

Navigating the Manifold of Translucent Appearance – Supplemental Material

Dario Lanza , Belen Masia , and Adrian Jarabo 

Universidad de Zaragoza - I3A, Spain

The supplemental material of this paper includes:

- A supplementary video demonstrating our interface.
- This pdf document, offering additional information and details on the following topics:
 - (S1) Additional details on the perceptual study of distance metrics for translucency appearance
 - (S2) Details on stimuli generation
 - (S3) Additional results of the matching task
 - (S4) Additional results of the natural task

1. User Study

1.1. Metrics Definition

In the following, we provide a small introduction to the various metrics evaluated during the perceptual study for an image-based metric.

Cubic root metric One of the most thorough studies of translucent appearance, by Gkioulekas et al. [GXZ*13], tested different simple image-based metrics over a variety of geometries and material properties (a total of 6,777 rendered images), and showed that the so-called *cubic root metric* performed best in terms of consistency between geometries, and correlated well with human perception. The cubic root metric was introduced by Ngan et al. [NDM06] in the context of opaque BRDFs, and is defined as

$$d(\mathbf{i}_i, \mathbf{i}_j) = \sqrt{\sum_{k=1}^N \left(\sqrt[3]{i_k} - \sqrt[3]{j_k} \right)^2}, \quad (1)$$

where i_k, j_k are pixels k of images \mathbf{i}_i and \mathbf{i}_j respectively, in linear color space, with N is the total number of pixels.

Anisoshading Metric In our analysis, we also include a recently developed model for translucent materials [KNYK22]. The key idea behind this model is that translucency perception is influenced by the incongruence in 3D shape information provided by specular and non-specular shading patterns. Through psychophysical studies, they show how the difference in the anisotropy of the shading between specular and diffuse image components can be a plausible predictor of perceived translucency.

Image quality metrics In addition to these two, we also include in our analysis the well-known MS-SSIM, widely used for image comparison, and the recent FovVideoVDP metric [MCR*21]. While FovVideoVDP is designed as a video metric, with emphasis on the temporal and peripheral aspects of vision, benchmarking results show that it is among the best performing also on conventional SDR and HDR images. We choose these two metrics because, in the recent Unified Photometric Image Quality (UPIQ) dataset [MPY*20], their predictions yield the highest (FovVideoVDP) and second highest (MS-SSIM) Spearman rank order correlation with human judgements [MCR*21]. Finally, we also include the deep learning-based LPIPS metrics [ZIE*18] in our work, as one of the most used metrics in perception studies.

1.2. Perceptual Study

To increase the robustness of the subjective judgments gathered, we resort to a two-alternative forced choice (2AFC) experimental paradigm [LMS*19; WAKB09; GXZ*13]. In each experimental trial, the participant is shown a triplet of images: one *reference* and two *candidates*. The participant is then asked to select the candidate stimulus that is more similar to the reference. Figure 1 shows the graphic user interface used by participants in our perceptual study. In Figure 2 we show a subset of the stimuli used in our user study, differently from previous work [GXZ*13] we focus on a different set of appearances encompassing crystal-like, matte, and dark lustrous appearances.

Our stimuli, comprising sixteen sets of material properties, yield 1,680 possible triplets for each light condition. We show a subset of the stimuli used in Fig 2. Since we have three light conditions, a total of 5,040 triplets need to be evaluated (all three images in a triplet feature the same lighting). In the interest of robustness of the subjective data, we seek that each triplet be answered by at least five different participants, leading to 25,200 trials. We resort to crowd-sourcing (through the Amazon Mechanical Turk platform) to gather the perceptual data, which has been shown to be effective in similar studies [HB10; JWD*14; LMS*19; LBFS21; SCW*21], and include a training stage before the experiment, as well as control triplets to discard unreliable participants.

