https://doi.org/10.2352/EI.2025.37.8.IMAGE-271

Identification of Cultural Artifacts using Deep Learning

Huajian Liu, Xiaoying Yang, Raphael Antonius Frick, Martin Steinebach
Fraunhofer Institute for Secure Information Technology (SIT), Darmstadt, Germany

Abstract

This work addresses the challenge of identifying the provenance of illicit cultural artifacts, a task often hindered by the lack of specialized expertise among law enforcement and customs officials. To facilitate immediate assessments, we propose an improved deep learning model based on a pre-trained ResNet model, fine-tuned for archaeological artifact recognition through transfer learning. Our model uniquely integrates multi-level feature extraction, capturing both textural and structural features of artifacts, and incorporates self-attention mechanisms to enhance contextual understanding. In addition, we developed two different artifact datasets: a dataset with mixed types of earthenware and a dataset for coins. Both datasets are categorized according to the age and region of artifacts. Evaluations of the proposed model on these datasets demonstrate improved recognition accuracy thanks to the enhanced feature representation.

Introduction

Cultural assets have enormous cultural, historical, artistic and scientific significance. From ancient artifacts to contemporary artworks, cultural assets are like mirrors reflecting different facets of human history, artistic innovations and the evolution of social norms [1]. Therefore, the protection of these cultural properties is of paramount importance. However, illicit trafficking of cultural property poses a direct threat to heritage preservation and has become an increasingly pressing issue [2-5].

In combating the trafficking of illicit cultural artifacts, it is a challenge task for law enforcement and customs and other involved authorities to identify the provenance of cultural artifacts. Since identifying the origin of cultural objects requires strong specialized expertise, and law enforcement and customs officials do not have the relevant know-how, they often need the help and support of relevant experts. However, experts are not readily available to provide professional support and do not have sufficient capacity to identify all suspicious objects in a professional manner.

AI-based artifact identification technologies can bridge this knowledge gap and provide immediate on-site assistance, allowing the law enforcement agencies and customs to make initial assessments of the provenance of artifacts immediately.

Identifying the geographical and temporal origins of cultural assets is an important step in the protection of cultural heritage. Such knowledge not only facilitates the academic study and interpretation of cultural objects, but also allows customs and law enforcement officials to make informed preliminary assessments of the cultural goods they encounter. Therefore, this step is essential to combating illicit trafficking in cultural property.

The goal of this work is to develop a deep learning model to identify the provenance of unknown cultural goods, including the chronological and geographical origin. The identification of archaeological artifacts poses new challenges because of their unique characteristics. Especially for those illegally excavated artifacts, they cannot be identified by image matching because they have never been documented.

Unlike general object recognition, identifying artifacts cannot be done by considering only their appearance features, such as shape and color, because similar-looking artifacts may be vastly different in age and geographic origin, and are unrelated to each other. On the contrary, even different types of cultural artifacts, if they come from the same age and the same cultural region, will often share similar cultural imprints and characteristics. Therefore, in order to correctly identify cultural artifacts, it is essential to take into account their age, geographic origin and their cultural background, etc., and to extract the features that are closely related to their age and cultural background, including shape, decoration, symbols, and so on.

Therefore, firstly, it is necessary to establish a targeted and appropriate network model for the identification of cultural artifacts. Most existing deep learning models are developed and trained for general image recognition and not specialized for identification of archaeological objects. Secondly, it is also required to create dedicated datasets of cultural artifacts for training and testing the model. To the best of our knowledge, there are no publicly available datasets of this type that can be used directly for model training and testing. Available data from museums are neither fully labeled nor uniformly tagged.

In our previous work [6], we created an artifact dataset, named SMB dataset, in which the artifact types were not categorized and all types of objects were mixed in each class. This means that each class contains a variety of artifacts from the same era and cultural context, including pottery, coins, sculptures, print material, portrait, clothes, and so on. However, due to the lack of comprehensive metadata, SMB dataset had to rely on museum collections as a proxy for actual geographic provenance, resulting in an underrepresentation of this key attribute. This also limited the predictive accuracy of models trained on this dataset for the geographic origin of cultural goods.

