INNOVATION–SANDBOX
for Artificial Intelligence

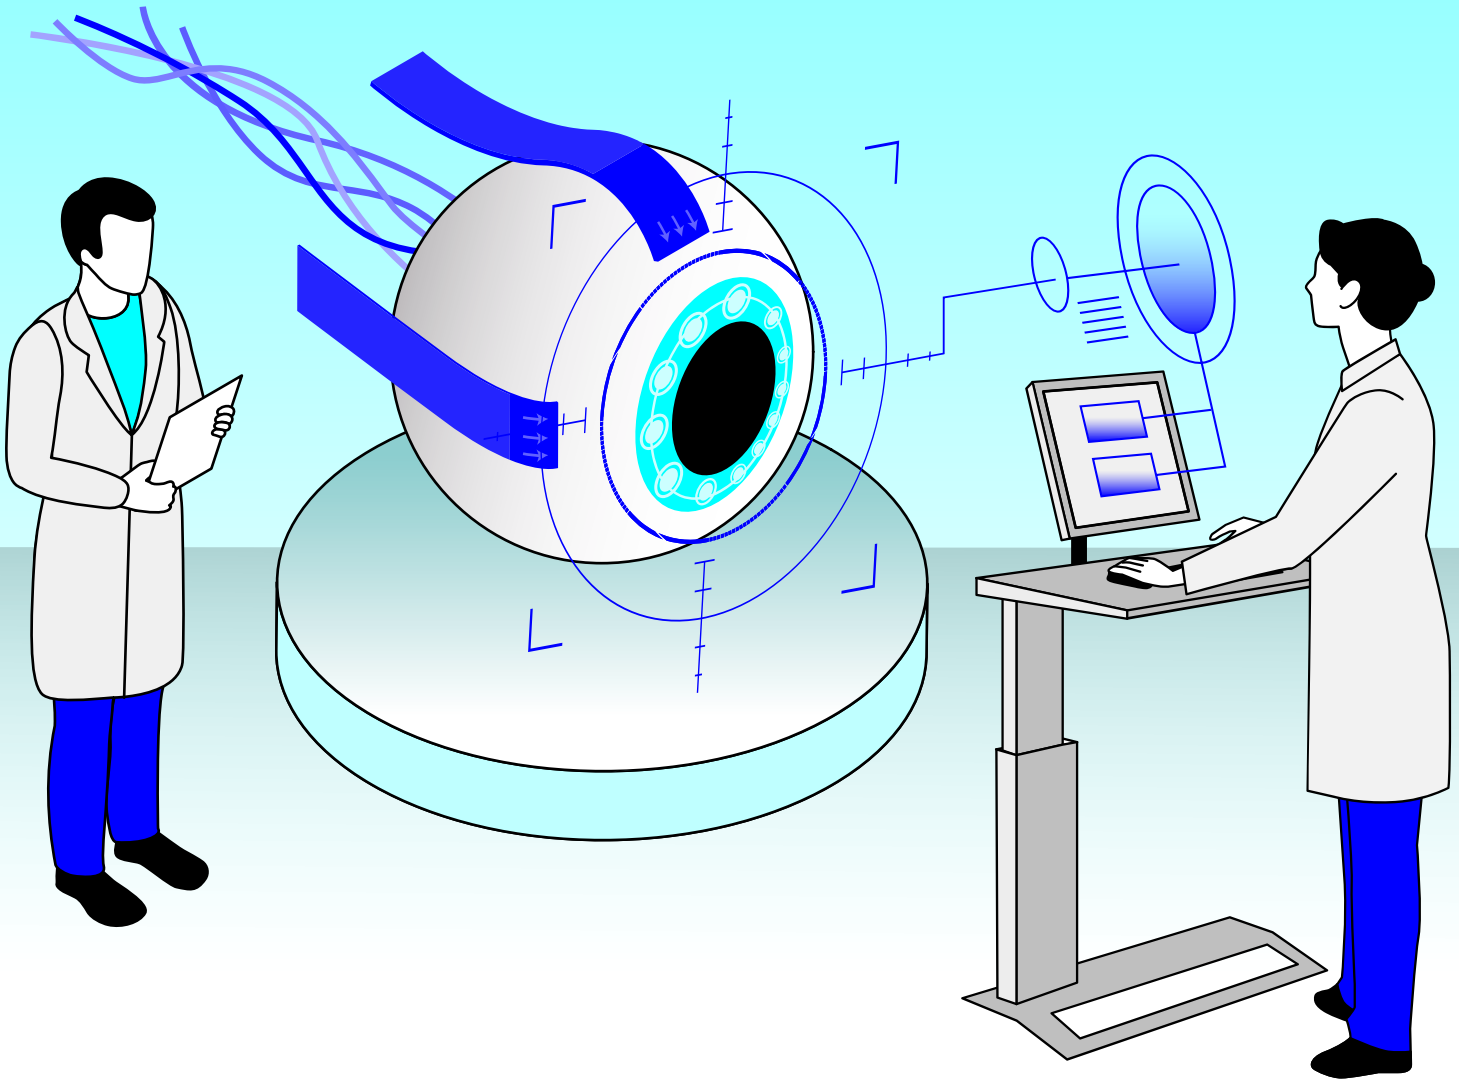

Digital Eye Clinic

Bringing AI diagnostics into practice

innovationsandbox.ai

metropolitan konferenz zürich

Kanton Zürich

Innovation Zurich

Diabetic retinopathy is one of the most frequent ocular diseases and poses a growing challenge due to rising patient numbers, staff shortages and increasing healthcare costs. Automating routine screening processes with artificial intelligence (AI) can reduce the workload of medical specialists and improve efficiency without compromising the quality of care. The Digital Eye Clinic project, carried out by Stadtspital Zürich and the Spross Research Institute for the Advancement of Ophthalmology, was part of the Canton of Zurich's Innovation Sandbox for AI. It explored the medical, technological, regulatory and ethical prerequisites for integrating AI diagnostics into clinical workflows. The project evaluated existing diagnostic providers and platforms within the current legal and technical framework of the city of Zurich. As no commercial vendor met the clinic's key criteria, the team developed and validated its own AI models based on open-source research frameworks. Alongside technological validation, the project examined critical issues in regulation, data protection and medical ethics. The results provide practical insights and best practices for hospitals and healthcare providers considering AI-assisted diagnostics. They emphasise the need for structured selection and implementation processes, open and interoperable architectures, early clinical involvement, and a careful assessment of whether to use open-source models or rely on external vendors. Beyond ophthalmology, the findings offer transferable lessons for the safe and effective adoption of AI in other areas of clinical practice.

Innovation Sandbox for AI

The project team produced this document as part of the Innovation Sandbox for AI. The Sandbox is a test environment for implementing AI projects from various sectors. This broad-based initiative from government, business and research promotes responsible innovation by ensuring the project team and participating organisations work together closely on regulatory issues and enable the use of new data sources.

More information

Contents

Project-Team

Prof. Matthias Becker, MD
Head of the Department of Ophthalmology and President of the Werner H. Spross Stiftung zur Förderung der Augenheilkunde*, Stadtspital Zürich

Gábor Márk Somfai, MD, PhD
Senior Consultant and Head of Research, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

Tahm Spitznagel, MD
Research Resident, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

Rui Santos
PhD, AI Translation Officer, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

Dávid Isztl
Machine Learning Engineer, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

Raphael von Thiessen
AI Sandbox Programme Manager, Canton of Zurich

Florian Lüchinger
Managing Director, Foresite AG

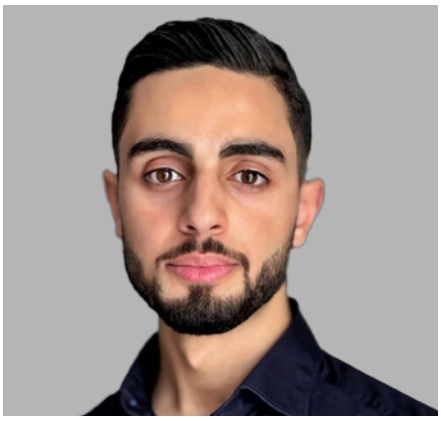

Amr Saad, MD, MBA
Resident, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

Stephanie Volz
Managing Director, ITSL, University of Zurich

* The foundation supports the Spross Research Institute

01. *Addressing diabetic retinopathy*

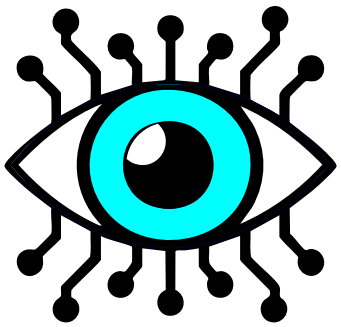

*Diabetic retinopathy (DR)** is one of the leading causes of blindness in industrialised countries and poses a major burden on society. In Switzerland, an estimated 430,000 people live with diabetes[1], and around 25% of them are affected by some degree of DR[2]. Regular eye examinations are essential to prevent vision loss including vision-threatening complications such as diabetic macular oedema or proliferative retinopathy. International guidelines recommend annual screenings for diabetic patients. These screenings are resource-intensive, yet only about one third of screened patients show signs of diabetic retinopathy and roughly 10% experience visual impairment[3]. Unfortunately, less than half of the population with known diabetes gets their annual screening exams, making room for late diabetic eye complications. This highlights the need for more efficient and targeted screening tools. This mismatch underscores inefficiency in today's system: specialists spend valuable time on routine examinations that could be supported by digital tools or even automated.

