


File: X51005568866.png SHA-256: af9046c2d8077aec... Size: 936x2753 px Analysed: 2026-05-12 00:28:38 UTC Model: FraudX v2-multi (ResNet-18, ep13, thr=0.08)



VERDICT — High risk

Likely forged

Multiple regions show strong evidence of editing.

99%

CNN forgery confidence

Patch CNN: High risk
forgery confidence 99%

VAT QR: No QR
no ZATCA QR detected

Key findings

- Multiple regions show strong evidence of editing.
- Showing the top 5 highest-confidence regions (highest score 100%). 40 of 426 text-bearing regions exceeded the model's threshold overall.
- Edit type predicted across the top 5 regions: 2 x copy-paste edits, 2 x deletions / erasures, and 1 x pixel-level edits.
- Targeted fields: 3 x Total / payment lines and 2 x Product lines.
- This file is an exact duplicate of one already analysed in this session.

VAT QR check

No ZATCA-format VAT QR detected on this receipt.

Document with suspicion heatmap (red = CNN-flagged)



Top suspicious regions

Each thumbnail is the 128x128 patch fed to the CNN, with its score and predicted edit type / field.

1

100%

CPI

Total/payment

2

100%

PIX

Product

3

100%

CUT

Total/payment

4

100%

CPI

Product

5

97%

CUT

Total/payment

Approach

ResNet-18 patch classifier (128x128, stride 64) trained on FINDIT2 (Tornes et al., ICDAR 2023). Two auxiliary heads classify the modification technique and the document field affected; the binary backbone is frozen so the headline patch precision (92.25%) is preserved while adding explainability. The image-level fraud score is the top-k mean of patch probabilities over text-bearing regions only (edge density ≥ 0.02). The Saudi VAT QR check validates the ZATCA Phase-1 TLV format and cross-checks the QR's declared totals against pytesseract OCR of the printed receipt body.

Reported numbers vs paper

METRIC	OURS	PAPER
Patch precision	92.25% (this work)	79.41% (paper OH-JPEG)
Patch F1 / AUC	91.79 / 0.97	—
Image-level F1 (top-k)	29.66	28.39 (paper ChatGPT-relaxed)
Mod head test acc	CPI 73 · CUT 100 · IMI 38 · PIX 55 · Other 14	
Entity head test acc	Total/payment 67 · Metadata 47 · Product 27 · Company 24	

Raw model output

- Patches scored: 602 Text-rich: 426 Flagged ≥ 0.08 : 40
- Top-k mean fraud score: 0.9928 CNN risk: HIGH Duplicate: 1.00 (exact_duplicate)
- Patch CNN flagged 40 of 426 text-rich patches (top score 1.00).
- Top 5 suspicious regions clustered around image coordinates (704,1792), (512,1984), (768,1792), (448,1984), (0,1984).
- Predicted modification mix across top regions: 2x CPI, 2x CUT, 1x PIX.
- Predicted entity types: 3x Total/payment, 2x Product.
- Document is an EXACT duplicate of a file already analysed in this session (resubmission).