In this work, we created two new datasets of archaeological objects by collecting, cleansing and processing publicly available museum data, because a large amount of accurately labeled data is essential to train and fine-tune a model for recognition of cultural artifacts. The two datasets are both type-specific but are of two different kinds. One contains different types of objects belonging to earthenware, while the other includes only coins. All objects in both datasets are labeled and classified based on the geographical origin and age of the object. In the Earthenware dataset, the geographical origin is fine-grained to modern countries, while in the Coin dataset, it is fine-grained to ancient empires or kingdoms such as the Western Roman Empire.

In addition, because archaeological artifacts often carry both textural and structural features and convolutional operations has limited ability to capture long-distance dependencies across image regions corresponding to structural features, the CNN model proposed in [6] has been improved by introducing a self-attention mechanism in the feature extraction at each level. Thus, the local features in the feature maps at each level are enriched with extensive contextual information and long-range dependencies, which improvs their representational quality in terms of the characteristics of archaeological artifacts.

This paper is organized as follows. Section 2 presents the new datasets. Section 3 proposes the improved deep learning model. Section 4 presents the evaluation results. The paper is concluded in Section 5.

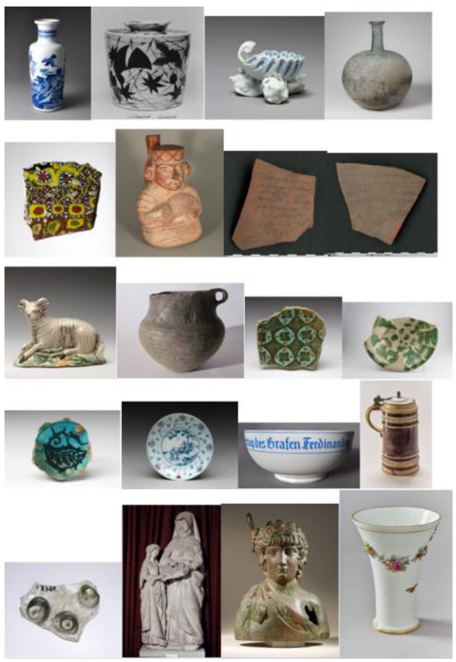

Figure 1. *Objects samples from the earthenware dataset*

New Datasets of Cultural Artifacts

The issues with the SMB dataset used in [6] centered on two main aspects: the first is that due to the lack of explicit geodata tagging in the artifacts' metadata, the taxonomy of museum collections is used to approximate the geographic origin of the artifacts belonging to each collection. The second is that each museum's collection covers a wide variety of artifacts, and there is too much variation in characteristics between different types of artifacts.

These deficiencies largely obscure the unique geographical characteristics of each class of artifacts that are critical to accurately identifying and predicting the geographic origin of different artifacts. As a result, models trained on these data are likely to tend to confuse artifacts from different geographic origins that happen to belong to the same museum collection, thereby undermining the ability of the model to make fine-grained predictions of geographic origin based on artifact characteristics alone.

To address this issue, we propose two new datasets in this work: Earthenware dataset and Coin dataset. All objects in both datasets are labeled and categorized according to the geographical and chronological origin of the object. Both datasets are type-specific, but the earthenware dataset contains a rich variety of different earthenware artifact types, while the coin dataset contains only one type of artifact, coins.

The earthenware dataset combines data crawled from *museum-digital.de*[1] and *metmuseum.org*[2], while the coin dataset gets its data only from *meseum-digital.de*. In order to create fully and accurately labeled datasets, the crawled data are first pre-processed and cleansed to eliminate incompletely labeled data, harmonize terminology used in different museums and collections.

Earthenware Dataset

The earthenware dataset represents a broad spectrum of antique artifact types made from earth, collectively referred to as earthenware. The artifact types include, for example, *plate, ceramic, vase, pot, porcelain, vessel, cup, bowl, ostracon, bottle, jug, dish, pottery, fragment, tile, glass, sculpture, jade, lacquer, gems, etc.* Figure 1 shows some randomly chosen object samples from the earthenware dataset. Data from *museum-digital.de* and *metmuseum.org* are mixed according to the geographical distribution and chronological origin of the artifacts.

