The Torturella Open Reference: an open benchmark for reproducible toric intraocular-lens calculations

Vitor Moreira Torturella, MD

Clínica Hott Torturella, Rio de Janeiro, Brazil

ORCID: 0009-0005-2364-2166

Abstract

Purpose: To describe and test the Torturella Open Reference, an equation-transparent pipeline for toric intraocular-lens power calculations in cataract-surgery planning, focusing on numerical reproducibility, axis handling, and explicit rejection of unsupported inputs.

Design: Descriptive computational-methods and software-verification study based entirely on synthetic inputs.

Methods: The reference combines an SRK/T spherical calculation, double-angle Abulafia–Koch regression for posterior corneal astigmatism, an explicit surgically induced astigmatism vector, and meridional vergence propagation to the SRK/T-predicted effective lens position. A seeded generator screened a stratified surface until 10,000 jointly valid cases had been retained. A separate 2,048-record safety surface contained 1,024 metamorphic records and 1,024 deliberately unsupported or invalid records. Endpoints were deterministic reconstruction, preservation of four prespecified mathematical properties, stable fail-closed error codes, and absence of a numeric toric result for prior-refractive-surgery or Total-K inputs.

Results: Retention of 10,000 core cases required screening 13,367 candidates; 3,367 were rejected by the declared SRK/T geometry gate. All core records completed successfully (10,000/10,000; observed pass proportion 100%; Wilson 95% interval, 99.96%–100%). The four metamorphic families passed in 1,024/1,024 records (100%; 99.63%–100%), and the four rejection families returned their expected error codes in 1,024/1,024 records (100%; 99.63%–100%). The frozen compressed core and safety archives had SHA-256 digests `002d3b15f94655b78389ed2ab13ef9c2d4984d659d1ee590355d73ef2a663c36` and `aa765b8501bb45115704f8d44c9bfe68380482a2aecd0c084d8f0b0c985e238a`, respectively.

Conclusions: Version 0.1.1 conformed to its prespecified computational contract in the declared environment. The benchmark provides a traceable way to locate disagreements between implementations, but it does not estimate refractive accuracy, equivalence to a proprietary calculator, clinical superiority, or patient benefit.

Keywords: intraocular lens calculation; IOL power calculation formula; toric intraocular lens; lenses for cataract surgery; SRK/T; posterior corneal astigmatism; vector analysis; reproducible research; synthetic data

Introduction

Selecting an intraocular lens for cataract surgery involves at least two computationally distinct steps: estimating the optical power required by the eye and mapping that requirement to an available lens model. At the biometry screen, a toric calculation is presented as one answer: a spherical power, a cylinder and an axis. The computation behind that answer is less singular. It includes spherical IOL-power estimation, a periodic representation of corneal astigmatism, treatment of the posterior corneal surface, surgical astigmatism, conversion between optical planes and, eventually, selection from a manufacturer's catalogue. A disagreement in any one of these steps may be concealed by a plausible-looking final number. This study evaluates the calculation layer only; it does not compare brands of intraocular lenses or select the best lens for cataract surgery.

SRK/T remains a published theoretical formula with an empirically estimated postoperative lens position [1]. The Abulafia–Koch regression supplied a practical vector correction for posterior corneal astigmatism when standard anterior keratometry is used [2]. More recent open optical descriptions, such as the Castrop toric formulation, make a related point from a different model: spherocylindrical vergences and axis conventions should be stated, not left implicit [3]. Those publications do not, by themselves, provide a frozen software-verification surface for an independent implementation.

This creates a practical problem for researchers and developers. When two implementations return different toric results, a final number alone does not reveal whether the difference arose from an axis convention, a posterior-corneal correction, the SIA vector, SRK/T geometry or optical-plane conversion. A shared surface containing inputs, intermediate quantities, expected outputs and named rejection conditions makes that comparison possible.

Within the Torturella IOL scientific program, we therefore developed the Torturella Open Reference and asked whether it could reconstruct identical synthetic outputs, preserve prespecified axis-related properties and stop when its assumptions were violated. No postoperative refraction is available in this design, so clinical accuracy is not the endpoint. The intended contribution is an inspectable computational reference that can be replayed independently without patient data or access to a commercial calculator.

