Neural Cages for Detail-Preserving 3D Deformations

Wang Yifan[1] Noam Aigerman[2] Vladimir G. Kim[2] Siddhartha Chaudhuri [2,3] Olga Sorkine-Hornung[1]

[1]ETH Zurich [2]Adobe Research [3]IIT Bombay

In the paper, we detailed two applications of our method: stock amplification and deformation transfer. In this document, we provide additional results for each application.

Stock amplification via deformation. In Table 1, 2 and 3 we show the entire deformation results from the *chair*, *table* and *car* categories. As mentioned in the main paper, the test sets consist of 100 pairs of unseen source and target shapes randomly sampled from ShapeNet dataset [2].

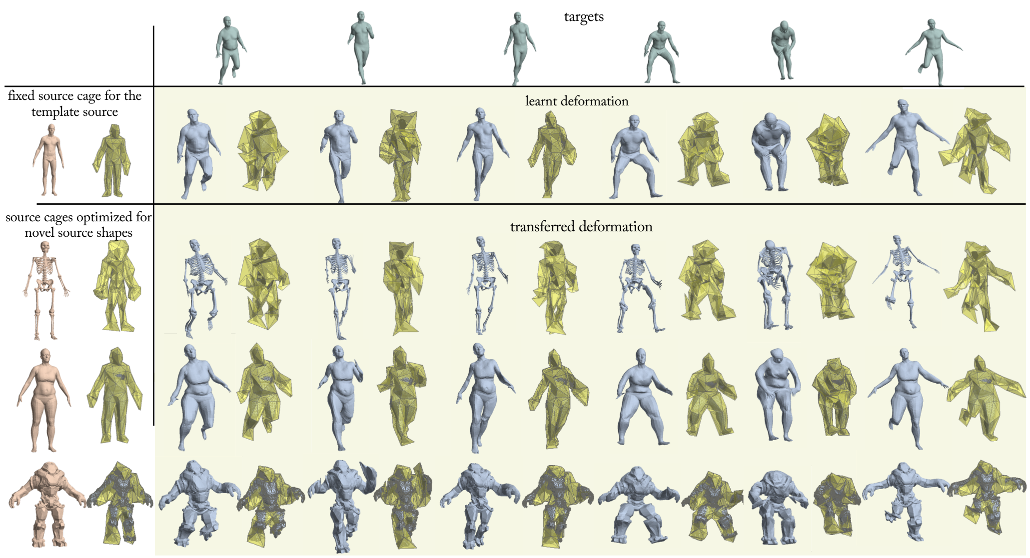

Figure 1: Visualization of the cages used in deformation transfer. Given a template source shape in a known pose with a manually created template cage (second row, left, brown), our deformation network translates the vertices of the template cage to match the source to a novel target shape in various poses (second row, blue). This cage-deformation can be transferred to a new character, provided a pose similar to that of the template source (row 3-5, left, brown). The deformed novel characters are shown in blue in row 3-5.

Deformation transfer. In Figure 1 we illustrate the cages used in the humanoid deformation examples. In Table 4 we provide additional results for 100 unseen target poses randomly sampled from the dataset provided by [3]. For each target pose, we show the deformed shape for the original source shape used during training as well as the transferred deformation for three novel source shapes, a woman from FAUST [1], a skeleton and a robot.

Two training variations are shown in Table 4. On the left, a rest-pose human is used as the source shape during training, whereas on the right a T-pose human is used. The first variation shows less distortion, but tends to underperform when large arm articulation is present in the target pose. The second variation shows more distortion on the arms but tends to match the articulation of the target poses slightly better, especially when the arms are up. The results illustrated in the main paper are generated using the first variation.

We think these distortions are related to the affine-invariance of MVC, and should be improved with Green Coordinates.

Table 1: The deformation results of the *chair* category.

