WHITEPAPER

Creating a veterinary pathology AI model

A step-by-step example of using Aiforia® Create to develop an AI model to segment the normal retina

Author:
Richard Fox, BVetMed, MRCVS, Dipl ECVP, Veterinary Pathologist at Aiforia

aiforia®
AI for image analysis

TABLE OF CONTENTS

Introduction

Veterinary pathology has long been an essential field in veterinary medicine, aiding in the diagnosis of diseases, infections, and cancers in animals. With the growing sophistication of medical technologies, there is a shift toward automating processes traditionally dependent on human interpretation. One such advancement is the development of artificial intelligence (AI) solutions to assist pathologists.

Aiforia® Create is a tool for designing and implementing AI-based histopathology models. To showcase its capabilities and prove it's not rocket science, we swiftly developed **a model to segment the normal retina from a rat**. However, this approach is easily adaptable to other species and applications.

The retina is a highly specialized, multi-layered neural tissue responsible for vision, and its accurate segmentation is crucial for studying both normal and pathological conditions. Establishing an AI-driven image analysis model for a normal rat retina lays the foundation for analyzing diseases like glaucoma-induced retinal atrophy and other degenerative diseases in the future. This whitepaper outlines the ease of use of Aiforia® Create to develop an AI model.

Planning the AI model

Before beginning the model development process, it is essential to define clear objectives and set up a structured approach. A suitable ground truth (guide) should also be established before starting notation on what needs to be annotated and what is not required (background or left out of training).

When we defined the purpose of our AI model, we listed that it should:

- **Accurately segment retinal layers** (e.g., ganglion cell layer, inner nuclear layer, outer nuclear layer, and retinal pigment epithelium).
- **Identify and quantify key retinal cell types**, including cell bodies of ganglion cells, bipolar cells (also Müller and Amacrine), and photoreceptors.
- **Exclude non-retinal tissues** from quantification.

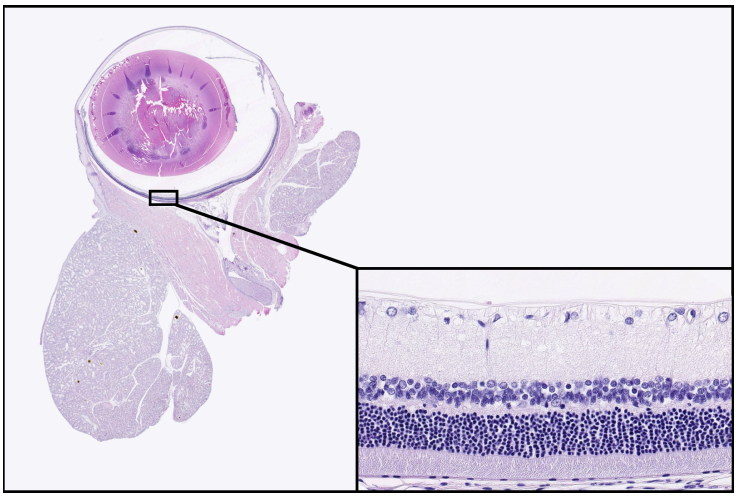

Histological section of typical HE-stained rat eye and extraocular tissues. Inset: high-power view of retinal layers and cell bodies.

Data collection for training

Sourcing histopathology images

To develop a comprehensive AI model, histological images need to be obtained.

The dataset should include high-quality histologic sections with consistent staining and section quality. In addition, it is helpful if the section orientation, staining, and scanner are consistent and reflect future data sets.

How many examples do I need?

- With limited variation in staining and samples (animal breed/species, age, etc.), **50-100 WSIs** are usually recommended to construct a basic model that is generalizable to your needs and will be suited to the analysis of new material.
- For testing and validation, the recommendation is to use approximately **10% of the number of cases** used in the training with any variability in tissues, staining, or scans included.

