DOI: 10.1007/978-3-031-51026-7_31
ACM DL: https://dl.acm.org/doi/abs/10.1007/978-3-031-51026-7_31

Towards Using Natural Images of Wood to Retrieve Painterly Depictions of the Wood of Christ's Cross

Johannes Schuiki[1] ♦ Miriam Landkammer[2] ♦ Michael Linortner[1]
Isabella Nicka[2] ♦ Andreas Uhl[1]

[1]Department of Artificial Intelligence and Human Interfaces, University of Salzburg, Austria
[2]Institute for Medieval and Early Modern Material Culture, University of Salzburg, Austria

Abstract

A painting, much like written text, allows future viewers to draw conclusions about the time of its creation or its painter. Also, when looking at a whole corpus of images instead a single instance, trends in painting can be analyzed. One particular trend originating in the 14th century is the transfer of the visual impression of real world materials onto paintings. One object, which is often depicted in paintings around that time being made from wood, is Christ's cross. Scarce research has been done in the direction of automatically analyzing painterly depictions of the wooden cross of Christ. Hence, this study walks a step towards automatic annotation of wooden crosses in paintings by evaluating three publicly available databases containing natural images of wood for their applicability to use their images as queries to retrieve painterly depictions of the wood of the cross. Experimental results underline the demand for further investigations.

Contents

This work has been partially supported by the Salzburg State Government project "How Material Came into the Picture: Exploring Cultural Innovations Interdisciplinarily with Artificial Intelligence (KIKI)."

1 Introduction

Throughout the 14th century, we can observe new trends in painting all over the Latin West. One of the most important of these was the imitation of the material of objects and things from nature. This trend culminated in the 15th century in a new realism which – stimulated in particular by the achievements of Early Netherlandish painting – brought with it a new peak in the preoccupation with the intricate rendering of surface qualities. With regard to the representation of materials, the roots and course of such a major and sustained cultural innovation have not been the subject of research on a European scale. Nor are there detailed analyses for the different ways, surface qualities were rendered in art of the time for a lot of materials. One important reason for this is that hardly any annotation data to depicted material are available, making even the formation of a research corpus a very difficult task. Grouping similar ways of artistic approach to surface qualities would allow to study both the various narrative functions and the proximity of artistic concepts. To fill this research gap, the KIKI project (title: "How the material got into the image: Exploring cultural innovation with DH and AI") develops automated recognition of material in historical visual media, to enable in-depth analyses of the representation of materials and surface qualities in painting across different geographical areas and time spans, and to apply a distant viewing approach [1], [2] that allows the recognition of underlying patterns and specific solutions in big image data.

In art history, the perspective of Materialikonologie – the analysis of the functions of the physical material of works of art – has been central to the study of medieval art since the late 1960s (most eminent: [3]). More recently, the concept of Materiality opened up the view to broader phenomena such as material interactions [4], thereby also drawing the attention to materials evoked, imitated or fictionalised by artists in images. Research has been conducted on various rendered materials, like rock (i.a.[5], [6]), marble (i.a. [7], [8] or hair, skin and (semi-)transparent materials [9], [10]. However, other materials, like rendered wood, have not been in the focus in a broader sense. The function of rendered materials in the narrative strategies of visual media is only touched upon by a few researchers, most notably Kim in [6] and [11].

The art historical part of KIKI therefore sheds light on how artists reflected or fictionalised materials through imitation in painting, and how they made depicted materials productive for the pictorial narrative. The research corpus consists primarily of works from today's Austria as well as neighbouring regions in the database REALonline. One of the challenges faced by art historical researchers is to cope with large amounts of – in this case visual – data. Thus, computer-aided annotation of big corpora of data is desired. While there exists literature on object retrieval [12] object detection [13], face retrieval [14] and object segmentation [15] in paintings, very scarce literature exists on detection and segmentation of painted materials.

Zuijlen et al. [16] conducted a study regarding the perception of painted materials and their attributes. They used existing segmentation software on a rather small set of images to segment 15 material categories. Later, the same group [17] published a material database named materials in paintings (MIP) along with bounding boxes encapsulating objects from the same 15 material categories used in their earlier study. They utilized Amazon's Mechanical Turk for the exhaustive manual annotation work. By fine tuning an existing object detection network using the manually annotated materials, the number of bounding boxes was artificially increased.

