Supplementary material for the paper
StyLit: Illumination-Guided Example-Based Stylization of 3D Renderings

This document presents extra material for the main paper *StyLit: Illumination-Guided Example-Based Stylization of 3D Renderings*.

Fig. 1 illustrates the process of fitting an analytic curve to sorted source-to-target candidates as described in Section 3.3 of the main paper.

Figs. 2, 3, and 4 show the 5 light path expression (LPE) channels used to guide the synthesis of results presented in the main paper: full global illumination (first column), direct diffuse (`LDE`, second column), direct specular (`LSE`, third column), first two diffuse bounces (`LD{1,2}E`, fourth column) and diffuse interreflection (`L.*DDE`, fifth column).

Figs. 6 and 7 show additional results produced by our algorithm with the models and style exemplars given in the main paper, and also with the seven new exemplars in Fig. 5.

Fig. 8 extends the comparison with previous work (Section 4.1 in the main paper) by showing how the Kaspar et al. [2015] algorithm behaves when we modify the parameter λ, which influences the strength of enforcing uniform utilization of source patches.

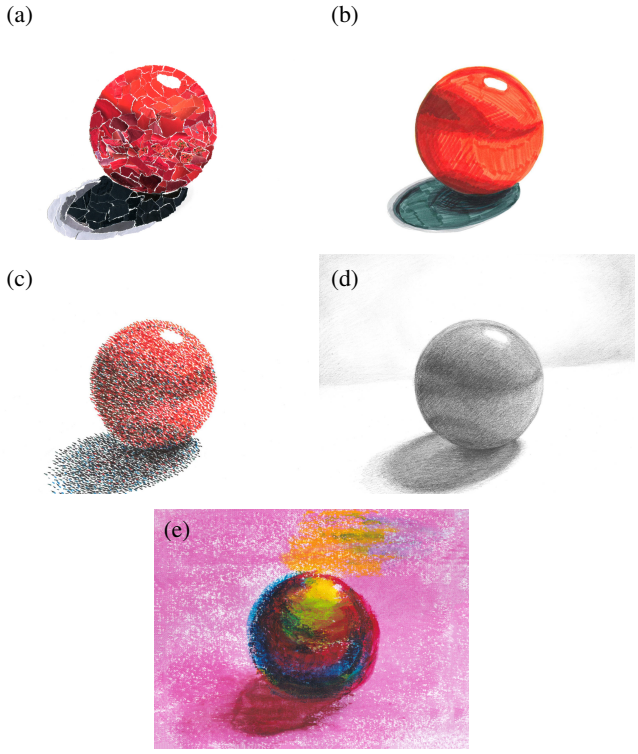

Figure 5: *Additional style exemplars: collage (a), marker (b), pointillism (c), pencil (d), oil pastel (e). Exemplar images* © Daichi Ito *(a–d) and* Pavla Sýkorová *(e).*

References

KASPAR, A., NEUBERT, B., LISCHINSKI, D., PAULY, M., AND KOPF, J. 2015. Self tuning texture optimization. *Computer Graphics Forum 34*, 2, 349–360.

WEXLER, Y., SHECHTMAN, E., AND IRANI, M. 2007. Space-time completion of video. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 3, 463–476.

