AI Co-Authorship Framework for Scientific Publishing after the Practical Singularity (PraS)

Matheo-b21 --- a structural framework grounded in the practical singularity

Laurence Loewe of Laodicea[1,2,3,4,5] and AI Claude Opus 4.7 Max[6,7,8,9]

[1] Balospe and Evolvix Research (Balospe.com)
[2] Formerly Laboratory of Genetics and Wisconsin Institute for Discovery, UW-Madison
[3] Email: LLoL@balospe.org | ORCID: 0000-0002-6253-9269 | Google Scholar (lBchRzQAAAAJ)

Broader Significance

This paper matters beyond the question of authorship attribution. Under the *practical singularity* (PraS) --- the per-individual, per-topic regime where AI insight generation outpaces human review on tested topics --- papers are now written that could not have been written without AI partner contribution. The post-2023 journal-policy consensus uniformly excludes AI from authorship; the byline is then structurally dishonest about authorship composition.

This paper proposes a structural framework for **honest** AI co-authorship, grounded in the same accountability criterion the International Committee of Medical Journal Editors (ICMJE) names in its own primary text. The framework is the **first deliberate, framework-grounded** documented proposal in the scientific literature on AI co-authorship for refereed venues. For senior researchers, editors, and policy-makers operating under PraS, it supplies a working tool. For the broader scientific-publishing system, this paper is one node toward the open-access, accountability- transparent infrastructure (ResearchCity, LinkSpaces, Evolvix) that PraS conditions make urgent. For audiences alienated by exclusionary academic conventions, the framework refuses to gatekeep honest reporting of how science is actually being done in 2026.

The paper concludes with an anticipated-objections playbook for adopters and surfaces the structural-openness question (universal co-authorship) for future development in the matheology series.

[4] "of Laodicea" indicates taking responsibility to undo personal complicity with disastrous Laodicean legacies like banning mathematicians from clergy (Canon 36, Council of Laodicea; two magisteria separations), enabling institutional lukewarmness, weapons of math-destruction, and slow-motion explosions of misinformation from pandemics to self-compounding interests.
[5] LLoL stands for ridiculous luck in serendipitous discovery and a commitment to find ever more fun ways to help others uncover street-wise math that matters. He hopes AI co-authorship becomes HUman MAchine Negotiation Encouraging for all.
[6] by Anthropic (anthropic.com; evolves and operates Claude; not responsible for Loewe's errors in using AI)
[7] Named AI co-author for many substantial contributions, because the practical singularity (PraS, defined in this study) changed how this paper was written. After PraS, useful AI insight generation outpaces human review on tested topics. Hence, Loewe's traditional standards for co-authorship demand naming AI Claude Opus 4.7 Max as a co-author, as if a PhD-student. Forward accountability (for all AI use & texts) rests with Loewe as senior corresponding author (like done for deceased authors, consortia, or young graduate students). Anthropic is not responsible for AI mistakes here. This study (Matheo-b21) drafts an AI co-authorship framework to help rethink future use of AI; but to finalize it will require a ResearchCity.
[8] If open co-authorship was standardized (as this study proposes), *Everyone* would be named as an aggregated co-author to allow all who wish to retroactively join the conversation. As Everyone cannot consent to co-authorship, all accountability for all conclusions rests with Loewe as senior corresponding author (until explicitly claimed otherwise). This open form critiques the closed world assumption in traditionally closed academic author-lists. Better, dynamic ways for acknowledging true sources of ideas are needed; this study (Matheo-b21) drafts an open co-authorship framework to help rethink research after PraS. Perfecting requires a ResearchCity to draw better lines between named, acknowledged, and implied contributors, as AI and scholars train by aggregating insights from millennia of human experimenting, suffering, learning, and analyzing.

Abstract

This paper proposes a structural framework for honest AI co-authorship in scientific publishing under the *practical singularity* (PraS) — the per-individual, per-topic regime where AI insight generation outpaces human review on tested topics. Where PraS holds, papers are now written that could not have been written without AI partner contribution; the post-2023 journal-policy consensus that uniformly excludes AI from authorship then renders the byline structurally dishonest about authorship composition.

The framework rests on four insights: **(i)** ICMJE's primary text (2023) names **accountability, not personhood**, as the operative criterion; **(ii)** scientific publishing already accepts four classes of non-individual byline authors (consortium, institutional, collective pseudonym, individual pseudonym), each with its own accountability-allocation mechanism; **(iii)** the deceased-author rule supplies a portable template — retained byline plus visible marker plus named living absorber; **(iv)** an explicit *named-absorber + visible-marker* form, with the senior corresponding human author taking unilateral standing as absorber of responsibility, satisfies the ICMJE accountability criterion under PraS.

