MARLUC

CASE STUDY · MARLUC SYNOPSIS

Thailand diabetic retinopathy screening AI

IMPLEMENTED 2018–2024 · HEALTH AND LIFE SCIENCES · THAILAND · BAND C · PARTLY PROVED

VERSION 1.0 · PUBLISHED 16 NOVEMBER 2026 · UPDATED 16 NOVEMBER 2026 · FACTS CHECKED 16 NOVEMBER 2026

A deep-learning model graded eye photographs in real time at primary care sites in Thailand's national diabetic retinopathy screening programme. A prospective study found it as accurate as retina specialists, while earlier field observation found many images rejected and slow connections.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

Frontier case: AI used for a scientific, medical or engineering breakthrough. MarLuc grades what has changed in practice so far.

AI used: Deep-learning image model grading retinal photographs for diabetic retinopathy.

Implemented: 2018–2024. Grade: band C, Process change, Partly proved.

Headline result, as reported: Prospective study at nine primary care sites: sensitivity for vision-threatening retinopathy 91.4% for the AI against 84.8% for retina specialists, with similar specificity

1 · The case

The facts, as reported by the 5 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Department of Medical Services, Thai Ministry of Public Health, with Rajavithi Hospital. Thailand's Department of Medical Services, part of the Ministry of Public Health, runs the national diabetic retinopathy screening programme. Rajavithi Hospital in Bangkok co-funded and led the clinical study with Google. Canvas block: 4. Operations and process. Also 6. Data and technology, 5. People and skills.

Business challenge

The challenge. Thailand has around 4.5 million people with diabetes and about 200 retina specialists, per MIT Technology Review, against a ministry goal of screening 60% of people with diabetes each year. Under the previous system, nurses photographed eyes and sent images to a specialist elsewhere, which could take up to 10 weeks.

How it was addressed. A deep-learning system developed by Google graded fundus photographs and gave referral recommendations in real time at nine primary care sites in the national programme. As a safety measure, regional retina specialists over-read every image. Results were compared with an adjudicated reference standard from fellowship-trained retina specialists. The Department of Medical Services then registered an implementation study to expand AI screening to 13 health districts, and in 2024 Google licensed the model to partners including Perceptra for use in Thai public hospitals. [1] [2] [3] [4] [5]

Who implemented it

Google (Model or cloud provider): Developed the model and co-funded the study.
Rajavithi Hospital (University or research lab): Co-funded and led the clinical study in Bangkok.
Perceptra (Software vendor): Licensed the model in 2024; a partnership with the Department of Medical Services for public hospitals is planned.

Technology (SEAM pillar E)

SEAM types: T8 Predictive / ML models. Deployment: Cloud APIs. Vendors and products: Google, Google deep-learning diabetic retinopathy model.

Where it was built in: Nurses photograph the retina at primary care sites; images are uploaded for grading and a referral recommendation returns in real time.

What changed

Before: images sent away for grading, with results taking up to 10 weeks. After, at the study sites: an AI grade and referral recommendation in real time, with specialist over-reading as a safety check. In the prospective study the AI was more sensitive than the specialist over-readers and equally specific. [1] [2] [3] [4] [5]

What it cost

Not reported. The cost-utility model tested AI prices of US$1 to US$4 per patient.

What it saved

Modelled, not measured: AI screening was associated with a societal cost reduction of about US$2.70 per patient and similar quality-adjusted life years compared with trained human graders, in a peer-reviewed cost-utility analysis. [1] [2] [3] [4] [5]

Governance and control

Retina specialists over-read each AI-graded image during the study. Google held CE mark clearance covering Thailand in 2020, per MIT Technology Review; licensing partners were to seek local regulatory approvals from 2024, per Google. Accountable owner for national rollout: the Department of Medical Services sponsors the expansion study.

S2 · Accountable owner · Claimed
The Department of Medical Services runs the national screening programme and sponsors the registered expansion study. [3] [5]

A1 · Safe use · Not reported
Not reported.

A2 · Human in the loop · Evidenced
The peer-reviewed study reports that regional retina specialists over-read each image as a safety mechanism. [1]

A3 · Monitoring and assurance · Evidenced
Performance was measured prospectively against an adjudicated reference standard, and field problems were studied by observation and interviews. [1] [4]

A4 · Vendor AI · Not reported
Google funded the study and supplies the model; contract terms and data arrangements are not reported. [1]

A5 · Agent controls · Not reported
No agents described.

