



CLINICAL RESEARCH STUDY BINDER

Prospective Evaluation of Retinal Imaging and Clinical Data for Prediction of Ocular and Systemic Disease

IRB	Sterling IRB • ID 16201
PROTOCOL	v1.4 • 21 April 2026
CONSENT IN FORCE	v1.0 — IRB-approved 27 April 2026
PRINCIPAL INVESTIGATOR	Sai Chavala, MD
SPONSOR	SpectOcular LLC
DESIGN	Prospective • observational • minimal risk
SITE	

Participation involves no additional testing.

It involves permission to release medical and imaging data so SpectOcular can develop new tests for whole-body health. Routine imaging obtained during standard eye care may be used for research. Study data are coded or de-identified to the extent practicable.

Binder issued 19 August 2026 • Controlled document

Consent only on the form in Tab 10. See Tab 02 — Document Register and Version Control.

SpectOcular LLC • 11615 Forest Central Drive, Suite 308, Dallas, Texas 75243
schavala@gmail.com • 214-233-6170 • spectocularhealth.com/clinical-research

SpectOcular Site Binder

Site	
Site number	01
Protocol	Sterling IRB #16201 v1.4 — observational, minimal risk
Issued	19 August 2026
Contact	Sai Chavala, MD · schavala@gmail.com

The three rules everything else follows

- 1. Participation never changes routine care.** No imaging is performed solely for research. Only images acquired as part of routine care are included.
- 2. Written research consent and HIPAA authorization are signed before research inclusion.** No exceptions, no retrospective consent.
- 3. A patient is enrolled once.** A follow-up visit is not a new enrollment — see below.

What is in this binder

	Document	Use it for
01	Site Operations Manual	Start here. This is the document you work from every day — both service lines, research, imaging, export, troubleshooting, and the checklist for the day research opens
02	Topcon Maestro Imaging Protocol	Acquisition standards — dilation, minimum image set, quality
03	SOP-003 — OCT Export and Imaging Data Handling	The controlled export procedure. Where anything disagrees with it, SOP-003 governs
04	Patient Research Information (English · Spanish)	Draft — pending IRB approval. Not for patient distribution.
05	Retinal Imaging Modalities	Explaining the imaging to a patient who asks
06	Approved IRB materials — ICF template, medical release	Consent packets. Approved version only
07	Site Staff Workflow Quick Guide — SUPERSEDED	Reference only. Written before the current system was built; it describes folders and an ID register that do not exist. Do not train from it

What is live, and what is not

	From 20 August
SpectOcular Patient Services	LIVE
Retina IQ consults	LIVE
Research enrollment	NOT LIVE. Do not consent anyone for research until you receive written confirmation that IRB approval has been granted — expected the week of 24 August

Manual **Part 13** is the checklist for the day that confirmation arrives. Work through it before enrolling anyone.

A follow-up visit is not a new enrollment

This is the single easiest thing to get wrong, and the hardest to undo.

A patient keeps **one Subject ID for the life of the study**. When they come back in three months, you do not fill in the enrollment form again and you do not generate a new Subject ID.

If a returning patient is enrolled a second time they get a second Subject ID and a second folder, and their follow-up scan can no longer be connected to their first one. The study exists to measure *change over time*, so that connection is the whole point. It cannot be repaired later without going back through identifiable records to work out that the two were the same person.

For a repeat visit: look the patient up, use their existing Subject ID, and file the new imaging under it.

What this site's imaging is

Topcon **color fundus photography and structural OCT**. This site does not have OCT-A.

That is a complete submission for this site. Nobody should be chasing OCT-A images here, and no one from SpectOcular should query their absence. If that ever happens, say so and it will be corrected at our end.

What the site never does

- Never perform imaging solely for research
 - Never de-identify anything — SpectOcular does that centrally
 - Never copy or synchronise the live Topcon database
 - Never alter files inside an exported package
 - Never put a patient name, MRN or date of birth in a folder or file name
 - Never route imaging to Retina IQ unless the enrollment record says the patient consented to it
-

If something goes wrong

Situation	What to do
The form errors or will not save	Stop. Do not keep re-submitting — that can create duplicate subject records. Call Sai
You think a patient may already be enrolled	Use the existing Subject ID. Never create a second
The export produces nothing, or looks wrong	Do not upload it. Write down what happened — the exam can be re-exported later
A patient asks a research question you cannot answer	Route it to the investigator or coordinator. Please do not improvise an answer
Anything concerning a patient's safety	Normal clinical pathway, immediately. The study never comes before care

Questions are always welcome, including after the fact. A question asked late is much cheaper than a record corrected late.

Document Register and Version Control

Site	_____ · Site # _____
Binder issued	19 August 2026
Protocol	Sterling IRB #16201 v1.4 — observational, minimal risk
Issued by	Sai Chavala, MD · Principal Investigator · schavala@gmail.com

The one rule

Consent only on the version in this binder.

If you believe a newer consent form exists — someone mentions it, emails it, or you see one elsewhere — **stop and call before consenting anyone.** A newer version is only in force once it has been issued into this binder through the procedure below.

This matters because an amendment is currently pending with the IRB. **Pending is not approved.** Until a replacement is issued to you in writing, the version in Section 06 is the only one that may be used.

Register — what is in force today

§	Document	Version	Status
00	Binder cover and contents	19 Aug 2026	Current
00b	This register	19 Aug 2026	Current
01	Site Operations Manual	19 Aug 2026	Work from this daily
02	Topcon Maestro Imaging Protocol	—	Current
03	SOP-003 — OCT Export and Imaging Data Handling	v1.0	Current. Where anything disagrees with it, SOP-003 governs
04	Patient Research Information (English · Spanish)	—	Draft — pending IRB approval. Not for patient distribution.
05	Retinal Imaging Modalities	—	Reference
06	IRB-approved materials	See below	Approved versions only

Section 06 — the approved set

Document	Status
16201.FINAL_IRB_PROTOCOL_v1.4.pdf	Approved protocol. The study you are participating in
16201.ICF.TEMPLATE.APPROVED	The consent form in force. Use this one.
16201.Study Material(s) - APPROVED.pdf	Approved participant-facing materials
Medical_Release_Form_v3 red line (1).APPROVED.pdf	Approved medical release

An amendment is pending — Spanish-language support, consent administration by clinical technicians under supervision, and additional data sources. **None of it is in force yet.** When it is approved you will receive a replacement pack.

Replacement procedure — how a new version reaches this binder

- 1. SpectOcular issues it.** Replacement pages come from the Principal Investigator, with a dated cover note saying what is being replaced.
- 2. Remove and destroy the superseded copy,** or mark it clearly and file it outside the working binder. A superseded consent form left in a binder will eventually be used.
- 3. File the replacement** in the same section.

4. Sign the acknowledgement below and return a copy to SpectOcular.
5. Nothing is consented on a version that did not arrive this way — not an emailed draft, not a printout from another source.

Why the paperwork: for every participant, the study must be able to state which consent version was in force on the date they signed. That record is built from this register and these acknowledgements. Without it, the enrollment cannot be evidenced later — no matter how correctly it was actually done.

Acknowledgement of receipt

Sign on issue, and again on each replacement.

Issue / replacement	Date issued	Received by (print)	Signature	Date signed
Initial binder — 19 Aug 2026				

If something goes wrong

Stop and call — do not improvise. Sai Chavala, MD · schavala@gmail.com · 214-233-6170

Call before proceeding if: a signature or date is missing · the patient has questions about risks, research use, or withdrawal · imaging would be done only for research · the wrong patient or eye may have been selected · an export contains another patient's information · you are unsure which consent version to use.

Nothing is lost by pausing an enrolment. A great deal is lost by an enrolment that cannot be evidenced.

SpectOcular Site Operations Manual

Site	Site # ____ — _____
Version	1.0 — 19 August 2026
Supersedes	Site Staff Workflow Quick Guide (retained for reference; where they disagree, this manual governs)
Protocol	Sterling IRB #16201 — observational, minimal risk
Support	Sai Chavala, MD · schavala@gmail.com

Read this first

This manual is written to be used without anyone from SpectOcular present. Training happens once, on 20 August. After that, this document is your reference — including for research enrollment, which you will start on your own once IRB approval arrives.

If something here does not match what you see on screen, **stop and ask**. Do not improvise a workaround. A question costs a phone call; a workaround costs a record that has to be unpicked later.

PART 0 — WHAT IS LIVE, AND WHAT IS NOT

Service	Status from 20 Aug	Governed by
SpectOcular Patient Services	LIVE	Practice Patient Services consent + BAA
Retina IQ consults	LIVE	Patient Services consent + the patient's separate Retina IQ permission
Research enrollment	NOT LIVE — do not consent anyone for research	Sterling IRB #16201, pending approval

Research enrollment is expected to open the week of 24 August. You will be told in writing. Until that message arrives:

- Do not sign anyone onto the research consent
- Do not use the research enrollment form except for practice with made-up data
- Do not export imaging into the research folders

Everything else in this manual is live and in use from day one. **Part 13** is the checklist you work through on the day approval arrives.

PART 1 — THE THREE PERMISSIONS

The most common source of error at a new site is treating these as one thing. They are three, and a patient can agree to any combination.

	What it permits	Required for
1. Practice Patient Services consent	SpectOcular holds the patient's record and supports their care	Everything. Nothing happens without this
2. Retina IQ	The patient's images are sent for a specialist consult to their optometrist	Retina IQ only
3. Research consent + HIPAA authorization	The patient's coded data supports the research study	Research enrollment only

Three rules that never bend:

1. **Participation never changes routine care.** No imaging is ever performed solely for research or for Retina IQ. Only images acquired as part of the patient's routine care are used.
2. **Signatures come before inclusion.** Every consent is signed and dated *before* the thing it permits happens. There is no retrospective consent.
3. **Retina IQ is a consult to the doctor, not a report to the patient.** It supports the supervising eyecare professional's decision. It is not a diagnosis delivered to the patient, and results are not given to the patient directly by SpectOcular.

PART 2 — IDENTIFIERS — READ THIS TWICE

Three different numbers travel with a patient. Confusing them is the single most damaging mistake available in this system.

Identifier	Looks like	Who assigns it	Contains a name?
Clinic patient ID / MRN	Whatever your practice uses	Your practice	Yes, in your system
SpectOcular Patient Number	01-0001 — sequential, four digits	The system, automatically	No
Research Subject ID	01-K7M2Q — random, five characters	The system, automatically	No

The Subject ID is random on purpose. It is not sequential, and it encodes nothing — not diagnosis, not sex, not the order people enrolled. The approved consent form tells patients exactly this. The characters I, L, O, 0 and 1 are never used, so there is no ambiguity when a code is written by hand.

⚠ One patient, one number, for the life of the study

A follow-up visit is never a new enrollment.

When a patient returns in three months, you do not fill in the enrollment form again and you do not generate a new Subject ID. You look up the one they already have.

If a returning patient is enrolled a second time they receive a second Subject ID and a second folder, and their follow-up imaging can no longer be connected to their first visit. The entire point of the study is measuring *change over time*, so that link is the asset. Once broken it cannot be repaired without going back into identifiable records to prove the two records were the same person.

The system now checks for this and will hand back the existing ID rather than creating a new one. **Do not rely on it.** Look the patient up first — the check is a safety net, not the procedure.

PART 3 — SPECTOCULAR PATIENT SERVICES • LIVE

Follow this for every patient joining SpectOcular Patient Services.

Before the patient arrives

4. Confirm the imaging planned is **clinically indicated as routine care**.
5. Have the Practice Patient Services consent and the intake worksheet ready.

With the patient

6. Give the patient time and privacy to read. Route clinical questions to the supervising eyecare professional.
7. Obtain the **Practice Patient Services consent** — every signature, every date.
8. Ask about **Retina IQ** and record the answer. Yes or no, it gets recorded.
9. Collect intake: demographics and preferred contact method; PCP name, clinic, address, phone and fax; preferred pharmacy; relevant specialists; diagnoses, medications, allergies, smoking history; cardiovascular, diabetes, kidney and neurologic history; eye history including surgery, laser, injections and prior imaging.

Registering the patient

10. Open the SpectOcular site and register the patient. The system assigns the next **SpectOcular Patient Number** (01-0002, 01-0003, ...) and creates their folder.
11. **Before the patient leaves:** re-check every signature and date, confirm PCP and pharmacy details are usable, and confirm the history is legible.

Returning patients: search first. If they already have a SpectOcular Patient Number, use it. Do not register them again.