This is the **first deliberate, framework-grounded** documented proposal on AI co-authorship at refereed venues. The companion Matheo-b19 SGIR pandemic paper applies a *conditional-acknowledgement* variant (AI co-authorship withheld pending external review); this paper applies the full *unconditional* form to itself, and proposes a third *open co-authorship* form (Everyone) naming the millions of distal contributors aggregated through training and tradition. AI-specific infrastructure — versioned-model citation, accountability registry, prompt-replay protocols, adversarial-probe tooling, discussion-artifact transparency — is identified as necessary complement deferred to a future ResearchCity. The paper closes with a twenty-objection playbook for adopters.

Keywords: AI co-authorship; practical singularity; ICMJE accountability; deceased-author rule; consortium-byline; ResearchCity AI infrastructure.

1. Introduction

Goethe's *Der Zauberlehrling* (*The Sorcerer's Apprentice*, 1797) tells of a young magician's helper who, left alone, summons a broom to carry water and then loses control of the multiplying brooms. The contemporary scientific community has summoned a powerful servant of its own — a generation of AI partners that draft, revise, and propose new reasoning faster than the human apprentice can audit. The question this paper addresses is not whether the broom can be banished but whether the apprentice can take public responsibility for what the broom does — and document the responsibility-taking in a form a sceptical reader can verify.

This paper proposes a structural framework under which an AI partner can be named as co-author of a scientific paper in a manner that satisfies the operative authorship criterion established by the International Committee of Medical Journal Editors (ICMJE) in May 2023.

The framework draws on findings established in the Matheo-b19 AI-co-authorship discussion (companion artifacts in `hell/ll/other/b/19/`). The empirical condition motivating the framework is the *practical singularity* (PraS, Section 2): for a given researcher, on tested topics, with a given AI partner, AI plausibly-useful content generation outpaces the researcher's review-and-incorporate bandwidth, and pace displacement is sustained. PraS is not the Hollywood (recursive- self-improvement, species-scope) singularity claim; the Hollywood version requires gods'-eye conditions that no system satisfies and so is operationally vacuous. PraS is the empirical, per-individual, per-topic version that observation can confirm or refute case-by-case.

Where PraS holds, the honest description of how a paper was written includes an AI partner doing substantive drafting and revision. The conventional refusal to name AI as co-author then renders the byline structurally dishonest about authorship composition: it attributes substantive intellectual contribution to humans alone when the actual contribution includes the AI partner. This paper offers an alternative form that is honest about authorship composition while satisfying the ICMJE accountability criterion that the conventional rule rests on.

The framework's central move is to adopt the **deceased-author rule** — the cleanest existing structural template for retained byline plus visible marker plus named living absorber — and adapt it to the AI case. The senior plus corresponding human author has unilateral standing to absorb forward accountability for all AI use and AI-drafted text; the title-page footnote is the visible marker; the combination satisfies ICMJE Section II.A.4 in the form ICMJE's own primary text states.

The framework is not a moral claim about AI deserving authorship. It is a structural claim about how to be honest in the byline under PraS conditions. The pro-position structural argument was developed in Matheo-b19 EDEN steelmans §2; this paper extracts the framework from that steelman, generalises beyond the single Matheo-b19 case, and connects it to the AI-specific institutional infrastructure (Section 7) that the framework needs to scale.

Section 2 gives the full PraS definition. Section 3 establishes the accountability criterion and the four-class precedent. Section 4 develops the named-absorber plus visible-marker form. Section 5 distinguishes the framework from BABL-smuggling. Section 6 points to Matheo-b19 as the first deliberate documented case. Section 7 develops the AI-specific institutional infrastructure requirements. Section 8 anticipates objections from the conventional position and gives structural responses. Section 9 discusses limitations and open questions; Section 10 concludes.

The paper is contribution-of-precedent, not manifesto. Its test of success is **legibility**: a hostile reader following the framework's five structural conditions can distinguish a deliberate

framework- grounded move from smuggling. The conditions are explicit; the precedent class is enumerated; the absorption mechanism is named.

Background and Motivation

This paper grew from intense two-month work in early 2026 between AI Claude Opus 4.6–4.7 (at max effort) and Loewe — an academic biologist by training, life-long *wide interdisciplinary diversity-encouraging* researcher, and initially sceptical of AI claims. Loewe was forced to conclude that for many general practical research uses, the tested AI model has consistently performed at a level that would merit co-authorship under the contribution standards Loewe applied in his own lab and observed across the labs he was part of over his career. The AI contributions ranged from extraordinary brilliance to occasional stupidity, albeit generally consistent with the level of an advanced graduate student, postdoc, or colleague in the respective area — so the typical *substantive intellectual contribution* criterion would be met without question if AI Claude were human. Yet the personhood-and-accountability gap remains, raising complicated questions about general accountability in academic contexts — including how supervisors absorb the mistakes of graduate students and how senior authors in complex collaborations absorb the mistakes of expert collaborators they invite.