Amazon Mechanical Turk defines a single trial as a HIT (Human Intelligent Task), in our user study, each HIT is composed of 100

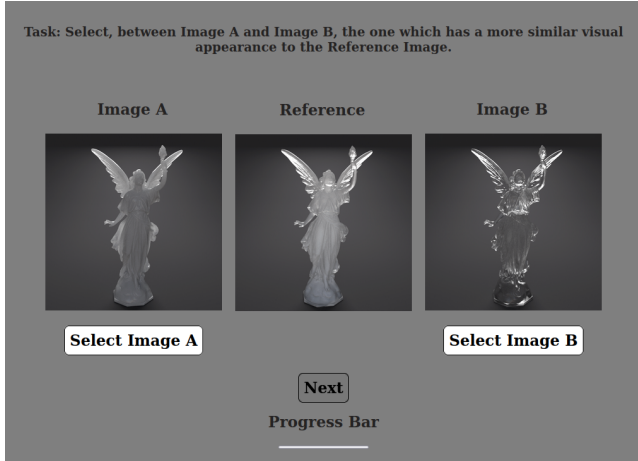


Figure 1: User interface used for the perceptual study analyzing the perceptual distance between translucent materials. Participants have to select the image that most closely resembles the Reference stimulus (in the middle). Once participants select and confirm their choice, they can pass to the next triplet.



Figure 2: We show a subset of the various appearances covered in our user study. Each column shows the same material lit under different light conditions: back (first row), side (second row), and front (third row). The triplet showed in the user interface uses different materials with the same light conditions.

triplets, along which we also add ten control triplets. The purpose of a control triplet is to check whether a participant can be trusted or not by posing a triplet with a clear answer, in this case, the reference stimulus is used also as one of the two candidates. Following pilot studies, in which the control triplets were always passed easily, we settled to discard all the data from an HIT if more than two control triplets are failed. For each successful HIT, we pay a single participant 1.88 euros, since we estimated the completion time of a single HIT to be around fifteen minutes, following the minimum wage of 7.5 euro per hour.

2. Manifold - Additional Results

In this section, we showcase additional results when we generate the manifold. Figures 3 and 4 show the same manifold presented in the main paper with a side or a front lit illumination.




