


File: X00016469619_zatca_bad.png SHA-256: cbab68c3d0b56354... Size: 439x1004 px Analysed: 2026-05-12 00:28:53
 UTC Model: FraudX v2-multi (ResNet-18, ep13, thr=0.08)



VERDICT — High risk

VAT discrepancy detected

0%

CNN forgery confidence

The VAT QR contradicts the printed receipt or violates ZATCA Phase-1 spec — review the QR check below.

Patch CNN: Low risk
forgery confidence 0%

VAT QR: Discrepancy
1 QR

Key findings

- No regions exceeded the model's forgery threshold.
- Showing the top 3 highest-confidence regions (highest score 69%). 3 of 81 text-bearing regions exceeded the model's threshold overall.
- Edit type predicted across the top 3 regions: 3 × deletions / erasures.
- Targeted fields: 2 × Company / header info and 1 × Product lines.
- This file is an exact duplicate of one already analysed in this session.
- VAT QR check: VAT number format — VAT# '123' violates ZATCA format (must be 15 digits, first and last digit '3').

VAT QR check

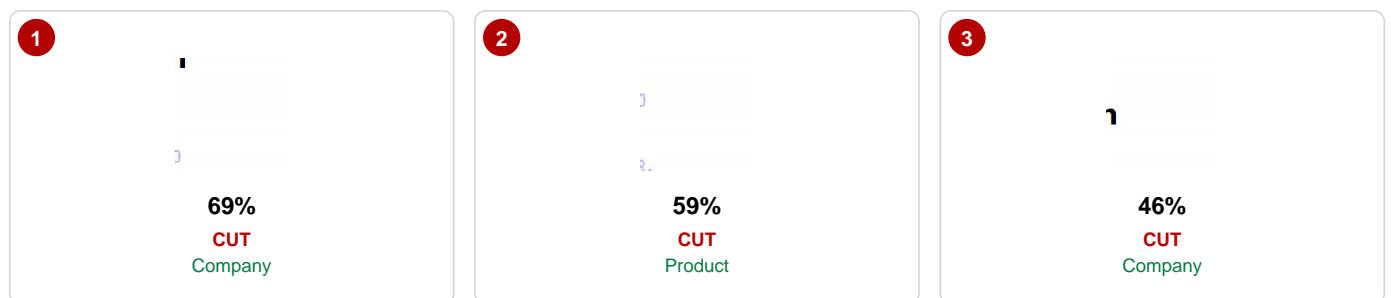
- Al-Othaim Markets · VAT# 123 · 61.83 SAR — FAIL
 - VAT# '123' violates ZATCA format (must be 15 digits, first and last digit '3').

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 (128×128, 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: 90 Text-rich: 81 Flagged ≥ 0.08 : 3
- Top-k mean fraud score: 0.0019 CNN risk: LOW Duplicate: 1.00 (exact_duplicate)
- Patch CNN flagged 3 of 81 text-rich patches (top score 0.69).
- Top 3 suspicious regions clustered around image coordinates (311,64), (311,128), (311,0).
- Predicted modification mix across top regions: 3× CUT.
- Predicted entity types: 2× Company, 1× Product.
- Document is an EXACT duplicate of a file already analysed in this session (resubmission).

ZATCA QR — decoded fields & full check list

seller_name: Al-Othaim Markets
vat_number: 123
timestamp: 2024-08-15T14:32:00Z
total: 61.83
vat: 123.65

[PASS] payload decodes

Base64 + TLV well-formed; 5 tag(s) found.

[PASS] all mandatory tags present

Tags 1 (seller), 2 (VAT#), 3 (timestamp), 4 (total), 5 (VAT) all present.

[FAIL] VAT number format

VAT# '123' violates ZATCA format (must be 15 digits, first and last digit '3').

[PASS] timestamp valid

2024-08-15T14:32:00Z (parses as ISO 8601).

[PASS] amounts numeric

Total 61.83 SAR · VAT 123.65 SAR.

[FAIL] VAT \leq total

VAT 123.65 exceeds total 61.83 — impossible.

[WARN] VAT rate \approx 15%

Implied rate 0.00% deviates 15.0pp from 15%. Could be mixed-rate / exempt items — review the line items.

[WARN] seller name on receipt

'Al-Othaim Markets' not found in OCR'd receipt text (token overlap 0%). OCR may have missed the header.

[PASS] invoice total matches receipt

QR total 61.83 SAR appears in the receipt (closest OCR value 61.83).

[WARN] VAT amount matches receipt

QR VAT 123.65 SAR not directly found in OCR — receipts often round or summarise VAT differently. Manual review suggested.