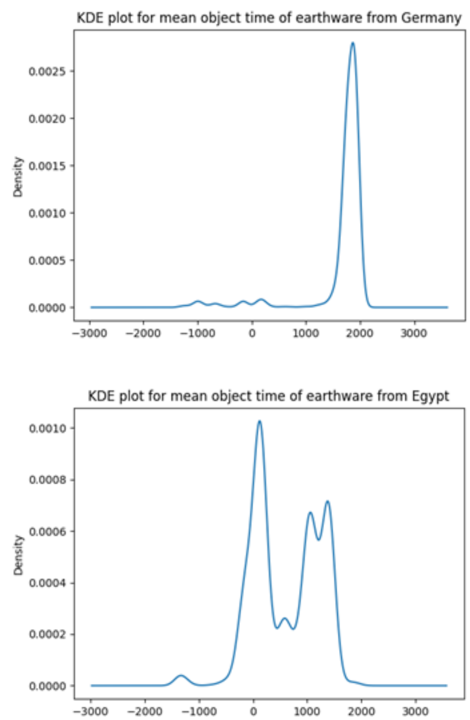

Figure 2*. KDE of mean object chronological time*

Artifacts from various countries are very unevenly distributed across the timeline. For most countries, the KDE (Kernel Density Estimation) of the mean object chronological time has unimodal or bimodal distribution. Figure 2 shows the KDE of Germany and Egypt, where Germany has a unimodal distribution with a total of 15,708 data points and Egypt has a bimodal distribution with 3,897 data points. These two countries have the largest number of data

[1] https://www.museum-digital.de/
[2] https://www.metmuseum.org/

points and therefore their data points are divided into multiple classes. For the remaining countries, the data points from each country constitute a single class.

With class balance in mind, the data points are first divided according to country, and then the data points within the same country are further divided according to chronology. We use *qcut* (quantile cut) to divide the dataset into classes, with each class containing approximately the same number of data points.

The earthenware dataset contains a total of 39,576 objects, which are divided into 18 classes as listed in Table 1. Geographically, the earthenware dataset covers 10 modern countries, and chronologically, the dataset spans the period from 1315 B.C. to 1400 A.D.

Table 1: Class Labels for Earthenware Dataset

Index	Label	Number of Objects
0	Germany1 [-1315, 1600)	1747
1	Germany2 [1600, 1735)	1876
2	Germany3 [1735, 1775)	2421
3	Germany4 [1775, 1825)	1825
4	Germany5 [1825, 1905)	2965
5	Germany6 [1905, 1914)	1681
6	Germany7 [1914, 1969)	2193
7	China [1500, 1900)	3304
8	Egypt1 [-1000, 550)	1636
9	Egypt2 [550, 2020)	2261
10	France [1000, 2000)	2528
11	Italy [-200, 500)	2562
12	UK [1700, 1920)	1527
13	Japan [1600, 1900)	2145
14	Peru [300, 400)	1695
15	Iraq [500, 1100)	1932
16	Iran [700, 1900)	3407
17	Syria [600, 1400)	1871

Coin Dataset

The coin dataset consists of coins of various types. The term coin is used to refer to a collection of different metal currency types issued throughout history, which include *coin, didrachm, as, sestertius, dime, denarius, cruiser*, *antoninianus, bullion, aureus, batzen, pegione*, and *drachma*. Figure 3 shows some randomly chosen sample coins from the coin dataset.

The classification of ancient coins should be based on an understanding of historical periods and the interconnection of geographical and historical dimensions. Historically, coins have had important political attributes in addition to their economic role. Notably, coins were used as propaganda instruments, often bearing the image of the ruler. This practice was particularly prevalent in making new emperors known or recognized by the population over a wide area. Thus, coins were widely disseminated throughout an empire or kingdom and, in addition to their economic use, contained a political message.

The dual role of coins as economic and political instruments is crucial for understanding their geographical and historical distribution. When classifying ancient coins, it is important to consider the historical context in which they circulated and not rely solely on contemporary geopolitical boundaries. Modern state territories are much more finely delineated than in previous historical periods. Applying modern geographic divisions to the identification of ancient or medieval coins will lead to inaccuracies and confusion in the classification of their geographic origins.

