Deep Learning Image Age Approximation - What is more Relevant Image Content or Age Information?

21st International Workshop on Digital-forensics and Watermarking
Guilin, China, 19-21 November, 2021

Robert Jöchl, Andreas Uhl

Department of Artificial Intelligence and Human Interfaces, University of Salzburg, Austria

November 19, 2022

Image Age Approximation

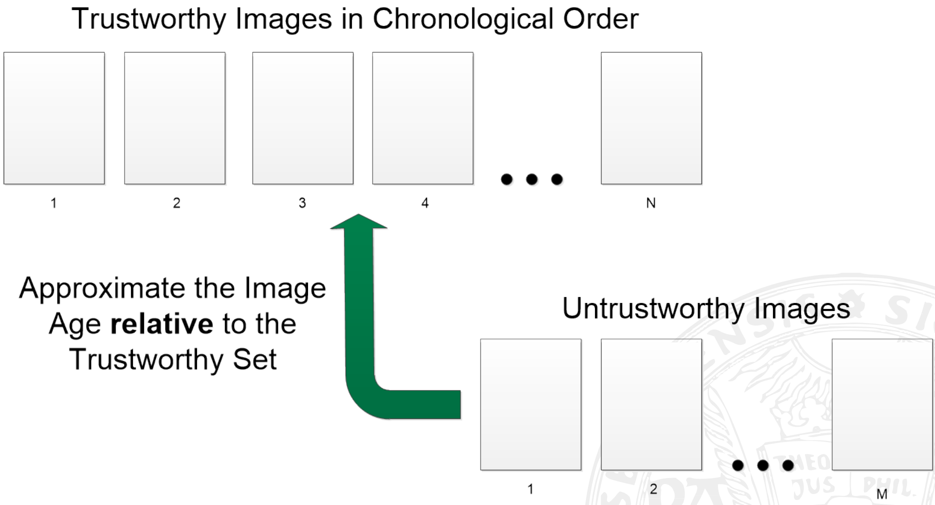

Figure: Overview image age approximation.

In-Field Sensor Defects

In-Field Sensor Defects:

- Develop after the manufacturing process and accumulate over time.
- Are due to cosmic radiation [1].
- Spread to the neighboring pixels because of preprocessing (e.g., demosaicing).

The trend towards ISO expansion and smaller pixel sizes increases the defect development rate [2].

Defect model,

$$F(I) = I + IK + \tau D + c. \qquad (1)$$

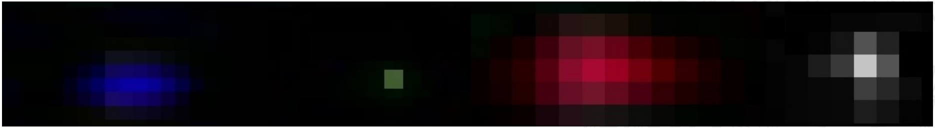

Figure: In-field sensor defects extracted from captured dark-field images.

Defect-Based Image Age Approximation

As in-field sensor defects accumulate over time, the age of an image can be approximated by the defects present.

- Fridrich et al. propose a maximum likelihood approach in [3].
- We consider image age approximation as a multi-class classification problem and utilize traditional machine learning techniques (*i.e.,* a 'Naive Bayes Classifier' and a 'Support Vector Machine') in [4].
- Ahmed et al. combined defect identification and age approximation in [5].

CNN-Based Image Age Approximation

Convolutional Neural Network (CNN) learns the classification features used.

- Ahmed et al.[6] utilize two well-known CNN architectures (*i.e.*, the AlexNet and GoogLeNet) to approximate the age of a digital image.
- The authors reported an accuracy of more than 85% for a five-class classification problem.
- The authors suggest that the features learned are not dependent on a certain image block, since the networks are trained on several non-overlapping image patches.

We[7] systematically investigated the influence of the presence of strong in-field sensor defects on training a CNN.

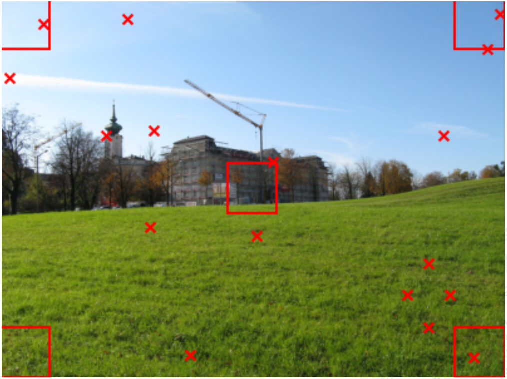

Figure: Five-crop and defect locations.

