

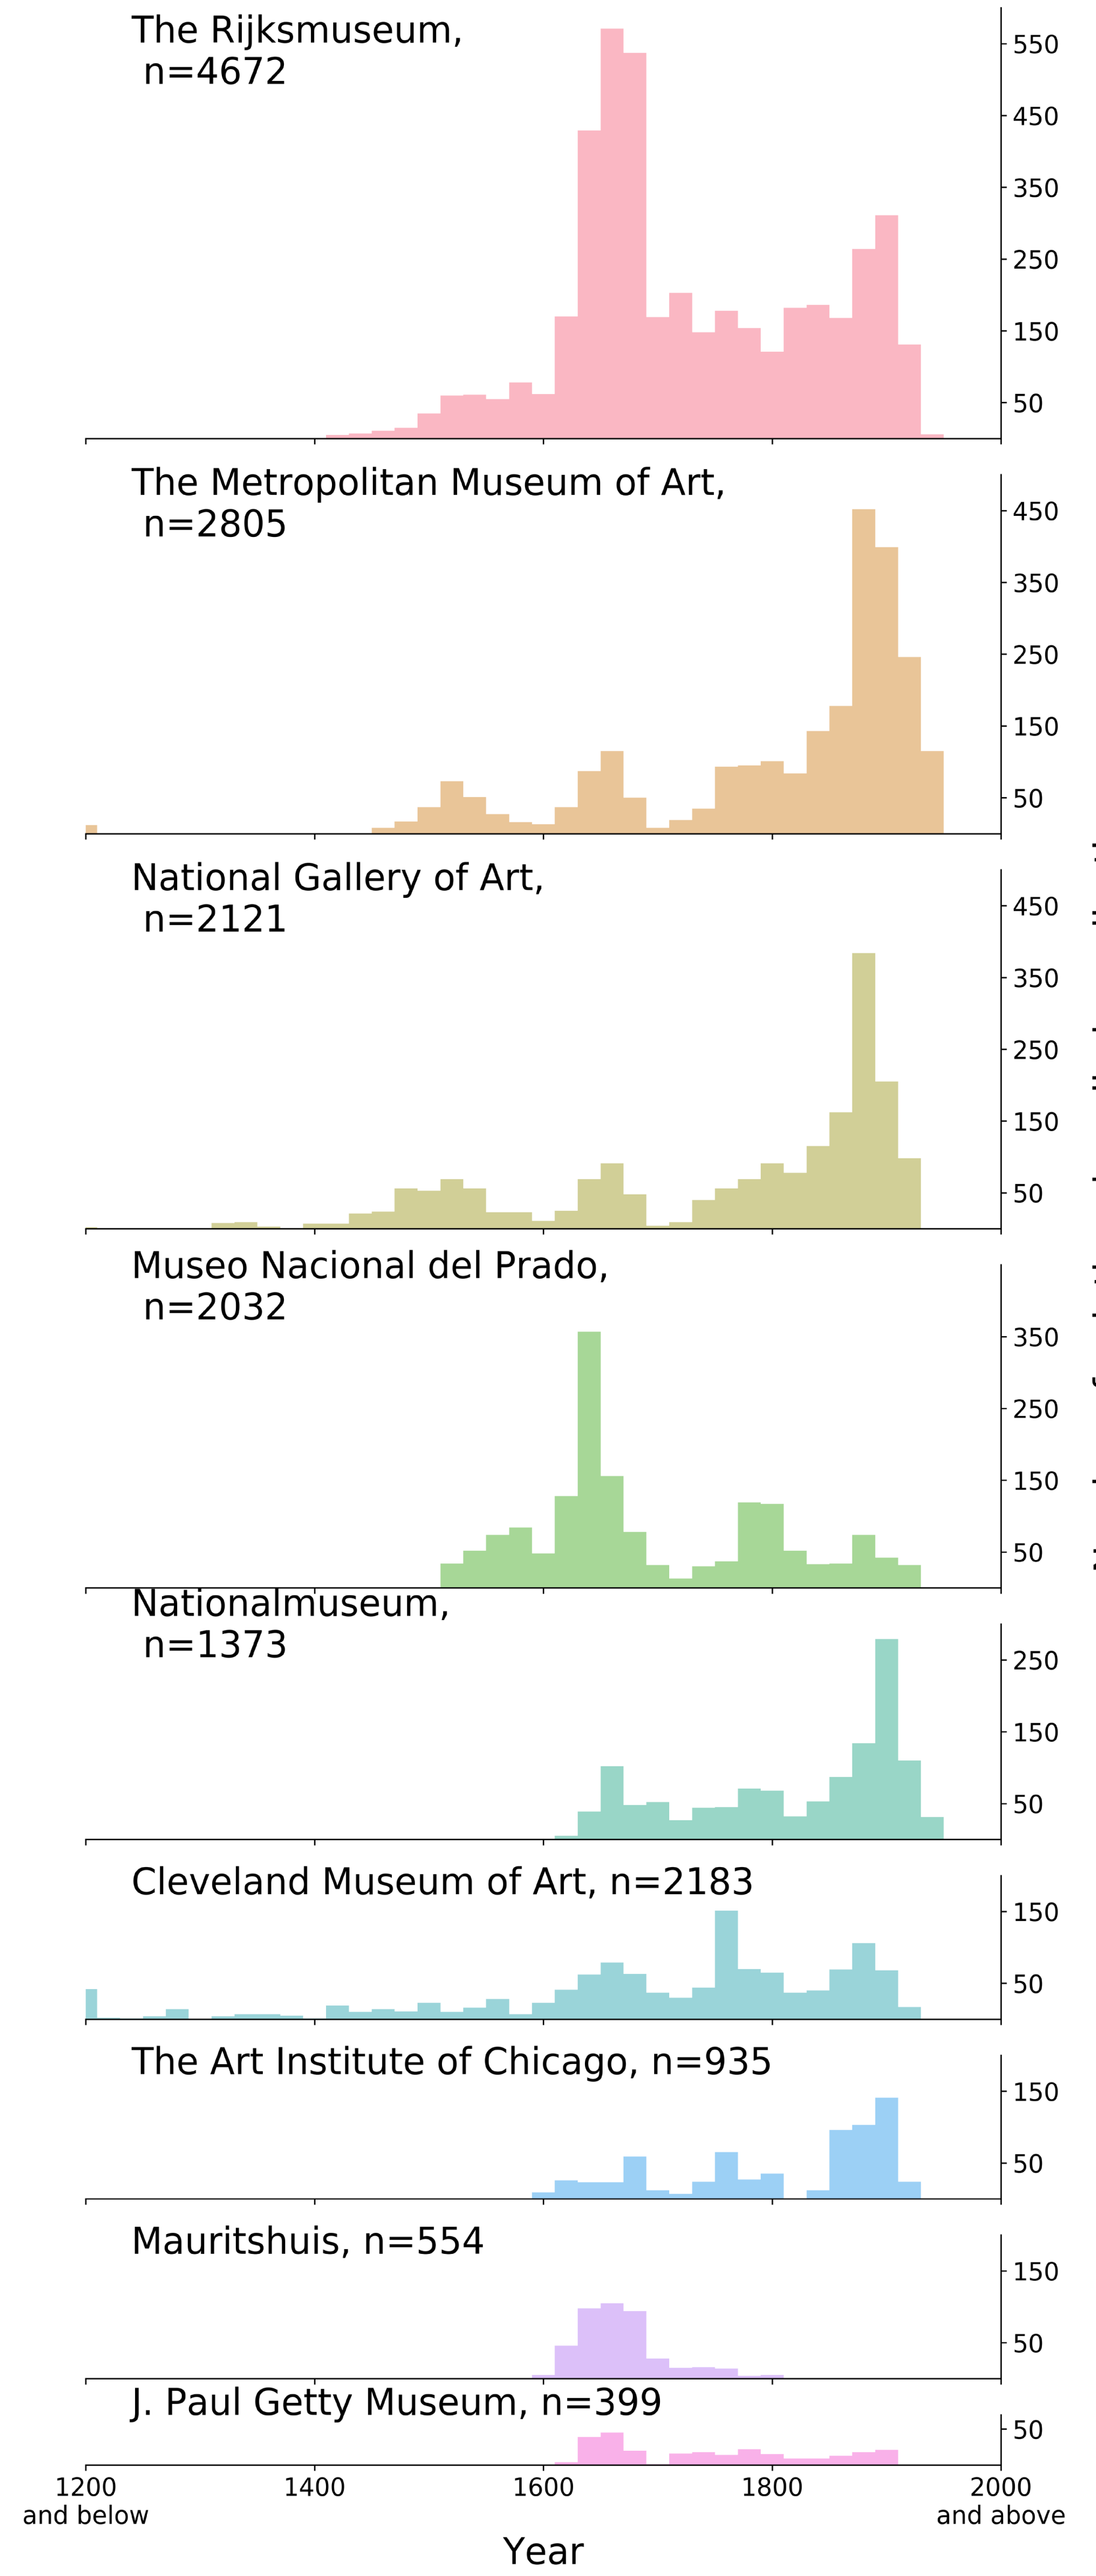
A Database of Painterly Material Depictions

