


File: X00016469676.png SHA-256: ebf4a481c57b461e... Size: 463x797 px Analysed: 2026-05-12 01:13:13 UTC Model: FraudX v2-multi (ResNet-18, ep13, thr=0.08)



VERDICT — Low risk

Looks clean

No regions exceeded the forgery threshold. No VAT QR present to cross-check.

4%

CNN forgery confidence

Patch CNN: Low risk
forgery confidence 4%

VAT QR: No QR
no ZATCA QR detected


Key findings

- No regions exceeded the model's forgery threshold.
- Showing the top 5 highest-confidence regions (highest score 70%). 7 of 73 text-bearing regions exceeded the model's threshold overall.
- Edit type predicted across the top 5 regions: 5 × deletions / erasures.
- Targeted fields: 4 × Product lines and 1 × Company / header info.
- This file is an exact duplicate of one already analysed in this session.

VAT QR check

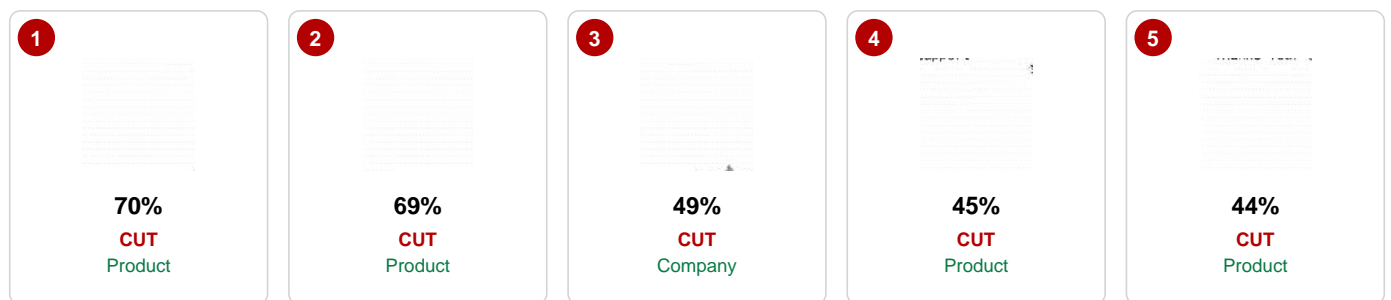
No ZATCA-format VAT QR detected on this receipt.

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.



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: 84 Text-rich: 73 Flagged ≥ 0.08 : 7
- Top-k mean fraud score: 0.0429 CNN risk: LOW Duplicate: 1.00 (exact_duplicate)
- Patch CNN flagged 7 of 73 text-rich patches (top score 0.70).
- Top 5 suspicious regions clustered around image coordinates (64,669), (0,669), (128,669), (128,640), (0,640).
- Predicted modification mix across top regions: 5x CUT.
- Predicted entity types: 4x Product, 1x Company.
- Document is an EXACT duplicate of a file already analysed in this session (resubmission).