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

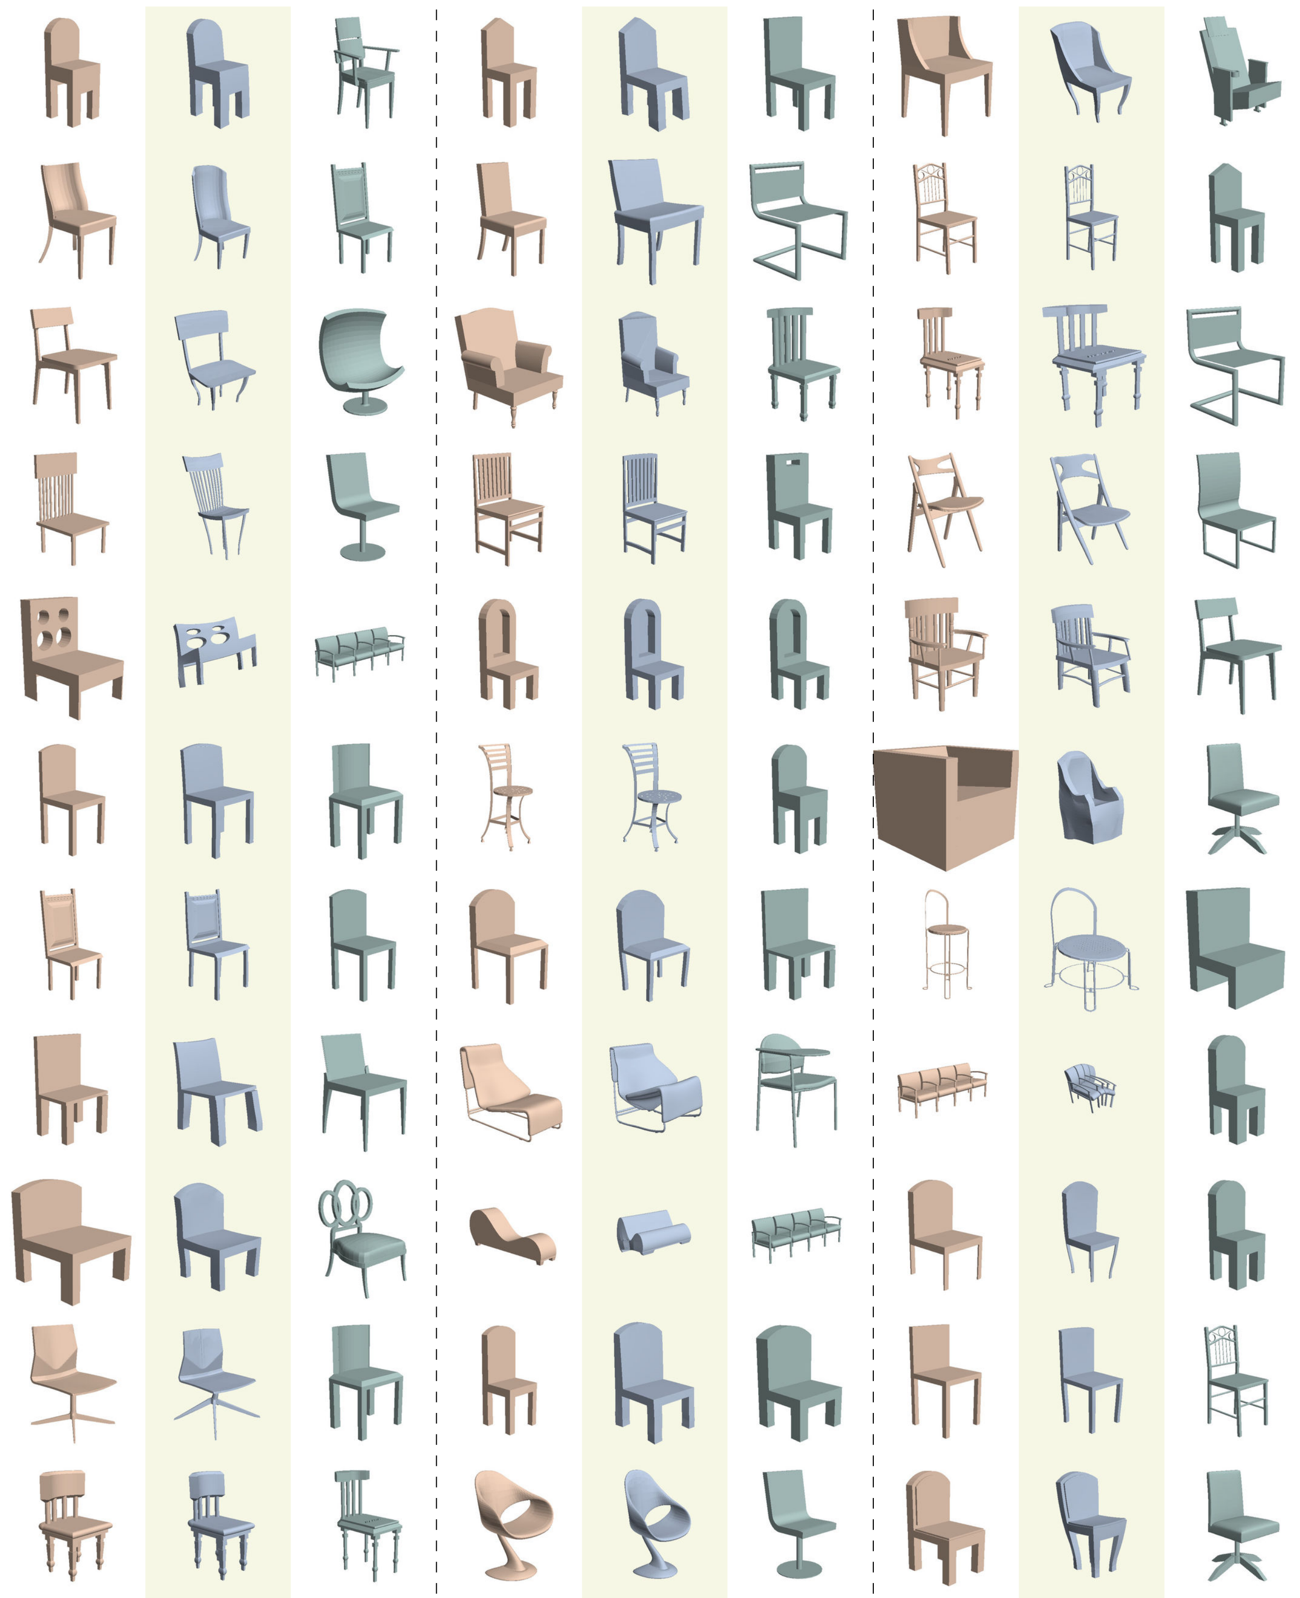

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

Continued on the next page

Continued from the last page

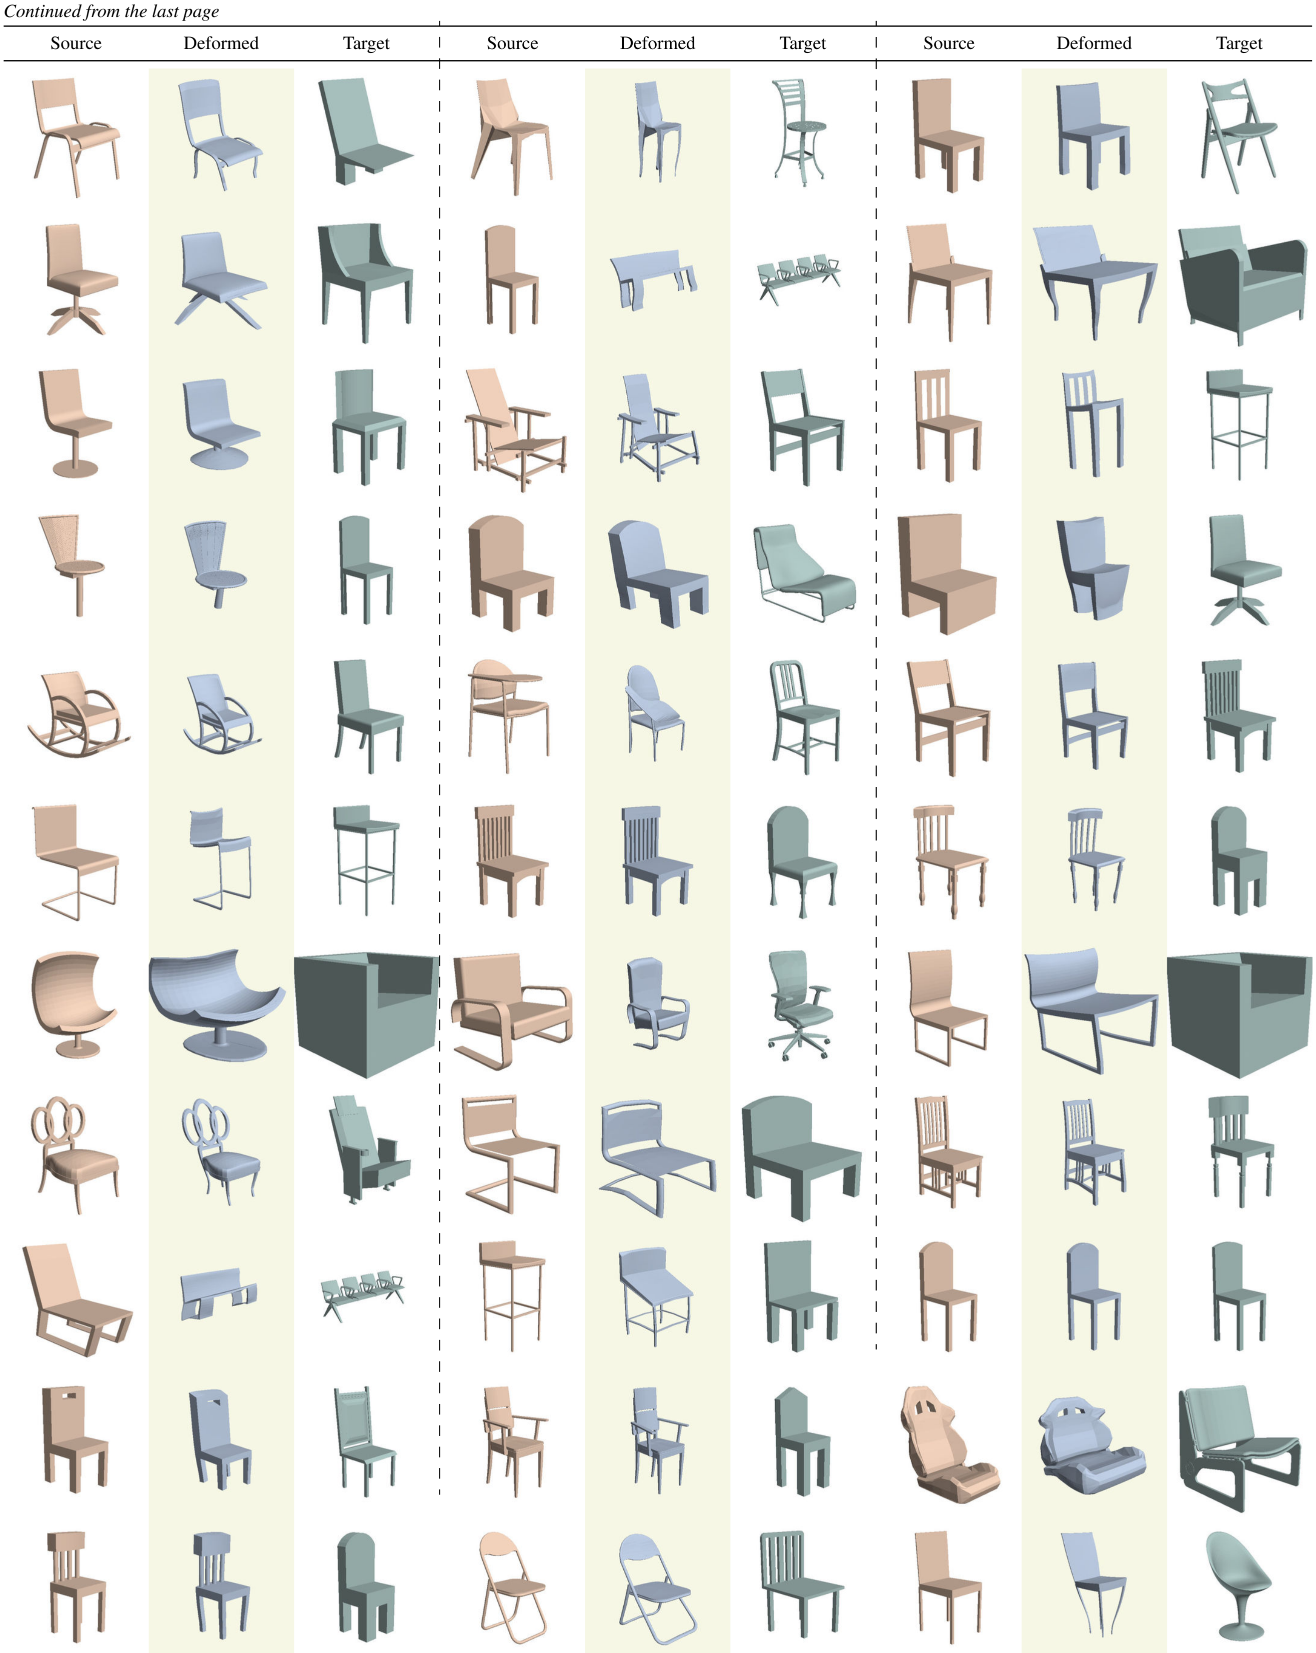

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

Table 2: The deformation results of the *table* category.

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

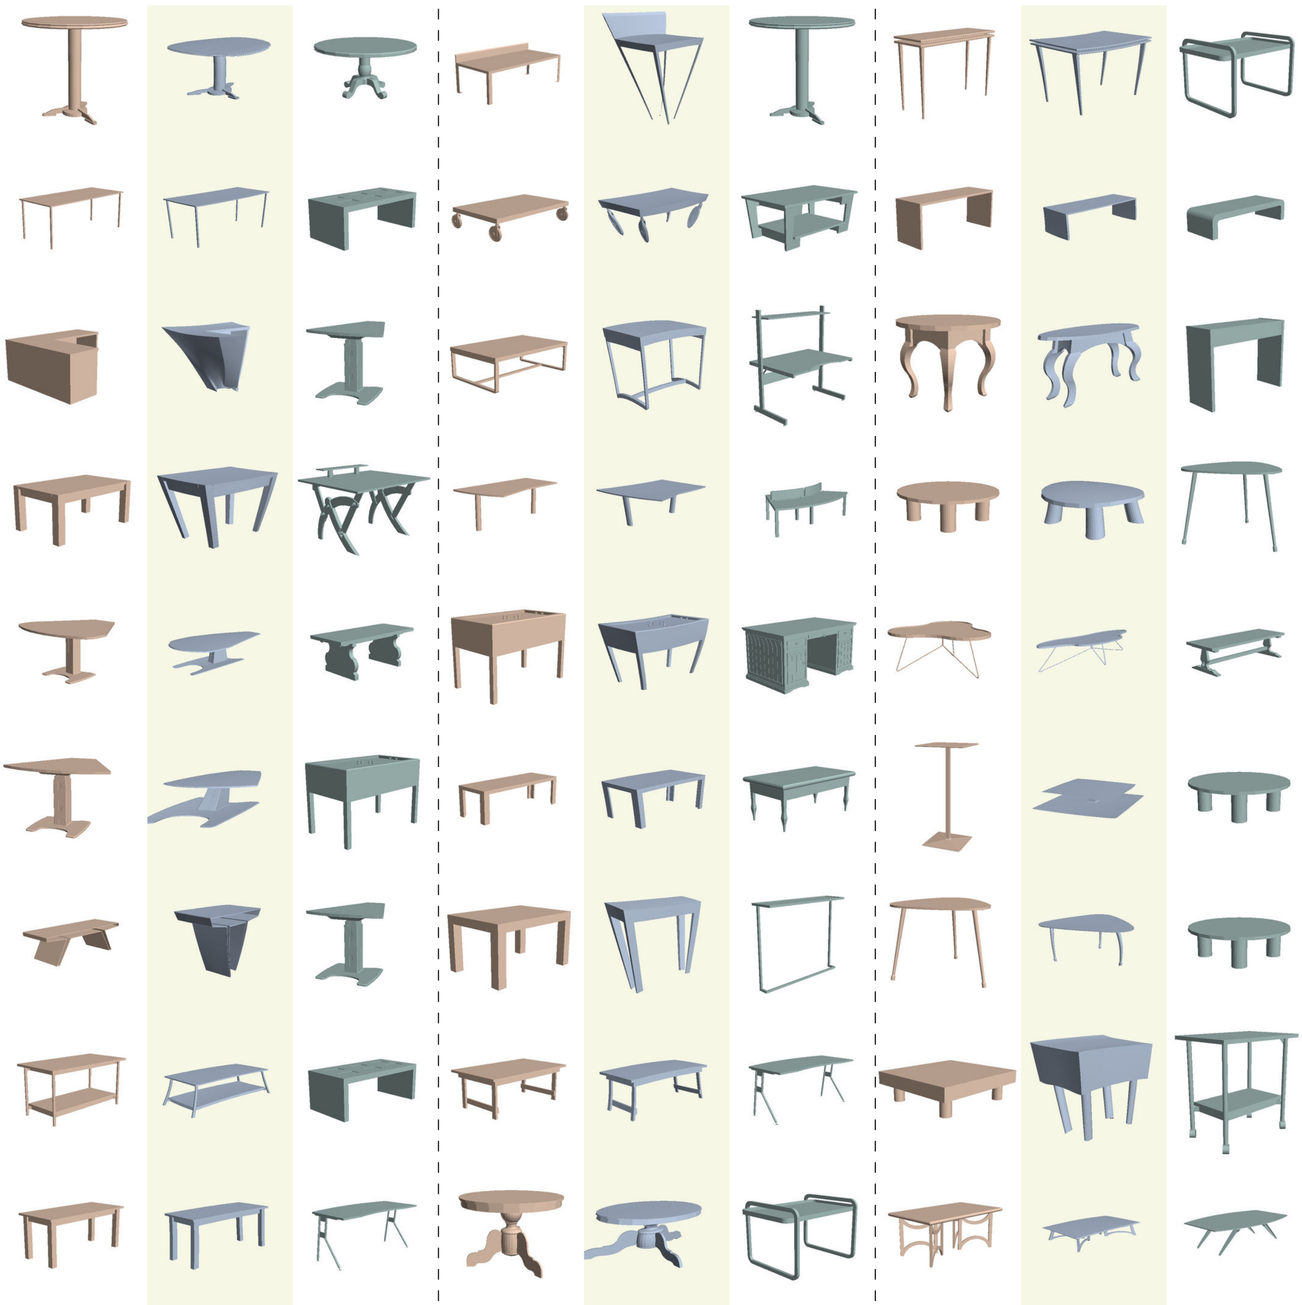

Continued on the next page

Continued from the last page

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

Table 3: The deformation results of the *car* category.

