



OPEN ACCESS

EDITED BY

Roshan Joy Martis,
Manipal Institute of Technology Bengaluru,
India

REVIEWED BY

Mohan Bhandari,
Samridhi College, Nepal
Lopamudra Das Ghosh,
Texas A and M University, United States

*CORRESPONDENCE

Imanol Pinto
✉ imanol.pinto.lopez@navarra.es

RECEIVED 17 December 2024

ACCEPTED 02 June 2025

PUBLISHED 19 June 2025

CITATION

Pinto I, Olazarán Á, Jurío D, De la Osa B,
Sainz M, Oscoz A, Ballaz J, Gorricho J, Galar M
and Andonegui J (2025) Improving diabetic
retinopathy screening using artificial
intelligence: design, evaluation and before-
and-after study of a custom development.
Front. Digit. Health 7:1547045.
doi: 10.3389/fdgth.2025.1547045

COPYRIGHT

© 2025 Pinto, Olazarán, Jurío, De la Osa,
Sainz, Oscoz, Ballaz, Gorricho, Galar and
Andonegui. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Improving diabetic retinopathy screening using artificial intelligence: design, evaluation and before-and-after study of a custom development

Imanol Pinto^{1*}, Álvaro Olazarán¹, David Jurío¹, Borja De la Osa¹, Miguel Sainz¹, Aritz Oscoz¹, Jerónimo Ballaz¹, Javier Gorricho^{2,3}, Mikel Galar^{3,4} and José Andonegui^{3,5}

¹Health Technology Services, General Directorate of Telecommunications and Digitalization (DGTD), Sarriguren, Spain, ²Health Outcomes Evaluation and Dissemination Service, Navarre Public Health Service (SNS-O), Pamplona, Spain, ³Navarre Institute of Health Research, IdiSNA, Pamplona, Spain, ⁴Institute of Smart Cities (ISC), Public University of Navarre (UPNA), Pamplona, Spain, ⁵Ophthalmology Service, University Hospital of Navarre, Pamplona, Spain

Background: The worst outcomes of diabetic retinopathy (DR) can be prevented by implementing DR screening programs assisted by AI. At the University Hospital of Navarre (HUN), Spain, general practitioners (GPs) grade fundus images in an ongoing DR screening program, referring to a second screening level (ophthalmologist) target patients.

Methods: After collecting their requirements, HUN decided to develop a custom AI tool, called NaIA-RD, to assist their GPs in DR screening. This paper introduces NaIA-RD, details its implementation, and highlights its unique combination of DR and retinal image quality grading in a single system. Its impact is measured in an unprecedented before-and-after study that compares 19,828 patients screened before NaIA-RD's implementation and 22,962 patients screened after.

Results: NaIA-RD influenced the screening criteria of 3/4 GPs, increasing their sensitivity. Agreement between NaIA-RD and the GPs was high for non-referral proposals (94.6% or more), but lower and variable (from 23.4% to 86.6%) for referral proposals. An ophthalmologist discarded a NaIA-RD error in most of contradicted referral proposals by labeling the 93% of a sample of them as referable. In an autonomous setup, NaIA-RD would have reduced the study visualization workload by 4.27 times without missing a single case of sight-threatening DR referred by a GP.

Conclusion: DR screening was more effective when supported by NaIA-RD, which could be safely used to autonomously perform the first level of screening. This shows how AI devices, when seamlessly integrated into clinical workflows, can help improve clinical pathways in the long term.

KEYWORDS

diabetic retinopathy, AI medical device, decision-support system, deep learning, before-and-after study

1 Introduction

Diabetic retinopathy (DR) is the leading cause of vision loss among the working-age population in developed countries (1), but the worst outcomes can be prevented with early detection and treatment. In fact, the implementation of DR screening programs is recommended by the American Diabetes Association (2) and the International Council of Ophthalmology (3).

DR screening is usually performed by trained personnel who are not necessarily ophthalmologists. They grade (visualize and assess) eye fundus photographs, called retinographies, which are taken with a non-mydriatic digital camera. These graders look for DR signs, such as hemorrhages, and refer (send) the patient to an ophthalmologist if necessary. Their primary goal is to refer patients who need evaluation by a specialist, who will then decide if treatment or further follow-up is necessary.

Given this workflow, automated methods can bring a more efficient and cost-effective DR screening (4, 5). Several AI-based medical devices (CE-marked or FDA-approved) are available for this purpose (6). These tools promise to eliminate or reduce the burden of manual grading.

However, the performance of AI-based medical devices often degrades when they are used outside the clinical sites from which their data originated (7, 8). A recent study (8) compared seven algorithms that were being used in clinics, and highlighted the need for prospective, interventional trials for commercialized products, as they measured a wide range of sensitivities (50.98%–85.90%). These interventional studies are not required to obtain a CE mark or FDA approval and are therefore very rare (9).

To make matters worse, AI-powered medical devices are often negatively affected by their environment: task sharing, user knowledge, infrastructure, integration, and socio-environmental factors are challenges that hinder their success (10, 11). This problem is exacerbated when the clinical protocol is already established before it is supported by AI. Sometimes it is simply not feasible to implement a generic AI tool into an ongoing DR screening program.

This is the case at the University Hospital of Navarre (HUN). This hospital in Spain has been offering a public DR screening program since 2015, and has been working with us to support it with AI. We collected HUN's DR screening requirements and found that none of the available CE-marked medical devices could be used without significant limitations and risks. Therefore, we developed a custom, AI-based DR screening tool for HUN: NaIA-RD.

After validating the performance of NaIA-RD using two private and six public datasets (12–17), we deployed it in July 2020, integrated into the Hospital Information System (HIS). It has been used for routine DR screening for more than three years. Using the data from this interventional prospective study, we compared how DR screening was performed before and after the deployment of NaIA-RD, measuring how the tool has influenced clinical decisions.

This paper makes two significant contributions to the literature: First, it measures the impact of an AI tool on real-world clinical decisions. Most published prospective studies

typically compare the AI tool's performance to that of manual graders (18–25), or they evaluate the tool's ability to reduce the burden of manual grading (26), but they do not assess how the tool has influenced clinician behavior. Second, this paper details a novel procedure for combining DR grading models with retinal image quality (gradability) models. Our proposed system selects the most clinically suitable image (field of view) and consistently provides independent DR and gradability scores. We found no other work describing how to integrate both assessments into a single AI system, although there are numerous publications dedicated to each topic individually (27, 28).

This paper is organized as follows. First, Section 2.1 details the DR screening process at HUN, before and after NaIA-RD's assistance, and Section 2.2 summarizes the hospital's requirements and how commercial AI devices do not meet them. Section 2.3 details NaIA-RD's development, from system design to neural network training, calibration, interpretation, image enhancement, and machine learning operations (MLOps). Then, Section 3 reports the results: First, in laboratory settings (Sections 3.1–3.3) and last, in real clinical settings (Section 3.4). Section 4 discusses NaIA-RD's performance, impact, and limitations. We draw our final conclusions in Section 5.

2 Materials and methods

2.1 DR screening at HUN

In 2015 a DR tele-screening program was set up at HUN. Since then, all patients assigned to the hospital and diagnosed with Type 2 diabetes have been scheduled for annual retinal imaging and screening. Over these years, the number of patients screened has steadily increased, reaching nearly 8,000 in 2023.

A team of four primary care general practitioners (GPs)—who received specific training (29, 30)—have remotely assessed retinal images using a centralized HIS. When they detected signs of referable DR or the eye fundus was non-gradable due to insufficient image quality, they referred (sent) the images (grouped as a study) to a second screening level (an ophthalmologist), who decided whether an on-site eye examination was necessary.

The following subsections further explain this DR screening protocol (Section 2.1.1) and how it has been supported by AI (Section 2.1.2).

2.1.1 The DR screening protocol

Figure 1 shows the screening process of HUN using the Business Process Model Notation (BPMN) standard. To model this process, we visited nurses, GPs, and ophthalmologists at the screening sites. Then, we validated our observations using anonymized hospital data. In summary, DR screening at HUN is performed in three main steps:

1. **Image taking (nurse).** A nurse usually takes two non-mydriatic fundus images: a macula-centered (central) and an optic-disc-centered (nasal) fundus field. However,

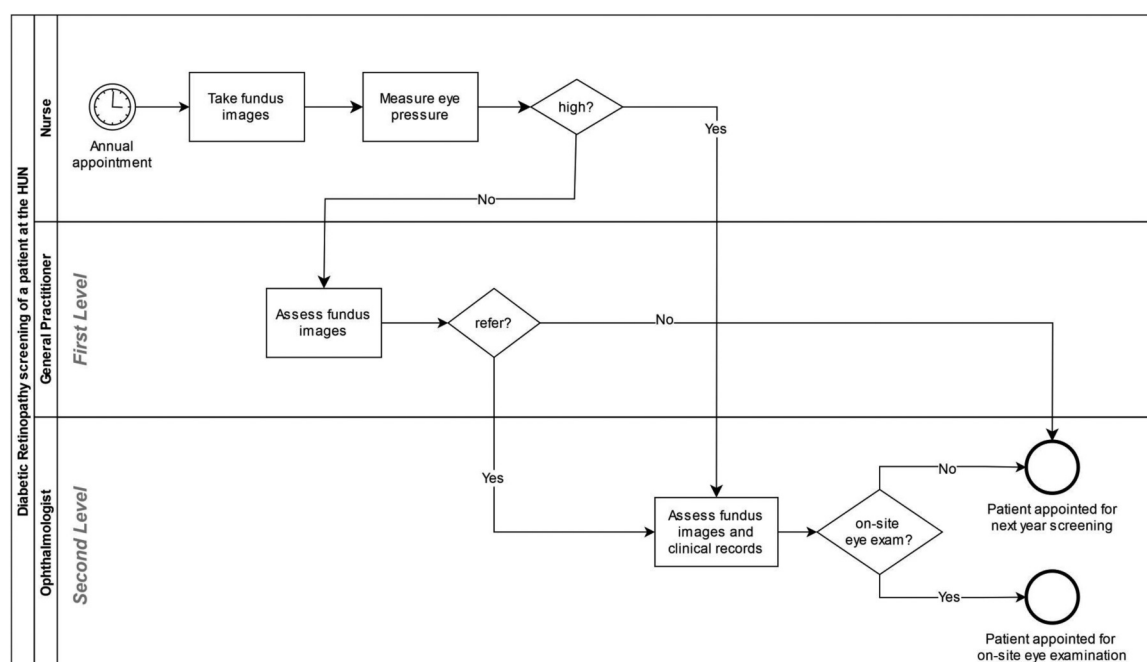


FIGURE 1

BPMN diagram of the process of screening a patient at HUN. Tasks are represented with squares, diamonds represent bifurcations, and circles represent start and end events. Patients appointed for an on-site eye examination abandon the DR screening program until an ophthalmologist decides otherwise.

additional images may be taken if necessary. These fundus images are uploaded to a centralized Picture Acquisition Server (PACS) as a study, which is composed of two eyes (left and right), and will be assessed by a first screening level. The nurse also measures the intraocular pressure.¹ If it is high, the nurse will refer the study directly to the second screening level.

2. **First screening level (GP).** A trained GP visualizes the images, grades DR following the International Clinical Diabetic Retinopathy (ICDR) severity scale (31), and decides whether the study should be referred to the second screening level. They should refer a study if it is not gradable or if it shows signs of more than mild DR (3). Their primary goal is to refer patients who need to be scheduled for an on-site eye examination.
3. **Second screening level (ophthalmologist).** An ophthalmologist grades the referred study following the ICDR scale (31)² and determines if the patient needs an on-site eye examination. This decision is based on the fundus images, patient history, and glycemic (hemoglobin A1c) measurements.

If the second screening level determines that an on-site eye examination is needed, the patient will leave the DR screening program (a retina specialist will examine the patient and decide if treatment or other outpatient care is needed). Otherwise, the patient will be scheduled for the next year's DR screening imaging session.

Regarding the imaging protocol, Figure 2a shows an example of a screened eye composed of the two non-mydratic fundus images that the nurse usually takes, which are often accompanied by a composite image. The composite does not add any new clinical information, as it is just a collage of the central and nasal images. However, as we show in Figure 2b, there is no guarantee that nurses will strictly follow this protocol. In fact, more images are often included in difficult cases, as it is done in other hospitals (11).