Preparing digital images

Histological slides must be digitized using high-resolution whole-slide scanners. Things to consider when preparing digital images:

- Consistent brightness and contrast minimize variability.
- Sections should ideally be free of artifacts such as tissue folds or uneven staining.
- Image format compatibility: Aiforia® Platform supports all main image formats.
- Uploading the images to the cloud: Aiforia® Platform provides several ways to upload images, such as web upload, software plugin, or automated image upload through software integration.

Annotation of features

Model design

The model is designed to differentiate the features we call classes. **A tissue area** is usually the first class added, which segments from the slide background. This speeds up analysis, as the background is ignored and won't be analyzed further. The next layer in the model tree design would segment **the retinal from all non-retinal tissues**, which won't be analyzed further. Essentially, this is carried on to segment **the different layers of the retina**.

In our example, we wanted to identify the following retinal layers:

- **Retinal Pigment Epithelium (RPE)** – Responsible for supporting photoreceptors
- **Photoreceptor Layer** – Contains rods and cones
- **Outer Nuclear Layer (ONL)** – Houses photoreceptor cell nuclei
- **Outer Plexiform Layer (OPL)** – Layer between photoreceptors and bipolar cells

- **Inner Nuclear Layer (INL)** – Contains bipolar, Müller, and Amacrine cell bodies
- **Inner Plexiform Layer (IPL)** – Layer between bipolar and ganglion cells
- **Ganglion Cell Layer (GCL)** – Ganglion cells transmit visual signals to the brain

And for retinal cells, we decided to identify the following cell types:

- **Ganglion cells** (large, located in the **GCL**)
- **Bipolar cells (also Müller, Amacrin, and others)** (in the **INL**, connecting photoreceptors to ganglion cells)
- **Photoreceptors** (small, located in the **ONL**)

Model structure

Areas to be segmented

➥ Tissue

➥ Retinal tissues vs non-retinal tissues

➥ Subclassification of retinal layers

As we would like to calculate the number and density of cell bodies within these segmented layers, we employ object detection as a class type.

Objects to be detected

- INL cell bodies (unless further subclassification of cell types is required)
- ONL cell bodies
- Ganglionic cells

The layer order can be structured to segment tissues into distinct areas, allowing object detection within the cellular retinal layers to be applied specifically in regions where cell body quantification is needed.

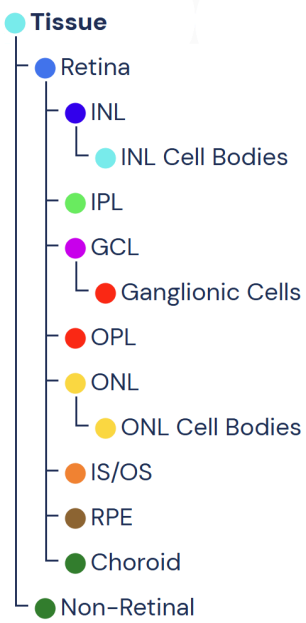

Supervised learning - manual annotation process

Using Aiforia® Create, researchers can manually annotate images by:

- Outlining representative examples of features.
- Labeling individual cells based on morphology and location.

Semi-supervised learning - using Annotation Assistant

Alternatively, after the model has undergone some training, you can utilize Aiforia's Annotation Assistant. This feature suggests potential examples of features and backgrounds, allowing the user to review and select them.

Area segmentation

Training regions are selected areas on the images that you use to teach the neural network (areas delineated with a black line). The neural network will collect data within these drawings. Small areas are good for consistency and training speed.

If areas inside the training region contain the feature you are looking for, you should label the area with an assigned class (represented by a colored line). Features that you want the AI to exclude (background) should be left without any class label.