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

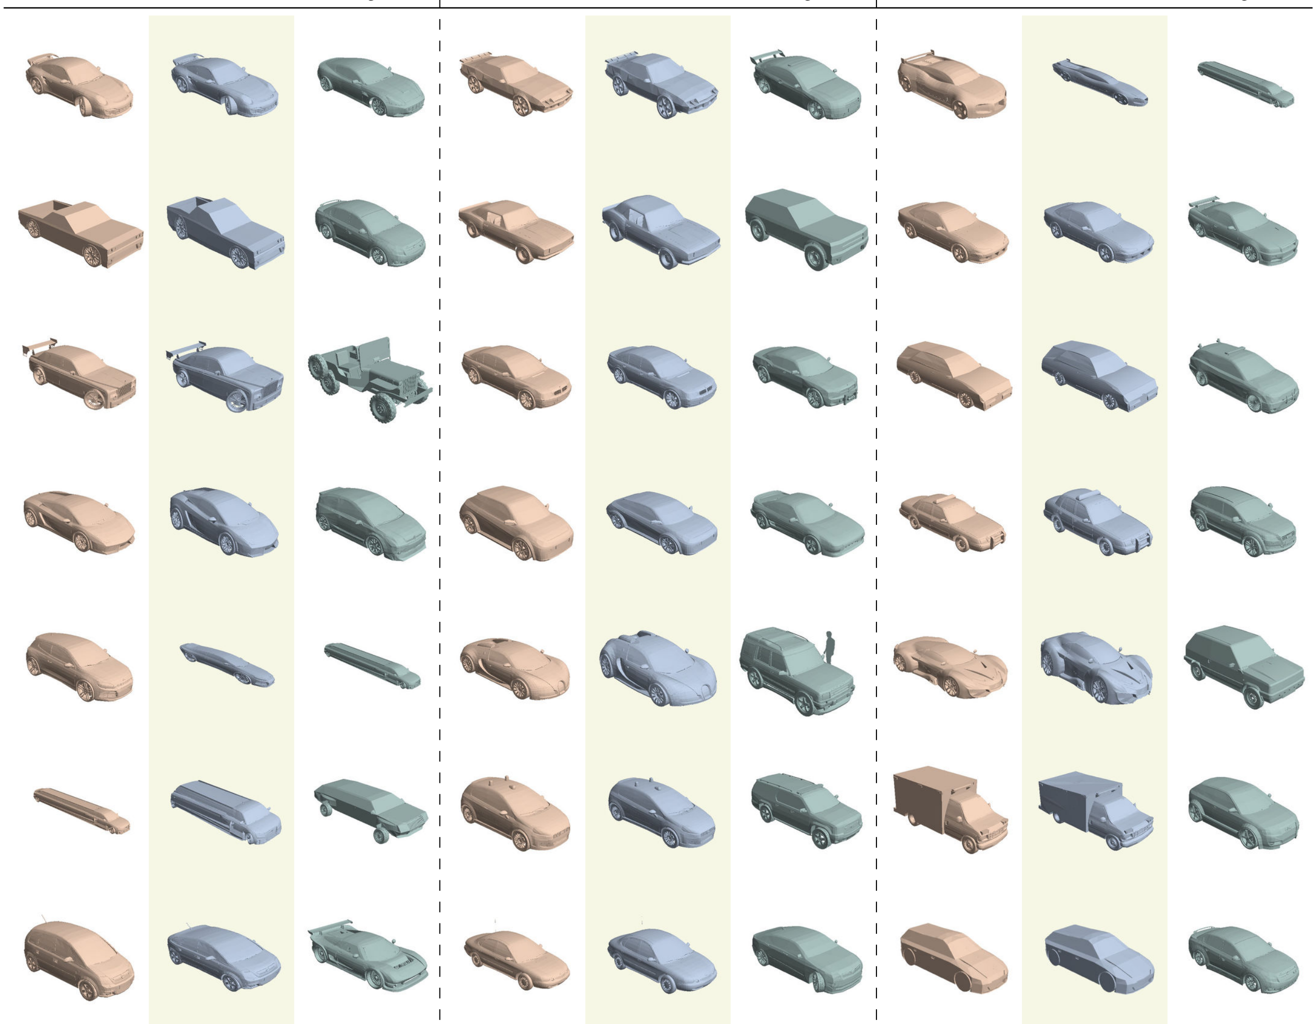

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

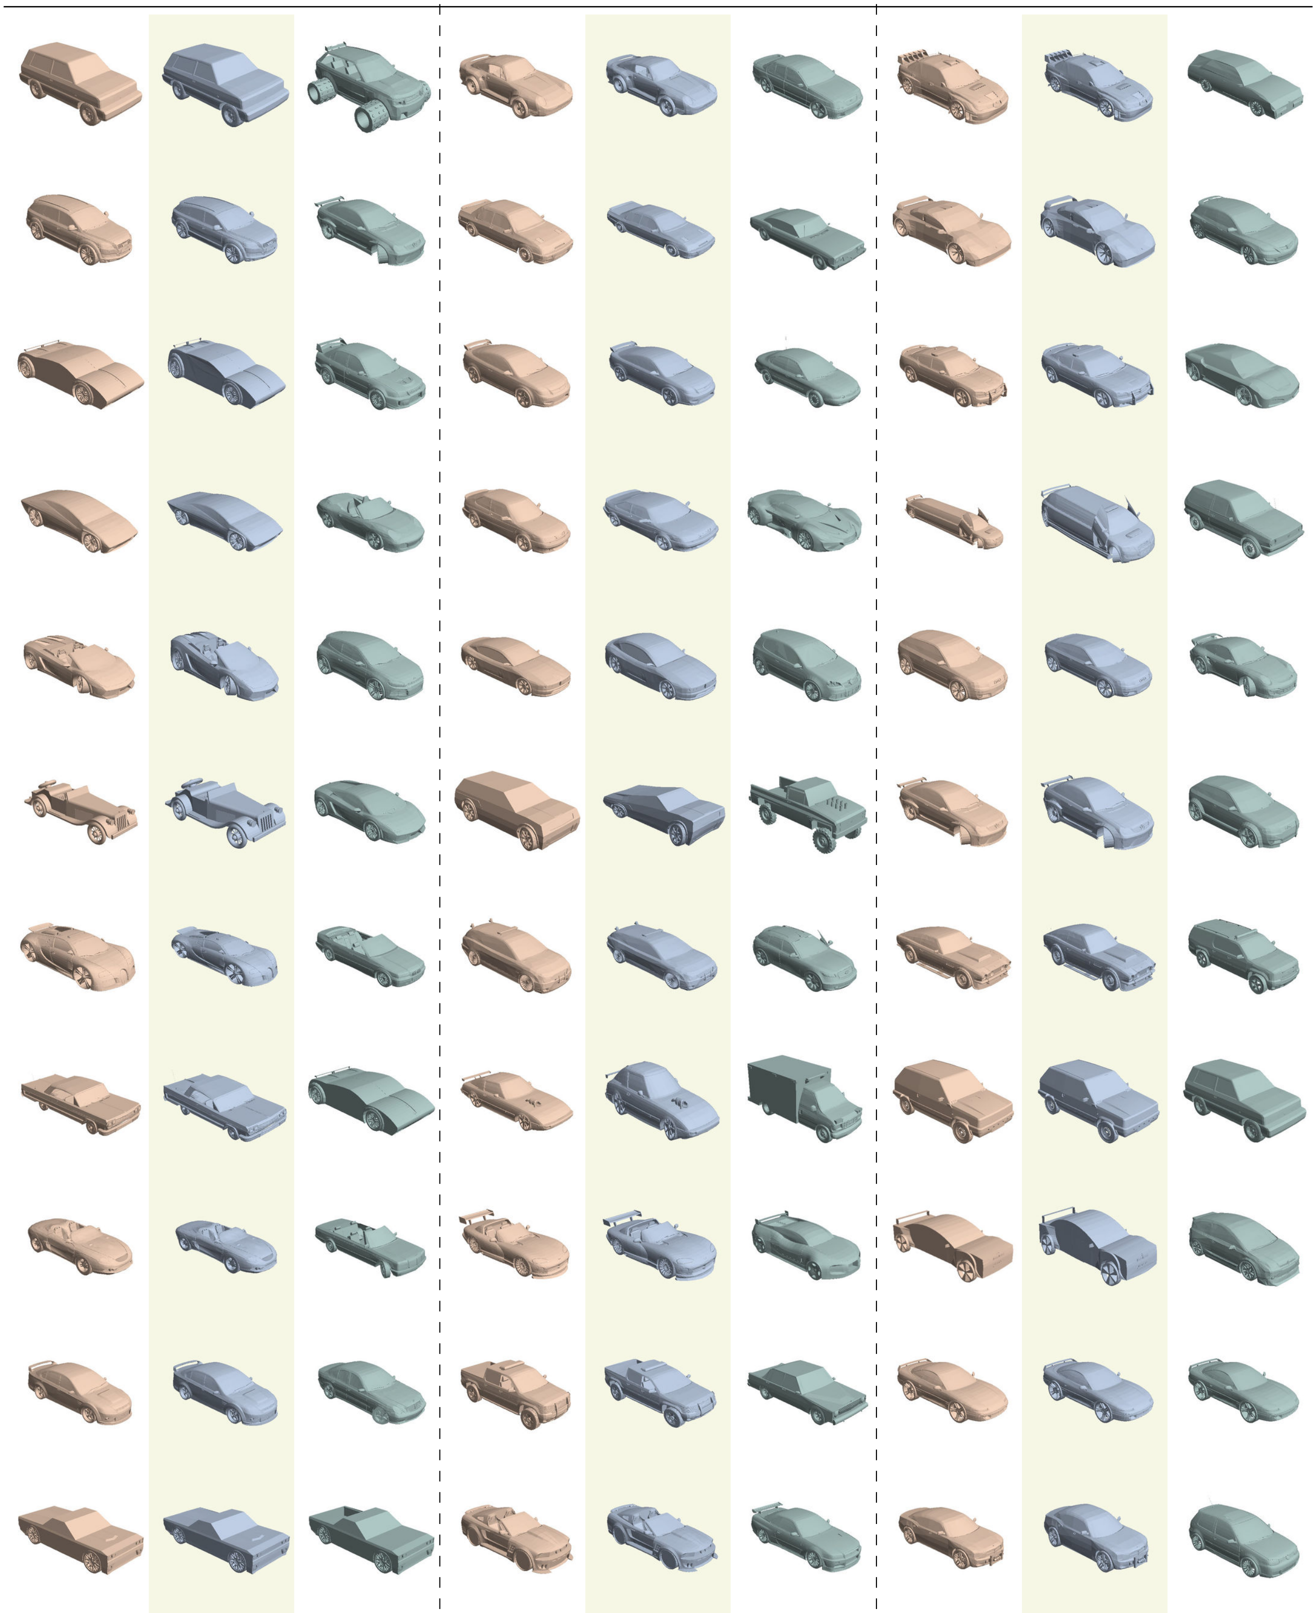

Continued on the next page

Continued from the last page

Source	Deformed	Target	Source	Deformed	Target	Source	Deformed	Target

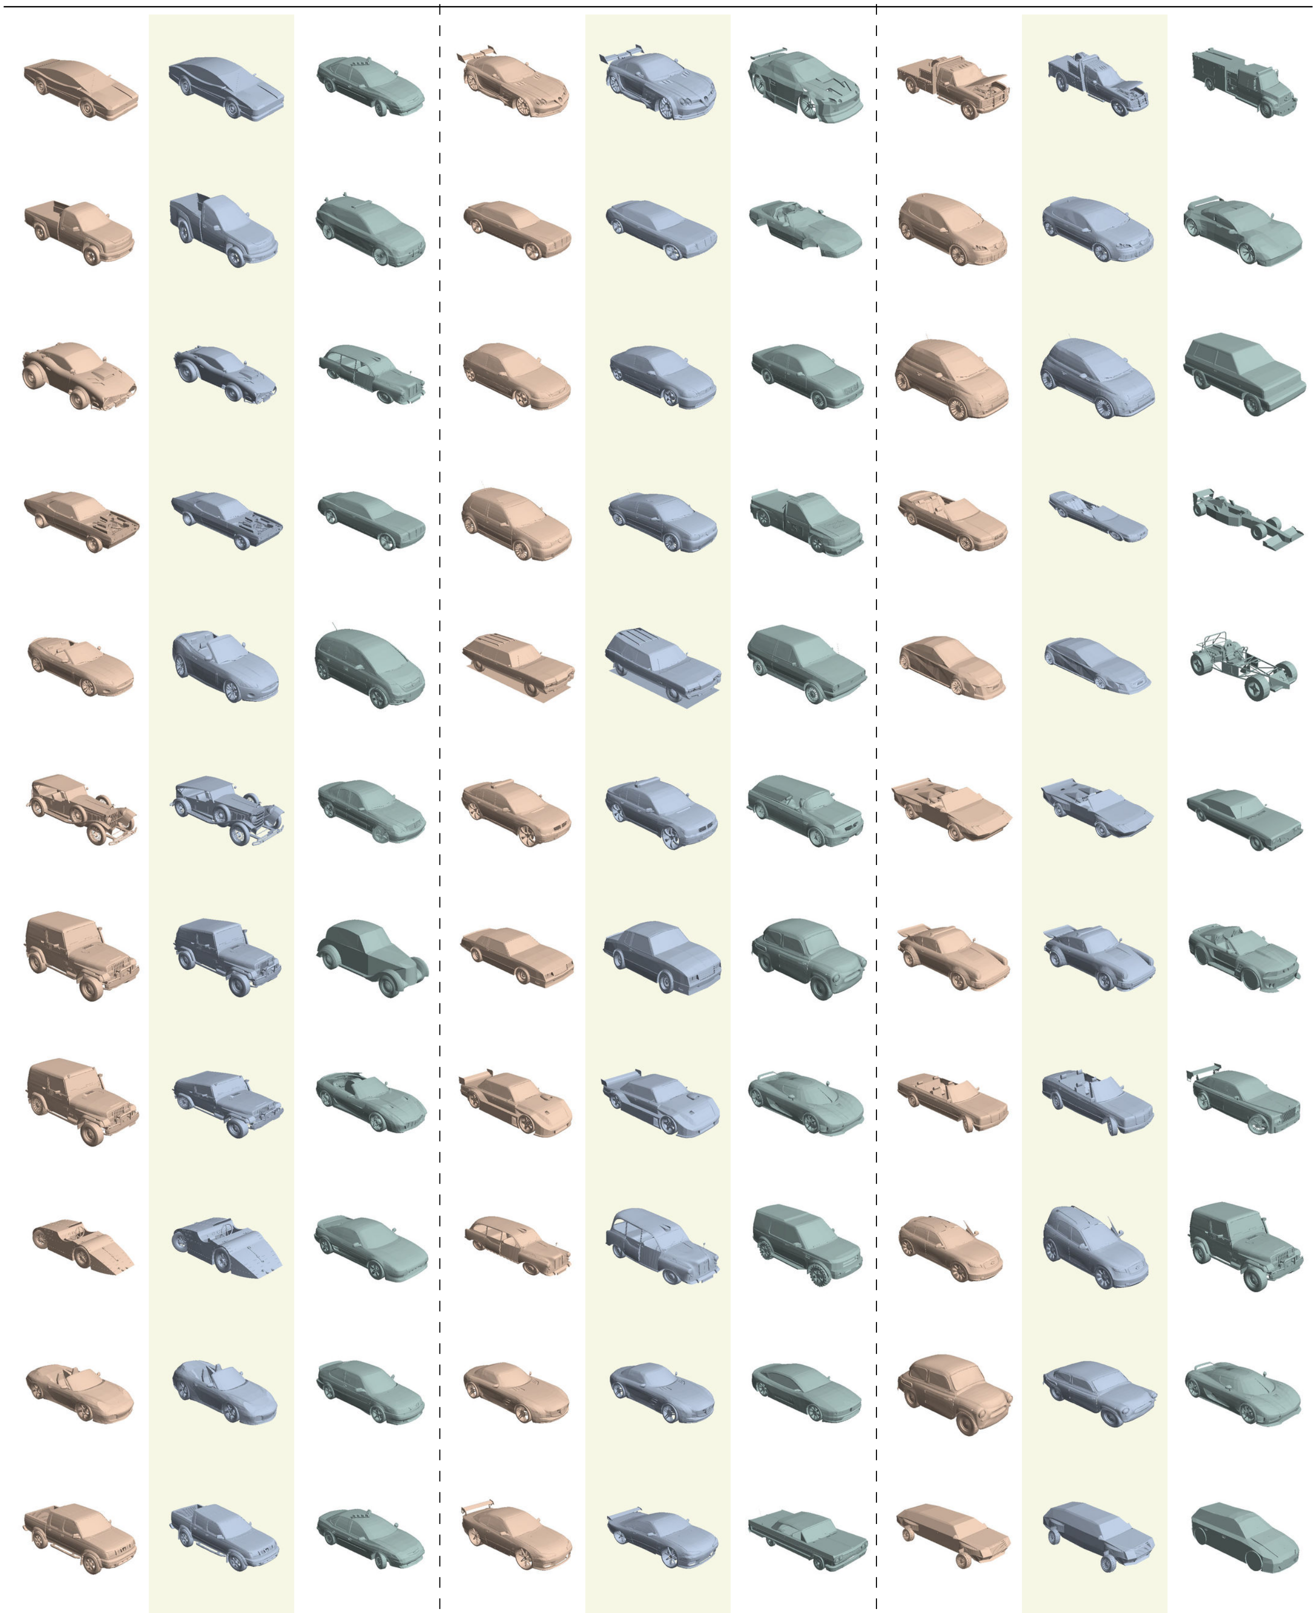

Continued on the next page

Continued from the last page