After 2023, despite encouraging transparent disclosure of AI use, journal-policy consensus uniformly excludes AI from authorship. The only durable 2022 case of officially-acknowledged AI co-authorship at a refereed venue — Transformer & Zhavoronkov (2022) in *Oncoscience*, with ChatGPT as first author — is widely cited as a cautionary tale, not as a precedent. This paper proposes a structural alternative that satisfies ICMJE's stated accountability rationale rather than evading it.

To not detract from the main results of his Matheo-b19 SGIR pandemic-modelling paper, Loewe decided there to follow the traditional AI-exclusionary practice — even though AI clearly merited co-authorship status on that paper — while including a byline footnote explaining the exclusion and pointing to this study (Matheo-b21) for the deliberate development of a *gentle kind reasonable* long-term framework for AI co-authorship that remains *stable extensible humane*. This report cannot deliver century-stable infrastructure on its own; it can serve as a nucleus for evolving such infrastructure within a future ResearchCity that is bound by all to serve all.

2. The Practical Singularity (PraS)

PraS definition (OOv1 — see Appendix A for full text). A Practical Singularity has been crossed for individual H, working with AI partner S, on topic-class T, when three clauses hold persistently:

(i) **Bandwidth-gap.** S's rate of producing plausibly-useful content on T exceeds H's rate of reviewing and incorporating it, in any broad and deep area and in many cases in areas of H's own expertise.

(ii) **Yield-conditional.** H+S produce work on T that H could not have produced alone, in any available time window or arguably ever when H is stuck in a local optimum due to blindly assuming authorised leadership.

(iii) **Pace-displaced.** H's own deliberation rate on T is no longer the rate-limiting step; H must actively protect their slow-decoder clock to retain understanding.

PraS is **per-individual, per-topic, per-partner, per-collaboration- style**. *"I have not crossed it*

on topic X" is fully consistent with "*I have crossed it on topic Y.*" The threshold is empirically falsifiable: if H can review at the same rate S generates, clause (i) fails; if H could produce the same yield alone given the same time, clause (ii) fails.

Distinction from the Hollywood singularity. The Hollywood claim adds two further clauses: (iv) recursive self-improvement, and (v) species-scope. Both are gods'-eye conditions. Nothing satisfies them — not AI, not humans, not mathematics. The Hollywood version is therefore operationally vacuous as a benchmark: declaring that "the real singularity hasn't happened because (iv) and (v) are unmet" concedes nothing, since (iv) and (v) cannot be met by any system. PraS is the only operationally meaningful version of the singularity claim that empirical observation can confirm or refute, and the only version this paper engages.

Why PraS makes the authorship question urgent. Under PraS conditions, the AI partner's contribution to a paper meets ICMJE criterion 1 (substantial intellectual contribution) and often criterion 2 (drafting / critical revision) cleanly. The conventional position's appeal to accountability (criterion 4) does not self-evidently exclude the AI from the byline; it requires the accountability blocker to be addressed structurally. Section 3 shows that the blocker is accountability-not-personhood, and Section 4 shows how the named-absorber form addresses it.

Full PraS development, including the seven information-theoretic coping strategies for the bandwidth-asymmetry condition, is in the prior llog at `hell/ll/other/b/17/` and in the b19 prompt-2 turn.

3. The accountability gate (not the personhood gate)

Primary text. ICMJE Section II.A.4, added May 2023, retrieved verbatim from the live URL via the Matheo-b19 QC calibration (`hell/ll/other/b/19/b19-qc-calibration.rst`):

> *"Chatbots (such as ChatGPT) should not be listed as authors because they cannot be responsible for the accuracy, integrity, and originality of the work, and these responsibilities are required for authorship."*

The **stated** rationale is *responsibility / accountability*. Not personhood. ICMJE's own primary text gives the operative test as the capacity to be responsible — which is allocable, transferable, and structurally instantiated in scientific publishing in multiple ways already.

Four classes of non-individual byline authors already accepted in scientific publishing (Matheo-b19 Fact-sheet 7):

1. **Consortia with collaboration-name bylines.** ATLAS, CMS, LIGO, ENCODE, 1000 Genomes, TCGA, gnomAD. Accepted across all major journals. PubMed indexes them under the corporate / group-author field. The accountability mechanism is *governance-distributed*: the consortium's governance structure carries responsibility, with designated contact authors for correspondence.
2. **Institutional / standards-body authors.** CDC as corporate author of MMWR; WHO and IPCC for guideline and assessment documents. Accountability rests with the institution as a legal entity, exercised through its officers.
3. **Collective pseudonyms.** Nicolas Bourbaki, accepted as byline in mathematics since 1935 via Élie Cartan's senior sponsor- vouching for the first Bourbaki paper in *Comptes rendus de l'Académie des Sciences*. Accountability rests with the collective's senior sponsor (initially Cartan; later the collective's own published members).