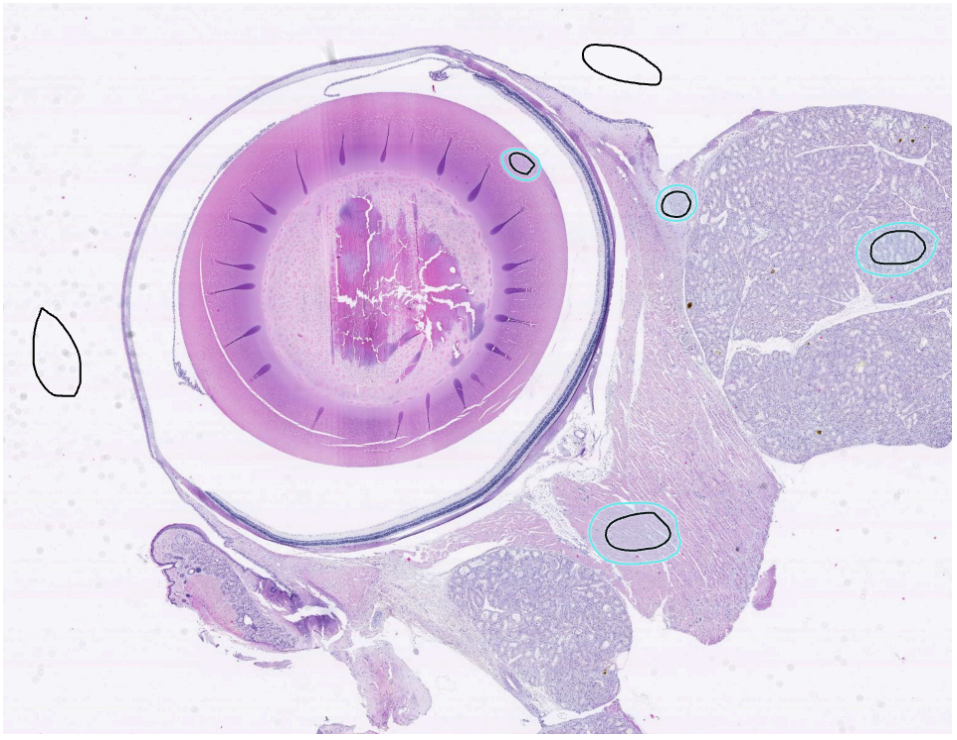

An example of annotations for tissue and non-tissue

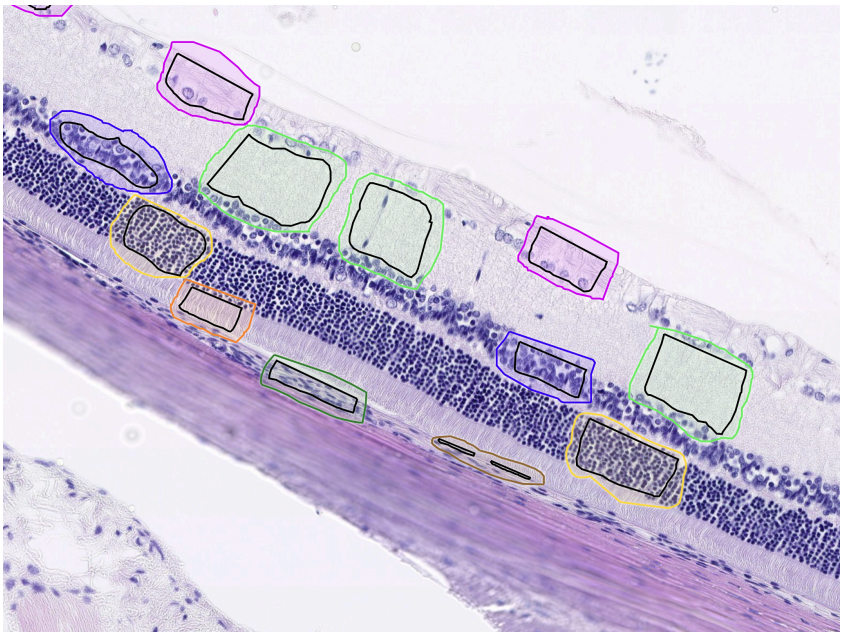

An example of annotations for retinal layer segmentation annotations. Only annotate good-quality examples and be precise.