At the same time, healthcare systems face a shortage of trained ophthalmologists and rising personnel costs. AI-based diagnostic tools offer an opportunity to relieve specialists, reduce costs and ensure preventive healthcare across the population. These systems use *deep-learning* models trained on large sets of retinal images to identify disease patterns with a high degree of accuracy and consistency. Automated screening solutions can triage patients, so that only those with suspected findings are referred for a consultation. In the long term, these kinds of solutions could also enable walk-in clinics on the high street , providing rapid, low-threshold access to eye screenings and facilitating earlier detection of disease.

«Diabetic retinopathy is a leading cause of preventable blindness.»

Tahm Spitznagel, MD – Research Resident, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

* Highlighted words are explained in more detail in the glossary.
[1] https://diabetesatlas.org/data-by-location/country/switzerland/
[2] Prevalence, incidence and future projection of diabetic eye disease in Europe: a systematic review and metaanalysis – PubMed
[3] Global prevalence and major risk factors of diabetic retinopathy – PubMed

01. Addressing diabetic retinopathy

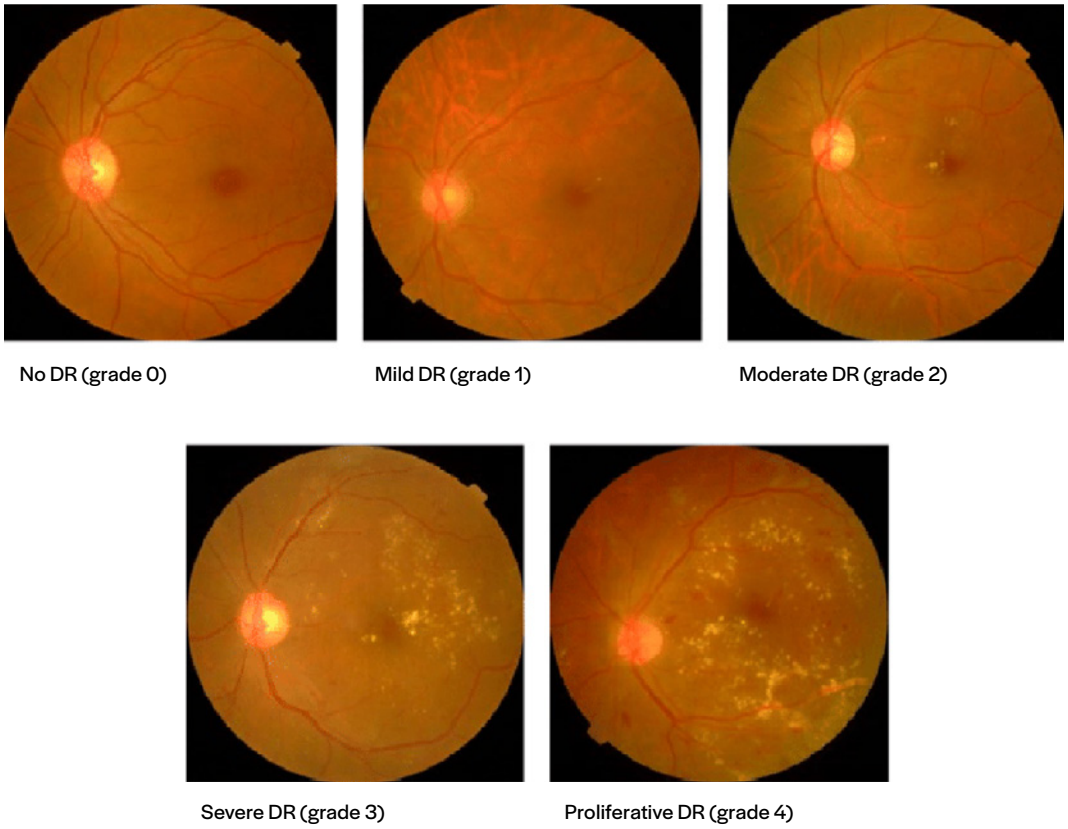

Figure 1: Representative images illustrating the spectrum from a healthy retina (grade 0), through the non-proliferative stages (grades 1–3), to proliferative diabetic retinopathy (grade 4) (source: APTOS dataset).

[4] Example: In Germany, the drugstore chain dm offers walk-in eye screenings with AI analysis and physician validation in cooperation with Skleo Health.

01. Addressing diabetic retinopathy

The Digital Eye Clinic project was part of the Innovation Sandbox for AI and explored the technical, regulatory, economic and ethical prerequisites for integrating AI diagnostic solutions into ophthalmology. Based on a concrete project proposal by Stadtspital Zürich and the Spross Research Institute for the Advancement of Ophthalmology, the project team developed insights on key questions: What could the patient journey of the future look like? How should a healthcare provider assess third-party diagnostic tools that are using AI at their core? What are the practical steps to develop an in-house diagnostics solution using *open-source AI models*? What are the most pressing legal and regulatory requirements?

The following chapters summarise the main findings of this work and aim to provide a foundation for researchers and practitioners in other medical fields that mainly rely on medical imaging for screenings and diagnosis, such as dermatology, radiology, and pathology, among others, to build upon.

02.
The medical and technological potential of AI

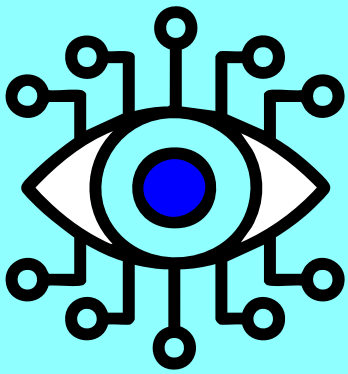

AI offers significant opportunities in ophthalmology. Algorithms can analyse large sets of retinal images quickly and consistently, supporting the detection and grading of DR with an accuracy comparable to that of medical specialists[5]. This makes AI particularly attractive for large-scale screening programs, where efficiency, scalability and reliability are essential. By automating routine checks, specialists can focus on the smaller group of patients who show disease progression.

Beyond eye diseases, the emerging field of ***oculomics*** demonstrates that retinal scans can provide insights into a person's overall health[6]. Research indicates that AI models can infer factors such as blood pressure, smoking status, body mass index, or metabolic health, and even predict early signs of systemic conditions like cardiovascular disease, multiple sclerosis, Parkinson's disease, dementia, or Alzheimer's disease. This suggests that the eye could serve as a «window into the body», offering new pathways for preventive and personalised medicine.

«Oculomics holds promise for the earlier detection, risk prediction and personalised treatment of ocular and systemic diseases.»
Gábor Márk Somfai, MD, PhD – Senior Consultant and Head of Research, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

[5] Artificial intelligence versus manual screening for the detection of DR: a comparative systematic review and metaanalysis
[6] Transformative applications of oculomics-based AI approaches in the management of systemic diseases: A systematic review