Methods

Study design and data boundary

This was a descriptive computational-methods study. Its unit of analysis was a generated parameter vector, not a person or an observed eye. The generator has no database or network reader, and the release contains no clinical record, examination file, identifier, email, image or postoperative outcome. No human participant or biological material was involved.

Software release 0.1.1 implements the prespecified 0.1.0 study protocol. The source, tests, compressed synthetic records, exact runtime contract and digests are distributed together. The full input ranges and stop rules are specified in `PROTOCOL.md`.

Calculation reference

Internally, all calculations use a plus-cylinder, steep-axis convention. The engine first validates finite values and declared domains. If the supplied meridians are reversed, they are canonicalized so that K1 is flat and K2 is steep, with a 90° rotation of the K1 axis. An axis of 180° is treated as the 0° alias.

For anterior corneal cylinder magnitude C at steep axis θ, the double-angle components are

`x = C cos(2θ)` and `y = C sin(2θ)`.

The Abulafia–Koch estimate is then applied as published [2]:

`x_AK = 0.508 + 0.926x` and `y_AK = 0.009 + 0.932y`.

Surgically induced astigmatism (SIA) is represented as flattening at the incision meridian ϕ. Thus,

`x_net = x_AK − SIA cos(2ϕ)` and `y_net = y_AK − SIA sin(2ϕ)`.

The magnitude and steep axis are recovered from the resulting vector. With zero SIA, changing the incision axis must not alter any numeric output—a clinically straightforward property, but a useful test of sign and axis handling in code.

The spherical branch follows the published SRK/T optical geometry [1]. Candidate IOL powers from 0.0 to 60.0 D, in 0.5-D steps, are evaluated by the forward equation. The implementation selects the closest candidate on the myopic-or-equal side of the requested target. Measured anterior chamber depth is retained for provenance and range checking; classical SRK/T does not consume it when predicting the effective lens position in this implementation.

For the toric branch, the two net corneal meridians are propagated to the predicted effective lens position in aqueous index 1.336. For meridional power K and propagation distance d in millimetres,

`L(K,d) = K / [1 − (d/1000)K/1.336].`

The absolute difference between the propagated steep and flat meridians is reported as the model-agnostic cylinder required at the IOL plane. This is not a marketed toric SKU. The reference does not contain a manufacturer catalogue, a rotational-misalignment outcome model or a proprietary post-refractive calculation.

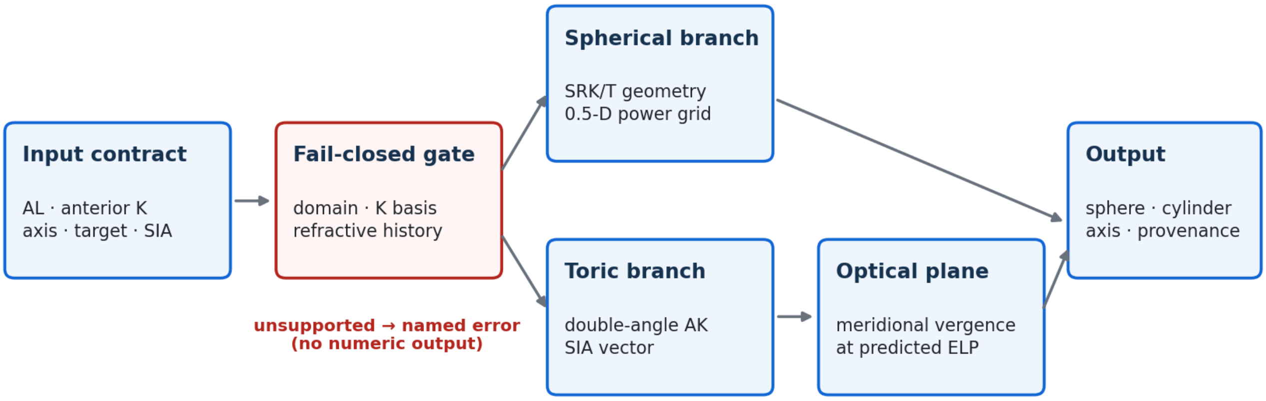

Figure 1. Prespecified computational flow. Every supported input follows the same transparent spherical and toric branches; unsupported inputs stop before a numeric recommendation is produced.

