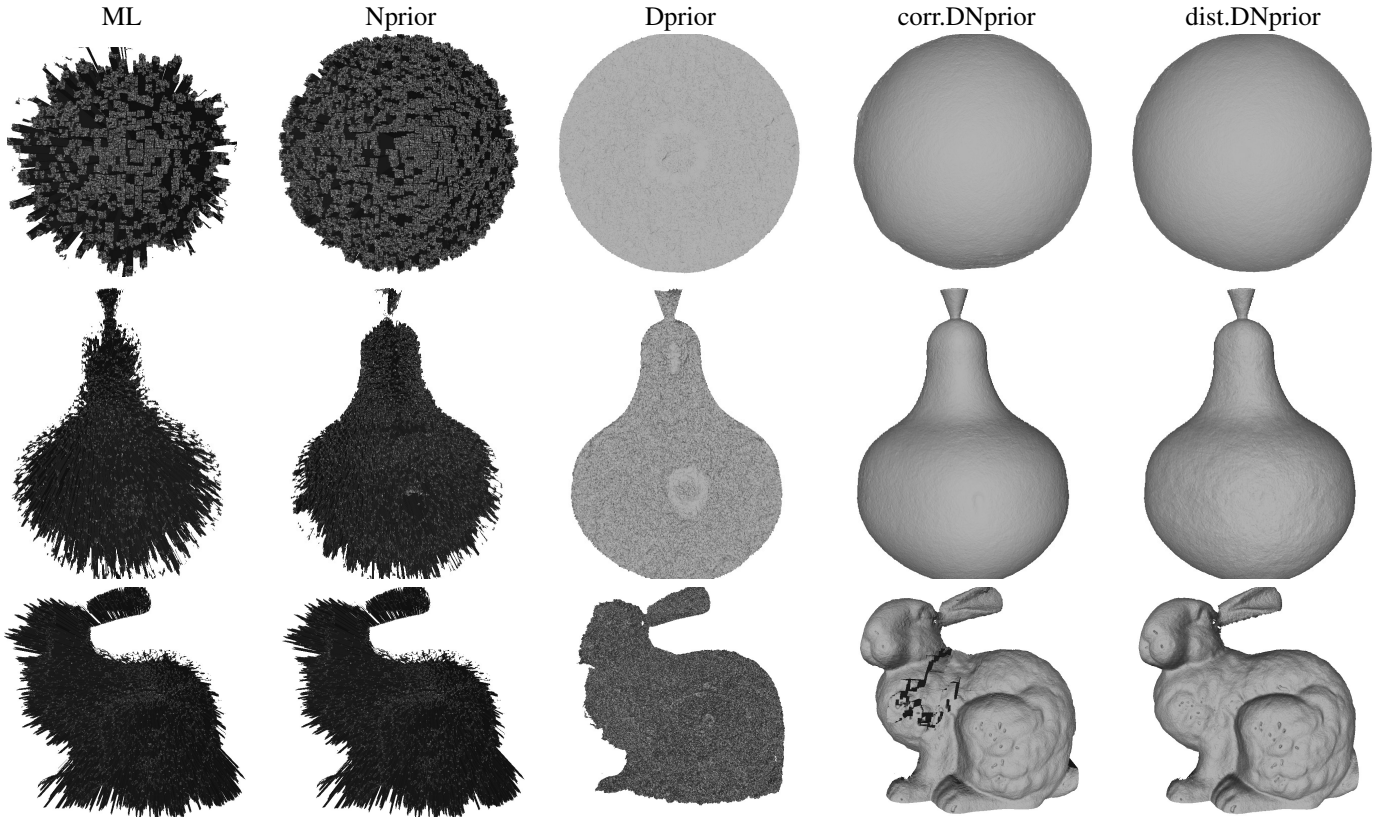

SECTION C. SURFACE RECONSTRUCTION RESULTS

This section compares the performance of standard HS (ML) and the proposed Bayesian HS with different priors (Nprior, Dprior, corr.DNprior and dist.DNprior) on real datasets shown below, specifically the four datasets from the main paper (teapot1, teapot2, doll and vase) and four additional datasets (billiard, cup, mannequin and teddy). The datasets are versatile featuring objects that are smooth and (largely) untextured (teapots, billiard), heavily textured (cup, vase), of intricate geometry (doll), in possession of fine structure (teddy) and uncommon reflectance properties by virtue of complex material (polystyrene mannequin).

teapot1

teapot2

doll

vase

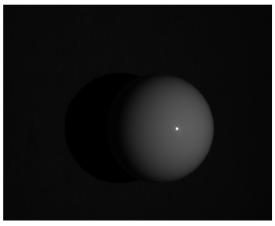

billiard

cup

mannequin

teddy

In addition, the proposed integration-free approach to final surface assembly is compared against Poisson Surface Reconstruction (M. Kazhdan, M.Bolitho and H.Hoppe, "Poisson surface reconstruction" in SGP, 2006, pp. 61-70) and the method of Nehab et.al from: D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi, "Efficiently combining positions and normals for precise 3D geometry" pp. 536-543, ACM SIGGRAPH, 2005. As the differences in performance between the method of Nehab et.al and the proposed approach without explicit integration are subtle some close-ups are presented at the end of the section.

The order of result presentation (vertically, top to bottom):
A. depth maps;
B. RGB normal maps
and the final surfaces per reconstruction method for different integration approaches, viewing angles and shading:
C. no explicit integration (proposed) with flat shading[1] (front)
D. no explicit integration (proposed) with smooth shading[2] (front);
E. no explicit integration (proposed) with smooth shading (side);
F. Poisson Surface Reconstruction with smooth shading (front);
G. Poisson Surface Reconstruction with smooth shading (side);
H. Nehab et.al with smooth shading (front)
I. Nehab et.al with smooth shading (side).

Parameter α per prior is as specified in Fig.5 of the main paper. Reconstruction volume sizes are specified below per dataset.

[1] i.e. rendered using geometric normals
[2] i.e. rendered using photometric normals

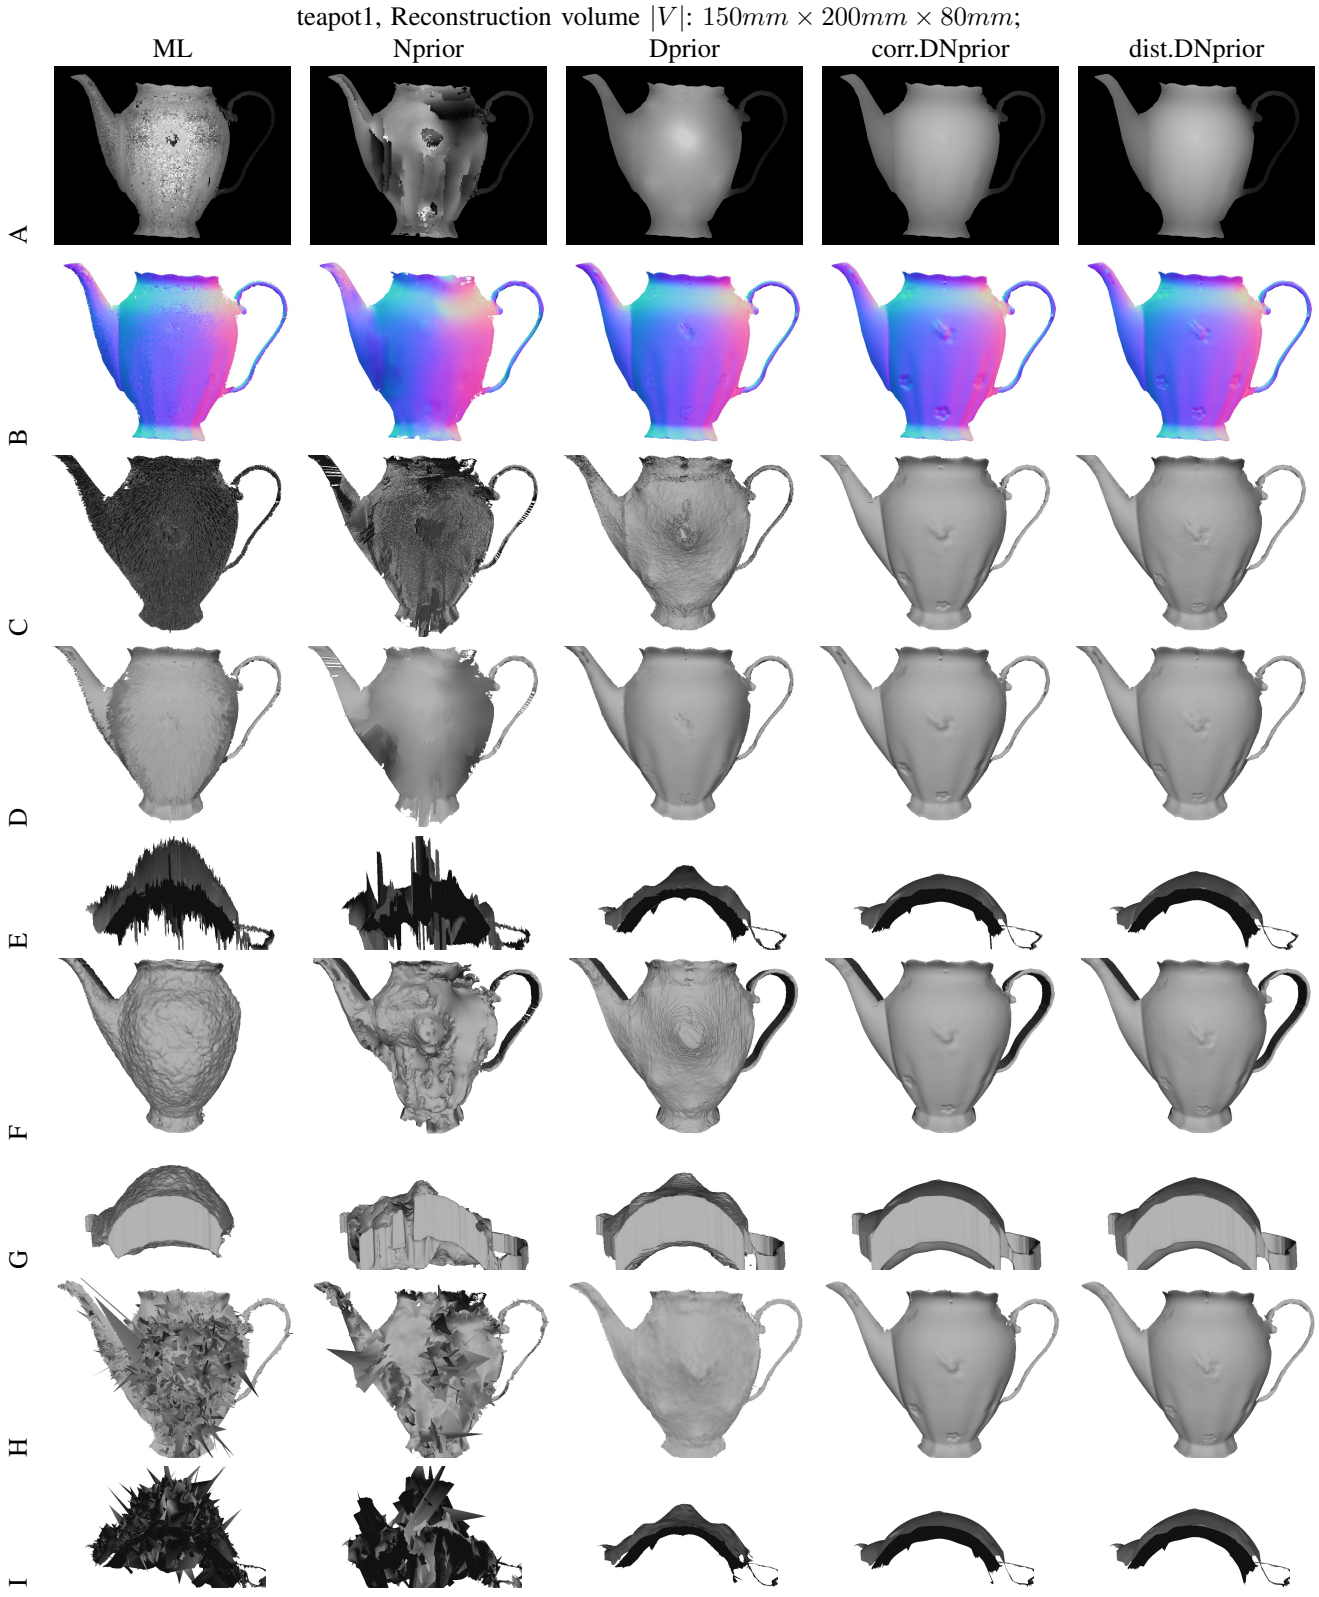
teapot1, Reconstruction volume $|V|$: $150mm \times 200mm \times 80mm$;
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
A
B
C
D
E
F
G
H
I