In an approach to close these gaps in both research disciplines, the art historical and computer vision, this study focuses on the painted representation of wood, especially wooden planks used to depict Christ's cross. Although the MIP database would contain wood as a material category, the annotated bounding boxes often contain background but no segmentation mask, hence deliver insufficient information for our experiments. As a first step towards detection of cross wood in paintings, the experiments within this research aim to answer the question whether patches of natural images of actual could be used as a query to retrieve patches of painted cross wood from a larger pool of image patches.

The remainder of this paper is structured as follows: Section 2 motivates why analyzing the depictions of crosses, especially Christ's cross, is of interest. Section 3 describes the experiments carried out in section 4 along with a detailed description of the used databases. Finally, section 5 summarizes the findings of the investigations.

2 Holy wood: The material of Christ's cross

In the Golden Legend, compiled around 1260, a dilemma – the fact that the Son of God died in a particularly ignominious, humiliating way on the cross [18, ch. 1] – is resolved: Christ's redemptive death makes the cross an instrument of salvation, which is demonstrated by the reversal of characteristics of the wood into their opposite (from cheap to precious, from ignoble to sublime, from dark – and without any beauty – to light, from malodourous to fragrant, etc.) [19, p. 554]. The passage shows how densely the notion of the cross of Christ was intertwined with its material in the Middle Ages and how the wood of the cross served as a connecting point for different ideas and forms of metaphor. The significance of the material is evident from "wood" being a common synonym for Christ's cross in Christian literature. Also, cross relics have played a role in the Latin West especially since the High Middle Ages [20]. Unlike the body of Christ, particles of the cross were accessible as material remnants of Christ's passion. Veneration of the cross formed part of several liturgical feasts and included rituals such as the pouring of wine or water over the relic of the cross, which had healing powers when drunk [21, pp. 18–23]. The *lignum crucis* was believed to have "absorbed" Christ's power through his blood and the touch of his limbs, and the cross could become a substitute for Christ's body. In both contexts, it is relevant that wood was frequently characterised as a living material, with veins and "humours" (body fluids), which, like the human body, could fall ill and be injured [22, pp. 252–254].

A link between the cross particles and the historical cross was provided by the immensely popular Legend of the Cross: the "biography" of the wood of the cross, which was in turn firmly embedded in the liturgical year [23]. The Legend of the Cross establishes a link to Old Testament time by tracing the origin of the wood back to a branch of the Tree of Life in paradise. It also reports on events and miracles that happened long after Christ's sacrifice (i.a. the finding and miraculous identification of the True Cross by the Empress Helena in the fourth century).

These ideas about the wood of the cross, as well as the liturgical and devotional practices revolving around it, made painters dedicate special attention to Christ's cross concerning not only its construction but also its texture. Much research has been done on tree-like or vegetal depictions of the cross as a prospect of eternal life, and establishing the link to the tree in the Edenic garden and the Tree of Life in the Book of Revelation [24], [25]. Depicted wood grains on crosses have however not been examined comparatively, which might partly result from the lack of data to systematically retrieve and compare all examples. The application of computer vision methods allows to examine the representation of the *lignum crucis* across a diverse corpus of images: In which pictorial contexts is the wood of the cross represented by the depiction of a grain? In which cases is there no texture, and where does the rendered materiality deviate decidedly from Christ's death-instrument (e.g. a golden cross)? A comparison of textures can provide an overview to the frequently depicted wood grains of Christ's cross, as well as outliers. On this basis further questions can be addressed, e.g. if a particularly abstract or unnaturally regular texture emphasizes the supernatural character of the wood, if there are textures that can be interpreted as references to other entities (e.g. Christ's body), if particularly explicit renderings of the wooden surface are due to the local context of use of the artworks (e.g. an existing cross relic) etc.

3 Experimental Setup

The conducted experiments aim to evaluate the similarity of painterly depictions of Christ's cross and the two thieves' crosses to natural images of actual wood. To do so, a simple query-based scheme is employed: Image patches containing real wood (query) are compared with image patches containing painted cross wood as well as patches containing other painted content (i.e. not cross wood). Since every comparison yields a similarity score, the closest N patches can be determined for every query. The main metric for the experiments in this study is the proportion of painted cross wood within the closest N patches.
Before comparison, images are cut into smaller, quadratic tiles (databases included in this study along with their corresponding tiling strategies are described in section 3.1) and features are extracted from every tile using well-known texture descriptors. To gain a similarity measure between two tiles, the Euclidean distance is calculated on their corresponding feature vectors for a certain feature extraction method. Hence, a lower distance indicates a higher similarity between two patches. Experiments in this work are evaluated using three tiling sizes (64×64 px, 96×96 px, and 128×128 px).