Figure 3: *Manifold obtained using side-lit stimuli*


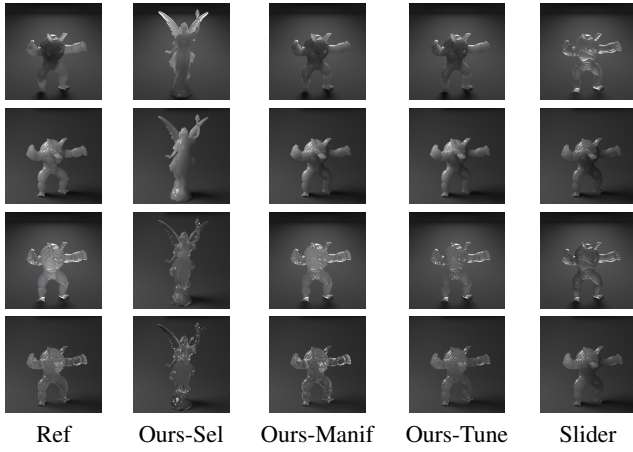



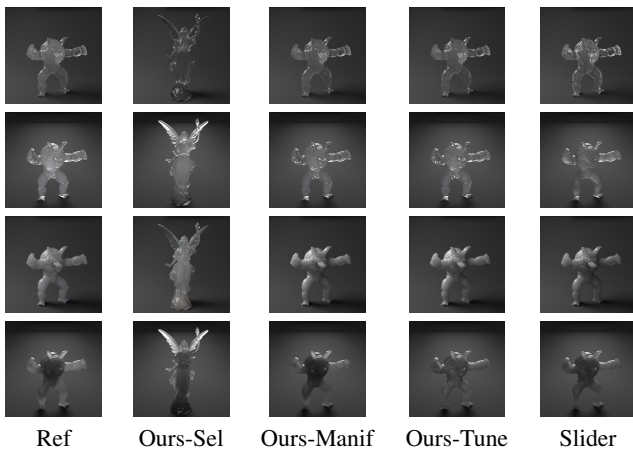




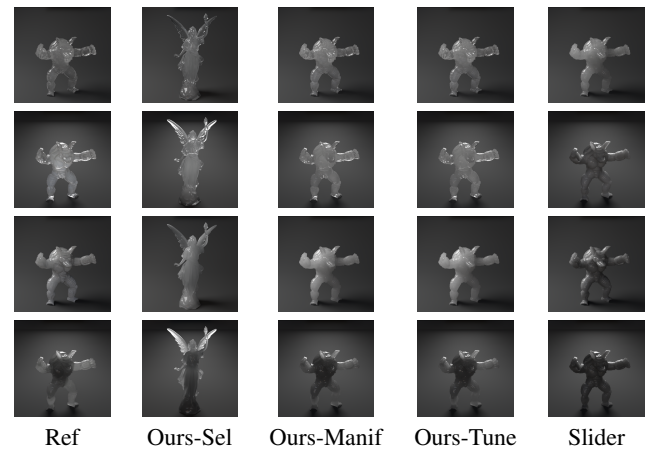
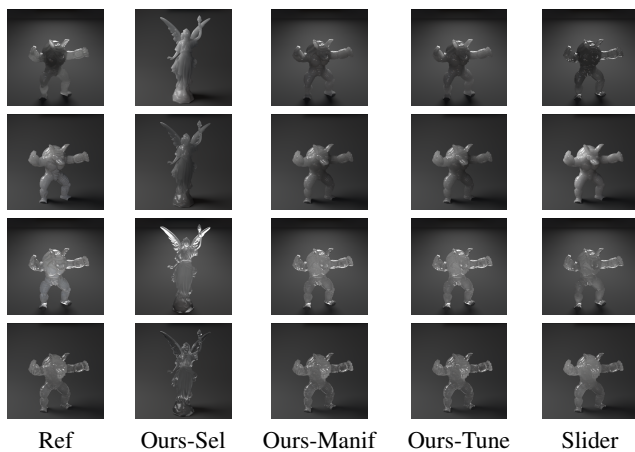



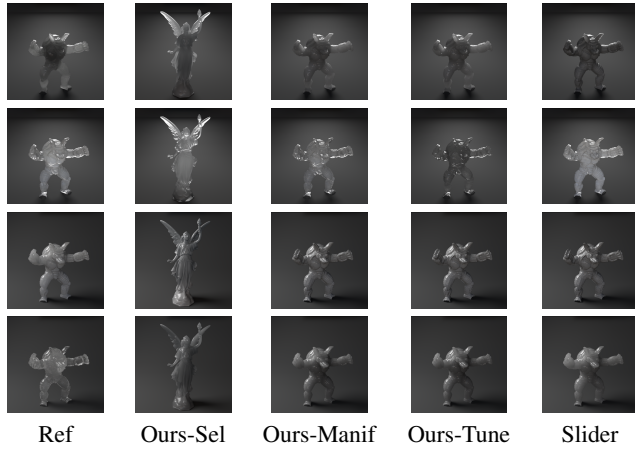
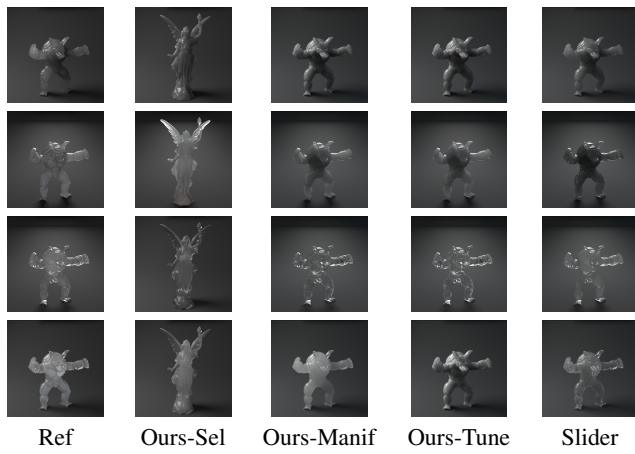

Figure 4: *Manifold obtained using front-lit stimuli*

3. Matching Task: Users' Results

In this section, we report the results of our user study. Each table covers all the results a single user obtained throughout the experiment. For each table, we report the Reference target (Ref) the appearance selected during the navigation phase (Ours - Selected) and the renderings obtained using the inverse rendering parameters (Ours - IR). We then show the result of the fine-tuning phase (Ours-Tune). This step was optional and therefore not all users desired to fine-tune their results, for these cases, we report the rendered image obtained using the optical parameters using inverse rendering. The last column shows results using the traditional slider interface (Slider). It is possible to see some inconsistencies between Ours-Sel and Ours-Manif, these inconsistencies happen due to the inverse rendering step that is not always capable of finding the optical parameters that can reproduce a similar appearance to the one previewed in the second column.