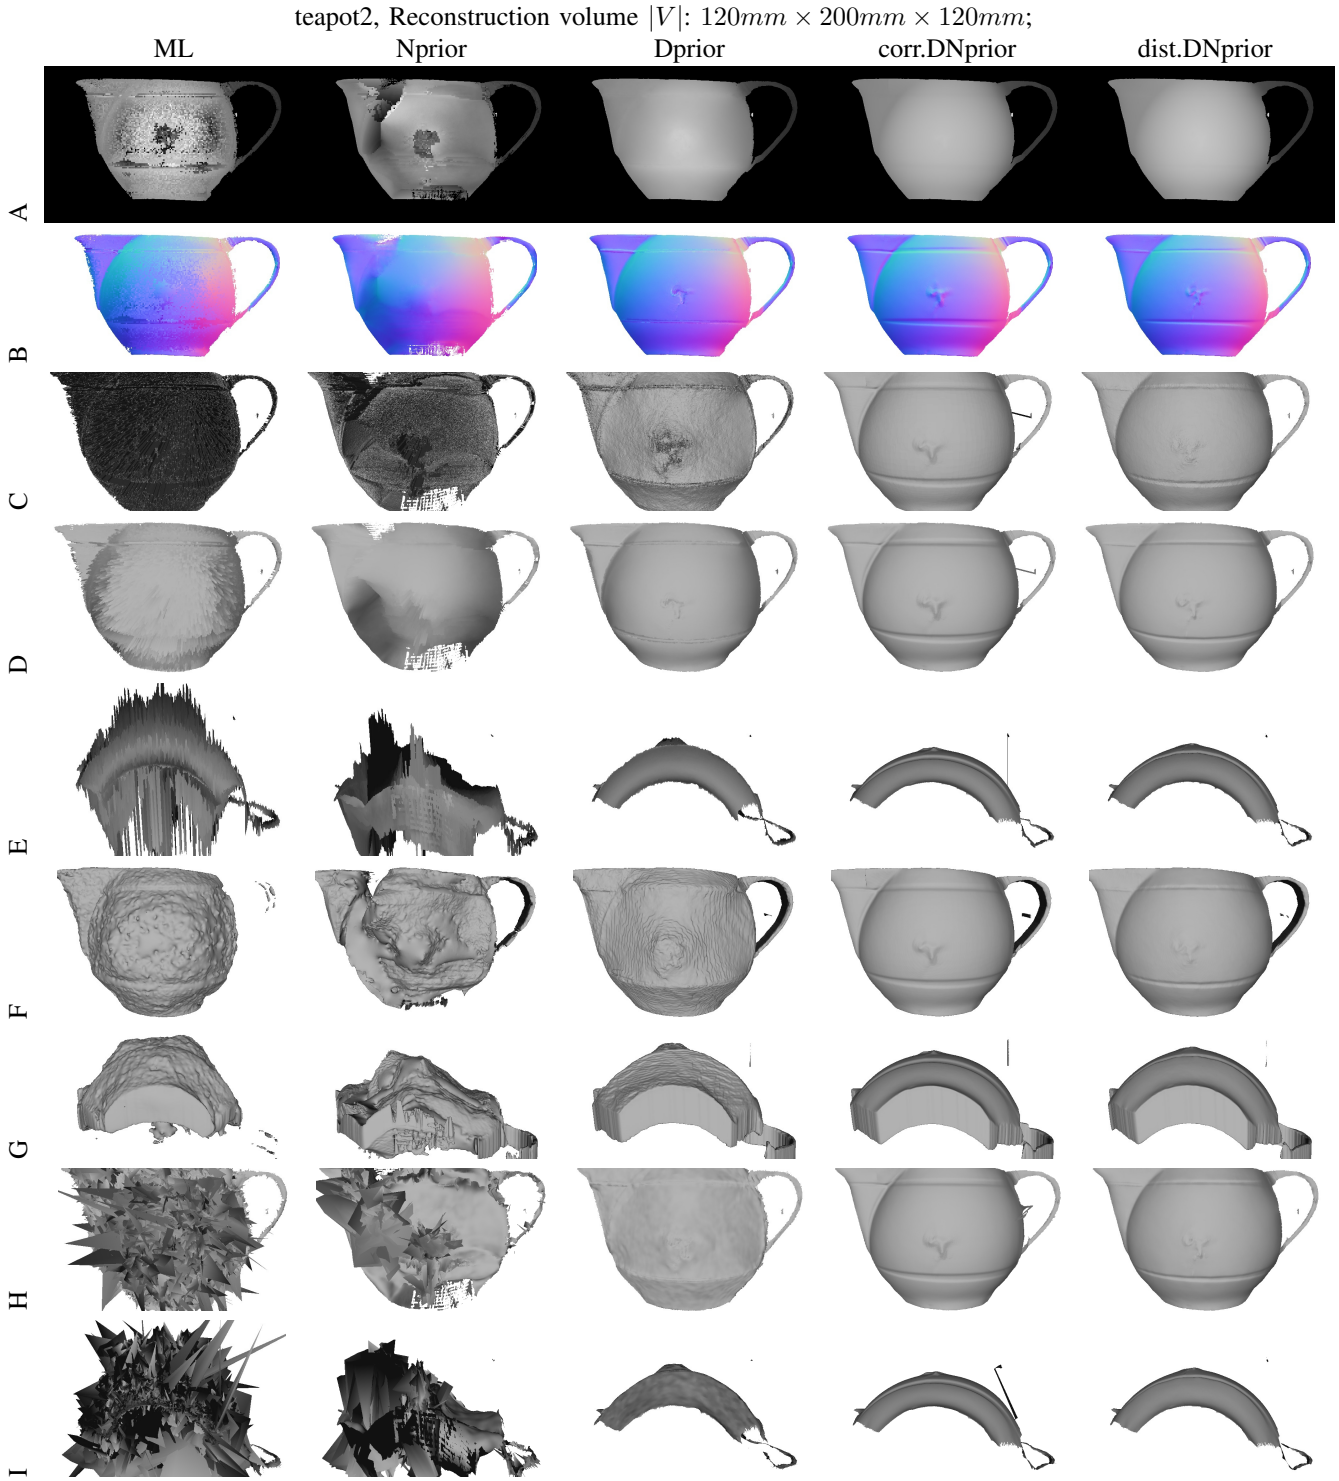
teapot2, Reconstruction volume $|V|$: $120mm \times 200mm \times 120mm$;
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
A
B
C
D
E
F
G
H
I

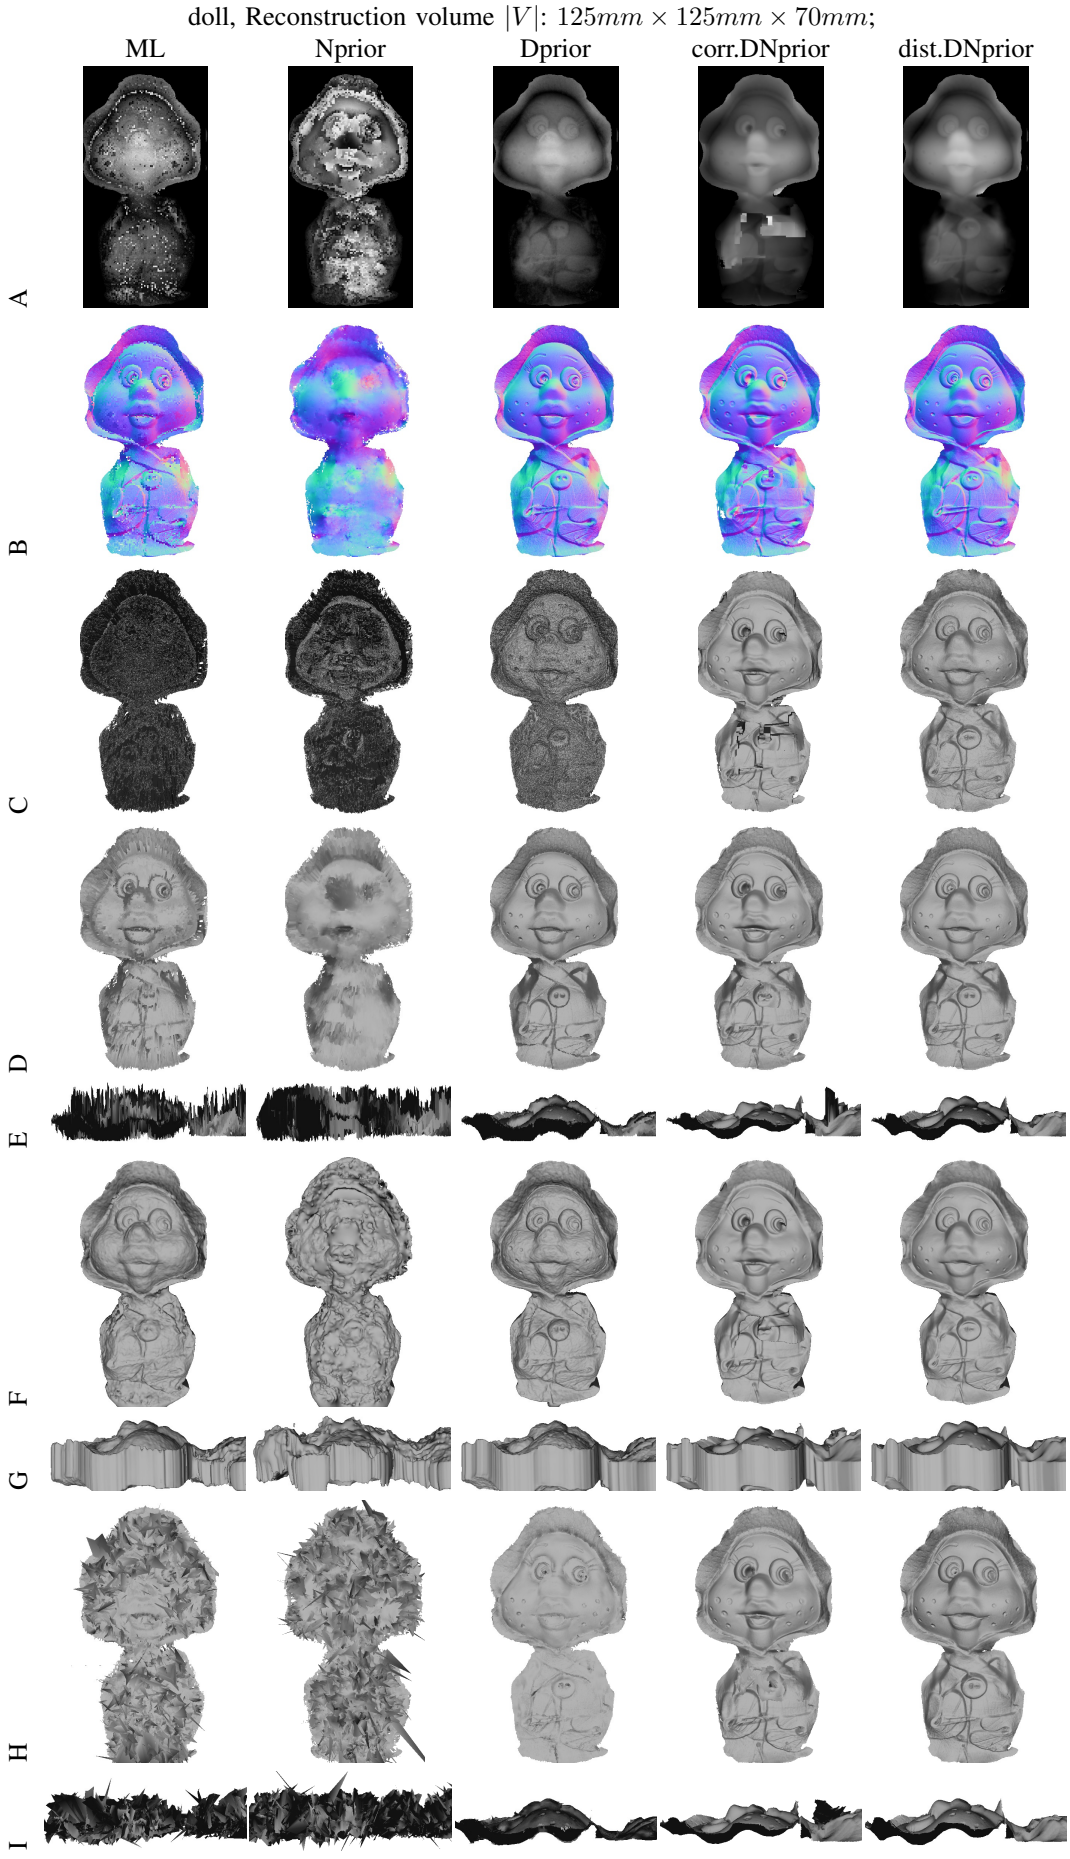
doll, Reconstruction volume $|V|$: $125mm \times 125mm \times 70mm$;
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
A
B
C
D
E
F
G
H
I

vase, Reconstruction volume $|V|$: $190mm \times 130mm \times 120mm$.

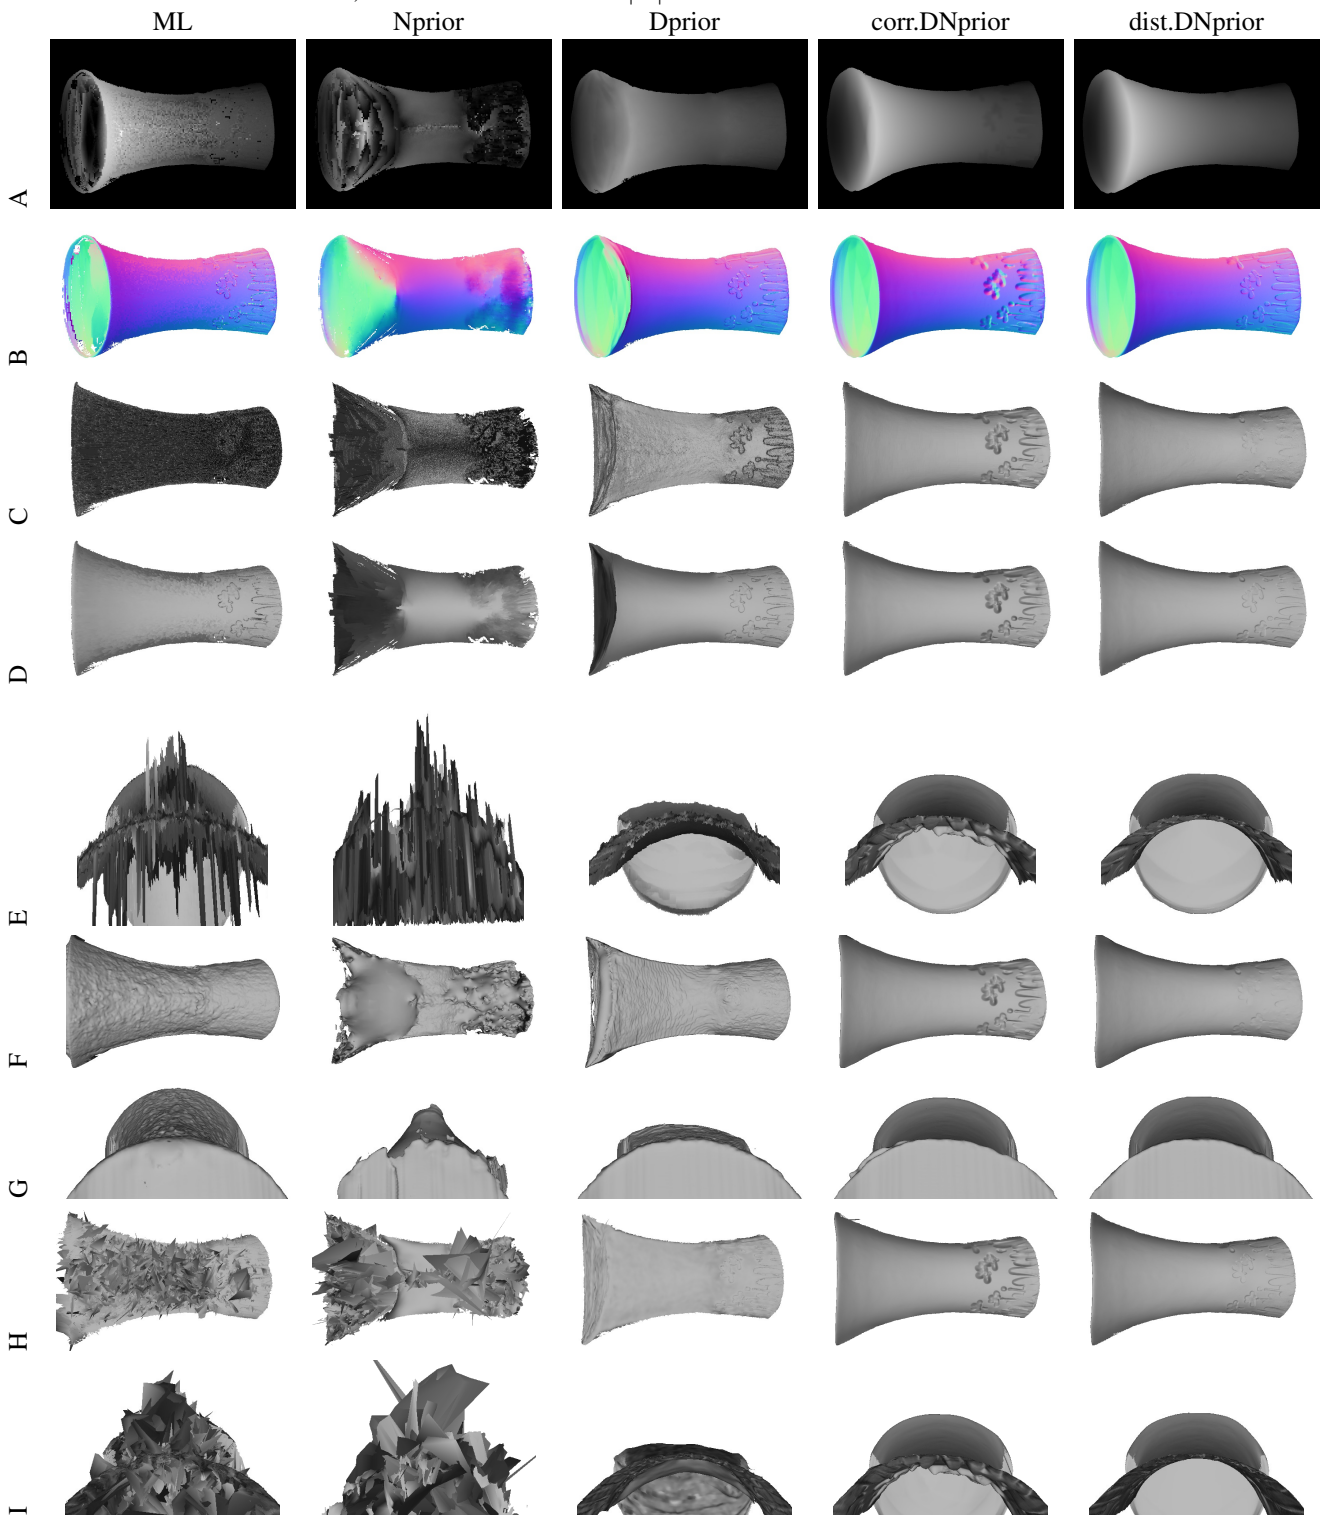

billiard, Reconstruction volume $|V|$: $60mm \times 60mm \times 35mm$.