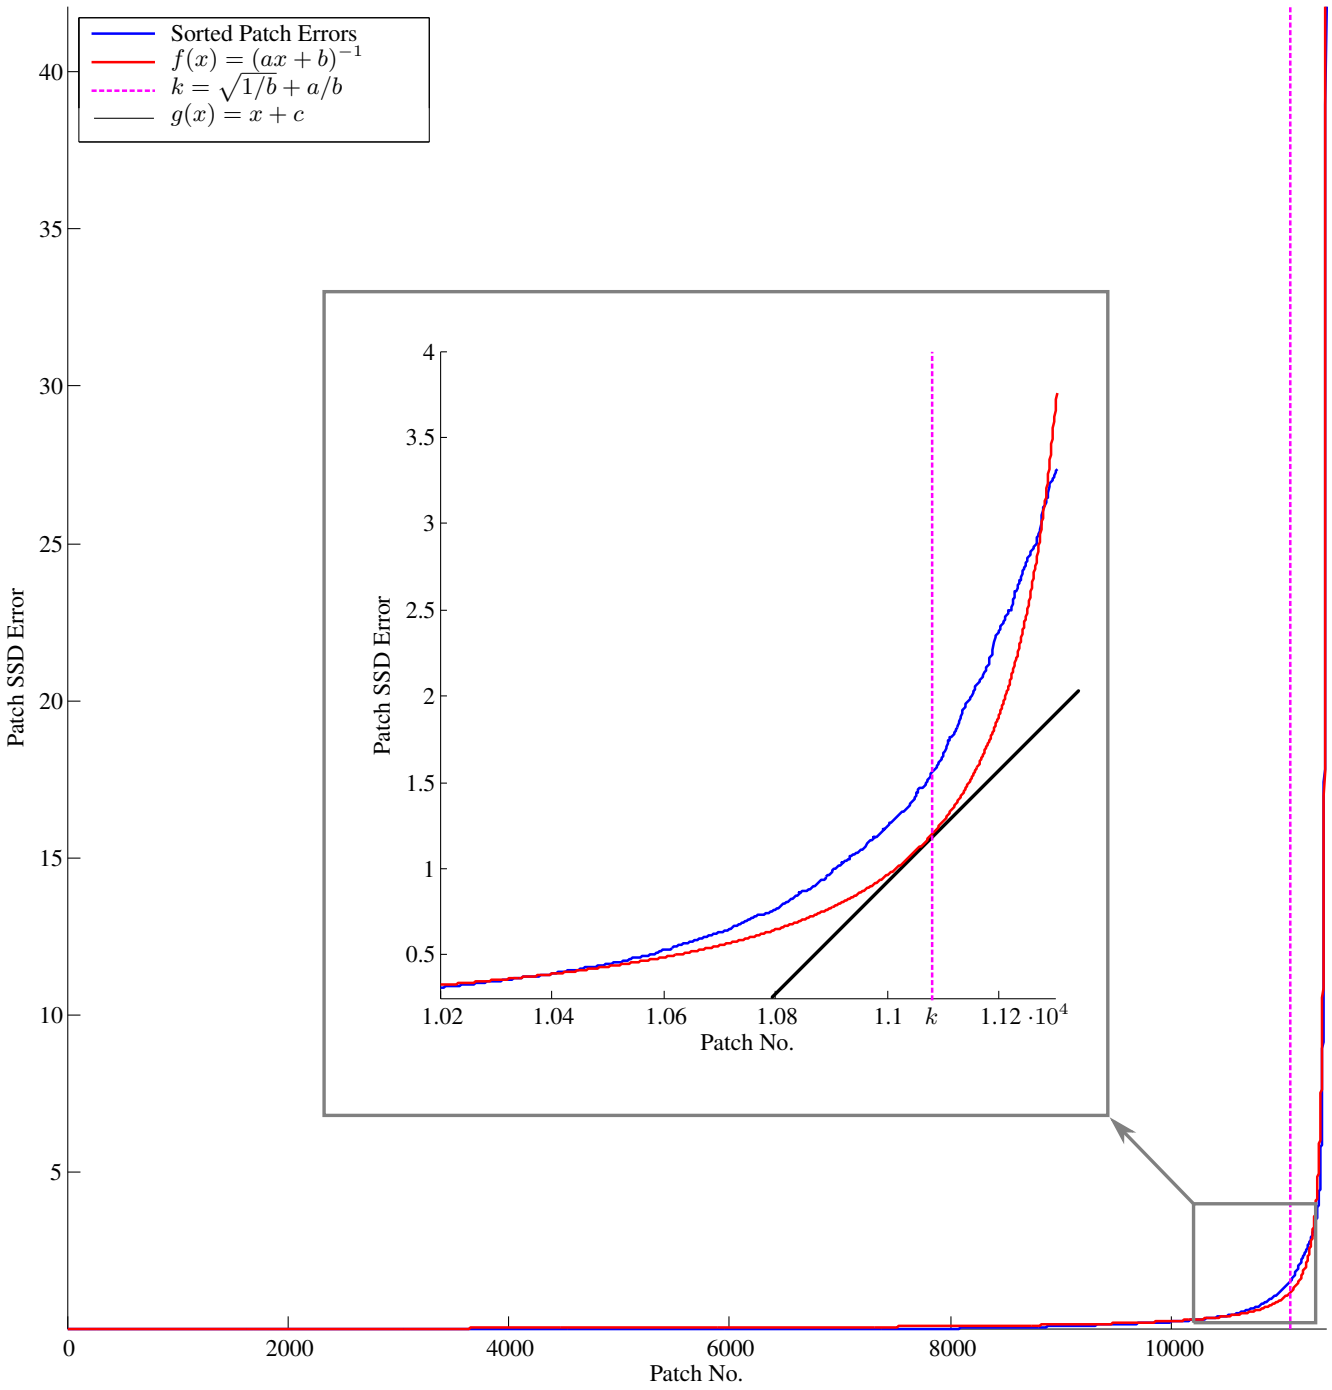

Figure 1: *Finding the feasible error budget T in source-to-target correspondences: All patch correspondences are sorted according to their matching error E (blue). We fit a hyperbolic function $f(x) = (ax + b)^{-1}$ (red) to these values and find the patch k where $f'(k) = 1$ (magenta line). The sum of all patch errors up to the k-th patch represents the error budget T.*