In total, seven distinct techniques for feature extraction are employed, briefly introduced in the following. The first six of the algorithms were successfully used in classical texture classification, image tampering detection and paper identification [26]: Dense SIFT [27] (SIFT descriptors applied in a grid), Local Binary Pattern (LBP) [28], Weber Pattern (WP) [29], Local Phase Quantization (LPQ) [30], Rotation Invariant - LPQ [31] and Binarized Statistical Image Features (BSIF) [32]. While most of the extracted features are used without adaptation, the features from SIFT further undergo a dimensionality reduction using principle component analysis and soft-quantization utilizing a Gaussian mixture model, hereby minimizing the spatial-dependent information within the extracted features. In addition to that, a pre-trained ResNet50 [33] is employed, where the fully connected layer is skipped, resulting in an 2048 dimensional feature vector.

The experiments in section 4 can be divided into four parts and the corresponding research questions defined as:

A) Is it generally feasible to employ image patches containing real wood as queries to successfully retrieve patches containing painted cross wood?

B) How does the behaviour change using patches of painted wood as query instead of real wood images?

C) Can image patches from other real objects (other than wood) be used to receive similar results as in experiment A?

D) Do variations in tiling size of query and test images have an effect on the results?

3.1 Databases & Tiling

Hereafter, the used databases and their corresponding tiling strategies are explained in detail. The number of resulting patches per database is presented in Table 1. Examples for each database can be seen in Fig. 1, Fig. 2 and Fig. 3.

Table 1: Number of available tiles per database and tile size.

	#Patches	Patch size [px] 64	96	128
Database	REALonline-Wood	1725	448	149
	REALonline-Non-Wood	1725	1725	1725
	Sharan14	3346	1255	535
	Elzaar21	6877	2955	1608
	Santos21	8598	3557	1364

REALonline : Unlike many other image databases of medieval and early modern visual cultural heritage, in REALonline [34], [35] all image elements are recorded by name, category (i.e. clothing), colour, form and (when possible) material. The database contains more than 28.000 digital and digitised photographs and more than 22.500 in records on works of arts or their components. There are more than 1.2 million high-level annotations on image elements, manually recorded over almost 50 years by database editors trained in a discipline of medieval or early modern studies. For this research, we have selected 115 images based on the query of the existing annotation data. We also restricted the works in the dataset to painting and drawing techniques, and according to whether they were produced in the 14^{th} or 15^{th} century.

To locate Christ's cross and the two thieves' crosses, images were annotated manually as can be seen in Fig. 1-B. The annotation mask separates regions where tiles for the REALonline-Wood can be sampled (white areas) and background regions (black) where tiles for the REALonline-Non-Wood can be extracted. In case of the REALonline-Wood data, quadratic patches are cropped from a every annotated polygonal shape. This is accomplished by first applying morphological erosion with a squared structure element of size $S \times S$ to the annotated bitmask. S here is the size of the desired patch resolution. Doing so, every remaining logical 1 constitutes a center point of a patch that is guaranteed to be within the annotated polygon. The first tile is cropped around the left-most center point. Afterwards, the area is removed from the annotated bitmask. This procedure is repeated until no further valid patch is available. The construct the REALonline-Non-Wood database, 15 tiles per image are randomly cropped at locations which do not overlap with the annotated wooden areas.

Flickr Material Database (Sharan14) [36] : The Flickr Material Database was constructed by [36] with the goal of capturing the natural range of material appearances. It comprises of 500 images from ten material categories. For the purpose of this work, only the images from the category *wood* are used. Since some wood images from this database also contain background, segmentation masks for the wooden areas are also available. In order to only draw samples stemming from the wooden areas, the same algorithm which is also used for tiling the REALonline-Wood data is employed. For experiment C (mentioned earlier this section), two other image categories (*stone* & *water*) from this database are used as query images.

