

Perceived Translucency at Different Spatial Scales in Color and Grayscale Images

Aqsa Hassan and Davit Gigilashvili

Department of Computer Science, Norwegian University of Science and Technology; Gjøvik, Norway

Abstract

Many objects and materials in our daily lives are translucent. Translucency is an important attribute of appearance together with color, gloss, and texture. However, it remains largely unexplored whether and how these attributes impact each other. While the vast amount of literature exists about color reproduction, very little is known whether color reproduction at the same time affects perceived translucency. A substantial part of the translucency perception research is conducted on grayscale stimuli, which raises the question whether their findings can be generalized to the chromatic world we live in. A previous work showed that translucency changes when the image is converted to grayscale. One potential explanation the authors offered was the easier recognition of familiar materials. In this work we conducted psychophysical experiments where four versions of the images of different objects were shown: cropped close-ups and full images both in color as well as grayscale. The observers had to classify materials in each image as transparent, translucent, or opaque. We hypothesized that cropping would make material identity more ambiguous, and hence, affect translucency. We observed that for some images, conversion to grayscale affects translucency, and this effect is usually stronger for cropped versions. However, this effect was not observed for some other images. Overall, the way color and cropping affect translucency was not systematic across the dataset, which opens up additional questions for future work to explain these cross-stimuli differences.

Introduction

Translucency is an important attribute of objects and materials. We daily encounter and interact with many translucent materials, such as our own skin, plastic, fabric, wax, jade, frosted glass, and a broad range of foodstuff (meat, fruits, cheese) and beverages (milk, beer, juice). Translucency can be interpreted both as an optical as well as perceptual property. The standard terminology of appearance (ASTM) defines translucency as “the property of a specimen by which it transmits light diffusely without permitting a clear view of objects beyond the specimen and not in contact with it” [1]. The human visual system (HVS) has a remarkable ability to distinguish opaque and translucent materials. Perception of translucency has significant practical implication for humans, which may explain its evolutionary basis. For instance, translucency may indicate the state of food, whether it is ripe or raw, edible or rotten and spoiled [2]. Our understanding of the exact mechanisms of translucency perception is limited, and much remains still to be learnt.

The CIE (International Commission on Illumination) names translucency among the four fundamental appearance attributes of a material, along with color, gloss, and texture. These attributes are not isolated and impact each other [3]. The abundance of translucent materials makes generation, reproduction, and communication of translucent appearance significant

in many fields, such as 3D printing, computer graphics (e.g. skin rendering), and aesthetic medicine [4]. Color reproduction has long been a central problem in color science, and a vast amount of literature exists about color management and reproduction [5, 6]. However, little is known about how color and its reproduction affect other appearance attributes and overall appearance of an object. For example, if a publication contains a color image of a translucent gelatin, will it still appear translucent if we print it out on a grayscale printer, or will it look more opaque in a grayscale reproduction? The majority of the studies investigating translucency perception use grayscale stimuli [4, 7]. Can their findings really generalize to the chromatic materials and chromatic world we live in?

An analogous question was recently investigated by Liao *et al.* [7]. They found an intriguing dichotomy when an RGB photograph is converted to grayscale. Some objects and materials that appear translucent in RGB images start looking opaque in the grayscale, while the exact opposite is true for some others, i.e. they may look opaque in RGB and translucent in grayscale. The authors do not offer a conclusive explanation for this phenomenon and discuss that recognition of a familiar object could play the role in this process. This raises another interesting question: does cream appear more translucent, because we first understand it is cream, and then based on our prior experience, we expect it to be translucent? A similar phenomenon is already known in color science, where the memory and associations with familiar objects and materials are known to play a role in color perception – referred to as *memory color* [8].

To answer these questions, we conducted psychophysical experiments. In addition to converting RGB images to grayscale, as done in [7], we also cropped the images in such a way to ensure that the object and material recognition was substantially more difficult. Cheeseman *et al.* [9] demonstrated that material recognition may be affected by the spatial scale. The contribution of this work is the following:

- We test whether conversion of an image to grayscale affects its translucency. We explore whether the findings by Liao *et al.* [7] are reproduced.
- Additionally, we introduce the third label *transparent*, since *opaque* and *translucent* are not inclusive and encompassing enough. “Transparent substances, unlike translucent ones, transmit light without diffusing it” [10].
- We investigate whether cropping an image in a way to complicate object and material recognition affects its translucency. This way we want to understand the impact of spatial scale and consequent material recognition on perceived translucency. To the best of our knowledge, this is the first study to address this.
- We also study the case where both color and full context information are missing (cropped grayscale images).
- We analyze the potential role of chromatic and achromatic components in each individual case.