Figure 3. *Coin samples from the coin dataset*

For example, a coin from a specific period of the Western Roman Empire found in today's Italy may be identical to another coin from the same period found in today's Egypt. This is because even though both regions are currently independent states, they were both part of the same empire during that historical period. The geographical classification of the dataset is therefore based on the historical territories of the empire or kingdom, rather than the modern state. This approach ensures a more accurate and historically contextualized analysis.

Figure 4 illustrates the distribution of the coins based on their countries of origin. The main countries of origin include Italy, Germany, Turkey, Greece, France, Egypt, Bulgaria and Poland. The four largest origin countries are Italy, Germany, Turkey and Greece. The analysis indicates that coins from Italy are mainly centered between 400 B.C. and 400 A.D., which roughly coincides with the period of the Roman Republic and the Western Roman Empire. Coins from Germany are mainly focused on two periods: 100 to 500 A.D. and 1100 to 2000 A.D. Coins from Turkey and Greece are concentrated between 500 BC and 500 AD.

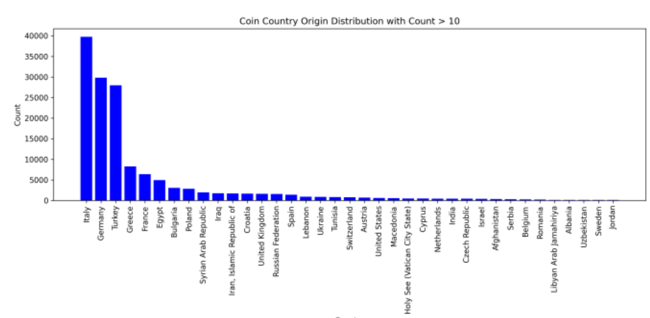

Figure 4. *Distribution of coins in the coin dataset based on origin countries*

Based on historical geography, the 126,218 objects in the coin dataset, covering the period from 500 B.C. to 2020 A.D., are divided into the following 6 empires and kingdoms: Ancient Greek (3 classes), Roman Republic and Western Roman Empire, (16 classes), Byzantine Empire (1 class), Islamic Iran and Iraq (1 class), Germanic Kingdoms (5 classes) and Poland (1 class). Ancient Greek includes Classical Period and Hellenic Period. Germanic Kingdoms include Merovingian and Carolingian Dynasties, House of Habsburg (Habsburg Monarchy), House of Habsburg (Holy Roman Empire), House of Habsburg-Lorraine (Holy Roman Empire) and various periods including Napoleonic and post-Napoleonic periods and the modern German states. Detailed class labels and the number of objects in each class are listed in Table 2.

Table 2: Class Labels for Coin Dataset

Index	Label	Number of Objects
0	Greek, Classical period [-480, -323)	6379
1	Greek, Hellenic period [-323, -281)	3496
2	Greek, Hellenic period [-281, -133)	6298
3	Roman, [-500, -146)	6145
4	Roman, [-146, -79)	5122
5	Roman, [-79, -27)	5227
6	Roman, [-27, 14)	3711
7	Roman, [14, 69)	5156
8	Roman, [69, 96)	5654
9	Roman, [96, 117)	2419
10	Roman, [117, 138)	5925
11	Roman, [138, 180)	4839
12	Roman, [180, 211)	4834
13	Roman, [211, 235)	4233
14	Roman, [235, 253)	4798
15	Roman, [253, 270)	4266
16	Roman, [270, 306)	5841
17	Roman, [306, 337)	7281
18	Roman, [337, 476)	5527
19	Byzantine Empire, [476, 1435)	2422
20	Islamic Iran and Iraq, [636, 1000)	2175
21	Germanic Kingdoms, [476, 1438)	6718
22	Germanic Kingdoms, [1438, 1640)	4774
23	Germanic Kingdoms, [1640, 1741)	3492
24	Germanic Kingdoms, [1741, 1806)	5433
25	Germanic Kingdoms, [1806, 2020)	2909
26	Poland, [1650, 1786)	2101

Proposed Scheme

Based on our previous work in [6], we propose an improved deep learning model to optimize the feature extraction and artifact recognition.