This screening protocol is based on the guidelines published by the International Council of Ophthalmology in 2017 (3). Many other DR screening programs around the world follow the same or similar guidelines (32–37), but each program has its own characteristics. For example, the NHS Diabetic Eye Screening Program (36) adds an arbitration grading step in case of disagreement between first- and second-level graders and always performs two-field mydratic photography. Closely related, the Scottish DR Screening Program uses a microaneurysm detection software prior to manual grading and performs single-field non-mydratic photography (34). On the other hand, the Singapore Integrated Diabetic Retinopathy Program centralizes a single human level of screening and is piloting

¹Intraocular pressure is measured for safety reasons, as ocular hypertension usually does not show any findings in fundus images.

²The ophthalmologist overrides the grade of the GP.




FIGURE 2

Two sample studies (left eyes only) from the DR screening program of HUN. Note that some fundus fields are often redundant. (a) Sample eye study composed by central, nasal and composite fundus fields. (b) Sample eye study composed by repeated fundus fields: central, nasal, central, OD down, no OD, and OD up (OD refers to Optical Disc).

SELENA+,³ a Deep Learning system that fully automates the first screening level (38–40).

2.1.2 AI assistance

Since July 2020, the DR screening protocol has been assisted by NaIA-RD. However, the screening protocol has not changed with the AI. In this new setting, GPs review the screening proposal of NaIA-RD before assessing the images, while the rest of the screening steps remain the same. NaIA-RD has enabled the following advanced features:

- The HIS shows NaIA-RD's motivated screening proposal for each study. This proposal is a referral (due to DR or non-gradable fundus) or non-referral recommendation. When the proposal is accepted, the clinical report is generated automatically.
- The HIS study worklist can be sorted by NaIA-RD's outputs, either by referability (a single probability score) or by category (non-referable, referable DR or non-gradable).
- Lesions detected by NaIA-RD are highlighted on the image when the AI recommends a referral due to DR.
- The original fundus image is enhanced by NaIA-RD.

2.2 Motivation for a custom AI development

When the use of AI to assist in DR screening at HUN was first considered, we gathered the requirements for a tool that could effectively support the process and explored CE-marked products that might meet these needs. However, we found none that fulfilled the criteria. As a result, HUN opted to develop NaIA-RD as an in-house AI solution. This section outlines the reasoning behind this decision. Section 2.2.1 describes the requirements, and Section 2.1.2 summarizes the commercial medical devices considered.

2.2.1 Requirements

We met diverse stakeholders from the hospital to collect a broad set of goals, expectations and restrictions: clinicians, ophthalmologists, nurses, IT engineers and managers were interviewed. For brevity, a detailed list of the collected requirements is provided in the [Supplementary Material](#). However, they can be summarized as follows. The AI tool should:

- Be compatible with the current DR screening protocol, patient groups and cameras, while allowing for future inclusion of Type I diabetic patients and new camera models.
- Assist with the fundus image assessment task of the first screening level, enabling task automation and worklist prioritization in the HIS through data-level integration.
- Support workflow orchestration, as well as monitoring of disease prevalence and retinal image quality.
- Offer interpretability and image enhancement features to facilitate human assessment.

³SELENA+ official web page: <https://www.synapse.sg/healthtech/health-ai/selena/>

2.2.2 Commercial medical devices

Following the requirements, we first explored the acquisition of a commercial solution. We evaluated six Class IIa CE-marked devices: IDx-DR,⁴ EyeArt,⁵ Retmarker,⁶ OpthAI,⁷ RetCad⁸ and SELENA+. A detailed comparison of these devices is provided in the [Supplementary Material](#). Note that we discarded products with insufficient public information, those requiring a non-standard 45° field camera, or those self-certified as Class I medical devices.

According to this comparison, we could not find any Class II CE-marked device that met the requirements without significant limitations and risks:

- No tool fully supported the current patient population and imaging protocol (including cameras and study format).
- Few tools offered a data-level integration API, and those that did, lacked a detailed gradability (image quality) score.
- Few tools provided interpretable results or enhanced fundus images.

As a consequence, HUN decided to request to the competent governmental units⁹ the development of NaIA-RD, whose details we present in the following sections.

2.3 Development of a custom AI tool for DR screening

In this section we describe the technical design of NaIA-RD, including its neural networks, datasets, calibration, interpretability, image enhancement and MLOps. For brevity, a summary of the development project life cycle can be found in the [Supplementary Material](#).

2.3.1 Architecture

We designed NaIA-RD as a modular system consisting of three neural networks and a software component that orchestrates them. The neural networks are the following:

1. *Field Classifier*: Identifies the field of view of a fundus image (described in [Section 2.3.1.1](#)).
2. *Gradability Classifier*: Determines if a fundus image is gradable (described in [Section 2.3.1.2](#)).
3. *DR Classifier*: Determines if a fundus image shows referable DR (described in [Section 2.3.1.3](#)).

A software component we call Orchestrator uses these neural networks to generate a DR screening proposal in response to an incoming eye screening request from the HIS. The request consists of a list of fundus images belonging to the same eye. [Figure 3](#) summarizes this architecture.

At runtime, the Field Classifier is executed to identify the fundus images that are the most similar to the central and nasal fundus fields. Then, the Gradability and DR Classifiers are executed: While DR is evaluated in both the central and nasal fields, gradability is only evaluated in the central field, as it is the most representative for assessing DR gradability (28, 41). If the eye screening request is formed by a single image, the Orchestrator assumes that it is a central fundus field, and evaluates both DR and gradability on that image.

The Orchestrator will return a referral proposal if either central or nasal DR output is positive¹⁰ (more than mild DR is detected) or if the central fundus field image is not gradable. The Orchestrator will return a non-referral proposal in all other cases. This behavior is consistent with international guidelines for DR screening, as a non-gradable eye fundus should always be referred (3).

Note that the Orchestrator always assesses DR even if an image is classified as non-gradable. Thus, the output scores of DR and Gradability Classifiers (which are independent of each other), are always included in the screening proposal.

2.3.1.1 Field classifier

The Field Classifier is a neural network that classifies each fundus image into 7 custom categories, listed in [Table 1](#) and previously illustrated in [Figure 2](#). We chose these categories because they adequately describe all images taken since the start of the screening program. Our categories do not require a distinction between the right and left eye (laterality), as the ETDRS imaging protocol does (42). In fact, the laterality of the eye is included as a tag in the DICOM file, so there is no need to use a model to identify the laterality of each fundus field.

NaIA-RD uses the Field Classifier assuming that all the images of the request belong to the same eye. However, it does not impose any restriction on the imaging protocol: it works with any number of fundus images, ignoring non-standard ones. When multiple central and nasal fields are taken in an eye, this method usually selects the best quality image per category, as they resemble the most to the ideal fundus field.

2.3.1.2 Gradability classifier

The Gradability Classifier is a neural network that classifies a central fundus field image as gradable/non-gradable, returning also a gradability score. This model produces a gradable (negative) output if at least 80% of the eye fundus is visible (28). The gradability score is always included in NaIA-RD proposals as an image quality measurement.

⁴IDx-DR: <https://www.healthvisors.com/idx-dr/>

⁵EyeArt: <https://www.eyenuk.com/en/products/eyeart/>

⁶Retmarker: <https://www.retmarker.com/morescreening/>

⁷OpthAI: <https://www.opthai.com/en/>

⁸RetCad: <https://retcad.eu/>

⁹The Navarre Public Health Service and the Health Technology Services of the Government of Navarre are the official IT service of HUN.

¹⁰If both central and nasal fundus fields show DR signs, the Orchestrator will return the most severe DR scores.

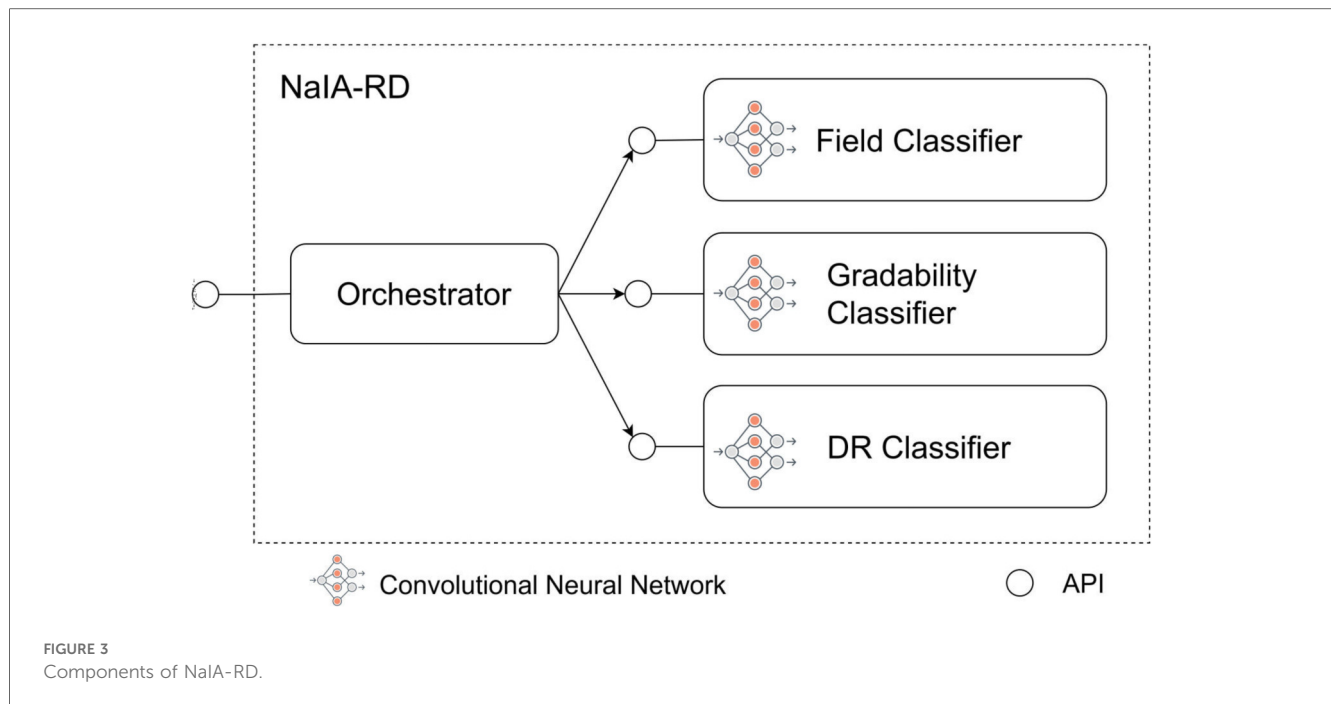


TABLE 1 Field selection model categories and their frequency (OD refers to optical disk).

Category	Description	Frequency	Samples
Central	ETDRS field 2, centered on the macula.	29.5%	28,634
Nasal	ETDRS field 1, centered on the optical disk.	29.5%	28,749
OD up	The image is centered on the inferior arcade, and the optical disk is visible.	0.5%	489
OD down	The image is centered on the superior arcade, and the optical disk is visible.	0.5%	482
No OD	Optical disk is not visible.	3.9%	3,769
Temporal	Temporal to the macula, while at least a portion of the optical disk is visible.	1.7%	1,626
Composite	Composition of multiple fields.	34.4%	33,464

2.3.1.3 DR classifier

The DR Classifier is a neural network trained as a binary classifier to detect more than mild DR signs based on the ICDR scale. It assesses DR as referable/non-referable, classifying the fundus image as referable if it shows more abnormalities than just microaneurysms (31).

This model returns a negative output when the input image is deemed non-referable due to DR. This can occur if the fundus is barely visible in a non-gradable image. However, if a single hemorrhage is detected in a poor quality image, the DR Classifier will produce a positive output. For this reason, the DR score returned by this model is always included in NaIA-RD proposals as a measure of the DR signs present in the eye.