Related Works

A comprehensive review on translucency perception can be found in [4], which summarizes the state of the art as follows: the HVS relies on spatial distribution of luminance intensities. Edges and thin parts are usually bright in translucent materials, as many photons manage to go through. Another area that may help with telling translucent and opaque objects apart are those that are in a shadow on an opaque object, such as concavities. Conversely, they are brighter on a translucent object, because unlike opaque materials, additional light reaches there via subsurface. Despite these observations, the exact calculations made by the HVS remain unknown.

When talking about color, it is important to emphasize that we mean a chromatic component of it. The majority of the studies on translucency perception use grayscale stimuli, and proposed cues to translucency are related to luminance intensity and distribution. Lightness is correlated with absorption and scattering coefficients, since highly absorbing materials are darker and highly scattering materials are lighter [11]. The work by Gigilashvili *et al.* [12] is one of the few exceptions where real physical objects were used. They noticed that lighter objects, both chromatic and achromatic, are often associated with translucency due to their “milky” appearance. Less is known on how chromatic information relates to perceived translucency. Fleming and Bülthoff [2] observed that saturation alone is not enough to produce translucent appearance, but saturation can enhance or modify translucency. Even when the mean saturation of two objects is equal, if the saturation and lightness (value of the HVS color space) are positively correlated, translucency is manifested as a warm glow, while if the correlation is negative, translucency looks icy and cold. Di Cicco *et al.* [13] also mention the correlation between translucency and saturation in still life paintings.

Chadwick *et al.* [14] reported an interesting case, where an observer with a color blindness of a cortical origin was able to scale translucency. The authors conclude that color and translucency processing happen in different parts of the brain, and these two processes are anatomically independent. Despite this discovery, the performance of the color normal observers deteriorated when they were shown grayscale images instead of RGB [14], which can be an indication that material recognition and high level associations could play the role in this process.

The most comprehensive study on this topic was conducted by Liao *et al.* [7]. They observed many label flips between color and achromatic conditions. In the majority of the images, the label flipped from *translucent* or *unsure* to *opaque*. However, there were considerable number of cases, where the flip was in an opposite direction: the objects previously labeled as *opaque* or *unsure* became *translucent* in grayscale. Among the potential explanations for this effect, the authors name recognition of familiar objects and materials, and the amount of spatial variation in lightness and chroma channels. In their subsequent work, the authors [15] introduce a deep generative model to produce synthetic images of translucent materials. Interestingly, color and translucency representation emerges in different parts of the latent space. Translucency can be manipulated in the middle layers, while color in the deeper layers independent of translucency.

Methodology

Dataset Preparation

The dataset of the experimental stimuli includes 236 images of everyday objects such as fruit, honey, seafood, plastics, crystals, soap, wax, jade, stone, and so on. The dataset was constructed using 59 original RGB photographs (royalty-free pho-



Figure 1: The examples of full-sized images from the dataset used in the psychophysical experiment with their grayscale, cropped, and grayscale cropped counterparts from top to bottom, respectively.

tos from the *Unsplash.com* stock photography database), and each image was shown in four different versions ($59 \times 4 = 236$): Full-sized color image (the original), full-sized grayscale image, cropped color image, and cropped grayscale image. Full-sized means that images are not being cropped. The images include a variety of backgrounds, e.g. some are segmented and put on a black background, some are on a white background, and some are in a full scene context. We cropped the images to make it difficult to identify a familiar object and a material it is made of. The size of full-sized images is 800×800 pixels, and that of cropped ones is 226×273 pixels. The examples of full-sized and cropped images are illustrated in Figure 1. To generate images in equal resolutions, we had to resize some photographs using OpenCV function *cv2.resize* with bicubic interpolation to minimize the artifacts. *cv.cvtColor* function was used for RGB to YIQ conversion to obtain a grayscale version. Liao *et al.* [7] observed no significant differences between different grayscale conversion methods. Therefore, we stick to a single method.