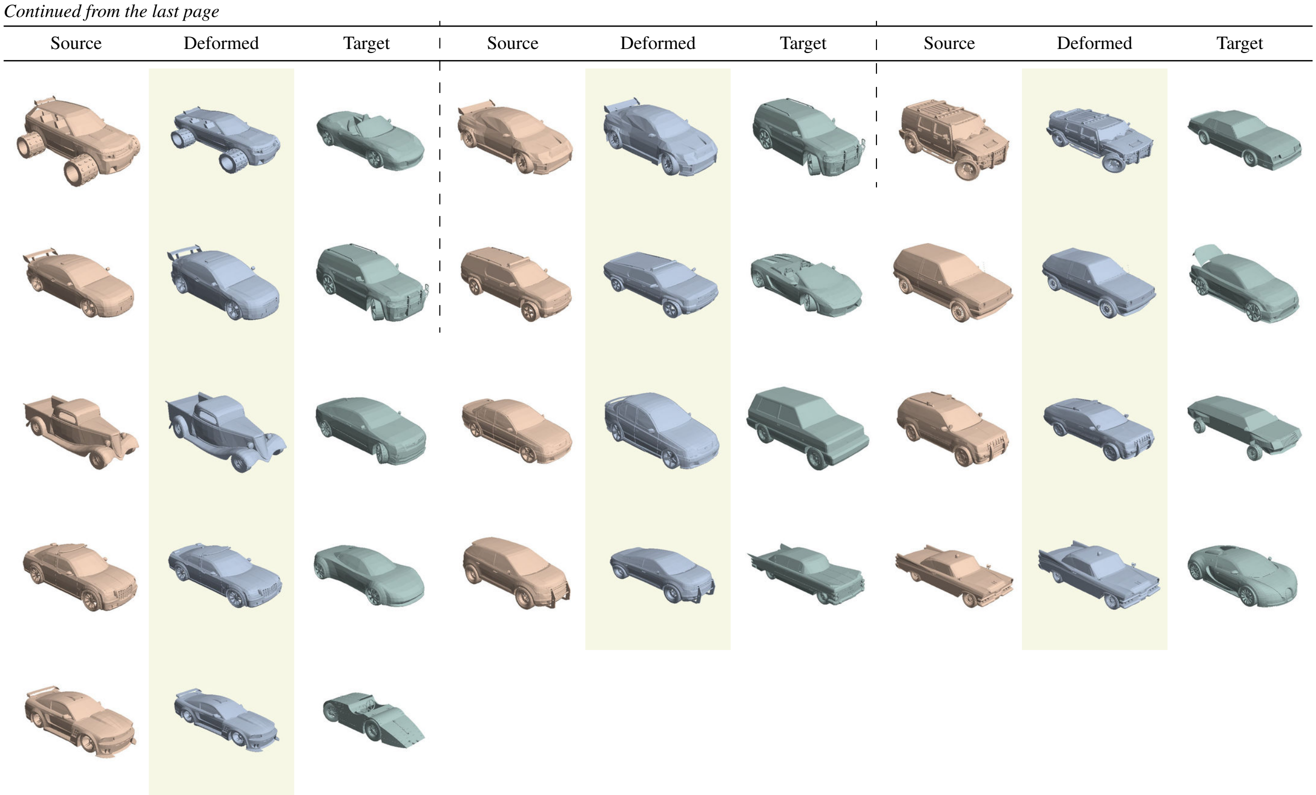

Table 4: Additional results for humanoid deformation and deformation transfer to new characters. The target poses are shown on the left in green, the template source (fixed during training) and the novel new sources are shown on the top in brown. We exhibit two training results, using a rest pose template source and using a t-pose template source, shown on the left and right half of the table respectively.