02. The medical and technological potential of AI

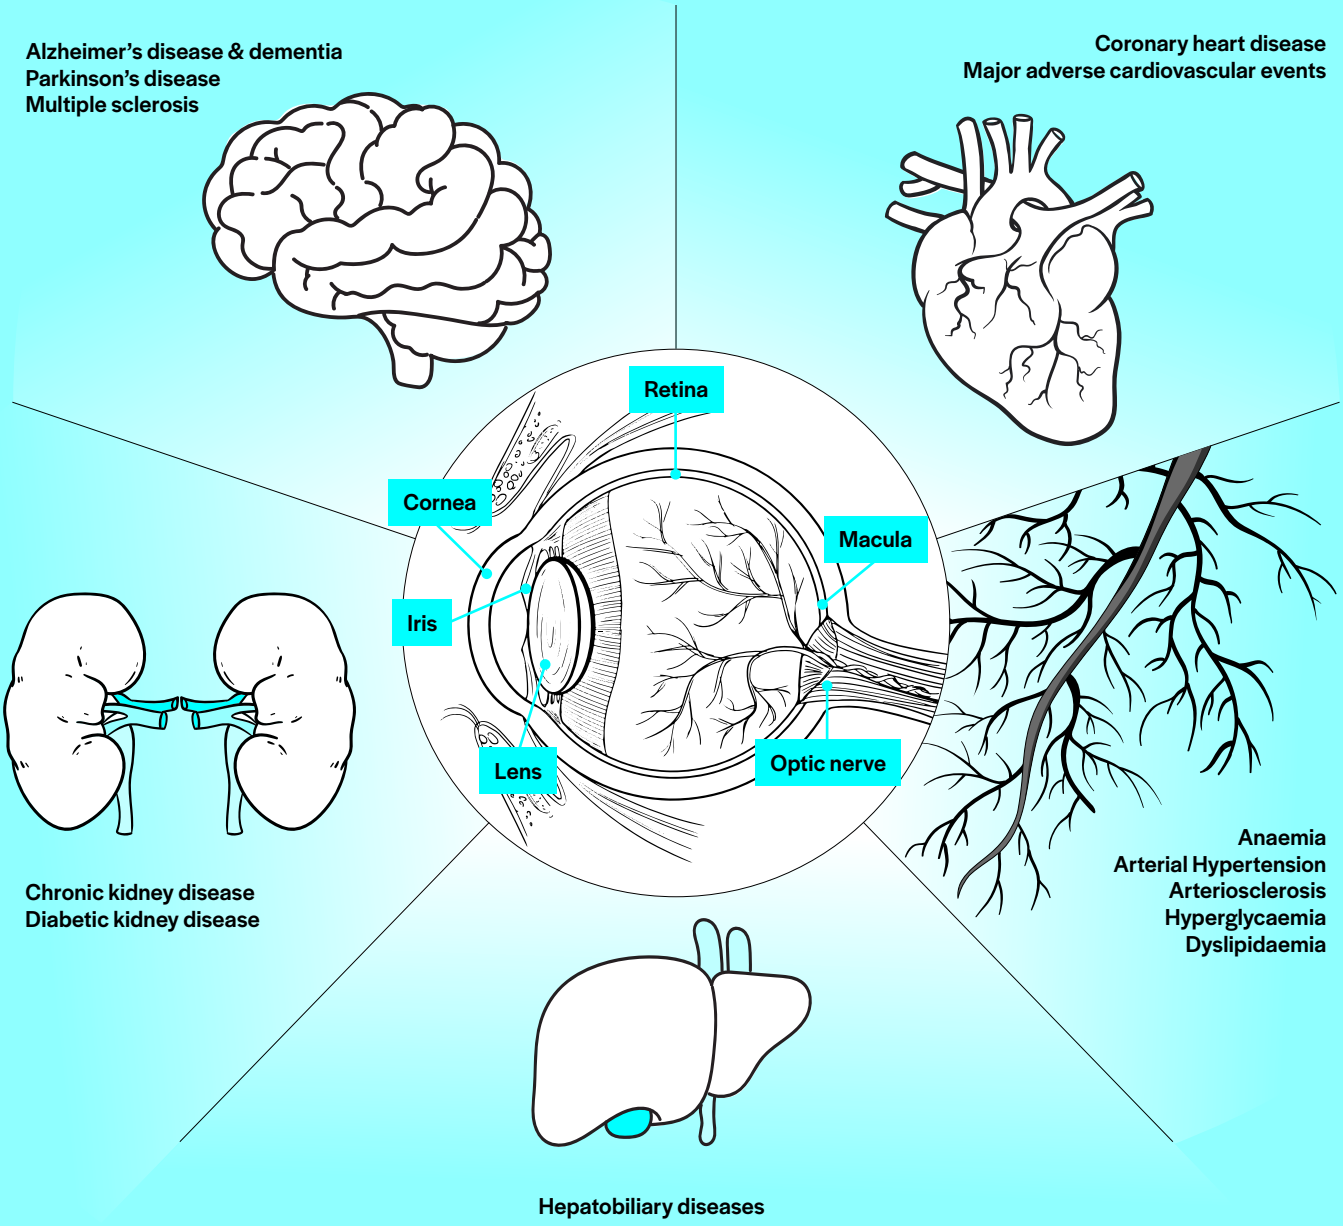

Figure 2: Schematic view of the eye as a window into systemic health based on Oculomics: Current concepts and evidence

02. The medical and technological potential of AI

Different imaging modalities play an important role. Two of the most standard techniques are:

- ***Colour fundus photography (CFP)***
  A 2D image of the retina is captured by a camera. It is cost-efficient, widely available in clinics and particularly well suited for deep-learning models used in DR screening.

- ***Optical coherence tomography (OCT)***
  A 3D imaging technique that captures high-resolution cross-sections of the retinal layers. It has become one of the most used modalities in ophthalmology worldwide.

At the same time, key limitations and open questions remain. Image quality varies across devices and strongly influences the reliability of AI results. Many algorithms are trained on limited datasets and may not yet be robust across diverse patient populations, raising concerns about generalisability and bias. Questions regarding data privacy and patient communication are also critical: Should healthcare providers inform patients about findings like as a ***retinal age gap*** (the difference between biological age inferred from the retina using deep learning and chronological age which aim to estimate overall ageing), and how might such information affect their wellbeing? Finally, it remains essential to emphasise that AI should complement – not replace – clinical decision-making.

In summary, AI diagnostics in ophthalmology hold great promise for prevention, early detection and efficient monitoring; however, their successful integration into practice will require certified tools, robust data quality, transparent human-AI collaboration, clear regulatory guidance and careful consideration of ethical issues.

03.
Designing the patient journey with AI diagnostics

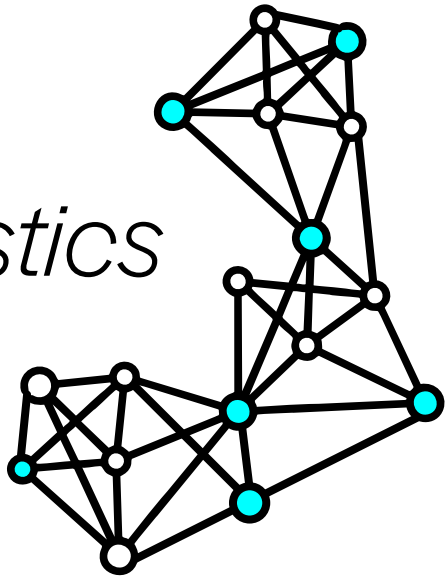

The Digital Eye Clinic team envisioned a patient journey that integrates AI diagnostics seamlessly into the clinical workflow. In this vision, a walk-in patient undergoes standard admission and imaging (either via OCT or CFP) as usual. AI algorithms then automatically process the captured images, which generates a preliminary diagnostic AI report. The ophthalmologist reviews the AI output, verifies or adjusts the findings (adding comments or corrections as needed), and then uses the report to guide the therapy. In addition, large language models (LLMs) can support clinicians by summarising the patient visit, prefilling the electronic health record, and generating referral letters or reports. This end-to-end process – from patient intake to image acquisition, AI analysis, clinical review and treatment decision – defines the future AI-enhanced patient journey.

The diagram above illustrates this patient journey and clinical workflow in detail. A patient's case is created in the hospital information system, and relevant diagnostic images (OCT scans or CFP) are captured and sent to the diagnosis management system for analysis. The AI system automatically starts analysing the images and produces an AI report (preliminary diagnosis). The clinician then checks the AI-generated report and can either approve it or modify/comment on the report before final approval. Throughout this workflow, images are stored in an image archive system, and the AI results (both unauthorised and authorised reports) are transferred into the clinic's systems. Once the report is finalised, the patient proceeds to therapy or further procedures as needed. This patient journey map demonstrates how AI diagnostics could be embedded in routine practice, maintaining a clear division between automated analysis and physician oversight.