Experimental Procedure

The experiment was hosted at the QuickEval [16] platform and was conducted under controlled laboratory conditions (discussed in the following section). In total, 236 images have been displayed to each observer, where observer’s task was to label each of the displayed image. Unlike Liao *et al.* [7], who used only two labels, either *translucent* or *opaque*, we used the third label *transparent* as well. Highly transparent objects and materials permit a clear view of the background, while translucent ones blur and often completely occlude it. According to the CIE, “*translucency is a subjective term that relates to a scale of values going from total opacity to total transparency*” [3]. Gigilashvili *et al.* [17] propose that translucency is part of the transparency-opacity continuum, and it is correlated with them with a Gaussian-like bell-shaped curve – the degree of translucency is gradually increasing, reaching its peak and then decreasing between transparency and opacity.

The sequence of the images displayed to each observer was the following: grayscale cropped, color cropped, grayscale full-sized, and finally color full-sized. The images were randomized within each category. The reasoning behind displaying the images in this particular order is related to giving away the least information to the observer about the object at the start of the experiments. For instance, when the grayscale cropped images are being displayed first, the observer has information neither on color nor the full object. This helps us avoid bias due to recogniz-

ing the cropped part from a full-scale image. Liao *et al.* [7] used different observers for color and grayscale experiments. In that case, they could not rule out the possibility of potential differences in life experiences and respective familiar material recognition between the two groups. Therefore, we took an intentional decision to have all images assessed by the same observers.

Display and Observation Conditions

The 32-inch PA32UCG HDR ASUS ProArt display has been used with a peak luminance of 1600 nits. The display has been operating in HDR PQREC.2020 mode while supporting 10 bits. The pixel resolution of the display is 3840×2160 . The display light was the only source of lighting in the room. The distance between the observer and the display was 60cm.

Observers

14 observers (5 female and 9 male) participated in this study, and each observer took approximately 35 minutes to complete the experiment. The average age was 28. 11 observers had previously participated in psychophysical experiments. All observers had either normal (20/20) or corrected-to-normal visual acuity. Each observer was given above-mentioned definitions of the terms. They went through a short training session on similar images to get acquainted with the task.

Results

χ^2 Test of Independence

First of all, we conducted a χ^2 statistical test of independence, which tests whether significant relationship exists between two categorical variables. We had the version of image as a categorical predictor variable with four options: color full, gray full, color cropped, and gray cropped; and observer's answer as a dependent categorical variable with transparent, translucent, and opaque as options. We tested whether the version of the image that is shown has a statistically significant relation with observers' answers. Against our expectations, the null hypothesis that the two are independent was not rejected ($\alpha = 0.05$; $DoF = 6$; $\chi^2 = 8.49$). However, this test takes all observations into consideration, while Liao *et al.* [7] have demonstrated that the effect varies substantially from image to image. Fig. 2 shows the average number of observers that changed their response between different versions (color full, gray full, color cropped, gray cropped) of the image. It confirms the observation made by Liao *et al.* and shows that while for some images the observers respond consistently across different versions, for others the majority of them change their answer. Therefore, we need to have a closer look at the data, identify unique cases and attempt to explain them. Interestingly, the 59 images are relatively uniformly distributed along the vertical axis in the plot, which indicates that our dataset includes a broad range of diverse images.

How many observers change their response?

Fig. 3 shows how many observers on average (considering all 59 images) changed their response between the two versions of the image. There are six possible pairs of versions to compare observers' responses for: Color Full (CF) and Gray Full (GF); Color Cropped (CC) and Gray Cropped (GC); Color Full (CF) and Color Cropped (CC); Gray Full (GF) and Gray Cropped (GC); Gray Full (GF) and Color Cropped (CC); Color Full (CF) and Gray Cropped (GC). For convenience's sake, we will use these two letter codes. For instance, CF-GC means that we compare their responses given to Color Full and Gray Cropped versions of the same original image. The figure shows that on av-

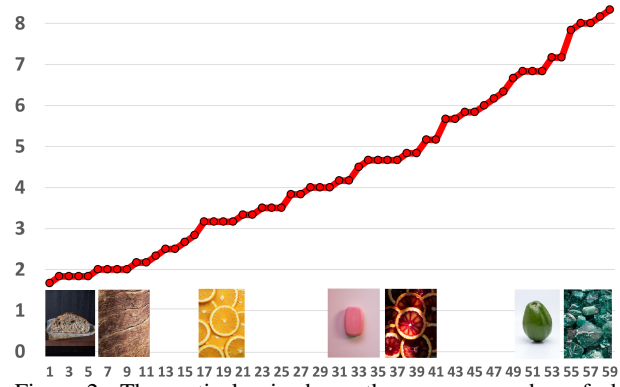


Figure 2: The vertical axis shows the average number of observers that changed their answer across the different versions of the image. The horizontal axis is each of the unique 59 images. The average was calculated as the sum of changes between all different combinations of the four versions divided by 6 (the number of combinations). The plot shows that observers' behavior depends on the image content. While they rarely change their opinion for some images, the change is frequent for others (see examples below). Sorted for readability.