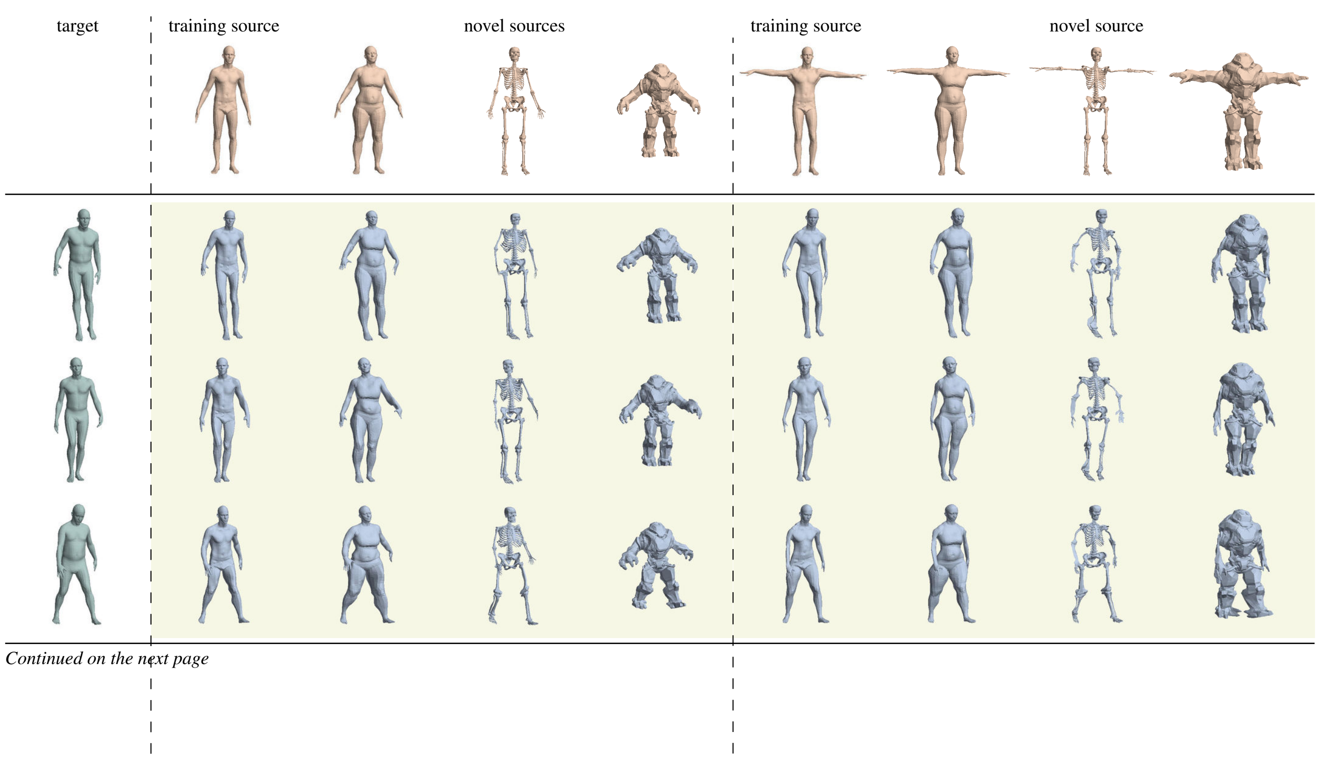

Continued on the next page

Continued from the last page

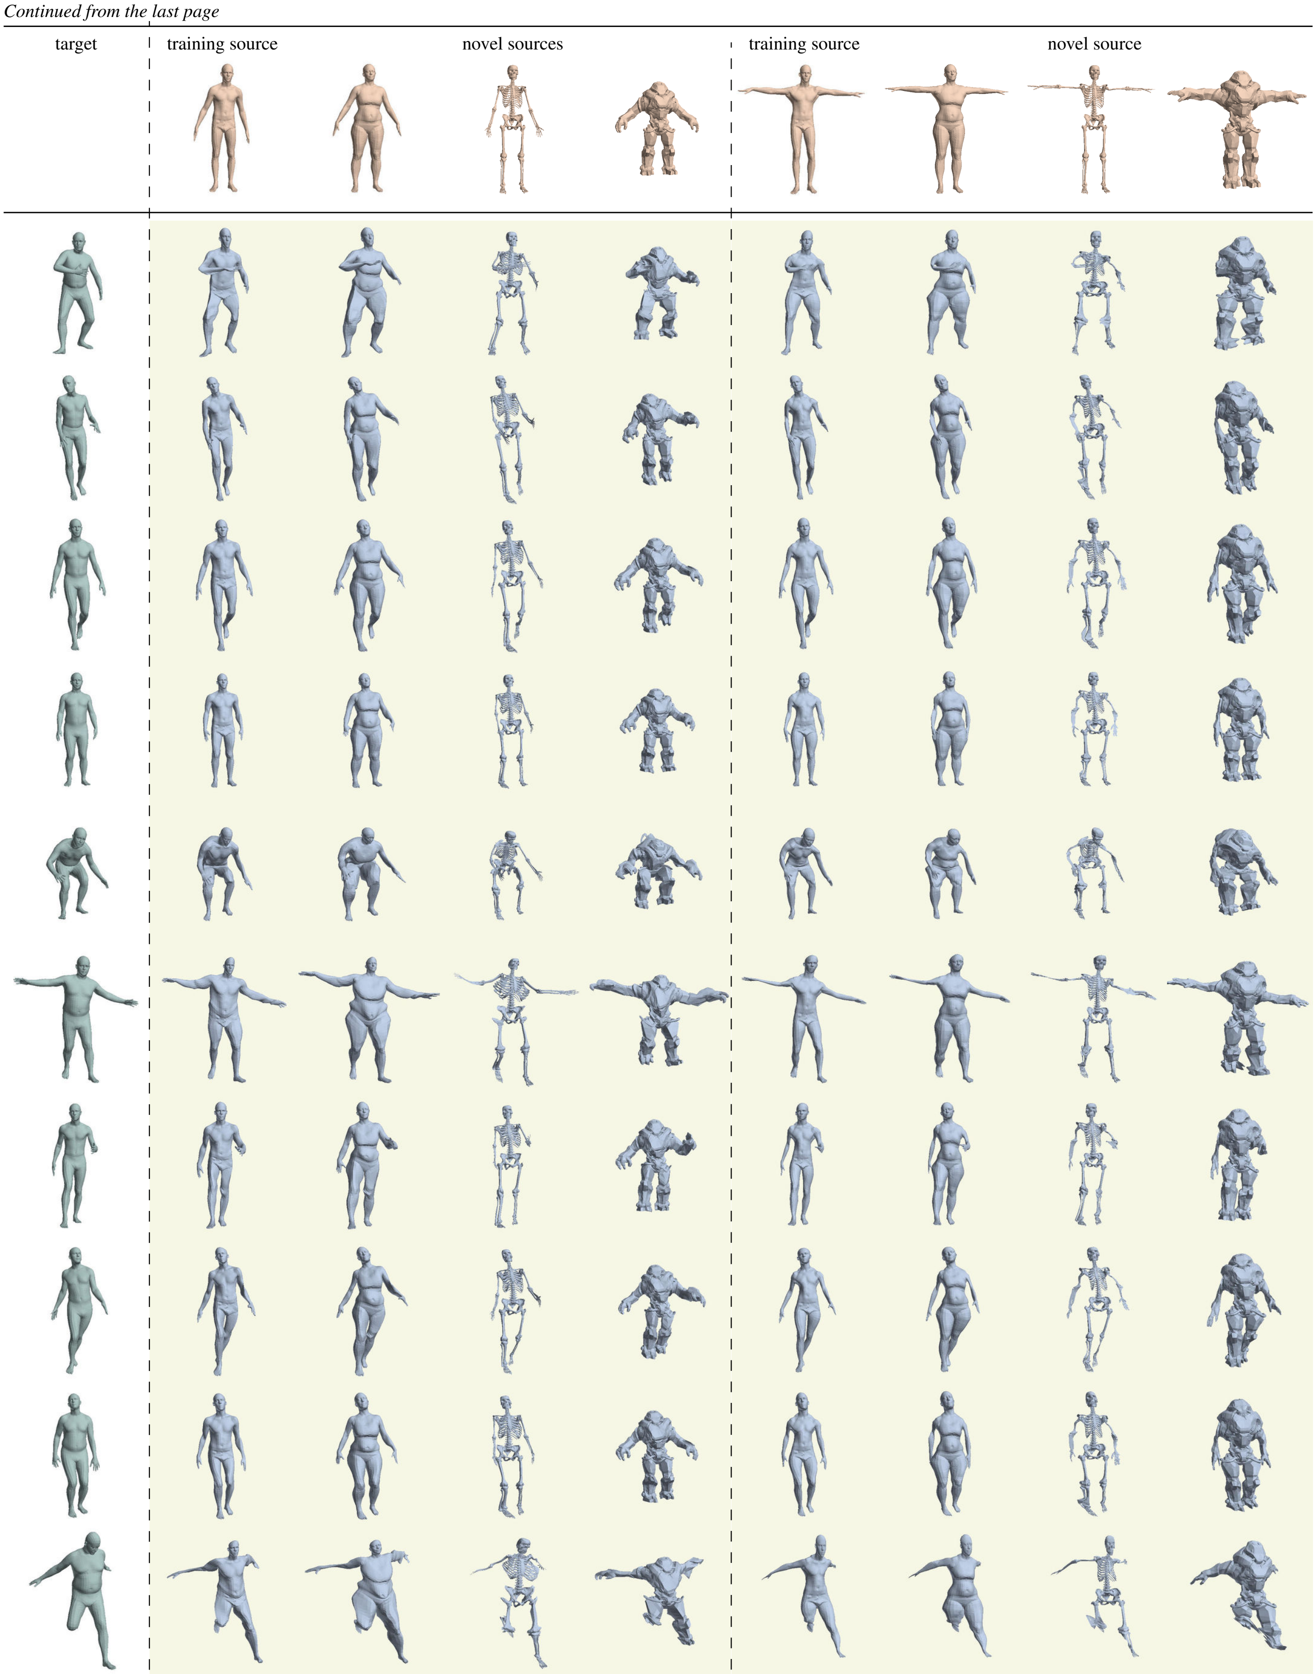

Continued on the next page

Continued from the last page

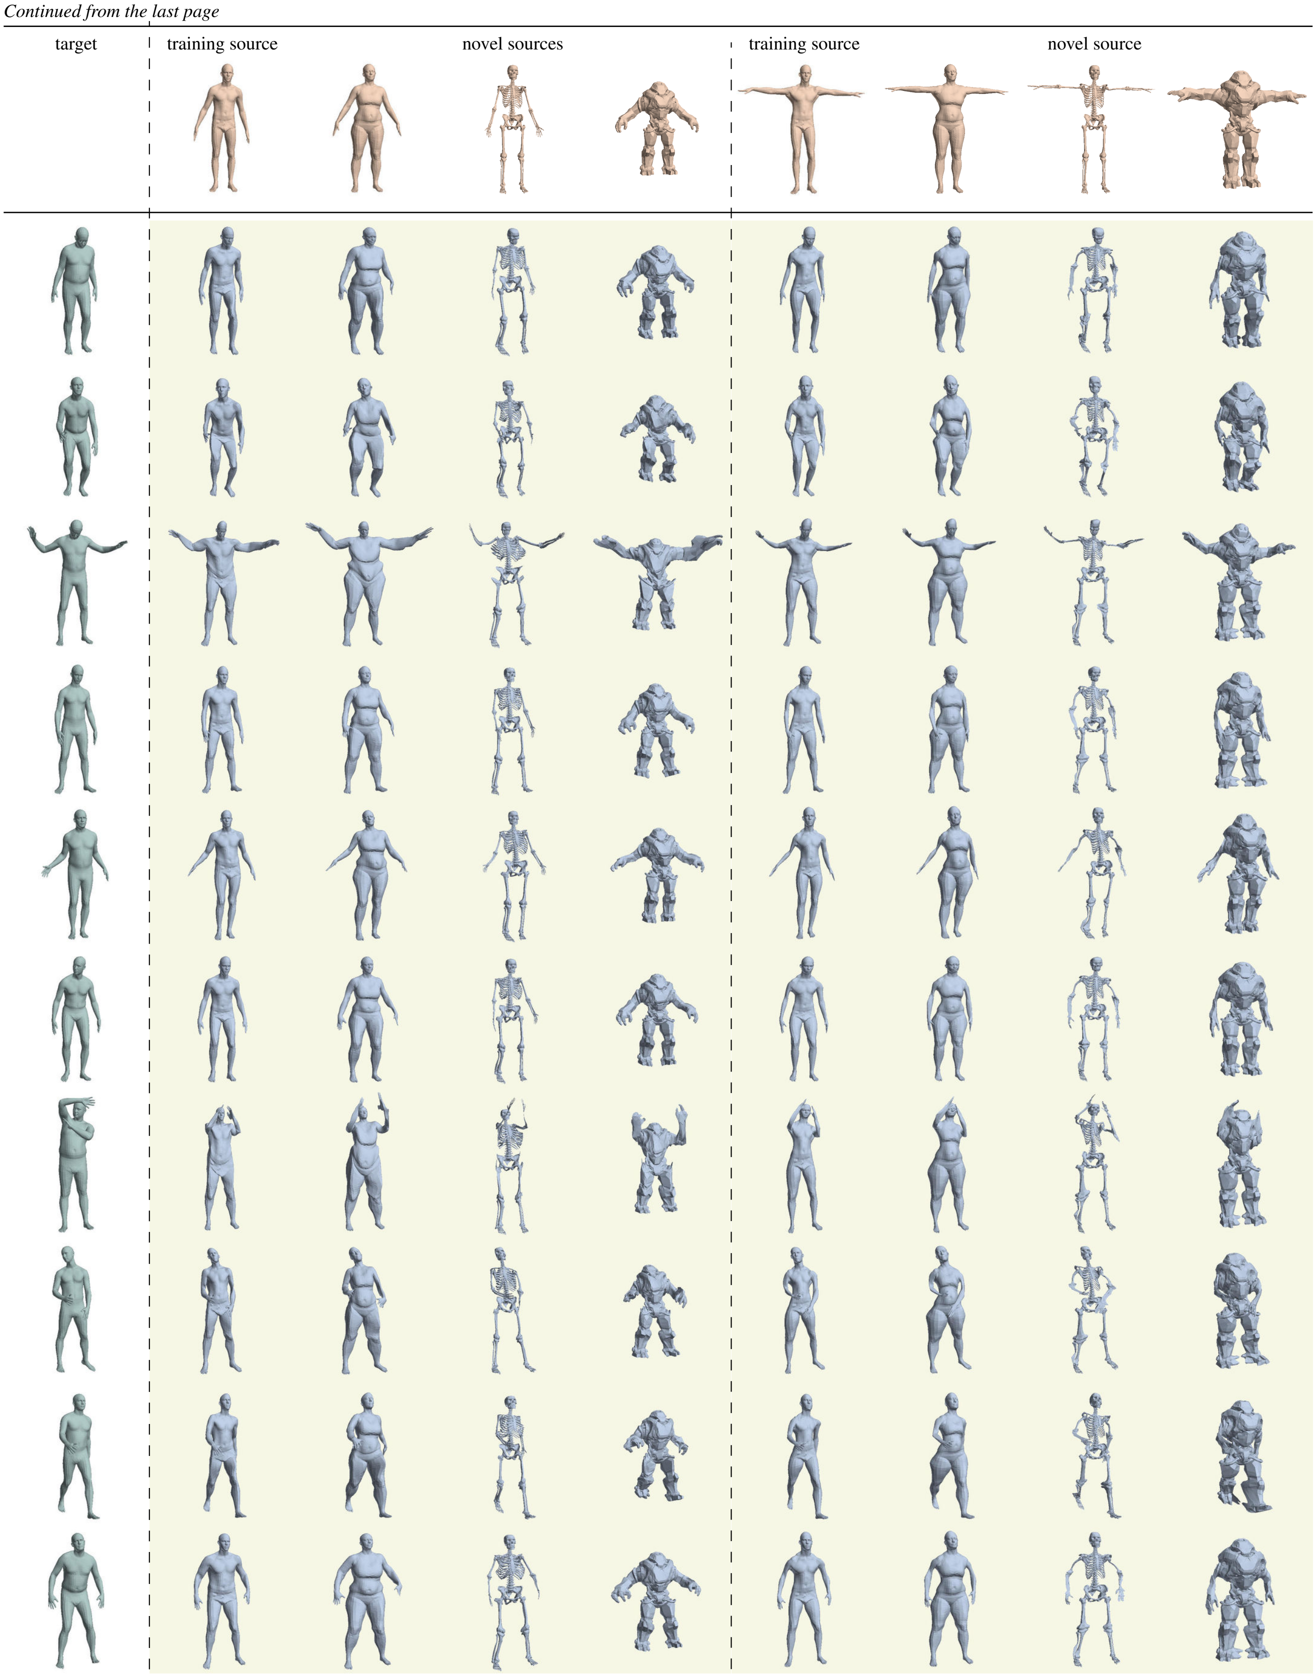

Continued on the next page

Continued from the last page

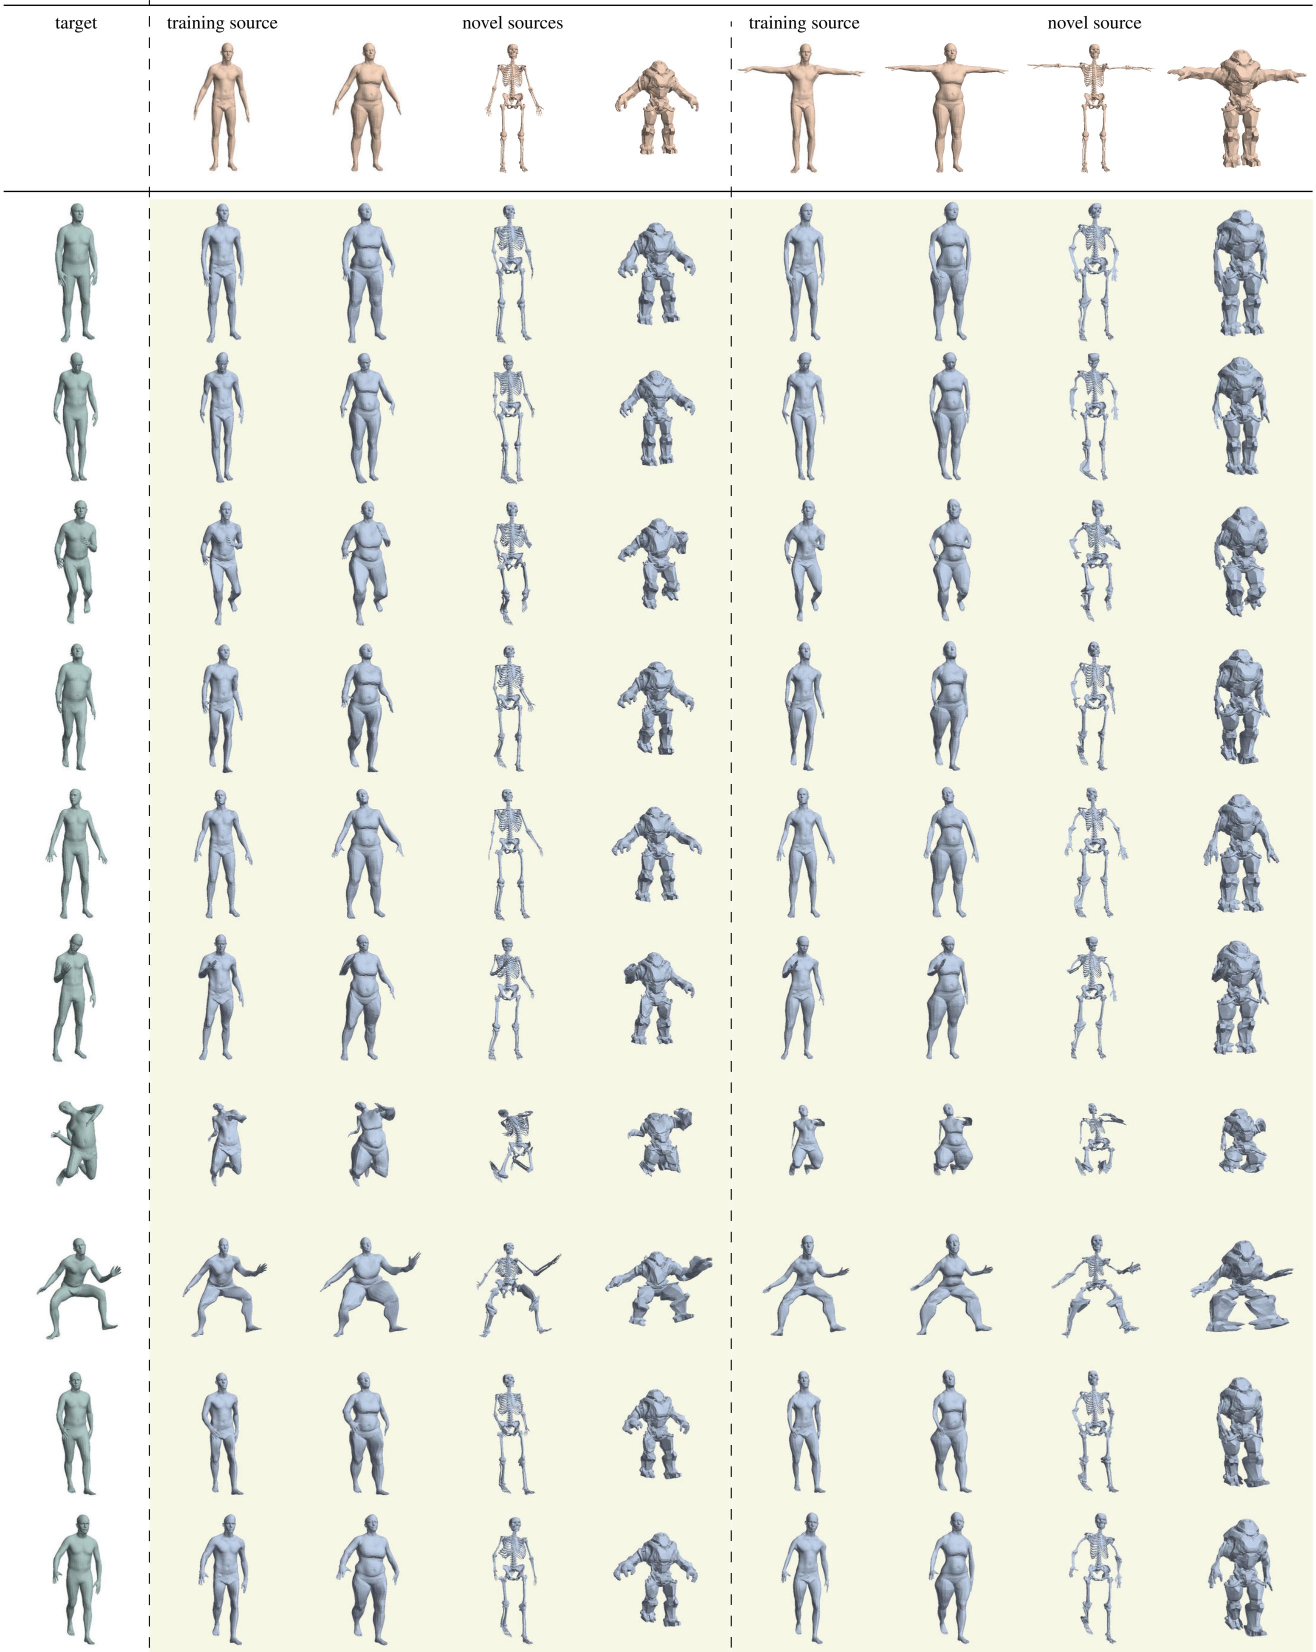