2.3.2 Training techniques

All three neural networks mentioned above (Field Classifier, Gradability Classifier, and DR Classifier) are ResNet34 Convolutional Neural Networks (43). We trained all of them using `fast.ai` v1 library (44), which provides a powerful wrapper of Pytorch. Each model uses square images of different sizes: 150 px, 400 px and 700 px for Field, Gradability and DR Classifiers, respectively. We used a common preprocessing step for all models: finding and cropping the eye fundus circumference before the image is resized. We used standard `opencv` library functions for this task.

We initially loaded all the models with pretrained ImageNet weights, and progressively increased image size in each training iteration using progressive resizing (45). We also used `fast.ai`'s Batch Normalization, Weight Decay, MixUp and LabelSmoothing as regularization techniques. Additionally, we used several random image augmentations, such as: brightness and contrast adjusting, rotating, warping and cropping.

We chose the best models based on their performance on the validation sets. The weights were saved at the end of each epoch when a maximum metric value was reached. For DR and Gradability Classifiers, we used the Area Under the Receiver Operating Characteristic Curve (AUROC, also known as AUC) metric, as a binary classifier metric that does not require a decision threshold. We obtained an AUC value of 0.979 for the DR Classifier, and an AUC value of 0.982 for the Gradability Classifier. For the Field Classifier, which is a multi-class model, we used the Cohen kappa metric rather than the AUC. This is due to kappa's superior simplicity for assessing multi-class classification problems (46). We obtained a Cohen kappa value of 0.976 in Field Classifier's validation set.

TABLE 2 Summary of datasets used for NaIA-RD development. Note that the Gold Standard was used to test multiple components. Also, some datasets are not used for training, validation, or testing. For example, OIA-DDR and EyeQ datasets are used for testing only.

Dataset	Source	Components	Classes	Image quantity		
				Train	Valid	Test
Fundus field	Private	Field classifier	7	36,317	29,968	29,928
Gradability	Private	Gradability classifier	2	8,744	8,592	–
DR	Private	DR classifier	2	17,336	8,592	–
Gold standard	Private	NaIA-RD	3	–	–	984
		DR classifier				
		Gradability classifier				
EyePACS (Kaggle) (12)	Public	DR classifier	2	77,787	–	10,906
APTOS (Kaggle) (13)	Public	DR classifier	2	1,831	–	1,831
Messidor-2 (14, 47)	Public	DR classifier	2	872	–	872
IDRiD (15)	Public	DR classifier	2	413	–	103
OIA-DDR (16)	Public	DR classifier	3	–	–	4,105
		Gradability classifier				
EyeQ (17)	Public	Gradability classifier	2	–	–	16,249

Then, we calibrated DR and Gradability Classifiers as we explain in Section 2.3.4. Using the calibrated outputs, we chose the decision thresholds that maximized the arithmetic mean of the recalls of the positive and negative classes. We found a best threshold of 0.1 for the DR Classifier, which gave a sensitivity of 90.29% and a specificity of 95.92% on the validation set; and a threshold of 0.5 for the Gradability Classifier, which gave a sensitivity of 85.78% and a specificity of 96.07% on the validation set.

2.3.3 Datasets

We used 10 datasets to develop NaIA-RD, both public and private, which are summarized in Table 2. The following subsections describe them, organized by the task they attempt to solve (field, DR, or gradability classification). The Gold Standard dataset, which we consider to be the clinical reference standard for NaIA-RD, is also presented.

2.3.3.1 Field classification

We used a private dataset to train, validate, and test the Fundus Field Classifier neural network. We classified more than 96,000 HUN images into one of the 7 fundus fields listed in Table 1. We detail the obtained test metrics in Section 3.1.

2.3.3.2 DR classification

We used the following public and private datasets to train, validate and test the DR classification model:

1. *Private datasets*: We created a private dataset (DR dataset), labeled by our engineering team, to train and select the best DR Classifier model (validation). We ensured that its 25,928 images were labeled as referable/non-referable based on visible DR signs, strictly following the ICDR grading standard (31). In addition, the DR Classifier was carefully tested using the Gold Standard dataset (Section 2.3.3.4).
2. *Public datasets*: We used 4 popular public datasets to train and test the DR Classifier: EyePACS from the Kaggle 2016 competition (12, 48, 49), APTOS from the Kaggle 2019 competition (13), Messidor-2 (14, 47) and IDRiD (15).

Additionally, we used the OIA-DDR dataset (16) for external validation, meaning that its images were used only to test the generalizability of the model.

We trained the DR Classifier jointly using the aforementioned public and private datasets, excluding the Gold Standard and OIA-DDR. Figure 4 shows that most of this training set comes from EyePACS, while private data is only 17.65% of the total. These datasets involve diverse patient populations, cameras and imaging conditions. Using multiple datasets for training DR classifiers is a common approach (27). However, we chose the DR Classifier with the best AUC in our private validation subset (0.979), as it represented the target data distribution.

We also tested the DR Classifier using the same public datasets we used for training. Specifically, we used the public test sets of EyePACS Kaggle 2016 and IDRiD. As we could not find any defined test set, we used a random 50% sample of APTOS Kaggle 2019 and Messidor-2 datasets. Since we found some incorrect labels in the APTOS dataset, we relabeled some *mild DR* images (9%) as *moderate DR* (referable), according to the ICDR severity scale (31).

Finally, we evaluated DR Classifier's generalization capabilities using the OIA-DDR dataset, which was totally excluded from the training process. OIA-DDR is formed by 13,000 labeled fundus images of 9,500 patients, which were taken in 147 hospitals using 42 different fundus camera models (Topcon D7000, Topcon NW48, Nikon D5200, Canon CR 2 and others). All images were labeled by four professional graders. The authors propose a random 30% of this data as a test set, from which we excluded 346 images labeled as non-gradable. Our final test set consisted of 3,759 images, where 1,690 were labeled as referable (45%). We used the non-gradable images to assess the Gradability Classifier, as we explain in Section 2.3.3.3.

All these public datasets follow the ICDR severity scale for DR grading, with DR grades ranging from 0 to 4 (0—No DR, 1—Mild, 2—Moderate, 3—Severe, 4—Proliferative DR). We binarized them as referable/non-referable DR, considering more than mild grades (grade > 1) as referable.




FIGURE 4
Source of the data used to train the DR classifier.

We detail the obtained metrics in [Section 3.3](#), while in [Section 4](#) we compare them with the results published in other works.

2.3.3.3 Gradability classification

We used the following public and private datasets to train, validate and test the gradability classification model:

1. *Private datasets*: We created a dataset consisting of more than 17,000 images from HUN (gradability dataset) to train and validate the Gradability Classifier. We labeled each image as gradable or non-gradable under close expert supervision.
2. *Public datasets*: To test the Gradability Classifier with external data, we used two public datasets that were totally excluded from training and validation: OIA-DDR ([16](#)) and EyeQ ([17](#)).

We trained and chose the best performing Gradability Classifier model using our private dataset. Our grading criteria were based on fundus visibility: a gradable fundus image should allow localization of small hemorrhages, especially in the macular area. We considered image focus, clarity, artifacts, macular visibility, and the gradable area, which had to reach 80% of the image ([28](#)). Our labellers were allowed to make brightness and contrast adjustments in order to ignore easily fixable image quality issues. Note that both the gradability and DR classifiers share the same

validation images, but the gradability training set is a subset of the DR training set. We respected these partitions to avoid any system-level bias.

Later, we evaluated the generalization capabilities of the Gradability Classifier using the OIA-DDR test set ([16](#)). OIA-DDR is a large public dataset introduced in [Section 2.3.3.2](#). This dataset consists of 346 non-gradable and 3,759 gradable images in which a DR severity grade has been assigned. Therefore, we used the non-gradable category to obtain a binarized dataset suitable for external validation.

Nevertheless, OIA-DDR is not specifically designed to assess retinal image quality. For this purpose, we used the Eye Quality Assessment Dataset (EyeQ) ([17](#)). In EyeQ, two experts labeled 28,792 retinal images from the EyePACS (Kaggle) ([48](#)) image set in three categories: *good*, *usable*, and *reject*. The *reject* category indicates that the image is not suitable for a reliable diagnosis. Therefore, we binarized the EyeQ test set based on this category, resulting in 3,215 non-gradable and 13,029 gradable images.

2.3.3.4 Gold standard

The Gold Standard is a private dataset we created to evaluate NaIA-RD as a black box prior to its deployment. It is labeled by expert ophthalmologists from HUN. We used it as a clinical

reference standard to evaluate the overall DR screening capabilities of the system, as well as two of its key components: DR and Gradability classifiers. We report the obtained metrics in [Section 3.2](#).

To create this dataset, we first agreed with HUN to measure an expected sensitivity of 80% with a confidence interval (CI) width of 10% and a confidence level of 95%. Given an estimated prevalence of referable eyes of 7%, we calculated a required dataset size of 1,265 eyes to measure the expected sensitivity ([50](#), [51](#)). However, we decided to reduce the required labeling effort by artificially increasing the prevalence (but maintaining the same statistical properties), resulting in a dataset size of 492 eyes with a prevalence of 18% referable eyes.

Therefore, we selected a sample of 492 eyes (retinographies, which we will refer to as *eyes* for simplicity) using anonymized clinical records (from March 2019 to August 2019), ensuring a prevalence of 18% of referable eyes. These eyes belonged to 205 different patients,¹¹ 66% males and 34% females (sex assigned at birth), with a mean age of 64.25 years (SD 14.87) at the time of the study. We excluded all these patients from all training and validation sets of NaIA-RD.

Three ophthalmologists from the hospital labeled each eye. Each expert provided a blind, independent label (non-referable, referable due to non-gradable fundus or referable due to DR), following the ICO guidelines and the ICDR scale as grading standard ([3](#), [31](#)). For each eye, the expert visualized all the fundus images taken by the nurse during the imaging session, along with an enhanced version created using the image enhancement technique detailed in [Section 2.3.6](#). After labeling was completed, we discarded eyes that had received three different votes (no consensus). 3 eyes were discarded by this procedure, resulting in a final Gold Standard dataset of 489 eyes.

We deployed NaIA-RD in a simulated production environment, obtaining its output for each eye of the dataset. Each eye was composed of the same real-world fundus images that the experts had labeled. Using this procedure, we compared the returned screening proposals with the expert labels and evaluated NaIA-RD in three different tasks: DR screening, DR classification, and gradability classification.

Task 1: DR screening. We evaluated NaIA-RD for binary DR screening (refer/not refer), without taking the motivation into account (DR or non-gradability). We compared the performance of NaIA-RD in three different ways using this data:

1. **Compared with the consensus of 3 ophthalmologists.** The main goal of the Gold Standard was to compare NaIA-RD

with the best possible clinical judgement, so we compared NaIA-RD's proposal per eye with the simple majority label of the experts (refer/not refer). As 3 ophthalmologists had graded all the eyes, no ties were possible.

2. **Compared with a single ophthalmologist.** We also wanted to compare NaIA-RD with an ophthalmologist performing the screening alone. To do this, we obtained the metrics of each Gold Standard labeler using the majority label of the other two ophthalmologists as the ground truth. We used NaIA-RD's outputs to break ties. Then, we compared NaIA-RD's metrics with those of each ophthalmologist.
3. **Compared with first-level screening GPs in real-world settings.** We compared both NaIA-RD proposals and GP decisions in the screening program with the majority label of the ophthalmologists. DR screening decisions had been made per patient, involving both eyes, so for this comparison we used the positive class (refer) if one of the eyes was considered referable (both the Gold Standard and NaIA-RD provided a label per eye). Due to missing data in clinical records, we had to use a subset of the Gold Standard for this comparison. The final dataset consisted of 122 screening decisions involving 244 eyes.

Task 2: DR classification. We evaluated NaIA-RD for referable/non-referable DR classification without considering gradability. Our goal was to test the DR Classifier model isolated from the Gradability Classifier model. Therefore, we first discarded eyes graded as non-gradable by simple majority (15 eyes), and we finally used the resulting majority label as the ground truth. We did not find any ties using this procedure. We binarized NaIA-RD outputs considering a positive class only if NaIA-RD outputted referable due to DR.