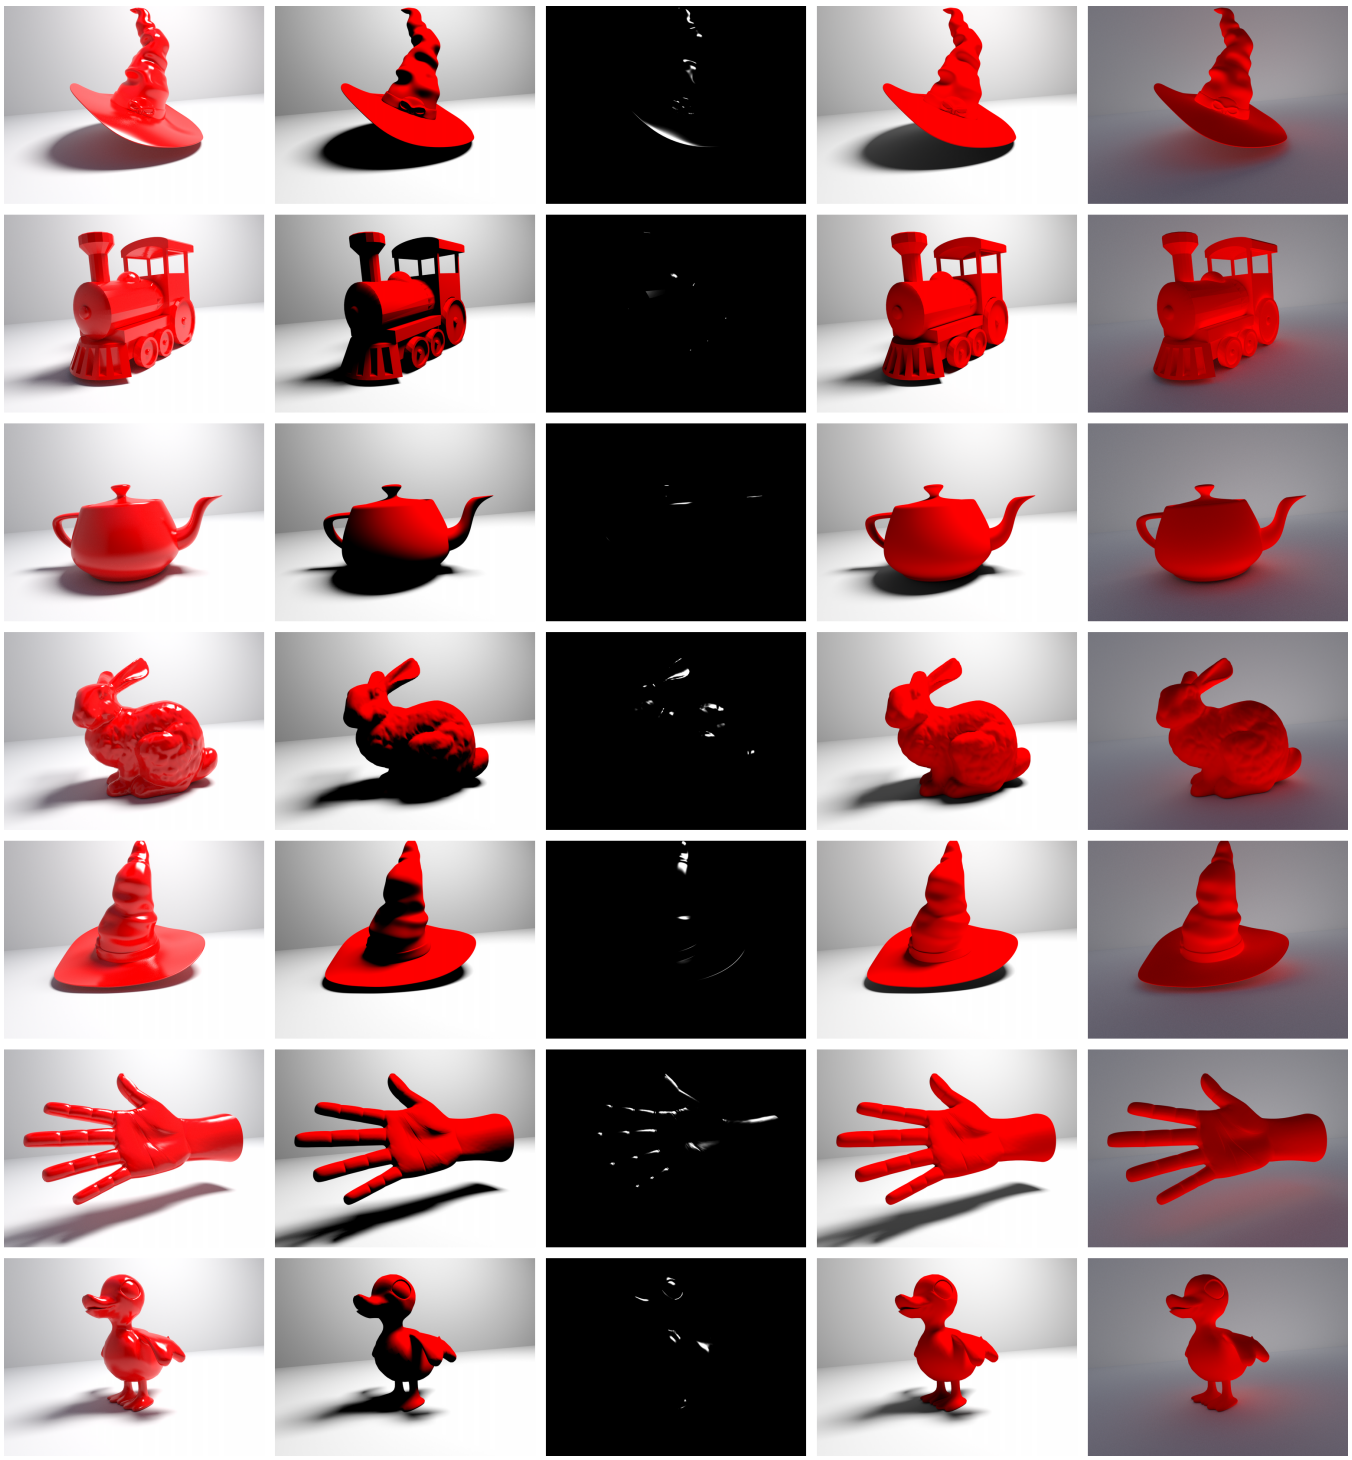

Figure 2: *LPE guiding channels. Columns from left ro right: full global illumination, direct diffuse (*`LDE`*), direct specular (*`LSE`*), first two diffuse bounces (*`LD{1,2}E`*) and diffuse interreflection (*`L.*DDE`*).*