Because high-level features alone are not expressive enough to describe the characteristics of artifacts, the CNN model in [6], named MaxAvgCat, extracts multi-level features for classification of artifacts, which combine high-level, intermediate-level and low-level features. As a CNN model, MaxAvgCat is more capable of capturing textural features than structural features.

However, it is observed in our work that cultural artifacts often have both textural and structural features and contain elements of different scales, viewpoints, and colors, such as decorative patterns and motifs. Figure 5 shows eight representative objects from our dataset. They all show varying degrees of mixed textural and structural characteristics. For example, in image a, the branches and flowers of the Chinese plum blossom represent a form of structure, and the ice floe-like blue motifs represent a homogeneous texture. The white branches and blossoms do not share the characteristics of typical monotone textures, such as grass or sky. Instead, they embody a unique pattern that contributes to the overall composition of the pot in a more structured manner.

Nevertheless, CNN networks have an inductive bias of locality, and convolutional operations capture only local information due to local receptive fields, which limits their ability to capture long-distance dependencies across image regions corresponding to structural features. This leads to a lack of rich contextual information and dependencies between locations in the feature maps, which produces inconsistent features for objects belong to the same class that in turn negatively impact the classification accuracy.

Figure 5. *Objects from our datasets that embody a mixture of texture and structure*

Therefore, we propose an improved model, named SA-MaxAvgCat, in which self-attention mechanisms are introduced into the feature extraction of each level as shown in Figure 6. After extracting the feature maps from the four ResNet blocks, the self-attention mechanisms are applied on each of the four feature maps. The global information captured by the self-attention modules helps the model to establish relationships between each pixel and other pixels. The local features in the feature maps of each level are enriched with extensive contextual information and long-range dependencies, thus improving their representational quality in terms of the characteristics of artifacts.

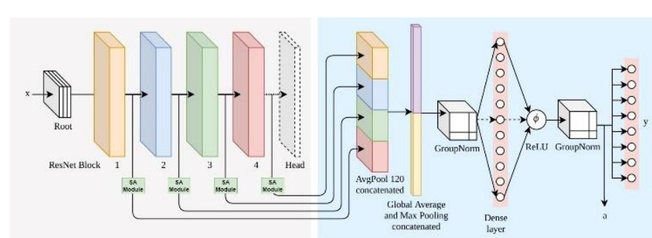

Figure 6. *SA-MaxAvgCat: MaxAvgCat model with self-attention*

Evaluation

The proposed new model is trained and evaluated on the two new datasets created in this work as well as the SMB dataset in [6]. Each dataset is split into train, validate and test sets in 8:1:1 ratio. The test results are compared with the original model in [6].

Data augmentation increases the variability and diversity of the training data and reduces overfitting. In contrast to [6] where Mixup[9] is used, in this work we propose to use CutMix[8] and random rotation (RR) instead.

Test Setup

The pretrained ResNet model BiT-M-R50x1 in [7] is used as backbone in the proposed model. The optimization used is SGD with a momentum of 0.9 and a batch size of 512. The learning rate is initially set at 0.003 and is reduced by a factor of 10 after 30%, 60%, and 90% of the whole 10000 training steps with 500 steps warmup.

In our tests, the following data augmentation techniques, CutMix[8], Mixup[9], random rotation (RR) and horizontal flip are evaluated. In the training phase, the input images are first rescaled to 512x512, then randomly cropped to 480x480. After that, random horizontal flipping with a probability of 0.5 and random rotation with a range of degrees (-90, +90) are applied.

Test Results

Table 3 lists the test results of our model and the original model in [6] on the earthenware dataset, which include the top1 and top5 accuracy, F1, precision and recall. Both models are evaluated with different combinations of data augmentation techniques, respectively, where NRR stands for No Random Rotation. For each case, two checkpoints are evaluated: the one with the lowest validation loss and the one with the highest top1 accuracy.