Figure 3: Left: The mean number of observers that changed their answer between the versions. Right: Standard deviation.

erage at least 5 out of 14 observers give a different response (either transparent, translucent, or opaque) when they assess CF and GC or GF and GC versions of the same image. This number is lowest between their judgments of CF and GF images. When comparing color and grayscale versions, the changes are more frequent if both are cropped (cf. CF-GF and CC-GC). When comparing full and cropped versions, the changes are more frequent when the both are grayscale (cf. CF-CC and GF-GC). The figure also reports standard deviation among 59 images, which is lowest for CF-GF and largest for CF-GC. This indicates that the dependence on image's unique context may be largest when both color and scene information are missing.

While Fig. 2 reports the number of observers that change their mind on average between all different combinations of image versions, it is worth having a closer look to each individual cases (59 images). Fig. 4 shows how many observers change their answer between a given pair of different versions. The results are sorted by the results for the case when the two versions differ most (the responses for a full color image and its grayscale cropped version; one has both color and full scene information, while the other lacks the both). It would have been intuitive to expect that this pair (CF-GC) would yield the largest differences – which is true on average (Figure 3); however, it is not always the case. For some images, observers' responses differ most between cropped color and full grayscale (GF-CC), or between the full and cropped versions of the grayscale image (GF-GC). Although the frequent response change across one pair of conditions may mean the frequent changes between all conditions, this is not al-




Figure 4: The figure illustrates how many observers changed their label between the versions of images. The vertical axis shows the number of observers (out of 14) that changed their answer between the two versions; the horizontal axis corresponds to individual 59 images. The results vary among images.

ways the case. For example, for the 59th image, the responses between CF and GC differed for 12 out of 14 observers, while 12 out of 14 assessed CF and CC consistently.

Interobserver Consistency

Liao *et al.* [7] noticed that observers are more consistent with one another when a color image is shown. Since we hypothesize that observers use material identification and respective knowledge about its properties, they are expected to be more consistent when it is easier to identify a material and object; and conversely, if the object and material are ambiguous, there is more room for individual interpretation. To check this hypothesis, we calculated Cohen’s $Kappa$ (κ), which is a widely used metric to quantify the response similarity between a pair of observers [18]. It is calculated as:

$$\kappa = \frac{P_{\alpha} - P_e}{1 - P_e} \quad (1)$$

where κ is the coefficient of interobserver consistency; P_{α} is observed proportionate agreement (the portion of their answers that match); and P_e a chance agreement. Negative κ means that there is no agreement or even a trend that the two observers answer differently. 0 to 0.20 corresponds to little, 0.21 to 0.40 mild, 0.41 to 0.60 moderate, 0.61 to 0.80 substantial, and 0.81 to 1 a very strong agreement (perfect for 1). A general trend is consistent with our expectation (see Fig.5): observers are more consistent with one another when full versions are shown than in case of cropped ones; and they are more consistent with color images. While the majority of observers agree to some extent (e.g. there is a substantial agreement among observers 9 through 13 in full color images), there are several observers, such as observers 4, 8, and 14 that disagree with the majority. A potential explanation can be the interpretation of the concepts of *transparency*, *translucency*, and *opacity*; as discussed by Gigilashvili *et al.* [17], the boundaries among the three concepts can be ambiguous.

Label Flips

Until now, we discussed how individual observers changed their response among the four versions of the same image. As shown in Fig. 5, there are several observers that often disagree with the rest that may create an impression that answer change is frequent. In order to consider a given image labeled as *transparent*, *translucent*, or *opaque*, Liao *et al.* [7] took a threshold of 60% agreement (although they had just two categories - translucent and opaque). For instance, if at least 60% of observers classify shown material as translucent, we assign it a label *Translucent*; if at least 60% answer that it is opaque, its label is *Opaque* and so on. However, if none of the three answers (two in Liao *et*

al. [7]) get at least 60%, then it is labeled as *Uncertain*. In our case, the nearest integer number of observers made the requirement at least 64% (9 out of 14).