Continued on the next page

Continued from the last page

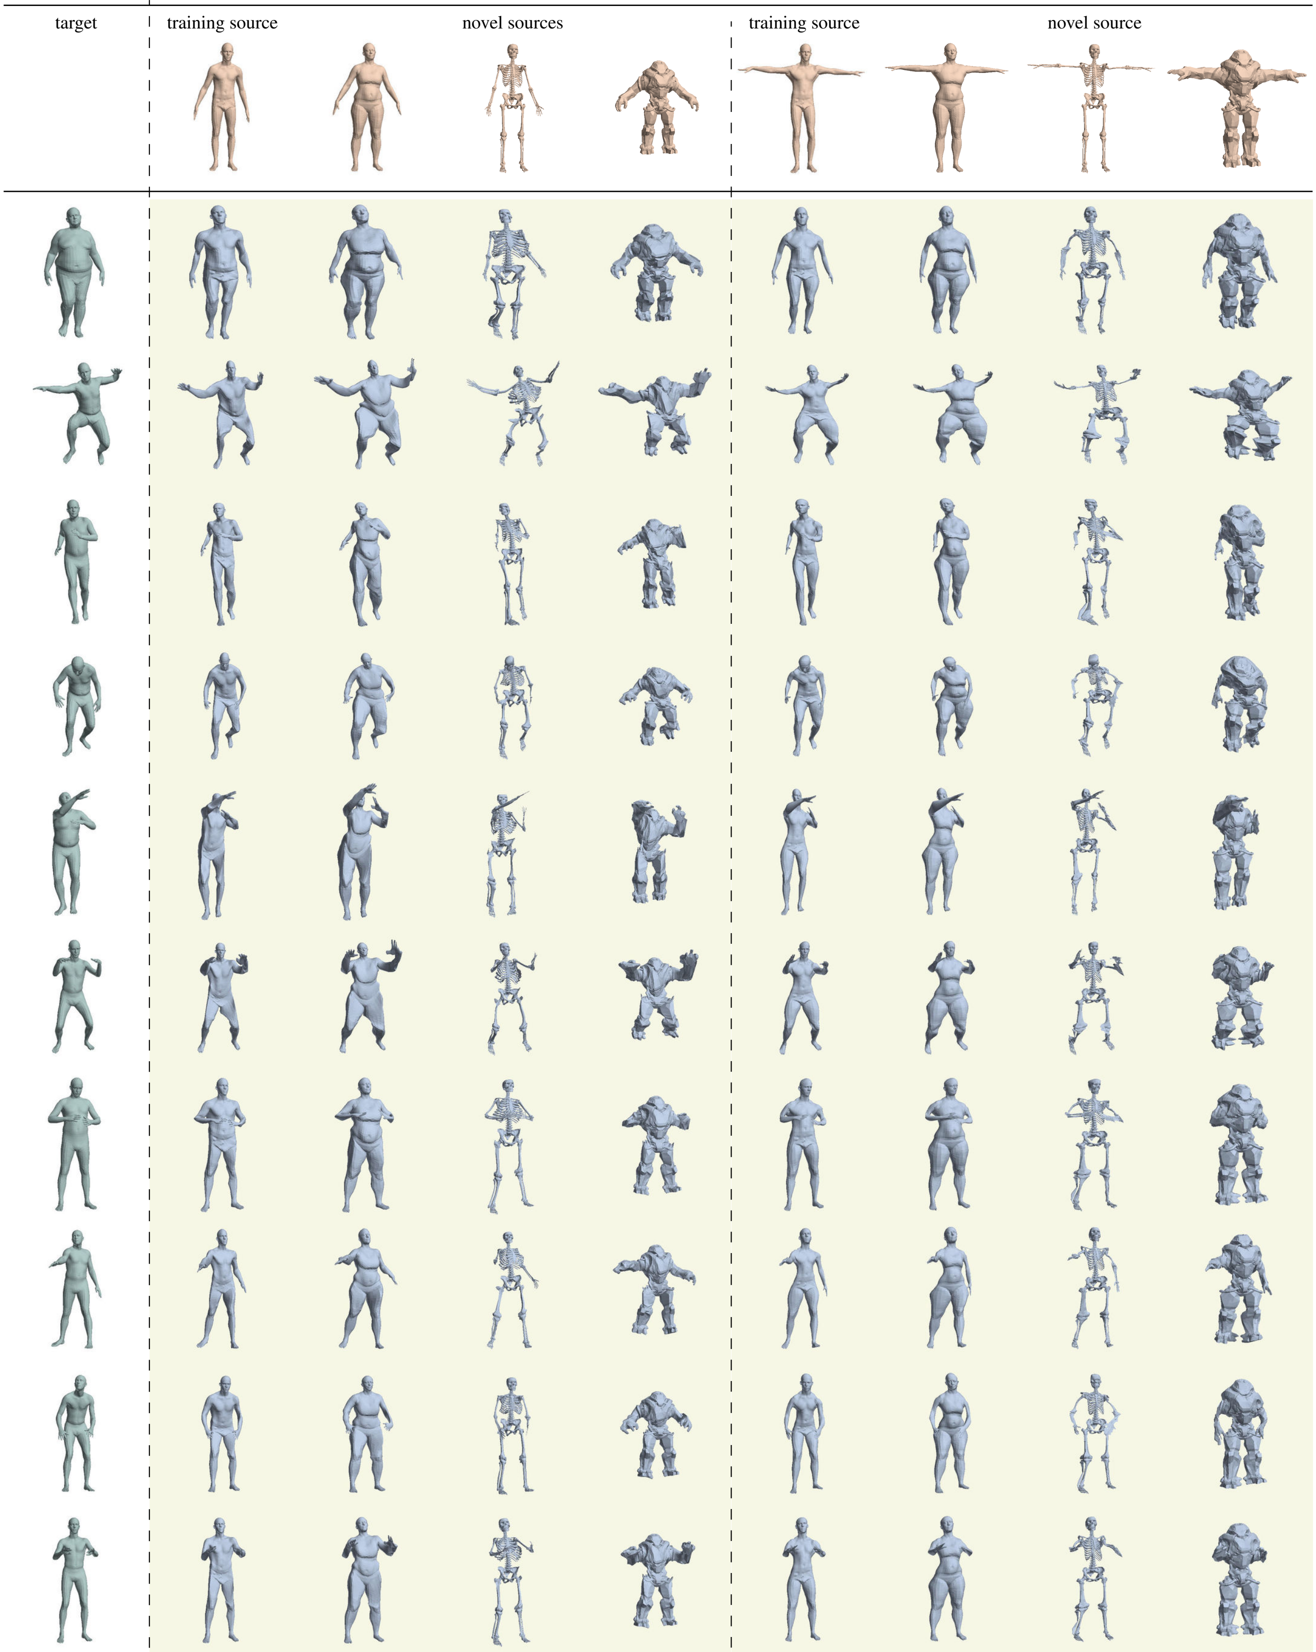

Continued on the next page

Continued from the last page

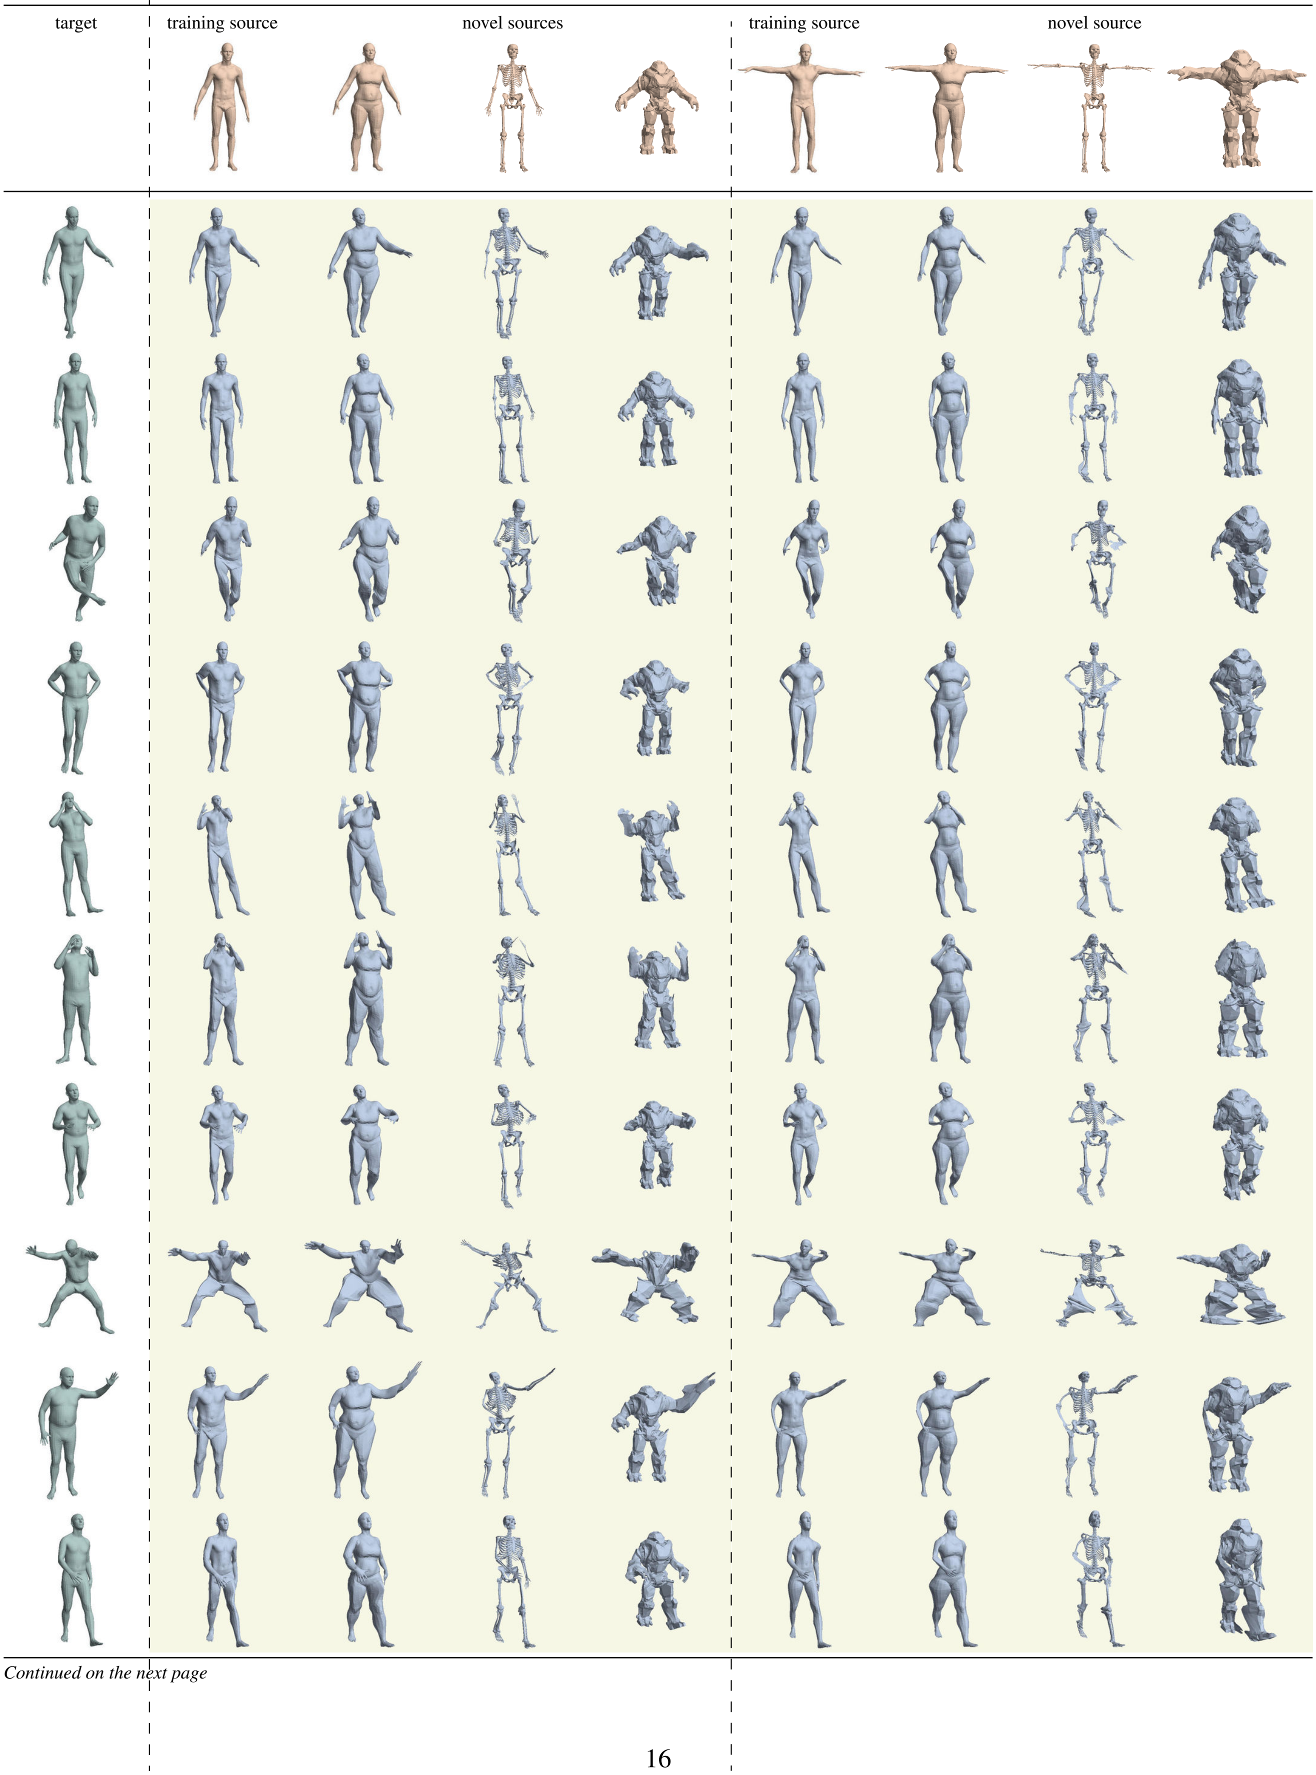

Continued on the next page

Continued from the last page

Continued on the next page

Continued from the last page