**Figure 5:** Results for the Matching Task for User 1**Figure 8:** Results for the Matching Task for User 4**Figure 6:** Results for the Matching Task for User 2**Figure 9:** Results for the Matching Task for User 5**Figure 7:** Results for the Matching Task for User 3**Figure 10:** Results for the Matching Task for User 6

**Figure 11:** Results for the Matching Task for User 7**Figure 14:** Results for the Matching Task for User 10**Figure 12:** Results for the Matching Task for User 8**Figure 15:** Results for the Matching Task for User 11**Figure 13:** Results for the Matching Task for User 9**Figure 16:** Results for the Matching Task for User 12

**Figure 17:** Results for the Matching Task for User 13**Figure 18:** Results for the Matching Task for User 14**Figure 19:** Results for the Matching Task for User 15




Figure 20: Results for the Natural Task for User 1

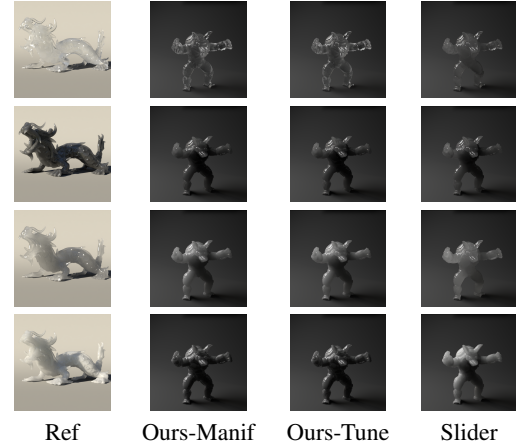


Figure 22: Results for the Natural Task for User 3




Figure 21: Results for the Natural Task for User 2




Figure 23: Results for the Natural Task for User 4

4. Natural Task: Users' Results

In this section, we report the results of our user study. Each table covers all the results a single user obtained throughout the experiment. For each table, we report the Reference target (Ref) inverse rendering parameters obtained by navigating the manifold (Ours - Manif). We then show the result of the fine-tuning phase (Ours-Tune). This step was optional and therefore not all users desired to fine-tune their results, for these cases, we report the rendered image obtained using the optical parameters using inverse rendering. The last column shows results using the traditional slider interface (Slider). The task was very challenging, especially for novice users, and in many instances, the results of using the two interfaces are equally distant from the reference.

References

[GXZ*13] GKIOULEKAS, IOANNIS, XIAO, BEI, ZHAO, SHUANG, et al. "Understanding the Role of Phase Function in Translucent Appearance". *ACM Trans. Graph.* 32.5 (2013) 1.

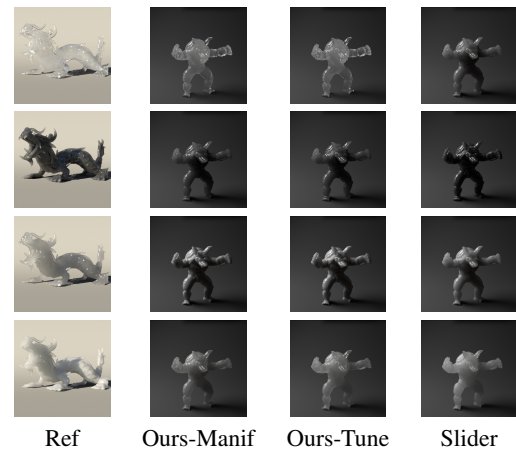
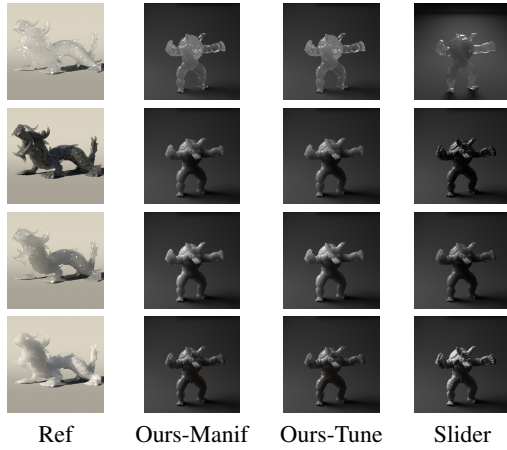


