


File: X00016469669.png SHA-256: e7922674d7241ce7... Size: 463x605 px Analysed: 2026-05-12 01:12:38 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.

30%

CNN forgery confidence



Key findings

- No regions exceeded the model's forgery threshold.
- Showing the top 5 highest-confidence regions (highest score 86%). 11 of 51 text-bearing regions exceeded the model's threshold overall.
- Edit type predicted across the top 5 regions: 3 × deletions / erasures and 2 × copy-paste edits.
- Targeted fields: 3 × Company / header info and 2 × Total / payment lines.
- This file is an exact duplicate of one already analysed in this session.

VAT QR check

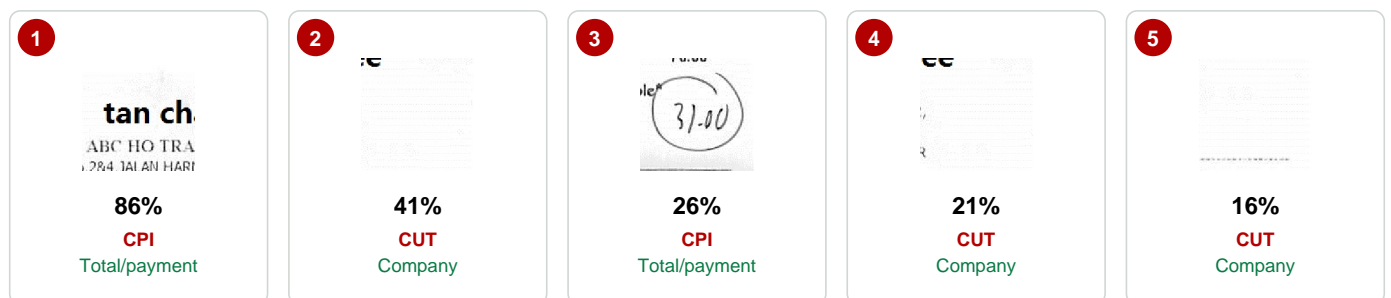
No ZATCA-format VAT QR detected on this receipt.

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.



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: 63 Text-rich: 51 Flagged ≥ 0.08 : 11
- Top-k mean fraud score: 0.2955 CNN risk: LOW Duplicate: 1.00 (exact_duplicate)
- Patch CNN flagged 11 of 51 text-rich patches (top score 0.86).
- Top 5 suspicious regions clustered around image coordinates (128,0), (335,64), (335,477), (320,64), (335,128).
- Predicted modification mix across top regions: 3x CUT, 2x CPI.
- Predicted entity types: 3x Company, 2x Total/payment.
- Document is an EXACT duplicate of a file already analysed in this session (resubmission).