Mitchell Van Zuijlen¹, Hubert Lin², Kavita Bala², Sylvia C. Pont¹, Maarten W.A. Wijntjes¹
¹Delft University of Technology, Perceptual Intelligence lab, ²Cornell University, Computer Science Department
 Corresponding Author: m.j.p.vanzuijlen@tudelft.nl - www.MitchellVanZuijlen.nl/

Painters depict materials by utilizing an implicit knowledge of human material perception. To enable studies of this implicit knowledge, we created a database of 20K paintings annotated with 300K bounding boxes of various depicted materials.

Collecting Paintings


First, we collected a set of 19325 paintings from online open-access galleries of nine internationally renowned art institutions. This was done through web scraping, where each individual image was downloaded with corresponding meta-data.



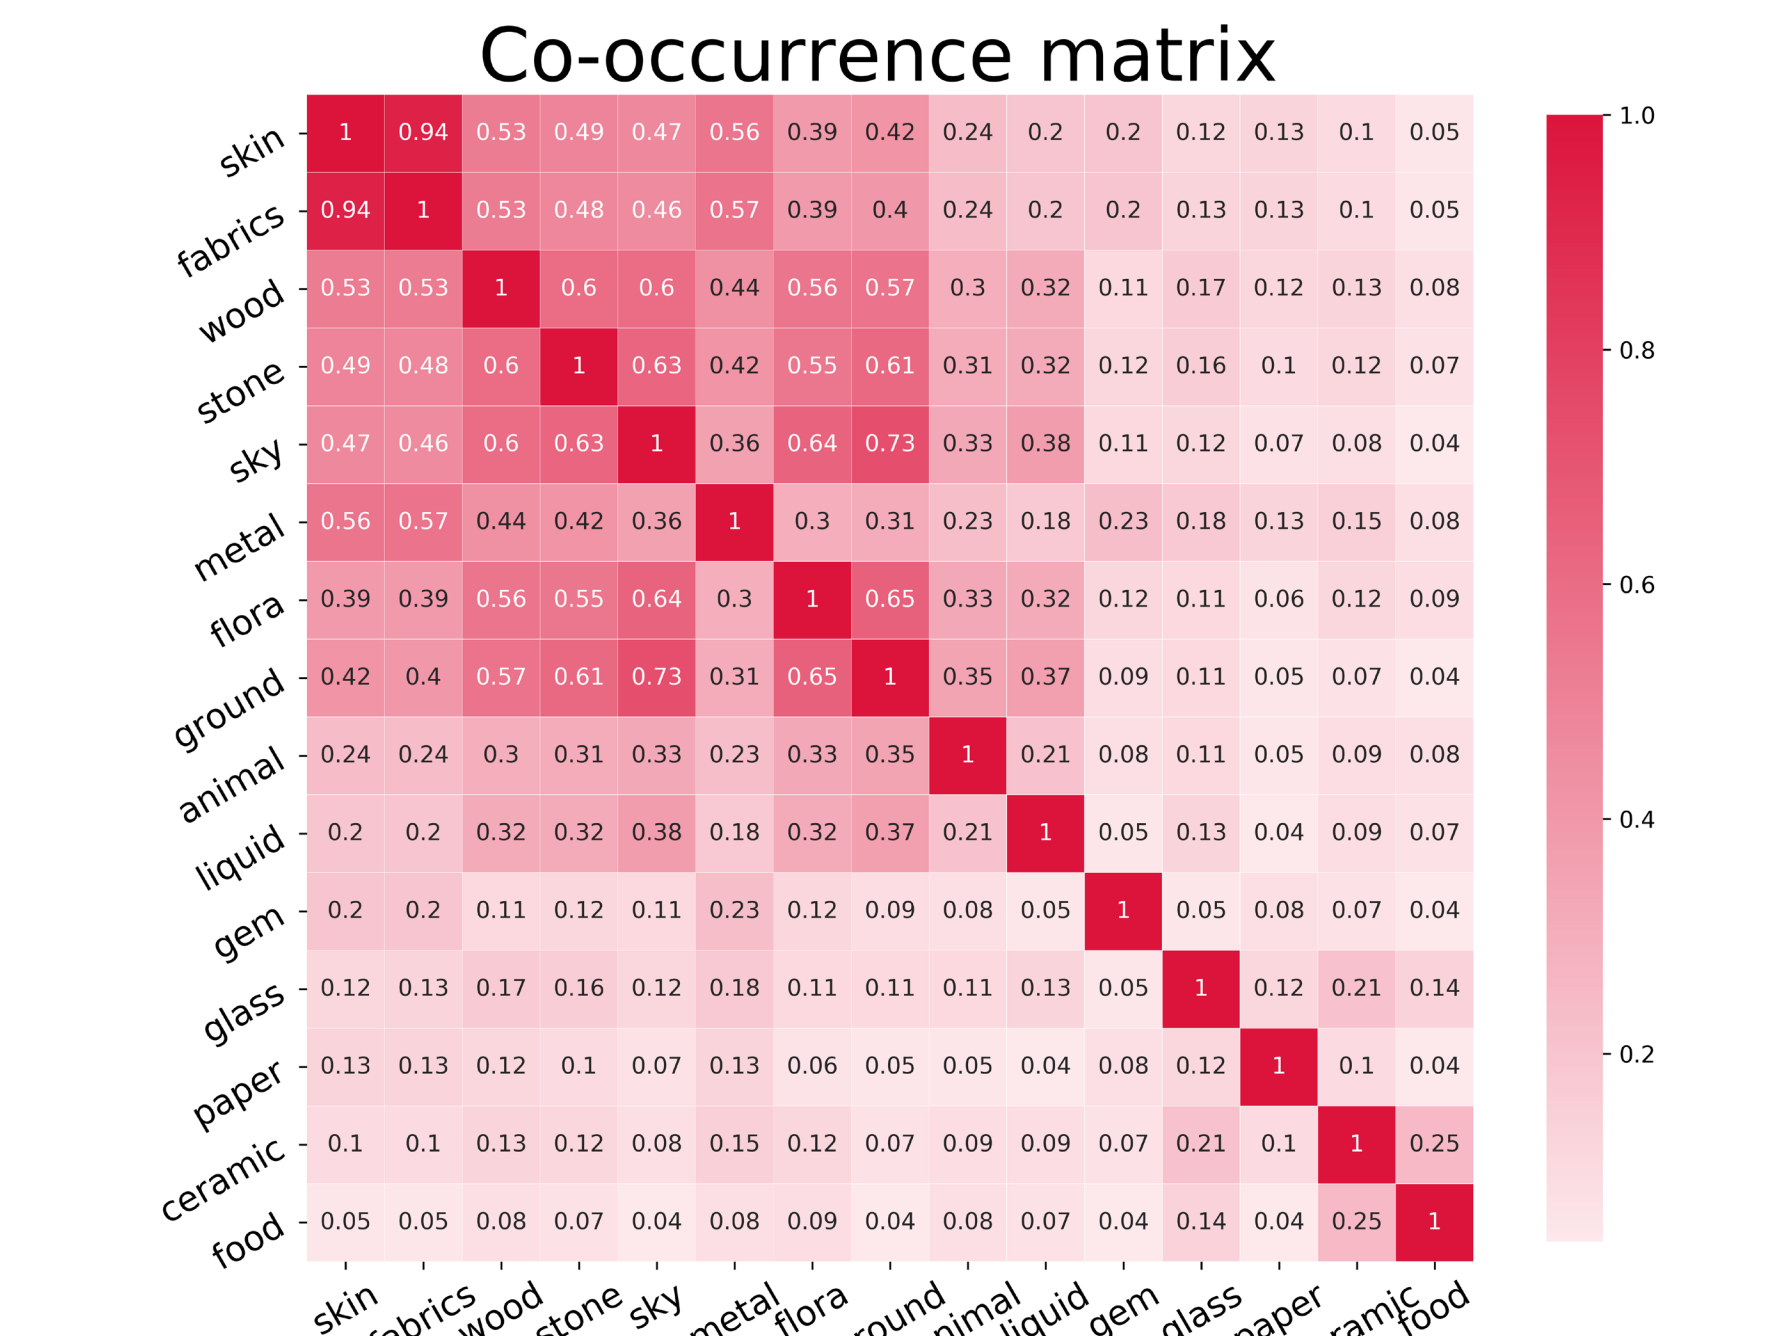
Individual galleries do not necessarily offer a representative sample of paintings throughout history. Here we can see that paintings are not evenly distributed over time for galleries.