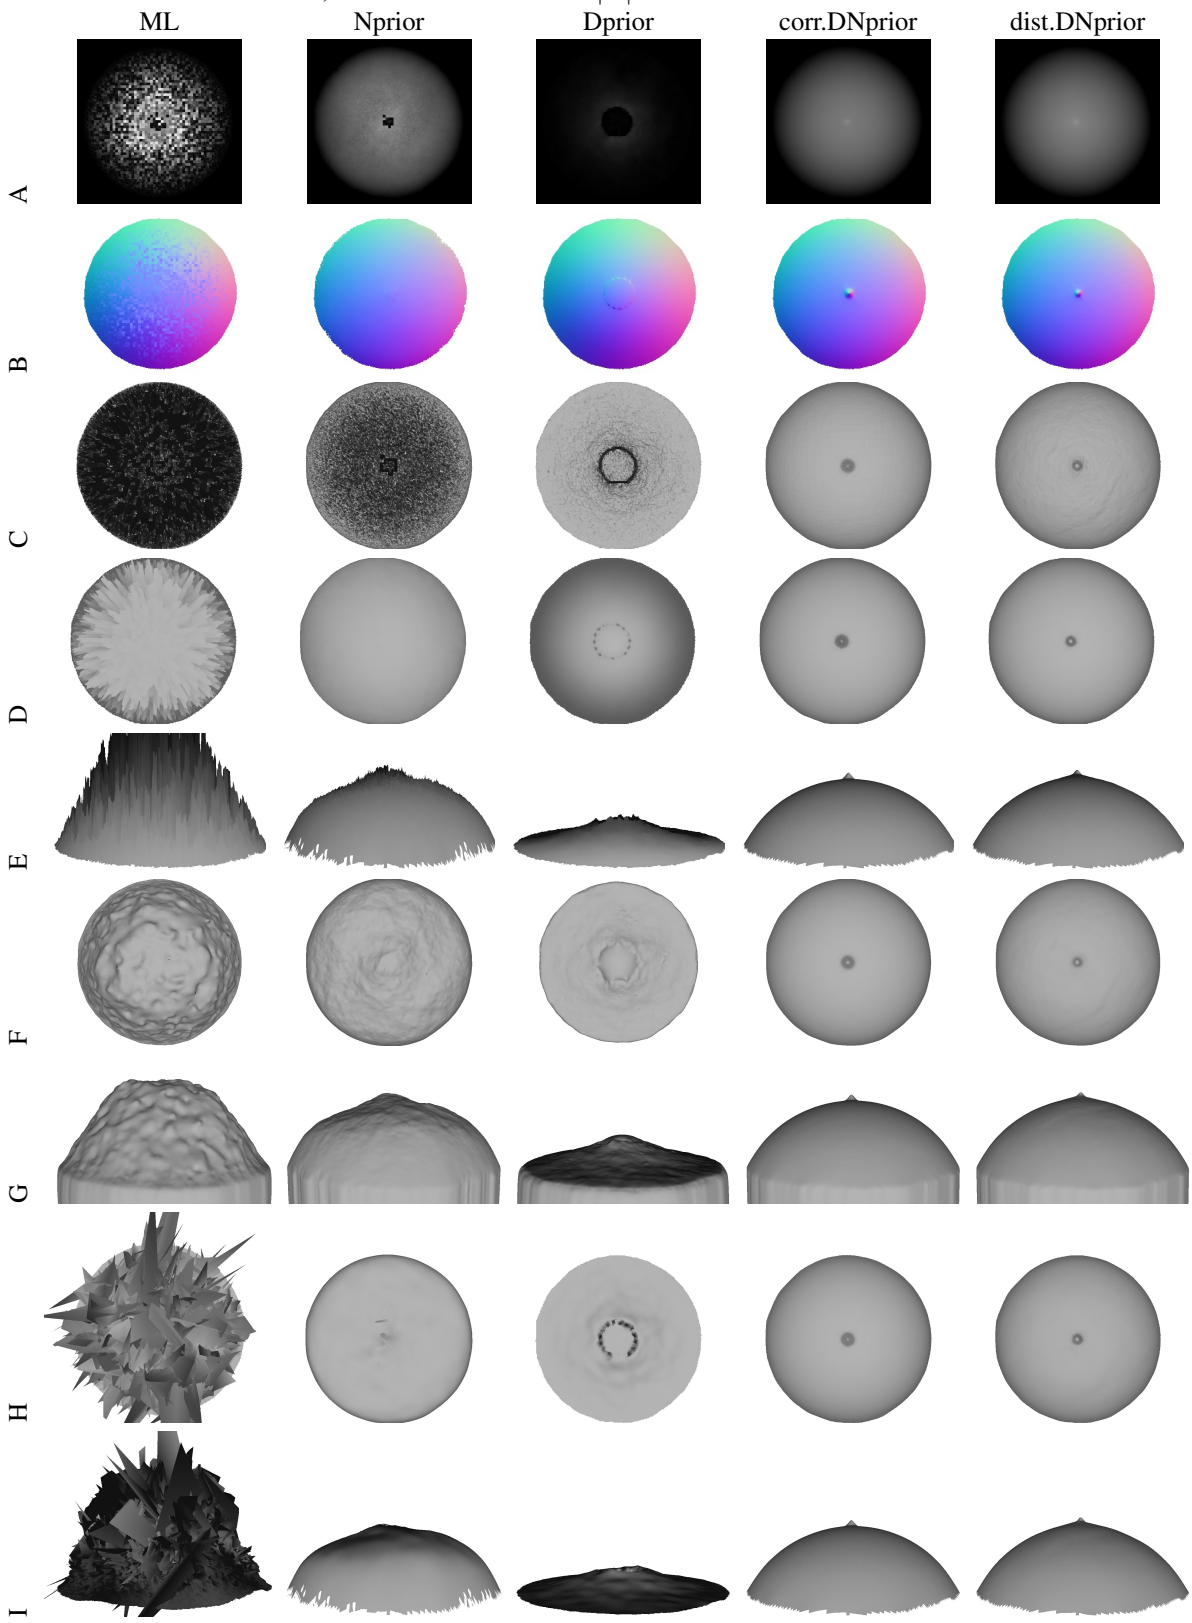

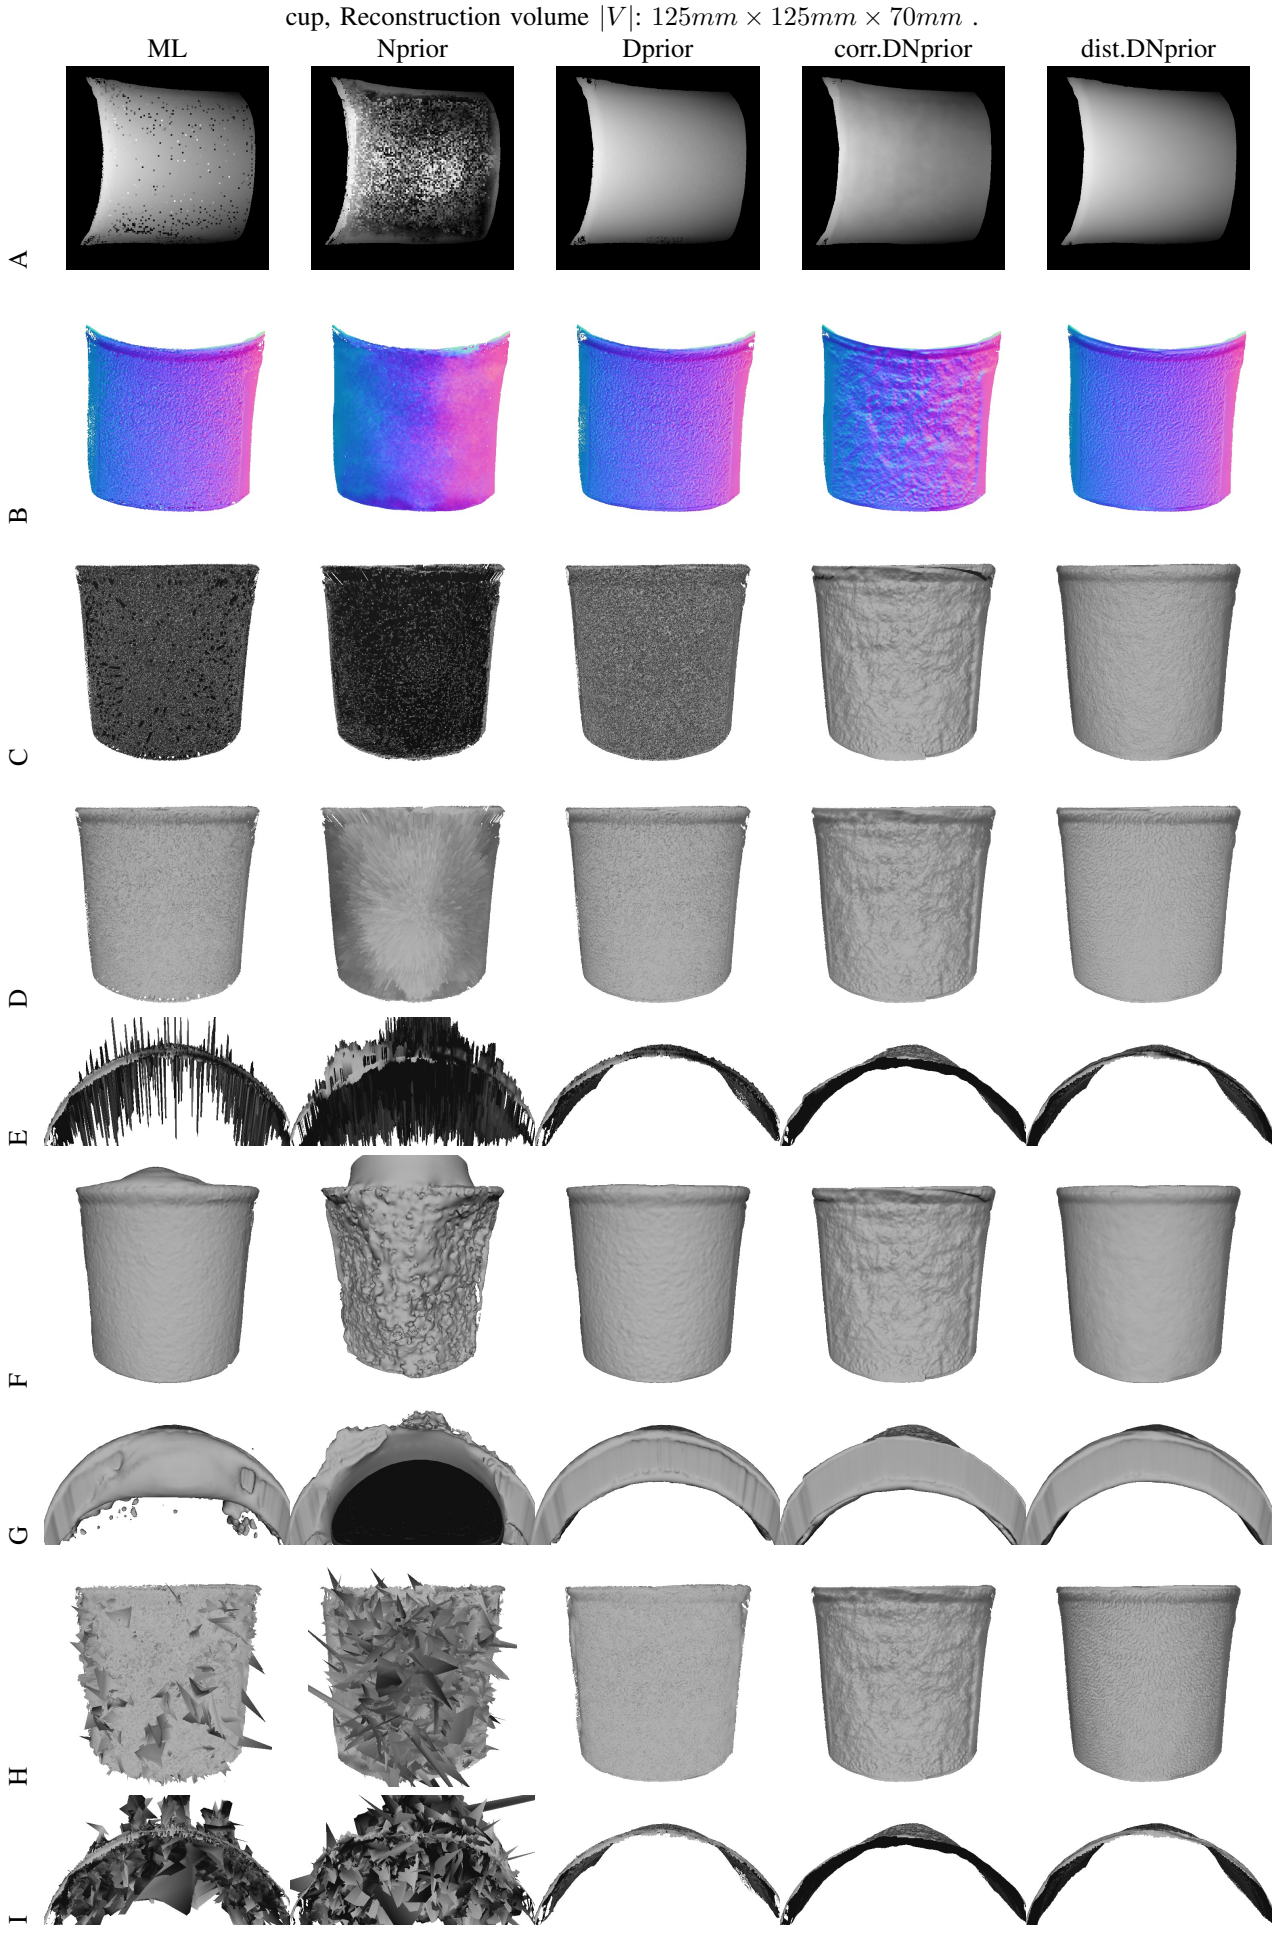
cup, Reconstruction volume $|V|$: $125mm \times 125mm \times 70mm$.
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
A
B
C
D
E
F
G
H
I

mannequin, Reconstruction volume $|V|$: $210mm \times 160mm \times 100mm$.