03. Designing the patient journey with AI diagnostics

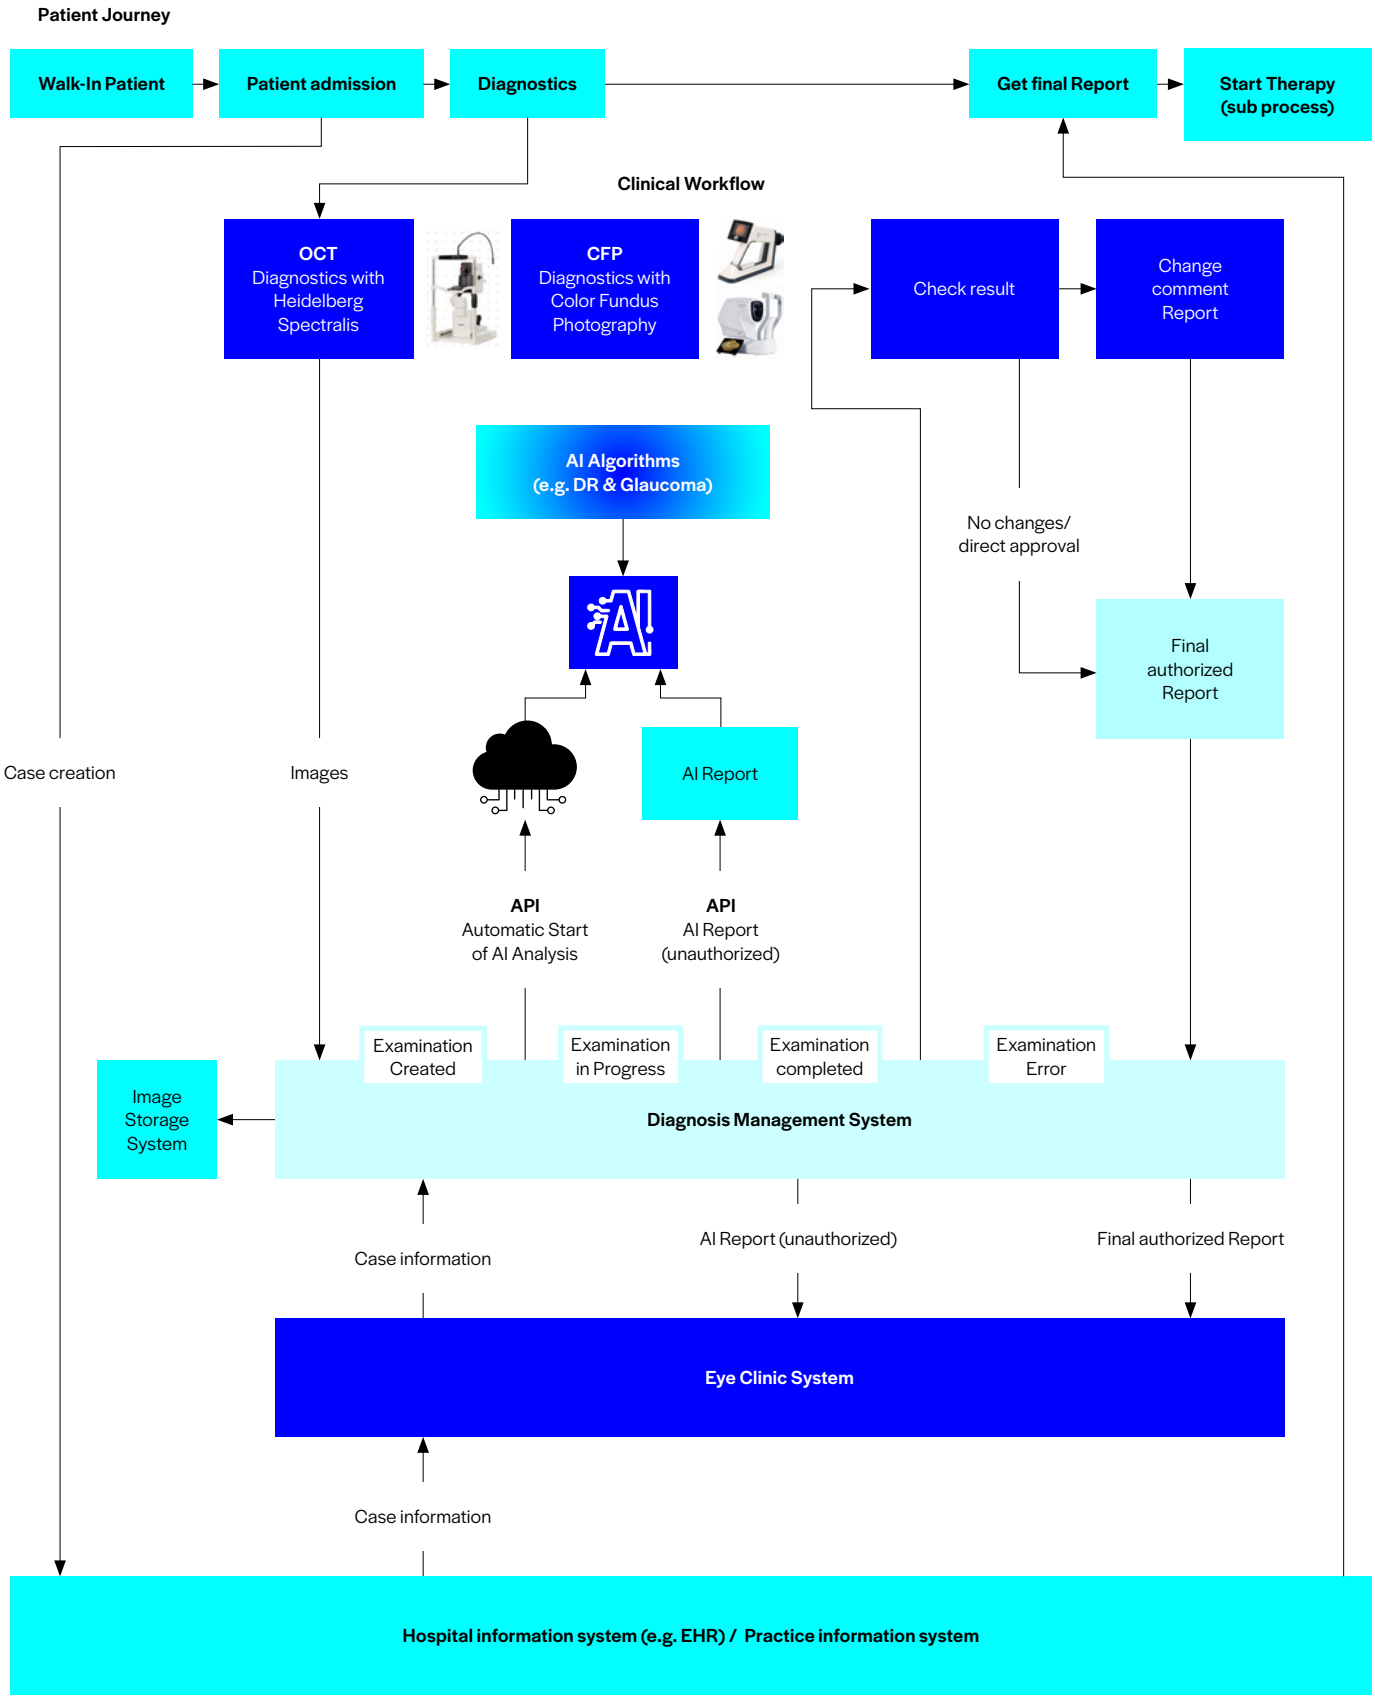

Figure 3: Patient journey and clinical workflow integrating AI diagnostics

04. *Assessment of diagnostic providers and platforms*

- Multiple diagnostic tools based on AI have obtained regulatory approval or certification since 2018, e.g., from the *Food and Drug Administration (FDA)* in the US and are now available commercially. The project evaluated various AI diagnostic providers and platforms to determine if any could be integrated for a *Minimum viable product (MVP)* implementation. A set of technological, regulatory and economic criteria was defined to compare these solutions:

Technological criteria
- Supported imaging modalities (CFP and/or OCT)
- Compliance with *DICOM* standards for image formats
- Overlay functions for AI results to ensure transparency of decision-making
- Autonomous decision-making vs. assistive support (fully automated system or diagnostic aid)
- Integration of in-house AI models to ensure extensibility of the platform
- Compatibility with existing system landscape (electronic health record, image management systems, etc.)

Regulatory criteria
- Regulatory approvals (e.g., *CE marking* under *medical device regulation (MDR)* in Europe, FDA clearance/approval in the US)
- Data governance: secure data storage, audit trails, traceability of results

Economic criteria
- Cost per scan or per analysis
- Scalability of the solution (can it handle increasing volumes)
- Contract models (licensing, pay-per-use, subscriptions, etc.)

Case examples of providers considered included five vendors[7] of AI diagnostic tools for DR. The project team assessed each against the above criteria.

Outcome of market analysis
Using the criteria, the project team conducted a market analysis to compare the leading AI diagnostic solutions. Table 1 summarises key features of four prominent providers. The detailed comparison excluded one provider as the company did not provide sufficient information in response to the market analysis request. The evaluation compared these providers for their AI integration capabilities, supported imaging modalities, regulatory approvals, device compatibilities and other relevant factors.

[7] Further information on the vendors is available from the Sandbox team upon request.