Fig. 6 shows the label distribution for each version of the images. Generally, transparent ones were the fewest, and the most images were labeled as *Opaque* or *Uncertain* (i.e. neither transparent, nor translucent or opaque got 64% of the observers’ responses). Both full as well as cropped color images received slightly less *Transparent* and *Translucent* labels when converted to grayscale. Interestingly, the number of opaque ones also decreases when cropping or converting to grayscale. This is explained by a substantial increase in *Unsure* category. For example, if for full color there were only 10 images out of 59 where none of the answers got 64%, this number is doubled for cropped gray images. This means that when we either crop or remove color, we remove significant cues for judging translucency.

Finally, Fig. 7 shows the distribution of the flip directions. Liao *et al.* [7] observed that there is no consistent direction of the flips. While some images flip from translucent to opaque when shown in grayscale, some others flipped from opaque to translucent. They hypothesize that this can be explained by image statistics and whether the gradient is more apparent in the lightness or chroma channel (of CIE LCh). Similar to our results, they also observed abundance of flips to and from *Unsure* category.

We can see that some flips, such as *Transparent to Opaque* or *Translucent to Transparent* never happened. In general, few flips involve *Transparent*. There can be two reasons for that: first, there are less transparent objects in the dataset; and second, the cues to transparency are more robust to losing color and cropping – you either see the background or you don’t, while translucency and opacity involve a more sophisticated analysis of spatial distribution of color gradient.

Direct flips between *Translucent* and *Opaque* are relatively rare. The most frequent flips are to and from the *Uncertain* category. *Opaque to Uncertain* is a frequent flip from cropped grayscale to full color or full grayscale, as well as from cropped color to full color. It can indicate that the gradients in the smaller region are initially considered part of the surface texture, but the observers may notice that the variation could be due to subsurface scattering when they see the full context. The flip from *Uncertain to Opaque* is less frequent. Translucency becomes more uncertain when a color version of a grayscale is seen, or a full version of a cropped color is seen. This may indicate that both full spatial scale as well as variation of chromatic information may be contributing to judging an object as translucent. However, these observations are limited to individual cases, and utmost care is required before generalizing them.

Discussion

Liao *et al.* [7] hypothesize that the flip direction may depend on whether the gradient and spatial variation is primarily present in the lightness or the chroma channel. In the former case, grayscale image keeps its translucency or even looks more translucent, while in the latter case, removal of the chromatic information flips the label from translucent to opaque. Similarly to Liao *et al.*, we converted the images to CIE L*Ch and separated L* and chroma channels. As also shown by Liao *et al.*, some materials exhibit more variation in the lightness channel, while others keep more visible gradient in chroma (see the examples in Fig. 8). Afterward, we calculated the standard deviation and the average magnitude of the gradient in each of these channels. The intensity variation neither in the lightness nor in the chroma channel turned out informative enough to predict either the num-




Figure 5: The heatmap of Cohen’s κ interobserver differences. Some observers, such as number 4, 8, and 14 stand out from the rest, while others are more consistent. The observers are more consistent for color than for gray and for full than for cropped images.

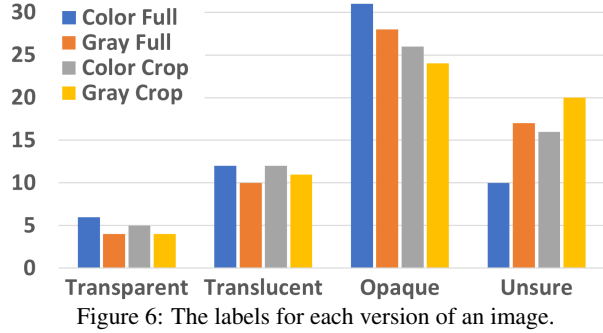


Figure 6: The labels for each version of an image.

ber of flips or the flip directions in a consistent manner. However, this analysis still revealed one important aspect that we did not account for before. It has shown that the RGB images that depict primarily achromatic objects have little chroma and nearly zero variation in the chroma channel. This kind of images were, indeed, consistently judged between RGB and grayscale versions.