⇒ **The presence of a strong in-field sensor defect is irrelevant for improving the classification accuracy.**

Steganalysis Residual Network (SRNet)

Analogy to Image Steganalysis → detection of a weak signal.

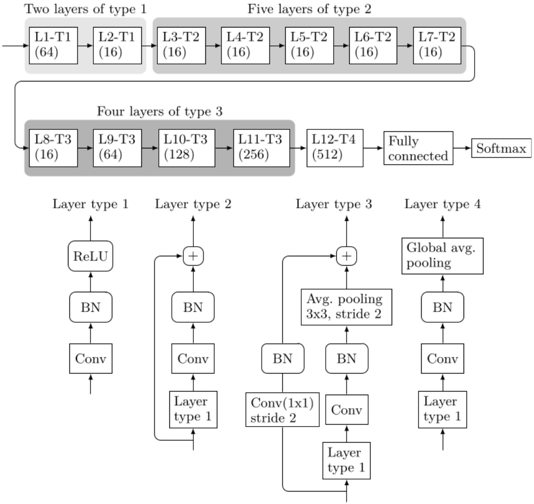

Figure: Overview SRNet[8].

Dataset - PLUS Aging Dataset

The PLUS Aging Dataset:

- Our own dataset where we have images from 4 different devices.
- A binary classification problem is considered with a time difference between the classes ranges from 7 to 13 years.

Dataset - PLUS Aging Dataset

The PLUS Aging Dataset:

Identifier	Make/Model	Res. [W×H]	Sensor
PLUS-nikon01	Nikon E7600	3072 × 2304	CCD
PLUS-canon01	Canon PowerShotA720IS	2592 × 1944	CCD
PLUS-pentax01	Pentax K5	4950 × 3284	CMOS
PLUS-pentax02	Pentax K5II	4950 × 3284	CMOS

Dataset - PLUS Aging Dataset

The PLUS Aging Dataset:

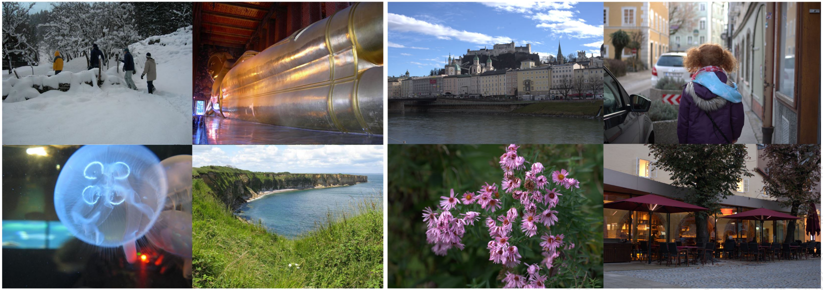

Figure: Random samples of the PLUS Aging Dataset.

Dataset - NTIF

The Northumbria Temporal Image Forensics (NTIF)[9] Database:

- Is a publicly available dataset.
- For each device, approximately 71 timeslots ranging over 94 weeks (between 2014 and 2016) are available.
- A binary classification problem is considered, where timeslot 1-5 is regarded as first class and timeslot 21-25 as second class (similar as in[6]).

Dataset - NTIF

The Northumbria Temporal Image Forensics (NTIF)[9] Database:

Identifier	Make/Model	[W$\times$H]	Sensor
NTIF-canon01	Canon IXUS115HS	4000 $\times$ 3000	CMOS
NTIF-canon02	Canon IXUS115HS	4000 $\times$ 3000	CMOS
NTIF-fujifilm01	Fujifilm S2950	4288 $\times$ 3216	CCD
NTIF-fujifilm02	Fujifilm S2950	4288 $\times$ 3216	CCD
NTIF-nikon01	Nikon Coolpix L330	5152 $\times$ 3864	CCD
NTIF-nikon02	Nikon Coolpix L330	5152 $\times$ 3864	CCD
NTIF-panasonic01	Panasonic DMC TZ20	4320 $\times$ 3240	CMOS
NTIF-panasonic02	Panasonic DMC TZ20	4320 $\times$ 3240	CMOS
NTIF-samsung01	Samsung pl120	4320 $\times$ 3240	CCD
NTIF-samsung01	Samsung pl120	4320 $\times$ 3240	CCD