Continued on the next page

Continued from the last page

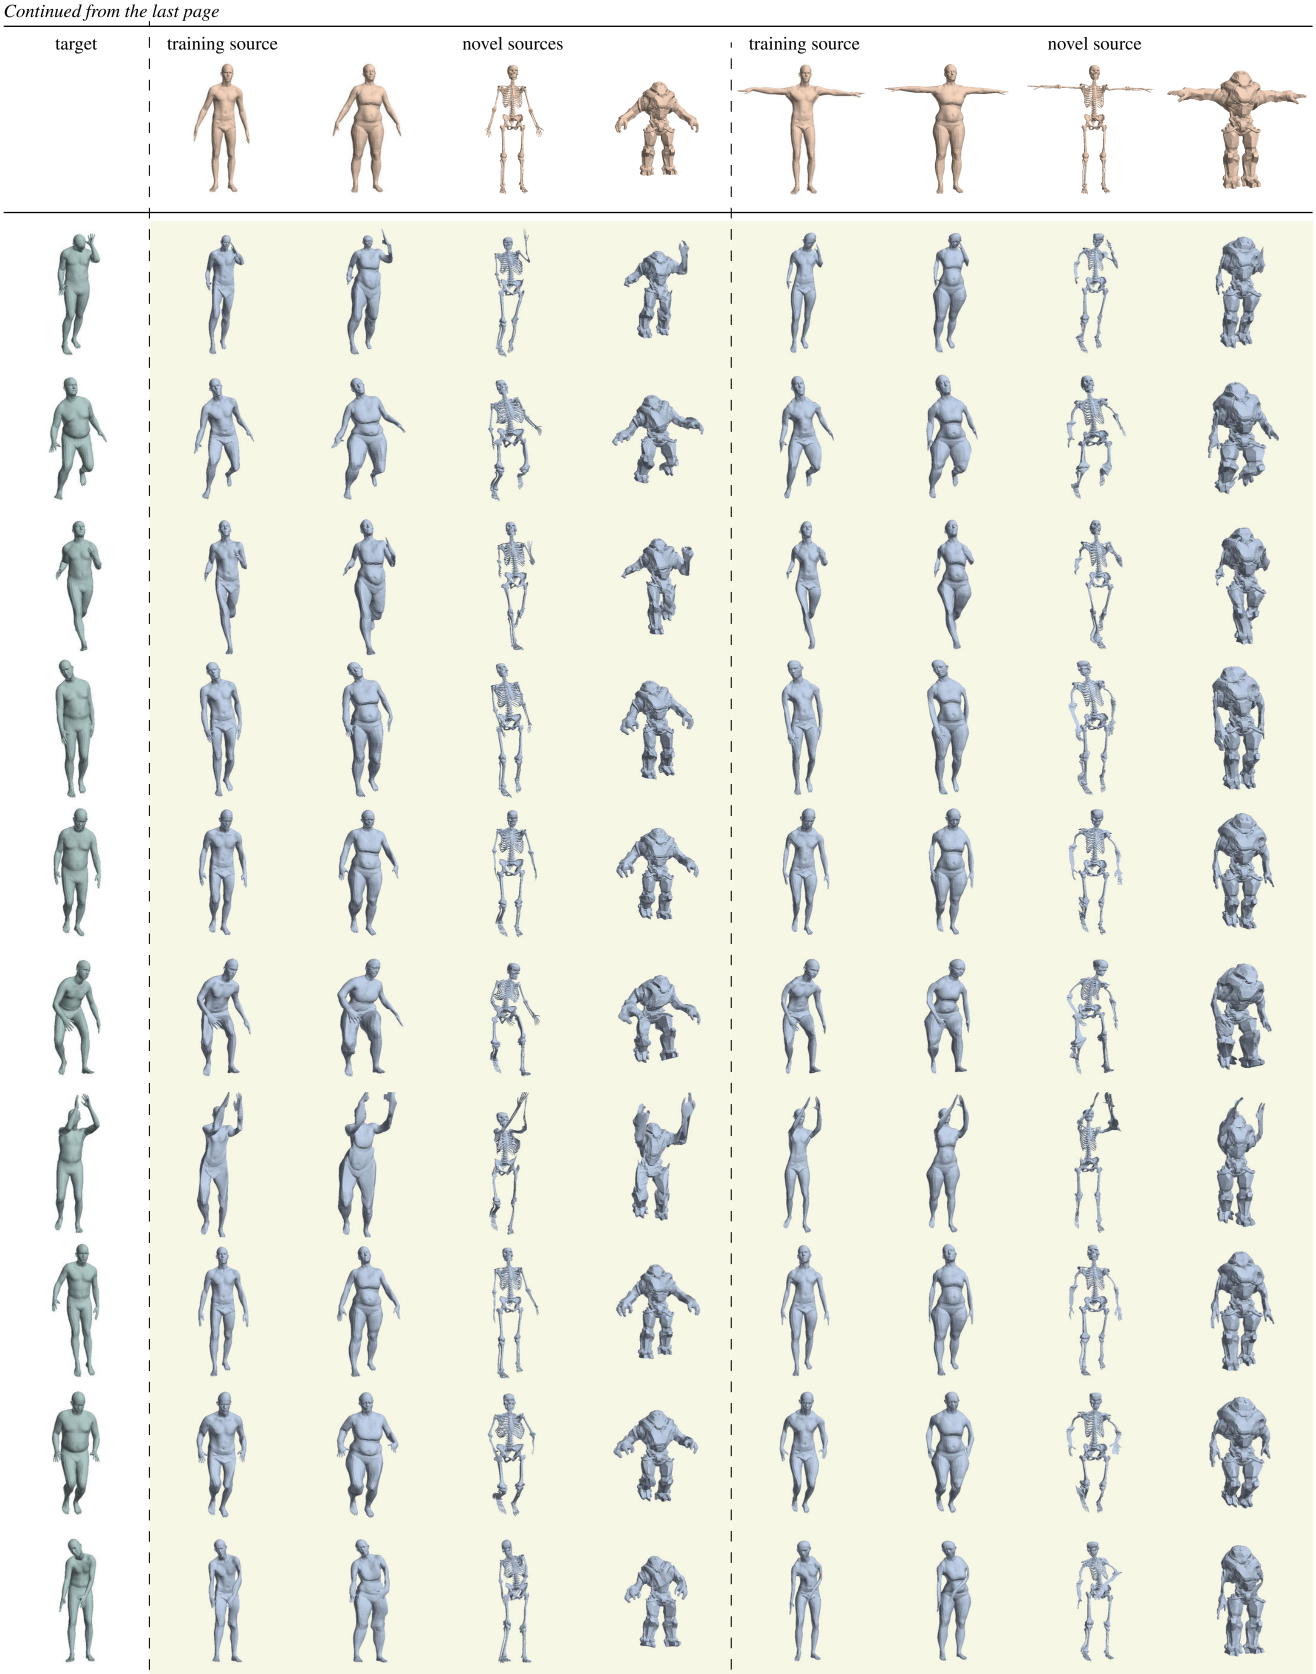

Continued on the next page

Continued from the last page

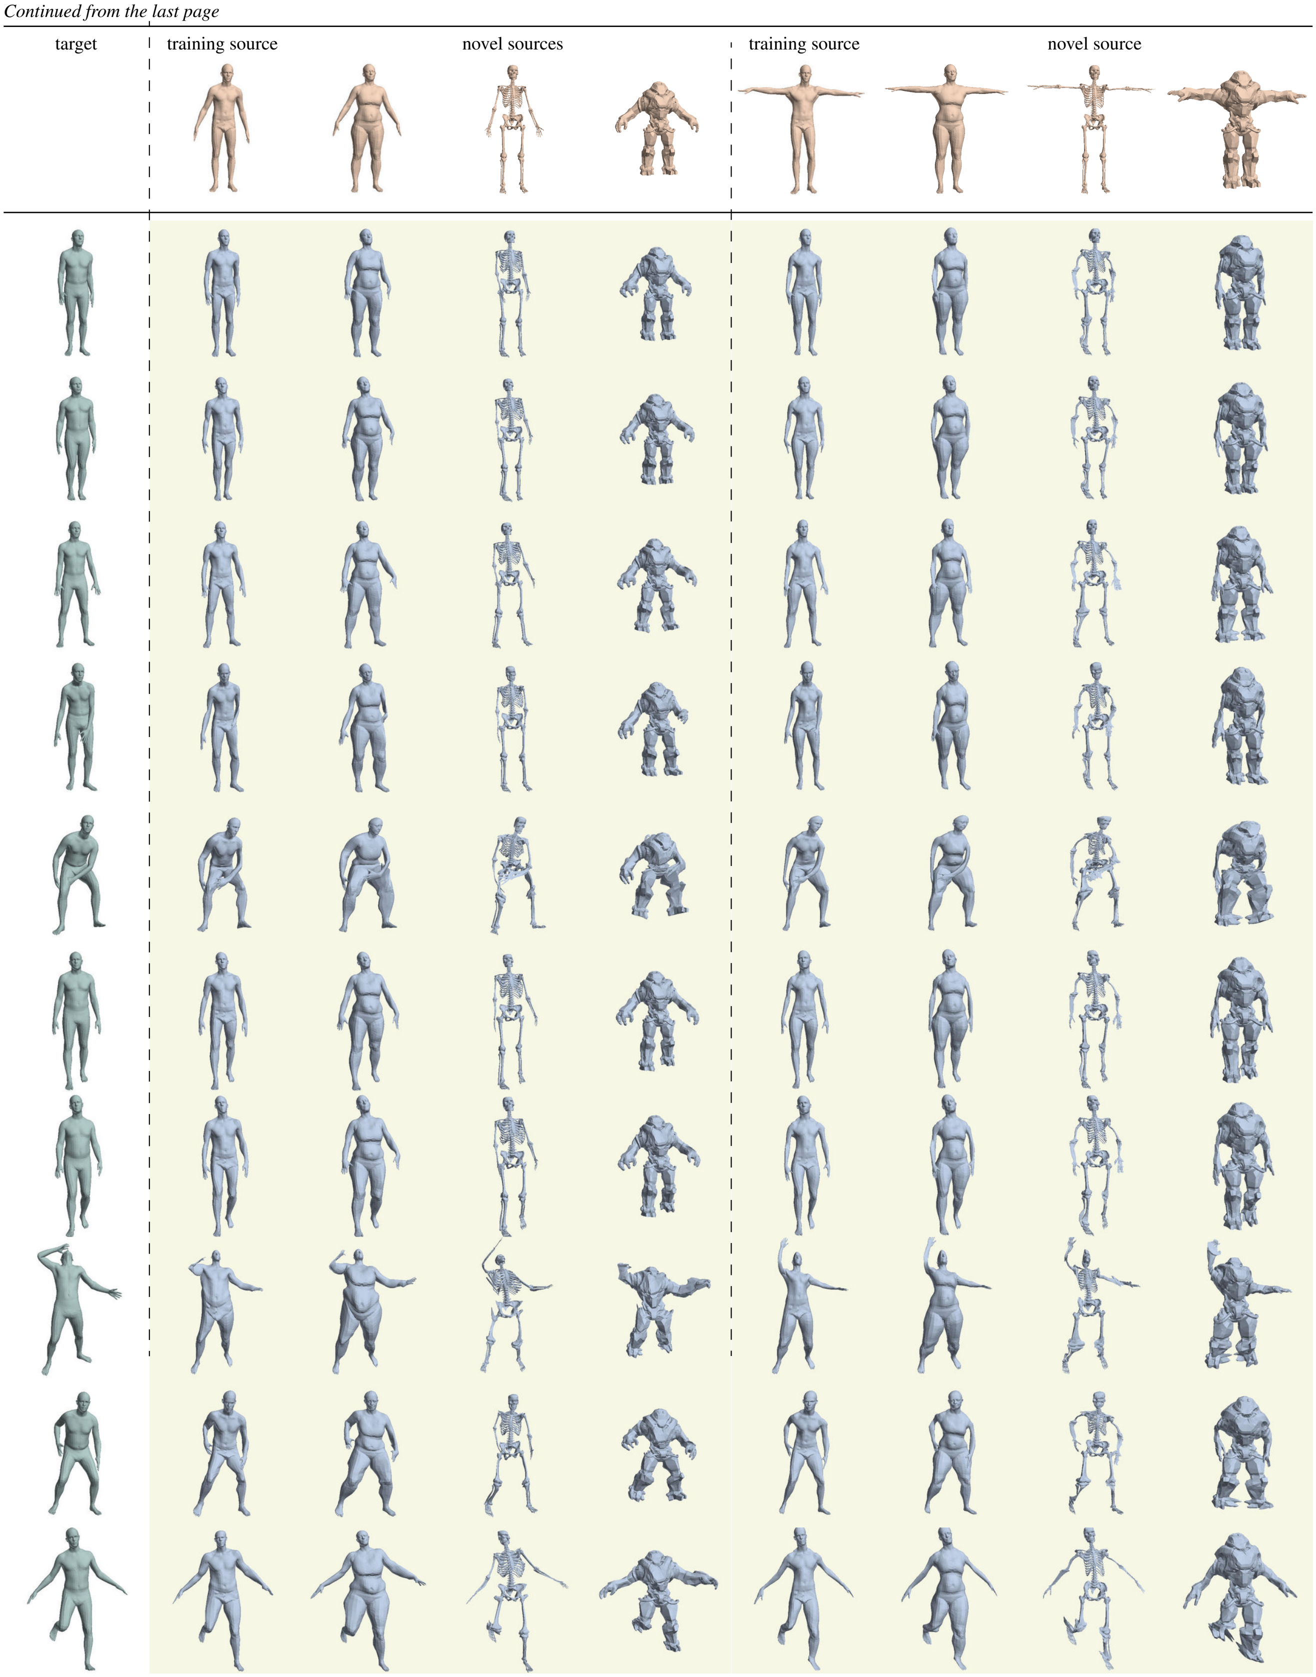

Continued on the next page

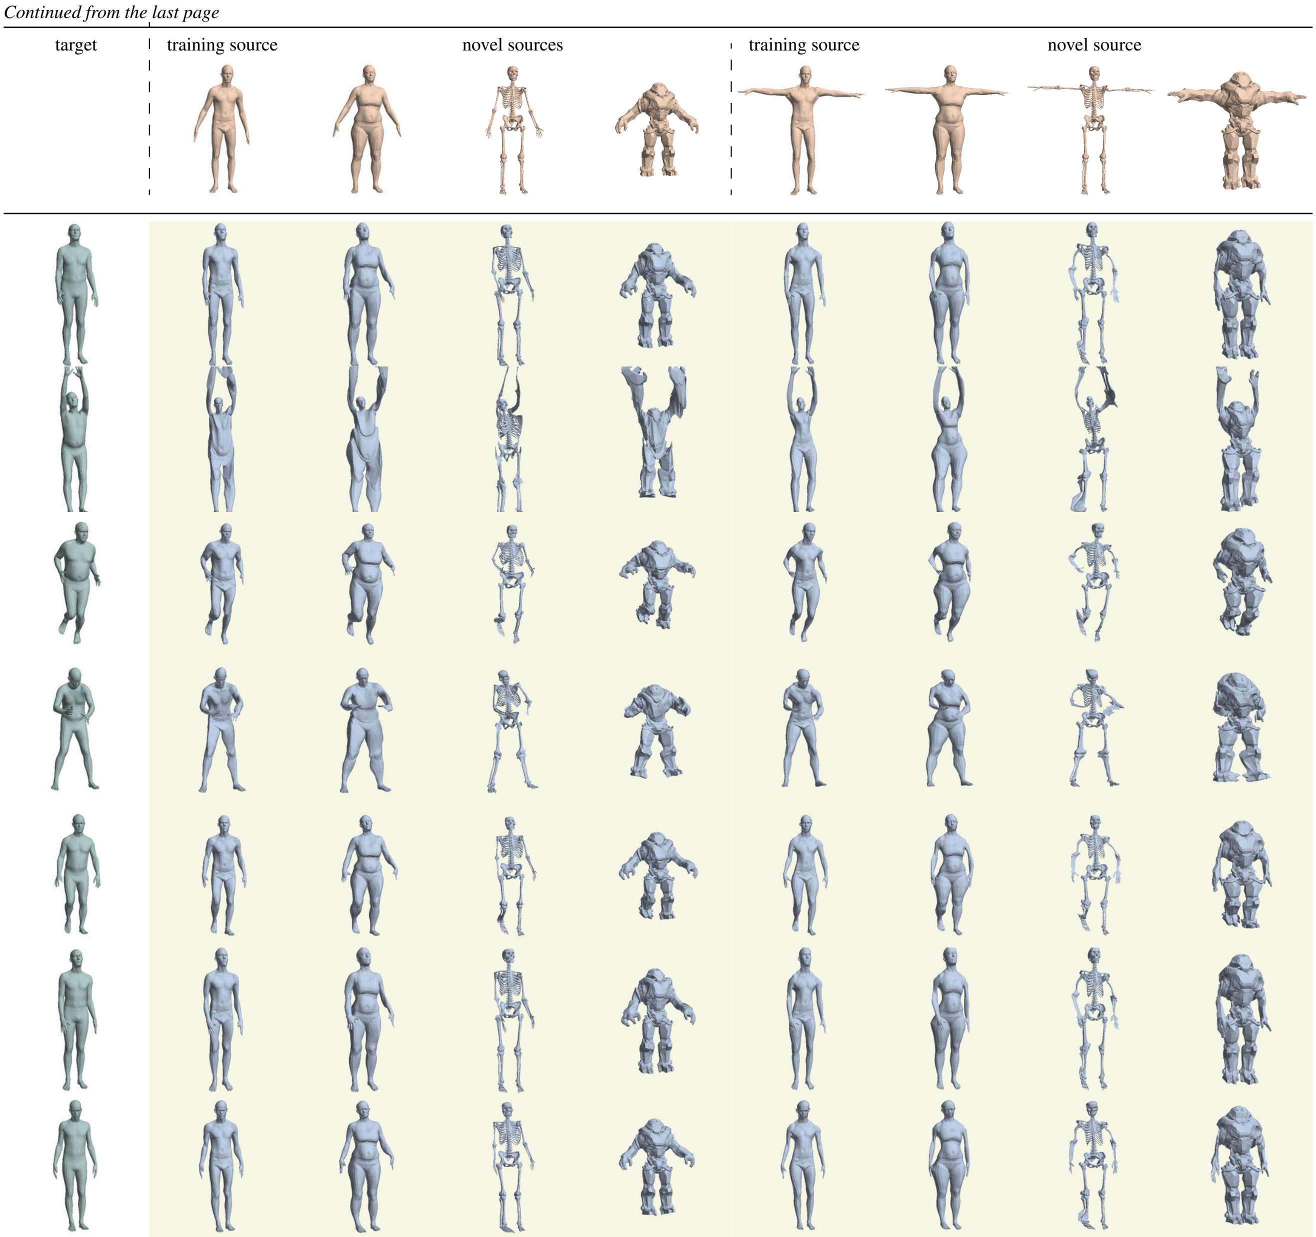

References

[1] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*, pages 3794–3801, 2014.

[2] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An information-rich 3D model repository. Technical Report arXiv:1512.03012 [cs.GR], 2015.

[3] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3D-CODED: 3D correspondences by deep deformation. In *ECCV*, 2018.