What turned out more informative in terms of predicting flips, is whether the material has a very unique characteristic texture that makes its recognition possible even from the cropped and grayscale versions. The example is illustrated in Fig. 9 (left and middle). It shows that avocado, whose texture can easily become ambiguous, has one of the most frequent flips across the conditions, while wood, whose texture is very characteristic and easy to spot both in the cropped as well as grayscale versions, was judged very consistently with minimal number of flips. This can be an additional indication to the role of material recognition.

If color and cropping do not contribute to translucency perception, ideally, the responses given by the observers should have been consistent across the four versions. Opposite to this, we observed that flips occur, and observers answer differently when they are shown different versions of an object. Further indication is the fact that when we remove color and change spatial scale, the observers become less consistent with one another, i.e. more ambiguity is created and random decisions are more often.

The flips happen when an image, either full or cropped, is converted to grayscale. This may indicate the importance of color. However, if we compare CF-GF and CC-GC flips, we will see that the flips are less frequent for full spatial scale (uncropped). This can be an indication that recognition of object’s and material’s identity makes the label more robust to the loss of color. High number of flips between full and cropped versions (CF-CC) of the same image may also indicate to the role of material recognition – however, here we cannot rule out that cropping itself does not remove critical low-level translucency cues. Cropping has stronger effect on grayscale images (GF-GC), because absence of color creates an additional layer of ambiguity. Finally, the biggest difference was observed between full color and cropped grayscale images (CF-GC), indicating that both color and spatial scale play a role. Fig. 9 (middle and right) illustrates the examples of flips from translucent to uncertain when

the image was cropped, and from opaque to translucent when it was cropped and converted to grayscale (refer to the caption). However, it is important to point out that despite these individual observations, there is no systematic manner in which color removal or cropping affect translucency – some flip in one direction, while others may flip in the opposite direction. Future work should examine individual cases to explain these discrepancies.

Limitations and Future Work

We acknowledge that this work is based on two assumptions: first, we do not remove essential low-level cues to translucency by cropping; and second, cropping makes object and material identification more difficult. These assumptions may not always hold. Keeping the translucency cues and making material identification more difficult are somewhat mutually exclusive tasks. If we leave a large region to keep more cues to translucency, it will be easier to identify object and material; conversely, if we crop too much, we risk inadvertent removal of translucency cues (e.g. see crystals in Fig. 9 (right)). Not all objects and materials are equally straightforward to mask the identity of. For instance, in Fig. 10, it is more difficult to recognize a kiwi from a close-up image of its flesh than it is for a strawberry, which has its unique and characteristic texture and seed composition.

Material recognition in each version of the images should be experimentally tested in future works. Instead of a categorical experiment, the magnitude of translucency should be also quantified by rating. Another experiment may use image scrambling instead of cropping to complicate material recognition while keeping some of the image statistics. Future works should also identify the critical regions that should be kept for ensuring that the image contains translucency cues even though it is not possible to recognize the material. This can be achieved in two ways: first, reverse-correlation study conducted by Nagai *et al.* [19] that used mosaic images composed of different combinations of spatial patches to detect translucency “hotspots”; second, eye-tracking as in [20] to identify where the observers look for cues.

Conclusions

We explored whether color affects translucency and what role material recognition plays in this effect. We used four versions of images: full color, full grayscale, cropped color, and cropped grayscale. The images were cropped in a way that material and object recognition were more difficult if not impossible. The observers had to label them either as transparent, translucent, or opaque. Their responses across the four versions were consistent for some images, while differed substantially for others. Conversion to grayscale affects translucency, but this effect is weaker at higher spatial scale (full object). The responses did not change in a systematic manner: cropping or grayscale conversion could increase translucency for some images and decrease it for others. This certainly merits a further study.





Figure 8: Some objects (top row) retain a characteristic gradient in the lightness channel (middle) and have relatively homogeneous chroma (right) channel (CIE L*Ch), while others (bottom row) have a visible gradient in the chroma channel.




Figure 9: The wood images were judged consistently and had one of the lowest number of flips, while the image of avocado had one of the highest number of flips. This can be explained by wood's texture, which is easier to detect both in cropped as well as grayscale images. The crystal image flipped from translucent to uncertain when it was cropped, while avocado flipped from opaque to translucent when it was cropped and converted to grayscale. In the latter case, the familiar object at a higher spatial scale may help observers to judge it as opaque. In the former case, the full image contains more crystals, and the one in the cropped version may not be the most translucent one.