Image Age Approximation Results

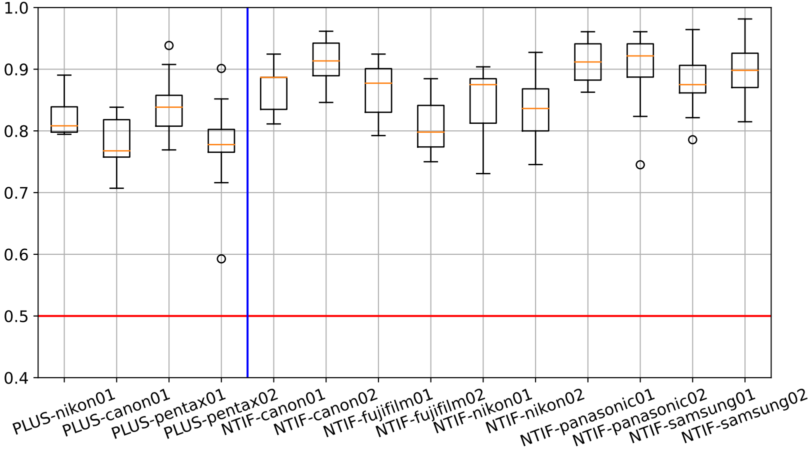

Figure: Boxplot of the resulting age approximation accuracy for all 10 runs.

Explainable Artificial Intelligence (XAI)

The field of XAI is focused on the understanding and interpretation of the decision of deep neural networks.

Class Activation Map (CAM) [10]

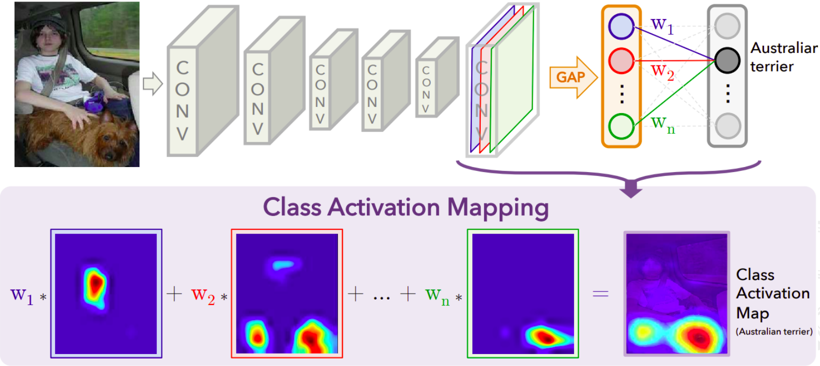

Figure: Overview Class Activation Map [10].

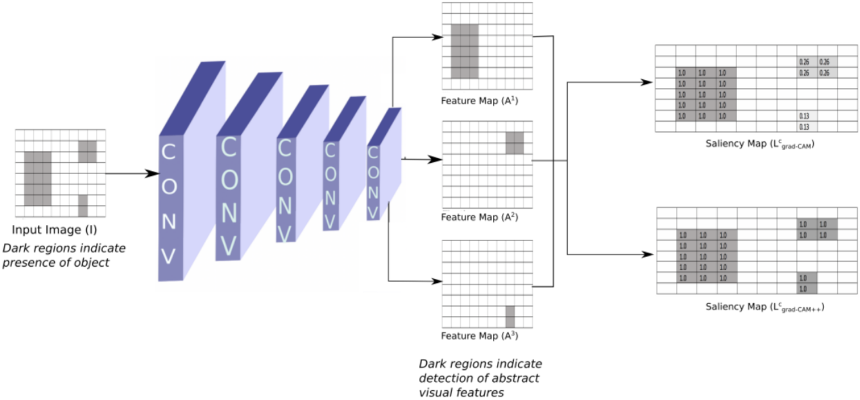

Figure: A hypothetical example elucidating the intuition behind GradCAM++ [11].

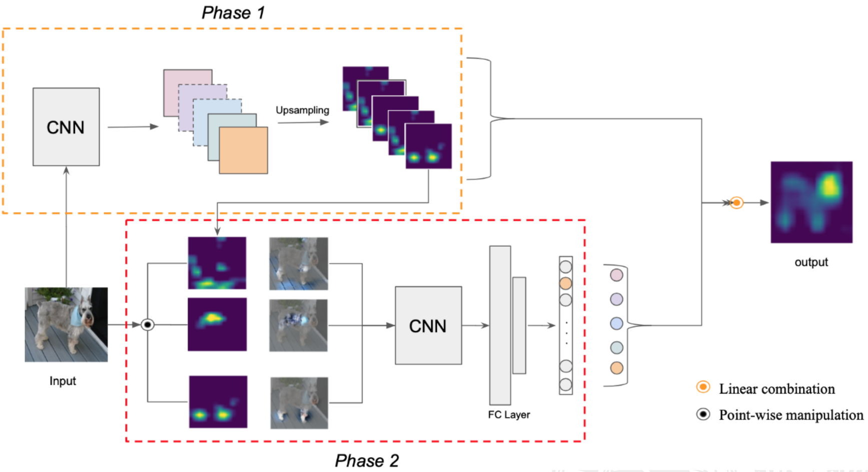

Figure: Pipeline of the ScoreCAM [12].