Plastic and Wood Pollution Dataset (Elzaar21) [1]: This dataset consists of 100 images from two classes. Similarly to the Flickr Material Database, only the class *wood* is used, resulting in 50 wood texture im-

[1] https://www.kaggle.com/datasets/abdellahelzaar/plastic-and-wood-pollution-texture-dataset

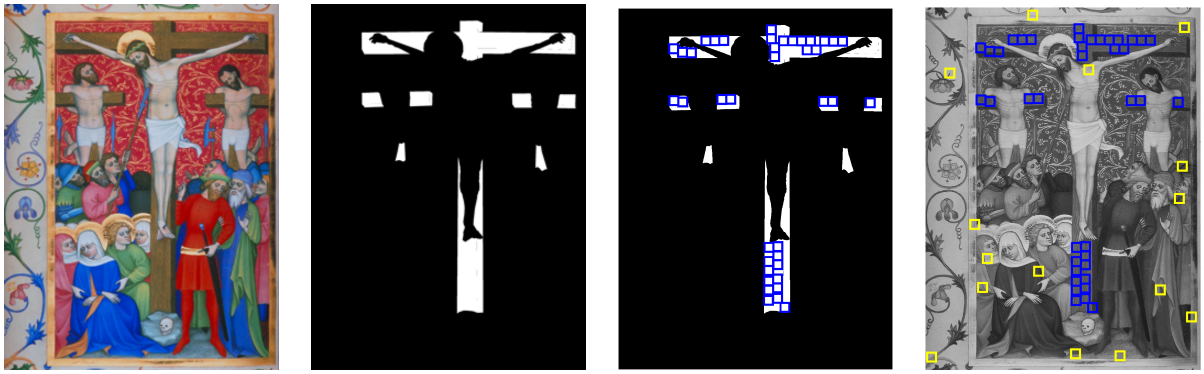

Figure 1: Patch extraction pipeline for the REALonline data: A) example image from database, B) manually annotated mask of wooden crosses of Christ and the thieves, C) found patches of certain patch size within wooden areas, D) wooden patches (red) and 15 randomly chosen non-wooden patches (yellow).

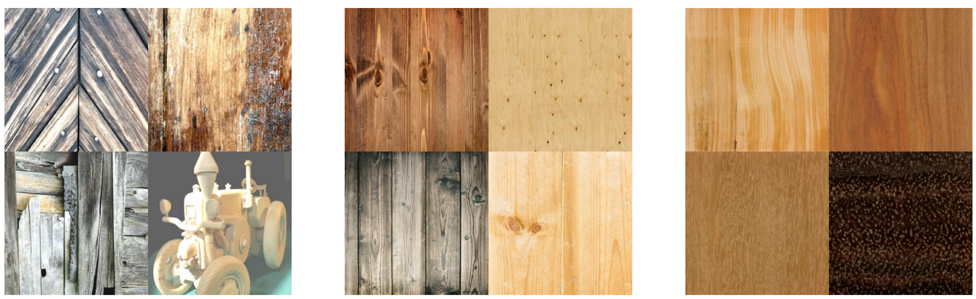

Figure 2: Examplary depiction (before tiling) of images from the three query databases (resized to quadratic shape to fit the grid).

ages of varying resolution. Tiling is accomplished by cropping adjacent patches without overlaps, starting on the upper left corner. Upon reaching the right image border, crops are continued underneath the just finished row of patches.

The Wood Database Images (Santos21) [2]: This database consists of images from the website https://www.wood-database.com/ acquired via web-scraping. Without modification, it comprises of 2105 images of varying resolution. For the purpose of automatically cropping patches from wooden areas, a subset of 1463 images has been selected by hand, only keeping images where solely wood is visible. Tiling is done in a similar manor as with the pollution dataset.

[2] https://www.kaggle.com/datasets/edhenrivi/the-wood-database-images

4 Experimental Results

This section contains the results from the experiments motivated in section 3. Although the REALonline-Non-Wood data set consists of 1725 patches for every tiling size (see Tab. 1), the number of patches are artificially reduced in a random fashion for the experiments where less patches from the REALonline-Wood are available in order for the pool of patches to be balanced. Note that data points within a plot always constitute the arithmetic mean over all patches from a set of query images.

Experiment A : Images from the databases containing real wood images are used as query to retrieve N closest samples from a pool of images combining all samples from REALonline-Wood and REALonline-Non-Wood . The proportion of retrieved wood patches within the N closest samples is analyzed and the results are depicted in Fig. 4.

Per row of plots, one particular real wood database

Figure 3: Randomly picked tiles of size 128×128 px from the 5 wood databases. Note that for the six classical texture features, images were converted to grey-scale prior to the experiments.

