Siamese neural networks for detecting banknote printing defects

K. Boria, A. Luciani, *S. Marchetti* and M. Viticoli
(Bank of Italy)

3rd IFC Workshop on Data Science in Central Banking
Oct. 18, 2023

Introduction

- Banknote production in the Eurosystem relies on strict quality controls throughout the process (ECB quality requirements, ISO standards);
- Similar to any manufacturing process, the printing of banknotes can give rise to various imperfections and defects;
- The number, type, and size of defects on banknotes are critical for the conformity of production batches (with related losses).

Introduction

- Banknotes exhibit variability in the position of elements, and in the shade and intensity of colors;
- While measurement of many parameters is carried out via automatic optical systems, the validation process cannot be fully automatized and requires highly trained staff (potential subjectivity factors);
- Artificial Intelligence (AI) systems are profitably used across industries to support and automatize quality control;
- Our work focuses on a denomination/defect pair to assess whether neural networks could improve process efficiency.

Data

- Since acquisition of high-resolution images is costly:
  - Focus on €50 banknotes of the *Europa* series;
  - Focus on flag-*bite* defects: lack of ink on a homogenous background, with varying shape, size, and position.
- We manually acquire 24-bit color images (RGB model) of banknotes with 2,656 x 1,467 pixels resolution and annotate them as *fit*/*unfit*.

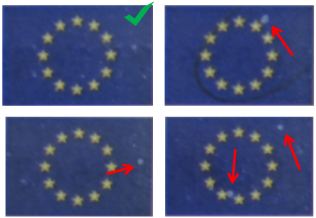

Neural Networks & Image Processing

- Previous research on banknotes analysis mostly resorts to convolutional neural networks (CNNs);
- Applications address recognition of denomination, serial number, currency, state of wear as well as counterfeit detection. On the production phase, we mention Pham et al., 2017; Ke et al., 2016.
- The inherent variability of banknotes requires training of complex CNN architectures, usually addressed via data augmentation.

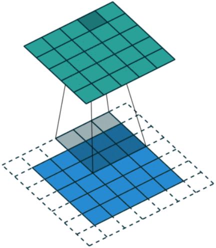

One-shot learning & Siamese Networks

- The traditional way: data augmentation
  - Manual: Overfitting of specific bite defects;
  - AI-enabled: Hallucination-prone behavior;
- The practical way: *one-shot learning* (Fei et al., 2006)
  - Emulates the ability of the human mind to associate entities based on similarity criteria.

VERIFICATION

SUPPORT SET

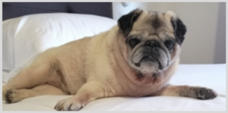

Model development & Training

- We build a Siamese neural network architecture (Koch et al., 2015) for the bite-detection task:
  - Each twin branch extracts a signal from its input, either a *verification* or a *support* image;
  - The more similar the signals, the most likely images belong to the same class.

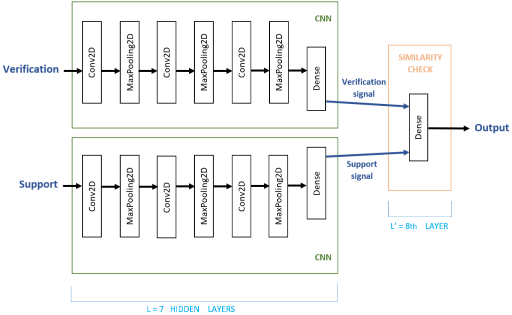

Model development & Training

We train our models on different resolutions of the input set: high, medium and low resolution + cropped flag detail:

- Training set: $N = 200$ images (100 bites, 100 fit; $C(N, 2)$ training pairs);
- We ensure convergence of the training process to global minima via drop-out, batch normalization, random re-initialization of weights;
- We control internal logics of networks via soft masks extracted from convolutional layers.

Empirical Results

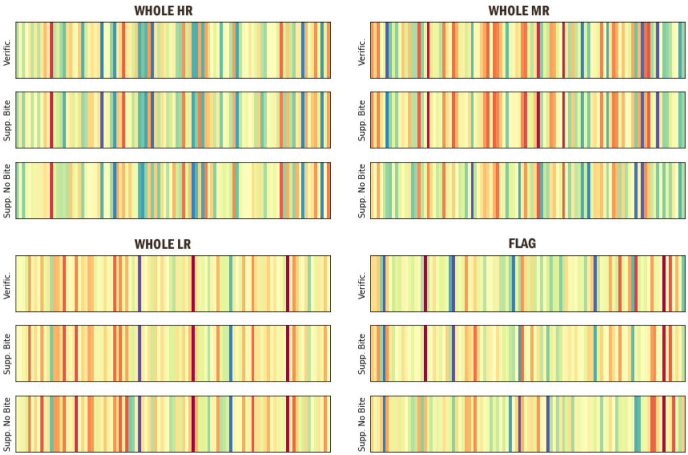

Empirical Results

	Recall				F1-Score			
	RND-1	SSIM-0	SSIM-1	SSIM-k	SSIM-k	SSIM-0	SSIM-1	SSIM-k
WHOLE HR	-	.75	.62	-	-	.67	.55	-
WHOLE MR	1.	.87	.87	.5	.84	.82	.93	.67
WHOLE LR	-	.62	.62	-	-	.55	.62	-
FLAG	-	-	.5	.98	-	-	-	-

Test set size: 46 images (14 bites; 32 fit);
RND-1: Support image is picked at random;
SSIM-j, j=0,1: the (j+1)th most similar image is picked as support;
SSIM-k: the k most similar images are picked as support, with k=10.

Empirical Results

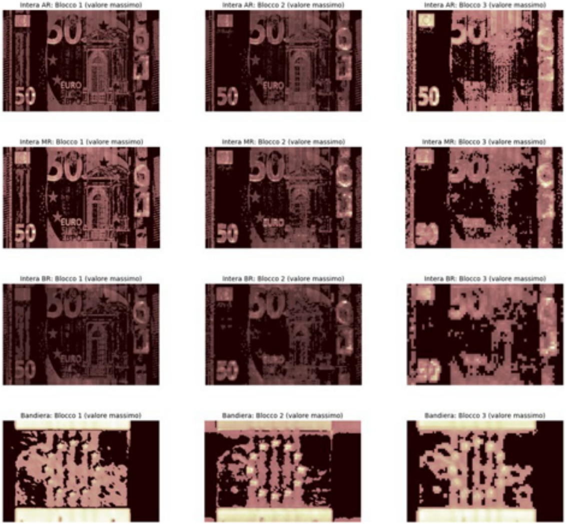

Conclusions & Next Steps

- Our exploratory study provides insights on the detection of banknote defects for quality control via one-shot learning;
- We find medium resolution input allows for greater accuracy based on different metrics and evaluation criteria;
  - Trade-off between resolution and attention scattering due to variability patterns of banknotes;
- We are able to gain insights on the internal logics adopted by the models via soft masks;
- Future work will extend the analysis to additional defects (minor bite type defects already tested) and denominations, and to additional model architectures.

Thank you!