04. Assessment of diagnostic providers and platforms

Criteria	Vendor 1	Vendor 2	Vendor 3	Vendor 4
Own AI integration	Yes	Yes	Yes	Yes
Third-party AI integration	No	No	No	No
Regulatory approval	FDA (USA), CE marked	FDA (USA)	FDA (USA)	EU MDR
Disease focus	*DR, AMD, glaucoma*	DR	DR	DR, AMD, glaucoma, etc.
Image modality	OCT images	OCT images	CFP images	CFP/OCT images
DICOM support	Yes	Yes	Yes	Yes
Supported devices	Heidelberg Spectralis, Zeiss Clarus	Topcon NW400 (proprietary camera)	Zeiss Clarus (fundus camera)	Heidelberg Spectralis/Zeiss
Swiss references	None	None	None	Concrete Swiss references available
Cloud/hosting	Cloud-based (AWS)	On-prem or private (AWS used for some services)	On-prem or private (AWS backend)	Yes/AWS

Table 1: Comparison of key criterion of selected AI diagnostic providers (for the abbreviations, see the glossary).

04. Assessment of diagnostic providers and platforms

As shown in Table 1, all evaluated platforms allowed for the integration of their own AI models to some degree, but none supported integrating third-party AI algorithms beyond what they provided themselves. All had at least some regulatory approval (FDA and/or CE/MDR certification) for DR detection. There were notable differences in disease coverage and imaging modalities: Vendor 1 and Vendor 4 support both OCT and fundus (CFP) images, whereas Vendor 3 focusses on fundus images and Vendor 2 focusses on fundus via their own camera system. All platforms claimed DICOM compatibility for image data storage.

One practical consideration was device support. Vendor 1 and Vendor 4 can interface with common ophthalmic OCT imaging devices like Heidelberg Spectralis OCT and Zeiss Clarus fundus cameras. In contrast, the solution from Vendor 2 is tied to a proprietary fundus camera (Topcon NW400), limiting flexibility. Vendor 3 primarily supports standard fundus camera inputs (e.g., Zeiss Clarus) and may not yet handle OCT, which was a drawback for a clinic wanting multimodal support.

Another important factor was local experience: Neither Vendor 1, Vendor 2 nor Vendor 3 had existing references or deployments in Switzerland at the time of evaluation, whereas Vendor 4 had local reference sites. This provided some confidence in Vendor 4's ability to understand local needs, but it had to be weighed against technical criteria.

Cloud and IT integration: Vendor 1 and Vendor 4 are designed as cloud-native solutions (both favouring deployment on AWS cloud). Vendor 2 and Vendor 3 can be deployed on-premises but also utilise cloud components (the slide indicated «No/AWS», suggesting they are not purely cloud-native but do use AWS in some capacity). For the hospital's MVP, on-premises capability was important due to data governance concerns, so fully cloud-only solutions were suboptimal.

In addition to technical features, cost and trial conditions were compared (Table 2):

04. Assessment of diagnostic providers and platforms

Provider	Pricing model	Trial availability
Vendor 1	Subscription tiers by volume: e.g., 125 scans/month at EUR189; 250 scans/month at EUR329; higher volumes negotiable.	Yes – offered a free trial (e.g., approx. 5,000 images for 6 months).
Vendor 2	Upfront device + subscription: Requires purchase of Topcon NW400 camera (approx. USD18,000), plus USD1,500/month for unlimited AI analyses.	No free trial – no trial program was offered (only direct purchase).
Vendor 3	One-time license + per-patient: One-time fee approx. CHF6,000–9,000, then approx. CHF 10 per patient analysed (unlimited images per patient).	Yes – trial available (e.g., a limited-time pilot license, approx. CHF 5,000/year).
Vendor 4	Setup + subscription: One-time setup EUR9,000, plus EUR1,330/month license fee.	Yes – pilot projects possible (often negotiated case-by-case).
Vendor 5	No pricing information – did not provide an offer	No data (no response)

Table 2: AI diagnostic provider pricing and trial offers

04. Assessment of diagnostic providers and platforms

Pricing models (as of 2025) are shown for different usage volumes, along with availability of trial programs.

From a cost perspective, Vendor 1 stood out for offering a flexible subscription with relatively low entry cost and even a free trial period. Vendor 3 required a moderate onetime investment but charged per patient, which could become expensive at scale. Vendor 2 had the highest upfront cost (requiring a proprietary camera purchase) and a substantial monthly fee, making it the most expensive option for the clinic's expected volume. The model of Vendor 4 involved significant setup fees and ongoing costs, reflecting its positioning as an enterprise platform.

During the analysis, Vendor 5 could not be fully evaluated because the company did not provide sufficient information (did not respond to inquiries). This lack of responsiveness raised concerns about transparency and support for this project.

Ultimately, none of the off-the-shelf providers perfectly met the clinic's requirements for the MVP, which was the ultimate reason for their exclusion.

«Only a holistic, well-structured vendor assessment can reveal if an AI diagnostic tool is fit for clinical use and institutional adoption.»
Raphael von Thiessen, AI Sandbox Programme Manager, Canton of Zurich

Some specific drawbacks noted were:

- **Vendor 2**
  Very high total cost of ownership (especially the mandatory camera), limited flexibility in supported hardware (only their camera), and no Swiss user base or references.

- **Vendor 3**
  No local presence or sales/support office in Switzerland, support for only CFP images (no OCT analysis capability), lack of Swiss references, and relatively outdated or less userfriendly interface.

- **Vendor 5**
  Unresponsive during evaluation, providing no data. This raised concerns about transparency, unclear regulatory status, and absence of local language support or references.

04. Assessment of diagnostic providers and platforms

Given these limitations, the evaluation concluded that no single AI vendor could satisfy the key MVP criteria (notably, **none could integrate the clinic's own AI model**, support both OCT and CFP modalities on the required devices and run on the preferred Azure/on-premise infrastructure). Tables 1 and 2 underscore these gaps. For example, no evaluated solution supported the Stadtspital Zürich custom AI integration, Zeiss Clarus fundus camera imaging, and an on-premise or Azure deployment all together. As a result, the project team decided that pursuing a proprietary solution was necessary. Building an in-house AI model and a custom platform became the chosen path.

05.
Development and validation of in-house AI models

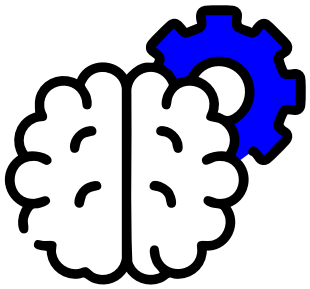

With no vendor platform fully suitable, the team moved to develop and test a proprietary AI model and system:

«By developing and deploying an in-house DR detection model, we can uphold the highest levels of data privacy and compliancy.»

Rui Santos, PhD - AI Translation Officer, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

- **Model choice**
  The MVP leveraged an existing research-based model adapted from an open-source foundation model (Meta AI's *DinoV2*). Specifically, the team used the Block Expanded *DINORET*[8] approach (an adaptation of a natural image foundation model for retinal imaging). This made it possible to build on state-of-the-art vision model architectures without starting from scratch.

- **Training and validation**
  The project team validated the custom model on various ophthalmic datasets to ensure it could detect diabetic retinopathy (and potentially other retinal findings) accurately. Furthermore, the project team pursued access to large datasets (e.g., applying for UK Biobank retinal images) to improve validation quality.

- **Advantages of a custom model**
  Developing the model in-house offered several benefits: the team had full control over model performance and transparency of the AI's decision process; the model could be optimised for the clinic's own patient cohort (reflecting local population characteristics); and the approach avoids vendor lock-in, potentially enabling reuse or extension of the model beyond this clinic (benefiting other clinics or research).

- **Challenges**
  The project also encountered challenges common to custom AI development. Obtaining annotated datasets for training was labourintensive; significant computing resources were required

[8] Block Expanded DINORET: Adapting Natural Domain Foundation Models for Retinal Imaging Without Catastrophic Forgetting

05. Development and validation of in-house AI models

for model training, *fine-tuning* and inference; and close clinical collaboration was needed to define the model outputs and evaluate its performance. Additionally, navigating governance and approval processes for AI in a clinical setting proved complex – e.g., determining if the in-house AI would be classified as a medical device under MDR and how to get ethics committee oversight (see chapter 7).

- **MVP setup**
  The initial MVP ran on a local PC workstation within the hospital network, first requiring *GPU*-enabled computation. Later on, the pipeline was optimize to be CPU-compute only. This setup had direct access to the hospital's data pool and allowed development of a *graphical user interface (GUI)* to be used by clinicians. The choice of a local deployment was deliberate in order to keep patient data on-site and to iterate quickly without cloud dependencies.