PART 4 — RETINA IQ • LIVE

Retina IQ is a specialist consult on the patient's imaging, delivered to the supervising eyecare professional.

Only for patients who have agreed to it. The recorded answer governs — not anyone's recollection of the conversation. If the record says No, the imaging supports their care and nothing else, however interesting the scan.

To submit a request you will need:

- The patient's SpectOcular Patient Number and clinic patient ID
- Date of the scan and which modalities were acquired
- The requesting professional and their contact details
- PCP details
- Relevant medical and ocular history
- **The clinical question.** The most useful field on the form — "rule out macular oedema" produces a better consult than "please review"

What happens next: the request and its imaging go to the interpreting physician; the report returns to the practice and is filed against the patient. Turnaround and billing follow the executed agreement.

What Retina IQ is not: it is not an emergency pathway. Anything requiring urgent attention goes down your normal clinical route immediately, and does not wait on a consult.

PART 5 — RESEARCH ENROLLMENT • NOT LIVE UNTIL IRB APPROVAL

Do not begin any part of this section until you receive written confirmation that IRB approval has been granted. Practising with invented patient details is encouraged. Consenting a real patient is not.

What the study is: an observational, minimal-risk study. It collects information and images that are already being generated by routine care. It changes nothing about how a patient is examined or treated.

Consent, in order

12. Confirm the patient already has the Practice Patient Services consent.
13. Give them the **Patient Research Information** sheet (English or Spanish) and time to read it.
14. Obtain the **research consent and HIPAA authorization** — the current approved version only. If you are unsure which version is current, ask before signing anyone.
15. Route every research question to the investigator. Please do not answer research questions from your own knowledge, however reasonable the answer seems.

Enrolling

16. Open the research enrollment form and complete it with the patient's details.
17. The form generates the **Subject ID. Write it on the consent packet** — this is the only durable link between the paper and the record.
18. Complete the medical-history and covariate fields.

"Don't know" is a real answer. Three of the history questions offer *Don't know*. Use it when that is the truth. Never record *No* to mean *we did not ask* — leaving it blank is correct when unknown. A false *No* is worse than a blank, because a blank is honest about being unknown and a *No* is not.

19. Confirm the record saved before the patient leaves.

Enrollment is once. See Part 2. For a returning subject you use their existing Subject ID.

PART 6 — IMAGING

This site: Topcon Maestro — colour fundus photography and structural OCT. This site does not have OCT-A.

That is a complete submission here. Nobody should chase OCT-A images at this site, and no one from SpectOcular should query their absence. If that ever happens, say so and it will be corrected at our end.

Acquisition standard (full detail in the Topcon Maestro Imaging Protocol, document 02):

- Imaging after standard-of-care dilation unless contraindicated, declined, or not clinically appropriate. Target pupil ≥ 4.0 mm where achievable; document if smaller
- Minimum set per eye: macula-centred fundus photo, disc-centred fundus photo, macular OCT cube, optic nerve / RNFL OCT cube
- Same device and software version wherever possible
- Retain poor-quality and ungradable scans **with their labels** — do not silently delete and repeat. A known-ungradable image is data; a missing image is a hole

And again, because it is the rule that gets broken: nothing is acquired solely for research or for Retina IQ.

PART 7 — EXPORT AND UPLOAD

Controlled procedure: **SOP-003** (document 03). This is a summary — SOP-003 governs.

Confirm first: the Topcon is running the **research installation**. A clinical installation produces exports at the wrong regulatory tier which cannot be used, and the error is silent until someone checks. If you are not certain which is installed, ask before exporting.

Export

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -allDcm -json -debug
```

then a second pass for the layer-thickness files, which are the study's primary measurement:

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -thicknessMap -debug
```

In the GUI, tick **New Folder for Each FDA File**. Without it, a batch writes everything into one directory and the exams cannot be separated afterwards.

Name the output folder with the Subject ID and the exam date, and nothing else:

```
03-HN7JZ_2026-08-18
```

Never put a patient name, MRN or date of birth in a research folder or file name. Folder names travel into backups, sync logs and screenshots, none of which carry the protections the contents get.

Upload to that subject's own imaging folder:

01_Sites / <Site folder> / Subjects / <SubjectID> / 05_Retinal_Imaging_ZIP

Then confirm the upload was recorded. An export sitting in Drive with no custody record is not considered filed (SOP-003 §6.4). If you do not see it recorded, say so — do not upload it a second time.

Let each export finish completely before moving or zipping it. Never copy or synchronise the live Topcon database, and never alter files inside an exported package.

PART 8 — FOLDER NAMES: WHY TWO RULES THAT LOOK CONTRADICTIONARY ARE BOTH RIGHT

You will notice something that looks like an inconsistency, and it is worth understanding rather than working around.

	Folder name	Why
Patient Services	Lastname_Firstname_MMDDYYYY — name and date of birth	Identified record, held under the BAA. It is their chart; it is supposed to be findable by name
Research imaging	03-HN7JZ_2026-08-18 — Subject ID and date only	Coded record. The whole point is that it carries no identity

Both are correct. The rule is not "never write a name anywhere" — it is "**never write a name in the research side.**" If you are working in Subjects, no names. If you are working in SpectOcular_Patient_Services, names are how the folder is found.

Research subject folders are created for you: 01_Intake, 02_Signed_Consents, 03_Medical_Records, 04_Labs, 05_Retinal_Imaging_ZIP, 06_Retina_IQ_Reports, 07_Invoices, 08_Audit_Log. Do not create your own alongside them.

PART 9 — WHAT THE SITE NEVER DOES

- Never perform imaging solely for research or for Retina IQ
- Never de-identify anything — SpectOcular does that centrally. You have no second command to run
- Never copy or synchronise the live Topcon database
- Never alter files inside an exported package
- Never put a name, MRN or date of birth in a research folder or file name
- Never route imaging to Retina IQ unless the record says the patient agreed
- Never re-enroll a returning patient

- Never answer a research question on SpectOcular's behalf — route it

PART 10 — TROUBLESHOOTING WITHOUT AN EXPERT PRESENT

What you see	What to do
The form errors, or will not save	Stop. Do not keep re-submitting — repeated attempts can create duplicate records. Note the time and any message, and call
You think the patient may already be registered	Search first. Use the existing number. Never create a second
A Subject ID was generated but the patient was already enrolled	Use the original ID. Report the duplicate so it can be cleaned up before data accumulates against it
The export produced nothing, or far fewer files than usual	Do not upload it. Note what happened. The exam can be re-exported later — nothing is lost
The export ran but you are unsure of the tier	Do not upload. Ask. An export at the wrong tier is unusable and better caught now
The upload does not appear as recorded	Do not upload again. Report it — a duplicate is harder to resolve than a missing record
Patient asks a research question you cannot answer	Route to the investigator. Do not improvise
Patient withdraws, or asks to be removed	Stop all further collection for them immediately, and report the same day
Anything concerning patient safety	Normal clinical pathway, immediately. The study never takes precedence over care

When in doubt, the safe action is always: stop, write down what you saw, and ask. Nothing in this study is urgent enough to justify a guess.

PART 11 — SUPPORT

Sai Chavala, MD	schavala@gmail.com
Central account	admin@spectocularhealth.com
Best for	Anything you are unsure about — including after the fact

Please report near-misses, not just errors. If something nearly went wrong, that is the cheapest possible signal that a step needs changing. Nobody is in trouble for reporting one.

PART 12 — SITE ACCESS

Site address	☐ _____ (to be completed at training)
Access code	SPECTIQ2026
Who may use it	Only staff on the delegation log in Appendix A

Do not share the access code outside the delegated staff, and do not save it where a patient could see it.

PART 13 — ACTIVATION: THE DAY IRB APPROVAL ARRIVES

Work through this in order. It should take under an hour.

- [] 1. Written confirmation received from SpectOcular that research enrollment is open. **Do not start on a verbal or an assumption.**
- [] 2. Confirm which ICF version is approved, and that the packets you hold match it. **If they do not match, stop — do not consent anyone on a superseded form**
- [] 3. Destroy or clearly mark any superseded consent packets so they cannot be picked up by mistake
- [] 4. Confirm every staff member who will consent or enroll is on the delegation log (Appendix A)
- [] 5. Run one enrollment with invented details end to end. Confirm the Subject ID appears and the record saves. Report the practice record so it can be removed
- [] 6. Confirm the Topcon research installation is the one in use
- [] 7. Enroll the first real subject, and **call the same day** to confirm the record arrived intact
- [] 8. For the first three subjects, check each one's record before the patient leaves

First week: reconcile the number of consent packets against the number of records. They must match. A mismatch found in week one is a conversation; found in month three it is an audit finding.

APPENDIX A — TRAINING RECORD AND DELEGATION LOG

Signed at training on 20 August 2026. Required by SOP-003 §11 — only staff listed here may perform these tasks.

Name	Role	Trained on	Signature	Date
		Patient Services · Retina IQ · Research · Imaging · Export		
		Patient Services · Retina IQ · Research · Imaging · Export		
		Patient Services · Retina IQ · Research · Imaging · Export		
		Patient Services · Retina IQ · Research · Imaging · Export		

Trainer

Name	_____
Signature	_____ Date _____

Confirmed at training — tick each once demonstrated, not merely described:

- ☐ Each person can state the three permissions and which one governs what
- ☐ Each person can explain why a follow-up visit is not a new enrollment
- ☐ Each person has registered a patient and seen the number assigned
- ☐ Each person has submitted a Retina IQ request
- ☐ Each person has run an export and named the folder correctly
- ☐ Each person knows the research gate is closed and what opens it
- ☐ Each person knows what to do when something fails — stop, record, call

Training was performed on invented data only. No real patient record was created for training purposes.

APPENDIX B — BENCH CARD

Print this page and keep it at the imaging station.

EVERY PATIENT

20. Imaging clinically indicated? → If no, stop
21. Practice Patient Services consent signed and dated

- 22. Retina IQ — asked, answered, recorded
- 23. Intake complete: PCP · pharmacy · contact · history
- 24. Register → SpectOcular Patient Number assigned
- 25. Re-check signatures and dates **before they leave**

RESEARCH — only after approval

- 26. Approved research consent + HIPAA signed
- 27. Enroll → Subject ID generated → **write it on the packet**
- 28. Returning subject? **Use their existing ID**

IMAGING

Dilate · macula fundus · disc fundus · macular OCT · disc OCT

No OCT-A at this site — expected, not missing

Nothing acquired solely for research

EXPORT

Research installation · -allDcm -json -debug · second pass -thicknessMap -debug

Folder name: SubjectID_YYYY-MM-DD — **no names, no MRN, no DOB**

Upload to that subject's 05_Retinal_Imaging_ZIP · confirm it was recorded

STOP AND CALL IF

Form errors · patient may be already enrolled · export empty or wrong tier · upload not recorded · patient withdraws · anything you are unsure about

Sai Chavala, MD · schavala@gmail.com

APPENDIX C — GLOSSARY

Term	Meaning
Coded	Identifiers replaced by a Subject ID, with the key held separately by SpectOcular. Still protected information
De-identified	Identifiers removed entirely. Done centrally by SpectOcular — never at the site
Subject ID	01-K7M2Q. Random. Research only
SpectOcular Patient Number	01-0001. Sequential. Patient Services
Linkage key	The restricted record connecting a Subject ID to a person. Access is limited and logged
DICOM	The standard imaging format the Topcon exports
.fda	Topcon's native file. The export tool reads it and produces DICOM
Regulatory tier	Which Topcon installation produced an export. This study requires the research tier
Gradable / ungradable	Whether an image can be assessed. Ungradable images are kept and labelled, never deleted

Topcon Maestro Retinal Imaging Protocol for AI/ML Dataset Creation

SpectOcular, LLC | Draft SOP / Protocol Insert

Prepared for standardized retinal image acquisition, quality control, metadata capture, and downstream AI/ML analysis.

Protocol Element	Standard
Device	Topcon Maestro2, Maestro3, or substantially equivalent Topcon Maestro platform
Dilation	All imaging should be performed after standard-of-care dilation unless contraindicated, declined, or otherwise not clinically appropriate
Minimum image set per eye	Macula-centered fundus photo; disc-centered fundus photo; macular OCT cube; optic nerve/RNFL OCT cube
Preferred expanded image set	3D Wide 12 x 9 mm OCT and macula OCT-A if available; disc OCT-A optional
Data preservation	Native device format and/or DICOM preferred; high-resolution image files retained; PDF reports secondary only
Quality control	Image quality grading, repeat imaging rules, and retention of poor/gradable/ungradable scans with labels
AI/ML handling	Patient-level dataset splitting, standardized metadata, de-identification/coding, and secure storage

Purpose

The purpose of this imaging protocol is to obtain standardized, high-quality, multimodal retinal images suitable for downstream artificial intelligence and machine learning analysis. The protocol is designed to maximize image gradability, reproducibility, and consistency across subjects, imaging sessions, technicians, and clinical sites.