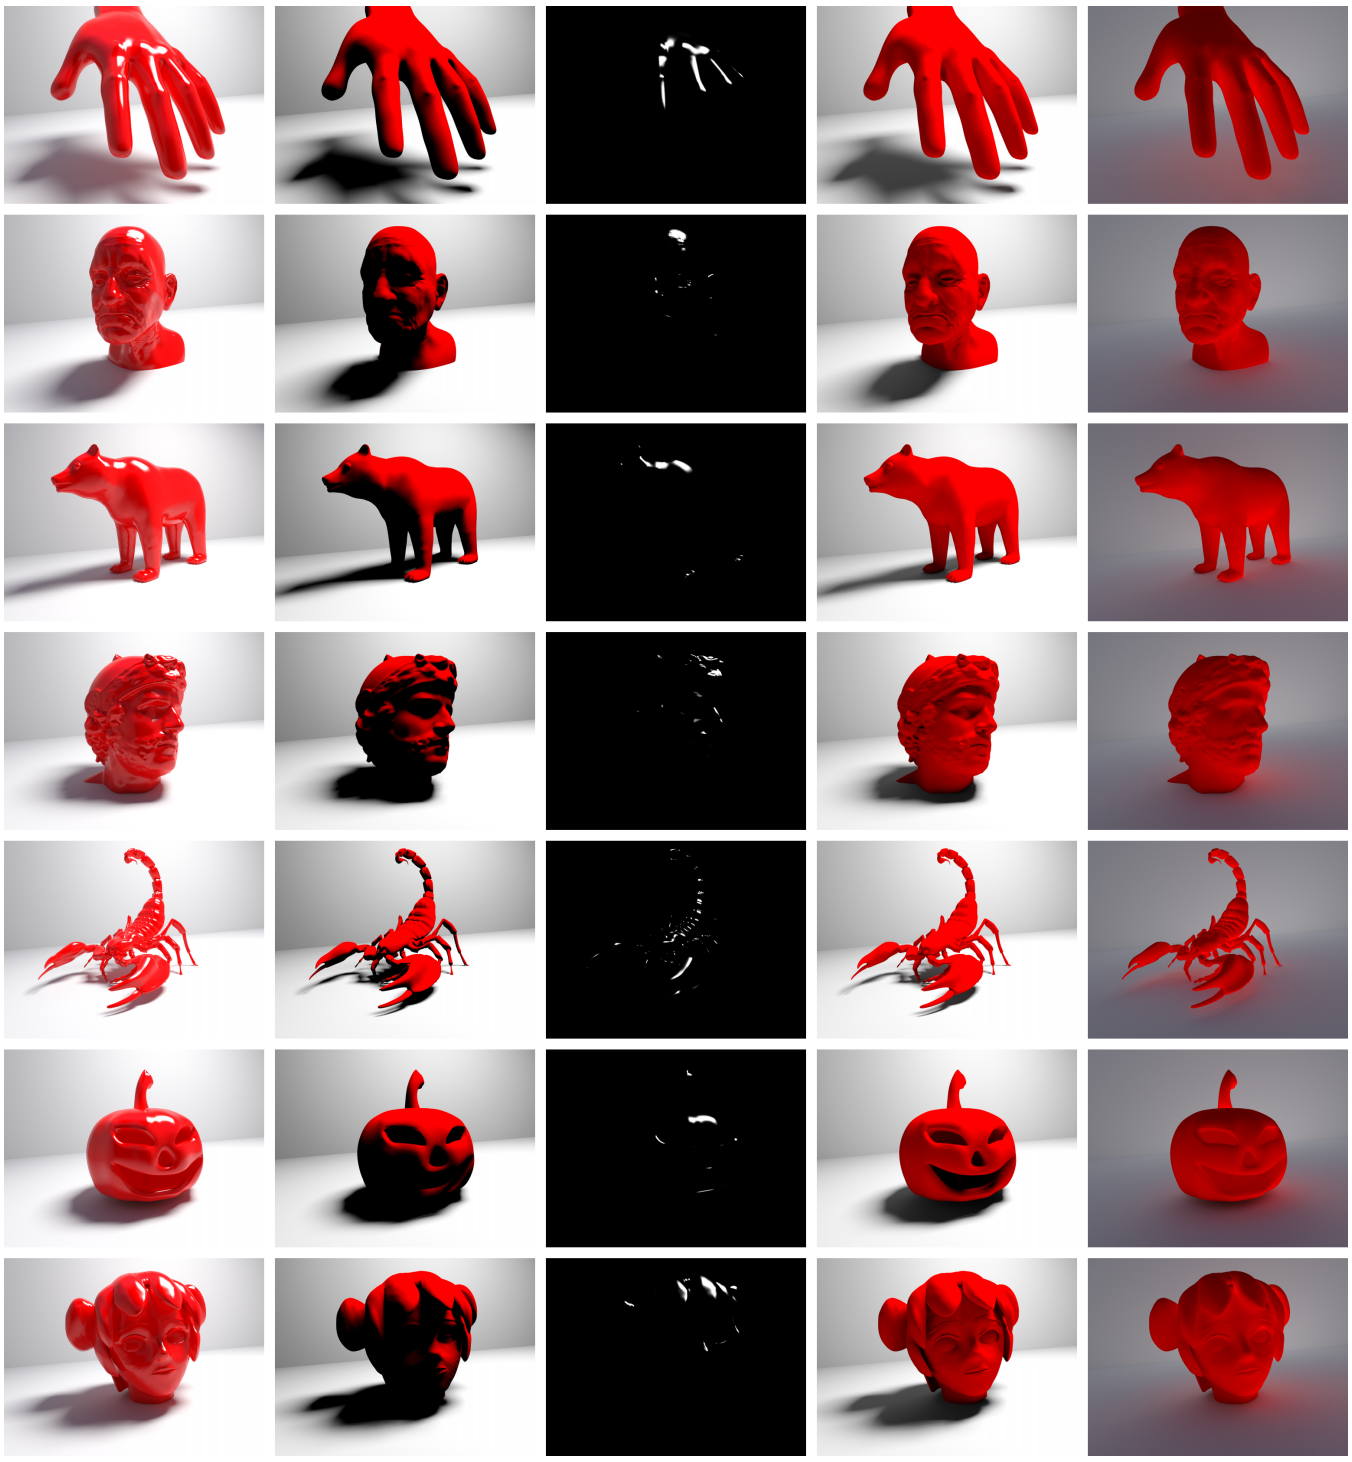

Figure 3: *LPE guiding channels. Columns from left ro right: full global illumination, direct diffuse (*`LDE`*), direct specular (*`LSE`*), first two diffuse bounces (*`LD{1,2}E`*) and diffuse interreflection (*`L.*DDE`*).*

