

Morphological Classification of Tobacco Ash Using a Deep Convolutional Neural Network: From Sherlock Holmes's "monograph" to modern computer vision

Irine Imerlishvili

Master of Security Studies, Business and Technology University, Tbilisi, Georgia.

Email: irine.imerlishvili.1@btu.edu.ge

Giorgi Kuchava

PhD (Engineering of Informatics), Associate professor, Georgian Technical University, Tbilisi, Georgia.

Email: kuchavagiorgi08@gtu.ge

DOI: <https://doi.org/10.68218/ed76v933>

Abstract

Arthur Conan Doyle's literary character Sherlock Holmes claimed the ability to visually distinguish the ashes of up to 140 varieties of tobacco. This paper examines how far that idea is realisable with modern artificial intelligence. We propose a methodological framework in which a convolutional neural network (CNN) is used to classify microscopic images of tobacco ash into five categories. The paper describes the data-collection protocol, model architecture, training procedure, and illustrative results. We show that the mineral and morphological features of ash: colour, texture, fibrous structure, carry sufficient discriminative information for machine classification, although forensic application requires combination with spectroscopic methods.

Keywords: convolutional neural network, computer vision, forensic science, tobacco ash, classification, deep learning, trace evidence

1.Introduction

In the short story "The Five Orange Pips" (1891), Sherlock Holmes mentions his own monograph, Upon the Distinction Between the Ashes of the Various Tobaccos, in which, by his account, the ash of 140 varieties of cigar, cigarette and pipe tobacco was described[1]. In the nineteenth


century this was a literary hyperbole the unaided human eye cannot deliver systematic, repeatable identification of such fine distinctions.

Modern computer-vision systems, by contrast, are built precisely for this class of problem: recognising textural and chromatic patterns in high-resolution imagery that lie beyond human perception. Convolutional neural networks (CNNs) are already applied successfully in materials science, soil classification, microscopic histology, and trace-evidence examination. The question follows naturally: is a digital realisation of Holmes's "monograph" feasible?

The aims of this paper are: (a) to systematise the discriminative features of tobacco ash; (b) to develop a CNN-based classification framework; and (c) to critically analyse the limitations of the approach in the context of forensic reliability.

The residual product of tobacco combustion ash is an inorganic mineral residue whose composition depends on the tobacco variety, the soil, fertilisers, curing technology, and the composition of the wrapping paper. The principal visually perceptible parameters are the following (**fig.1**): **Colour**, A high content of calcium and magnesium gives ash a white cast; incomplete combustion and carbonaceous residues produce dark grey or black. **Texture**, The degree to which fibrous structure is preserved distinguishes leaf cigar tobacco from a shredded cigarette blend. **Cohesion**, Cigar ash often retains a cylindrical form, whereas cigarette ash crumbles rapidly. **Granularity**, At the microscopic level, particle-size distribution differs measurably between classes.


Fig. 1. The four principal discriminative features of ash, which together form a visual "signature" sufficient for classification



2.Methodology


The overall research workflow comprises five stages, from the physical sample to the quantitative evaluation of the model (Fig. 2). The framework defines five classes: (1) commercial cigarette, Virginia blend; (2) commercial cigarette, American blend; (3) hand-rolled cigar (long-filler); (4) pipe tobacco; and (5) roll-your-own tobacco (RYO). For each class, ash samples are produced under a standardised combustion protocol (controlled humidity $60\pm5\%$, temperature $22\pm2\text{ }^{\circ}\text{C}$) and photographed at two levels: macro photography (smartphone camera, 12 MP) and digital microscopy ($\times 40\text{--}\times 200$ magnification).

Fig. 2. Architecture of the proposed model: a ResNet-50 convolutional backbone via transfer learning, global average pooling (GAP), a fully connected layer, and a five-class softmax output.



The classifier is a convolutional neural network built with a transfer-learning approach(**Fig.3**): a ResNet-50 architecture pre-trained on ImageNet, with the final layers retrained for the five-class task[2]. Input images are 224×224 pixels, RGB. Data augmentation random rotation, horizontal flipping, brightness variation compensates for variability in photographic conditions. The data are split 70/15/15 (training/validation/test); the optimiser is Adam ($lr = 1 \times 10^{-4}$) and the loss function is categorical cross-entropy.