- **Implications**
  Going with an open-source, selfdeveloped model had positive implications, such as avoiding ongoing license fees, eliminating dependency on external vendors and fostering internal know-how. However, it also came with negative implications: the clinic must take on the maintenance and iterative improvement of the model, which requires sustained effort and specialised expertise that a vendor solution would normally provide. In summary, the team had to balance cost and flexibility benefits against the burden of longterm support and development.

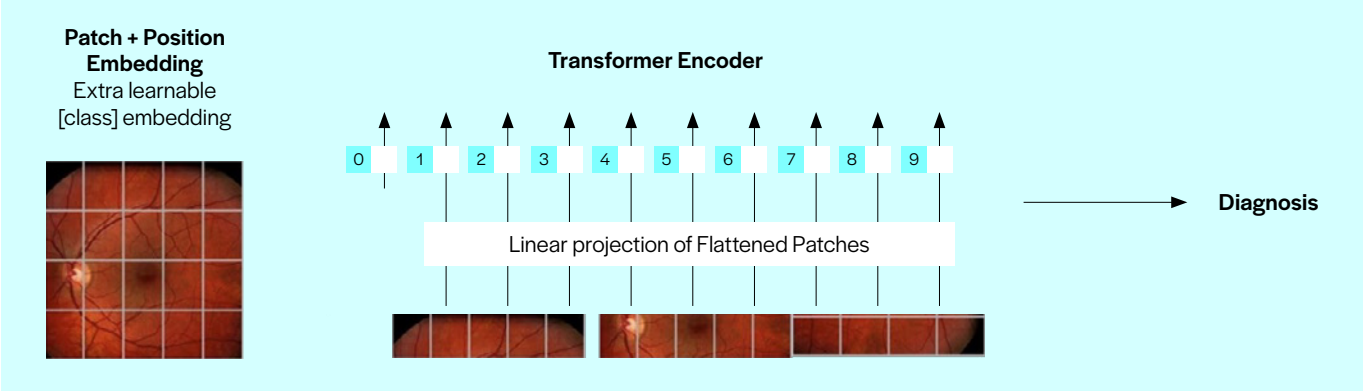

Figure 4: Simplified schematic of a *Vision Transformer (ViT)* architecture applied to diabetic retinopathy classification using colour fundus images.

06. *Concept for testing phase*

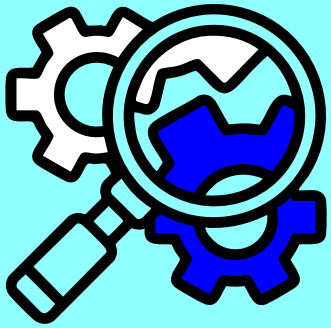

After a cycle of concepts and an initial development phase including testing, the MVP prototype system was up and running. The following section summarises the results from the concept phase.

The MVP concept included a range of basic but essential functionalities to support the diagnostic workflow:

- Secure user authentication (with username/password and two-factor login) to control access to the system.
- A patient examination management module that allows staff to create new exam records, attach patient info and images, run AI analysis, and track the status of each exam.
- Integration of the AI analysis pipeline: The system must send OCT or CFP images to the AI model, run the inference, and retrieve results automatically. This must be done through a simple interface where the user selects images and an AI model (e.g., DR detection) and clicks «Run Examination».
- Results and reporting: After analysis, the AI's findings (e.g., detection of DR severity or normal vs. abnormal results) are stored and displayed. Each exam entry in the system shows a summary of the AI result (e.g., «Mild non-proliferative retinopathy» or «Normal macular structure») and allows the user to open a detailed report.
- Error handling: If the AI model fails or encounters an error, the system marks the exam with an «error» status (with an indication such as «Model not running» or «Error in analysis»), so it can be flagged for technical review.
- Data storage and export: All images and exam results are saved with-in the system. Basic DICOM compliance is maintained so that data can be exported or integrated with other hospital systems in the future.

MVP software and user interface
A significant part of the MVP after the concept phase was the creation of a graphical user interface (GUI) to make the AI tool usable in the clinical workflow. The GUI was developed as a web-based application accessible on the hospital's network. As it was a prototype, it was intentionally kept simple and secure.

Prototype evaluation: To assess the MVP's performance, an initial evaluation method was defined. This included comparing the AI's diagnostic outputs with a medical second opinion for each case to calculate sensitivity and specificity and track metrics like false positives/negatives. The team also measured technical robustness (e.g., average processing time per case, system uptime) to ensure functionality in a clinical environment. In addition, feedback (a clinical review) was gathered from the ophthalmologists using the system: they provided input on the usefulness of the AI reports, practical applicability in daily opera-

06. Concept for testing phase

tions, and required adjustments to the workflow or interface. These insights formed the basis for further improvements, including additional validation cycles or final regulatory steps before deployment.

Lessons learned: The prototyping phase yielded important insights. There were clear trade-offs between third-party solutions and proprietary models: while building our own solution offered flexibility, it also meant assuming full responsibility for performance and support. Most importantly, medical acceptance proved crucial: even a well-performing AI is only valuable if the clinical team trusts and adopts it. Continuous involvement of end users (the clinic's doctors and staff) during prototyping is essential to foster trust in the AI and improve system usability.

«Machine learning in ophthalmology paves the way for safe, accurate, and accessible ocular disease prediction.»

Dávid Isztl - Machine Learning Engineer, Department of Ophthalmology and Spross Research Institute, Stadtspital Zürich