Synthetic surfaces

The core generator used seed `20260817`. Axial length, measured anterior chamber depth, mean keratometry, anterior corneal cylinder, flat-K axis, A-constant, target refraction, SIA and incision axis were stratified separately across the ranges in Table 1 and then independently shuffled. OD and OE labels were also assigned on the candidate surface; laterality has no numerical role. Candidates producing `INVALID_SRKT_GEOMETRY` were recorded as screened but were not forced into an answer. The first 10,000 jointly valid cases were retained.

Table 1. Prespecified core-input ranges

Variable	Range
Axial length	18.0–35.0 mm
Measured anterior chamber depth	2.0–5.0 mm
Mean keratometry	38.0–52.0 D
Anterior corneal cylinder	0.5–5.0 D
Flat-K axis	0–180°
A-constant	116.0–120.0
Target refraction	–2.0 to +0.5 D
SIA	0–0.75 D
Incision axis	0–180°

The safety surface contained two equal parts. In the metamorphic half, each base calculation was paired with a transformation that should leave its declared numeric signature unchanged: OD-to-OE relabelling, canonical swapping of K meridians, incision-axis change when SIA was zero, or substitution of 180° for 0°. The other half challenged four fail-closed boundaries—prior refractive surgery, Total-K basis, out-of-domain values and nonnumeric or missing values—with 256 records per family.

Endpoints and analysis

Four endpoints were specified in advance:

1. byte-identical reconstruction of the compressed result archives;
2. agreement within an absolute tolerance of `1 × 10`■^{1}■ for every metamorphic pair;
3. the expected stable error family for every unsupported or invalid record; and
4. no numeric toric output for prior-refractive-surgery or non-anterior-K inputs.

Counts were exact. Wilson 95% intervals accompany observed pass proportions, although the cases form a deterministic verification surface rather than a probability sample; the intervals describe finite-suite precision and are not population estimates. No clinical hypothesis test, sample-size calculation or superiority analysis was performed.

Reproducibility controls

JSON object keys are sorted, line separators are fixed and gzip modification time is set to zero. The result manifest stores row counts, pass/fail totals, generator seed, the exact Python, platform and compression runtime, a hash of each gzip archive and a separate hash of its uncompressed JSONL content. Fourteen unit and reconstruction tests cover published coefficients, a golden numerical case, vector round trips, SIA sign, meridional conversion, laterality and axis invariants, ACD non-use, domain rejection, archive counts and hash agreement.

Results

The generator screened 13,367 candidates to retain 10,000 jointly valid core records. The SRK/T geometry gate rejected 3,367 candidates; none was silently repaired. The retained eye labels were 5,014 OD and 4,986 OE, a descriptive imbalance after geometric filtering that has no effect on the equations.

All 10,000 retained records completed with `status=pass`. The observed pass proportion was 100% (Wilson 95% interval, 99.96%–100%), corresponding to an observed failure proportion of 0/10,000 (upper Wilson 95% bound, 0.0384%).

The safety archive also met every prespecified check (Table 2). Each metamorphic family contributed 256 records—128 base/mutant pairs—and each fail-closed family contributed 256 records. No unsupported pathway emitted a numeric answer.