Task 3: Gradability classification. Analogously, we evaluated NaIA-RD for gradable/non-gradable classification without considering DR. Eyes with a non-gradable simple majority vote were taken as non-gradable (15 eyes), otherwise they were taken as gradable. We binarized NaIA-RD outputs considering a positive class only if NaIA-RD outputted referable due to non-gradability.

2.3.4 Calibration

Models whose scores are to be used for human decision making or automation should be calibrated ([52](#)). This is the case for NaIA-RD, where the output scores will be interpreted by clinicians and the HIS. Therefore, we calibrated both the DR and Gradability classifiers, while the Field Classifier did not need this feature. We have addressed model calibration as the next step after selecting the best model [this is called *post-hoc* calibration ([52](#))].

To train two separate DR and Gradability calibrators, we first tried using their respective training sets, but they performed poorly. So, we used their entire validation sets for this purpose. We chose the Beta Calibration ([53](#)) technique for the DR classifier, and Isotonic Regression ([54](#)) for the Gradability Classifier, based on systematic cross-validation experiments on the validation sets. The estimated calibration error, mean calibration error and Brier Score of the uncalibrated models were

¹¹When we randomly selected an eye for the Gold Standard, we added the corresponding fellow eye (if present) to complete an entire study. We will use the term study to refer to an object formed by fundus images of two fellow eyes. Following this procedure, some patients were represented in the Gold Standard with multiple studies due to the random nature of eye selection, but we ensured that the same study was not included twice.

0.017248, 0.299778, 0.023504 (DR) and 0.114798, 0.278866, 0.063906 (Gradability), respectively (mean values), and the calibrated models obtained mean values of 0.005958, 0.155110, 0.021976 (DR) and 0.010228, 0.116718, 0.045477 (Gradability), respectively. More details along with calibration curves are included in the [Supplementary Material](#).

Unfortunately, presenting two separate calibrated probabilities in a screening proposal may be difficult to interpret. It would be more convenient to combine the DR and gradability scores into a single number. Additionally, the DR and gradability scores should be self-explanatory, meaning it should not be necessary to know the decision threshold to understand them. For example, it would be counterintuitive to suggest a referral for possible DR with only a 10% probability (where the DR referral threshold is set to 0.1).

Therefore, the scores returned by NaIA-RD are not the calibrated probabilities directly. Instead, given a new decision boundary t' , NaIA-RD returns a transformed referral score s , where $s < t'$ is given for non-referral proposals, and $s \geq t'$ for referrals. We used $t' = 0.5$ in order to make the screening score more intuitive. Given a probability p , a threshold t and a decision boundary t' , we have defined a transformation function $f(p, t, t')$ as follows:

$$f(p, t, t') = \begin{cases} t' \cdot \left(1 + \frac{p-t}{1-t}\right), & \text{if } p > t \\ t' \cdot \left(1 + \frac{p-t}{t}\right), & \text{otherwise} \end{cases}$$

This transformation has the property of preserving the order of the input probabilities p , applying a different linear transformations for each case. After transforming the DR and gradability scores, NaIA-RD generates a single DR screening score, using the highest transformed value from both classifiers in its proposal.

2.3.5 Interpretability

The DR classifier generates heatmaps for positive predictions to improve interpretability. We understand model interpretability as the intuitive mapping between inputs and outputs, as described in (55). Therefore, we used a technique called Integrated Gradients (56) to be able to highlight small lesions on the input image. This technique provides an accurate mapping of input image pixels without modifying the original neural network by approximating the integral of the gradients of neural activations. Integrated Gradients was successfully used for DR screening interpretability in (57). The authors concluded that heatmaps can increase the confidence and accuracy of human graders, but also their grading time, so heatmaps should only be used for positive DR predictions.

NaIA-RD uses this attribution technique to provide interpretability: First, the Integrated Gradients algorithm is applied to obtain an attribution mask (Gauss-Legendre integral approximation method is executed for 20 forward steps using the `captum` library¹²). Then, the attribution mask pixels are clustered using the OPTICS algorithm (58) (`sklearn`'s OPTICS class is

used, with an *epsilon* maximum distance between two cluster points of 23 and a *min_samples* minimum cluster size of 4). Finally, NaIA-RD returns the center coordinates and radii of the clusters, which the HIS overlays on the image as standard DICOM circumference annotations. These circumferences usually highlight DR signs such as hemorrhages. In [Figure 5](#) we show an example that illustrates this process.

2.3.6 Image enhancement

In addition to the DR screening proposal, NaIA-RD provides an enhanced image of the eye fundus to ease human interpretation. In particular, it returns two enhanced central and nasal images per request, and the HIS stores them in the DICOM study. In this way, clinicians can benefit from image enhancement using any DICOM viewer.

In [Figure 6](#) we show an example of how NaIA-RD enhances a challenging fundus image. The enhanced image is computed in three steps:

1. The eye fundus is cropped from the original image.
2. The dynamic range of the image is linearly extended after clipping intensity values below 1% and above 99% percentiles, respectively.
3. Contrast is further increased applying Contrast Limited Adaptive Histogram Equalization (CLAHE) (59, 60) in each RGB channel. We use `opencv`'s `createCLAHE` function for this last computation.

We have found that extending the dynamic range (step 2) before applying CLAHE (step 3) gives better results than applying CLAHE first.

2.3.7 MLOps

NaIA-RD was developed in Python, with each component providing a REST API and running in its own Docker container (see [Figure 3](#)). All code was written in Jupyter notebooks, using the `nbdev`¹³ library to implement a literate programming paradigm (61). We extensively unit tested all code and components, and a sanity check job was run periodically to notify if any significant deviation was detected in the last month's data.

3 Results

In this section we evaluate NaIA-RD using the datasets we introduced in [Section 2.3.3](#). The performance of NaIA-RD is compared with that of experts, and its ability to generalize is assessed. These results were important to HUN, as the decision to deploy NaIA-RD was based on them.

Additionally, this section presents the before-and-after study we conducted. This study compares the screening decisions made at HUN before (retrospectively) and after (prospectively) the

¹²Captum library: https://captum.ai/api/integrated_gradients.html

¹³nbdev library: <https://nbdev.fast.ai/>




FIGURE 5

Annotation process of NaIA-RD using Integrated Gradients, as a mechanism of increasing interpretability. NaIA-RD provides pixel attributions as circumference coordinates, which the HIS stores in the DICOM object as standard DICOM annotations. (a) Original image. (b) Pixel attribution. (c) Annotations.




FIGURE 6

Example of fundus image enhancement by NaIA-RD. Notice how two hemorrhages near the macula are more visible in the enhanced image. Another smaller bleedings are also more evident in the inferior arcade. (a) Original. (b) Enhanced.

deployment of NaIA-RD, measuring its real-world impact on clinicians and patients.

When appropriate, we report the area under the ROC curve (AUC), the sensitivity and the specificity, as similar works do (18, 22, 26, 62–64). We also use the Cohen kappa score to measure agreement (65). All confidence intervals are calculated with a 95% of confidence level using bootstrapping (66).

3.1 Field classification

The Field Classifier obtains a weighted multi-class AUC of 99.62% on the test set (20,858 fundus images that were not used for optimization nor model selection).

3.2 Gold standard

As described in Section 2.3.3.4, we used our private Gold Standard to evaluate NaIA-RD on three different tasks: DR screening, DR classification, and gradability classification. Table 3 shows the overall results of NaIA-RD on these tasks, while Tables 4, 5 provide a deeper comparison between NaIA-RD and individual ophthalmologists as well as first-level screening GPs. The most relevant results for each task are summarized below.

Task 1: DR screening. The following comparisons are performed:

1. **Compared with the consensus of 3 ophthalmologists.** According to the first row of Table 3, NaIA-RD achieves a

TABLE 3 NaIA-RD metrics on different tasks using our private Gold Standard. Ground truth labels are obtained as the majority label of three ophthalmologists. For the DR classification task, non-gradable images are excluded. All confidence intervals are calculated with a 95% of confidence level using bootstrapping (66).

ID	Task	Target question	Ground truth positives	Kappa	AUC	Sensitivity	Specificity
1	DR screening	Is referable to the ophthalmologist?	27.61% 135/489	0.818 0.764–0.869	0.979 0.969–0.989	92.5% 88.1–96.3	92.4% 89.5–94.9
2	DR classification	Are there observable signs of more than mild DR?	25.3% 120/474	0.852 0.801–0.901	0.986 0.977–0.993	91.6% 86.6–95.8	95.2% 92.9–97.2
3	Gradability classification	Is non-gradable for DR screening?	3% 15/489	0.423 0.229–0.602	0.754 0.684–0.822	53.0% 26.6–80.0	97.2% 95.5–98.5

TABLE 4 DR screening performance comparison per ophthalmologist vs. the majority label of the other two ophthalmologists of the Gold Standard and NaIA-RD. All confidence intervals are calculated with a 95% of confidence level using bootstrapping (66).

Reader	Ground truth positives	Kappa	Sensitivity	Specificity
Ophthalmologist 1	28.66% 141/489	0.835 0.781–0.89	89.4% 84.2–94.3	94.7% 92.2–96.8
Ophthalmologist 2	30.47% 149/489	0.805 0.753–0.861	78.4% 71.4–84.6	98.2% 96.8–99.4
Ophthalmologist 3	27.81% 136/489	0.789 0.736–0.844	93.3% 88.9–97.1	90.0% 86.9–92.3
NaIA-RD	27.61% 135/489	0.818 0.764–0.869	92.5% 88.1–96.3	92.4% 89.5–94.9

TABLE 5 First-level screening general practitioners (GPs) and NaIA-RD metrics on a subset of the Gold Standard. Real-world screening decisions from four different GPs are used as labels. Each screening decision involves two eyes (right and left eye). All confidence intervals are calculated with a 95% of confidence level using bootstrapping (66).

Reader	Ground truth positives	Kappa	Sensitivity	Specificity
GPs	24.59% 30/122	0.432 0.241–0.626	43.2% 26.6–63.3	94.4% 89.1–98.9
NaIA-RD	24.59% 30/122	0.794 0.674–0.913	93.6% 83.3–100.0	91.3% 85.3–96.7

TABLE 6 NaIA-RD metrics on external datasets assessing DR classification.

Dataset	Reference	Test partition	Size	Relabeled	AUC	Kappa
EyePACS (Kaggle)	(12, 48, 49)	Public test set	10,906	No	0.964	0.818
APTOS (Kaggle)	(13)	50% random	1,831	Yes ^a	0.999	0.962
Messidor-2	(14, 47)	50% random	872	No	0.988	0.877
IDRiD	(15)	Public test set	103	No	0.991	0.879
OIA-DDR	(16)	Public test set without non-gradable images	3,759	No	0.957	0.760

^aWe relabeled 165 images (9%) from *mild* to *moderate* DR using the ICDR (31) scale.

sensitivity and specificity greater than 92%, with a Cohen kappa of 0.81 [strong agreement (65)].

2. **Compared with a single ophthalmologist.** In Table 4 we can observe that NaIA-RD is the only grader with sensitivity and specificity above 91%, and its Cohen kappa score (0.818) is higher than the kappa score of the two thirds of the ophthalmologists.
3. **Compared with first-level screening GPs in real-world settings.** Table 5 shows a significantly lower sensitivity with a much wider confidence interval for GPs (26.6%–63.3%) than for NaIA-RD (83.3%–100%). Cohen kappa scores show a weak agreement (0.432) for GPs, while NaIA-RD shows a moderate-strong agreement (0.794).

Task 2: DR classification. The second row of Table 3 shows that the performance of NaIA-RD on this task is superior to its performance on Task 1: DR screening (the obtained AUCs are 0.986 and 0.979, respectively).

Task 3: Gradability classification. The third row of Table 3 shows a gradability classification specificity of 97.2% (95.5–98.5). However, only 15 eyes (3%) were graded by the ophthalmologists as non-gradable: this resulted in a wide confidence interval for the sensitivity measurement (26.6–80.0), revealing a limitation of this dataset.

3.3 External validation

3.3.1 DR classification

Table 6 shows the results of the DR Classifier over the test partitions of several public datasets. The obtained AUCs range from 0.957 to 0.999. Note that the public test sets of EyePACS, IDRiD, and OIA-DDR can be used directly for comparison with previous works. In the discussion section (Section 4.1), we make this comparison. However, recall from Section 2.3.3.2 that some EyePACS and IDRiD partitions were used for training, so these metrics should be used with caution.