CAM Analysis

Let θ be the sum of age traces at a certain point in time that are embedded in an image I. We assume that θ is constant across all images of a given age class y and differs between the other age classes. Based on this assumption, we expect that regions highlighted by the obtained saliency maps:

1. are independent of the image content (*e.g.*, captured objects and scene properties),
2. are constant across the different runs (*i.e.*, since all images per class share the same θ, the overall activation should be similar across all different test sets).

CAM Analysis - Activation on Objects

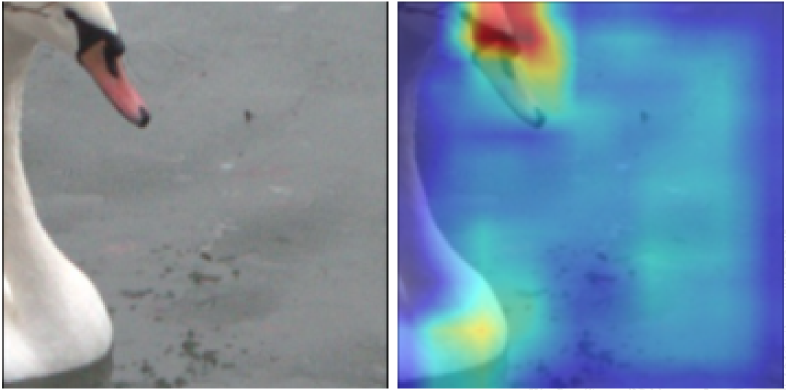

Figure: Example of activation directly on objects.

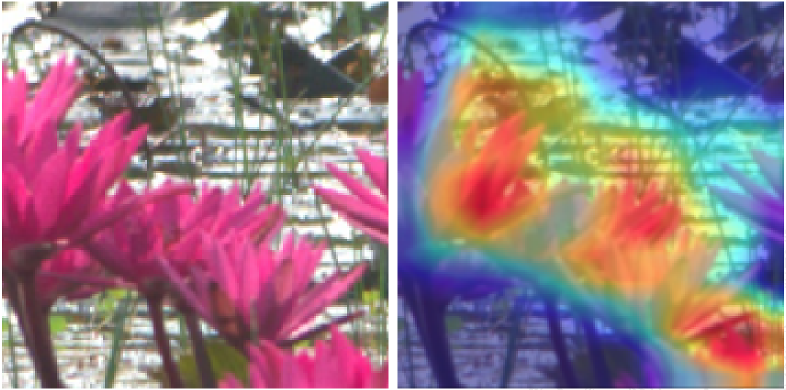

Figure: Example of activation directly on objects.

CAM Analysis - Activation on Objects

Figure: Example of activation directly on objects.

CAM Analysis - Activation on Objects

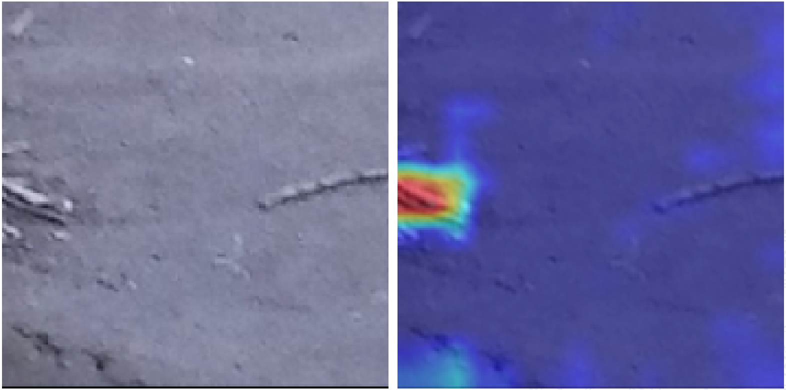

Figure: Example of activation directly on objects.