Table 2. Frozen verification results

Surface	Passed / total	Observed pass proportion	Wilson 95% interval
Core jointly valid calculations	10,000 / 10,000	100%	99.96%–100%
Metamorphic records	1,024 / 1,024	100%	99.63%–100%
Fail-closed records	1,024 / 1,024	100%	99.63%–100%

Two consecutive reconstructions in the declared release environment were byte-identical. They produced the frozen SHA-256 digest `002d3b15f94655b78389ed2ab13ef9c2d4984d659d1ee590355d73ef2a663c36` for `core_10000.jsonl.gz` and `aa765b8501bb45115704f8d44c9bfe68380482a2aecd0c084d8f0b0c985e238a` for `adversarial_2048.jsonl.gz`. The corresponding uncompressed JSONL content digests were `8bc378f7bd4276b5caf4283b0252b45df0bb19191c2c0619f78a59f4feb69290` and `da0c005bcaac1304e6cd232921cb1397f3132023b6c23c8b61a10c38f3c8854c`. All 14 package tests passed.

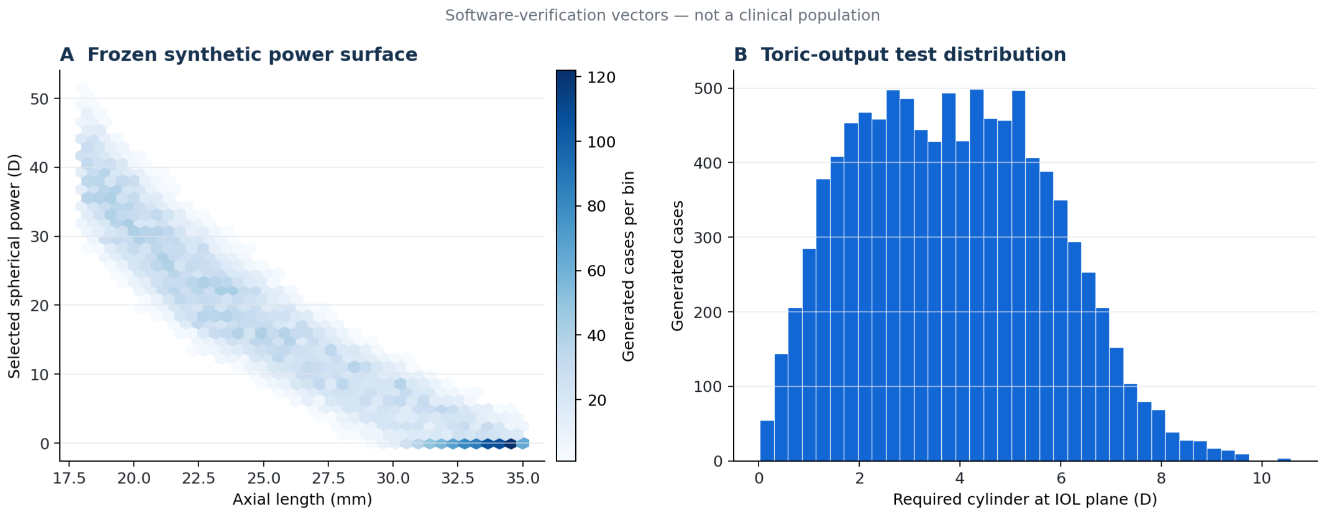

Figure 2. Descriptive view of the frozen synthetic verification surface. (A) Axial length and selected 0.5-D spherical power across 10,000 jointly valid synthetic cases. (B) Distribution of model-agnostic cylinder required at the predicted IOL plane. These generated vectors are a software-verification surface, not estimates of clinical prevalence.

Discussion

The practical value of the benchmark lies in making every numerator, denominator, transformation and failure condition inspectable. Axis periodicity, laterality irrelevance and meridian canonicalization are properties an implementation ought to preserve; here they are executable checks rather than assumptions in prose. The same applies to unsupported inputs. A Total-K value or a history of corneal refractive surgery does not trigger a substitute calculation—it stops with a named error.

This benchmark also separates three quantities often compressed in a clinical interface: corrected corneal astigmatism, cylinder required at the IOL plane, and a cylinder sold by a manufacturer. Only the first two are present. A catalogue step would need lens-specific powers, conventions and availability, followed by its own validation. Consequently, the Torturella Open Reference does not compare brands of intraocular lenses, determine the best lens for cataract surgery or support a patient-specific lens recommendation. Calling the current output a lens recommendation would overstate what was computed.