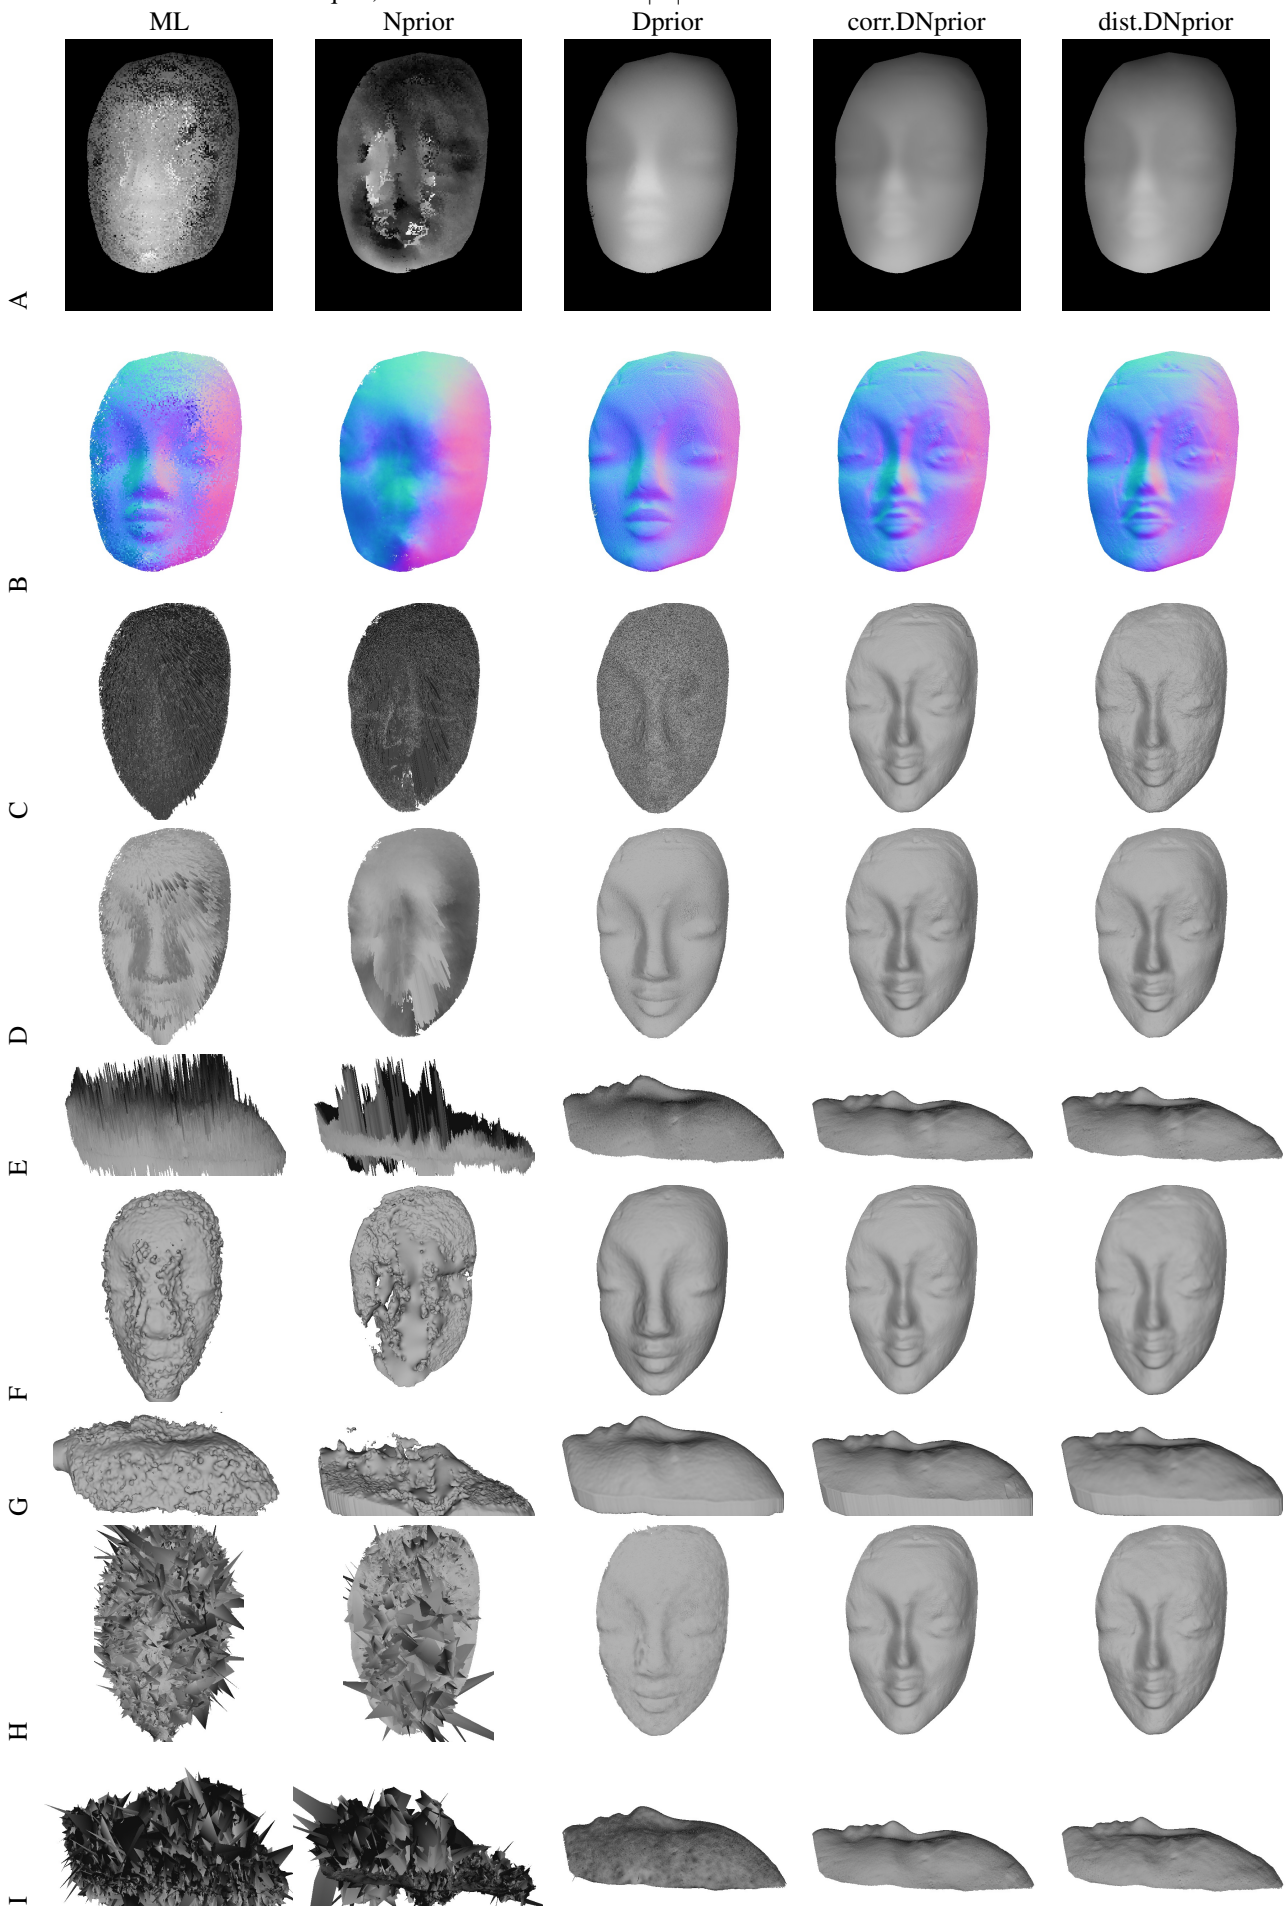

teddy, Reconstruction volume $|V|$: $170mm \times 180mm \times 80mm$.

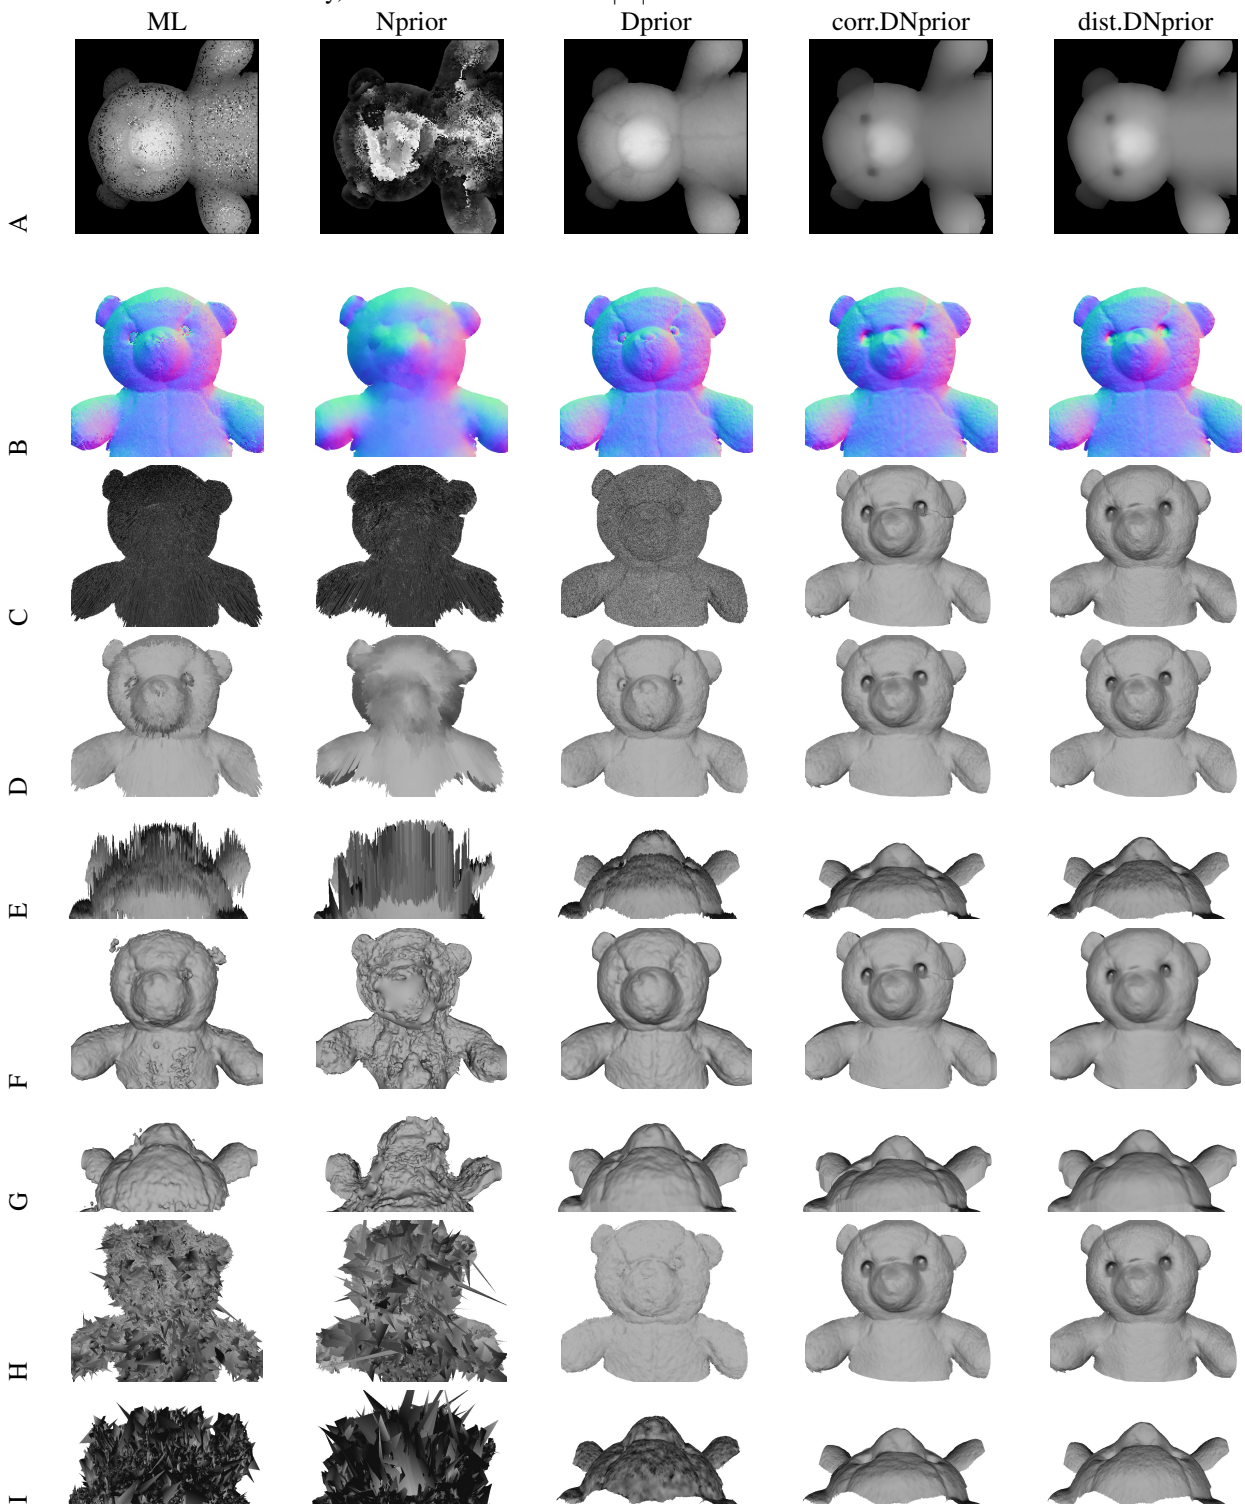

This is a closer look at the performance of the method of **Nehab et.al (left-hand side)** vs. **no explicit integration (right-hand side)** (proposed). Both methods perform very well. On careful inspection of the close-ups however one will notice that the proposed integration-free approach produces somewhat sharper features e.g. eyebrows, eye sockets, face-to-bonnet transition. This is manifested through deeper feature shading in the meshes and more pronounced normal field fluctuation across the surface (i.e. sharper colour variations around the features) in the rgb normal maps of the proposed method.