Collecting Material

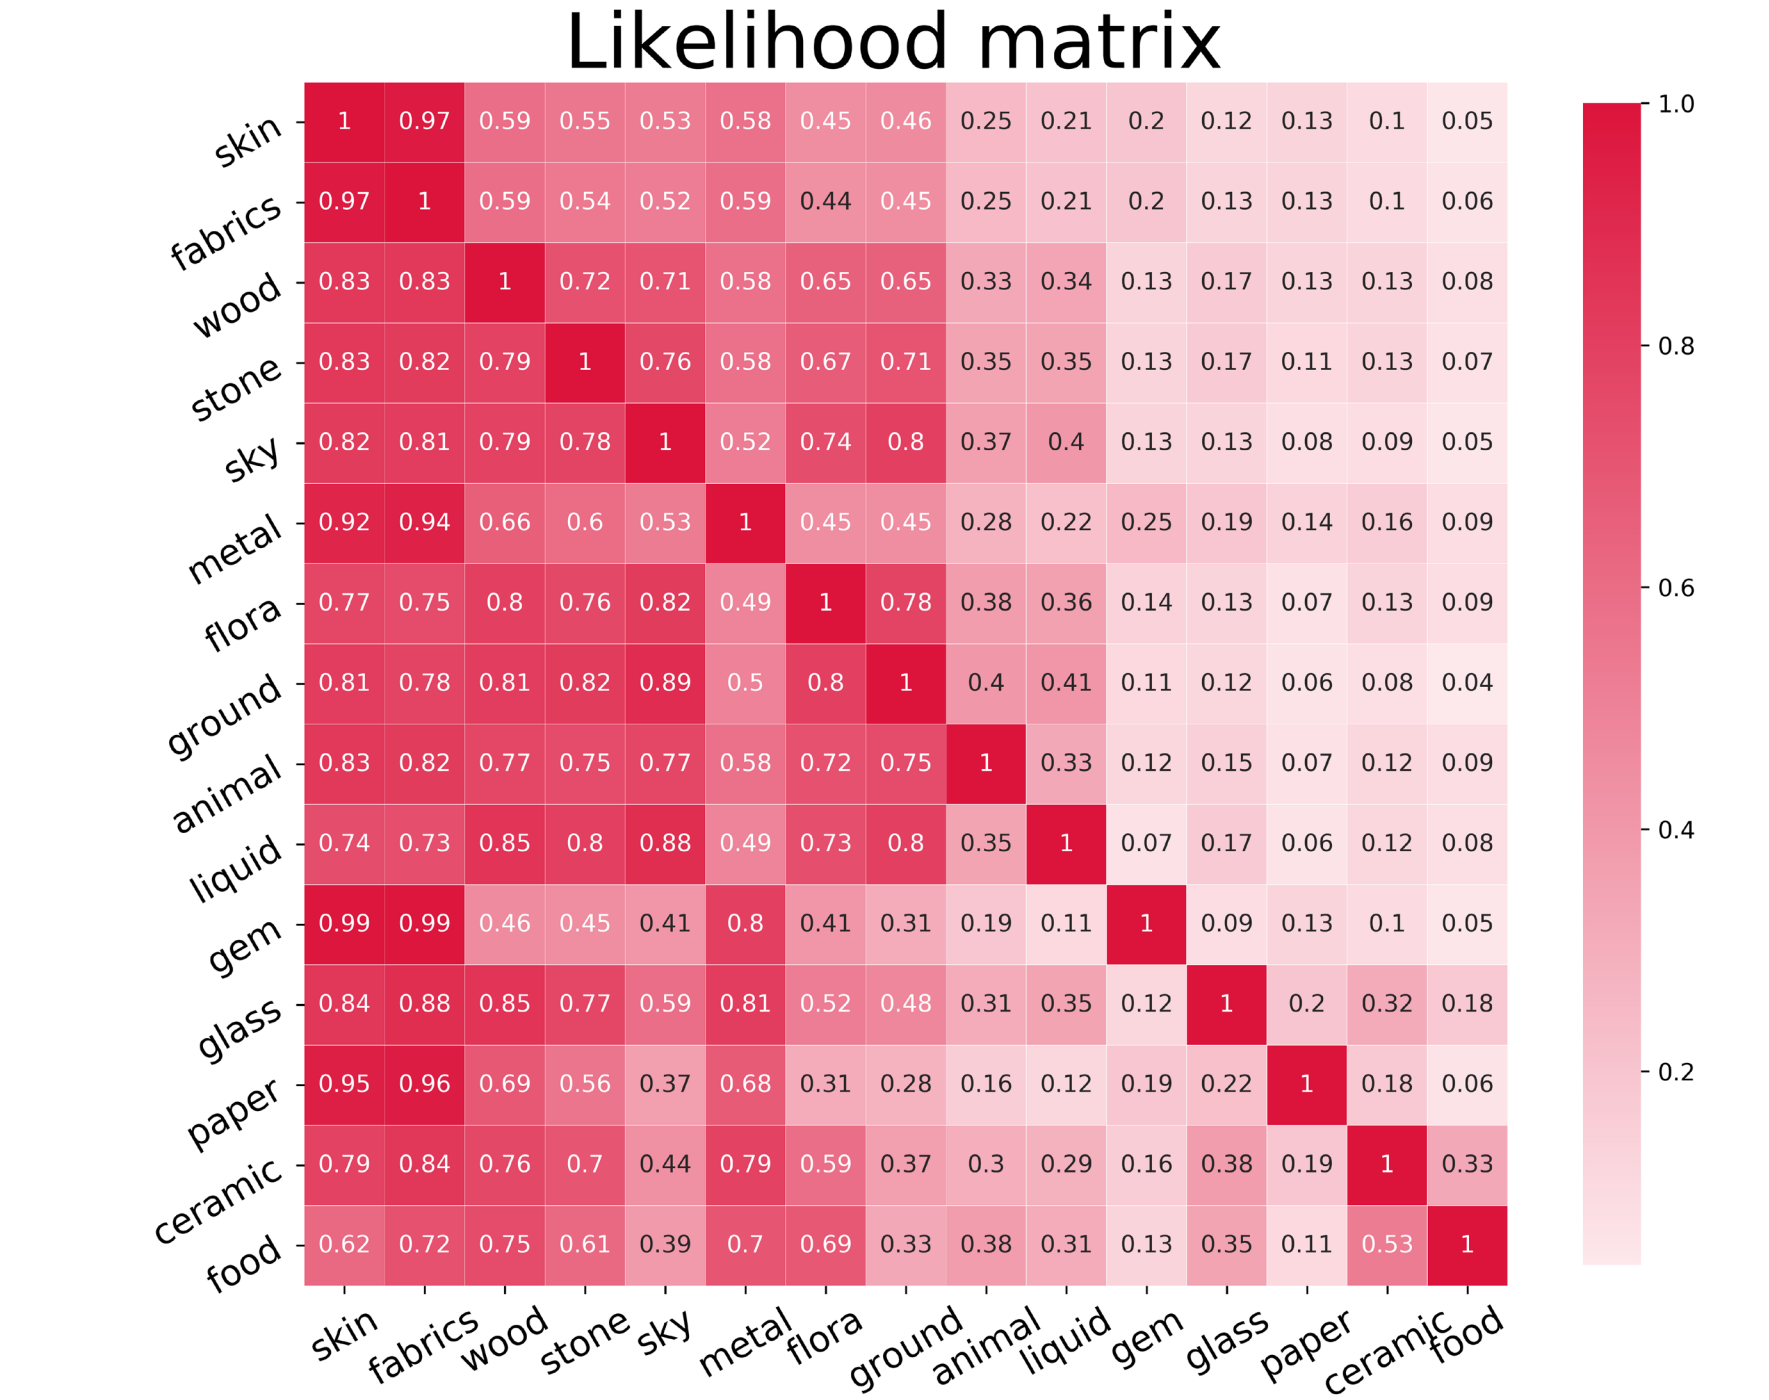
For 15 material categories, we asked human annotators from Amazon Mechanical Turk (AMT) to indicate the presence of each material within each painting. Each painting was annotated for each material by at least 5 participants. If annotators achieved 80% agreement on a painting/material combination, we would mark that painting to depict the target material.



The proportion of paintings in our dataset that depict at least one instance of each material. For example, skin and fabric are depicted in +/- 80% of our paintings, while food is only present in about 5% of the paintings



A matrix showing the rate of co-occurrence between each pair of materials, calculated as the ratio between the number of paintings where both materials are present divided by the number of paintings where either material are present.