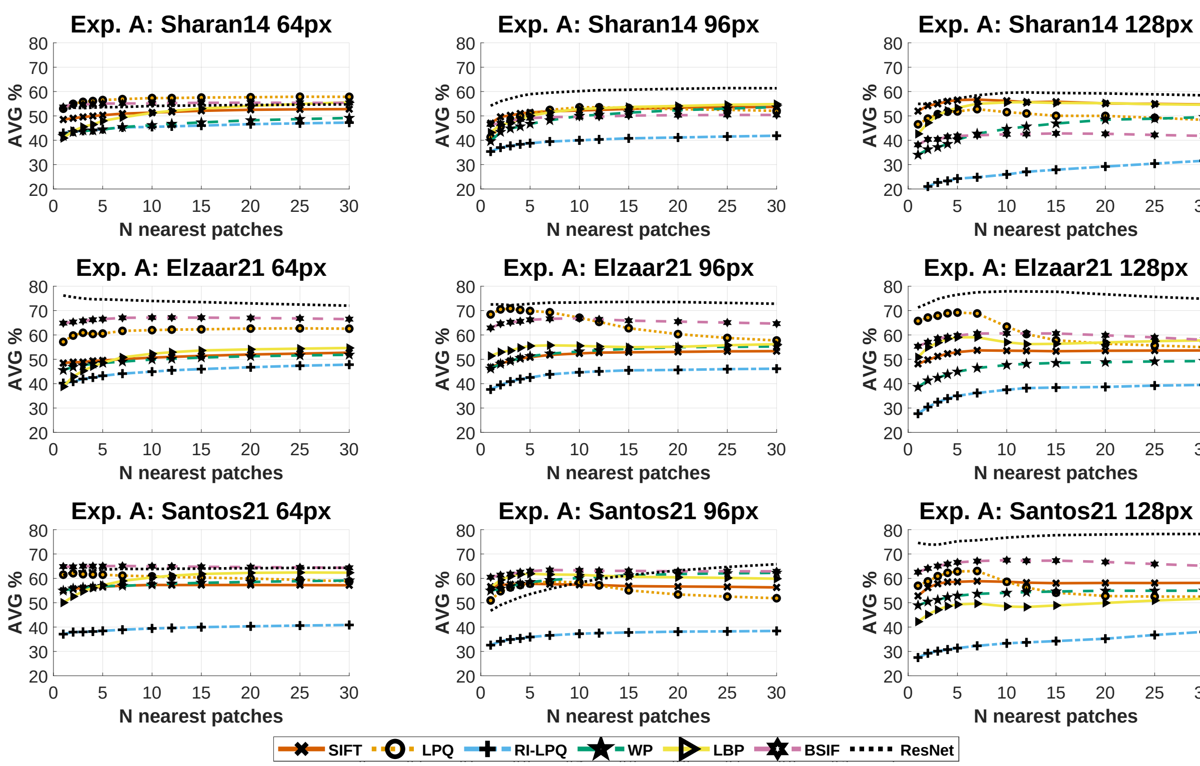

Figure 4: Results experiment A.

was used for the query patches. Tiling size increases from left to right. LPQ and BSIF often perform better than SIFT and LBP. RI-LPQ is guessing at best (i.e. always near or below 50%). Query images from the Sharan14 database (first row) generally result in guessing (below 60% on average). The ResNet features often perform best. Using image patches from Elzaar21 or Santos21 as queries, retrieval performance can reach up to almost 80%.

Experiment B : Since results from the first experiment seem rather inconclusive, the second experiment uses a subset of the REALonline-Wood database as query images. In order to still include all tiles in the results, a 2-fold cross validation is employed. The plots (see Fig. 5) show better results for increasing tile size, however the ResNet and SIFT seems to be perform best for any size.

Results indicate, at least for small values of N, tiles containing painted cross wood can indeed be used to find other tiles containing also painted cross wood rather than other painted contents. This experiment can be viewed as a sanity check whether the used feature descriptors are actually capable of capturing the structure within the tiles of painted cross wood.

Experiment C : This experiment is meant to evaluate the outcome when different natural images (other than wood) are used as queries. Two databases are tested. Results are depicted in Fig. 6. For Sharan14-stone (upper row) can be said that, on average, at best 50% of image tiles with highest similarity contain painted wood. Especially for the case 128 px, where results were best in experiment A, results tend to contain a higher amount of non-wooden content. Different results can be observed when looking at Sharan14-water (lower row), where the average paint wood hit-rate even reaches 70%.