On the contrary, OIA-DDR was totally excluded from training. In this dataset, NaIA-RD obtained an AUC of 0.957, a Cohen kappa of 0.76 (moderate agreement), and a sensitivity and specificity of 93% and 84%, respectively.

3.3.2 Gradability classification

On the OIA-DDR test set, the Gradability Classifier obtained an AUC of 0.928, a Cohen kappa of 0.213 (minimal agreement), and sensitivity/specificity of 100% and 62%, respectively. When assessing the EyeQ test set, the Gradability Classifier obtained an AUC of 0.938, Cohen kappa of 0.63 (moderate agreement), and sensitivity/specificity of 87% and 86%, respectively.

3.4 Before-and-after study

Four trained GPs worked as the first DR screening level of HUN from March 2015 to December 2023. They visualized NaIA-RD's proposals since the 1st of July 2020, when NaIA-RD was deployed. We have compared their screening decisions with the screening proposals of NaIA-RD, both before (retrospectively) and after (prospectively) they started using NaIA-RD. Clinical decisions and NaIA-RD grades were available as electronic health records. For the period prior to the deployment of NaIA-RD, we obtained the corresponding AI grades by calling NaIA-RD's API.

The mean patient age at the time of the eye study was 66.70 years (SD 12.65) before NaIA-RD was deployed, with 60.67% males and 39.32% females (sex assigned at birth). After NaIA-RD was implemented, the demographics remained very similar, with a mean age of 66.84 years (SD 13.33) and 60.48% males and 39.51% females.

The volume of data that supports this study is illustrated in [Figure 7](#), with a histogram of the number of screened studies¹⁴ per year. This number has increased since the start of the screening program in 2015, with the exception of 2023, when the number of screened patients slightly decreased. A NaIA-RD grade is available for the majority of the screened studies (81% before deployment and more than 99% after), but some NaIA-RD grades are missing due to image retrieval errors (3,685 grades in total).

Using this data, we have examined the evolution of the DR screening program over time. We have analyzed the following aspects.

3.4.1 Referral decisions vs NaIA-RD proposals

To examine the relationship between decisions and AI grades in terms of volume, in [Figure 8a](#) we compare the percentage of studies that GPs referred to the second screening level with the percentage that NaIA-RD graded as referable. Prior to its deployment in 2020, NaIA-RD would have proposed to refer slightly more studies than those referred by GPs (median¹⁵ difference of 2.74%). However, this proportion increases in 2020 and afterwards (median difference between proposals and actual referrals of 9.16%). We observe an anomaly in 2020 and 2021, when the proportion of referral proposals by NaIA-RD reaches 34.1% and 35%, respectively. Importantly, this anomaly only affects NaIA-RD: it did not lead to an increase in referral decisions by GPs, even though they were supported by NaIA-RD after July 2020. The anomaly ends in 2022, when the proportion of NaIA-RD referral proposals decreases to 24.6% (2017 level), and reaches a minimum in 2023 (18% studies). GPs' decisions also show this trend, as they refer fewer studies in 2022 and 2023 than ever before.

3.4.2 Non-gradable vs. referable DR

To try to explain the 2020-2021 anomaly, we have used the outputs of NaIA-RD to compare the evolution of the disease and the image gradability. [Figure 8b](#) shows the percentage of referable studies due to possible DR and the percentage of referable studies due to non-gradability. While the proportion of possible DR studies has remained quite constant over time (between 11.6% and 17.2%), the proportion of non-gradable studies increases notably in 2020 and 2021 (from 12.2% in 2019 to 20.9% in 2021), while this proportion drops to historical minimums in 2023 (7%).

3.4.3 Patients requiring on-site eye examination

To analyze the detection capability of the DR screening program, [Figure 8c](#) shows the annual proportion of screened studies that resulted in an on-site eye examination. These eye exams were appointed by the second screening level, if the patient needed it, after the first-level screening GP (or the nurse due to high intraocular pressure) had referred the study. Patients left the DR screening program when this appointment occurred (see [Section 2.1](#) for more details). Note that the appointment proportion is calculated relative to the total number of studies screened.

We observe that the appointment proportion decreases from 2015 to 2019, reaching a minimum of 2.58%. However, coinciding with the deployment of NaIA-RD, this trend was broken in 2020, with a peak of 9.29% in 2021, consistent with the observed non-gradability anomaly. In 2022 and 2023, despite the fact that the first screening level referred fewer studies than ever before (see [Figure 8a](#)), the second screening level appointed more patients for eye examinations than in 2018-2019, achieving eye examination rates similar to the early years of the screening

¹⁴We will define a study as an object consisting of digital retinographies, where each eye is represented by multiple fundus images. NaIA-RD grades each eye separately, and the global NaIA-RD screening proposal is referable if either eye is referable. Both the first and second screening levels made their decisions after visualizing both eyes (the entire study). See [Section 2.1](#) for more details.

¹⁵We used the median instead of the mean to minimize the influence of outliers.




FIGURE 7
Graded study quantity comparison: GPs vs NaIA-RD.

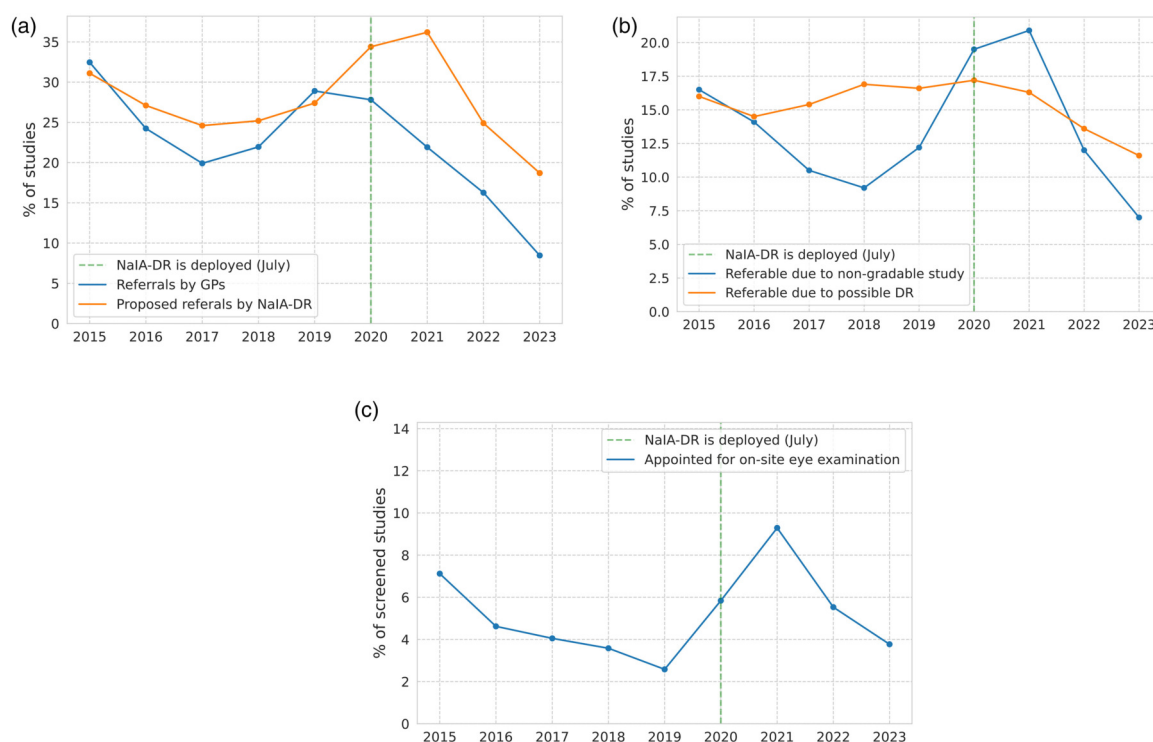


FIGURE 8
Evolution of the DR screening program of the HUN over time. (a) Referral decisions vs. NaIA-RD proposals. (b) Non-gradable vs. referable DR (NaIA-RD outputs). (c) Percentage of patients scheduled for on-site eye exam.

program. Specifically, the mean on-site eye examination proportion increased 1.5 times (from 3.08% to 4.65%), while 13.3% fewer patients were referred (mean referral rate) in 2022–2023 compared to 2018–2019.

3.4.4 Agreement between GPs and NaIA-RD

To analyze the influence of NaIA-RD on clinical decisions, Figure 9a shows the Cohen's kappa between first-level screening decisions and NaIA-RD outputs, a metric that measures the level of agreement between both. Note how the level of agreement

more than doubles after the system was deployed (going from a kappa score of 0.2 in 2019 to a kappa score of 0.48 in 2023). To delve deeper, Figure 9b shows the kappa score per first-level screening GP. We observe that some GPs had different screening criteria compared with others—they screened random studies from the same patient population, so we attribute these differences to screening criteria. Criteria differences occur particularly between two groups of GPs: 1, 4 and 2, 3. Interestingly, almost all GPs (except GP 4, who left the screening program in 2021 with few screened studies) clearly increased




FIGURE 9

Level of agreement between first-level screening general practitioners (GPs) and NaIA-RD's grades. (a) Level of agreement between NaIA-RD and all GPs. (b) Level of agreement between NaIA-RD and each GP.

TABLE 7 NaIA-RD metrics compared to appointed on-site eye examinations per general practitioner (GP). All GPs took their decision after viewing NaIA-RD's proposal, once NaIA-RD was in use.

GP	PA ^a	NA ^b	Kappa	Studies	Referred	On-site eye exams ^c
1	0.234	0.999	0.307	10,543	6.3%	3.8%
2	0.612	0.980	0.664	6,902	16.7%	8.8%
3	0.866	0.946	0.812	5,089	29.2%	13.3%
4	0.291	0.985	0.322	428	11.7%	6.1%

^aPA, Positive agreement. NaIA-RD proposes referral and the GP refers.

^bNA, Negative agreement. NaIA-RD proposes non-referral and the GP does not refer.

^cOn-site eye exams: Proportion of screened studies that resulted in an on-site eye examination appointment by the second screening level. A higher value represents a higher sensitivity of the GP for detecting target patients.

their agreement with NaIA-RD's proposals after its deployment. We would like to highlight the behavior of GP 1, who changed from an increasingly divergent agreement (reaching a minimum kappa score of 0.05 in 2019) to a much higher and constant level of agreement (kappa scores between 0.20–0.32) since the deployment of NaIA-RD. Anyway, GP 1's agreement was still minimal after this change.

3.4.5 NaIA-RD metrics compared to appointed on-site eye examinations per GP

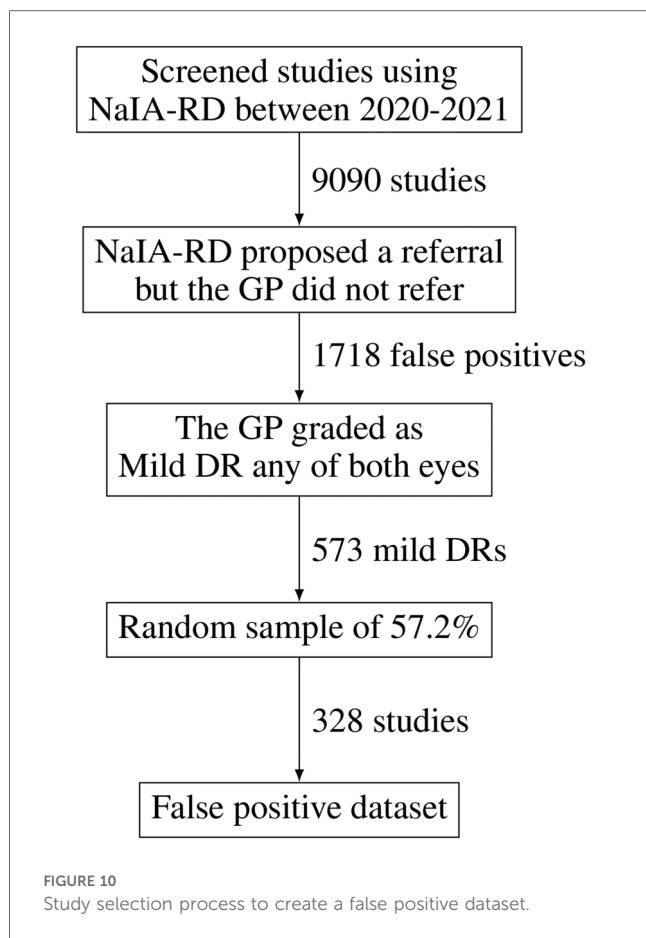
To further explore the impact of NaIA-RD, in Table 7 we compare NaIA-RD's post-deployment metrics (2020–2023) with the proportion of referrals and appointed on-site eye examinations for each first-level screening GP. We obtained NaIA-RD's positive agreement (PA), negative agreement (NA) and Cohen's kappa for each GP. PA and NA measure the proportion of NaIA-RD's referral and non-referral proposals that are agreed by the GP, respectively. Each row of this table takes the corresponding GP's decisions as the ground truth. The last three columns show the number of screened studies by each GP and two proportions: the rate of referred studies and the rate of studies that resulted in an on-site eye examination appointment. Both proportions are calculated relative to the total number of studies screened by each GP.