References

- [1] ASTM E284-17 Standard Terminology of Appearance, ASTM International, West Conshohocken, PA, 2017.
- [2] Roland Fleming and H. Bülthoff, Low-level image cues in the perception of translucent materials, *ACM Trans. Appl. Percept.*, 2, 3, pp. 346–382 (2005).
- [3] CIE 175:2006 A framework for the measurement of visual appearance, International Commission on Illumination, 2006.
- [4] Davit Gigilashvili, J. B. Thomas, J. Y. Hardeberg, and M. Pedersen, Translucency perception: A review, *J. Vis.*, 21, 8:4, pp. 1–41 (2021).
- [5] Phil Green and L. MacDonald, *Colour Engineering: Achieving Device Independent Colour*, John Wiley & Sons, 2011.
- [6] Robert R. W. Hunt, *The Reproduction of Colour*, John Wiley & Sons, 2005.
- [7] Chenxi Liao, M. Sawayama, and B. Xiao, Crystal or jelly? Effect of color on the perception of translucent materials with photographs of




Figure 10: Some materials retain a characteristic, easily detectable texture even when cropped and converted to grayscale (such as a strawberry on the left), while others, such as a kiwi, are challenging to recognize if color or full context is missing.

- real-world objects, *J. Vis.*, 22, 2:6, pp. 1–23 (2022).
- [8] Christoph Witzel and T. Hansen, Memory effects on colour perception, *Handbook of Color Psychology*, pp. 641–665 (2015).
- [9] Jacob Cheeseman, R. Fleming, and F. Schmidt, Scale ambiguities in material recognition, *iScience*, 25, 103970, 17 pages (2022).
- [10] Walter Gerbino, C. I. Stultiens, J. M. Troost, and C. M. de Weert, Transparent layer constancy, *J. Exp. Psychol.: Hum. Percept.*, 16, 1, pp. 3–20 (1990).
- [11] Alice C. Chadwick, G. Cox, H. E. Smithson, and R. W. Kentridge, Beyond scattering and absorption: Perceptual unmixing of translucent liquids, *J. Vis.*, 18, 11:18, pp. 1–15 (2018).
- [12] Davit Gigilashvili, J.-B. Thomas, M. Pedersen, and J. Y. Hardeberg, On the appearance of objects and materials: Qualitative analysis of experimental observations, *JAIC*, 27, pp. 26–55 (2021).
- [13] Francesca Di Cicco, M. W. Wijntjes, and S. C. Pont, If painters give you lemons, squeeze the knowledge out of them. A study on the visual perception of the translucent and juicy appearance of citrus fruits in paintings, *J. Vis.*, 20, 13:12, pp. 1–15 (2020).
- [14] Alice Chadwick, C. Heywood, H. Smithson, and R. Kentridge, Translucence perception is not dependent on cortical areas critical for processing colour or texture, *Neuropsychologia*, 128, pp. 209–214 (2019).
- [15] Chenxi Liao, M. Sawayama, and B. Xiao, Unsupervised learning reveals interpretable latent representations for translucency perception, *PLOS Comp. Biol.*, 19, 2:e1010878, pp. 1–31 (2023).
- [16] Khai Van Ngo, J. J. Storvik, C. A. Dokkeberg, I. Farup, and M. Pedersen, Quickeval: a web application for psychometric scaling experiments, in *Proc. IQSP XII*, 9396. SCIA, pp. 1–13 (2015).
- [17] Davit Gigilashvili, J. B. Thomas, J. Y. Hardeberg, and M. Pedersen, On the nature of perceptual translucency, in *Proc. of MAM2020*, pp. 17–20 (2020).
- [18] Jacob Cohen, A coefficient of agreement for nominal scales, *Educ. Psychol. Meas.*, 20, 1, pp. 37–46 (1960).
- [19] Takehiro Nagai, Y. Ono, Y. Tani, K. Koida, M. Kitazaki, and S. Nakauchi, Image regions contributing to perceptual translucency: A psychophysical reverse-correlation study, *i-Perc.*, 4, 6, pp. 407–428 (2013).
- [20] Davit Gigilashvili, A. S. Sole, S. D. Nath, and M. Pedersen, Exploring the role of caustics in translucency perception — an eye tracking approach, in *Proc. Int'l. Symp. on El. Img.*, 34, pp. 222:1–222:6 (2022).