Experiment D : This experiment is meant to investigate the effect of scaling. Tiles are cropped at a certain size and then scaled to another tile size. The experiment can be separated into two parts: First, only the query tiles are scaled and compared to un-scaled tiles from the REALonline databases (see Fig. 7-left). Second, tiles from query databases and REALonline databases are cropped with different tile sizes but then scaled to the same size in order to make them comparable (see Fig. 7-right). Note that for the experiment D, only the BSIF feature descriptor is used exemplary on the Sharan14 database. In order to compare the results to the results from experiment A, the pink dash-dotted line with hexagrams in the first row of Fig. 4 needs to be looked at. For better visualization, these data points are included in Fig. 7 as dashed lines. The colors correspond to one certain (initial) tile size for the painted tiles. The results from experiment A stagnate on around 55%, 50% and 42% for 64 px, 96 px and 128 px, respectively. A mostly similar behaviour can be observed for experiment D. Thus, it can be concluded that scaling entire databases does not improve the ability to retrieve painted wood.

5 Conclusion

This study analyzed the usability of natural images containing real wood to filter image patches containing painted cross wood from a pool of images where also other painted motives are depicted. Several experiments were conducted using a three wood databases and varying patch sizes. It was found that experimental results often show a random characteristic and also the behaviour is very similar to using a different material as query. Experiments using patches of painted wood as query showed that the employed feature descriptors should be principally sufficient to capture the textural characteristics from the depicted wood. However, it appears the gap between the domains of natural wood and painted wood, at least for the used databases, is too high to deliver reliable results.

Future work will include an attempt to overcome the gap in domain by trying to model the domain shift using e.g. image to image (I2I) translation. Also paint to paint results hold the potential of improvement by resorting to more sophisticated learning based schemes. Furthermore, the annotation database will be extended. Once enough data is available, instance segmentation of crosses using state of the art segmentation tools will be done.

References

[1] T. Arnold and L. Tilton, "Distant viewing: Analyzing large visual corpora", *Digital Scholarship in the Humanities*, vol. 34, no. Supplement_1, 2019.

[2] H. Klinke, "The digital transformation of art history", in *The Routledge Companion to Digital Humanities and Art History*. Routledge, 2020, ch. 3.

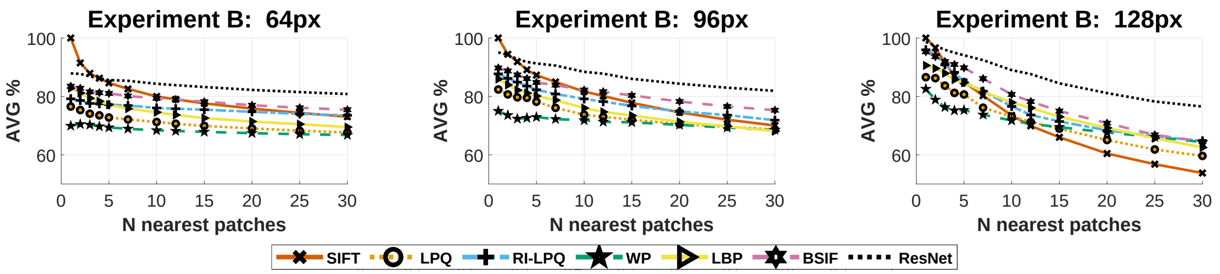

Figure 5: Results experiment B. Using tiles containing painted wood as query.