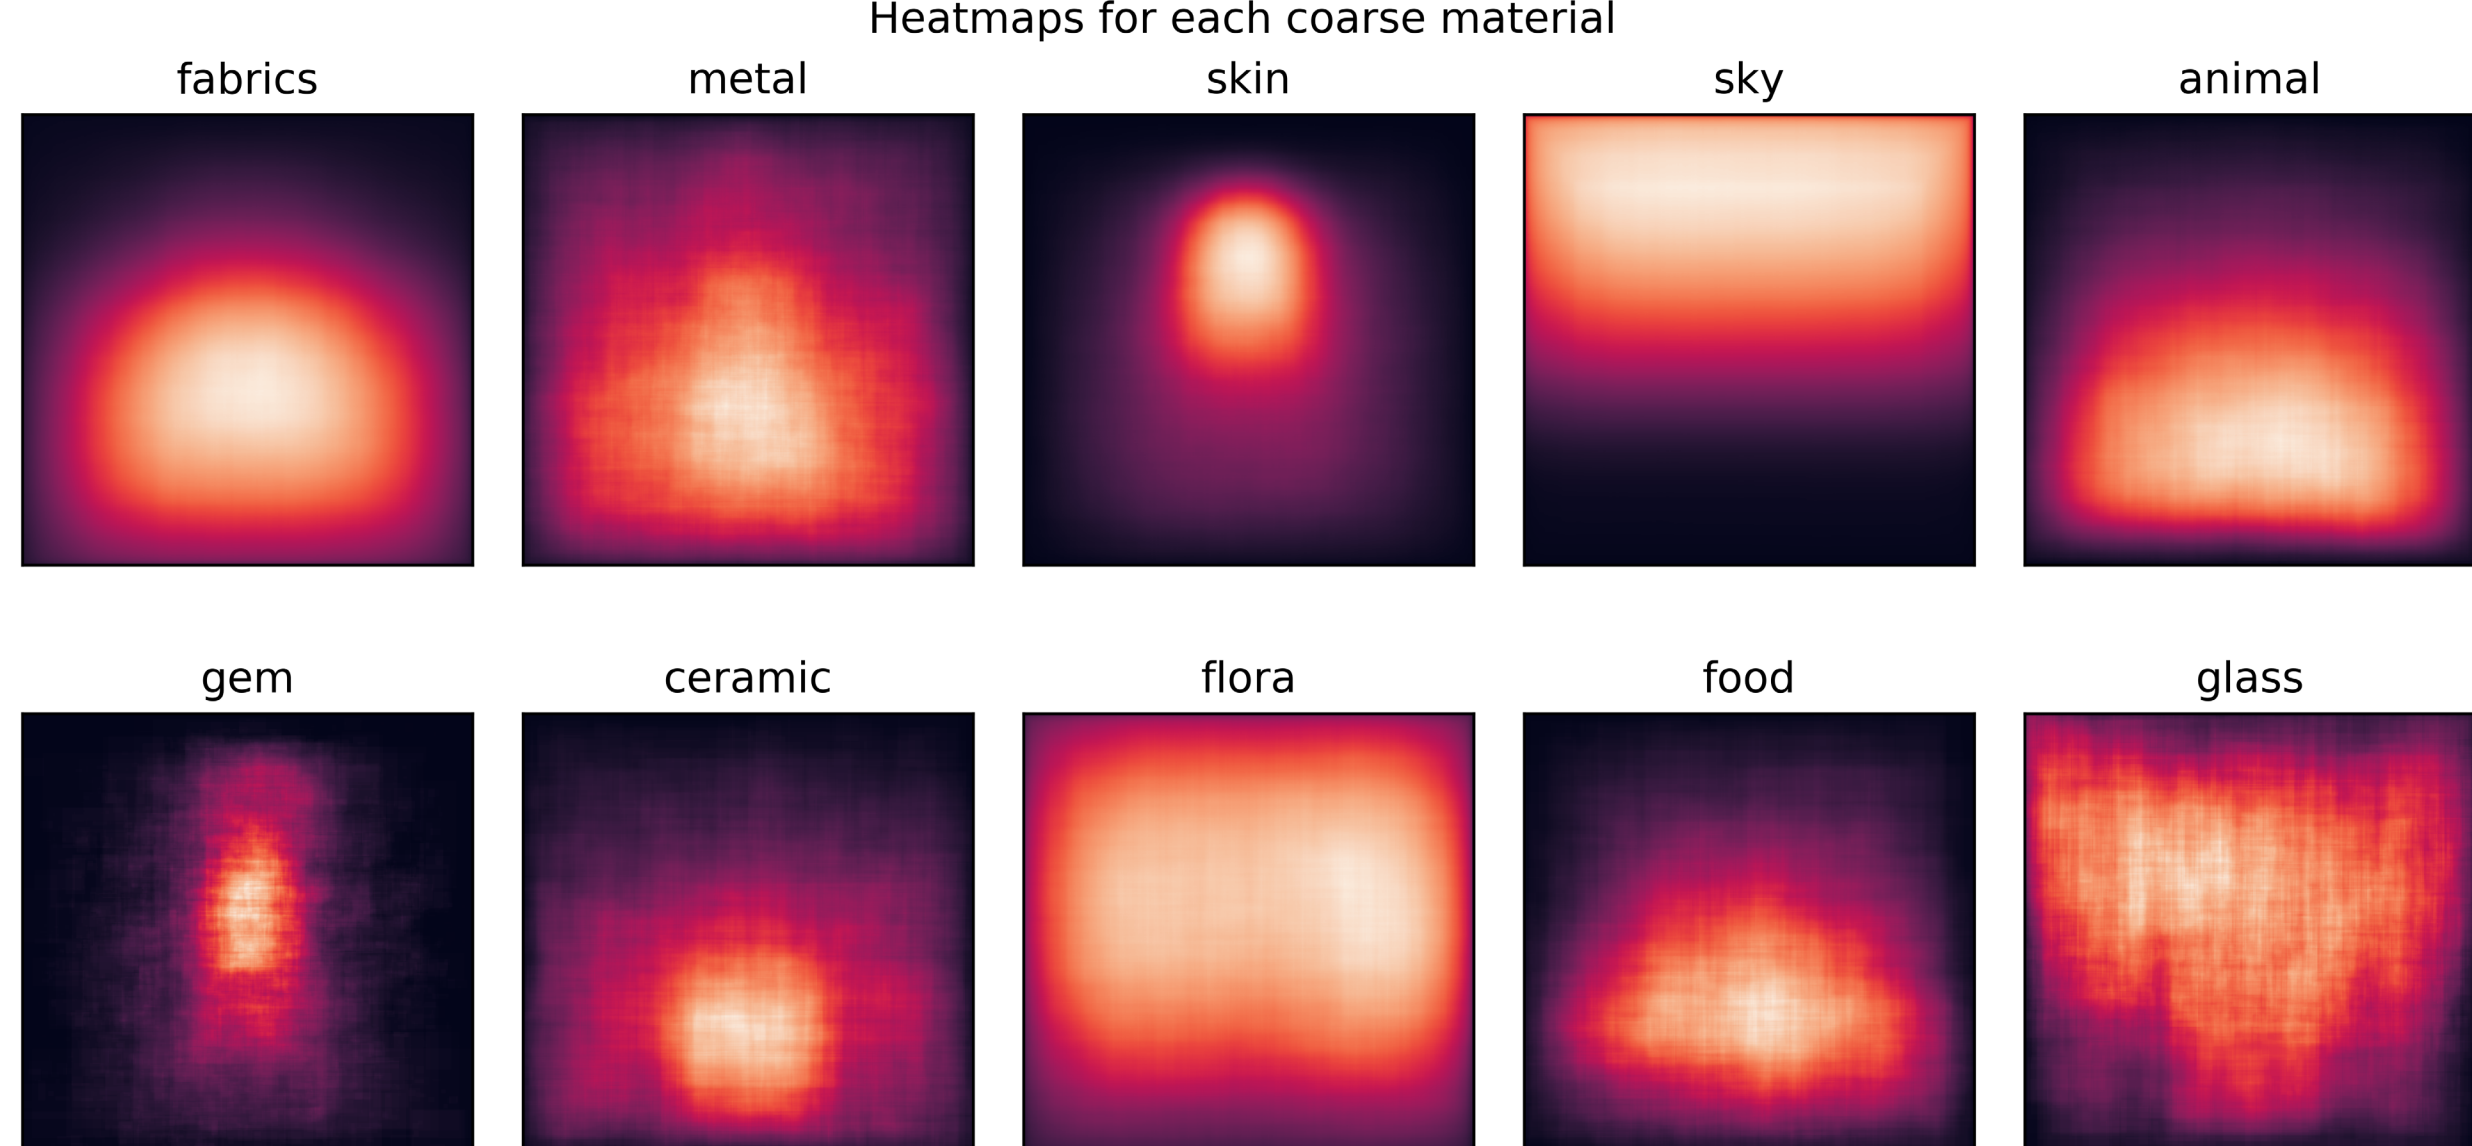
This matrix visualizes the influences a material has on the likelihood of finding another material within the same painting, that is, if one material on the y-axis is present, then how does this impact the presence of other materials on the x-axis? For example, if gemstones are depicted, then skin is depicted in 99% of the cases. However, if skin is depicted, then gemstones are depicted in only 20% of the cases. Calculated as the number of paintings where both materials are present, divided by the number of paintings that contains only one of the materials.