Retinal imaging data will be used for research and AI/ML model development, including potential analysis of ocular disease, retinal structure, retinal vascular features, and systemic disease risk. The imaging protocol is intended to support creation of an AI-ready dataset by standardizing image acquisition, documenting image quality, and preserving device-generated imaging metadata.

Large ophthalmic AI datasets require consistent image acquisition, quality control, and separation of training, validation, and testing datasets. Technical reviews of deep learning in ophthalmology emphasize the importance of preprocessing, gradability assessment, standardized datasets, and independent testing to ensure model generalizability. The AI-READI manual similarly describes a multimodal retinal imaging protocol designed to generate an AI-ready dataset and includes OCT, OCT-A, fundus photography, standardized retinal imaging checklists, and secure data transfer workflows.

Required Device

Primary imaging platform: Topcon Maestro2 or Topcon Maestro3 OCT system, or substantially equivalent Topcon Maestro platform capable of color fundus photography, spectral-domain OCT, and OCT angiography if available.

All imaging should be performed using the same device model and software version whenever possible. If multiple devices or software versions are used, the device model, serial number, software version, scan protocol, operator, site, and date/time of acquisition must be recorded as metadata.

Required Dilation Standard

All imaging should be performed after standard-of-care pupillary dilation unless dilation is medically contraindicated or declined by the patient.

Dilation should be performed according to the site's routine clinical standard of care. Imaging should begin only after adequate dilation is achieved. For SpectOcular, the preferred standard is bilateral dilation before Topcon Maestro imaging.

Minimum target pupil size: ≥ 4.0 mm if achievable, with documentation if smaller. If dilation is inadequate, the technician should wait additional time, recheck fixation and alignment, and repeat acquisition when feasible.

Acceptable reasons for nondilated or partially dilated imaging should be documented, including narrow angle concern, allergy, patient refusal, physician decision, or other clinical contraindication.

Imaging Sequence

Images should be acquired in a consistent order for each subject to reduce protocol variability. Before acquisition, confirm the patient or coded study ID, laterality, date/time of imaging, technician/operator ID, device model, device serial number, software version, dilation status, approximate pupil size, lens status if known, presence of major media opacity if known, and whether artificial tears were used before imaging.

The participant should be positioned comfortably with forehead and chin stabilized. The patient should be instructed to blink before each scan and then keep the eye open during acquisition.

Minimum Required Image Set Per Eye

At a minimum, the following should be obtained for both eyes **after standard-of-care dilation**.

Color Fundus Photography

- Acquire high-quality color fundus photographs for each eye.
- Required views per eye: macula-centered color fundus photo and optic disc-centered color fundus photo.
- Preferred field: use the Maestro's standard color fundus field, typically posterior pole coverage. If a single fundus photo captures both disc and macula clearly, it may be retained, but separate macula- and disc-centered images are preferred for AI/ML analysis.
- Quality requirements: sharp focus on retinal vessels; adequate illumination without significant underexposure or overexposure; macula visible and centered for macula image; optic nerve visible and centered for disc image; minimal eyelid/lash artifact; no major blink artifact; no severe motion artifact; no major defocus or haze obscuring vascular arcades.

Macular OCT Cube

- Acquire a macular OCT volume scan for each eye.

- Recommended Maestro scan: 3D Macula 6 x 6 mm OCT cube using the highest practical scan density available within routine workflow; prefer a protocol equivalent to 512 x 128 or comparable high-density cube when available.
- Required quality elements: fovea centered; full retinal thickness visible; no major motion artifact; no scan truncation; signal strength/quality score documented; segmentation lines retained for review but raw B-scan data preserved; repeat scan if the fovea is decentered, signal is poor, or motion artifact prevents layer analysis.
- Macular OCT is critical because retinal layer thickness and OCT-derived structural features have been linked to ocular and systemic health in large cohort analyses.

Optic Nerve / RNFL OCT

- Acquire optic nerve head and retinal nerve fiber layer imaging for each eye.
- Recommended Maestro scan: 3D Disc / optic nerve cube, including peripapillary RNFL analysis if available, with the optic nerve centered.
- Required quality elements: optic disc centered; circumpapillary scan not truncated; RNFL segmentation present and reviewable; signal strength documented; no major motion artifact; repeat if segmentation failure or decentration is present.
- This imaging is important for glaucoma-related model development and for systemic biomarker analysis involving inner retinal layers, RNFL, ganglion cell complex, and neurodegenerative risk features.

OCT Angiography, If Available

- If the Topcon Maestro unit has OCT-A capability, acquire OCT-A for each eye.
- Recommended minimum OCT-A scans: macula OCT-A 6 x 6 mm; optional but preferred optic nerve/disc OCT-A 4.5 x 4.5 mm or 6 x 6 mm, if available.
- Quality requirements: minimal motion artifact; acceptable signal strength; no major segmentation failure; superficial and deep capillary plexus outputs preserved if available; vessel density maps and en face images exported when available; raw volume data retained when possible.

Preferred Expanded Image Set Per Eye

For a stronger AI/ML dataset, the preferred SpectOcular standard should include the following:

Modality	Required?	Scan/View
Color fundus	Yes	Macula-centered
Color fundus	Yes	Disc-centered
OCT	Yes	3D Macula 6 x 6 mm
OCT	Yes	3D Disc / RNFL cube
OCT	Preferred	3D Wide 12 x 9 mm posterior pole
OCT-A	Preferred if available	Macula 6 x 6 mm
OCT-A	Optional	Disc OCT-A
Repeat scans	Preferred	Repeat one high-quality macula OCT and OCT-A when feasible for test-retest/reproducibility

The 3D Wide 12 x 9 mm scan is particularly attractive for AI/ML because it captures broad posterior pole anatomy in a single standardized volume, which may help future models evaluate macula, nerve, posterior pole contour, staphyloma/myopia features, vascular shadows, and structural disease patterns.

Repeat Imaging Rules

Repeat imaging should be performed during the same visit when feasible if any of the following are present:

- Poor signal strength
- Fovea or optic nerve decentered
- Motion artifact affecting central macula or disc
- Blink artifact
- Eyelid or lash obstruction
- Segmentation failure
- Image truncation
- Poor focus
- Media haze severe enough to obscure vascular or retinal layer detail
- OCT-A projection/motion artifact limiting vessel analysis

For each eye and modality, the best-quality image should be designated as the primary research image, but repeat images should also be retained because they may be useful for AI robustness, quality modeling, and reproducibility analysis.

Image Quality Grading

Each image or scan should receive a quality designation. For AI/ML purposes, poor or ungradable images should not be automatically deleted. Instead, they should be retained with quality labels because they may support gradability and image-quality model development.

Grade	Designation	Definition
1	Excellent	Clear, centered, high signal, no meaningful artifact
2	Good	Minor artifact but fully interpretable and suitable for AI/ML
3	Adequate	Usable but with mild/moderate limitations
4	Poor	Limited interpretability; may be usable only for some analyses
5	Ungradable	Not suitable for disease model training or clinical feature labeling, but retain for image-quality/gradability modeling

Minimum Metadata to Capture

For each imaging session, capture the following metadata:

- Study participant ID / coded ID
- Eye: OD or OS
- Date and time of imaging
- Imaging site
- Device model
- Device serial number
- Software version

- Scan protocol name
- Scan dimensions and resolution
- Operator/technician ID
- Dilation status
- Approximate pupil size if available
- Signal strength / image quality score
- Laterality confirmation
- Lens status if known
- Media opacity if present
- Artificial tears used: yes/no
- Repeat scan performed: yes/no
- Reason for repeat or failed image
- Whether image is primary or repeat
- Any protocol deviation

This metadata should be exported with the image files whenever possible.

File Export and Data Preservation

For AI/ML analysis, retain the most complete native data possible.

Preferred export hierarchy

1. Native Topcon raw/proprietary imaging files, if accessible
2. DICOM files, preferred for interoperability and metadata preservation
3. High-resolution PNG or TIFF for fundus images if DICOM is not available
4. JPEG only as secondary copy, not as the sole research file when avoidable
5. PDF reports may be stored for clinical review but should not be the primary AI training format

For OCT and OCT-A, retain

- Raw B-scan volume data
- En face images
- Thickness maps
- Segmentation outputs
- OCT-A vessel maps if available
- Device-generated reports as secondary documentation

De-identification and Naming Convention

All exported files should be named using a coded study identifier and should not include patient name, date of birth, medical record number, or other direct identifiers in the file name.

Recommended naming convention: STUDYID_VISIT_EYE_MODALITY_SCAN_DATE_VERSION

Example: SPT000123_V1_OD_OCT_MAC6x6_2026-04-26_PRIMARY

A separate linkage file connecting the coded study ID to patient identity should be stored securely, access-restricted, and maintained separately from the AI/ML image dataset.

AI/ML Dataset Handling

Images should be categorized by participant ID, eye, visit, modality, scan protocol, quality grade, disease labels if available, clinical diagnosis labels, systemic disease labels, available laboratory/vital sign metadata, and longitudinal outcome availability.

For supervised model development, dataset splitting should occur at the patient level, not image level, so that images from the same patient do not appear in both training and testing sets.

Suggested Protocol Language for IRB Submission

Retinal imaging will be performed using a Topcon Maestro OCT platform or substantially equivalent FDA-cleared retinal imaging device as part of routine clinical care. To optimize image quality and standardization for downstream AI/ML analysis, all imaging should be performed after standard-of-care pupillary dilation unless dilation is medically contraindicated, declined by the patient, or otherwise not clinically appropriate. Imaging will include, at a minimum, bilateral color fundus photography, macular OCT, and optic nerve/RNFL OCT, with OCT angiography obtained when available. Images will be acquired using standardized scan protocols and saved with associated metadata, including laterality, scan type, device model, software version, date/time of acquisition, operator, dilation status, and image quality metrics.

Images will be reviewed for technical quality, including focus, illumination, centration, signal strength, motion artifact, blink artifact, media opacity, and segmentation accuracy when applicable. Repeat imaging may be performed during the same routine clinical encounter when image quality is insufficient and repeat acquisition is clinically feasible. Poor-quality and ungradable images may be retained with appropriate quality labels for research purposes, including development of image gradability and quality-control algorithms.

Image files will be exported in the most complete format available, preferably native device format and/or DICOM, with high-resolution image files retained when applicable. PDF reports may be stored as secondary documentation but will not serve as the primary format for AI/ML model development when raw or DICOM imaging data are available. All research datasets will be coded or de-identified to the extent practicable and stored on secure, access-controlled systems.

Practical SpectOcular Minimum Standard

6. Dilated imaging of both eyes
7. Macula-centered fundus photo
8. Disc-centered fundus photo
9. Macula OCT cube
10. Optic nerve/RNFL OCT cube
11. Macula OCT-A if available
12. Repeat any scan below quality threshold
13. Export DICOM/native files whenever possible
14. Capture full metadata and image-quality grade
15. Split future AI datasets by patient, not by image

References

16. Gulshan V, Peng L, Coram M, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA. 2016;316(22):2402-2410.

17. Ting DSW, Peng L, Varadarajan AV, et al. Deep learning in ophthalmology: The technical and clinical considerations. *Prog Retin Eye Res.* 2019;72:100759.
18. Zekavat SM, Doroodgar Jorshery S, Rauscher FG, et al. Phenome- and genome-wide analyses of retinal optical coherence tomography images identify links between ocular and systemic health. *Sci Transl Med.* 2024;16:eadg4517.
19. Matthies DS, Edberg JC, Baxter SL, et al. A Manual of Procedures for the Generation of the AI-Ready and Exploratory Atlas for Diabetes Insights (AI-READI) Database. *medRxiv preprint.* 2026.
20. Chang J, Ko A, Park SM, et al. Association of Cardiovascular Mortality and Deep Learning-Funduscopy Atherosclerosis Score derived from Retinal Fundus Images. *Am J Ophthalmol.* 2020;217:121-130.
21. Islam MM, Poly TN, Walther BA, Yang HC, Li YC. Artificial Intelligence in Ophthalmology: A Meta-Analysis of Deep Learning Models for Retinal Vessels Segmentation. *J Clin Med.* 2020;9(4):1018.