CAM Analysis - Activation on Objects

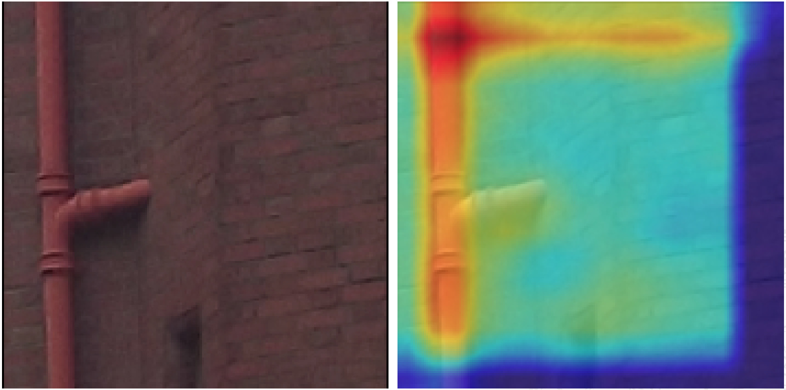

Figure: Example of activation directly on objects.

CAM Analysis - Activation on Shrub-, Tree-like Structures

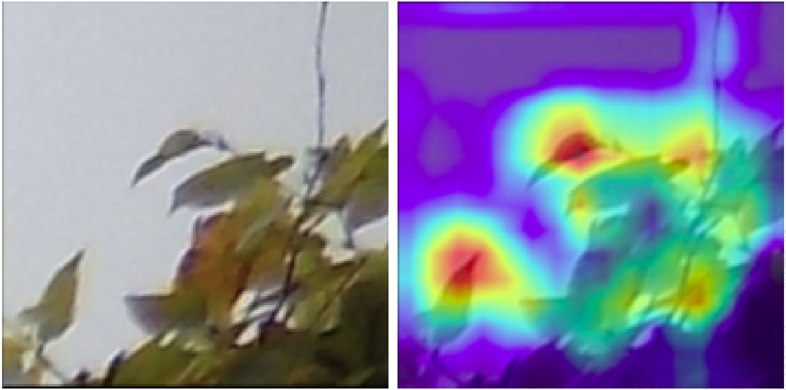

Figure: Example of activation directly on shrub-, tree-like strutures.

CAM Analysis - Activation on Shrub-, Tree-like Structures

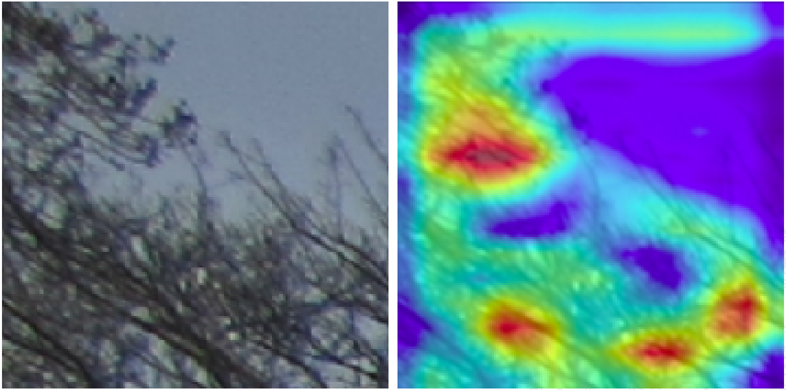

Figure: Example of activation directly on shrub-, tree-like strutures.

CAM Analysis - Activation on Shrub-, Tree-like Structures

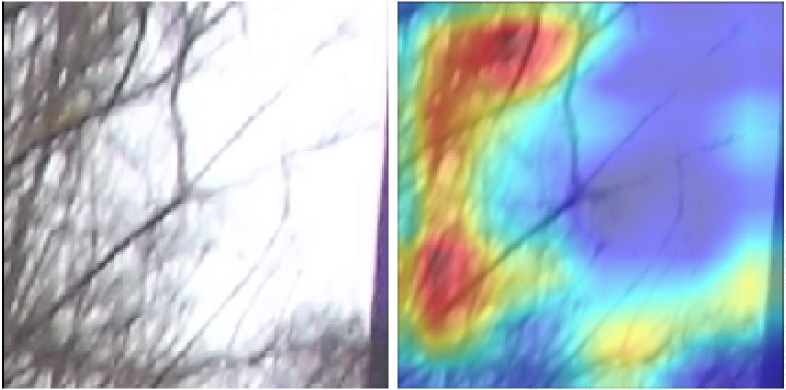

Figure: Example of activation directly on shrub-, tree-like strutures.