Object annotations

For this feature type, each distinct object is marked with an adjustable circular annotation based on its size. The training area is again confined within the designated black-delineated region. In this example, ganglionic cell bodies are annotated with an object size of 14µm, while smaller sizes are assigned to the cell bodies in the ONL and INL.

Ganglionic cell body object annotations example, 14um. Other cell nuclei are excluded in this example.

How to optimize annotations for the best performance

Region

- Keep the annotation size manageable
- Label the most unambiguous examples first
- Annotate regions that look variable
- Be sure to define the background

Object

- In the beginning, choose the most unambiguous examples to label
- Be sure the object detector is centered
- Use only one object size for a single class
- Be sure to define the background

Model training

As we develop the model by adding annotations, we need to initiate frequent periodic training to enable the CNNs to learn the features. This process doesn't have to be completed all at once—you can focus on one feature at a time. This approach makes the process more manageable, allowing analysis results to be isolated from parent layers rather than splitting projects into separate features and later recombining them. However, if multiple annotators work simultaneously, the flexibility remains to accommodate different workflows.

The best approach is to add some annotations and then train, analyze, and review. This helps you see whether your annotations have improved the model quickly.

When you train a model, you define how many **iterations** you want to use. Iterations determine how much time you give the AI to learn the features; the longer the training time, the more it will look for examples. An iteration describes the number of times a batch of data passed through the algorithm. In the case of neural networks, that means the forward pass and backward pass. So, you complete an iteration every time you pass a batch of data through the neural network.

Example training workflow

- **Training 1 (Start):** Annotate ~20 annotations per class → Train → Use results to guide new annotations
- **Training 2:** Use Training 1 results to guide where new annotations are needed → Add ~20-50 new annotations per class → Run 2nd training
- **Training 3 & beyond:** Repeat the above step as many times as needed until the AI model performs at the expected level

Post-training analysis

After each training round, you can use the model to analyze the tissues for performance (are there areas not identified in the analysis or false positive or false negative results) and verify if your annotations agree with model analyses.

Example

The following images are examples of post-training analyses for each feature, as annotated above. In this example, there were 2 hours of training with 2,300 iterations, along with a few hours of annotations preceding.