Data show that GPs who agreed more with NaIA-RD (Cohen's kappa) identified a higher proportion of patients that were

appointed for an on-site eye examination. For example, GP 3, with a kappa of 0.812 and an appointment rate of 13.3%, more than triples the appointment rate of GP 1, who shows a kappa of 0.307 and an appointment rate of 3.8%. However, with more than 10,543 screened studies, GP 1 has screened almost as many studies as the rest of the GPs combined, which explains the weak kappa values shown in Figure 9a. Finally, we observe high NAs for all GPs (0.946 or more), indicating that most of the NaIA-RD's non-referral proposals are followed. On the other hand, compared with GPs 2 and 3, the low PA of GPs 1 and 4 suggest a less sensitive referral criterion. This has resulted in lower referral rates, but also lower rates of patients appointed for on-site eye examinations. The global PA and NA are 0.493 and 0.981, respectively, with a global Cohen's kappa of 0.554.

3.4.6 Error analysis

Due to the observed low PA, an expert ophthalmologist blindly labeled a sample of false positives of the 2020–2021 period, creating what we call a false positive dataset. In Figure 10 we detail the study selection process we followed to create it. We wanted to review the most difficult and meaningful studies (potentially mild cases), so we selected studies that had been graded by the GP as mild non-proliferative DR. The ophthalmologist classified 93% of the studies as referable (306 out of 328), matching NaIA-RD's proposals. Thus, GPs correctly contradicted NaIA-RD's referral proposals in only 7% of the studies.



After reviewing a sample of NaIA-RD's false positives, we analyzed all false negatives from 2020–2023 period, which were only 180 studies (less than 1% of the 22,962 total screened)—studies that NaIA-RD did not propose to refer but the GP referred. The second screening level had graded these false negatives according to the ICDR scale (31), with the following results: 83.33% (150) were classified as “no apparent retinopathy,” 6.67% (12) as “mild non-proliferative DR,” 1.11% (2) as “moderate non-proliferative DR,” 0.56% (1) as “severe non-proliferative DR,” 0% (0) as “proliferative DR,” and 10.66% (29) as “not gradable.” An ophthalmologist reviewed the one study that was graded as severe and could not find any DR signs on the images. Therefore, we attributed the grade to a data error.

3.4.7 NaIA-RD as first screening level

To assess whether NaIA-RD could reduce current human supervision, we analyzed what would have happened if it had performed the first level of screening autonomously—without GP supervision—since 2020. During this period, 22,962 studies were screened in HUN, and 14.62% were referred by GPs assisted by NaIA-RD. NaIA-RD would have referred 26.85% of these, 1.84 times the current referred study quantity. However, the studies that NaIA-RD deemed non-referable would not have required human grading, which were the 73.15%, and the referable studies would have been visualized only once (and not twice as in the current setup). Overall, it would have allowed a 4.27 times

workload reduction (6,125 study visualizations instead of 26,318 current visualizations).

4 Discussion

We have found that NaIA-RD is a sensitive DR screening tool that has had a positive influence on the trained GPs who have screened patients from March 2015 to December 2023 at HUN. Our before-and-after study shows that these GPs are more capable of identifying patients who require on-site eye examinations since they started using NaIA-RD (see Figure 8c). In fact, GPs were unable to identify a single case of sight-threatening DR missed by NaIA-RD since its deployment. This suggests that NaIA-RD could be safely and effectively used for autonomous, first-level DR screening at HUN.

In this section we will discuss the performance of NaIA-RD in laboratory settings (Section 4.1), its impact on the screening program of HUN (Section 4.2), the convenience of buying or developing such a system (Section 4.3), and the limitations of this work (Section 4.4).

4.1 NaIA-RD's performance

With a Cohen's kappa of 0.818, NaIA-RD shows a strong agreement with the majority label of three ophthalmologists according to our private Gold Standard, a metric comparable to the performance of a single ophthalmologist. A subset of this dataset also shows (retrospectively) that NaIA-RD is clearly more sensitive than a trained GP working in real world settings.

In terms of Gold Standard sensitivity and specificity, NaIA-RD achieved values of 92.5% (95% CI, 88.1–96.3%) and 92.4% (95% CI, 89.5–94.9%), respectively. These results clearly exceed the superiority endpoints of 85.0% and 82.5% proposed by Abramof et al. in (18) for a DR screening medical device, as discussed by the authors with the FDA.

The goal of NaIA-RD is to improve the clinical pathway for DR screening at HUN. Therefore, it is not expected to be the most accurate DR screening tool that could be used in any other hospital. However, we have tested it on several public datasets dedicated to DR grading and retinal image quality assessment. Results suggest that NaIA-RD correctly assesses retinal images from diverse data distributions and patient populations.

With respect to DR assessment, we have tested NaIA-RD on EyePACS (Kaggle) (12, 48), APTOS (Kaggle) (13), Messidor-2 (14), and IDRiD (15), obtaining AUCs above 0.96 in all of them. These results do not seem to be far from the state of the art: For example, using an ensemble of five models trained on the Kaggle EyePACS training set, Papadopoulos et al. (67) achieve an AUC of 0.961 on the public test set of Kaggle EyePACS, and 0.976 on Messidor-2. Trying to replicate the well-known work of Gulshan et al. (68) using only public data, Voets et al. (69) obtained an AUC of 0.951 on the Kaggle EyePACS public test set and 0.853 on Messidor-2 with an ensemble of ten Inception-V3 models. Voets et al. argued that they struggled to reproduce the original

algorithm due to differences in the datasets used. Note that NaIA-RD was trained and tested on a subset of these datasets, so these metrics should be compared with caution. To test the generalization capabilities of the DR Classifier, we used OIA-DDR, which was completely excluded from training. An AUC of 0.957 and a sensitivity and specificity of 93% and 84% were obtained, but we could not find any other work with comparable results.

Regarding gradability assessment, we evaluated NaIA-RD's generalization capabilities using OIA-DDR (16) and EyeQ (17). On OIA-DDR, the Gradability Classifier achieved perfect sensitivity, correctly identifying all non-gradable images, albeit with a 38% false positive rate. However, it demonstrated higher specificity on EyeQ, with sensitivity and specificity values of 87% and 86%, respectively. Recall that EyeQ, unlike OIA-DDR, is a dataset specifically designed to assess retinal image quality.

While we explored alternative retinal image quality datasets like DeepDRiD (70) and DRIMDB (71), they proved unsuitable for our purposes. DeepDRiD offered a potentially relevant overall quality class but suffered from inconsistent labeling, an issue acknowledged by its authors (70). DRIMDB, on the other hand, provided cropped retinal images, incompatible with NaIA-RD.

These results show that NaIA-RD is more sensitive than specific for both DR and gradability assessment. This behavior ensures safety in a tool intended to be used as a first-level screening device. However, we also value its performance as competitive: in a recent multicenter study (8), the best performing AI device had 80.47% and 81.28% of sensitivity and specificity on an external dataset (a performance comparable to human graders).

4.2 NaIA-RD's impact

We have presented a before-and-after study showing the impact of NaIA-RD on a real screening program. Results show that NaIA-RD has influenced the decisions of the first-level screening GPs, biasing them towards system's proposals. We interpret this bias as desirable, as GPs with a higher level of agreement with NaIA-RD were better at identifying patients who needed an on-site eye examination by a specialist (they were more sensitive). In fact, the rate of on-site eye examinations increased since the use of NaIA-RD, breaking a downward trend in the pre-NaIA-RD period (2015–2019 years in Figure 8c).

Nevertheless, prospective results show that some first-level screening GPs such as 1 and 4 (for whom NaIA-RD shows low PAs in Table 7) could have improved their sensitivity by more frequently following NaIA-RD's referral proposals. These GPs minimized the amount of studies that the second screening level received, but they were less able to identify patients who should have been referred compared with other GPs. The impact of these decisions is relevant, as GP 1—who only followed the 23% of referral proposals (PA 0.234)—screened almost as many studies as the rest of the GPs combined. We discarded a malfunction of NaIA-RD with a false positive analysis dataset, in which an expert ophthalmologist labeled a sample of the

contradicted referral proposals: the expert agreed with NaIA-RD in 93% of the studies. Given that NaIA-RD was correct most of the time, we hypothesize that the observed variability in GP behavior may stem from a lack of standardized DR screening protocols and varying levels of trust in NaIA-RD and AI-based tools.

On the other hand, NaIA-RD's non-referral proposals were consistently accepted by GPs (high NAs on Table 7), with only 180 non-referral proposals (<1% of the total screened) not followed. Among these, only 2 undetected moderate DR cases were found. NaIA-RD has not missed any worse DR case since its deployment (at least cases that GPs did refer). This high sensitivity raises the question of whether total human supervision is necessary. It appears that NaIA-RD could safely perform the first level of screening autonomously, similar to other commercial devices such as IDx-DR, EyeArt, Retmarker (26) or SELENA+ (39). In this scenario, NaIA-RD would refer non-gradable and more than mild DR studies directly to the second screening level, reducing the study visualization workload by 4.27 times.

The before-and-after study also shows how GPs lowered their image quality requirements when there was an abnormal proportion of non-gradable fundus images in 2020 and 2021: they tried to find DR signs even in images that they had previously considered as non-gradable. We interpret this flexible behavior as a normal human adaptation to new circumstances (possibly related to COVID-19 pandemic), which avoided an excess of work in the second screening level. We believe that this illustrates how subjective fundus gradability assessment is, and how interesting a dynamic gradability adjustment feature may be for a DR screening algorithm. Also, this finding seems to justify the need for grading DR even when the image quality is poor, a feature that very few commercial devices offer.

4.3 Buy or develop?

In Europe, any medical device is subject to the EU 2017/745 regulation¹⁶ and the specific laws of the country where it is developed (in Spain it is regulated by the Royal Decree 192/2023¹⁷). Any AI-based software that makes clinical recommendations must be approved by a Notified Body before it can be used for any purpose other than research. There are two ways to obtain this authorization: a CE mark, which allows the device to be marketed, or an in-house authorization, which is restricted to local use. The latter option is a valid regulatory pathway in Europe under certain circumstances. According to EU 2017/745, this allows healthcare institutions to manufacture

¹⁶EU 2017/745: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32017R0745&from=ES#d1e1058-1-1>

¹⁷Royal Decree 192/2023: <https://www.boe.es/boe/dias/2023/03/22/pdfs/BOE-A-2023-7416.pdf>

medical devices for internal use if no equivalent device is commercially available.

We recommend considering a custom development if the healthcare institution has the required resources and does not find a commercial device that satisfies their requirements. However, the complexity of such development should not be underestimated: the access to target domain knowledge, data, software and IT infrastructure are essential. Dedicated AI and software development teams are needed, who must work closely with clinicians and the hospital. The regulatory work is also unavoidable, even in the case of an in-house authorization.

4.4 Limitations

This work has some important limitations and knowledge gaps that need to be considered.