close up of reconstructed mesh: doll view nr. 1, DNprior, smooth shading

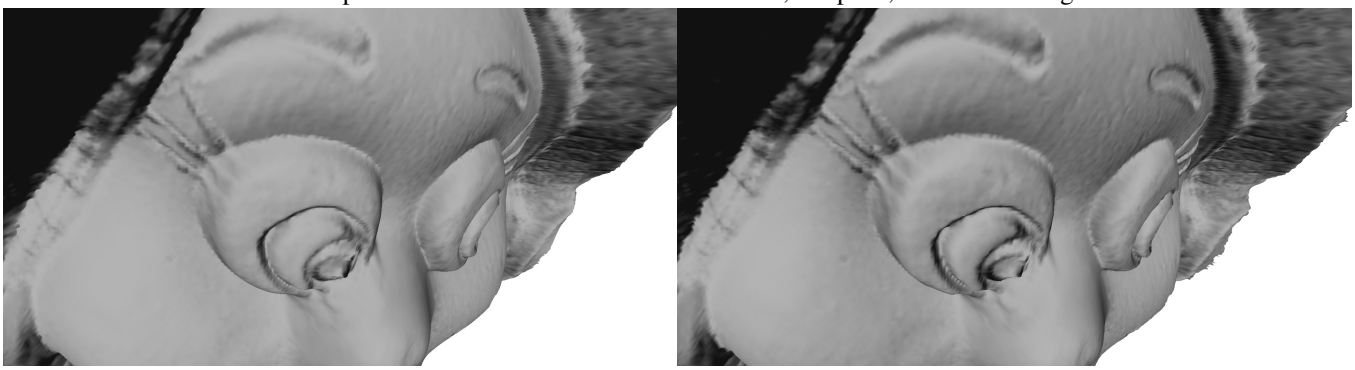

RGB normal map of the close-up

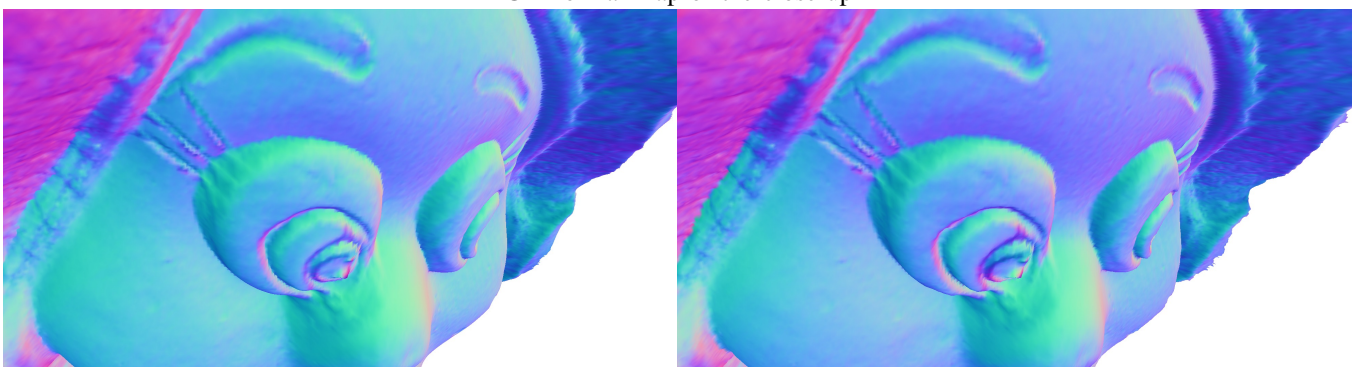

close up of reconstructed mesh: doll view nr. 2, DNprior, smooth shading

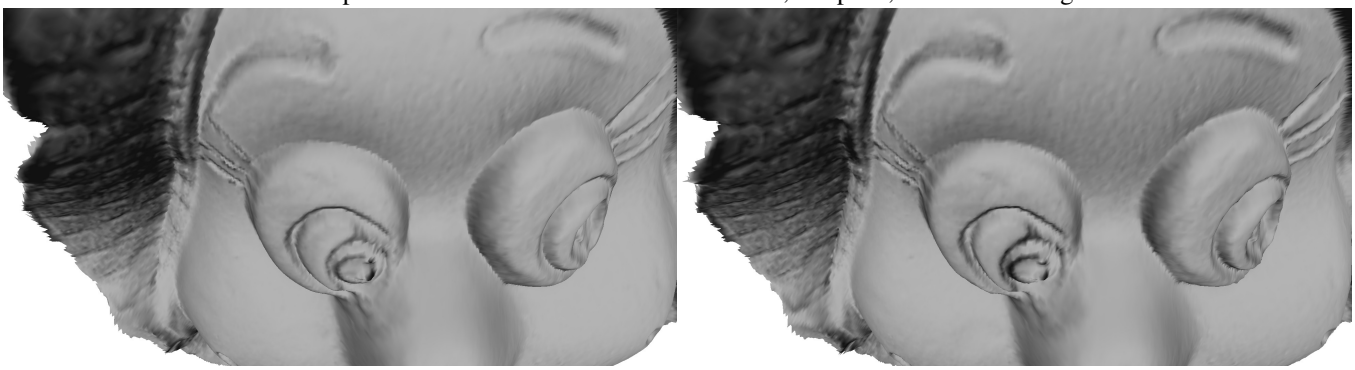

RGB normal map of the close-up

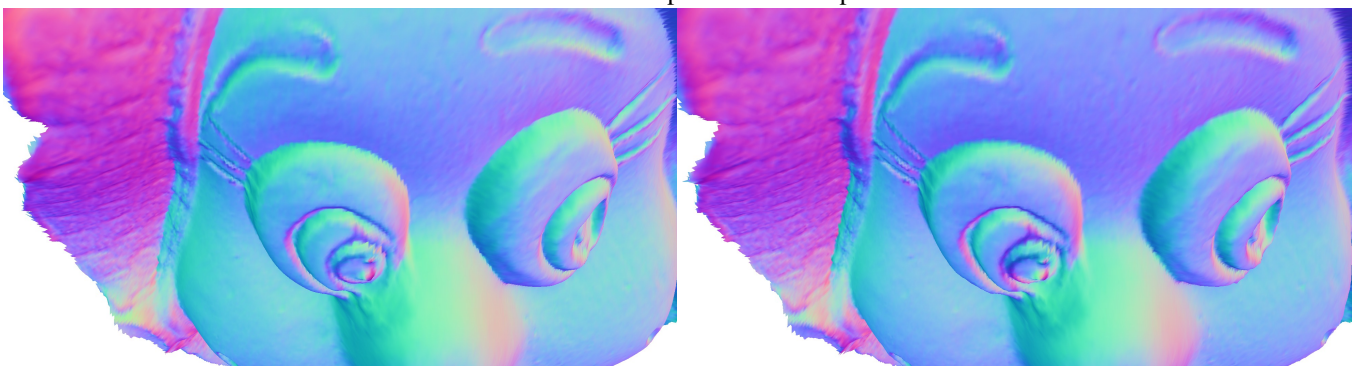

SECTION D. SURFACE INTEGRATION FROM SURFACE NORMALS

This section compares the performance of five state-of-the-art direct normal field integration methods for surface assembly from the reconstructed point cloud. The integrated point clouds are reconstructed by standard HS (ML) and the proposed Bayesian HS with different priors (Nprior, Dprior, corr.DNprior and dist.DNprior).

The integration methods (vertically, top to bottom) are:
diffusion tensor (Diffusion);
energy minimisation (EM);
Frankot-Chellappa (FC);
least squares (LS);
M-estimator (M-est);
For further details of the integration methods and their implementation the reader is referred to:
A. Agrawal, R. Raskar, and R. Chellappa, “Surface reconstructions from a gradient field?”, pp.578-591, ECCV, 2006.

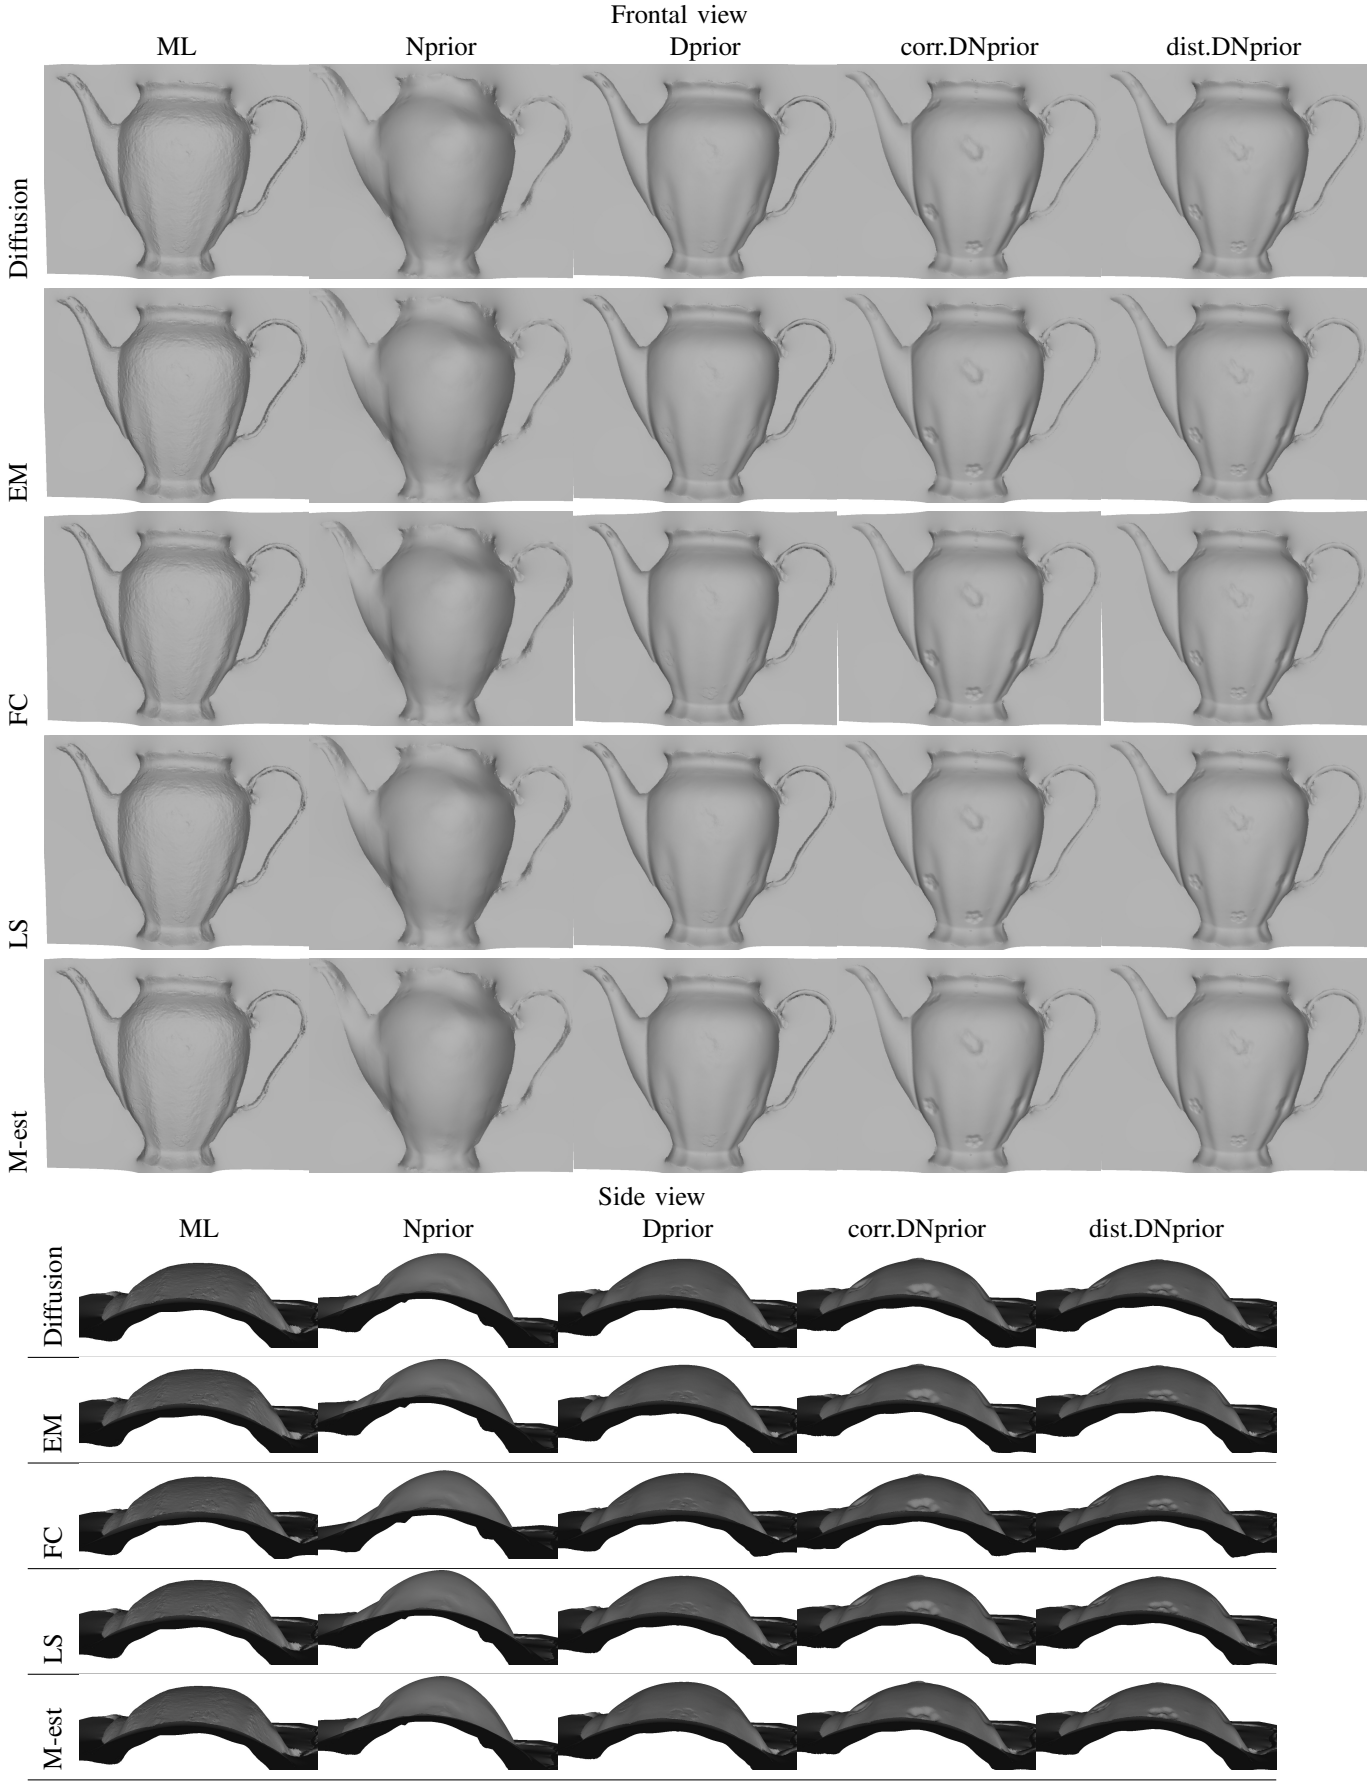
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