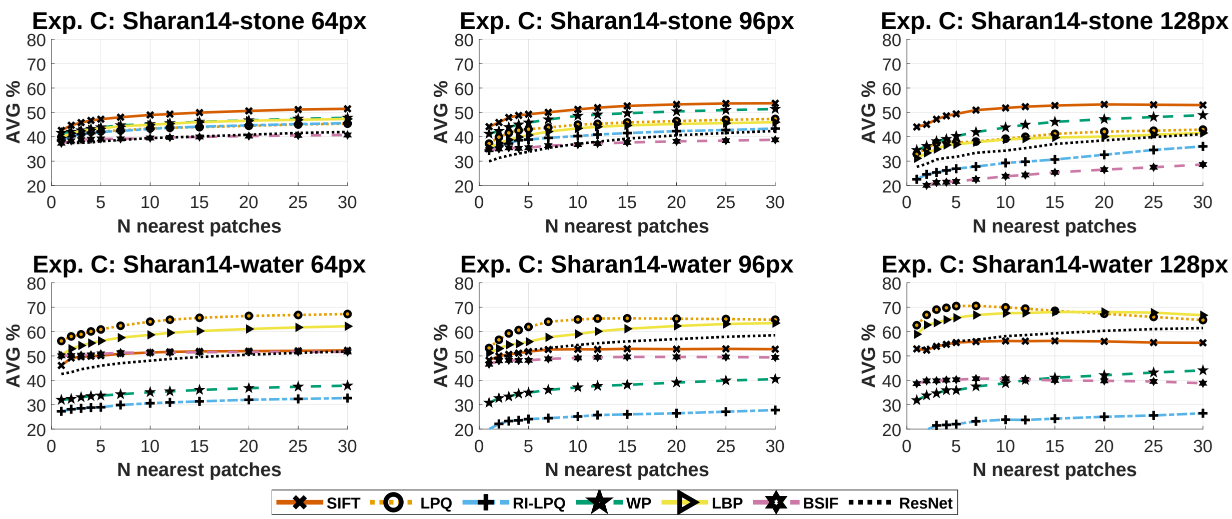

Figure 6: Results experiment C. Using images containing stone and water as query.

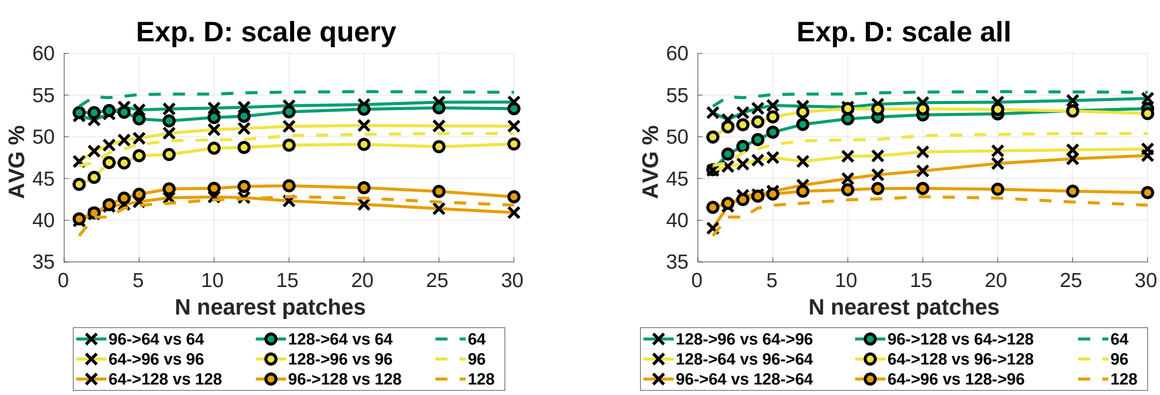

Figure 7: Results experiment D. The legend shows the initial crop size and the final size after resizing. The first part (before vs) describes the cropping & scaling for the the query tiles, while the second part (after vs) belongs to the REALonline databases.

[3] T. Raff, “Die sprache der materialien”, Zugl.: Augsburg, Univ., veränd. Habil.-Schr., 1991, Dt. Kunstverl, 1994.

[4] A. Kumler, “Materials, materia, “materiality””, in *A Companion to Medieval Art*. John Wiley & Sons, Ltd, 2019, ch. 4.

[5] J. Tripps, “”Nimm Steine, rauh und nicht gereinigt.” Cennino Cennini und das Gebirge als kunsttheoretisches Studienobjekt in der Malerei des Trecento”, in *Fantasie und Handwerk. Cennino Cennini und die Tradition der toskanischen Malerei*, 2008.

[6] D. Y. Kim, “Stonework and Crack in Giovanni Bellini’s St. Francis in the Wilderness”, in *Steinformen: Materialität, Qualität, Imitation*, De Gruyter, 2019.

[7] I. Mestemacher, “Marmor, gold und edelsteine”, PhD thesis, De Gruyter and Walter de Gruyter GmbH & Co. KG, 2021.

[8] J. Kölle, “Versteinertes Blut. Heilig-Blut-Säulen in der flämischen Malerei um 1500”, in *Steinformen: Materialität, Qualität, Imitation*, De Gruyter, 2019.

[9] A.-S. Lehmann, “Small hairs: Meaning and material of a multiple detail in the ghent altarpiece’s adam and eve panels”, English, in *Van Eyck Studies*. Leuven University Press; Uitgeverij Peeters, 2016.

[10] M. Lehman Anne-Sophie; Bol, “Painting skin and water: Towards a material iconography of translucent motifs in early netherlandish painting”, in *Rogier van der Weyden in context*, 2012.

[11] A. Degler, *Parergon: Attribut, Material und Fragment in der Bildästhetik des Quattrocento*. Verlag Wilhelm Fink, 2015.