Drafting note: This document is intended as a clinical research SOP/protocol insert and should be reviewed by the study PI, imaging lead, IRB/regulatory advisor, and legal counsel before final submission or operational deployment.

SOP-003 — OCT Export and Imaging Data Handling

Document	SOP-003
Title	OCT Export, Imaging Data Handling and Custody
Version	1.1 — awaiting Medical Director signature
Drafted	2026-08-18 · v1.1 2026-08-20: site references made generic
Owner	Medical Director, SpectOcular LLC
Effective	Not effective until signed by the Medical Director (§12)
Review cycle	Annually, or on any change to device fleet, export tool version, regulatory tier, or storage architecture
Supersedes	None — first controlled imaging-export procedure
Applies to	DICOM OCT Export Tool v2.1.2.14 (Topcon), User Instructions Rev E

Read before signing. The five decisions behind this procedure were reviewed and settled on 18 August 2026 and are recorded in **Appendix C** with their reasoning. The implementation item they left open — selection and validation of the Stage 2 de-identification tool (§7.2) — was closed on 18 August 2026 and is recorded in **VAL-001**. One confirmation against real device output remains outstanding, and until it is done no imaging leaves the covered Workspace account (§7.2). \$5.5 (automated export) and \$6.5 (follow-up visits) were added on 18 August 2026 and are recorded in Appendix C. Nothing here is legal, regulatory or billing advice.

1. PURPOSE

To define how OCT and OCT-A examinations are exported from Topcon devices, where the resulting data is stored, how it is coded, how custody is recorded, and how long each copy is kept — so that any image used in a read or a model can be traced back to the exam that produced it.

The absence of this procedure is the gap this document closes. As of drafting, exports are produced by a tool whose output content changes with the licence installed on the machine, and there is no written rule saying which licence to use, where the output goes, or how it is verified.

2. SCOPE — AND THE DISTINCTION THAT GOVERNS EVERYTHING

The same .fda file can produce two materially different DICOM outputs depending on the licence installed. They are **not** interchangeable.

	Clinical read	Research grade / model development
Export tier	US (no licence) — <i>not deployed, see §4</i>	Research — what SpectOcular runs
Retinal layers	5 — ILM, RNFL/GCL, IPL/INL, OS/RPE, BM	10 — adds GCL/IPL, OPL, ELM, IS/OS, CSI
OCTA metrics (vessel density, FAZ)	Not available	Available — <i>but see §3.1: not applicable to the current fleet</i>
Segmentation SOP Instance UID	...7.1	...7.3
Regulatory status	FDA-cleared subset	Not FDA-cleared

A Research-tier export must never be used to produce a clinical read, and must never be filed in a patient's clinical record. The additional layers and the OCTA metrics are outside the FDA clearance. Using them for care is a regulatory problem, not a preference. This mirrors SOP-001 §2: the same images and the same people, two activities that must not share a workflow.

The tier is verifiable after the fact — see §8.3. There is no need to trust anyone's memory about which licence was active.

3. WHAT THE TOOL PRODUCES

Input is a .fda file from a Topcon device (3D OCT-1, Maestro2, DRI OCT Triton). Output is written to a directory of your choosing.

Principal DICOM objects:

Object	Filename suffix	Contents
OP	...2.1.dcm	Fundus photograph
OPT	...1.1.dcm	B-scan tomography
OPTBSV	...3.1.dcm	Unprocessed OCTA flow volume
OPTENF	...6.y.dcm	En-face slabs
SEG (HMSEG)	...7.y.dcm	Layer segmentation, heightmap — y encodes the tier

Derived, non-DICOM outputs used for analysis: oct.json, layer_001...010.json, grid.json and faz.json (OCTA metrics), thickness binaries (Retina.dat, RNFL.dat, GCL+.dat, GCL++.dat), and rendered images (bscan_*.jpg, color_001.jpg, ir_001.jpg, shadowgram.jpg, projection.jpg).

3.1 The current fleet is structural only — read this before §5

The current pilot fleet operates **Topcon colour fundus and structural OCT**. There is no OCT-A.

Consequences, so nobody spends time chasing outputs that cannot exist:

- **Expected objects:** OP (fundus), OPT (B-scan), SEG (layer segmentation)
- **Will not be produced:** OPTBSV (...3.1.dcm, OCTA flow volume), OPTENF (...6.y.dcm, en-face slabs)
- **Will not be produced:** grid.json, densityMapGrid*.png, faz.json, fazEnface.png — the OCTA metrics in Appendix B
- Their absence is **correct behaviour on a structural scan, not a failed export**

The -octa and -enfaceSlabs switches inside -allDcm are inert on a non-angio .fda, so the standard parameter set in §5.2 needs no change. Leaving them in means the same command keeps working unchanged if an angio-capable device is added later.

This does not weaken the case for the Research tier. On a structural-only fleet the tier still governs **layer segmentation — 10 layers versus 5** — which is the scientific payload for retinal biomarker work. GCL/IPL, OPL, ELM, IS/OS and CSI are available only above the US tier and are exactly the boundaries a layer-thickness model depends on. §2 and §4 stand as written.

Record the exact Topcon model and software version for each site in the equipment log. Model determines which outputs are possible at all (Appendix B), and it is not safe to infer it from the fact that a scan succeeded.

3.2 Device scope — this SOP covers Topcon only

The export tool reads .fda, Topcon's proprietary format. **This procedure therefore applies to Topcon devices only.** Zeiss, Heidelberg, Nidek and Optovue each export DICOM by their own route and are outside its scope.

Two consequences:

1. **Procedural.** Admitting a non-Topcon device requires its own export procedure — an amendment to this SOP or a sibling document — before any imaging from it enters the corpus.
2. **Scientific, and the more important of the two.** Layer segmentation algorithms differ between manufacturers, so thickness values are **not directly comparable across devices**. Pooling them into one biomarker dataset without modelling the device effect is a recognised confound. While N is small, a homogeneous fleet is worth more than a larger heterogeneous one.

Device eligibility is governed by the qualified device list referenced in SOP-001 §5. A device not on that list is quarantined, not exported.

4. REGULATORY TIER AND LICENSING

SpectOcular runs the Research tier only. There is no clinical installation of this tool.

Clinical imaging work is done in IMAGENet 6, Topcon's clinical software. This utility exists in the SpectOcular workflow solely to produce study data, so no site has any reason to hold a US-tier installation.

Installation	Licence	Where
Research	Research tier, licence key from Topcon	Enrolling sites and SpectOcular
Clinical (US, no licence)	—	Not deployed. Add only if a documented need for clinical DICOM from this tool arises, and record it here first

A site holding only one installation cannot select the wrong one. The §8.3 check then confirms intent rather than catching a mistake.

The licence is a 32-byte hex key in DicomOCTExport.exe.license, tied to the machine's Hardware ID. Obtain it by running the tool once with no arguments to display the Hardware ID, then requesting the key from your Topcon representative.

The consequence of the Research tier is that its output is not FDA-cleared (§2). No export from this tool may be filed in a patient's clinical record or used for a care decision. If a clinician needs imaging for care, it comes from IMAGENet 6.

5. EXPORT PROCEDURE

Performed either by the **SpectOcular site agent** (§5.5, the normal case once a site is provisioned) or by trained site staff. The manual procedure below is what the agent automates, and remains the documented method for any site where the agent is not installed.

5.1 Before exporting

3. Confirm the subject is enrolled and has a **Subject ID** (format NN-XXXXX) recorded on the signed consent packet.
4. Confirm patient demographics on the device match the clinic record. **Demographics edited in IMAGENet 6 after acquisition do not propagate back into the .fda file** (Tool Instructions, Appendix B). If they were corrected downstream, the export will carry the stale values unless overridden per §7.2.
5. Confirm which installation you are running — clinical or research (§4).

5.2 Running the export

Either interface is acceptable. The GUI (ProcessFDABatchFiles.exe) is preferred for routine work; the command line is preferred for batch and for reproducibility.

Standard research export — default parameter set:

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -allDcm -json -debug
```

- -allDcm — equivalent to -dcm -segDcm -octa -enfaceSlabs; produces OP, OPT, OPTBSV, OPTENF and SEG
- -thicknessMap — retinal thickness binaries (Retina.dat, RNFL.dat, GCL+.dat, GCL++.dat). The study's endpoints are layer-thickness biomarkers, so these are captured directly rather than derived from the segmentation object afterwards
- -json — layer and scan metadata used for analysis
- -debug — writes a detailed log to the output directory; this is the export's audit record and is retained

Verify -thicknessMap on the first real scan. The tool's parameter compatibility matrix (Tool Instructions §3) lists exclusions that could not be read reliably from the published PDF. If it proves incompatible with -allDcm, run it as a second pass into the same output folder rather than dropping it:

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -thicknessMap -debug
```

Record which form was used in the equipment log, and amend this SOP to match.

In the GUI, tick **New Folder for Each FDA File**. Without it, batch output from multiple exams is written into one directory and cannot be separated afterwards.

Do not use -HarmonyExport true unless the Glaucoma Module specifically requires it; it substantially increases output volume.

OCTA metrics (-densityMapGrid, -fazData) emit only for the device, scan size and resolution combinations in **Appendix B**. Their absence for other scans is expected behaviour, not a failure.

5.3 Output folder naming

Name the output folder with the **Subject ID and the exam date only**:

03-HN7JZ_2026-08-18

Never put patient name, MRN or date of birth in a folder or file name. Folder names propagate into backups, sync logs and screenshots, none of which are covered by the controls applied to the contents.

A subject scanned twice in one day — both eyes acquired as separate exams, or a repeat after a blink artefact — produces two identically named bundles, and two same-named files in one Drive folder cannot be told apart afterwards. A six-character **scan discriminator** is therefore appended:

03-HN7JZ_2026-08-18_a1b2c3.zip

The discriminator is a digest of the exam's SOP Instance UIDs. It carries no patient information and is deterministic, so re-exporting the same exam yields the same name and overwrites rather than duplicates. The date stays ahead of it so a plain alphabetical sort of a subject's imaging folder is a visit timeline.

5.4 What may be exported and uploaded — bounded by consent and by equipment

Two independent limits apply to every upload. Neither is optional and neither substitutes for the other.

1. The subject's consent. A subject consents to SpectOcular Patient Services, and separately and optionally to Retina IQ. Those are different permissions.

Subject consented to	May be exported and uploaded	May be routed to Retina IQ
Research only	Yes — study corpus	No
Research + Retina IQ	Yes	Yes
Declined	Nothing. Do not export	No

The enrolment record carries RetinalQRequested (Yes/No) against the subject. **That field, not the operator's recollection, governs routing.** If it reads No, the imaging supports the study and nothing else, however clinically interesting the scan.

2. The equipment at the site. A site uploads what its devices actually produce, and nothing is expected of it beyond that. A structural-only site (§3.1) uploads OP, OPT and SEG; the absence of OCTA objects is not an incomplete submission and must not be queried as one. Site capability is recorded in the equipment log and is the reference for what a complete upload looks like *at that site*.

Where the two limits disagree, the narrower governs.

5.5 Automated export and upload — the SpectOcular site agent

Where installed, the export and upload steps in §5.2–§5.3 and §6.1 are performed unattended by the **SpectOcular site agent**, a scheduled task on the IMAGEnet PC that runs nightly after clinic hours. It runs the same tool with the same parameters and produces the same output; it removes the operator from the export and upload, not the requirements from the procedure.

It reaches out; nothing reaches in. The agent makes outbound HTTPS connections only. There is no inbound firewall rule, no VPN into the practice network, no standing remote-access account for SpectOcular on clinic hardware, and no Google credential on site hardware. This is a scope control as much as a network one: remote access into an IMAGEnet PC would expose the practice's entire record, whereas the agent can only ever act on exams belonging to subjects the enrolment record says have consented.

Default deny. An exam is uploaded only if it matches **exactly one consented subject**. Because the export tool operates on .fda files rather than on patients, and no IMAGEnet path can be relied on to say whose scan it is, attribution is made from the identity inside the produced DICOM, after export. Anything else is quarantined — the staged copy destroyed on the clinic PC, and a reason code carrying no patient identifier reported centrally:

Situation	Outcome
Matches one consented subject	Upload
Patient not on the consented worklist	Quarantine — no-consented-subject-match
Two consented subjects share a clinic patient ID	Quarantine — ambiguous-subject-match
Two patients present in one export folder	Quarantine — mixed-identity-within-export
Wrong regulatory tier (§4)	Quarantine — wrong-regulatory-tier
Backend unreachable	No export runs at all

The last row is deliberate. An unattended machine must never decide for itself what it may upload while the consent record is unreachable.

