



Thai Vehicle Make and Model Recognition: A Benchmark on Today's Traffic

Service version v3.0.0-20260828 · Measured 28 August 2026 · iApp Technology Co., Ltd., Bangkok, Thailand

Executive summary. Version 3.0 of the iApp Thai Vehicle License Plate OCR service replaces its vehicle make/model recognition engine with one built for the current Thai car market. On 853 real held-out photographs covering all 73 vehicle models prevalent in today's Thai traffic — including 2020–2026 releases and Chinese electric-vehicle brands — v3.0 identifies the vehicle brand correctly on **99.4%** of images (previous engine: 38.5%) and the exact make and model on **99.18%** (top-5: 100.0%), while returning an answer on **100%** of images (previous engine: about 50%). License plate, province, color, and orientation reading are unchanged from version 2.0. The API contract is fully backward compatible: every new output is an additive field.

1 Introduction

The iApp Thai Vehicle License Plate OCR service (LPR) reads Thai license plates from photographs and reports the plate number, registered province, and the vehicle's make, model, body type, color, and orientation. It is operated by iApp Technology Co., Ltd. as a public cloud API at api.iapp.co.th and as a self-contained on-premise deployment.

The service's plate-reading component performs well and is kept unchanged in version 3.0. Its make/model component, however, was built for the market of its era: its recognition catalog predates the recent wave of new Thai-market releases and Chinese electric-vehicle brands, so an increasing share of the vehicles on Thai roads today could not be identified — by design of its catalog, not by defect. Version 3.0 replaces the make/model engine with a deep-learning classifier and a recognition catalog rebuilt for the current market: 73 models across 17 brands, including 2020–2026 releases and electric-vehicle brands such as BYD, MG, GWM, and Neta.

2 Method

- **Test set.** 853 real Thai-market vehicle photographs, held out and never used in training. All 73 supported models (17 brands) are represented. Images vary in angle, color, lighting, and background — everyday capture conditions, not studio shots. The set contains real vehicles and license plates and is withheld under Thai PDPA obligations.
- **Identical inputs.** Both engine generations were scored on exactly the same images, so the comparison is like for like.
- **Metrics.** Top-1 = the single best prediction names the correct make and model; top-5 = the correct model appears among the five most-likely candidates; coverage = the share of images for which the engine returns any make/model answer at all.
- **Catalog-aware split.** 310 of the 853 images (27 of the 73 models) are vehicles introduced after the previous engine's catalog was built. To separate classifier quality from catalog coverage, each headline metric is also reported restricted to the models that exist in the previous engine's own catalog.

3 Results

Metric	Previous engine (v2.0, 2023 catalog)	v3.0
Vehicle brand (make) correct	38.5%	99.4%
Brand — restricted to models in the v2.0 catalog	41.8%	99.3%
Make + model correct (top-1)	10.4%*	99.18%
Make + model within top-5	—	100.0%
Coverage (returns an answer)	≈50%	100%

Table 1: Head-to-head results on the 853-image held-out benchmark, identical inputs to both engines.

* **A fair reading of the previous engine’s numbers.** This benchmark reflects *today’s* traffic: 310 of the 853 images are vehicles introduced after the version 2.0 catalog was built — 2020–2026 releases and Chinese electric-vehicle brands — which that engine cannot identify by design. Version 2.0 was built for the market of its era; the road changed. Version 3.0’s improvement comes from both a stronger classifier and a recognition catalog rebuilt for the current market: even restricted to models the previous catalog does contain, brand accuracy rises from 41.8% to 99.3%.

3.1 Progression across engine generations

The version 3.0 engine is the fifth internal generation of the make/model classifier. Each generation was measured on a held-out test set as coverage and training data grew:

Generation	Models covered	Top-1	Top-5
1	35	95.05%	97.69%
2	64	97.84%	99.41%
3	64	98.24%	99.61%
4	73	97.44%	99.49%
5 (deployed in v3.0)	73	99.18%	99.79%

Table 2: Internal generation history of the make/model classifier on held-out evaluation sets.

4 New Capabilities in Version 3.0

- **Fuel type and EV detection.** Every identified vehicle reports its powertrain (EV, Hybrid, Plug-in Hybrid, Petrol, Diesel), inferred from the identified make and model through a maintained powertrain table — never guessed from pixels. Models sold with several powertrains deliberately return multiple with the full list of options, because the powertrain of such a model is not determinable from a single photograph.
- **Confidence and verification.** Responses carry a make/model confidence score, the top-5 candidate list with probabilities, and an independent brand identification from the vehicle’s logo with its own confidence — allowing integrators to set their own acceptance thresholds.
- **Neighbor-country plates.** Malaysia (production), Myanmar (beta), and Laos (detection-only), reported through a `country_support_level` field.
- **Robust close-up framing.** Brand and model remain correct and deterministic on tight grille close-ups typical of parking-gate and village-entrance cameras — the same vehicle produces the same answer on every scan.

5 Production Performance

Metric	Value
Median end-to-end processing time	0.5 s per image
Sustained throughput	3.1 images/s (≈11,000 images/hour)
Supported inputs	JPEG, PNG (up to 2 MB)

Table 3: Measured on the deployed v3.0.0-20260828 production service, 28 August 2026, at request concurrency 4.

6 Compatibility and Rollout

The version 3.0 API contract is fully backward compatible: the endpoints and request format are unchanged, and every existing response field keeps its exact name and format. All new outputs are additive fields that existing integrations can simply ignore. Version 3.0 is served on a dedicated preview endpoint from 28 August 2026; the main endpoints switch to the version 3.0 engine on 11 September 2026 with no action required from integrators.

7 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. The service runs entirely on on-premise-capable infrastructure — no image is sent to any external service — and a fully self-contained deployment is available in which no photograph leaves the customer’s infrastructure.

8 Reproducibility

The benchmark method in Section 2 is fully specified: any party can score this service, or any competitor, by submitting their own labeled vehicle photographs and comparing the returned make and model against ground truth. The benchmark set itself is withheld because it contains real vehicles and license plates. Enquiries may be addressed to 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