First, we have measured NaIA-RD's performance on the target data distribution only at eye level, not at image level. We could not measure how performance differed between macula-centered and optic disc-centered fundus images. The private Gold Standard does not provide a label per fundus image, but rather a label per eye. DR screening is performed at patient level, so the before-and-after study does not provide this separate measurement either.

Second, the comparison between NaIA-RD and a single ophthalmologist in the Gold Standard may be biased. We only had three ophthalmologists available for labeling, so when we compared each ophthalmologist's label to the other two, we used the grade of NaIA-RD to break ties.

Third, we could not test the performance of the Gradability Classifier using an adequate private dataset. Due to the low prevalence of bad quality images, the private Gold Standard is not appropriate for measuring the sensitivity of the eye fundus Gradability Classifier. Therefore, we tested it using external datasets.

Fourth, we did not test NaIA-RD's performance in population subgroups, nor did we perform a systematic bias analysis. Although the overall results are promising, such an analysis should be performed to discard any bias due to imaging site, sex, gender, age or pre-existing patient conditions.

Fifth, with respect to the generalization ability of NaIA-RD, it should be noted that the DR Classifier was trained on a subset of the most popular public datasets (see [Figure 4](#)). Therefore, its metrics may be biased when evaluating these datasets. On the other hand, the metrics we obtained evaluating OIA-DDR and EyeQ (which were not used in training) are promising, but further testing is needed to assess the generalization capabilities of NaIA-RD's models.

Sixth, our comparison of commercial devices is based solely on public information available prior to 2024—we did not contact any vendors. The lack of a detailed public database of CE marked devices in Europe did not facilitate our research ([72](#)). We therefore advise against using our product overview as a purchasing guide.

Seventh, due to the lack of a control group in our before-and-after study, we cannot conclusively attribute the observed

improvements in the DR screening program solely to NaIA-RD. It is possible that other factors influenced the change in GP behavior. A different type of clinical study, such as a randomized controlled trial, may have provided stronger evidence.

Eighth, we did not compare NaIA-RD with other commercial devices in terms of accuracy. This comparison was beyond the scope of this work: our goal was to provide a capable (not the most accurate) AI tool for DR screening at HUN (not other hospitals) that met the requirements that commercial devices did not.

Ninth, due to copyright and privacy concerns, this work does not publish any code or dataset. Although this fact limits its reproducibility, we believe that other researchers can use this work to develop their own in-house AI medical device—given that they have their own data.

This work also leaves some important questions unanswered. First, we weren't able to assess patient benefit: Did patients have better DR outcomes after NaIA-RD was in use? A randomized controlled trial seems necessary to answer this key question. In addition, we could not clarify why there are some first-level screening GPs who contradict the system much more than others.

5 Conclusion

We have presented an AI-powered medical device, called NaIA-RD, customized to the needs of a hospital (HUN). We have detailed the entire development process with a level of detail that should allow other researchers to follow our approach using their own data. We have proposed a novel system design that combines DR grading and retinal image quality in a single, flexible AI device. We have measured its performance on private and public datasets, achieving competitive metrics, and assessed its real-world impact with a before-and-after study that is unprecedented in the literature.

We have found that NaIA-RD increased the sensitivity of the GPs who conformed the first screening level: these clinicians were influenced by the AI tool, increasing the proportion of patients who were scheduled for eye examination by the second screening level—which were the target patients to be detected. However, GPs showed a very heterogeneous screening criteria, and NaIA-RD did not influence them equally.

We have also observed that GPs adapted their DR screening decisions to seasonal anomalies, such as a sudden deterioration in image quality. When this occurred, GPs lowered their standards for fundus gradability. This clinical adaptability implies that there is a friction between best practices, which are strictly followed by AI tools, and their real-world applicability. To address this issue, it appears that DR screening tools should be adjustable to local needs. Calibrated decision thresholds should help achieving this to some extent, as long as AI tools separate disease from image quality assessment.

In spite of their lack of context awareness, we conclude that AI tools can improve and homogenize DR screening, while reducing the burden of manual grading: NaIA-RD could safely be used to perform first level screening autonomously, reducing the

workload by up to 4.27 times, at the expense of receiving more studies in the second screening level (1.84 times more). Therefore, our future work will be directed towards reducing the level of supervision of NaIA-RD, accompanied by new features to facilitate the task of the second screening level. In our opinion, an AI-generated ICDR (31) grade should be useful to automate clinical reports in a supervised manner.

We believe NaIA-RD exemplifies how an AI medical device can drive long-term clinical improvements. Its seamless integration into the clinical workflow appears to be a key factor in its success—potentially as important as its accuracy. Healthcare institutions should be aware that commercial AI tools are typically designed to work with specific patient populations, cameras, and imaging protocols, offering limited options for adaptation and customization. This can lead to challenges when incorporating them into existing clinical workflows. In such cases, developing an in-house solution may be a good alternative, provided the institution has the necessary resources.

Data availability statement

The raw data supporting the conclusions of this article are available from the authors upon reasonable request and subject to approval by the competent healthcare authority in accordance with applicable legal and ethical regulations.

Ethics statement

This work was approved by the Health Care Ethics Committee of the Hospital University of Navarre. All clinical records and DICOM studies were anonymized before being used, and no individual patient clinical history was accessed at any time. Therefore, no patient consent was needed for this work.

Author contributions

IP: Conceptualization, Data curation, Formal analysis, Investigation, Software, Writing – original draft. AO: Conceptualization, Data curation, Formal analysis, Investigation, Software, Writing – review & editing. DJ: Conceptualization, Data curation, Formal analysis, Investigation, Software, Writing – review & editing. BD: Conceptualization, Formal analysis, Investigation, Writing – review & editing. MS: Writing – review & editing. AO: Investigation, Project administration, Writing – review & editing. JB: Project administration, Writing – review & editing. JG: Resources, Supervision, Writing – review & editing. MG: Conceptualization, Methodology, Supervision, Writing – review & editing. JA: Conceptualization, Data curation, Resources, Supervision, Validation, Writing – review & editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. NaIA-RD has been fully funded by the Government of Navarre.

Acknowledgments

NaIA-RD would not exist without the invaluable work of many people. From the University Hospital of Navarre (HUN) and the Navarre Public Health Service (SNS-O), we would like to express our special thanks to Javier Turumbay, Elena Manso, Alejandro Dávila, Marcos Mozo, Fermín Bruque, María Jesús Esparza and Begoña Martínez. From the General Directorate of Telecommunications and Digitalization (DGTD), we would like to highlight the work of Jokin Sanz, Adrián Errea, Carlos Romero, Ainhoa Pagola, the Electronic Health System Team (HCI), the Digital Imaging Team, the Infrastructure Team and the Architecture Group (GAT).

Conflict of interest

IP, AO, DJ, BD, MS, AO, JB and JG were employed by the Government of Navarre.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fdgth.2025.1547045/full#supplementary-material>

SUPPLEMENTARY VIDEO 1

Summary of the manuscript: Improving diabetic retinopathy screening using artificial intelligence.

