iApp Technology · Technical Whitepaper WP-2026-02

Thai National ID Card OCR: Field-Level Accuracy Benchmark

Photographed-copy service, v1.1.1-20260822 · Evaluation executed 22 August 2026 · iApp Technology Co., Ltd., Bangkok, Thailand

Executive summary. We evaluated the iApp Photographed Thai National ID Card OCR service (v1.1.1-20260822) against a leading global cloud OCR service on 33 Thai national ID cards with human-verified field-level ground truth — 594 individual field observations. The iApp service achieved **97.1%** field-level exact match versus 96.3% for the cloud baseline, with the 13-digit national ID number at **100%** on both. In production, the service processes a card in 0.7 s (median) and sustains more than 10,500 cards per hour.

1 Introduction

The Photographed Thai National ID Card OCR API extracts structured fields — the 13-digit national ID number, names in Thai and English, dates, and the address — from photographed copies of Thai national ID cards, and detects whether the copy has been signed. It is operated by iApp Technology Co., Ltd. as a public cloud API at api.iapp.co.th and as a fully self-contained on-premise deployment.

For identity documents, aggregate character accuracy is the wrong yardstick: a single wrong digit in an ID number is a failed extraction regardless of how many other characters were read correctly. This benchmark therefore measures *field-level exact match* against human-verified ground truth, per field, and compares the iApp service against a widely deployed commercial cloud OCR service, referred to here as *Cloud OCR A*.

2 Dataset and Ground Truth

The evaluation draws on iApp's proprietary labeled corpus of 1,000 Thai national ID cards, in which every card is paired with 27 columns of human-verified field-level truth. Thirty-three cards were sampled for the deployment-gate evaluation, giving 594 scored field observations.

The corpus contains real identity documents with real names, addresses, and ID numbers. It is therefore **not published**, in line with Thai PDPA obligations and iApp's data-handling policy. The methodology in Section 3 is specified completely so that the measurement procedure itself is transparent and repeatable on any comparable labeled corpus.

3 Methodology

- **Fairness.** Both systems process byte-identical card images. The two evaluation candidates are identical builds of the production service differing *only* in the OCR engine invoked, so the engine is the only variable; card localisation, cropping, and field parsing are shared.
- **Scoring.** Per field: exact match against the human-verified truth, plus mean character error rate. Dates are normalised to ISO 8601 (`YYYY-MM-DD`) before comparison, so calendar-format differences do not count as errors.
- **No engine-as-referee.** Accuracy is scored against independent ground truth, not against the other engine — scoring against a baseline engine penalises a system precisely when it corrects the baseline's mistakes, which occurred repeatedly in this data.

4 Results

Field	**iApp ID Card OCR**	Cloud OCR A
National ID number (13 digits)	**100.0%**	**100.0%**
Address (full)	**93.9%**	87.9%
Address: house number	**97.0%**	—
Address: subdistrict	97.0%	**100.0%**
All fields (overall)	**97.1%**	96.3%

Table 1: Field-level exact match against human-verified ground truth, 33 cards / 594 field observations, deployment-gate evaluation of the shipped build. Best result per row in bold; — where the baseline figure was not recorded.

4.1 Where the difference comes from

Cloud OCR A showed a systematic weakness on Thai house numbers: across multiple cards it dropped or corrupted the leading digit of compound house numbers (reading `63/160` as `53/160` or `3/160`), which pulls its full-address score down to 87.9%. The iApp service reads these correctly.

The iApp service's residual errors are scattered one-off misreadings rather than a pattern; each occurred on a single card out of thirty-three. One error class is worth flagging for risk assessment: a subdistrict name was once misread as a visually similar sibling name. Rare, but semantically significant — which is why the service returns a per-field confidence score that integrators can use to route low-confidence extractions to review.

5 Production Performance

Metric	Value
Median end-to-end processing time	0.7 s per card
Sustained throughput	2.9 cards/s (≈10,500 cards/hour)
Maximum input file size	2 MB
Supported inputs	JPEG, PNG, HEIC/HEIF

Table 2: Production service performance, measured on the deployed v1.1.1-20260822 service, August 2026.

6 Deployment and Data Security

The public cloud API is GDPR- and PDPA-compliant: uploaded images are processed in memory and are not retained after the response is returned. For organizations with stricter requirements, the service is available as a fully self-contained on-premise deployment in which no card data leaves the customer's infrastructure.

7 Reproducibility

The scoring protocol, normalisation rules, and candidate construction are specified in full in Section 3; the ground-truth corpus is withheld solely because it contains real identity documents. Organizations wishing

to validate the results on their own labeled cards, or to enquire about the identity of the baseline service, may contact sale@iapp.co.th.

About iApp Technology. iApp Technology Co., Ltd. is a Thai artificial-intelligence company providing OCR, document understanding, speech, and large-language-model services for the Thai language.
Web: iapp.co.th · E-mail: sale@iapp.co.th · Tel: +66 86 322 5858