Collecting bounding boxes

Next, we selected 15 skilled AMT annotators to annotate more than 300k bounding boxes of the materials. This allowed us to identify the spatial location within paintings where materials are depicted.



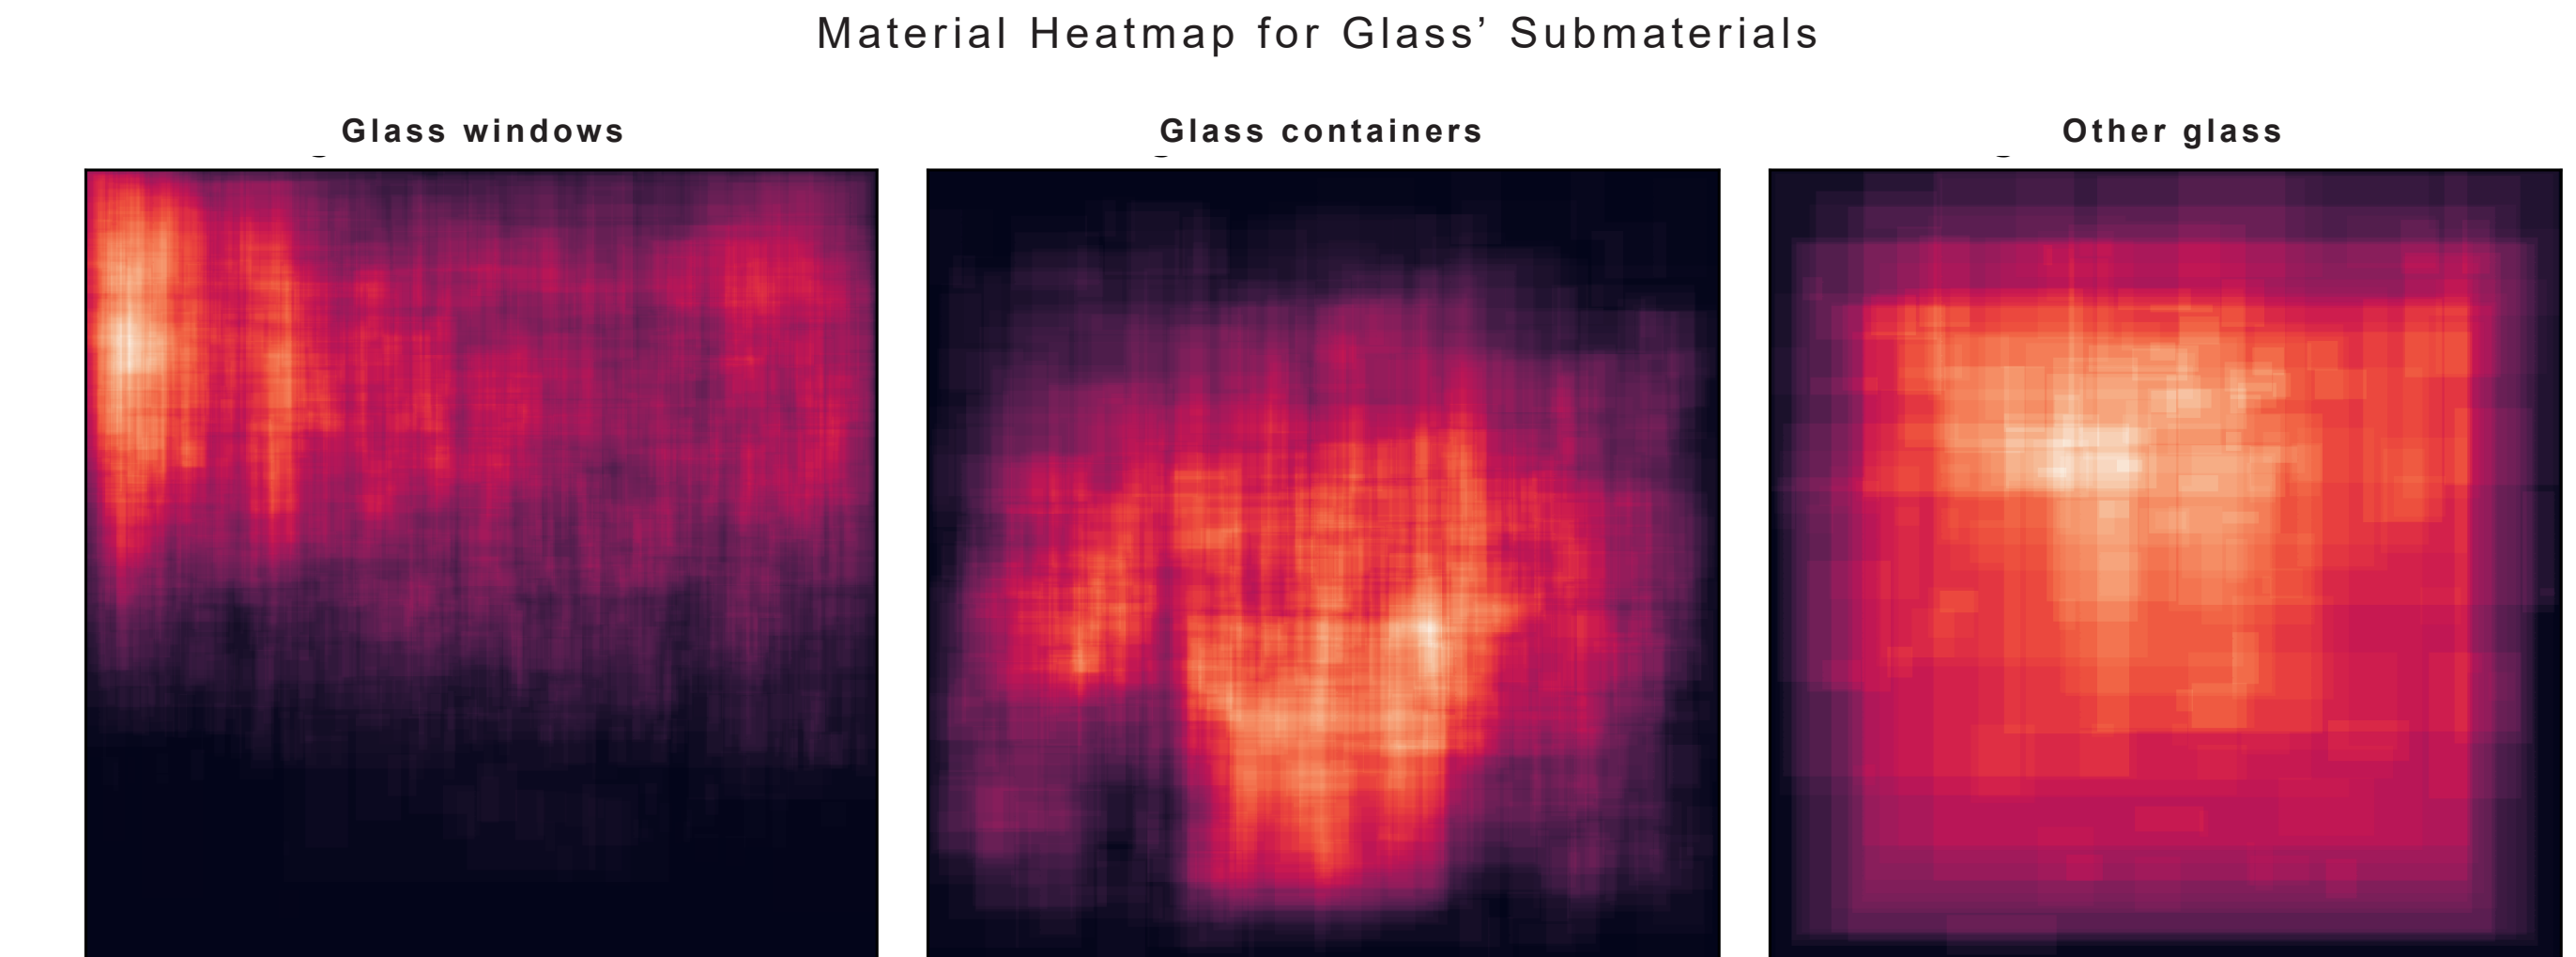
Two bounding boxes for fabrics are combined into a material heatmap. Overlapping regions becomes brighter. With hundreds of bounding boxes, the color gradients become smooth, with the brighter areas marking the spatial location where materials are more likely to be depicted.