The agent does not de-identify. It uploads Copy 1, identified, as produced (§6.1). SpectOcular continues to build the coded and de-identified copies centrally (§6.2, §7). Automating the site's upload does not move the de-identification duty onto sites.

Sites without the agent are not second-class. Some practices will not permit third-party software on a clinical PC. The manual procedure in §5.2–§5.3 remains fully supported for them.

Installation, preflight and operating procedure: [tools/site-agent/docs/INSTALL_WINDOWS.md](#). Design and decisions: [tools/site-agent/docs/ARCHITECTURE.md](#).

6. WHERE THE DATA GOES — TWO COPIES, DELIBERATELY DIFFERENT

This is the section to read carefully. The answer is two copies, and the value depends entirely on them not being identical.

A straight duplicate is worse than a single copy: it doubles the PHI footprint and buys nothing. The reason for a second copy is that the two serve different purposes with **different retention schedules under Protocol §14**.

6.1 Copy 1 — the subject folder (identified, source)

Destination: the enrolled subject's existing Drive folder,

01_Sites / Site NN - <Practice> / Subjects / <SubjectID> / 05_Retinal_Imaging_ZIP /

- Exported DICOM **as produced**, demographics intact
- Zipped per exam, named as §5.3
- This is the source record. It is the copy that supports a clinical question, a re-read, a query about what was acquired

- Governed by the identifiable-data retention schedule (§9)

6.2 Copy 2 — the SpectOcular research corpus (coded)

Destination: a SpectOcular-controlled corpus outside the site folder tree, organised by cohort rather than by site.

- Demographics **replaced with the Subject ID**, and identifying DICOM tags removed, per §7
- Contains no patient name, MRN or date of birth
- This is the copy that feeds analysis, grading and model development
- Governed by the de-identified-dataset retention schedule (§9)

Made by SpectOcular, not by the site. The site performs one export (§6.1) and uploads it. SpectOcular produces the coded and de-identified copies centrally from that upload. Site staff carry no second command and no de-identification responsibility — every extra step at the point of care is a step that gets skipped on a busy clinic day, and de-identification is not a task to distribute across sites.

6.3 Why two, stated plainly

Retention forces it. Protocol §14 says identifiers and linkage keys are destroyed when no longer required, while *"de-identified datasets used for model development may be retained beyond this period."* Those are different clocks. If the imaging exists only inside a subject folder keyed to the linkage system, you cannot honour the first without destroying the second. The separation has to exist before it is needed, because it cannot be created retrospectively once the key is gone.

Access forces it. Subject folders sit in a site's folder tree. If a site is ever granted access to its own subjects, the research corpus must already be somewhere else.

Analysis forces it. A model cannot be trained on data scattered across per-subject folders. The corpus needs its own organisation, and building it later by walking the subject tree re-introduces the identifiers you were trying to leave behind.

6.4 Recording custody

Every export is recorded in the master database Document_Uploads sheet with Subject ID, document type, file name, upload timestamp and uploader — written by the site agent where installed, and by hand otherwise. Agent uploads additionally record the SHA-256 digest, object count and exam date in Imaging_Uploads. The agent does not delete its local copy until the custody record exists, so a failure at this step costs a retry rather than the data. An export that exists in Drive but not in that sheet has no custody record and is not considered filed.

6.5 Repeat imaging and follow-up visits

A subject is enrolled once and keeps one Subject ID for the life of the study. A follow-up visit is not a new enrolment.

For a repeat scan the technician does nothing beyond scanning the patient as usual. There is no form to complete, no data to re-enter, and **no new Subject ID to generate**. The agent attributes the exam to the existing Subject ID by matching the clinic patient identifier against the consented worklist, and the visit lands in the same subject folder as the baseline:

Subjects / 03-HN7JZ / 05_Retinal_Imaging_ZIP /

03-HN7JZ_2026-08-18_a1b2c3.zip baseline

03-HN7JZ_2026-11-20_7f9e21.zip 3-month

03-HN7JZ_2027-02-14_c40d8b.zip 6-month

Re-enrolling a returning subject is an error, and a costly one. The enrolment form generates a fresh random code each time it is used. A subject enrolled twice acquires a second Subject ID and a second folder, their follow-up becomes unlinkable to their baseline, and the longitudinal change the study exists to measure is lost. It cannot be repaired afterwards without re-identifying both records to establish that they were the same person. Where a visit must be recorded manually, **look the subject up first and reuse their existing ID**.

The intervals survive de-identification. Stage 2 shifts dates by a fixed per-subject offset rather than deleting them (§7.2, VAL-001), so the spacing between a subject's visits is preserved into the research corpus while the calendar dates are destroyed. That is what makes the series usable for longitudinal analysis and model development after the identifiers are gone.

7. CODING AND DE-IDENTIFICATION — TWO STAGES

Coding and de-identification are different things and are performed at different points. Conflating them is the most common way a dataset ends up described as something it is not.

	Stage 1 — coded	Stage 2 — de-identified
Patient name, MRN, DOB	Replaced with Subject ID	Absent
Key linking ID to person	Exists — Patient_Linkage, MFA-restricted	Destroyed, or withheld from the recipient
Device serial, institution, exact scan dates, DICOM UIDs	Still present	Removed or shifted
Status	Still PHI. BAA applies	Outside HIPAA if done completely
May be disclosed externally	Only under BAA or DUA	Yes

7.1 Stage 1 — coded, at intake

Applied when SpectOcular builds Copy 2:

-patientData: ^<SubjectID>^^^^^^

The switch replaces PatientID, name fields, gender and date of birth. Parameter order is fixed; see Tool Instructions §1.1.

This is where most people stop, and stopping here is the mistake. After the switch, the DICOM still carries Study/Series/SOP Instance UIDs, device manufacturer, model and serial number, institution name, and acquisition date and time. Dates finer than a year are themselves Safe Harbor identifiers — a scan date plus a clinic's appointment book can re-identify a subject without any name at all. Burned-in annotation, where present, is untouched.

A Stage 1 corpus is **coded data under the BAA**. It is not de-identified and must not be described as such.

7.2 Stage 2 — de-identified, before any external release

Required before the corpus, or any subset of it, leaves the covered Workspace account — to a collaborator, a statistician, a contract reader, or a regulatory submission.

The step is a DICOM de-identification pass against **DICOM PS3.15 Annex E**, applying at minimum the Basic Application Level Confidentiality Profile with:

- **Retain Longitudinal Temporal Information** — **modified dates**, so intervals between visits survive while absolute dates do not. Longitudinal change is the scientific signal; the calendar date is not
- **Retain Device Identity: off** — manufacturer, model and serial removed
- **Retain UIDs: off** — UIDs regenerated, with the mapping held only inside the covered account
- **Clean Pixel Data: on** where any burned-in annotation is possible

Tool selected and validated — 18 August 2026. The record is **VAL-001**

(03_Quality_SOPs/VAL-001_Deidentification_Tool_Validation_2026-08-18), which carries the selection rationale, the profile, the method and the full result. In summary:

Tool	pydicom 2.4.4 driven by a written SpectOcular profile module
Profile	SpectOcular Basic + ModifiedDates, v1.0 — PS3.15 codes 113100, 113107, 113108
Source	tools/deid/spectocular_deid.py, SHA-256 d3dbf130a13babefb90ca16085f5efbe29a693f895da94c9 fe5fde617fcc3813
Validated	59 checks, 59 passed, 0 failed — including a residual scan confirming twelve known source values absent from every output element

pydicom with a written profile was chosen over dcm4che and the RSNA Anonymizer on auditability rather than capability: the profile is source code, so it is version-controlled, diffable and hash-pinned, and what a past run did is recoverable from the repository rather than from shell history or a screenshot of a settings pane. Any change to the profile invalidates the cited hash and requires re-validation (VAL-001 §5).

One confirmation remains open. VAL-001 validates the profile implementation against a synthetic file built to carry every identifier a Topcon export is known to write; no real .fda export exists yet. **On the first real export, before that scan or any other leaves the covered account,** dump its full tag set and confirm no identifying tag is present that the profile does not handle (VAL-001 §6). Until that confirmation, **no imaging leaves the covered account,** and the corpus is described as coded.

7.3 Why build Stage 2 now

Protocol §13 permits de-identification *"after initial data collection and quality assurance processes are completed"*, so deferring it is compliant. It is nonetheless cheaper now: retrofitting de-identification across a grown corpus means reprocessing everything, re-verifying it, and reconciling two generations of UIDs. Defining it while the corpus is empty costs a decision and a validation run.

8. VERIFICATION — DO NOT SKIP

Performed by the operator for every export, before the source .fda is considered archived.

8.1 Completeness

Confirm the expected DICOM objects are present (§3). An .fda with no B-scans yields OP only; this is expected. Anything else missing is a failed export.

8.2 Log review

Open the result_*.log written by -debug. Any error or warning is investigated before the export is filed.

8.3 Tier confirmation — the check that cannot be faked

Inspect the segmentation object filename. The final numeric component encodes the regulatory tier:

- ...7.1.dcm → US tier
- ...7.2.dcm → OUS tier
- ...7.3.dcm → **Research tier**

A study export must read 7.3. A clinical export must read 7.1. **If the suffix does not match the intended use, the export is quarantined and re-run on the correct installation** — it is evidence that the wrong licence was active, and any downstream use is suspect.

8.4 Identifier check on the research copy

Confirm the coded copy carries the Subject ID and no patient name, MRN or date of birth. Spot-check one DICOM header per batch.

This confirms **Stage 1 only**. It does not confirm de-identification — device serial, institution and acquisition dates survive Stage 1 by design (§7.1). Stage 2 has its own validation, performed once against the chosen tool rather than per batch (§7.2).

9. RETENTION AND DESTRUCTION

Per Protocol §14:

	Retention
Source .fda files	With the device/site per clinical record practice
Copy 1 — subject folder, identified	Minimum 6 years after study completion; destroyed with identifiers and linkage keys when no longer required
Copy 2 — research corpus, coded (Stage 1)	Retained with the identifiable-data schedule while the linkage key exists — it is still PHI
Copy 2 — de-identified (Stage 2, §7.2)	May be retained beyond that period for model development and validation, once de-identified and the key destroyed or withheld
Export logs	With the copy they document

The condition on the second row is the whole point of the two-copy model. Retaining a coded corpus indefinitely while the key still exists retains identifiable data by another name.

10. FAILURE MODES TO WATCH

Symptom	Likely cause	Action
Segmentation suffix 7.1 on a study export	Research licence not installed or wrong installation used	Quarantine, re-run on research installation
Batch output all in one folder	New Folder for Each FDA File not ticked	Re-run; do not attempt to separate by hand
Demographics wrong in DICOM but right in IMAGENet 6	Downstream edit did not propagate to .fda (Appendix B)	Re-export with -patientData:
No OCTA metrics, no OPTBSV/OPTENF	Structural-only device (current fleet, §3.1), or scan configuration outside Appendix B	Expected; not a failure
-bscanAVG produced DICOM but no smoothing	Relative path used for the .fda	Re-run with a full absolute path
Export in Drive, absent from Document_Uploads	Custody step skipped	File the record before the export is used
Imaging routed to Retina IQ for a research-only subject	RetinalIQRequested not checked before routing (§5.4)	Stop the routing; the enrolment record governs, not recollection
Site queried for OCTA objects it cannot produce	Site capability not consulted (§3.1, §5.4)	No action at the site; correct the expectation
Every agent export quarantines as wrong-regulatory-tier	Clinical (US/OUS) installation in use, not Research (§4)	Reinstall the research installation; no data is lost, exports re-run
Every agent export quarantines as no-consented-subject-match	Nobody enrolled at that site yet, or the clinic patient ID on the device differs from the one captured at enrolment	Correct the enrolment record; do not relax the matching
Agent reports ambiguous-subject-match	Two enrolled subjects carry the same clinic patient ID	A data problem at the site. Fix the enrolment; never resolve it by picking one
A site stops appearing in the fleet report	Agent stopped, PC replaced, token revoked, or endpoint protection quarantined the executable	Investigate within one working week — a quiet site is indistinguishable from a site with no patients
A returning subject appears under a second Subject ID	Re-enrolled instead of looked up (§6.5)	Stop before more visits accrue; merging requires re-identifying both records
Request to send imaging outside the Workspace account	Stage 2 profile validated (VAL-001), but not yet confirmed against real device output (§7.2)	Refuse until the first real export has been tag-dumped and confirmed against the profile per VAL-001 §6

11. TRAINING AND DELEGATION

Only personnel trained on this SOP perform exports. Training is recorded in the individual's study file and includes: tier distinction (§2), the two-copy model (§6), and the verification steps (§8). Delegation is recorded on the site delegation log per SOP-002 §10.