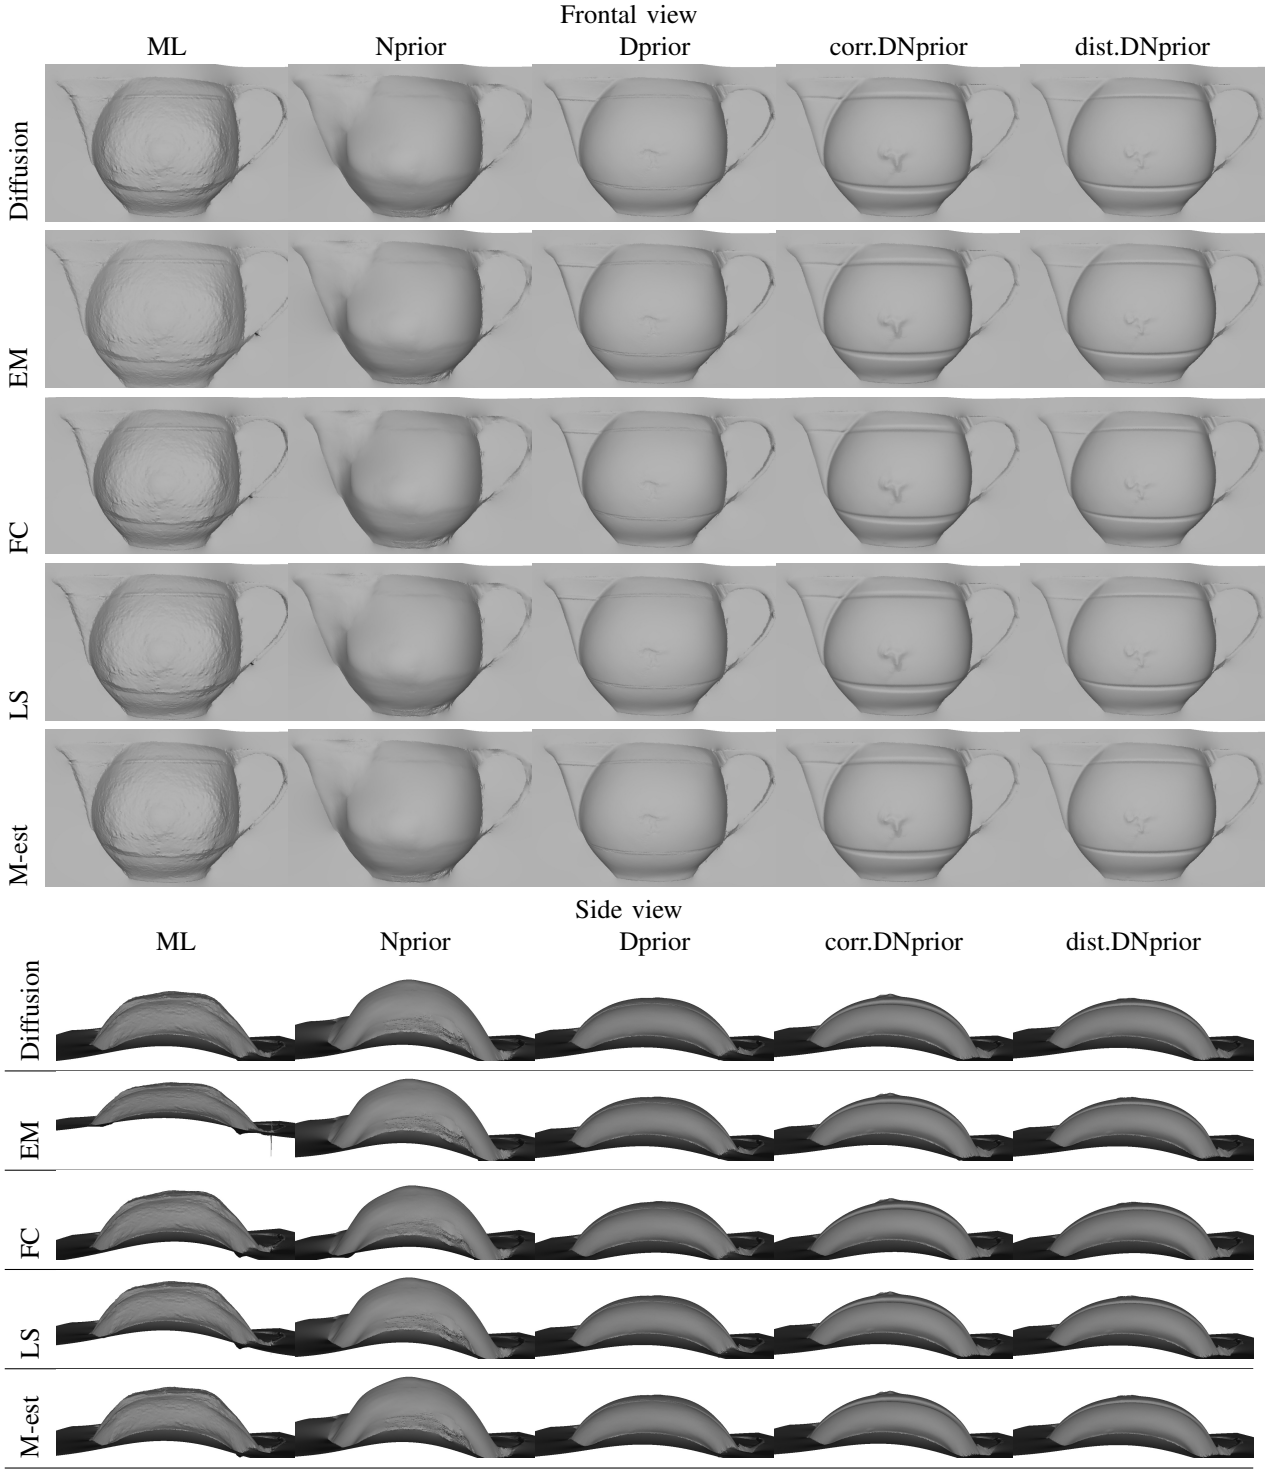
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

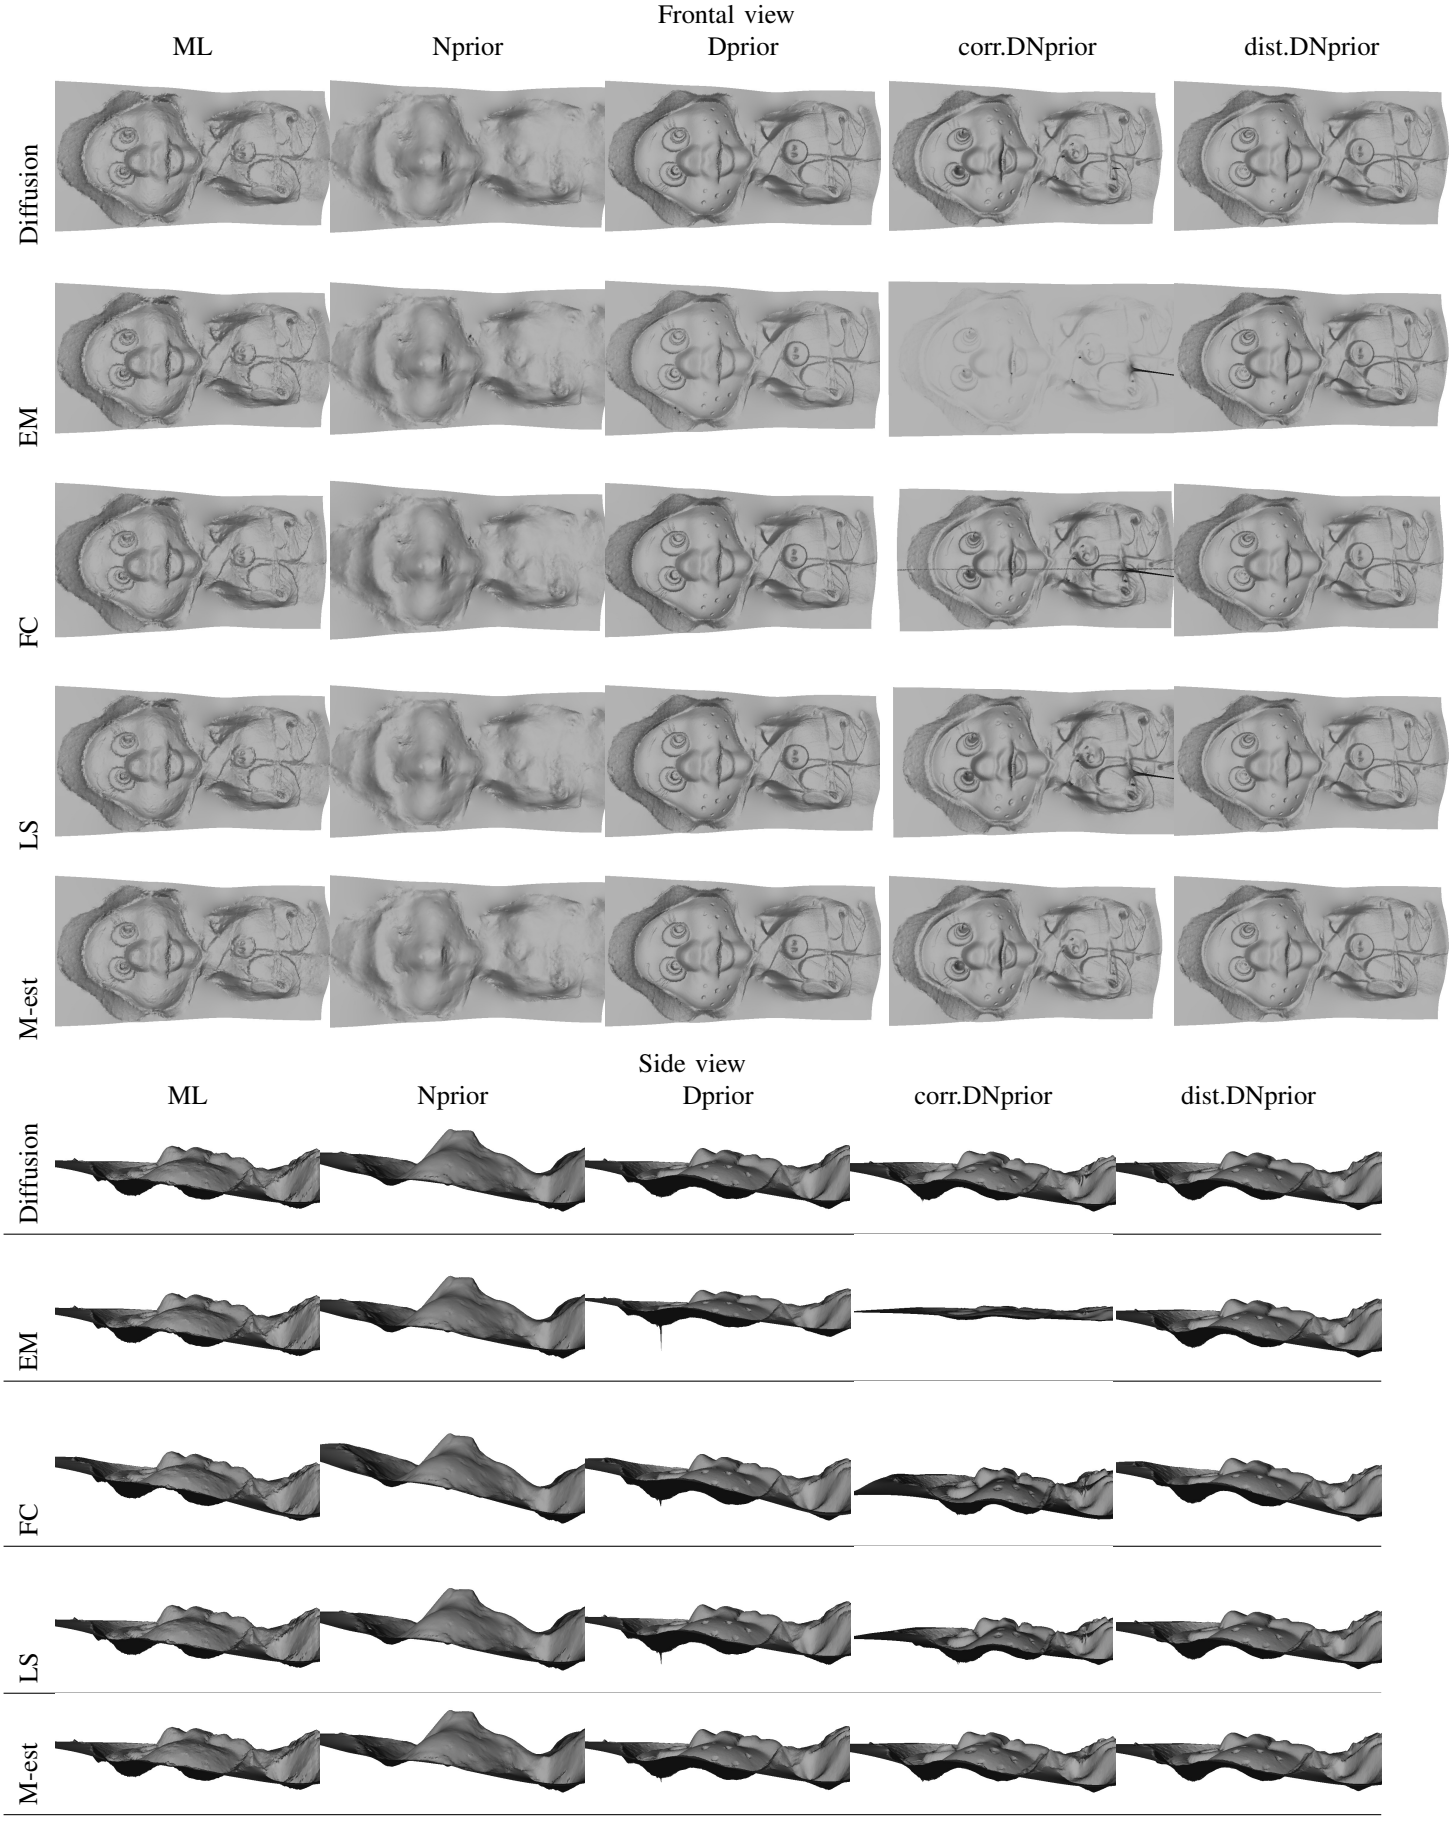
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

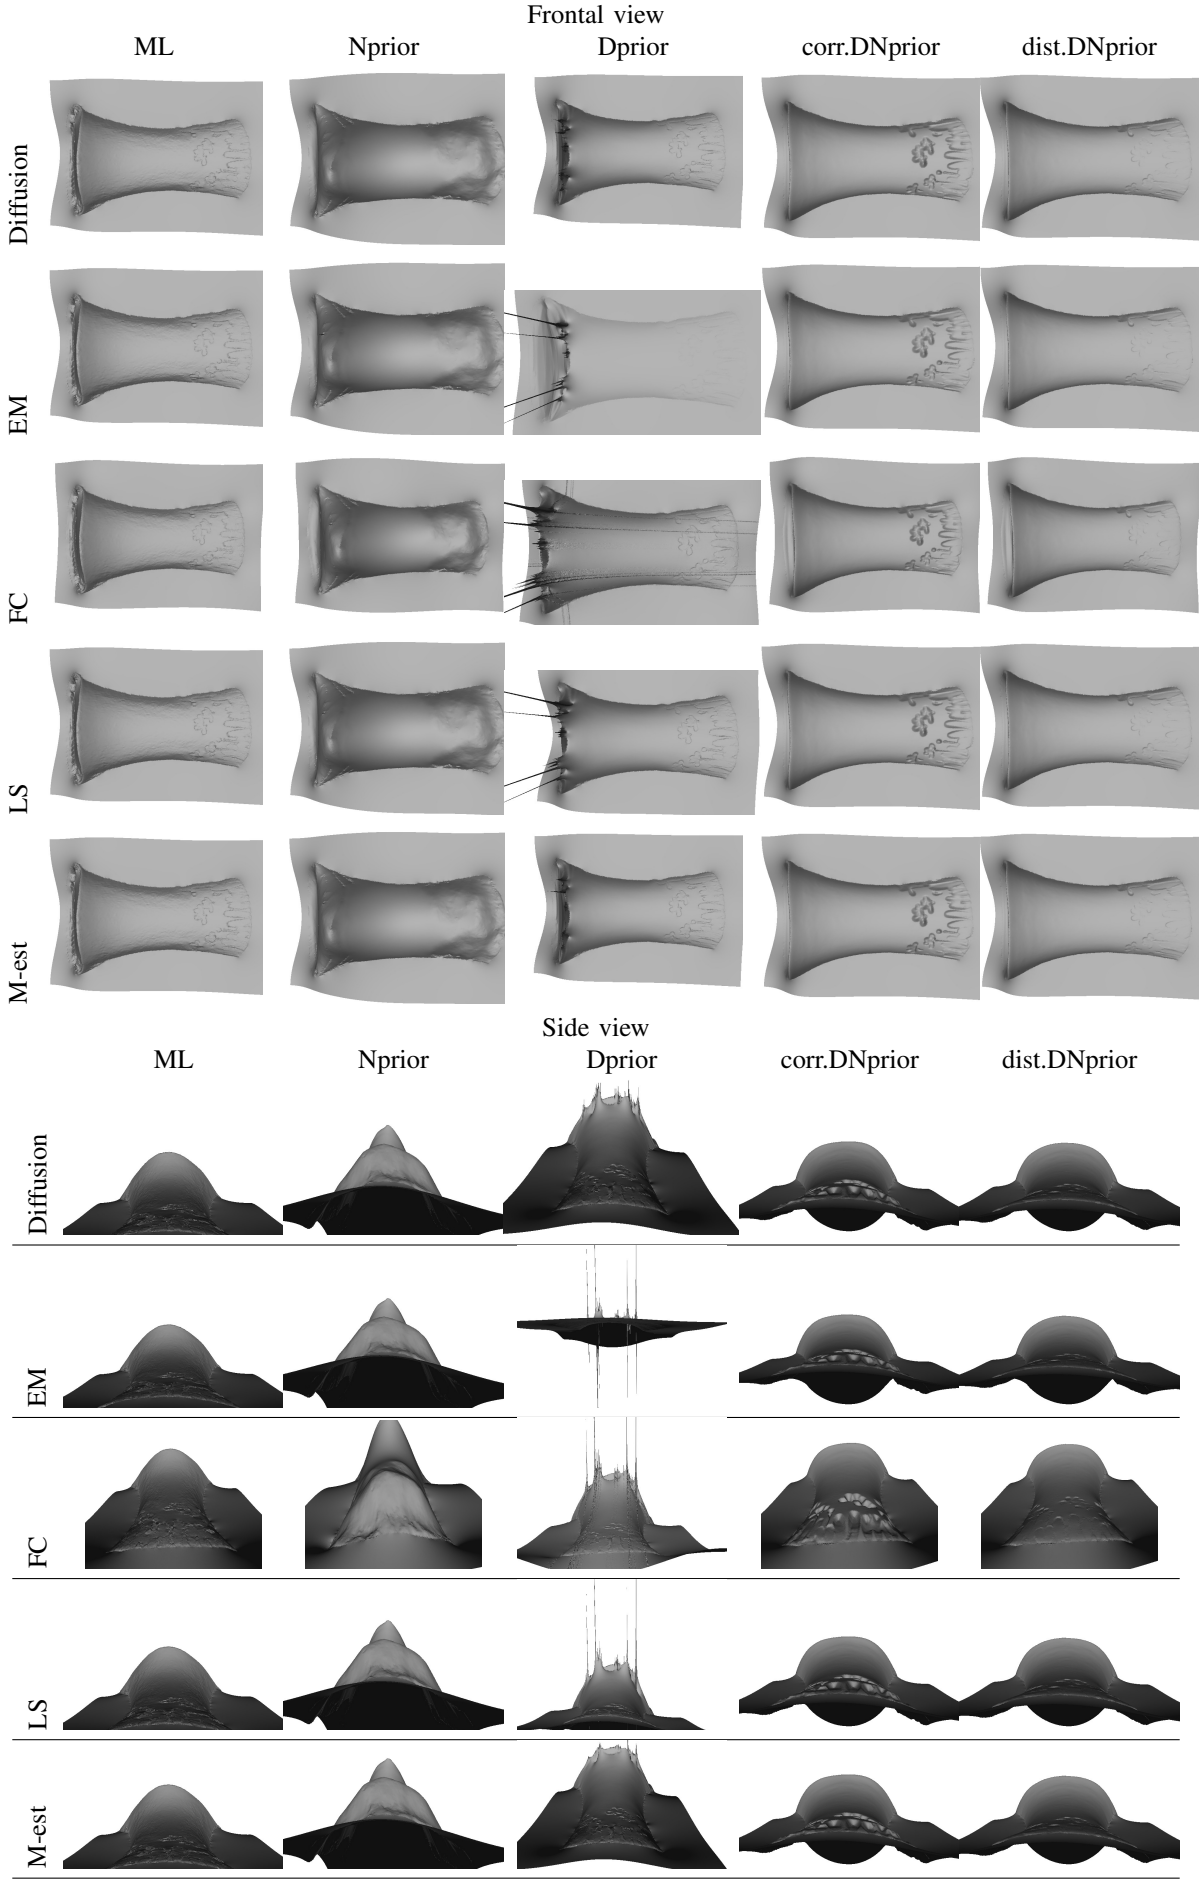
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