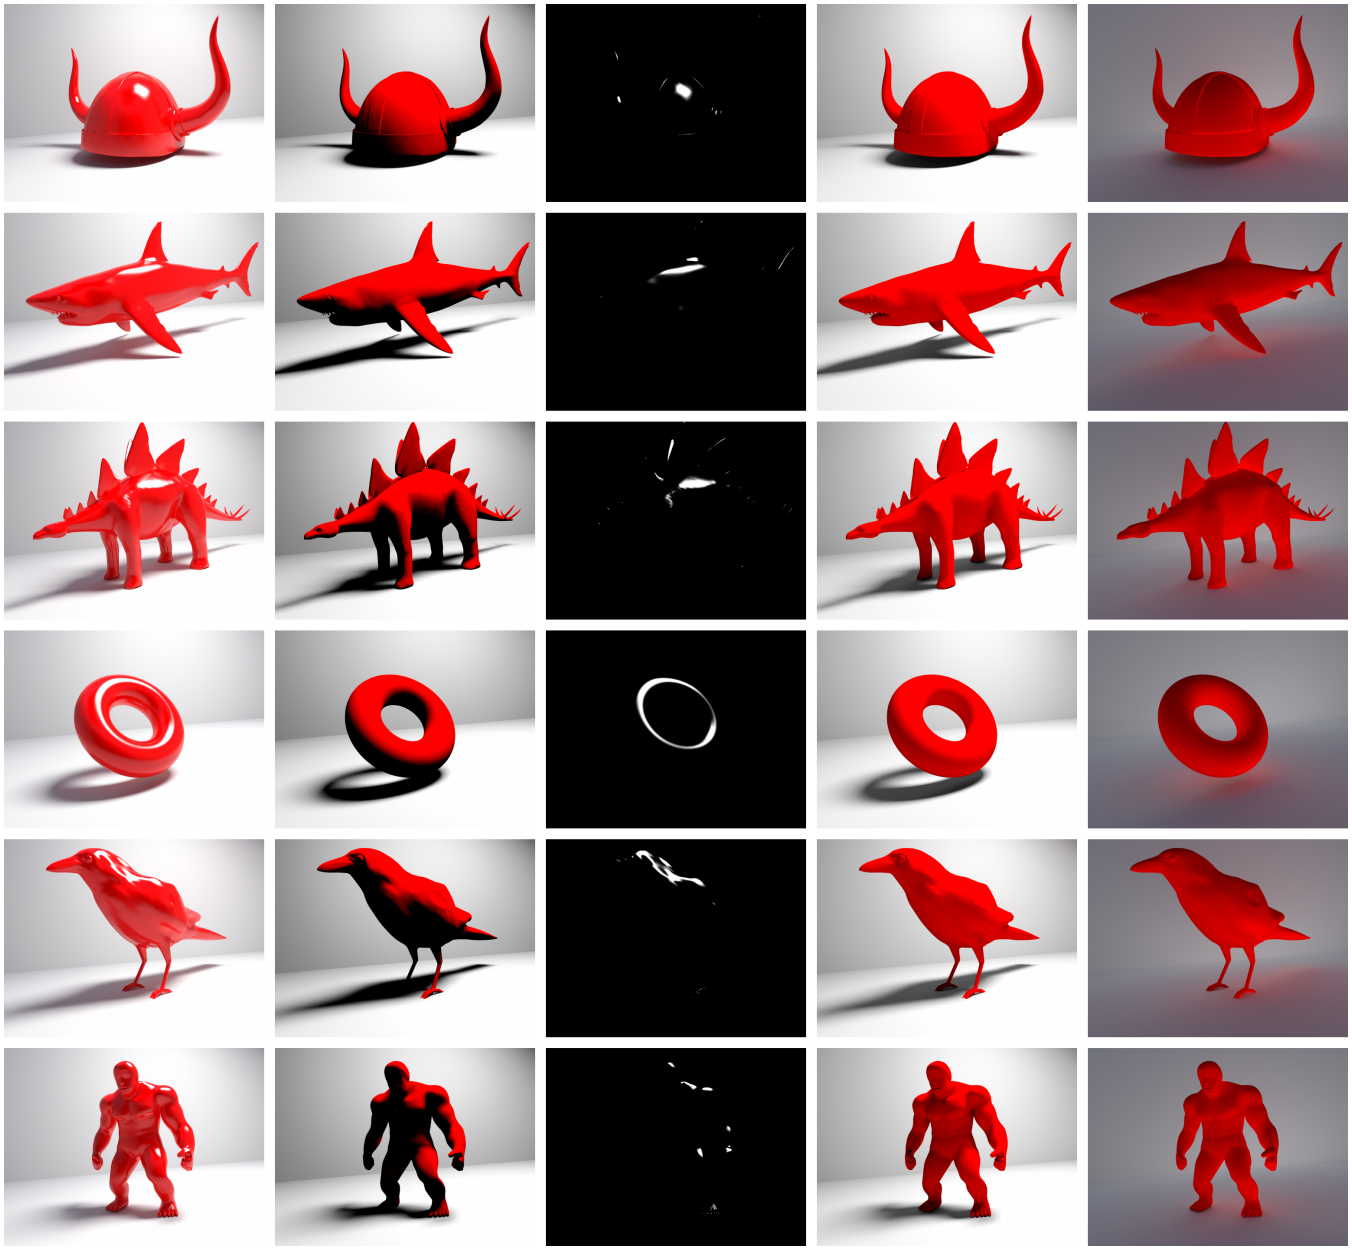

Figure 4: *LPE guiding channels. Columns from left ro right: full global illumination, direct diffuse (`LDE`), direct specular (`LSE`), first two diffuse bounces (`LD{1,2}E`) and diffuse interreflection (`L.*DDE`).*