12. APPROVAL

This SOP is not effective until signed.

Medical Director _____

Date _____

APPENDIX A — COMMAND REFERENCE

Purpose	Command
Study export, standard (site)	DicomOctExport.exe "<in>.fda" "<out>" -allDcm -thicknessMap -json -debug
Thickness as a second pass, if incompatible	DicomOctExport.exe "<in>.fda" "<out>" -thicknessMap -debug
Coding for the corpus (SpectOcular, §7.1)	add -patientData:^^<SubjectID>^^^^^^
Discover Hardware ID	run DicomOctExport.exe with no arguments
Reusable parameter set	-config "<path>\parser.config", editing OptionSwitch in <appSettings>

No clinical export command is listed: SpectOcular does not deploy the US-tier installation (§4).

Mutually exclusive: -dcm / -straightDcm; -dcm / -bscan. Full matrix in Tool Instructions §3.

APPENDIX B — WHEN OCTA METRICS ARE AVAILABLE

Vessel density — Maestro2 and 3D OCT-1, macula, superficial, ETDRS grid, 3.0/4.5/6.0 mm at 320×320; DRI OCT Triton additionally 6.0 mm at 512×512; Triton disc RPC RNFL 4.5 mm at 320×320.

FAZ — DRI OCT Triton only, macula, superficial, 3.0×3.0 mm, 320×320.

Not available on the US tier at any configuration.

Not applicable to the current device fleet at all — the pilot fleet is structural-only (§3.1). This appendix applies from the point an angio-capable device (e.g. DRI OCT Triton, or an angio-enabled Maestro2 / 3D OCT-1) enters the study.

APPENDIX C — DECISIONS SET ON 2026-08-18, FOR REVIEW

All five were reviewed and settled on 18 August 2026. They are recorded here so the reasoning survives the people who made it.

6. **Research tier only; no clinical installation (\$4).** Confirmed. Clinical imaging is done in IMAGENet 6, so this tool serves the study alone. A site holding one installation cannot select the wrong one, which turns \$8.3 from error-detection into confirmation. Revisit only if a documented need for clinical DICOM from this tool arises.
7. **Two copies, identified and coded (\$6).** Confirmed. Driven by Protocol \$14's differing retention clocks — a single copy cannot satisfy both, and identical duplicates satisfy neither. Refined on review: **the site performs one export and SpectOcular builds the coded copy centrally**, so no de-identification duty is distributed to sites.
8. **Stage 2 de-identification defined now, not deferred (\$7.2).** Changed on review. The original default was to describe the corpus as coded and add de-identification before first external transfer. Building it now was judged cheaper than retrofitting across a grown corpus. Tool selected and validated on 18 August 2026 — pydicom with a written, hash-pinned profile, 59/59 checks passed (VAL-001). The corpus remains coded and nothing leaves the covered account until the profile is also confirmed against a real Topcon export.
9. **-allDcm -thicknessMap -json -debug (\$5.2).** Amended on review to add -thicknessMap: the study's endpoints are layer-thickness biomarkers, so the thickness binaries are captured directly rather than derived later. Compatibility with -allDcm is to be verified on the first real scan, with a second-pass fallback documented.
10. **OCTA switches retained despite a structural-only fleet (\$3.1).** Confirmed. Inert on non-angio files, and keeping them means the command does not change when an angio device joins. Scope of an upload is bounded instead by consent and site equipment (\$5.4).
11. **Automated export and upload at the site (\$5.5), added 18 August 2026.** The manual procedure was written first and remains the fallback. Automation was added because the export tool is a command-line binary and therefore schedulable, and because the alternative — remote access into each clinic PC — does not scale past a handful of sites and exposes the practice's whole record rather than the consented subjects' exams. The division of labour in \$6.2 is unchanged: the site still performs one export of Copy 1, and SpectOcular still builds the coded and de-identified copies centrally.
12. **One Subject ID for the life of the study (\$6.5), stated explicitly 18 August 2026.** Raised on review: the enrolment form mints a fresh random code on every use and has no returning-subject lookup, so a technician re-enrolling a follow-up visit would create a second Subject ID and silently sever the subject's time series. Under the agent a repeat scan requires no enrolment at all. A backend guard that returns the existing Subject ID instead of minting a new one is drafted and awaiting deployment.

Implementation item carried out of this review, now closed: selection and validation of the Stage 2 de-identification tool (\$7.2) — completed 18 August 2026, recorded in VAL-001. The one task it leaves behind is confirmation of the profile against the first real Topcon export (VAL-001 \$6). Nothing may be disclosed externally until that confirmation is done.

APPENDIX D — REFERENCES

- DICOM OCT Export Tool, User Instructions Rev E, v2.1.2 (Topcon), 14 November 2025
- DICOM OCT Export Tool Description v2.1.2 Rev L
- SpectOcular Study Protocol v1.4, §13 Data Handling and Confidentiality, §14 Data Retention and Destruction
- SOP-001 Retinal Image Grading, §2 (clinical vs research separation)
- SOP-002 Reader Qualification and Certification, §10 (delegation logs)
- DICOM PS3.15 Annex E — de-identification profiles
- VAL-001 — De-identification tool selection and validation record (§7.2)
- tools/site-agent/docs/ARCHITECTURE.md — site agent design (§5.5)
- tools/site-agent/docs/INSTALL_WINDOWS.md — site agent installation and operation (§5.5)

About the SpectOcular Research Study

A simple guide to what this research is, and what taking part means for you

Your eye care team is inviting you to take part in a research study. Taking part is completely your choice. This page answers the questions patients ask most often. Your team is happy to walk through any of it with you.

● What is the purpose of this research?

The retina — the light-sensing tissue at the back of your eye — is the one place in the body where doctors can look directly at tiny blood vessels and nerve tissue without any surgery. Researchers are studying whether patterns in routine retinal images, analyzed with artificial intelligence, can help predict and monitor the health of vital organs such as the **heart, brain and nervous system, and kidneys**.

● Will this benefit me personally?

No — there is no direct medical benefit to you from taking part. The study is observational: it does not diagnose you, does not change your care, and its investigational results are not returned to you or used to guide your treatment. If researchers notice a potentially important finding, it may be shared with your treating eye care professional.

What your participation does do is help build knowledge that may allow doctors to better **predict and monitor systemic health in future patients** — potentially catching risk earlier, for many people down the road.

● What will I have to do? Are there extra tests?

Nothing beyond your normal visit. The study uses only the retinal images already taken as part of your routine eye care — color fundus photos and OCT scans. No additional imaging, appointments, medications, or procedures are performed just for research.

● Will my data be kept private and anonymous?

Your privacy is protected carefully. Your images and health information are labeled with a **randomly assigned study code**, not your name. The list connecting your identity to that code is kept separately, with access limited to authorized study personnel, and information is stored in secure, access-controlled systems.

Note: The privacy and consent details of this study are subject to review and approval by an Institutional Review Board (IRB), an independent committee that protects the rights and welfare of research participants. Final privacy protections will follow the IRB-approved consent and HIPAA authorization documents you will be asked to sign.

● Is taking part voluntary?

Yes. Participation is entirely voluntary. You may decline, or withdraw at any time, and your decision will **not** affect your care or your relationship with your eye care team in any way.

● Can I speak with someone from the SpectOcular research team before I decide?

Yes. Before you sign, you will be offered the opportunity to speak by phone or video with **licensed SpectOcular research personnel** to ask any questions about the study, or call 214-233-6170 directly. This is available whether or not you choose to take part, and asking questions does not commit you to participating.

Questions? Please ask your eye care team or the study coordinator before you decide. You will receive the full informed consent and HIPAA authorization documents to review and sign if you choose to take part.

SpectOcular — The Retinal Intelligence Platform

Patient information sheet. This document is for general information and does not replace the IRB-approved informed consent.

Version 1.0 (draft)

Pending IRB approval · Effective _____

Acerca del Estudio de Investigación de SpectOcular

Una guía sencilla sobre qué es esta investigación y qué significa participar para usted

Su equipo de atención oftalmológica lo está invitando a participar en un estudio de investigación. Participar es completamente su elección. Esta página responde las preguntas que los pacientes hacen con más frecuencia. Su equipo con gusto le explicará cualquier parte de esta información.

● ¿Cuál es el propósito de esta investigación?

La retina — el tejido sensible a la luz en la parte posterior de su ojo — es el único lugar del cuerpo donde los médicos pueden observar directamente diminutos vasos sanguíneos y tejido nervioso sin necesidad de cirugía. Los investigadores están estudiando si los patrones en las imágenes retinianas de rutina, analizados con inteligencia artificial, pueden ayudar a predecir y monitorear la salud de órganos vitales como el **corazón, el cerebro y el sistema nervioso, y los riñones**.

● ¿Esto me beneficiará a mí personalmente?

No — no hay ningún beneficio médico directo para usted por participar. El estudio es observacional: no lo diagnostica, no cambia su atención médica, y sus resultados de investigación no se le entregan a usted ni se utilizan para guiar su tratamiento. Si los investigadores observan un hallazgo potencialmente importante, este puede compartirse con su profesional de la vista tratante.

Lo que sí hace su participación es ayudar a generar conocimiento que en el futuro pueda permitir a los médicos **predecir y monitorear mejor la salud sistémica de otros pacientes** — lo que podría ayudar a detectar riesgos más temprano, para muchas personas en el futuro.

● ¿Qué tendré que hacer? ¿Hay pruebas adicionales?

Nada más allá de su visita normal. El estudio utiliza únicamente las imágenes retinianas que ya se toman como parte de su atención oftalmológica habitual — fotografías de fondo de ojo a color y tomografías de coherencia óptica (OCT). No se realizan imágenes adicionales, citas, medicamentos ni procedimientos solo para fines de investigación.

● ¿Mis datos se mantendrán privados y anónimos?

Su privacidad se protege cuidadosamente. Sus imágenes e información de salud se identifican con un **código de estudio asignado de forma aleatoria**, no con su nombre. La lista que vincula su identidad con ese código se guarda por separado, con acceso limitado al personal autorizado del estudio, y la información se almacena en sistemas seguros con acceso controlado.

Nota: Los detalles de privacidad y consentimiento de este estudio están sujetos a revisión y aprobación por una Junta de Revisión Institucional (IRB), un comité independiente que protege los derechos y el bienestar de los participantes de investigación. Las protecciones finales de privacidad se registrarán por los documentos de consentimiento y autorización HIPAA aprobados por el IRB que se le pedirá que firme.

● ¿Participar es voluntario?

Sí. La participación es completamente voluntaria. Usted puede rechazar participar, o retirarse en cualquier momento, y su decisión **no** afectará de ninguna manera su atención médica ni su relación con su equipo de atención oftalmológica.

● ¿Puedo hablar con alguien del equipo de investigación de SpectOcular antes de decidir?

Sí. Antes de firmar, se le ofrecerá la oportunidad de hablar por teléfono o video con **personal de investigación licenciado de SpectOcular** para hacer cualquier pregunta sobre el estudio, o llamar directamente al 214-233-6170. Esto está disponible ya sea que decida participar o no, y hacer preguntas no lo compromete a participar.

¿Preguntas? Por favor, consulte a su equipo de atención oftalmológica o al coordinador del estudio antes de decidir. Si elige participar, recibirá los documentos completos de consentimiento informado y autorización HIPAA para revisar y firmar.

Retinal Imaging on the Topcon Maestro 3D

Color fundus • OCT • OCT-A — what each captures and why the study collects it

The Topcon Maestro 3D acquires all three modalities below in a single, automated, non-contact session. Use this sheet to understand what each scan shows and to answer patient questions. Reminder: **only imaging acquired as part of routine clinical care is included — nothing is captured solely for research.**

Color Fundus Photography

CFP

WHAT IT CAPTURES

A wide, true-color photograph of the surface of the retina — the optic nerve head, macula, and the retinal blood vessels — in a single non-mydratic (typically no-dilation) shot.