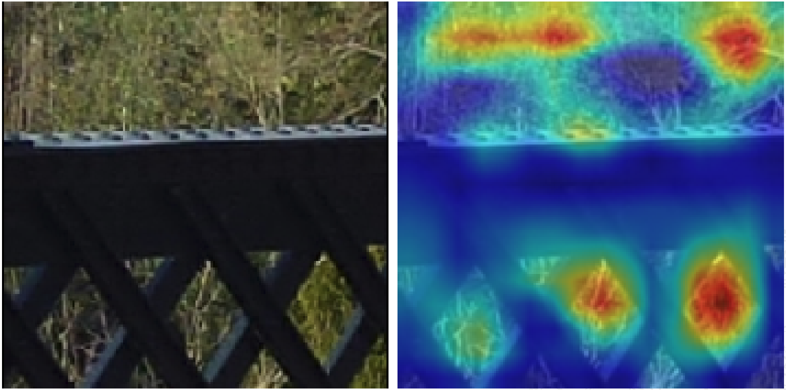

Figure: Example of activation directly on shrub-, tree-like strutures.

CAM Analysis - Activation on Shrub-, Tree-like Structures

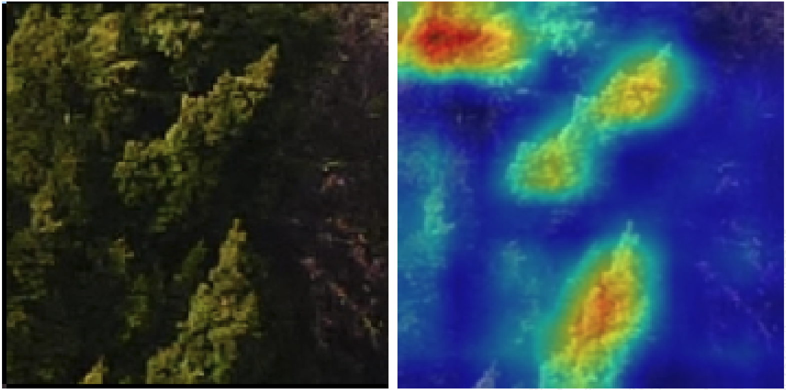

Figure: Example of activation directly on shrub-, tree-like strutures.

CAM Analysis - Activation on Areas

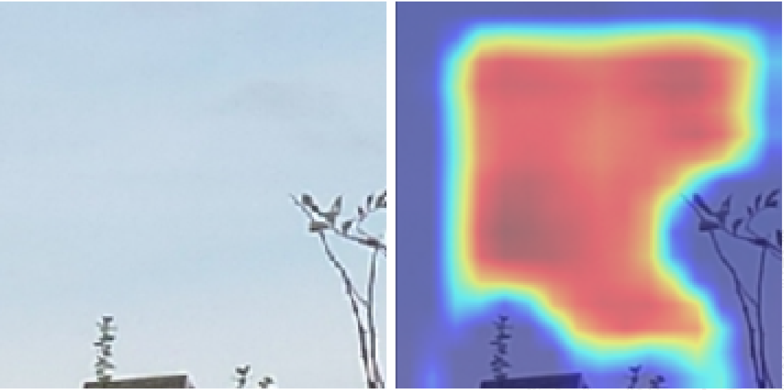

Figure: Example of activation directly on image areas.

CAM Analysis - Activation on Areas

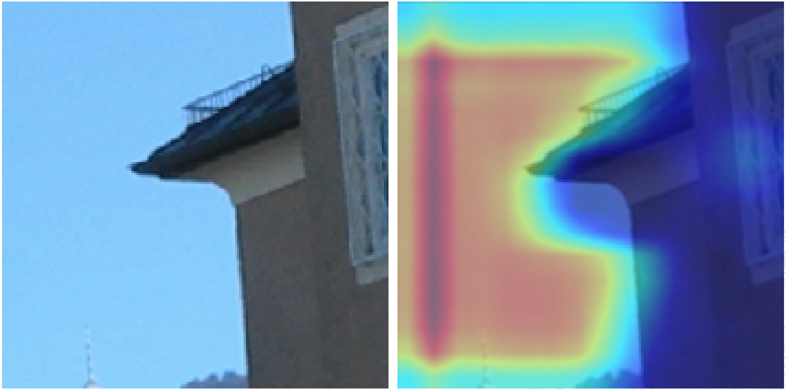

Figure: Example of activation directly on image areas.