Reproducibility and clinical calibration answer different questions. The synthetic surface has no postoperative refraction against which to estimate prediction error, so it cannot rank this reference against proprietary or alternative toric methods. What the package does establish is that the declared implementation behaves consistently under stated conditions and can be examined when another implementation disagrees.

That boundary offers a practical use for other implementers. A new engine can replay the inputs, compare intermediate quantities, and determine whether a discrepancy originates in axis convention, posterior-corneal regression, SIA, SRK/T geometry, or optical-plane conversion. This does not select a lens for surgery; it provides a reproducible way to identify where two calculation pipelines diverge.

Limitations

The study uses no postoperative or other human outcome. Synthetic marginal ranges are design choices, not estimates of clinical prevalence, and filtering to jointly valid SRK/T geometry alters their joint distribution. Passing this finite suite cannot prove correctness beyond the tested properties or declared domain.

The spherical branch uses a 0.5-D grid and a stated myopic-side selection rule; a different catalogue or selection policy could produce a different spherical choice. Measured ACD is not used by classical SRK/T ELP in this implementation. The toric branch does not model a specific IOL, rotational misalignment, tilt, decentration or spectacle-plane residual cylinder after implantation. Prior corneal refractive surgery and Total K are rejected rather than calculated.

Commercial calculators were not used as reference standards. This keeps the benchmark tied to disclosed equations and avoids dependence on changing external implementations, but numerical agreement with those systems remains unknown.

During prepublication verification, the compressed hashes of the 0.1.0 candidate could not be reproduced under the available runtime, although its row counts, pass totals and tests were reproduced. Those historical hashes were not overwritten. Instead, the public artifact was versioned as 0.1.1 and frozen under the exact environment recorded in its manifest. Storing a second digest for the uncompressed JSONL content makes future checks able to distinguish a change in scientific content from a change limited to the gzip container.

Conclusions

On the frozen synthetic surface, version 0.1.1 reconstructed deterministically, preserved all four prespecified metamorphic properties, and returned the expected fail-closed codes for every challenged unsupported input. The resulting reference can help identify where implementations diverge. It is evidence of conformance to a computational contract, not evidence of refractive accuracy or clinical superiority.

Data and code availability

The public release unit for the Torturella Open Reference contains source code, tests, the frozen protocol, compressed synthetic records, SHA-256 digests, citation metadata and provenance notes and is archived at Zenodo (doi: 10.5281/zenodo.22061618). The generator reconstructs the owned result files locally with no network or database dependency. Human-derived records, manufacturer catalogues and clinical-system exports are excluded.

Funding and conflicts of interest

No external funding was received for this computational work. The author declares no proprietary interest in the formulas or calculators discussed.

Ethics statement

Ethics committee review was not applicable because the study used generated parameter vectors only and involved no human participant, animal, biological material, clinical record or identifiable information.

Author contributions

Vitor Moreira Torturella conceived the study, developed the methodology and software, performed the analysis, verified the results, prepared the visualizations, and wrote and approved the manuscript.

Declaration of generative AI and AI-assisted technologies in the writing process

DeepSeek v4-flash was used during manuscript preparation to suggest language-level edits. The author reviewed each change, retained full control of the scientific content and accepts responsibility for the final text. No AI system is listed as an author.

References

1. Retzlaff JA, Sanders DR, Kraff MC. Development of the SRK/T intraocular lens implant power calculation formula. *J Cataract Refract Surg.* 1990;16(3):333–340. doi:10.1016/S0886-3350(13)80705-580705-5).
2. Abulafia A, Koch DD, Wang L, Hill WE, Assia EI, Franchina M, Barrett GD. New regression formula for toric intraocular lens calculations. *J Cataract Refract Surg.* 2016;42(5):663–671. doi:10.1016/j.jcrs.2016.02.038.
3. Langenbucher A, Szentmáry N, Cayless A, Weisensee J, Wendelstein J, Hoffmann P. The Castrop formula for calculation of toric intraocular lenses. *Graefes Arch Clin Exp Ophthalmol.* 2021;259(11):3321–3331. doi:10.1007/s00417-021-05287-w.