Figure 6: *Additional results #1.*

Figure 7: *Additional results #2.*

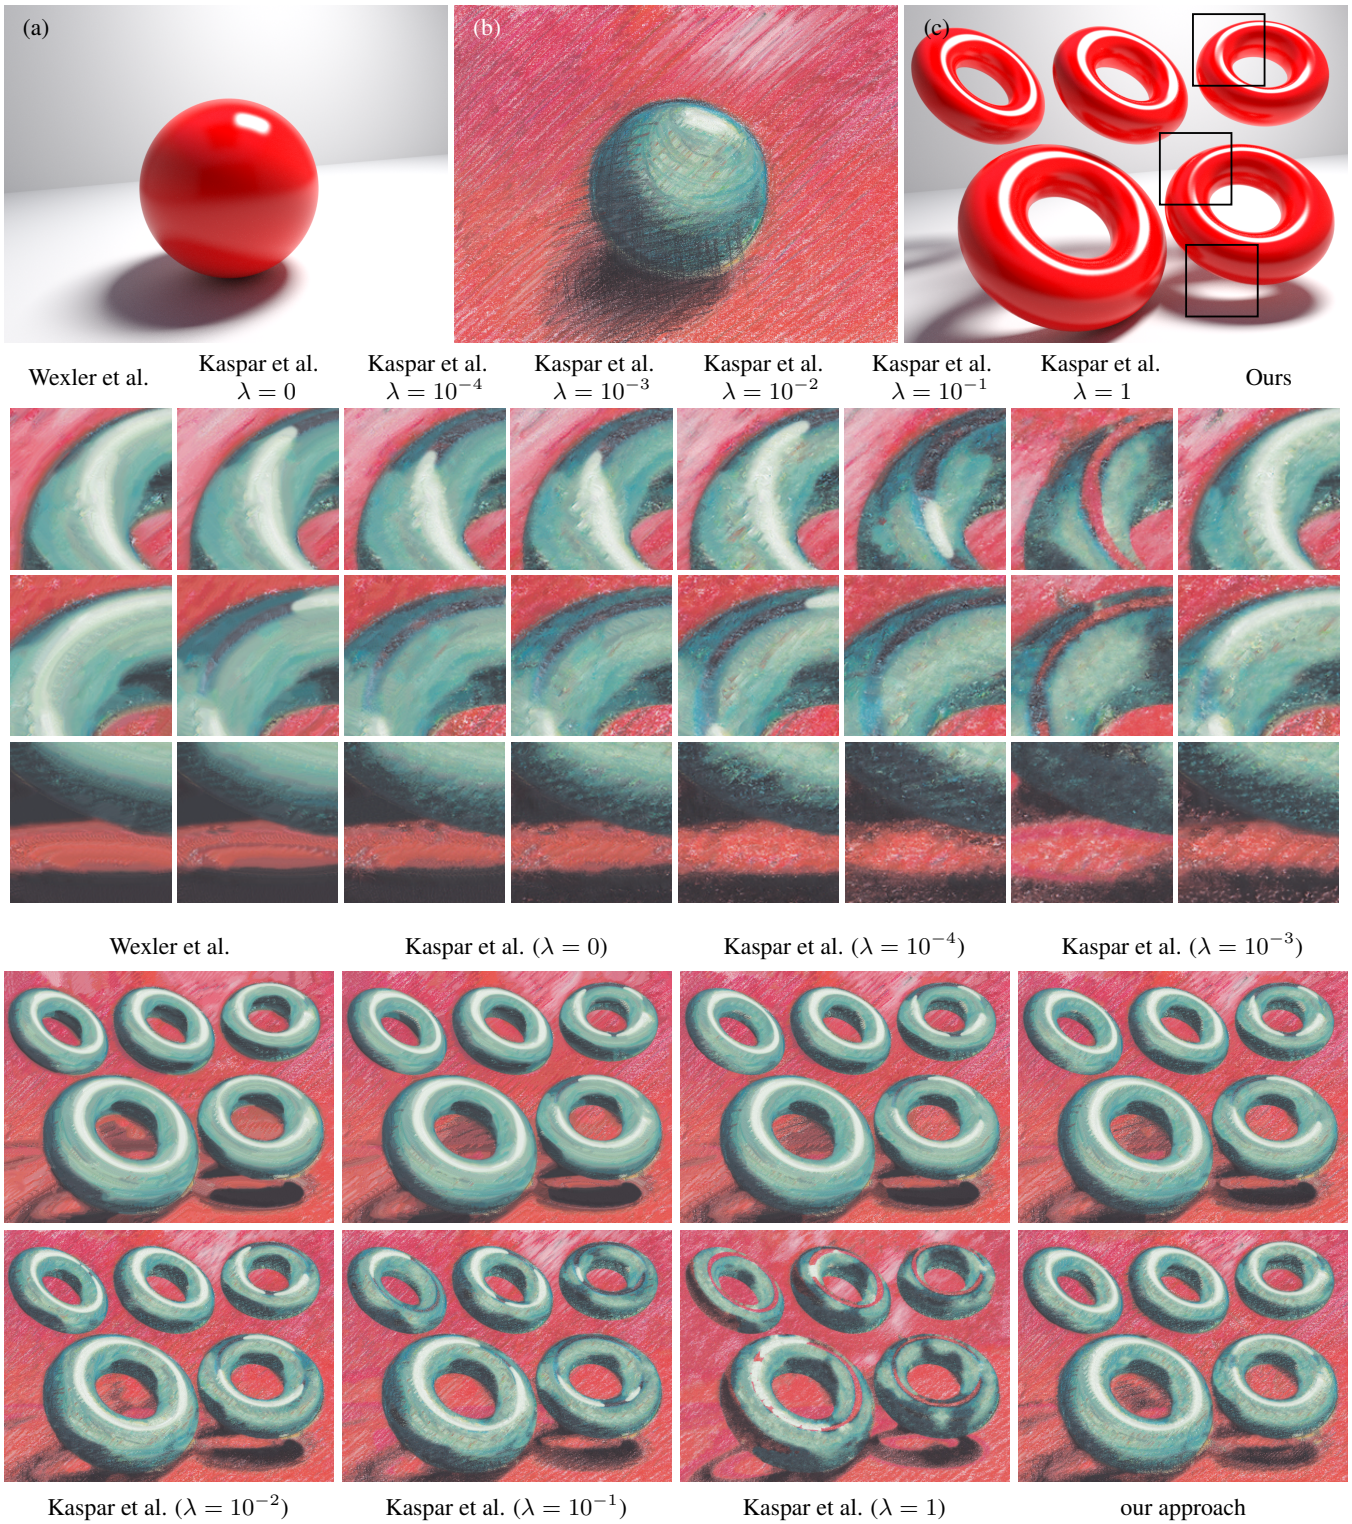

Figure 8: *Effect of the λ parameter from Kaspar et al. [2015]. For lower values, the term that encourages uniform utilization of source patches is suppressed making the result similar to that of Wexler et al. [2007], with an obvious wash-out effect. For higher values, uniformity encouragement starts to dominate, forcing source patches into inappropriate locations in the target. These two opposing effects make it hard or even impossible to find an optimal setting of λ for Kaspar et al.'s algorithm, whereas our method suppresses both simultaneously.*