WHY THE STUDY COLLECTS IT

The retinal vessels are the only blood vessels in the body that can be photographed directly and non-invasively. Their caliber, branching, and appearance carry signals linked to **cardiovascular and systemic vascular health**, making the color photo a foundational input for the AI analysis.

Optical Coherence Tomography

OCT • STRUCTURAL

WHAT IT CAPTURES

A high-resolution cross-sectional (“optical biopsy”) scan that images the individual layers of the retina in depth, including the nerve fiber and ganglion cell layers and the macula, at micron-level detail.

WHY THE STUDY COLLECTS IT

Retinal layer thickness reflects **neural tissue** — a direct extension of the central nervous system. Structural OCT provides quantitative measures relevant to **brain aging, neurodegeneration, and stroke risk**, complementing what the surface photo cannot show.

OCT Angiography

OCT-A • IF AVAILABLE

WHAT IT CAPTURES

A dye-free map of blood flow through the fine capillary networks of the retina and choroid, imaging the microvasculature layer by layer without any injection.

WHY THE STUDY COLLECTS IT

OCT-A visualizes the smallest vessels — the microcirculation — where early damage from **diabetes, kidney disease, and vascular disease** often first appears. When clinically acquired, it adds a rich microvascular signal to the analysis. *Collected only when clinically indicated and available on the device.*

At export

Include color fundus, structural OCT, and OCT-A (when available) in the raw export — DICOM when supported, otherwise the most complete vendor-supported raw format. Do not alter files inside the exported package, and never copy or synchronize the live Topcon database.

⚠ Clinical care comes first

These modalities are part of routine care. Research use does not change acquisition, and AI outputs are investigational — they do not guide care and are not communicated to patients.

STUDY PROTOCOL

Version: 1.4

Date: April 21, 2026

Principal Investigator: Sai Chavala, MD

Site: Retina of North Texas

1. Study Title

Prospective Evaluation of Retinal Imaging and Clinical Data for Prediction of Ocular and Systemic Disease

2. Background and Rationale

The retina is a highly vascularized and metabolically active neural tissue that reflects systemic microvascular and neurologic health. Retinal imaging has long been used clinically to assess conditions such as hypertension, diabetes, vascular disease, and common ocular disorders. Recent large-scale studies have demonstrated associations between retinal structural features, including retinal layer thickness and microvascular characteristics, and systemic conditions such as cardiovascular disease, stroke, neurodegenerative disorders, and overall mortality risk (1,2,5).

In parallel, advances in artificial intelligence and deep learning have enabled automated extraction of clinically meaningful features from retinal images, with strong performance in detection of diabetic retinopathy and broader ophthalmic image analysis (3,4). These technologies create an opportunity to develop scalable, non-invasive biomarkers for both ocular and systemic disease.

Despite these advances, much of the published work has relied on retrospective datasets. There remains a need for prospective, real-world data collection linking retinal imaging with clinical information to support robust model development and validation. This study is designed to address that gap using retinal images obtained during routine care together with associated clinical data.

3. Study Objectives

3.1 Primary Objective

- Develop machine learning models to predict cardiovascular and neurologic disease risk from multimodal retinal imaging and associated clinical data.

3.2 Secondary Objectives

- Develop machine learning models to detect ocular diseases including macular degeneration, glaucoma, and diabetic retinopathy.
- Evaluate associations between retinal imaging features and systemic and ocular disease.
- Assess feasibility of prospective collection of high-quality retinal imaging and linked clinical data for future biomarker development.

4. Study Design

This is a prospective, observational, minimal-risk study. No investigational treatment will be administered, and participation in the study will not alter clinical care, diagnosis, treatment, or follow-up recommendations.

5. Study Population

5.1 Inclusion Criteria

- Adults aged 18 years and older.
- Able to provide informed consent in English.
- Undergoing retinal imaging as part of routine clinical care.

5.2 Exclusion Criteria

- Unable to provide informed consent.
- Retinal image quality insufficient for analysis.
- Any circumstance in which the investigator determines study participation is not appropriate.

6. Target Enrollment and Sample Size Justification

Up to 5,000 subjects will be enrolled. This sample size is intended to provide sufficient retinal imaging and linked clinical data for machine learning model development, training, validation, and testing, while accounting for incomplete datasets, poor image quality, withdrawals, and variability in disease prevalence. A large data set is necessary to refine machine learning algorithms.

7. Recruitment and Enrollment Workflow

Potential participants will be identified during routine clinical care. Eligible adults may be approached by study personnel, including clinical staff or research coordinators, either prior to or following their clinical evaluation. Participants will be approached in a private and appropriate setting, and sufficient time will be provided to review the consent form and ask questions. Enrollment will occur only after informed consent and HIPAA authorization have been obtained. Relevant clinical data may be included in the study dataset and obtained from the participant's healthcare providers. Participation in the study will not affect clinical care, and individuals may decline participation without impact on their treatment or relationship with their provider. No recruitment advertisements, diaries, study cards, or retention materials are required for this protocol at the time of submission.

8. Study Procedures

Participants will undergo retinal imaging, including non-invasive fundus photography and optical coherence tomography (OCT), with or without OCT angiography, as part of routine clinical care. Clinical data may also be collected from the medical record, including demographics, medical history, diagnoses, medications, vital signs, and laboratory data if available. The participants healthcare providers may be contacted to obtain this data.

No subject will be asked to undergo retinal imaging solely for research purposes. All imaging utilized in this study will be obtained as part of routine clinical care, and participation in the study will not result in any additional procedures or changes to clinical management.

Data collection will occur prospectively at the time of routine clinical encounters. In addition, clinical data may be abstracted from the medical record either as a one-time extraction or longitudinally over time to capture relevant outcomes and disease progression, when available. No additional study visits or procedures will be required.

9. Artificial Intelligence System Description and Regulatory Framework

The investigational system is a machine learning–based clinical decision support tool designed to analyze retinal imaging data, including fundus photography, optical coherence tomography (OCT), and OCT angiography, in conjunction with structured clinical data such as demographics, medical history, vital signs, and laboratory values. The system is intended to generate probabilistic risk estimates for ocular and systemic diseases and is investigational in nature. The system will not be used to guide real-time clinical decision-making, diagnosis, or treatment and will not alter standard clinical care. Analyses will be retrospective in the sense that model development and evaluation will occur after acquisition of the clinical data and imaging.

The model architecture is based on deep learning methodologies, including convolutional neural networks (CNNs) for automated feature extraction from retinal images. These models are capable of learning hierarchical image representations without manual feature engineering. Imaging-derived features are integrated with structured clinical variables using multimodal data fusion techniques, and a prediction layer generates probabilistic outputs representing estimated disease risk. Model development follows a supervised learning approach using datasets consisting of retinal images paired with clinical outcomes derived from physician diagnoses, medical records, and laboratory data. Model performance is evaluated using standard statistical metrics, including area under the receiver operating characteristic curve, sensitivity, and specificity, and will be further assessed through prospective validation as data are collected.

Consistent with FDA draft guidance on artificial intelligence and machine learning-enabled Software as a Medical Device (SaMD), the system is being developed within a controlled lifecycle framework. This includes predefined processes for data curation, model training, validation, and performance monitoring. Model versioning and updates will be managed under a structured change control process, and updated models will not be deployed for clinical use without appropriate validation and, if applicable, regulatory review. The current study does not involve deployment of the model in a clinical setting and does not evaluate clinical performance for regulatory clearance.

The system is designed to function with **human oversight (human-in-the-loop)**, such that all clinical interpretation and decision-making remain under the authority of licensed eyecare professionals and physicians. Outputs from the system are intended to be interpretable probabilistic risk scores rather than

deterministic diagnostic conclusions. The system does not provide autonomous diagnosis, does not initiate treatment recommendations, and does not replace clinician judgment.

From a regulatory perspective, the investigational system is being developed as a potential Software as a Medical Device (SaMD) and may be subject to future FDA regulatory pathways, including 510(k) clearance or De Novo classification. However, within the context of this study, the system is used solely for research purposes and retrospective or prospective analysis of collected data. As such, the system does not currently function as a clinical diagnostic device and does not introduce additional risk to participants.

Overall, the use of the AI system in this study is limited to research, model development, and hypothesis generation. The system does not influence patient care, and its outputs are not communicated directly to patients. This design ensures that the study remains minimal risk while enabling the development and validation of predictive models aligned with FDA expectations for future clinical translation.

The data collected in this study may be used to support future regulatory submissions to the U.S. Food and Drug Administration (FDA), including but not limited to submissions related to Software as a Medical Device (SaMD) under 510(k) or De Novo pathways. Any such use will occur only after appropriate validation and in compliance with applicable regulatory requirements. The current study is not intended to establish clinical safety or effectiveness for regulatory clearance and does not involve clinical deployment of the AI system. The AI system described in this study meets the definition of Software as a Medical Device (SaMD) and is considered an investigational device; however, it is not utilized in a clinical setting and does not influence patient care.

10. Handling of Findings and Incidental Findings

Any findings, including incidental findings, will be communicated only to the treating eyecare professional and not directly to the subject by the research team. This approach is intended to ensure that any imaging-based observations are interpreted together with the full clinical picture by the eyecare professional before communication to the subject. AI-generated outputs in this study are investigational and will not be used to guide clinical care, but potential findings will be reviewed by the research team and communicated to the treating eyecare professional when warranted. These outputs will not be returned to patients. Analysis of AI outputs will occur separately from clinical workflows and solely for research and model development purposes.

11. Risks

This study presents no more than minimal risk to participants. The primary risk is potential loss of confidentiality associated with collection and storage of health information. No additional physical risk is introduced because retinal imaging used in this protocol is obtained as part of routine clinical care.

12. Potential Benefits

There is no guaranteed direct benefit to participants. Potential societal benefits include improved future detection and understanding of ocular disease and systemic disease through development of non-invasive retinal biomarkers and machine learning models.

13. Data Handling and Confidentiality

Study data will be de-identified or coded to the extent practicable and stored on secure, HIPAA-compliant systems. Access will be limited to authorized study personnel. Data may be used for research, model development, validation, and related scientific activities consistent with the informed consent and HIPAA authorization.

Study data will be coded at the time of collection, and each participant will be assigned a unique study identifier. Direct identifiers will be stored separately from study data. A linkage key connecting study identifiers to participant identity will be maintained in a secure, access-restricted environment by the Principal Investigator or designated study personnel.

To the extent practicable, datasets used for analysis and model development will be de-identified in accordance with HIPAA standards prior to use. De-identification will occur after initial data collection and quality assurance processes are completed.

All data will be stored on secure, HIPAA-compliant systems with access limited to authorized study personnel. Data may be used for research, model development, validation, and related scientific activities consistent with the informed consent and HIPAA authorization.

14. Data Retention and Destruction

Study data will be retained for a minimum of six (6) years following study completion or longer if required for regulatory, scientific, or legal purposes. De-identified datasets used for model development may be retained beyond this period to support ongoing research and validation for FDA regulatory approval.

Identifiers and linkage keys will be retained only as long as necessary to support study objectives and regulatory compliance. When no longer required, identifiable data and linkage keys will be securely destroyed using institutionally approved methods.

15. Regulatory Status

The artificial intelligence (AI) system described in this study meets the FDA definition of a medical device as Software as a Medical Device (SaMD) and is therefore considered an investigational device within the context of this research. This study qualifies for exemption from Investigational Device Exemption (IDE) requirements under 21 CFR 812.2(c)(3), as it is a **diagnostic device study** that meets the following criteria: A) The investigational device is not used as the sole basis for diagnosis or treatment decisions; B)

The investigational device does not introduce additional risk to subjects; C) The results of the investigational device are not used to guide clinical management; D) The study does not involve invasive sampling or procedures beyond routine clinical care. All data are collected as part of routine clinical care, and no invasive or non-routine procedures are performed.

The AI system is used solely for research purposes, including model development and evaluation. Outputs generated by the system are not returned to clinicians or patients and are not used to influence diagnosis, treatment, or clinical decision-making. All retinal imaging utilized in this study is obtained using FDA-cleared, commercially available devices as part of routine clinical care. No investigational hardware devices are used. Based on these factors, the study is considered IDE-exempt and does not require submission of an Investigational Device Exemption to the FDA.

16. Informed Consent Process

Written informed consent and HIPAA authorization will be obtained from all participants before inclusion in the study. Remote consent procedures may be used when appropriate. The study will enroll only adults with capacity to consent and will not enroll minors or adults unable to provide informed consent under this protocol. E-Consent may be used for this study.