[12] E. Crowley and A. Zisserman, “The state of the art: Object retrieval in paintings using discriminative regions”, 2014.

[13] E. Crowley and A. Zisserman, “The art of detection”, vol. 9913, 2016.

[14] E. Crowley, O. Parkhi, and A. Zisserman, “Face painting: Querying art with photos”, British Machine Vision Association, 2015.

[15] N. Cohen, Y. Newman, and A. Shamir, “Semantic segmentation in art paintings”, *Computer Graphics Forum*, vol. 41, no. 2, 2022.

[16] M. van Zuijlen, S. Pont, and M. Wijntjes, “Painterly depiction of material properties”, *Journal of Vision*, vol. 20, 2020.

[17] M. J. P. van Zuijlen, H. Lin, K. Bala, S. C. Pont, and M. W. A. Wijntjes, “Materials in paintings (mip): An interdisciplinary dataset for perception, art history, and computer vision”, *PLOS ONE*, vol. 16, no. 8, 2021.

[18] R. M. Jensen, *The Cross: History, Art, and Controversy*. Harvard University Press, 2017.

[19] Jacobus de Voragine, *The golden legend: Readings on the saints*, 2012 Edition. Princeton University Press, 2012.

[20] G. Toussaint, *Kreuz und Knochen: Reliquien zur Zeit der Kreuzzüge*. Reimer, 2011.

[21] L. van Tongeren, “Ein heilsames Zeichen: die Liturgie des Kreuzes im Mittelalter”, in *Das Kreuz*, ser. Regensburger Studien zur Kunstgeschichte, Schnell & Steiner, 2013.

[22] M. Pastoureau, “Introdulehmction à la symbolique médiévale du bois”, in *Le bois et la ville*, ser. Cahiers de Fontenay, Ecole Normale Supřieure de Fontenay Saint-Cloud, 1991.

[23] B. Baert, *A heritage of holy wood: The legend of the true cross in text and image*, ser. Cultures, beliefs and traditions. Brill, 2004, vol. 22.

[24] C. Irvine, “The iconography of the cross as the green tree”, in *The Edinburgh companion to the Bible and the arts*, Edinburgh Univ. Press, 2014.

[25] B. E. Kitzinger, *The Cross, the Gospels, and the Work of Art in the Carolingian Age*. Cambridge University Press, 2019.

[26] C. Kauba, L. Debiasi, R. Schraml, and A. Uhl, “Towards drug counterfeit detection using package paperboard classification”, in *Advances in Multimedia Information Processing – Proceedings of the 17th Pacific-Rim Conference on Multimedia (PCM’16)*, ser. Springer LNCS, vol. 9917, 2016.

[27] D. G. Lowe, “Distinctive image features from scale-invariant keypoints”, *Int. J. Comput. Vision*, vol. 60, no. 2, 2004.

[28] T. Ojala, M. Pietikainen, and D. Harwood, “Performance evaluation of texture measures with classification based on kullback discrimination of distributions”, in *Proceedings of the 12th IAPR International Conference on Pattern Recognition*, vol. 1, 1994.

[29] G. Muhammad, “Multi-scale local texture descriptor for image forgery detection”, in *Industrial Technology (ICIT), 2013 IEEE International Conference on*, 2013.

[30] V. Ojansivu and J. Heikkilä, "Blur insensitive texture classification using local phase quantization", in *International Conference on Image and Signal Processing*, 2008.

[31] V. Ojansivu, E. Rahtu, and J. Heikkila, "Rotation invariant local phase quantization for blur insensitive texture analysis", in *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, 2008.

[32] J. Kannala and E. Rahtu, "Bsif: Binarized statistical image features", *Proceedings of the 21st Int. Conf. on Pattern Recognition (ICPR2012)*, 2012.

[33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition", in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.

[34] I. Nicka, "Object links in/zu bildern mit realonline analysieren", ger, in *Object Links – Dinge in Beziehung*. 2019.

[35] I. Matschinegg, I. Nicka, C. Hafner, M. Stettner, and S. Zedlacher, *Daten neu verknoten: Die Verwendung einer Graphdatenbank für die Bilddatenbank REALonline*, ser. DARIAH-DE working papers. Niedersächsische Staats- und Universitätsbibliothek Göttingen, 2019, vol. 31.

[36] L. Sharan, R. Rosenholtz, and E. H. Adelson, "Accuracy and speed of material categorization in real-world images", *Journal of Vision*, vol. 14, no. 10, 2014.