

Model Card

Use Case: Automated transcription and structuring of user-provided identity documents for internal operational efficiency

1. System / Model Identifier

System Name: Assistive OCR Pipeline for ID Data Transcription

System Version: 0.5

System Type: ML-assisted document transcription pipeline

- Image preprocessing
- Document type classification (EfficientNet CNN)
- Segmentation (YOLO Object Detection — CNNs)
- OCR and text extraction (DL-based OCR)
- Field detection and parsing (NER + rules)

Lifecycle Status: Beta

Responsible Legal Entity: Kwantu Ltd

GitHub: https://github.com/kwantu/document_transcription

2. Regulatory Classification

EU AI Act Risk Category: Non-High-Risk

Applicable Articles: Articles 52 (transparency), general principles

Annex III Applicability: Assessed as not applicable

Human Oversight: Mandatory in all cases

Automation Level: Assistive only; no decision substitution

See annex for justification of this assessment.

3. Intended Purpose and Use

The system is intended **solely** to perform **preliminary transcription and structuring** of user-provided identity documents to improve internal operational efficiency.

The system:

- Does not verify identity
- Does not assess document validity
- Does not determine sufficiency for KYC
- Does not influence eligibility, approval, or rejection

4. Target Users

- Trained onboarding and compliance analysts

5. Explicit Out-of-Scope Uses

- Identity verification or authentication
- Automated KYC decisions
- Beneficiary acceptance, rejection, or throttling
- Biometric identification or categorisation
- Law enforcement or surveillance

6. System Description (Pipeline-Level)

Inputs

- User-submitted images or scans of identity documents
- Accepted formats: JPEG, PNG, PDF

Processing Stages

1. Image quality assessment and normalization
2. Document classification model (EfficientNet CNN)
3. Segmentation (YOLO Object Detection — CNNs)
4. OCR text extraction (deep learning OCR engine)
5. Field extraction and normalization (NER + deterministic rules)
6. Confidence scoring per extracted field

Outputs

- Document type label
- Structured metadata fields
- ID photograph
- Confidence scores per field
- Processing flags (e.g., low image quality)
- Optional verbose processing images (e.g., rotation algorithms, YOLO labelling, text preprocessing)

Explainability

- Confidence scores surfaced for each extracted field
- OCR bounding boxes available for human review
- Rule-based extraction logic documented

7. Role in Business Process

The system operates as an **optional first attempt** at transcription.

Process guarantees:

- A human analyst performs identity verification **in all cases**
- AI output is not required to complete onboarding
- Failure or inaccuracy does not delay or block service access
- Human-only processing is the guaranteed fallback path

The system replaces **manual data entry**, not regulated judgment.

8. Explainability and Transparency

- OCR bounding boxes displayed alongside source image
- Field-level confidence scores shown to operators
- Rule-based extraction logic documented

No automated explanations are presented to customers.

9. Training Data

Data Sources

- ID document datasets from programme data
- Trained on a batch of 500 new style ID cards and 500 old style ID booklets

Geographic Scope

- South Africa only (unless re-trained)

Temporal Coverage

- 2021–2024

Data Governance

- GDPR lawful basis: Legal obligation (KYC/AML)
- Personal data minimization applied
- No storage of raw images beyond defined retention window
- Sensitive attributes not explicitly inferred

Known Limitations

- Reduced accuracy on damaged or low-resolution documents
- Flash and lens flare affects accuracy

10. Evaluation and Performance

Evaluation Dataset

- Validation set covering all supported document types
- Separate stress tests for low-quality images

Key Metrics

KPI	Comments	Old style ID books	New style ID cards
Document classification accuracy	The accuracy of the EfficientNet model in identifying the class of ID document	100%	100%

OCR character accuracy rate (CAR)	The accuracy of the OCR in correctly identifying the characters in the ID number	71.8%	94.41%
Field-level extraction F1 score (Document Number)	The accuracy of the OCR in correctly extracting the entire ID number	66.67%	75.76%

Robustness and Bias Assessment

- Performance evaluated across old style identity booklets and new style identity cards in South Africa
- No material degradation observed across gender-coded names

11. Risk Management and Mitigations

Identified Risks

- Misclassification of document type leading to incorrect extraction
- OCR errors resulting in incorrect personal data
- Over-reliance by operators on extracted values
- Data leakage or excessive retention of identity documents

Mitigations

- Confidence thresholds triggering manual review
- Secure, time-bound storage and encryption of images
- Periodic review of supported document templates

12. Human Oversight

Oversight Measures

- Side-by-side display of original document and extracted fields
- Ability to correct or override all outputs
- Operator training on system limitations and error patterns

Escalation Procedures

- Low-confidence or anomalous outputs routed to senior reviewers

13. Post-Deployment Monitoring

Monitoring Activities

- Accuracy tracking by document type
- Error rate monitoring for new document versions
- Drift detection on OCR confidence distributions

Review Frequency

- Monthly performance review
- Quarterly risk and compliance review

Triggers for Retraining or Suspension

- Field extraction accuracy drop >5%
- Introduction of new national ID formats
- Regulatory or legal changes affecting identity processing

14. Ethical and Legal Considerations

Privacy Impact

- Risk assessment completed
- Data retention strictly limited to regulatory minimums

Fundamental Rights Considerations

- Risk of exclusion due to document quality or language mitigated via human review
- Accessibility guidance provided to end users

Compliance Notes

- Designed to meet EU AI Act Articles 9–15
- Aligned with POPIA, GDPR, AMLD, and internal AI governance policies

15. Caveats and Warnings

This system does not determine identity validity or customer eligibility on its own. All outputs are probabilistic and must be reviewed by a trained human operator. Use without documented human oversight or outside approved jurisdictions constitutes regulatory non-compliance.