References

- Fong DS, Aiello L, Gardner TW, King GL, Blankenship G, Cavigliaro JD, et al. Diabetic retinopathy. *Diabetes Care*. (2003) 26:99–102. doi: 10.2337/diacare.26.2007.S99
- American Diabetes Association. Standards of medical care for patients with diabetes mellitus. *Diabetes Care*. (2003) 26(Suppl 1):s33–50. doi: 10.2337/diacare.26.2007.s33
- Wong TY, Sun J, Kawasaki R, Ruamviboonsuk P, Gupta N, Lansingh VC, et al. Guidelines on diabetic eye care: the international council of ophthalmology recommendations for screening, follow-up, referral, and treatment based on resource settings. *Ophthalmology*. (2018) 125:1608–22. doi: 10.1016/j.ophtha.2018.04.007
- Scotland GS, McNamee P, Fleming AD, Goatman KA, Philip S, Prescott GJ, et al. Costs and consequences of automated algorithms versus manual grading for the detection of referable diabetic retinopathy. *Br J Ophthalmol*. (2010) 94:712–9. doi: 10.1136/bjo.2008.151126
- Tufail A, Kapetanakis VV, Salas-Vega S, Egan C, Rudisill C, Owen CG, et al. An observational study to assess if automated diabetic retinopathy image assessment software can replace one or more steps of manual imaging grading and to determine their cost-effectiveness. *Health Technol Assess*. (2016) 20:1–72. doi: 10.3310/HTA20920
- Grzybowski A, Brona P, Lim G, Ruamviboonsuk P, Tan GS, Abramoff M, et al. Artificial intelligence for diabetic retinopathy screening: a review. *Eye*. (2020) 34:451–60. doi: 10.1038/s41433-019-0566-0
- Wu E, Wu K, Daneshjouri R, Ouyang D, Ho DE, Zou J. How medical AI devices are evaluated: limitations and recommendations from an analysis of FDA approvals. *Nat Med*. (2021) 27:582–4. doi: 10.1038/s41591-021-01312-x
- Lee AY, Yanagihara RT, Lee CS, Blazes M, Jung HC, Chee YE, et al. Multicenter, head-to-head, real-world validation study of seven automated artificial intelligence diabetic retinopathy screening systems. *Diabetes Care*. (2021) 44:1168–75. doi: 10.2337/DC20-1877
- Nagendran M, Chen Y, Lovejoy CA, Gordon AC, Komorowski M, Harvey H, et al. Artificial intelligence versus clinicians: systematic review of design, reporting standards, and claims of deep learning studies. *BMJ*. (2020) 368:m689. doi: 10.1136/bmj.m689
- Farić N, Hinder S, Williams R, Ramaesh R, Bernabeu MO, van Beek E, et al. Early experiences of integrating an artificial intelligence-based diagnostic decision support system into radiology settings: a qualitative study. *J Am Med Inform Assoc*. (2023) 31:24–34. doi: 10.1093/JAMIA/OCAD191
- Beede E, Baylor E, Hersch F, Iurchenko A, Wilcox L, Ruamviboonsuk P, et al. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*; Honolulu, HI, USA. New York, NY: Association for Computing Machinery (2020). p. 1–12. doi: 10.1145/3313831.3376718
- Data from: Kaggle competition: diabetic retinopathy detection (2016). Available at: <https://www.kaggle.com/c/diabetic-retinopathy-detection> (Accessed October 23, 2023).
- Data from: APTOS 2019 blindness detection in Kaggle (2019). Available at: <https://www.kaggle.com/competitions/aptos2019-blindness-detection> (Accessed October 23, 2023).
- Krause J, Gulshan V, Rahimy E, Karth P, Widner K, Corrado GS, et al. Grader variability and the importance of reference standards for evaluating machine learning models for diabetic retinopathy. *Ophthalmology*. (2018) 125:1264–72. doi: 10.1016/j.ophtha.2018.01.034
- Porwal P, Pachade S, Kamble R, Kokare M, Deshmukh G, Sahasrabudhe V, et al. Data from: Indian diabetic retinopathy image dataset (IDRiD) (2018). doi: 10.21227/H25W98
- Li T, Gao Y, Wang K, Guo S, Liu H, Kang H. Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Inf Sci (Nij)*. (2019) 501:511–22. doi: 10.1016/j.ins.2019.06.011
- Fu H, Wang B, Shen J, Cui S, Xu Y, Liu J, et al. Evaluation of retinal image quality assessment networks in different color-spaces. In: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 11764 LNCS. (2019). p. 48–56. doi: 10.1007/978-3-030-32239-7_6
- van der Heijden AA, Abramoff MD, Verbraak F, van Hecke MV, Liem A, Nijpels G. Validation of automated screening for referable diabetic retinopathy with the IDx-DR device in the hoorn diabetes care system. *Acta Ophthalmol (Copenh)*. (2018) 96:63–8. doi: 10.1111/AOS.13613
- Ruamviboonsuk P, Krause J, Chotcomwongse P, Sayres R, Raman R, Widner K, et al. Deep learning versus human graders for classifying diabetic retinopathy severity in a nationwide screening program. *NPJ Digit Med*. (2019) 2:25. doi: 10.1038/S41746-019-0099-8
- Gulshan V, Rajan RP, Widner K, Wu D, Wubbels P, Rhodes T, et al. Performance of a deep-learning algorithm vs manual grading for detecting diabetic retinopathy in India. *JAMA Ophthalmol*. (2019) 137:987–93. doi: 10.1001/jamaophthalmol.2019.2004
- Ipp E, Liljenquist D, Bode B, Shah VN, Silverstein S, Regillo CD, et al. Pivotal evaluation of an artificial intelligence system for autonomous detection of referable and vision-threatening diabetic retinopathy. *JAMA Network Open*. (2021) 4:e2134254. doi: 10.1001/JAMANETWORKOPEN.2021.34254
- Lim JI, Regillo CD, Sadda SR, Ipp E, Bhaskaranand M, Ramachandra C, et al. Artificial intelligence detection of diabetic retinopathy. *Ophthalmol Sci*. (2023) 3:100228. doi: 10.1016/j.xops.2022.100228
- Heydon P, Egan C, Bolter L, Chambers R, Anderson J, Aldington S, et al. Prospective evaluation of an artificial intelligence-enabled algorithm for automated diabetic retinopathy screening of 30000 patients. *Br J Ophthalmol*. (2021) 105:723–8. doi: 10.1136/BJOPHTHALMOL-2020-316594
- Skevas C, Weindler H, Levering M, Engelberts J, van Grinsven M, Katz T. Simultaneous screening and classification of diabetic retinopathy and age-related macular degeneration based on fundus photos—a prospective analysis of the RetCAD system. *Int J Ophthalmol*. (2022) 15:1985. doi: 10.18240/IJO.2022.12.14
- Meredith S, van Grinsven M, Engelberts J, Clarke D, Prior V, Vodrey J, et al. Performance of an artificial intelligence automated system for diabetic eye screening in a large English population. *Diabet Med*. (2023) 40:e15055. doi: 10.1111/DME.15055
- Ribeiro L, Oliveira CM, Neves C, Ramos JD, Ferreira H, Cunha-Vaz J. Screening for diabetic retinopathy in the central region of Portugal. Added value of automated “Disease/no disease” grading. *Ophthalmologica*. (2014) 233:96–103. doi: 10.1159/000368426
- Atwany MZ, Sahyoun AH, Yaqub M. Deep learning techniques for diabetic retinopathy classification: a survey. *IEEE Access*. (2022) 10:28642–55. doi: 10.1109/ACCESS.2022.3157632
- Lin J, Yu L, Weng Q, Zheng X. Retinal image quality assessment for diabetic retinopathy screening: a survey. *Multimed Tools Appl*. (2020) 79:16173–99. doi: 10.1007/s11042-019-07751-6
- Andonegui J, Serrano L, Eguzkiza A, Berastegui L, Aliseda DJ, Gaminde I. Diabetic retinopathy screening using tele-ophthalmology in a primary care setting. *J Telemed Telecare*. (2010) 16:429–32. doi: 10.1258/jtt.2010.091204
- Andonegui J, Zurutuza A, Arcelus MD, Serrano L, Eguzkiza A, Auzmendi M, et al. Diabetic retinopathy screening with non-mydratric retinography by general practitioners: 2-year results. *Prim Care Diabetes*. (2012) 6(3):201–5. doi: 10.1016/j.pcd.2012.01.001
- Wilkinson CP, Ferris FL, Klein RE, Lee PP, Agardh CD, Davis M, et al. Proposed international clinical diabetic retinopathy and diabetic macular edema disease severity scales. *Ophthalmology*. (2003) 110:1677–82. doi: 10.1016/S0161-6420(03)00475-5
- Huemer J, Wagner SK, Sim DA. The evolution of diabetic retinopathy screening programmes: a chronology of retinal photography from 35 mm slides to artificial intelligence. *Clin Ophthalmol*. (2020) 14:2021. doi: 10.2147/OPHTH.S261629
- Cavallerano AA, Conlin PR. Technology for diabetes care and evaluation in the veterans health administration: Teleretinal imaging to screen for diabetic retinopathy in the veterans health administration. *J Diabetes Sci Technol*. (2008) 2:33. doi: 10.1177/193229680800200106
- Zachariah S, Wykes W, Yorston D. The Scottish diabetic retinopathy screening programme. *Community Eye Health*. (2015) 28:s22. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4944112/>
- Nguyen HV, Tan GSW, Tapp RJ, Mital S, Ting DSW, Wong HT, et al. Cost-effectiveness of a national telemedicine diabetic retinopathy screening program in Singapore. *Ophthalmology*. (2016) 123:2571–80. doi: 10.1016/j.ophtha.2016.08.021
- Scanlon PH. The english national screening programme for diabetic retinopathy 2003–2016. *Acta Diabetol*. (2017) 54:515. doi: 10.1007/S00592-017-0974-1
- Pereira AMP, de Lima Neto FB. Five regions, five retinopathy screening programmes: a systematic review of how Portugal addresses the challenge. *BMC Health Serv Res*. (2021) 21:756. doi: 10.1186/S12913-021-06776-8
- Ta AWA, Goh HL, Ang C, Koh LY, Poon K, Miller SM. Two Singapore public healthcare AI applications for national screening programs and other examples. *Health Care Sci*. (2022) 1:41–57. doi: 10.1002/hcs2.10
- Miller SM. Tracing the twenty-year evolution of developing AI for eye screening in Singapore: a master chronology of SiDRP, SELENA+ and EyRis. *Res Collect Sch Comput Inf Syst*. (2023):1–32. Available at: https://ink.library.smu.edu.sg/sis_research/7833
- Ting DSW, Cheung CYL, Lim G, Tan GSW, Quang ND, Gan A, et al. Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. *JAMA*. (2017) 318:2211–23. doi: 10.1001/JAMA.2017.18152
- Lee JC, Nguyen L, Hynan LS, Blomquist PH. Comparison of 1-field, 2-fields, and 3-fields fundus photography for detection and grading of diabetic retinopathy. *J Diabetes Complicat*. (2019) 33:107441. doi: 10.1016/J.JDIACOMP.2019.107441

42. Grading diabetic retinopathy from stereoscopic color fundus photographs—an extension of the modified airline house classification: ETDRS report number 10. *ETDRS Report* (1991). doi: 10.1016/j.ophtha.2020.01.030
43. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Tuytelaars T, Li F-F, Bajcsy R, editors. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; Las Vegas, NV, USA. Piscataway, NJ: IEEE (2016). p. 770–8. doi: 10.1109/CVPR.2016.90
44. Howard J, Gugger S. Layered API for deep learning. *Information*. (2020) 11:108. doi: 10.3390/info11020108
45. Howard J, Gugger S, O'Reilly Media Company Safari. *Deep Learning for Coders with Fastai and PyTorch*. Sebastopol, CA: O'REILLY (2020).
46. Ben-David A. About the relationship between roc curves and cohen's kappa. *Eng Appl Artif Intell*. (2008) 21:874–82. doi: 10.1016/J.ENGAPPAL.2007.09.009
47. Data from: Michael D. Abramoff's web page, where messidor-2 dataset can be downloaded (2023). Available at: <https://medicine.uiowa.edu/eye/abramoff> (Accessed October 23, 2023).
48. Cuadros J, Bresnick G. EyePACS: An adaptable telemedicine system for diabetic retinopathy screening. *J Diabetes Sci Technol*. (2009) 3:509. doi: 10.1177/193229680900300315
49. Data from: Full EyePACS dataset in kaggle (2023). Available at: <https://www.kaggle.com/datasets/benjaminwarner/resized-2015-2019-blindness-detection-images> (Accessed October 23, 2023).
50. Buderer NMF. Statistical methodology: incorporating the prevalence of disease into the sample size calculation for sensitivity and specificity. *Acad Emerg Med*. (1996) 3:895–900. doi: 10.1111/J.1553-2712.1996.TB03538.X
51. Hajian-Tilaki K. Sample size estimation in diagnostic test studies of biomedical informatics. *J Biomed Inform*. (2014) 48:193–204. doi: 10.1016/j.jbi.2014.02.013
52. Filho TS, Song H, Perello-Nieto M, Santos-Rodriguez R, Kull M, Flach P. Classifier calibration: how to assess and improve predicted class probabilities: a survey. *Mach Learn*. (2023) 112:3211–60. doi: 10.48550/arxiv.2112.10327
53. Kull M, Filho TS, Flach P. Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In: Singh A, Zhu J, editors. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR (2017). Proceedings of Machine Learning Research; vol. 54. p. 623–31.
54. Zadrozny B, Elkan C. Transforming classifier scores into accurate multiclass probability estimates. In: *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. (2002). p. 694–9.
55. Linardatos P, Papastefanopoulos V, Kotsiantis S. Explainable AI: a review of machine learning interpretability methods. *Entropy*. (2021) 23:18. doi: 10.3390/E23010018
56. Sundararajan M, Taly A, Yan Q. Axiomatic attribution for deep networks—integrated gradients. In: *34th International Conference on Machine Learning, ICML 2017*. (2017). Vol. 7. p. 5109–18.
57. Sayres R, Taly A, Rahimy E, Blumer K, Coz D, Hammel N, et al. Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy. *Ophthalmology*. (2019) 126:552–64. doi: 10.1016/j.ophtha.2018.11.016
58. Ankerst M, Breunig MM, Kriegel HP, Sander J. OPTICS: ordering points to identify the clustering structure. *ACM SIGMOD Rec*. (1999) 28:49–60. doi: 10.1145/304181.304187
59. Pizer SM, Amburn EP, Austin JD, Cromartie R, Geselowitz A, Greer T, et al. Adaptive histogram equalization and its variations. *Comput Vis Graph Image Process*. (1987) 39:355–68. doi: 10.1016/S0734-189X(87)80186-X
60. Zuiderveld K. VIII.5. – Contrast limited adaptive histogram equalization. In: Heckbert PS, editor. *Graphics Gems*. Boston, MA: Academic Press (1994). p. 474–85. doi: 10.1016/B978-0-12-336156-1.50061-6
61. Knuth DE. *Literate Programming (Center for the Study of Language and Information Publication Lecture Notes)*. Stanford, CA: Center for the Study of Language and Information Publications (1992). Vol. 384.
62. Abràmoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *npj Digit Med*. (2018) 1:1–8. doi: 10.1038/s41746-018-0040-6
63. Shah A, Clarida W, Amelon R, Hernaez-Ortega MC, Navea A, Morales-Olivas J, et al. Validation of automated screening for referable diabetic retinopathy with an autonomous diagnostic artificial intelligence system in a Spanish population. *J Diabetes Sci Technol*. (2020) 15:655–63. doi: 10.1177/1932296820906212
64. Quéllec G, Lamard M, Lay B, Guilcher AL, Erginay A, Cochener B, et al. Instant automatic diagnosis of diabetic retinopathy (2019). *Arxiv* [Preprint].
65. McHugh ML. Interrater reliability: the kappa statistic. *Biochem Med (Zagreb)*. (2012) 22:276. doi: 10.11613/bm.2012.031
66. Mooney CZ, Duval RD, Duvall R. *Bootstrapping: A Nonparametric Approach to Statistical Inference*. Thousand Oaks, CA: SAGE (1993).
67. Papadopoulos A, Topouzis F, Delopoulos A. An interpretable multiple-instance approach for the detection of referable diabetic retinopathy in fundus images. *Sci Rep*. (2021) 11:1–15. doi: 10.1038/s41598-021-93632-8
68. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. (2016) 316:2402–10. doi: 10.1001/jama.2016.17216
69. Voets M, Möllersen K, Bongo LA. Reproduction study using public data of: development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *PLoS One*. (2019) 14:e0217541. doi: 10.1371/JOURNAL.PONE.0217541
70. Liu R, Wang X, Wu Q, Dai L, Fang X, Yan T, et al. Deepdrid: Diabetic retinopathy-grading and image quality estimation challenge. *Patterns*. (2022) 3:100512. doi: 10.1016/J.PATTERN.2022.100512
71. Sevik U, Köse C, Berber T, Erdöl H. Identification of suitable fundus images using automated quality assessment methods. *J Biomed Opt*. (2014) 19:046006. doi: 10.1117/1.JBO.19.4.046006
72. Muehlematter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–20): a comparative analysis. *Lancet Digit Health*. (2021) 3:e195–e203. doi: 10.1016/S2589-7500(20)30292-2