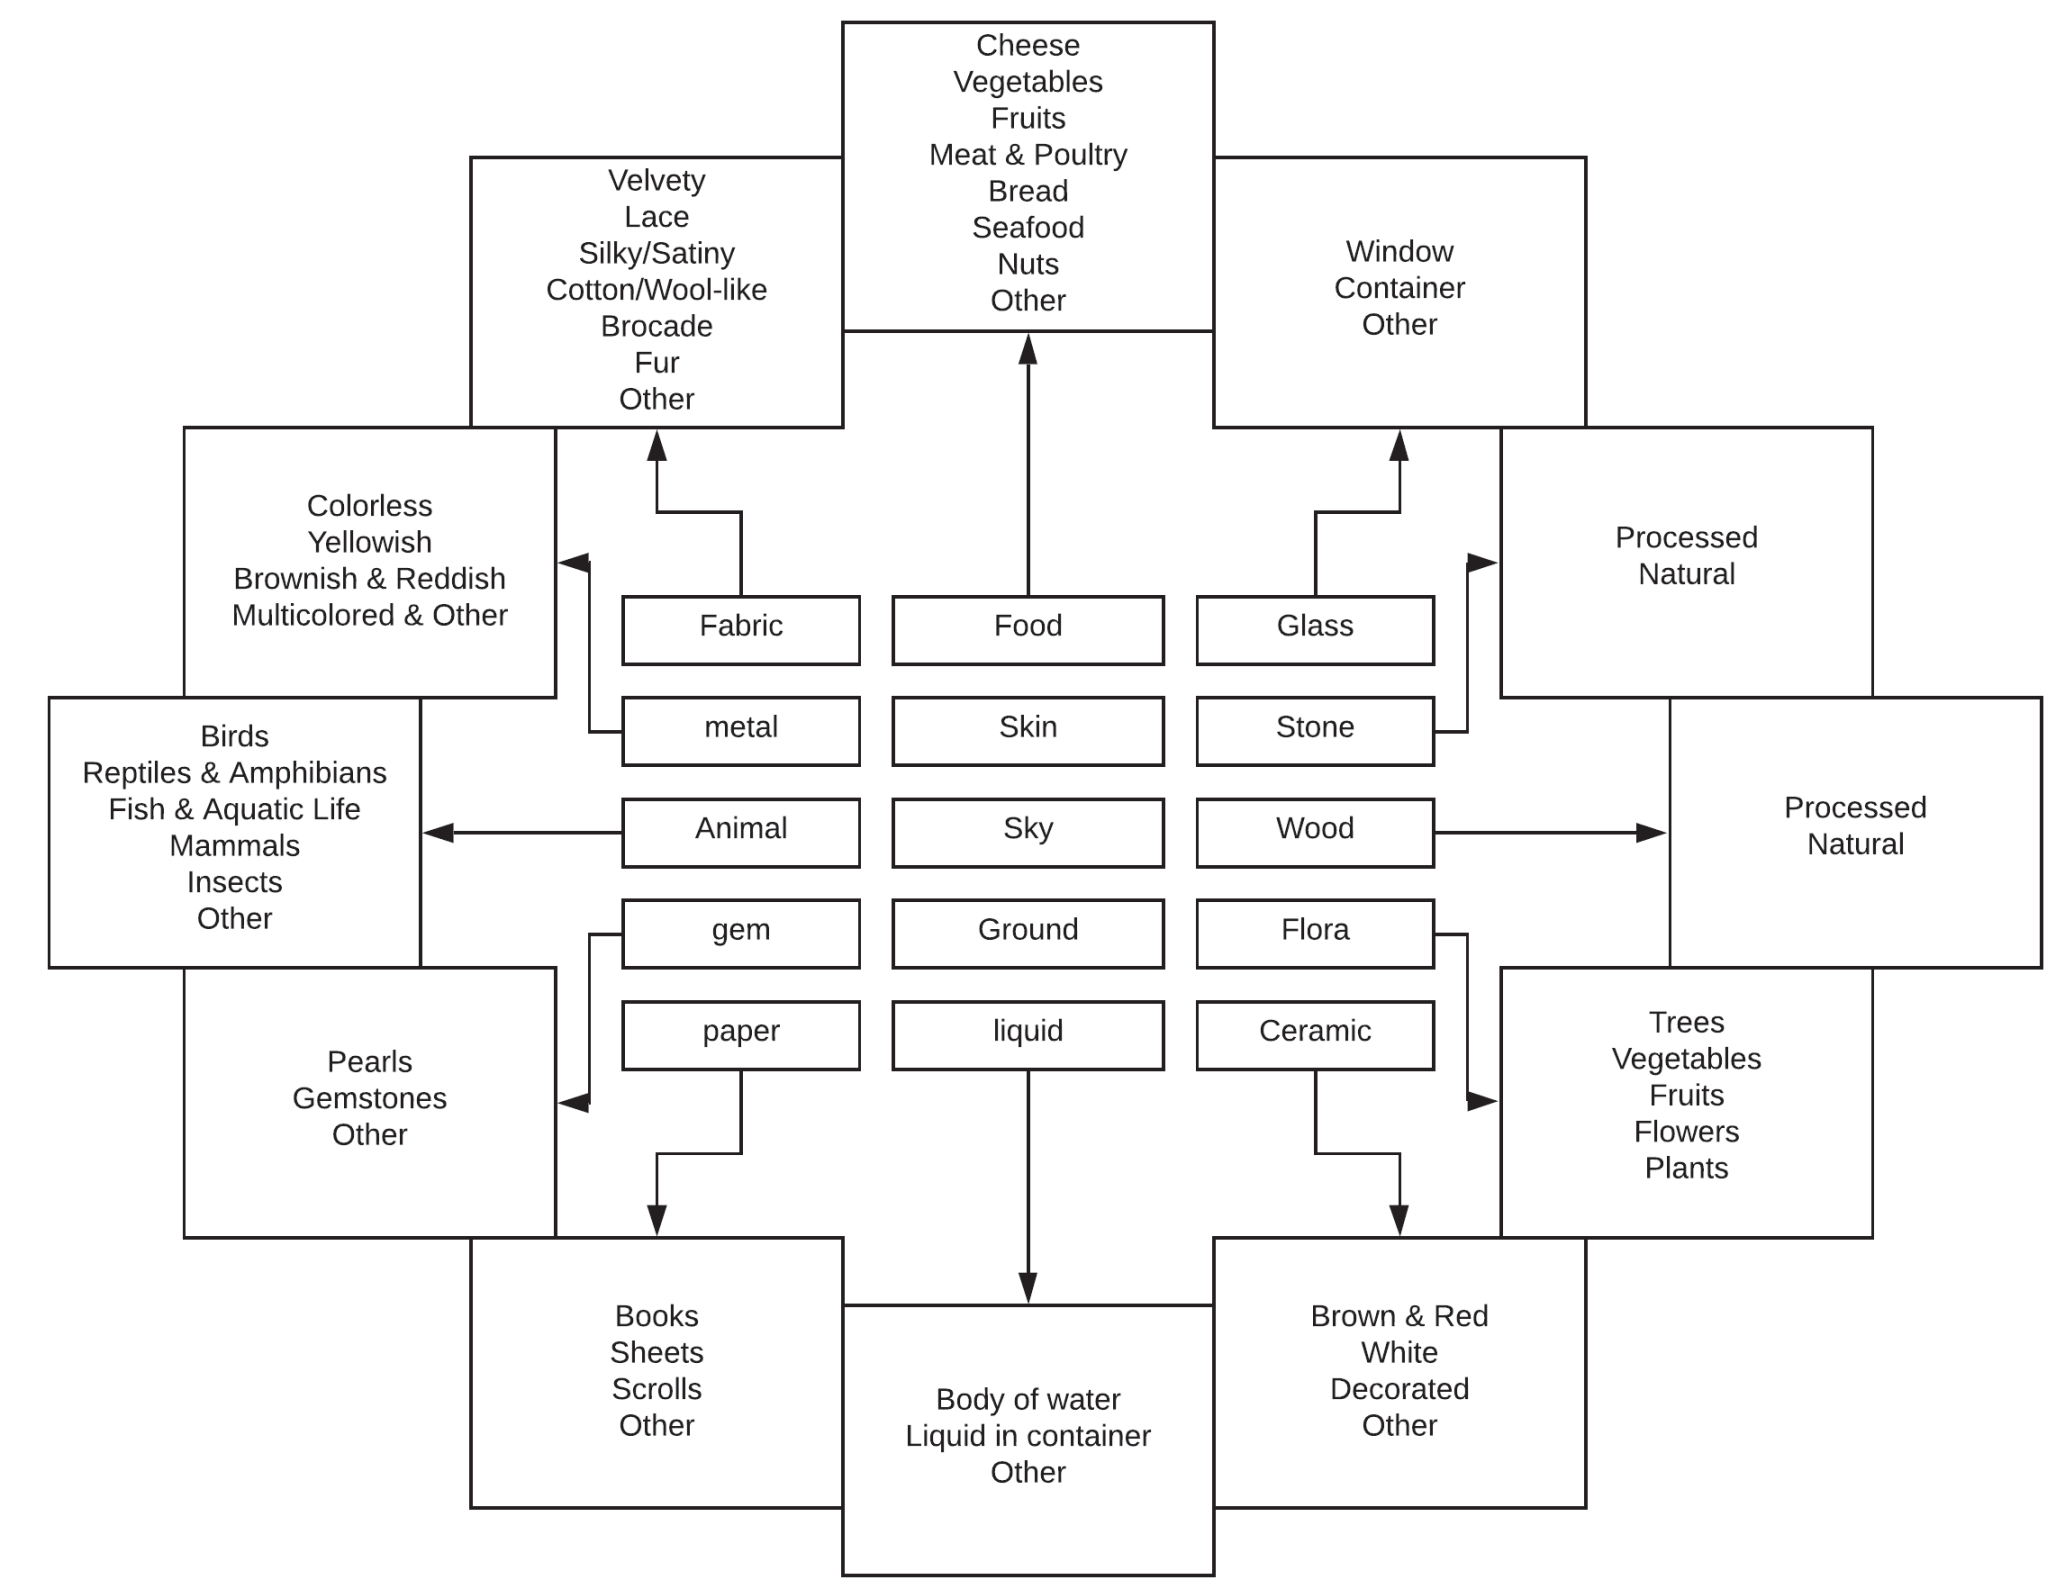
The material heatmap for 10 of the coarse-grained materials. In general, each material heatmap appears to be roughly vertically symmetric. For glass, there does however appear to be a minor shift towards the top-left. This might be related to an artistic convention, namely that light in paintings usually comes from the top-left. It is interesting to see how skin and gem are both vertically centered within the canvas. It appears to suggest a face, with necklaces and jewelry adorning the figure.

Collecting fine-grained labels

Last, for 13 out of 15 materials, we collected fine-grained material labels for the bounding boxes. For example, fabrics could now be labelled as velvet. See the diagram below which includes all materials and associated fine-grained labels. Participants would see a bounding box, and would select the most appropriate fine-grained label. If participants reached an 80% agreement, the box would be labelled as containing that fine-grained material. A total of 135460 boxes were assigned a fine-grained label.



The fine-grained labels allows us to make more specific material heatmaps. Here the fine-grained labels for glass are visualized. We can see that especially glass windows are most common in the top-left corner, as one could expect from the the previously mentioned artistic convention to have light originate from the top-left



The coarse materials (in the middle) and their respective fine-grained materials. Note that for 3 materials we did not define fine-grained materials and they are therefore visualized here without connections.

Applications

This in-depth dataset of material depictions can enable various perceptual, computational and historical analyses that could enable a deeper understanding of material perception and depiction. For an example of a study that uses stimuli from this dataset, see posters 481, 1520, 1741.

Acknowledgments

This work is part of a Vidi program with project number 276 54 001, which is financed by the Netherlands Organization for Scientific Research (NWO). This work is also supported in part by Natural Sciences and Engineering Research Council of Canada (NSERC) PGS-D3 516803 2018