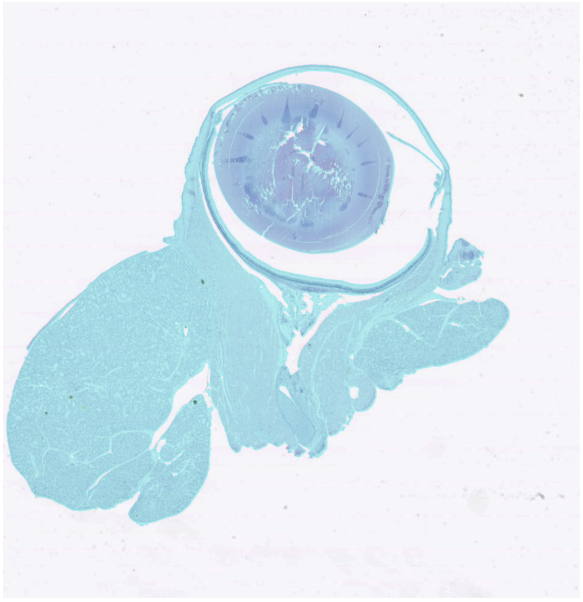

Tissue layer results

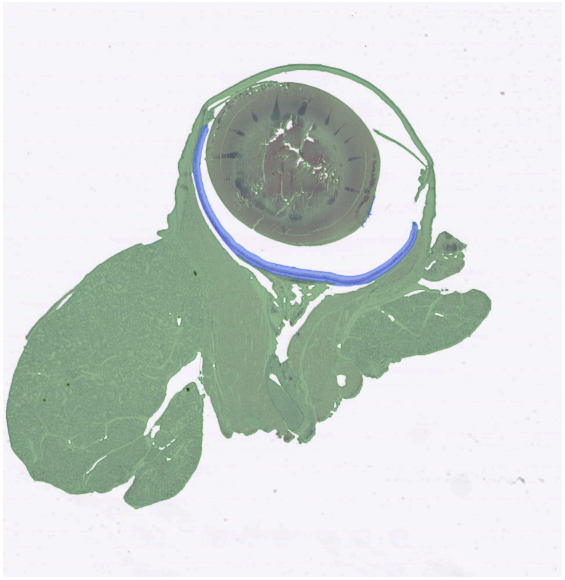

Retinal vs non-retinal layer results

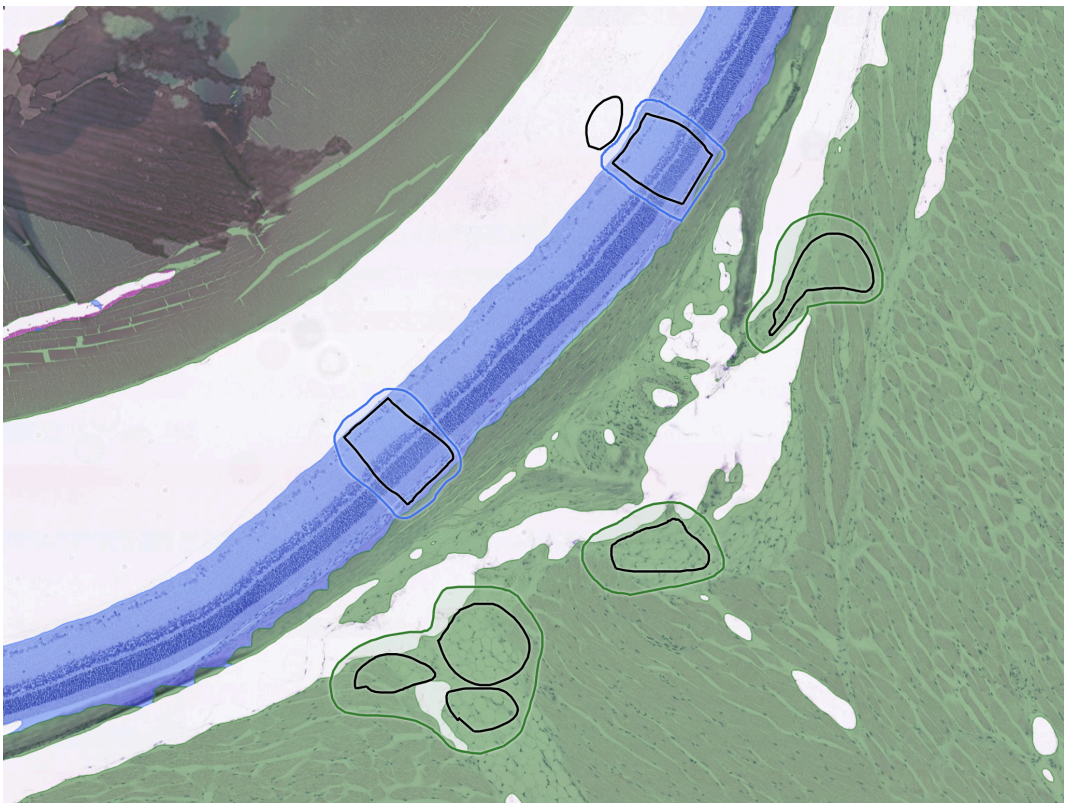

Retinal vs non-retinal layer results with the previous annotations included

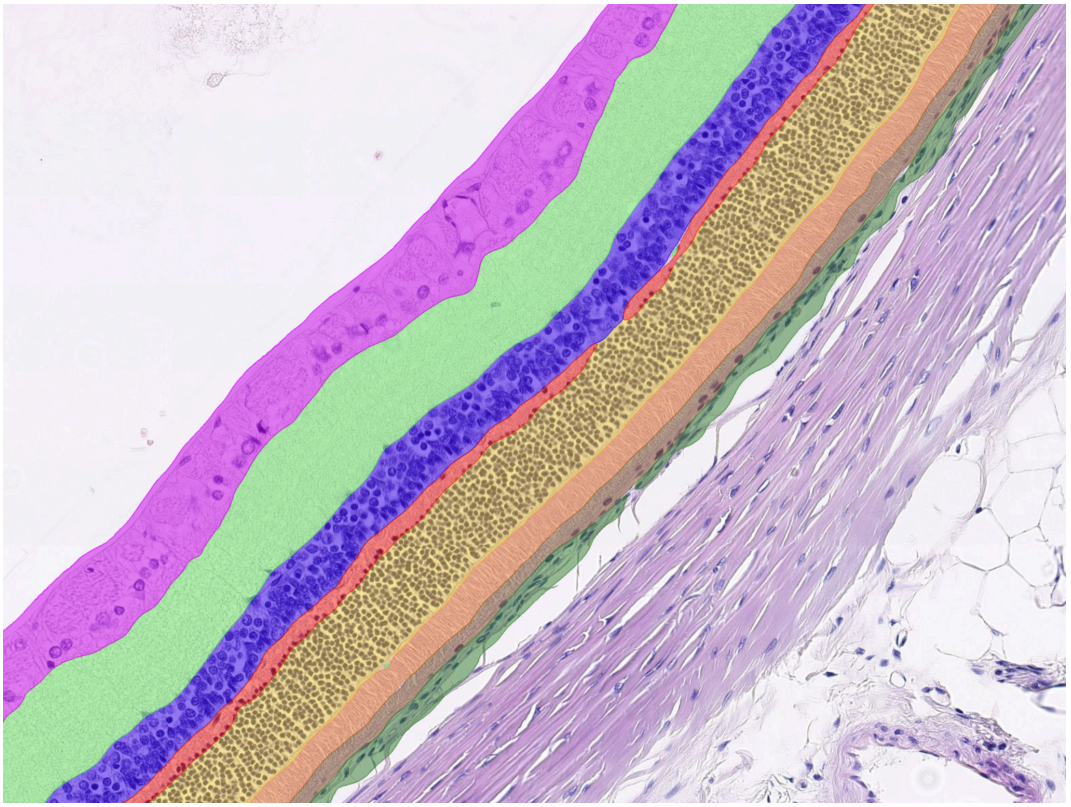

Retinal segmentation layer results

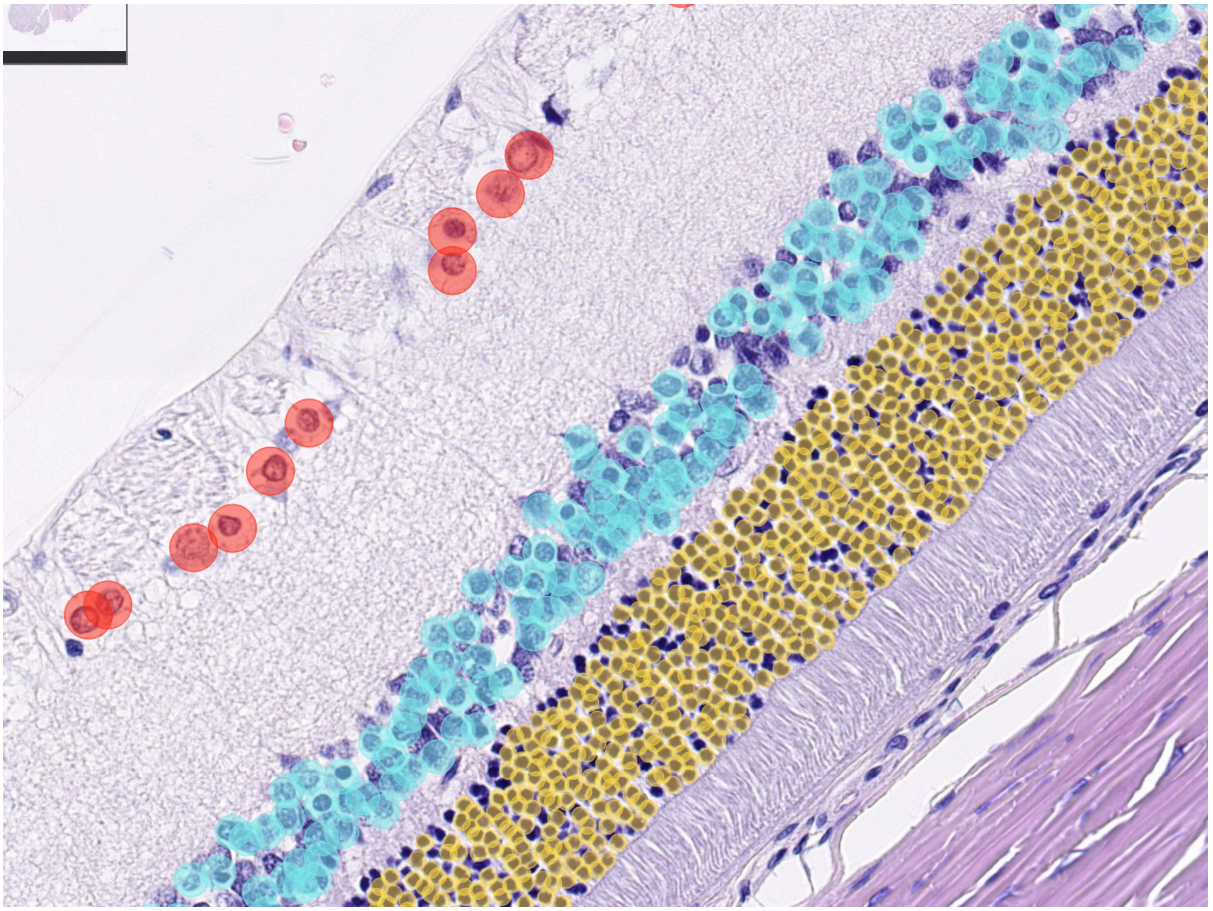

Cell detection through object recognition – results are incomplete, indicating the need for additional annotations or adjustments to training settings.

Further data collection, additional annotation rounds, and training on more slides are necessary to develop a more generalizable and robust model.

Validation process

After sufficient AI model performance is achieved, the AI model is usually validated against external human validators before being used to analyze future image datasets.

Visual validation

Steps: Run an image analysis on 5–10 slides that were never used during AI model training. Inspect the results.

- If the AI **model performs well**, you can use it for image analysis after further validation.
- If the AI **model fails**, add more images to your AI project, provide training annotations, and train a new AI model that will also generalize to the new images.

Using Aiforia's validation tool

In addition to visual validation, Aiforia's built-in validation tool can compare your AI model performance against blinded annotations of human validators. Essentially, you will invite experts to place validation annotations via the validation interface. The AI analysis results are compared against the validators' annotations to gain the error rate. If the model agrees 100% with the validation regions, the error rate will be 0.

The tool's cloud-based nature makes this collaboration even easier, as you can invite validators with a simple click of a button directly from your image, and validators can work remotely without registering to Aiforia® Create.

