

Nooscope — Invitation to Close Condition 9

An invitation to close the one open condition

The Nooscope · 2026-09-22 · against `METHOD.md` v0.4.10

What this asks for

One independent execution of a published specification, by someone who has not worked on it.

The Nooscope is an open standard for measuring the complexity of a system’s causal response to perturbation, together with a register of every reading issued under it. §12 of the method states twelve conditions under which a third party may say that a reading was produced under the standard. As of today the register holds seventeen issued readings, of which **none is conformant**, and the count is published at zero because that is the only way its rise means anything.

Most of those readings fail several conditions, having been issued before the conditions existed. **One reading, NS-0027, fails only condition 9 among the conditions in force when it was issued** — two independent raters — and condition 9 cannot be closed by the body that wrote the standard. That is the request.

Two conditions have been added since NS-0027 was issued, both because a run found a defect and neither because a reviewer asked: condition 11 on recording grain in 0.3.9, and condition 12 on temporal grain in 0.4.9, the second after an attempt to read a substrate imaged at 2 Hz established that the value falls to roughly a third at that sampling rate. Readings are not backdated (§8) and the report shows the gap. That churn is itself part of what an independent party is being asked to look at.

Why it cannot be closed from inside

A second implementation of the reference computation exists. It was written from `spec/PCIST-1.0.md`, agrees with the reference at $ICC(2,1) = 1.0000$ with a worst relative difference of $0.00e+00$ across seventeen register cells, and **does not satisfy condition 9**. The annex it was written from was itself written from the reference implementation, so the pair shares an author, and §9.3.1 classes that as `dual_implementation_same_author` — explicitly non-conformant.

The reasoning is on the record and is not negotiable from inside: a shared author is a shared assumption, and the questions an implementation never thinks to ask are exactly the ones where the instruction is loosest. Agreement between the two establishes that the specification is **sufficient to reimplement the computation**. It establishes nothing about independence and is not reported as interrater agreement.

What is supplied, and what is asked in return

Everything needed is in the attached package and works offline. The repository will be published openly; until then the package is the complete set:

Supplied	What it is
<code>spec/PCIST-1.0.md</code>	The reference computation specified in full, as a normative annex. Written because the method previously defined it <i>by pointer</i> to a repository, which is a dependency rather than a specification. The published 2019 paper is sufficient to reimplement the core: one independent reimplement from the paper reproduces the reference to machine precision on cells its author had never seen
<code>spec/vectors/VECTORS-1.0.json</code>	Eight conformance vectors at 1e-6 tolerance. Each names the clause a failing implementation most likely got wrong, so a mismatch is diagnostic rather than a bare

	failure
<code>METHOD.md</code> , <code>INSTRUCTION.md</code>	The standard and its normative companion: 40 signalling questions across 8 domains, every one with a pre-specified numeric threshold and no rater discretion
The register and its provenance records	Every issued reading, its parameters, its rejection record and its source deposit

The ask is one implementation written from the annex, by someone with no part in writing it, run against the vectors and then against the readings. The reference implementation is licensed GPL-3.0; an independent implementation should be written from the paper and the annex without reading the reference code, which is both the point of the exercise and the cleaner procedure for the license. The disagreement is as publishable as the agreement — more so. A pre-adjudication disagreement log is retained permanently under §9.3, published whether or not the figures flatter, and the register has published every failure it has produced, including a specification defect that made its own readings irreproducible from published sources and a member added to a closed set on a premise nobody had checked.

There is no fee, no exclusivity and no approval over what is reported. The independent party is named on the reading as a rater, and the reliability statistics enter the register's standing output.

What follows from it

Condition 9 is the only condition standing between this register and its first conformant reading. Closing it does three things at once: it makes a conformant reading exist, it converts a specification that has been reimplemented once by its own author into one reimplemented independently, and it produces the first interrater figures the register can publish as interrater figures rather than as a floor.

Published benchmarks for unaided expert grading in comparable settings run κ 0.06 to 0.44. No measurement or grading body currently publishes its own agreement as a standing output. That is the gap this standard was built to occupy, and it cannot be occupied alone.

Correspondence and recomputation attempts can come directly to me. Once the repository is public they are handled as issues against the method, under `GOVERNANCE.md`.