06. Concept for testing phase

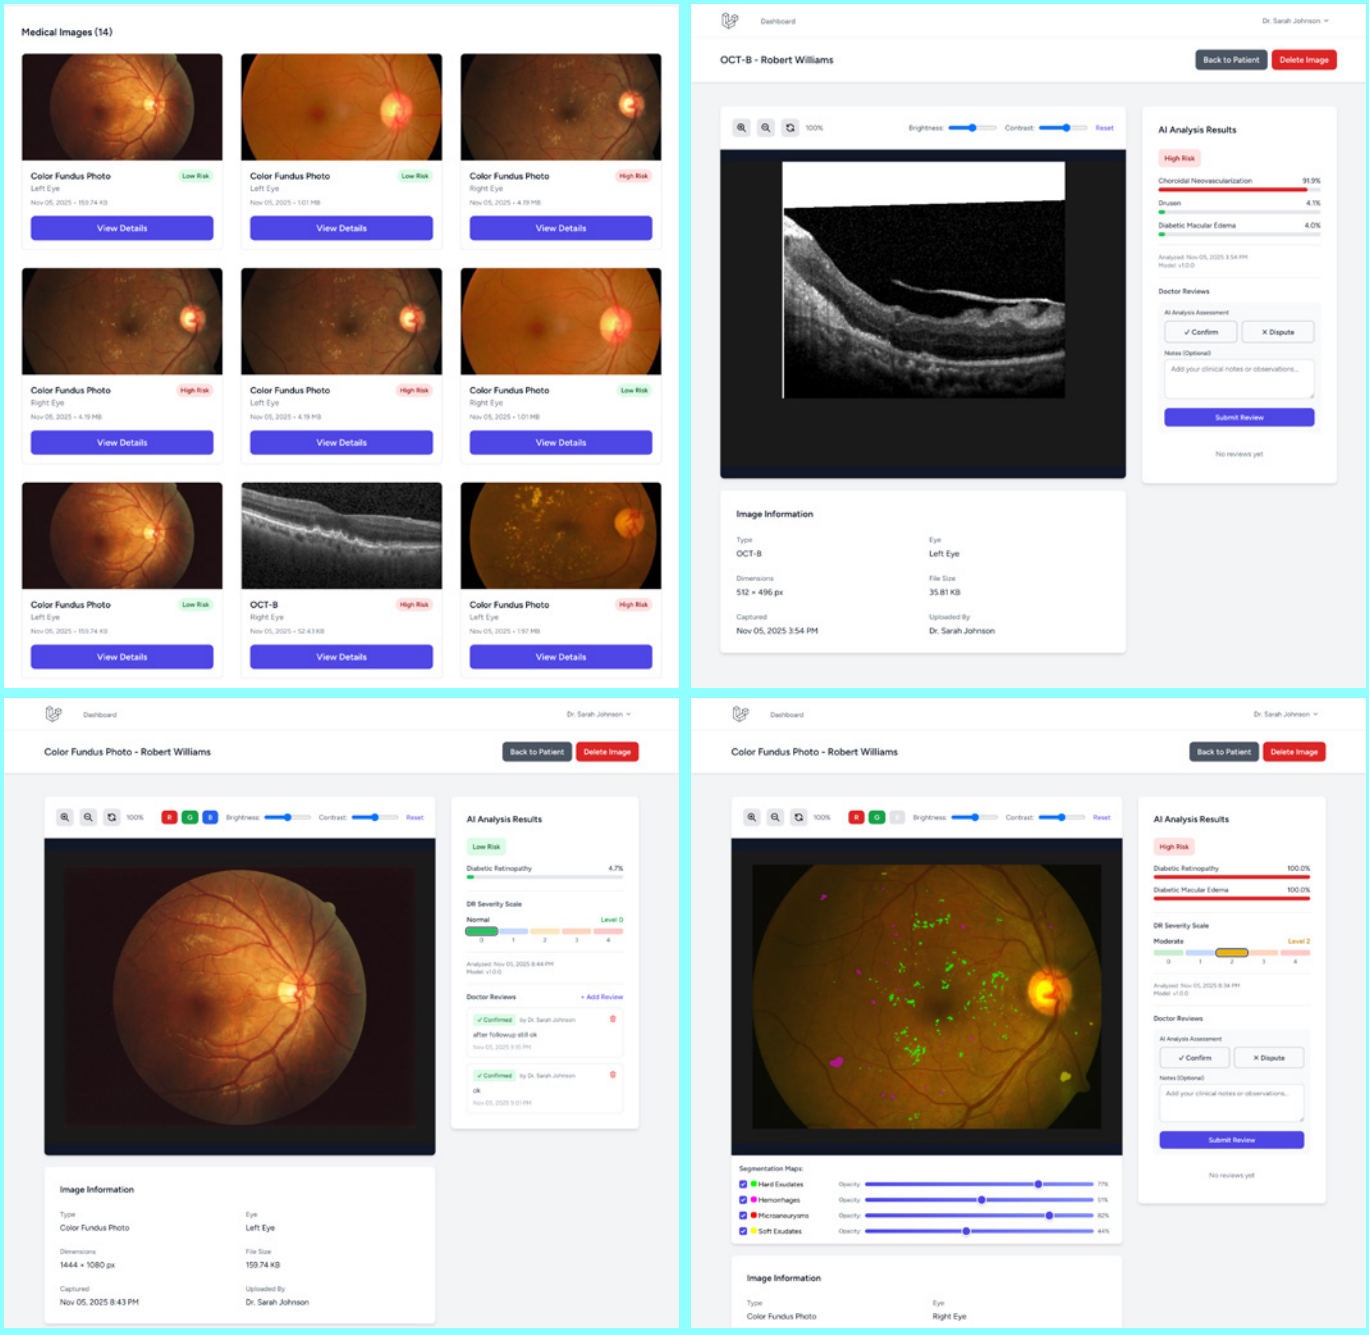

Figure 5: Actual deployed MVP interface – creating a new examination, image selection, AI model choice and model prediction.

07.
Moving towards the Digital Eye Clinic

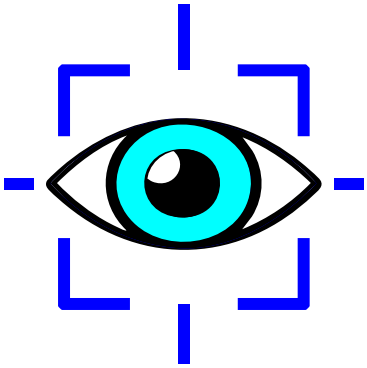

The MVP is only the first milestone on the path to implementing AI diagnostics in clinical practice through the Digital Eye Clinic. The next phase will focus on advancing the diabetic retinopathy module toward clinical use by addressing all operational, technical and regulatory requirements.

In parallel, additional uses such as ***glaucoma*** detection are being onboarded to extend the system's applicability and clinical value. Given the open and interoperable architecture, the platform will continuously evaluate and integrate the best-performing AI models for each diagnostic use case.

Cloud-based deployment is the current target state, enabling scalable, secure and easily maintainable operations. Edge solutions are also being evaluated and tested for local deployment. In the long term, AI diagnostics could also support routine eye screenings in primary care or community settings, improving early detection and accessibility. Finally, the platform's modular setup allows for future expansion, which would make the AI diagnostics solution available to additional clinics and research partners.

It is important to note that the vision of the Digital Eye Clinic goes beyond AI diagnostics – it aims to leverage AI across the entire clinical value chain, for example through AI-assisted medical documentation and automated patient triage and scheduling. In all applications, technology serves the well-being of the patient, with a human expert always in the loop to ensure safety, trust and clinical acceptance.

08. *Regulatory and ethical requirements*

Any AI diagnostic solution in healthcare must navigate complex regulatory and ethical considerations[9]. The project identified key requirements and considerations in this domain:

- **Medical device regulations:** The AI diagnostic tool falls under the purview of the Medical Devices Act and European Union Medical Device Regulation (EU MDR) if it influences clinical decisions. The team needs to determine which regulatory class the AI would fall under (likely Class IIa or IIb for diagnostic support) and what steps (documentation, risk analysis, etc.) are required to use it, even internally.

- **In-house use vs. commercial use:** A special consideration was whether the in-house developed AI, used only within the hospital, could be exempt from certain regulatory requirements (the so-called «in-house exception»). Generally, using a selfdeveloped AI within a single healthcare institution might not trigger full MDR compliance as a marketed device would, but the boundaries are nuanced. The project consulted with regulatory experts to clarify the threshold at which even an in-house tool might be considered a medical device needing regulatory clearance (for example, if outcomes are used for patient treatment decisions, it likely qualifies as a medical device, even if it is not sold).

- **Data protection and privacy:** Strict compliance with data protection laws is necessary, including adhering to internal Information Security and Data Privacy (ISDS) processes. Using cloud services with patient data introduces additional scrutiny. For the MVP, all patient images and data were kept on-premises; if cloud storage (e.g., Microsoft Azure) were to be used in future, measures like encryption, deidentification and contractual data processing agreements would be mandatory. Potential restrictions regarding the US Cloud Act would also need to be addressed.

- **Ethical and patient safeguards:** Before deploying AI diagnostics in a clinical setting, it is essential to ensure that patient rights, privacy and safety are protected. This includes transparent communication to patients that AI is being used in an assistive role, that the physician remains responsible for the final diagnosis, and that data use complies with applicable legal and ethical standards.

- **Bias and fairness:** AI models can inherit biases from their training data, making it crucial to evaluate whether the model's performance is consistent across different patient subgroups (e.g., ethnicity, gender, age). The training data origin was considered: Did it represent the local patient population? For instance, if the model is

[9] Another sandbox report on «AI in medical documentation» provides a more detailed description of the regulatory requirements for AI solutions in the healthcare sector.

08. Regulatory and ethical requirements

primarily trained on images from another region, it might underperform on the target demographics. These issues of explainability and fairness need to be continuously documented, and plans must be made to retrain or recalibrate the model if any bias is detected.

- **Transparency and accountability:** The project emphasised making AI decisions transparent and traceable. Each AI report stored the underlying images and analysis steps, so a human expert could retrospectively interpret why the AI might have made a certain assessment (e.g., via heat-maps or highlighting of image areas if available). Clinicians will remain the final decision-makers, and the workflow is designed so that the AI assists rather than diagnosing automatically to ensure a human in the loop for patient safety. Clearly defining who is responsible if the AI makes an error (the oversight clinician, the development team, etc.) is also part of the ethical governance to avoid ambiguity in accountability.

In summary, regulatory and ethical compliance were treated as a core component of the project from the outset, not as an afterthought. This proactive approach helped build trust among stakeholders (clinicians, patients, regulators) and laid the groundwork for eventual scaling of the solution beyond the pilot phase.

09. *Conclusions and best practices*

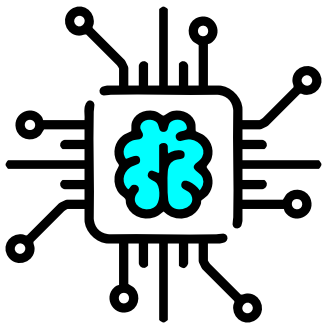

As part of the Innovation Sandbox for AI, the Smart Eye Clinic project demonstrated both the potential and the challenges of integrating AI into ophthalmology practice. In conclusion, while AI technology for detecting eye diseases is maturing rapidly, successful adoption in a clinical setting requires more than just algorithms – it demands careful consideration of workflow integration, user experience, regulatory compliance and stakeholder buyin. The project's experiences led to several best-practice insights that are applicable to other ophthalmology clinics (and other medical domains) looking to implement similar AI diagnostic tools.