As seen in Table 3, the model SA-MaxAvgCat achieves the best result with the checkpoint saved at the highest top 1 accuracy. For the different data enhancement techniques, the SA-MaxAvgCat model obtains better results in all metrics when using CutMix than when using Mixup, which demonstrates CutMix regularized the model better than Mixup. The last two models in Table 3 compare the results with and without random rotation. For the checkpoint saved at the lowest validation loss, the NRR version outperforms the RR version in all metrics, while for the checkpoint saved at the highest top 1 accuracy, the RR version performs better in all metrics.

Table 3: Comparison of results on Earthenware dataset

lowest loss ckpt highest top1 ckpt	top1	top5	f1	precision	recall
MaxAvgCat w/ Mixup_NRR	0.7816 0.8060	0.9654 0.9570	0.7379 0.7609	0.7590 0.7860	0.7374 0.7558
MaxAvgCat w/ CutMix_RR	0.7903 0.7972	0.9603 0.9633	0.7448 0.7715	0.7445 0.7470	0.7472 0.7372
SA-MaxAvgCat w/ Mixup_RR	0.7623 0.7999	0.9633 0.9540	0.7182 0.7557	0.7543 0.7773	0.7113 0.7510
SA-MaxAvgCat w/ CutMix_RR	0.7858 0.8111	0.9639 0.9678	0.7391 0.7662	0.7649 0.7928	0.7328 0.7605
SA-MaxAvgCat w/ CutMix_NRR	0.7873 0.7900	0.9681 0.9666	0.7438 0.7476	0.7659 0.7692	0.7392 0.7428

Table 4 shows the evaluation results of different models on the coin dataset. Our model SA-MaxAvgCat with CutMix achieves better results than the model MaxAvgCat in all metrics, especially in top1 accuracy. As in the case of the earthenware dataset, the results of SA-MaxAvgCat are always better when using CutMix than when using Mixup. However, the results of the last two models in Table 4 are very similar, which indicates that random rotation has little impact on the coin dataset.

Table 4: Comparison of results on Coin dataset

lowest loss ckpt highest top1 ckpt	top1	top5	f1	precision	recall
MaxAvgCat w/ Mixup_NRR	0.8124 0.8164	0.9789 0.9772	0.8104 0.8149	0.8163 0.8211	0.8068 0.8112
MaxAvgCat w/ CutMix_RR	0.7983 0.7977	0.9789 0.9797	0.7997 0.7989	0.8052 0.8042	0.7966 0.7960
SA-MaxAvgCat w/ Mixup_RR	0.7965 0.7975	0.9788 0.9789	0.7958 0.7960	0.8015 0.8025	0.7929 0.7930
SA-MaxAvgCat w/ CutMix_RR	0.8268 0.8267	0.9821 0.9819	0.8264 0.8262	0.8293 0.8292	0.8245 0.8242
SA-MaxAvgCat w/ CutMix_NRR	0.8263 0.8269	0.9819 0.9819	0.8261 0.8264	0.8299 0.8300	0.8235 0.8240

Table 5 gives the test results on the SMB dataset used in [6]. Similar results can be observed as for the earthenware and coin datasets. Our model SA-MaxAvgCat with CutMix but without random rotation outperforms all other combinations except in top5 accuracy. When random rotation is applied, the CutMix version of SA-MaxAvgCat always performs better than the Mixup version. The NRR version of SA-MaxAvgCat outperforms the RR version in all metrics, implying that random rotation has a negative effect on the SMB dataset.

Table 5: Comparison of results on SMB dataset

lowest loss ckpt highest top1 ckpt	top1	top5	f1	precision	recall
MaxAvgCat w/ Mixup_NRR	0.8224 -	0.9873 -	0.7692 -	0.7789 -	0.7628 -
MaxAvgCat w/ CutMix_RR	0.8191 0.8191	0.9863 0.9863	0.7670 0.7670	0.7836 0.7836	0.7574 0.7574
SA-MaxAvgCat w/ Mixup_RR	0.8148 0.8182	0.9833 0.9835	0.7612 0.7648	0.7764 0.7775	0.7532 0.7567
SA-MaxAvgCat w/ CutMix_RR	0.8253 0.8247	0.9837 0.9837	0.7756 0.7737	0.7884 0.7860	0.7679 0.7665
SA-MaxAvgCat w/ CutMix_NRR	0.8323 0.8335	0.9865 0.9863	0.7833 0.7819	0.7923 0.7910	0.7780 0.7763