17. Study Oversight and Data Monitoring

The study will be overseen by the Principal Investigator, who is responsible for ensuring compliance with the protocol, regulatory requirements, and data protection standards. Given the minimal risk nature of the study, a formal Data Safety Monitoring Board is not required. Data quality and integrity will be monitored by the study team through periodic review of data completeness, accuracy, and security. Any protocol deviations or adverse events will be documented and reported in accordance with institutional and IRB requirements.

18. References

1. Zekavat SM, Jorshery SD, Rauscher FG, Horn K, Sekimitsu S, Koyama S, Nguyen TT, Costanzo MC, Jang D, Burt NP, et al. Phenome- and genome-wide analyses of retinal optical coherence tomography images identify links between ocular and systemic health. *Sci Transl Med*. 2024 Jan 24;16(731):eadg4517.
2. Poplin R, Varadarajan AV, Blumer K, Liu Y, McConnell MV, Corrado GS, Peng L, Webster DR. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat Biomed Eng*. 2018 Mar;2(3):158-164.
3. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, Venugopalan S, Widner K, Madams T, Cuadros J, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016 Dec 13;316(22):2402-2410.
4. Ting DSW, Peng L, Varadarajan AV, Keane PA, Burlina PM, Chiang MF, Schmetterer L, Pasquale LR, Bressler NM, Webster DR, et al. Deep learning in ophthalmology: The technical and clinical considerations. *Prog Retin Eye Res*. 2019 Sep;72:100759.

5. Cheung CYL, Ikram MK, Chen C, Wong TY. Imaging retina to study dementia and stroke. Prog Retin Eye Res. 2017 Mar;57:89-107.

17. Principal Investigator Signature

Principal Investigator Signature: _____

Date: _____

PARTICIPANT INFORMED CONSENT FORM AND AUTHORIZATION TO USE AND DISCLOSE MEDICAL INFORMATION

STUDY TITLE: Prospective Evaluation of Retinal Imaging and Clinical Data for Prediction of Ocular and Systemic Disease

STUDY DOCTOR: <First Name> <Middle Name> <Last Name>, <Suffix>

STUDY SITE: <Company Name>
<Address>
<City, State, ZIP>

TELEPHONE: <Telephone>
<Telephone #2, if applicable>

SPONSOR: SpectOcular, LLC

You are being asked to participate in a research study. Your participation in this research study is strictly voluntary, meaning that you may or may not choose to take part. To decide whether or not you want to be part of this research, the purpose, procedures, risks and possible benefits of the study are described in this form so that you can make an informed decision. This process is known as informed consent. This form also explains how your medical information will be used and who may see it. You may have a copy of this form to review at your leisure or to ask advice from others.

The study doctor or study staff will answer any questions you may have about this form or about the study. Please read this document carefully and do not hesitate to ask anything about this information. After reading the consent form, if you would like to participate, you will be asked to sign this form. You will be given a signed copy of your consent form to take home and keep for your records.

Purpose of the Study

You are being asked to participate in a research study to determine whether images of the eye (retina) and related health information can be used to better understand and predict overall health, including heart and brain conditions, as well as detect certain eye diseases. About 5,000 people aged 18 years and older will participate in this research study. Your participation in this study will last for 3 years.

Procedures

If you agree to participate, images of your eye taken as part of your routine clinical care may be used for research. Your medical history and other health information may also be reviewed and obtained from your medical doctors and other healthcare providers involved in your care. You will not undergo any additional imaging or procedures for research purposes.

Risks

The primary risk is loss of confidentiality of your personal health information. No additional medical or physical risks beyond routine care are expected.

Benefits

You will not receive direct benefit from participation. This research may help improve future detection of eye and diseases that affect the whole body.

Alternatives to Participation

This study is for research purposes only so your alternative is to not participate in this study.

New Information

You will be informed in a timely manner if new information that may influence your willingness to continue participation in the study becomes available.

Costs and Compensation

There is no cost to you for participation in this study. You will not receive any compensation for participation in this study.

Artificial Intelligence Disclosure

Your data may be used to develop artificial intelligence (AI) models. These models are for research purposes only and will not be used to guide your medical care.

Data Ownership and Commercial Use

By participating, you agree that your de-identified data, including images, clinical information, and any derived data, may be used for research, development, and potential commercial purposes. You will not receive financial compensation from any future commercial use of this data.

Voluntary Participation

Participation in this study is voluntary. You may choose not to participate or may withdraw your participation at any time without penalty or loss of benefits to which you are otherwise entitled, and without affecting your medical care.

You may withdraw from the study at any time by notifying the study team using the contact information provided below either:

- In writing (for example, email or mailed notice), or
- By contacting the study team by phone

If you choose to withdraw, no new information will be collected about you for research purposes after the date of withdrawal. The effective date of withdrawal will be the date the study team receives your request.

However, information and retinal images that have already been collected and used as part of the research, including data that has been de-identified or incorporated into aggregated datasets, analyses, or machine learning models, cannot be removed or withdrawn.

Data that has been de-identified and integrated into research datasets or analytical models may continue to be used for research purposes and cannot be linked back to you.

Your participation may be stopped without your consent by the study doctor or the FDA for any reason. For example, your participation may be stopped:

- if it is deemed to be in the best interest of your health and welfare.

Employees as Research Subjects

Regarding employees who participate in the study, there will be no privilege given for participating in this research study. Likewise, there will be no penalty or drawbacks associated with not participating. Participation is strictly voluntary and will not be tied to any preferential treatment or promotion within the company.

Confidentiality and Authorization to Collect, Use and Disclose Your Medical Information

As a part of this research, records that contain information or data about you and your health may be collected and used. These records may identify you and will be kept as confidential as possible. To the extent permitted by applicable laws and regulations, the records identifying you will not be made publicly available.

Under the privacy laws, you have the right to decide who can use your protected health information (called PHI). By signing this form, you authorize the use and disclosure of your protected health information ("PHI") for research purposes as described in this study.

The information that may be used or disclosed includes:

- Retinal images obtained during your clinical care
- Medical history
- Diagnoses
- Medications
- Laboratory results
- Vital signs
- Imaging studies (including radiology such as CT, MRI, and X-ray)
- Cardiac testing (including echocardiograms and stress tests)
- Procedure reports (including colonoscopy, endoscopy, and other procedures)
- Other relevant health information

This information may be obtained from:

- Your eye care providers
- Your primary care physician
- Other healthcare providers involved in your care
- Hospitals, laboratories, and diagnostic service providers

The study team will know your identity; however, your records will be labeled with a code that is randomly assigned to you. The research staff are the only people who will have this code and its key.

The following groups may review and use your study information. They may review your study information to make sure that it is correct. They may also review your information to make sure that the study is being conducted properly.

- SpectOcular LLC and its research team
- Sterling Institutional Review Board (IRB)
- Authorized research personnel and collaborators
- Other doctors, health care professionals or research staff who are involved in the study

Your study information may be released to the groups listed above. If your study information is reviewed by these people, they may need to see your entire medical record; it is possible that your Social Security number may be included in the records reviewed. Because of this, it cannot be assured that your confidentiality will always be protected. It is possible that your information will be shared (re-disclosed) in a way that it would no longer be protected. However, this access to your records will be granted without violating your confidentiality to the extent permitted by applicable laws and regulations. By signing this form, you are authorizing this access to your records.

The results of the study, including your information, may also be presented at meetings or in articles written about the study (publications). If the results of the study (including your research or health information) are published, your identity will remain confidential.

The purpose of this authorization is to conduct research, including the development and validation of analytical models and artificial intelligence technologies related to eye and whole body health.

This authorization does not have an expiration date.

You may revoke (take away) this authorization at any time by contacting the study team in writing. However, any information already collected and used as part of the research, including data that has been de-identified or incorporated into datasets, analyses, or machine learning models, cannot be withdrawn.

As a rule, the researchers will not continue to use or disclose information about you, but will keep it secure until it is destroyed. Sometimes, it may be necessary for information about you to continue to be used or disclosed, even after you have canceled your permission or the study is over.

Examples of reasons for this include:

- To avoid losing study results that have already included your information.
- To provide limited information for research, education, or other activities. (This information would not include your name, social security number, or anything else that could let others know who you are.)
- To help government officials or the Institutional Review Board make sure that the study was conducted properly.

Information disclosed as part of this research may be subject to redisclosure and may no longer be protected under federal privacy regulations.

Contact Information

If you have questions about this study, please contact the study team at SpectOcular:

SpectOcular LLC
11615 Forest Central Drive, Suite 308
Dallas, Texas 75243
E-mail: info@spectocularhealth.com
Phone: 214-233-6170
Fax: 214-241-4947

If you have questions regarding your rights as a research participant, or if you have questions, concerns, complaints about the research, would like information, or would like to offer input, you may contact the Sterling Institutional Review Board Regulatory Department at telephone number 1-888-636-1062 (toll free) or info@sterlingirb.com.

Participant Statement and Authorization

I have read or have had read to me the Participant Informed Consent Form and Authorization to Use and Disclose Medical Information and I agree to participate voluntarily in this study. I give my permission to the study doctor to use and disclose my protected health information as described in this consent form.

I will receive a signed copy of this form.

All my questions have been answered.

I have not waived any of my legal rights by signing this document.

Printed Name of Participant

Signature of Participant

Date

Printed Name of Person Explaining Consent

Signature of Person Explaining Consent

Date

Signature of Witness

Date

*A witness is not required unless the participant is unable to read (such as blind or illiterate). If a witness is present, however, the witness must observe the entire informed consent process, including participant signature.

AUTHORIZATION FOR RELEASE OF MEDICAL RECORDS

I hereby authorize any physician, hospital, clinic, laboratory, radiology center, or other healthcare provider to release my medical records to Dr. Sai Chavala, MD, SpectOcular LLC, for research and healthcare analytics purposes.

Patient Information

Name: _____

Date of Birth: _____

Street Address: _____

City, State, Zip: _____

Information to be Released

All of the following information is requested and may be released, including but not limited to Medical history and physician notes, Laboratory results, Medications and diagnoses, Retinal and ophthalmic imaging, Radiology images and reports (e.g., CT, MRI, X-ray), Cardiac testing (e.g., echocardiogram, stress tests), Procedure reports (e.g., colonoscopy, endoscopy), and Vital signs and other clinical data. Sensitive information including mental health, substance use, and HIV-related information where permitted by law

Recipient

SpectOcular LLC
11615 Forest Central Drive, Suite 308
Dallas, Texas 75243
E-mail: info@spectocularhealth.com

Purpose

The purpose of this authorization is for research, data analysis, and development of healthcare algorithms.

Redisclosure

Information disclosed pursuant to this authorization may be subject to redisclosure and may no longer be protected under federal privacy regulations.

Expiration

This authorization does not expire unless revoked in writing.

Revocation

I understand I may revoke this authorization at any time in writing, except to the extent that action has already been taken.

Patient Signature

Patient Signature: _____

Date: _____

AUTHORIZATION FOR RELEASE OF MEDICAL RECORDS

I hereby authorize any physician, hospital, clinic, laboratory, radiology center, or other healthcare provider to release my medical records to Dr. Sai Chavala, MD, SpectOcular LLC, for research and healthcare analytics purposes.

Patient Information

Name: _____

Date of Birth: _____

Street Address: _____

City, State, Zip: _____

Information to be Released

All of the following information is requested and may be released, including but not limited to Medical history and physician notes, Laboratory results, Medications and diagnoses, Retinal and ophthalmic imaging, Radiology images and reports (e.g., CT, MRI, X-ray), Cardiac testing (e.g., echocardiogram, stress tests), Procedure reports (e.g., colonoscopy, endoscopy), and Vital signs and other clinical data. Sensitive information including mental health, substance use, and HIV-related information where permitted by law

Recipient

SpectOcular LLC
11615 Forest Central Drive, Suite 308
Dallas, Texas 75243
E-mail: info@spectocularhealth.com

Purpose

The purpose of this authorization is for research, data analysis, and development of healthcare algorithms.

Redisclosure

Information disclosed pursuant to this authorization may be subject to redisclosure and may no longer be protected under federal privacy regulations.

Expiration

This authorization does not expire unless revoked in writing.

Revocation

I understand I may revoke this authorization at any time in writing, except to the extent that action has already been taken.

Patient Signature

Patient Signature: _____

Date: _____