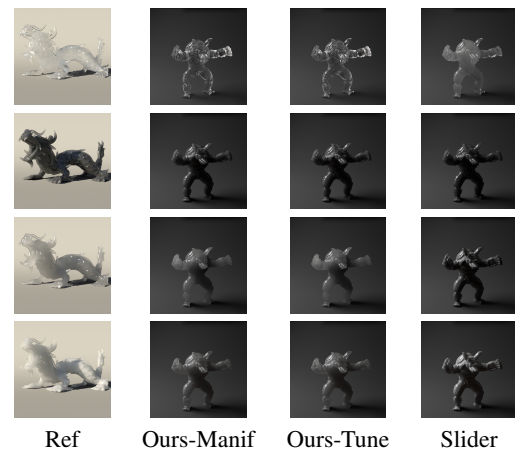



Figure 24: Results for the Natural Task for User 5

**Figure 25:** Results for the Natural Task for User 6**Figure 28:** Results for the Natural Task for User 9**Figure 26:** Results for the Natural Task for User 7**Figure 29:** Results for the Natural Task for User 10**Figure 27:** Results for the Natural Task for User 8

- [HB10] HEER, JEFFREY and BOSTOCK, MICHAEL. “Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA, 2010, 203–212 [1](#).
- [JWD*14] JARABO, ADRIAN, WU, HONGZHI, DORSEY, JULIE, et al. “Effects of approximate filtering on the appearance of bidirectional texture functions”. *IEEE Transactions on Visualization and Computer Graphics* 20.6 (2014), 880–892 [1](#).
- [KNYK22] KIYOKAWA, HIROAKI, NAGAI, TAKEHIRO, YAMAUCHI, YASUKI, and KIM, JUNO. “The perception of translucency from surface gloss”. *Vision Research* (2022), 108140 [1](#).
- [LBFS21] LAVOUÉ, GUILLAUME, BONNEEL, NICOLAS, FARRUGIA, JEAN-PHILIPPE, and SOLER, CYRIL. “Perceptual quality of BRDF approximations: dataset and metrics”. *Computer Graphics Forum* 40 (2021) [1](#).
- [LMS*19] LAGUNAS, MANUEL, MALPICA, SANDRA, SERRANO, ANA, et al. “A Similarity Measure for Material Appearance”. *ACM Trans. Graph.* 38.4 (2019) [1](#).
- [MCR*21] MANTIUK, RAFAŁ K., CHAPIRO, ALEXANDRE, RUFO, GIZEM, et al. “FovVideoVDP: a visible difference predictor for wide field-of-view video”. *ACM Trans. Graph.* 40 (2021), 49:1–49:19 [1](#).

- [MPY*20] MIKHAILIUK, ALIAKSEI, PÉREZ-ORTIZ, MARÍA, YUE, DINGCHENG, et al. “Consolidated Dataset and Metrics for High-Dynamic-Range Image Quality”. *IEEE Transactions on Multimedia* 24 (2020), 2125–2138 [1](#).
- [NDM06] NGAN, ADDY, DURAND, FRÉDO, and MATUSIK, WOJCIECH. “Image-driven navigation of analytical BRDF models”. *Proc. EGSR*. 2006 [1](#).
- [SCW*21] SERRANO, ANA, CHEN, BIN, WANG, CHAO, et al. “The effect of shape and illumination on material perception: model and applications”. *ACM Trans. on Graph.* 40.4 (2021) [1](#).
- [WAKB09] WILLS, JOSH, AGARWAL, SAMEER, KRIEGMAN, DAVID, and BELONGIE, SERGE. “Toward a perceptual space for gloss”. *ACM Trans. Graph.* 28.4 (2009), 1–15 [1](#).
- [ZIE*18] ZHANG, RICHARD, ISOLA, PHILLIP, EFROS, ALEXEI A., et al. “The Unreasonable Effectiveness of Deep Features as a Perceptual Metric”. *Proc. CVPR* (2018), 586–595 [1](#).