Best-practice recommendations for AI diagnostics
Based on our lessons learned, we have outlined a set of transferable success factors, common challenges and recommended approaches for implementing AI diagnostics in an eye clinic:

- **Structure the selection and implementation process:** Approach AI integration as a structured project. Start with a needs assessment and clear objectives (e.g., «reduce screening burden for DR by 30%»). Use well-defined criteria to evaluate solutions. This structured approach helps compare options objectively and ensures that all stakeholders know the roadmap. Best practice: Break the project into phases (market research, pilot, validation, deployment) and have decision gates at each phase.
- **Prefer open, interoperable architectures:** Favour solutions that can integrate with your existing systems and adapt over time. Open standards (like DICOM for imaging, FHIR for health records) and modular architectures will prevent any issues later. Interoperability avoids vendor lock-in and allows you to plug in new AI models, in-house developed diagnostic tools or alternative imaging devices as they become available. In this case, lacking interoperability (e.g., a vendor not supporting the current camera setup) was a dealbreaker. Best practice: Insist on compatibility and available APIs when evaluating AI platforms.
- **Ensure clinical involvement from the outset:** Engage medical specialists and clinical staff early and often. Their buy-in is crucial for success – an AI tool will only be used if the end users trust it and find it improves their workflow. In this project, having doctors codesign the workflow and interface was key to driving adoption. They provided feedback on everything from which results are clinically relevant to how the AI output should be presented. Best practice: Form a core user group of clinicians to pilot the AI and iterate on its use.
- **Evaluate open-source vs. vendor solutions:** When adopting AI tools, it is important to decide whether to adapt an open-source model or purchase a commercial solution. Open-source models can often be tailored to local needs and offer more transparency, but they require technical expertise and support structures for safe use.

09. Conclusions and best practices

Vendor solutions, on the other hand, are typically easier and quicker to deploy, yet may be less flexible and can involve higher recurring costs. Best practice is to systematically compare both options in terms of accuracy, integration, sustainability and longterm value for the clinical setting.

- **Data protection as a challenge for deployment:** Strict data protection rules often slow down AI deployment in practice, especially when using cloud solutions for sensitive health data in a public hospital. While on-premise setups can simplify compliance and keep sensitive data under local control, they require significant IT resources. Best practice: Address data protection requirements from the start and evaluate early on whether cloud or on-premise is the more feasible path.

- **Address regulatory issues early:** Engaging with the institution's regulatory experts from the beginning will smooth out deployment later. Navigate the question of MDR classification for your AI tool early on – even if it's an internal tool, document its intended use and risk mitigation. Ensure patient consent and transparency when AI is used in their care (even in a pilot). This avoids roadblocks where a promising pilot cannot transition to practice due to compliance issues. Best practice: Treat the regulatory strategy as part of the project plan, not separate or after the fact, and keep documentation (audit trails, performance evaluations) that regulators may require.

- **Leverage transferability to other specialties:** Although the project's use case was ophthalmology, many of the approaches can be generalised. Designing the system with a flexible back end (for example, the ability to plug in a different AI model and dataset) means the solution could be extended to, e.g., radiology, dermatology or pathology in the future. This foresight can make your investment more valuable. Best practice: Build AI infrastructure (data pipelines, databases, interfaces) in a general way when possible, so that adding another AI for a different medical condition would be easier down the line.

09. Conclusions and best practices

Key «dos and don'ts» in clinical AI implementation: Finally, we can summarise a few critical dos and don'ts:

- Do ensure robust validation of the AI with real-world patient data before relying on it – treat initial results with caution and involve clinicians in verifying outputs.
- Do invest in training and change management for your staff; even the best AI is ineffective if end users don't know how to use it or don't trust it.
- Don't underestimate the importance of the user interface and workflow integration – a poorly integrated tool will be ignored, no matter how good the algorithm is.
- Don't assume one size fits all; be ready to customise the AI's decision thresholds or retrain it on your patient population to improve its performance and acceptance.

In conclusion, the Digital Eye Clinic project showcased a successful proof-of-concept for AI-assisted ophthalmic diagnostics but equally highlighted that technology is just one piece of the puzzle. Organisational readiness, clinician engagement, and a clear focus on patient benefit are just as important. By following the best practices outlined above – from structured planning and interoperable design to proactive stakeholder involvement – other clinics can better navigate the journey of AI integration and reap the benefits of this promising technology in a safe and effective manner.

Case studies from the Innovation Sandbox for AI
The Digital Eye Clinic project by the Stadtspital Zurich and the Werner H. Spross Foundation for the Advancement of Ophthalmology served as a case study within the Innovation Sandbox for AI. The organisation submitted a pro-ject proposal to the AI Sandbox in summer 2024. The project team worked on the implementation of AI diagnostics in ophthalmology between September 2024 and November 2025. The content of this report was developed based on this specific case study.

Glossary

Age-related macular degeneration (AMD)
A common eye disease that affects the macula, the central part of the retina responsible for sharp vision. AMD is a leading cause of vision loss in older adults and can progress in «dry» or «wet» forms.

Artificial intelligence (AI)
Computer systems designed to perform tasks that normally require human intelligence, such as recognising patterns or making decisions. In medicine, AI can support diagnosis and screening.

CE marking
A certification mark that indicates a product's conformity with European health, safety and environmental protection standards. For medical devices, including AI-based software, CE marking under the EU Medical Device Regulation (MDR) is required before they can be marketed in the European Economic Area (EEA).

Colour fundus photography (CFP)
A imaging technique that uses a specialised camera to capture colour photographs of the retina. It is widely used for screening and monitoring several eye diseases and, at times, systemic diseases as well.

Deep learning
A subfield of machine learning that uses multi-layered artificial neural networks to automatically learn patterns from large amounts of data. In healthcare and ophthalmology, deep learning is widely applied to medical images (e.g., retinal scans, skin lesion analysis, neuroradiology) for tasks such as detection, classification and prediction.

Digital imaging and communications in medicine (DICOM)
An international standard for storing, transmitting and managing medical images. It ensures interoperability between different imaging systems.

Diabetic retinopathy (DR)
A diabetes-related eye disease caused by damage to the blood vessels in the retina. It can lead to vision loss if untreated.

DINORET
An AI model derived from DINOV2, adapted specifically for retinal image analysis. It is used in research for ophthalmic diagnostics.

DINOV2
An advanced open-source computer vision model developed by Meta AI. It is widely used for image recognition tasks, including medical imaging.
Food and Drug Administration (FDA)
The regulatory authority in the United States responsible for approving and overseeing drugs, medical devices and diagnostic tools. For AI-based diagnostics, the FDA grants clearance or approval depending on the risk class and regulatory pathway.

Glaucoma
A group of eye diseases that damage the optic nerve, often due to increased intraocular pressure. It can lead to irreversible vision loss if untreated. AI-based diagnostic tools can assist in early detection by analysing optic nerve and retinal imaging data.

Graphics processing unit (GPU)
A type of computer processor originally designed for graphics rendering. Today, GPUs are essential hardware accelerators for training and running AI models due to their high parallel processing power.

Graphical user interface (GUI)
A visual interface that allows users to interact with software through icons, menus and graphics instead of text commands. It improves usability in clinical systems

Glossary

Medical Device Regulation (MDR)
The European Union regulation (EU 2017/745) that governs the safety and performance of medical devices, including AI-based software. MDR requires risk classification, conformity assessment and CE marking before a device can be placed on the European market.

Minimum viable product (MVP)
A basic version of a new product with only essential features, created to test functionality and gather feedback. It helps reduce development risk.

Oculomics
An emerging field that studies how ocular imaging data, such as retinal scans, may reveal information about ocular and overall health. It links ocular imaging features to systemic conditions.

Optical coherence tomography (OCT)
A non-invasive imaging method that provides detailed cross-sectional images of the retina. It is a standard tool in modern ophthalmology.

Retinal age gap
The difference between a person's biological age and the age predicted by AI from their retina. A larger gap may signal underlying health risks.

Vision Transformer (ViT)
A deep-learning model that applies transformer architecture to image analysis. In oculomics, ViTs process retinal images by dividing them into patches and using attention mechanisms to detect and classify ocular and systemic diseases such as diabetic retinopathy.

Impressum

Publisher
Division of Business and Economic Development, Canton of Zurich
Metropolitan Area Zurich Association
Innovation Zurich

Project conception and coordination
Raphael von Thiessen
Standortförderung Kanton Zürich
8090 Zürich
raphael.vonthiessen@vd.zh.ch

Concept in collaboration with
Stephanie Volz
Isabell Metzler
Patrick Arnecke
Markus Müller

Authors
Raphael von Thiessen
Dávid Isztl
Rui Santos
Gábor Márk Somfai
Tahm Spitznagel
Florian Lüchinger

Design
here we are gmbh, here-we-are.ch

Publication
This report is published exclusively in digital format and is available in German and English.

Project-Steering
- Office for Economy, Canton of Zurich
- Statistical Office, Canton of Zurich
- Department for Digital Administration, Canton of Zurich
- Office for Economy, Canton of Schwyz
- Zurich Metropolitan Area Association
- ETH AI Center
- Center for Information Technology, Society, and Law (ITSL), University of Zurich
- ZHAW entrepreneurship