CAM Analysis - Activation on Areas

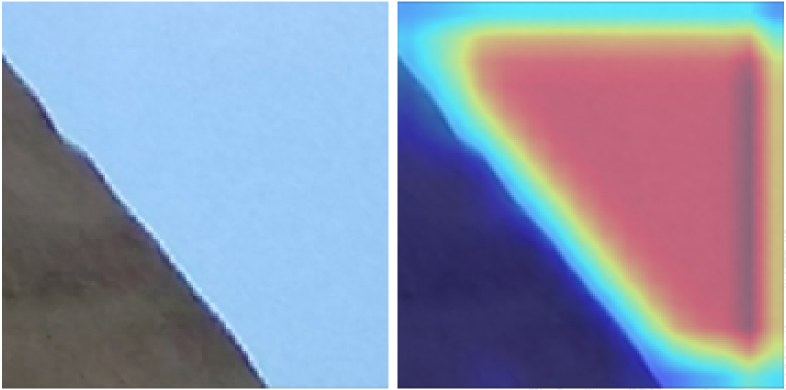

Figure: Example of activation directly on image areas.

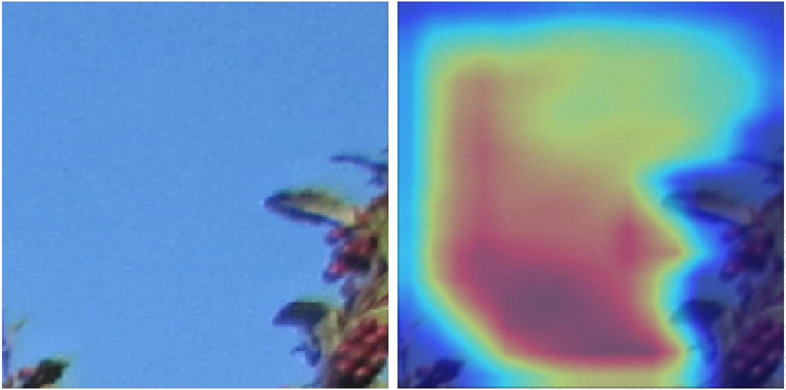

Figure: Example of activation directly on image areas.

CAM Analysis - Activation on Areas

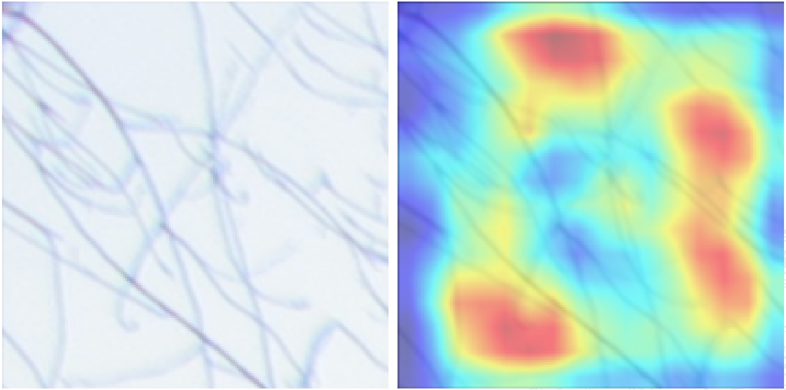

Figure: Example of activation directly on image areas.

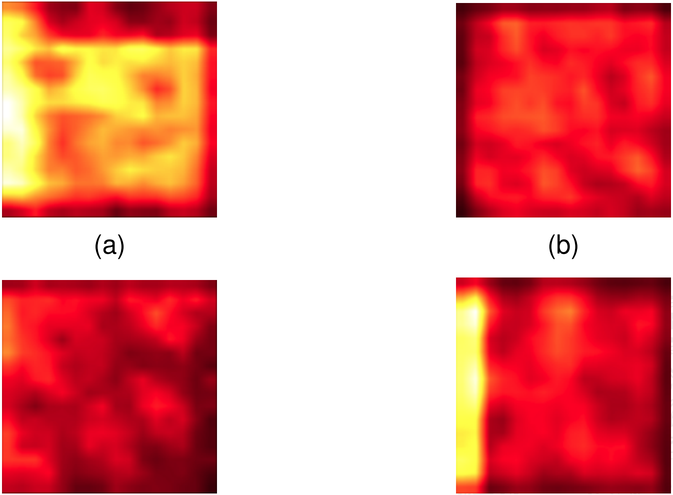

(a) (b) (c) (d)

Figure: Examples of superimposed activation of correctly predicted image patches of a given run.

CAM Analysis

Based on these observations, it is unlikely that a standard CNN trained on regular scene images would exploit solely age-related features to determine the age class.