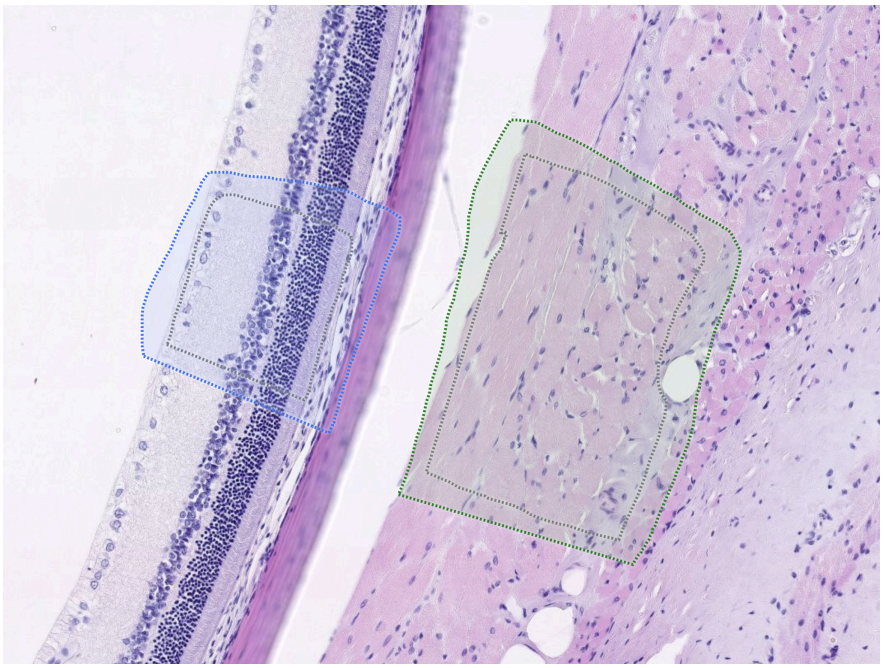

Example of hatched validation annotations (retina vs non-retinal)

This generates measurable performance metrics that can be analyzed and utilized. Validation results help assess the model's performance.

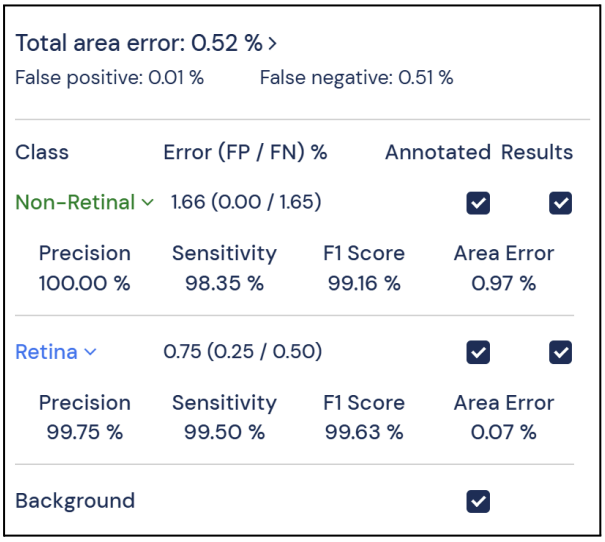

Ensuring model quality and future scalability

Continuous monitoring

- The AI model should be periodically tested with new samples to ensure consistency.
- Performance should be re-evaluated whenever new staining methods or improved imaging technologies are introduced (quality control).

Preparing for future expansion of an AI model

Once the AI model is created, it can be used as a basis to expand to other use cases flexibly.

In our example, we can elaborate on the identification of the normal retinal structures to include glaucoma-induced changes, such as:

- Thinning of retinal layers (notably the ganglion cell layer and inner plexiform layer)
- Ganglion cell loss (a hallmark of glaucoma progression) and loss of the INL and ONL cell bodies (decrease in density or area or both could be assessed)

This requires analyzing pathologically affected tissues (i.e., atrophic retinas) with the current validated model. Then, if needed, the model can be fine-tuned with further annotations to include the new pathology in any needed new data.

Conclusion

Developing an AI model with Aiforia® Create - in this example, to quantify the normal retina - provides a strong foundation for future diagnosis and research on retinal diseases, including disease processes like glaucoma-induced atrophy. By carefully collecting and annotating high-quality histopathological images, training a deep-learning model, and implementing quality control measures, researchers can enhance the accuracy and efficiency of tissue analysis.

Learn more about Aiforia® Create: https://www.aiforia.com/aiforia-create