A6 · Standards and regulation · Evidenced
MIT Technology Review reported in 2020 that Google had CE mark clearance covering Thailand; licensing partners were to seek local regulatory approvals from 2024. [4] [5]

What worked, what did not, what is unproven

WORKED

– Accuracy at the point of care: sensitivity 91.4% for vision-threatening retinopathy against 84.8% for specialist over-readers, with similar specificity.

– Real-time results replaced waits of up to 10 weeks for grading at the study sites.

– Modelled cost savings against human graders, with similar health outcomes.

DID NOT

– In an earlier field study in 11 clinics, more than a fifth of images were rejected as too poor to grade, and patients were sent to another clinic on another day, per MIT Technology Review.

– Slow internet connections delayed cloud processing in several clinics, per the same report.

– The study authors said socioenvironmental factors and workflows must be considered when implementing at scale.

UNPROVEN

– Results of the 13-district expansion study: its registry record still shows not yet recruiting, last updated in 2022.

– Patient outcomes, such as prevented vision loss, from AI-led screening.

– How many Thai patients are now screened with the model: no country figure was found in sources opened.

Timeline

Dec 2018 Prospective real-time screening study starts at nine primary care sites [1]

Mar 2020 Study recruitment ends after 7,940 patients [1]

Apr 2020 MIT Technology Review reports field problems in 11 Thai clinics: rejected images and slow connections [4]

Aug 2022 Department of Medical Services registers an implementation study to expand AI screening to 13 health districts [3]

Oct 2024 Google licenses the model to partners including Perceptra, with a planned partnership with the Department of Medical Services for public hospitals [5]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: C, Process change. Proof: Partly proved. Level of change: Process change.

Process change (Level 2, band C). At the study sites the screening process was rebuilt around a real-time AI grade, with specialists checking rather than grading first.

Graded C: at the study sites, screening moved from photographs sent to distant specialists to a real-time AI grade and referral decision at the point of care, with accuracy proved prospectively. National scale and patient outcomes are not yet reported.

The SEAM view and lessons learnt

Accuracy was shown in a prospective study; the harder work sat at the join with the clinic: image quality, connectivity and what happens to patients whose images are rejected. For larger organisations, test the full workflow in real settings, including the failure path, before scaling an accurate model.

1. Design for the cases the AI rejects as well as the ones it grades.
2. Check connectivity and image quality at each site before go-live.
3. Keep a specialist safety check until local performance is confirmed.

Evidence

Score 3 of 4. Independent peer review: a prospective interventional cohort study in The Lancet Digital Health, funded by Google and Rajavithi Hospital; a peer-reviewed cost-utility model with Thai ministry, university and Google authors; critical field reporting by MIT Technology Review.

Evaluation method. Prospective interventional cohort: AI grades compared with an adjudicated reference standard and with specialist over-readers; separate decision-tree and Markov cost-utility model.

Sources disagree. The Lancet Digital Health study covered nine primary care sites; the field study reported by MIT Technology Review covered 11 clinics, a separate piece of work. Google's 2024 post says the model had supported more than 600,000 screenings worldwide, without a breakdown by country.

Status. Ongoing: licensed to Perceptra and other partners from October 2024, with a planned

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/thailand-dms-rajavithi-diabetic-retinopathy-ai

[1] Google Research: publication page and abstract, Real-time diabetic retinopathy screening by deep learning in a multisite national screening programme (The Lancet Digital Health, 2022)
https://research.google/pubs/real-time-diabetic-retinopathy-screening-by-deep-learning-in-a-multisite-national-screening-programme-a-prospective-interventional-cohort-study

[2] Ophthalmology and Therapy: Cost-utility analysis of deep learning and trained human graders for diabetic retinopathy screening in a nationwide program (2023)
https://link.springer.com/content/pdf/10.1007/s40123-023-00688-y.pdf

[3] ClinicalTrials.gov: NCT05527535, Expansion of integrated AI solution for diabetic retinopathy screening in Thailand (registry record)
https://clinicaltrials.gov/api/v2/studies/NCT05527535

[4] MIT Technology Review: Google's medical AI was super accurate in a lab. Real life was a different story (April 2020)
https://www.technologyreview.com/2020/04/27/1000658/google-medical-ai-accurate-lab-real-life-clinic-covid-diabetes-retina-disease/

[5] Google: How AI is making eyesight-saving care more accessible in resource-constrained settings (October 2024)
https://blog.google/technology/ai/arda-diabetic-retinopathy-india-thailand/

Version 1.0 · Published 16 November 2026 · Updated 16 November 2026 · Facts checked against sources 16 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/thailand-dms-rajavithi-diabetic-retinopathy-ai

VERSION HISTORY

v1.0 16 November 2026 First published