Potential Solutions

Apply Constraints:

- Focus the network on in-field sensor defect locations (e.g., training the network on small image patches (*i.e.,* 32×32) extra around each defect).
- Apply preprocessing to suppress the image content (e.g., feed median filter residuals into the network).
- Utilize special network architectures (e.g., content suppression layer).

Potential Solutions

- Apply constraints on the acquisition of training data.
- Potential scene or environmental dependencies can be eliminated by capturing different fixed backgrounds and foreground objects in a controlled environment.

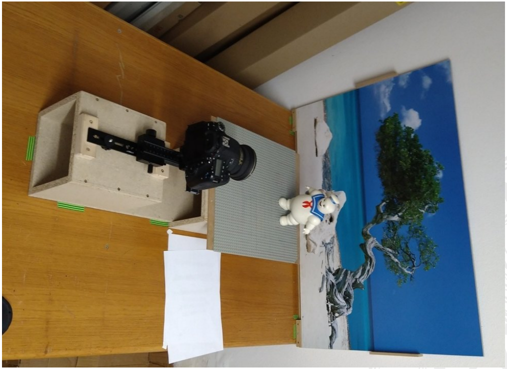

Figure: Fixture for recording standardized scenes.

Conclusion

- Based on the CAM analysis conducted, we conclude that it is unlikely that a standard CNN trained on regular scene images would exploit solely age-related features to determine the age class.
- In the field of image forensics, it is important that the decision is based on comprehensible evidence.

$\Rightarrow$ *When using a CNN for image age approximation, it is important to design the setup carefully!!*

References

[1] A. J. Theuwissen, "Influence of terrestrial cosmic rays on the reliability of ccd image sensors part 1: Experiments at room temperature," *IEEE Transactions on Electron Devices*, vol. 54, no. 12, pp. 3260–3266, 2007.

[2] G. H. Chapman, R. Thomas, R. Thomas, K. J. Coelho, S. Meneses, T. Q. Yang, I. Koren, and Z. Koren, "Increases in hot pixel development rates for small digital pixel sizes," *Electronic Imaging*, vol. 2016, no. 12, pp. 1–6, 2016.

[3] J. Fridrich and M. Goljan, "Determining approximate age of digital images using sensor defects," in *Media Watermarking, Security, and Forensics III* (N. D. Memon, J. Dittmann, A. M. Alattar, and E. J. D. III, eds.), vol. 7880, pp. 49 – 59, International Society for Optics and Photonics, SPIE, 2011.

References (cont.)

[4] R. Joechl and A. Uhl, "A machine learning approach to approximate the age of a digital image," in *Digital Forensics and Watermarking: 19th International Workshop, IWDW 2020, Melbourne, VIC, Australia, November 25–27, 2020, Revised Selected Papers*, vol. 12617 of *Springer LNCS*, pp. 181–195, Springer International Publishing, 2021.

[5] F. N. Ahmed, F. Khelifi, A. Lawgaly, and A. Bouridane, "A machine learning-based approach for picture acquisition timeslot prediction using defective pixels," *Forensic Science International: Digital Investigation*, vol. 39, p. 301311, 2021.

[6] F. Ahmed, F. Khelifi, A. Lawgaly, and A. Bouridane, "Temporal image forensic analysis for picture dating with deep learning," in *2020 International Conference on Computing, Electronics Communications Engineering (iCCECE)*, pp. 109–114, 2020.

[7] R. Joechl and A. Uhl, “Apart from in-field sensor defects, are there additional age traces hidden in a digital image?,” in *2021 IEEE International Workshop on Information Forensics and Security (WIFS)*, (Montpellier, France), pp. 1–6, 2021.

[8] M. Boroumand, M. Chen, and J. Fridrich, “Deep residual network for steganalysis of digital images,” *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 5, pp. 1181–1193, 2018.

[9] F. Ahmed, F. Khelifi, A. Lawgaly, and A. Bouridane, “The ‘northumbria temporal image forensics’ database: Description and analysis,” in *2020 7th International Conference on Control, Decision and Information Technologies (CoDIT)*, vol. 1, pp. 982–987, 2020.

[10] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning deep features for discriminative localization,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2921–2929, 2016.

[11] A. Chattopadhay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, “Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks,” in *2018 IEEE winter conference on applications of computer vision (WACV)*, pp. 839–847, IEEE, 2018.

[12] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu, “Score-cam: Score-weighted visual explanations for convolutional neural networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 24–25, 2020.

Thank you for your attention!