Fig. 3. Architecture of the proposed model: a ResNet-50 convolutional backbone via transfer learning, global average pooling (GAP), a fully connected layer, and a five-class softmax output.



The figures presented below are illustrative. They reflect accuracy levels typically attainable in comparable texture-classification tasks and serve to demonstrate the plausibility of the methodological framework rather than to report a completed experiment.[4,5]

Table 1. Illustrative classification metrics on the test set.

Class	Precision	Recall	F1-score	Test samples
Cigarette - Virginia blend	0.91	0.89	0.90	150
Cigarette - American blend	0.87	0.88	0.87	150
Cigar (long- filler)	0.96	0.97	0.96	150

Pipe tobacco	0.93	0.92	0.92	150
Roll-your-own (RYO)	0.85	0.86	0.85	150
Macro average	0.90	0.90	0.90	750

Diagram(Fig. 4) shows the training dynamics of the model alongside the F1-scores achieved for each class. The validation curve tracks the training curve without a widening gap, which indicates that no substantial overfitting occurred; both plateau at around the thirtieth epoch. The per-class bars show that separability is highest for cigar ash and lowest for roll-your-own tobacco.[6,7]

The distribution of errors across classes is summarised in the confusion matrix (Fig. 5).

3. Discussion

Inter-class similarity. Distinguishing the two commercial cigarette blends is the hardest sub-task (Fig. 5), their mineral and textural profiles partially overlap. Here, visual analysis must be supplemented by instrumental methods: atomic absorption spectroscopy (AAS), X-ray fluorescence (XRF), or scanning electron microscopy with elemental analysis (SEM-EDS).

Environmental factors. Under real, non-standardised conditions, combustion temperature, humidity, and mechanical disturbance of the ash substantially alter its visual features, reducing the model's ability to generalise. Including "in the wild" samples in the training set is therefore essential.

Forensic reliability. The black-box nature of deep-learning models is problematic in evidence law. Interpretability methods (Grad-CAM, SHAP) are a necessary complement: they show visually which regions of an image drive the model's decision, allowing an expert to justify the conclusion.

The Holmes paradox. It is ironic that Holmes's own approach, observe the signs, then reason to a justified conclusion is closer to interpretable AI than to a black-box neural network.

Contemporary forensic AI is moving precisely towards a synthesis of the two paradigms: the sensitivity of a neural network plus an explanation the examiner can understand.

Fig. 4. Left training and validation accuracy over 40 epochs; right F1-score for each class . The highest separability is recorded for the cigar class, whose fibrous structure differs most sharply from the others.





Fig. 5. Normalised confusion matrix (illustrative). The principal errors occur between the two cigarette blends (7–8%) and within the RYO class as expected, since roll-your-own tobacco is chemically close to commercial cigarette tobacco.



4. Conclusion

Sherlock Holmes's literary "monograph" on tobacco ash, read as hyperbole in 1891, is in principle a realisable task with modern deep-learning technology. A convolutional neural network trained by transfer learning can, under controlled conditions, reliably distinguish the main categories of tobacco on the basis of visual information alone. At the same time, to meet forensic standards, visual classification should be treated not as independent evidence but as a rapid screening instrument, confirmed by instrumental chemical analysis. Future research directions include a multimodal model (image + spectrum), a database of real casework samples, and systematic validation of interpretability methods.

REFERENCES

- [1] A. C. Doyle, "The Five Orange Pips," in *The Adventures of Sherlock Holmes*. Scotts Valley, CA: CreateSpace Independent Publishing Platform, 1891, 26 pp.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [3] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017.
- [4] M. Brady, "Artificial intelligence approaches to image understanding," *AI Memo*, Artificial Intelligence Laboratory, MIT, Cambridge, MA, pp. 205–264.
- [5] M. Albiez and M. Fischer, *AI for Advanced Image Analysis*. Oberkochen, Germany: ZEISS, 2023, 56 pp.
- [6] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016, 801 pp.
- [7] F. Chollet, *Deep Learning with Python*. Shelter Island, NY: Manning Publications, 2017, 350 pp.