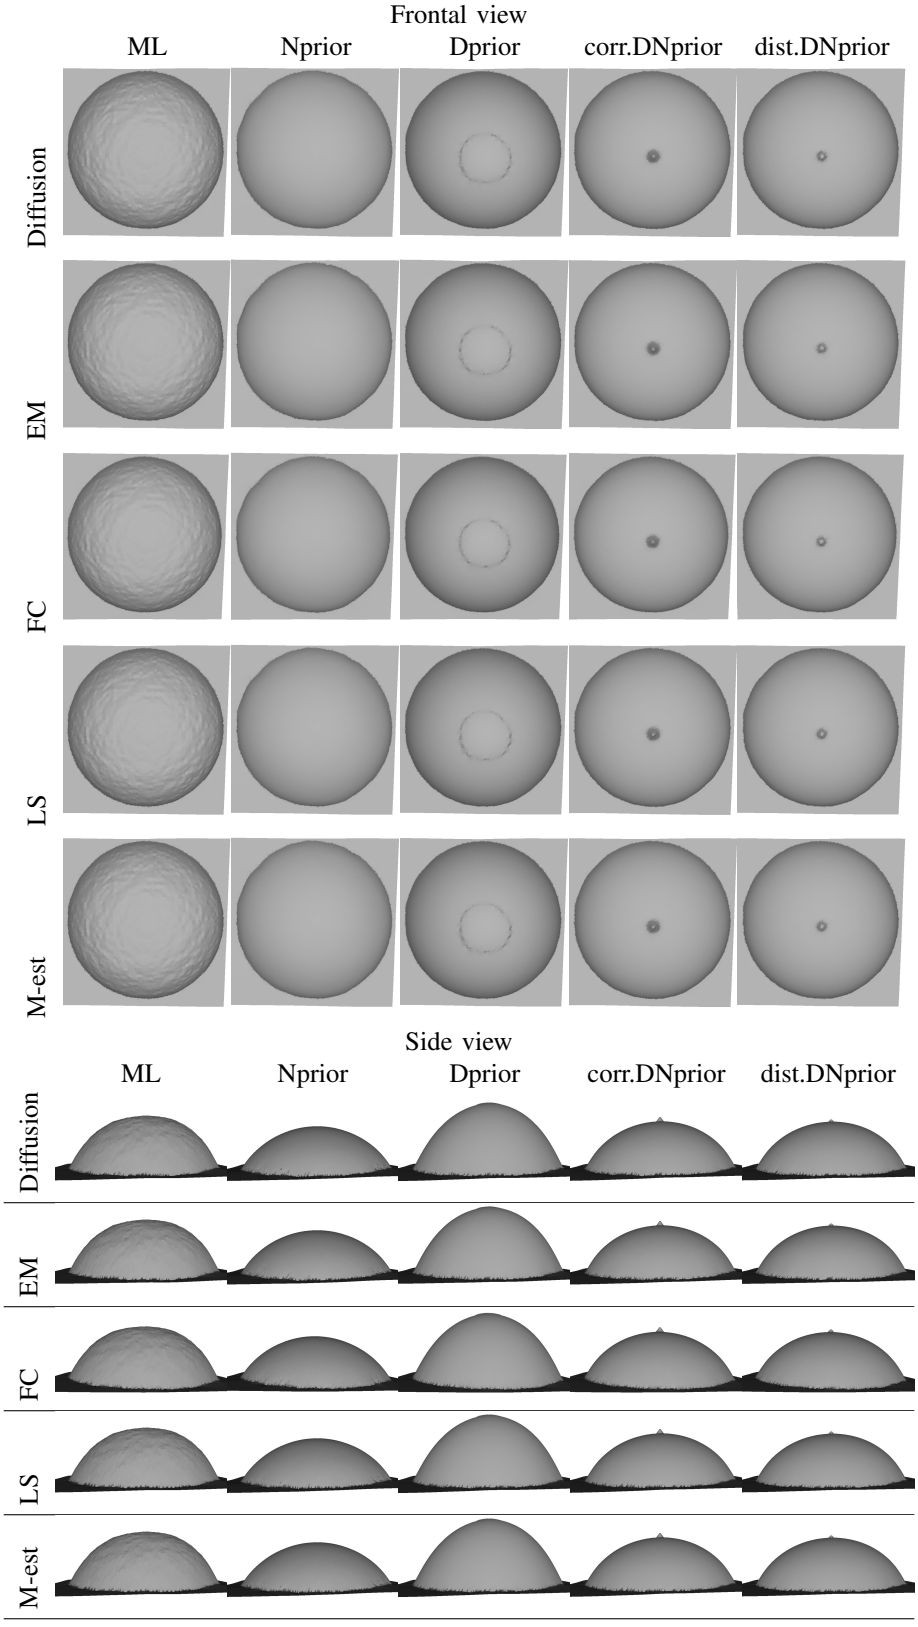
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

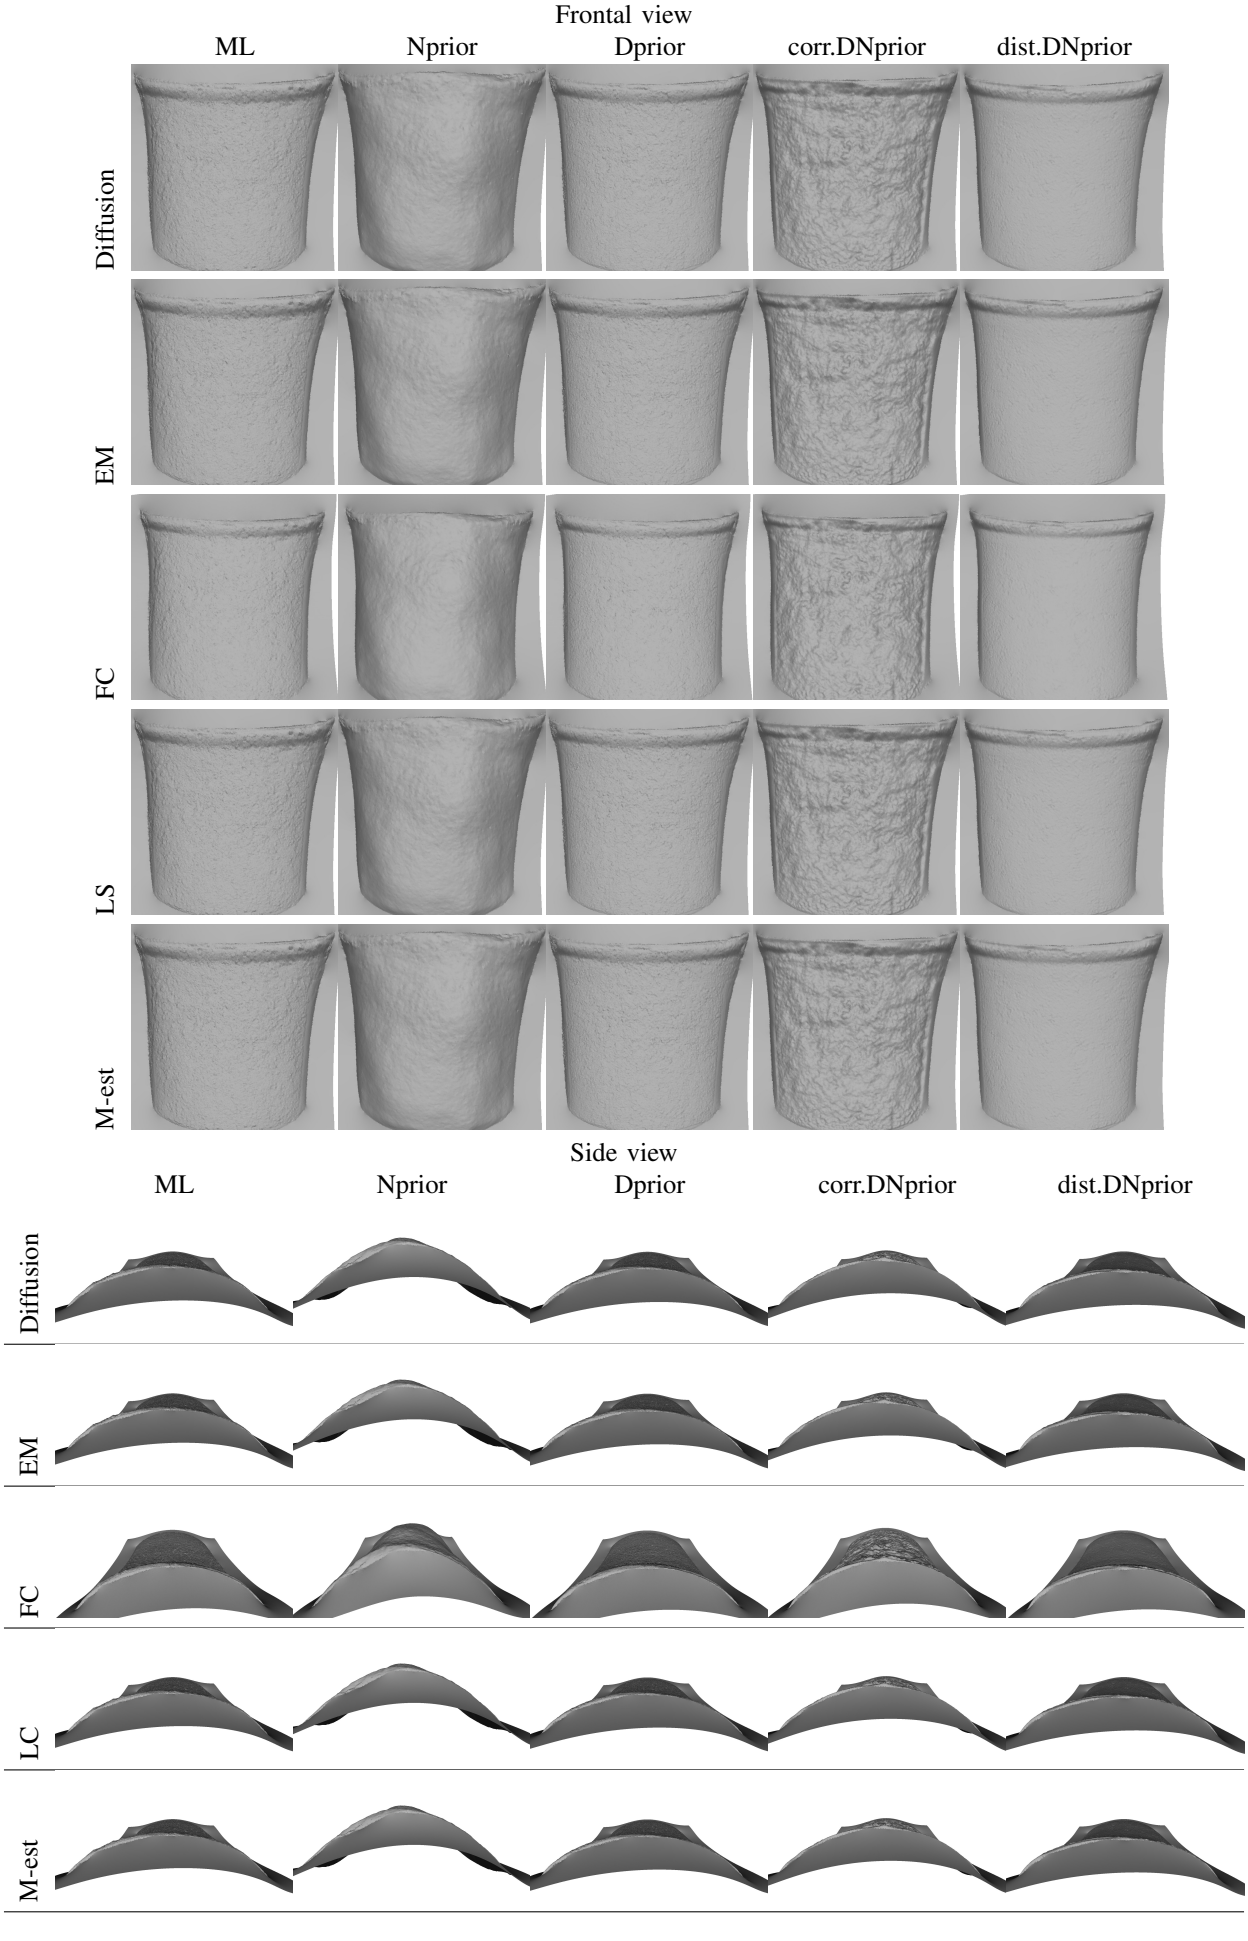
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LC
M-est