Conclusion

In this work, we first developed two new artifact datasets for training and testing of deep learning models specifically for identification of cultural artifacts, including a dataset with mixed types of earthenware and a dataset with various coins. In addition, we proposed an improved deep learning model, which used a pretrained ResNet model as backbone and was fine-tuned for artifact identification using the created datasets through transfer learning. The proposed model combines different levels of features in the new model head. Based on the fact that cultural artifacts often have both textural and structural characteristics, self-attention mechanisms are introduced into the extraction of the feature maps of each level so as to enrich local features with extensive contextual and long-range dependencies, thus improving their representational quality. In the evaluation using different datasets, the proposed model with self-attention achieved higher accuracy rate in recognition of cultural artifacts compared to the previous work.

Acknowledgements

The Federal Government Commissioner for Culture and the Media is funding the project with up to 500,000 euros from funds of the Federal Government's national AI strategy.

References

[1] Astrid Erll: Memory and Culture: A Semiotic Model. In Memory in Culture; Sara, B.Y., Translator; Palgrave Macmillan: London, UK, 2011.

[2] Neil Brodie and Isber Sabrine: The illegal excavation and trade of Syrian cultural objects: a view from the ground. In: Journal of Field Archaeology 43.1, pp.74–84, 2018.

[3] Neil Brodie: Virtually gone! The Internet market in antiquities. In Proceedings of the 6th International Conference of Experts on the Return of Cultural Property. Seoul: Overseas Korean Cultural Heritage Foundation, pp. 190–204, 2017.

[4] Relics from Syria not destroyed by ISIS sold on black market. https://www.cbsnews.com/video/relics-from-syria-not-destroyed-by-isissold-on-black-market/, September 2015. Accessed: February 3, 2025. 2023.

[5] Rachel Shabi: Looted in Syria – and sold in London: the British antiques shops dealing in artefacts smuggled by Isis. https://www.theguardian.com/world/2015/jul/03/antiquities-looted-by-isis-end-up-in-london-shops, July 2015. Accessed: February 3, 2025.

[6] Simon Bugert, Huajian Liu, Waldemar Berchtold, Martin Steinebach: Cultural assets identification using transfer learning. In Proc. IS&T Int'l. Symp. on Electronic Imaging: Imaging and Multimedia Analytics at the Edge, pp. 273-1 - 273-4, 2022.

[7] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby: Big transfer (BiT): General visual representation learning. In Proc. of ECCV, pp. 491-507, 2020.

[8] Sangdoo Yun et al.: CutMix: Regularization strategy to train strong classifiers with localizable features. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6022–6031, 2019.

[9] Hongyi Zhang et al.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (ICLR), 2018.

Author Biography

Huajian Liu received his B.S. and M.S. degrees in electronic engineering from Dalian University of Technology, China, in 1999 and 2002, respectively, and his Ph.D. degree in computer science from Technical University of Darmstadt, Germany, in 2008. He is currently a senior research scientist at Fraunhofer Institute for Secure Information Technology (SIT). His major research interests include information security, digital watermarking, robust hashing and digital forensics.

Raphael Antonius Frick is a research associate at Fraunhofer SIT / ATHENE Center in Darmstadt and a PhD candidate at Technische Universität Darmstadt. In the "Media Security \& IT Forensics" department, he researches new detection methods for AI-generated and manipulated audiovisual data, as well as techniques for identifying disinformation in social media across multiple modalities. Additionally, he investigates the possibilities of synthetic data to achieve other security goals such as anonymity and to improve the robustness and performance of deep learning classifiers.

Prof. Dr. Martin Steinebach is the manager of the Media Security and IT Forensics division at Fraunhofer SIT. From 2003 to 2007 he was the manager of the Media Security in IT division at Fraunhofer IPSI. He studied computer science at the Technical University of Darmstadt and finished his diploma thesis on copyright protection for digital audio in 1999. In 2003 he received his PhD at the Technical University of Darmstadt for this work on digital audio watermarking. In 2016 he became honorary professor at the TU Darmstadt.