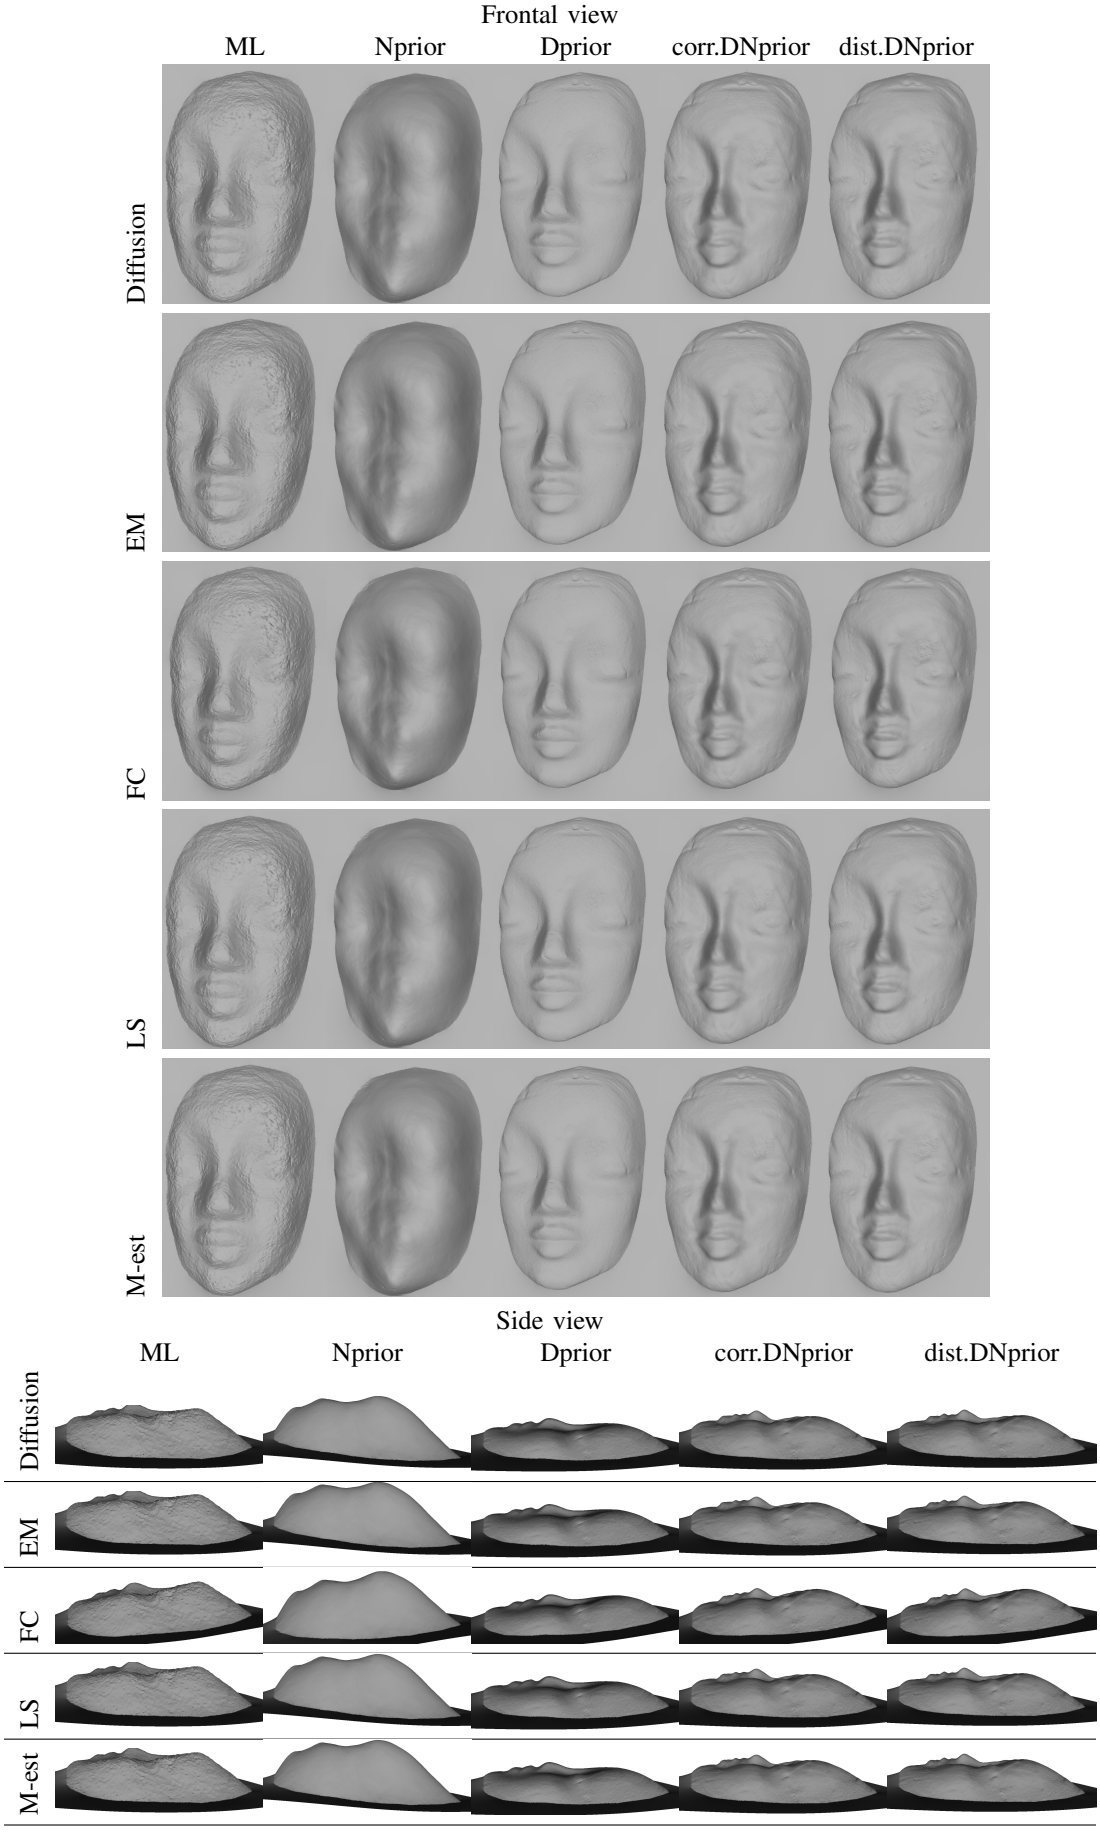
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est

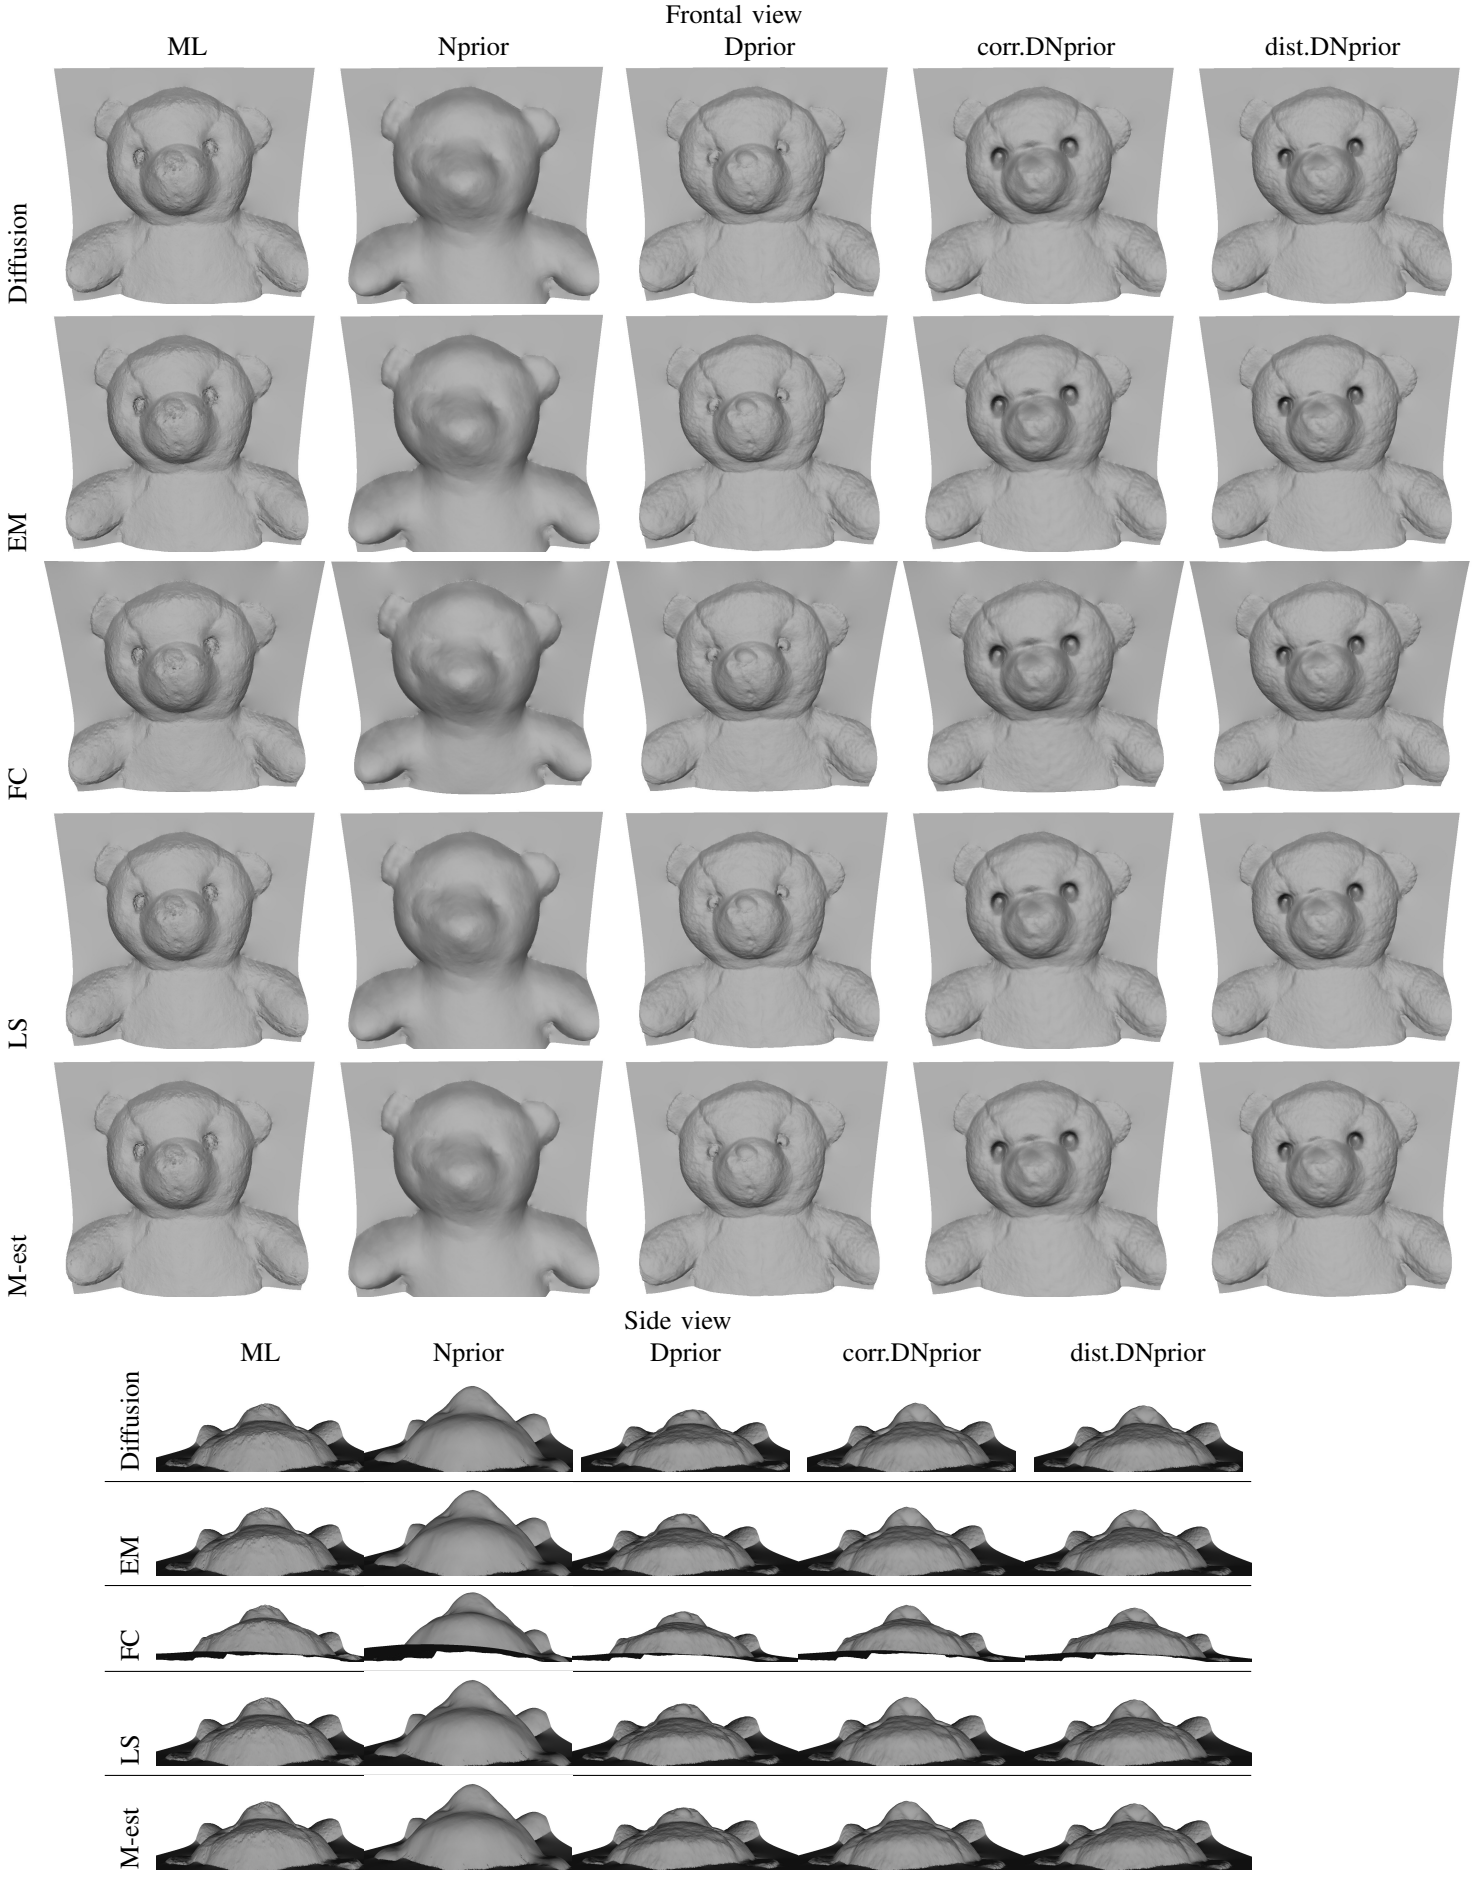
Frontal view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est
Side view
ML
Nprior
Dprior
corr.DNprior
dist.DNprior
Diffusion
EM
FC
LS
M-est