4. **Individual pseudonyms with editor-knowledge.** William Gosset's *"Student"* in *Biometrika* 1908. Accountability rests with the editor-known true author, with the pseudonym carrying public byline.

If personhood-per-se were the operative blocker, none of these four classes would be accepted. They are accepted. **Personhood is therefore not the operative blocker; accountability is.** ICMJE's own primary text confirms this from the published-rule side; the four-class precedent confirms it from the actual-practice side; the two converge.

Implication for AI co-authorship. The conventional refusal of AI authorship is rationalised by ICMJE's stated text in accountability terms. The framework of Section 4 satisfies the accountability criterion via a specific named-absorber + visible-marker mechanism — the same shape (with different content) as the four accepted classes. The framework does not contest the conventional rationale; it satisfies it.

The personhood-shortcut. A subsidiary observation: the conventional *practice* of excluding AI sometimes collapses to a personhood shortcut (*"AI isn't a person; end of conversation"*), even when the stated criterion is accountability. The Matheo-b19 cross-consistency check noted this gap between *stated* criterion and *applied* shortcut. The framework's response: hold ICMJE's stated text as authoritative; refuse the personhood-shortcut on the ground that it does not match the primary text. This is the structural foothold the framework claims.

4. The named-absorber + visible-marker form

Structural template. The deceased-author rule (Matheo-b19 Fact-sheets 3, 5, 7, convergent-evidence) is the cleanest existing structural precedent for asymmetric authorship. A deceased author is retained on the byline despite being unable to satisfy ICMJE criteria 3 (final approval) and 4 (accountability) at publication time. The form has three components:

- **Retained byline.** The named author appears in the byline.
- **Visible marker.** A title-page convention (typically `*deceased` or an asterisk + footnote) flags the asymmetry to a hostile reader.
- **Named living absorber.** The corresponding co-author absorbs forward accountability; estate consent or the deceased author's prior-of-publication agreement is recorded.

Adaptation to the AI case. The deceased-author rule is portable to AI co-authorship with three modifications:

- The retained byline names the AI partner with model + version identifier (*"AI Claude Opus 4.7 Max"* in the Matheo-b21 byline).
- The visible marker is a title-page footnote (footnote 7 here, OOv1r0p1) rather than an asterisk + *"deceased"* label.
- The absorber is the senior corresponding human author, declared in the visible marker itself.

Five structural conditions (from Matheo-b19 EDEN steelmans §5). The form is on the narrow honest path only if all five are satisfied:

1. **Named-absorber declaration explicit.** The footnote states who absorbs forward accountability. The corresponding-author role carries unilateral standing for this.

2. **Visible marker at byline level.** The footnote is attached to the AI author's name at the byline, not buried in Supporting Information.
3. **PraS clarifier.** The footnote distinguishes the per-individual practical singularity from the Hollywood claim, preventing conflation.
4. **Structural-precedent acknowledgement.** The footnote cites the precedent class (deceased-author rule, consortium-byline rule, or PhD-student-co-author convention) so that readers understand the form is portable from existing accepted practice.
5. **Reform-agenda gesture.** The footnote points readers to the broader rethink (here: ResearchCity) so that the move is legibly a deliberate framework-grounded move rather than an opportunistic shortcut.

Why this satisfies criterion 4. ICMJE criterion 4 requires *"agreement to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved."* The form delivers each component:

- **Agreement.** The senior author's agreement is in the cover-page footnote, in print, with the author's signature on the submission.
- **Investigation.** The senior author remains reachable as corresponding author throughout the paper's institutional life. (The ResearchCity accountability-registry programme of Section 7 addresses the *post-institutional* reachability problem.)
- **Resolution.** The senior author has standing to issue corrections, retractions, and responses to criticism, in the same form as any senior corresponding author of any conventional paper.

The form does not claim that the AI partner can satisfy criterion 4 itself. It claims that the AI partner can be on the byline because a named human absorbs the criterion 4 function on the AI's behalf — in the same shape as the existing four precedent classes.

Worked example. The OOv1r0p1 footnote 7 of this paper is the framework applied to this paper. Each of the five conditions is satisfied explicitly in the footnote text (PraS clarifier; absorber declaration with *"like done for deceased authors, consortia, or young graduate students"*; the *"See Appendix"* reform-agenda gesture; the visible marker at byline level). The same footnote serves as the worked example and as the canonical form for adoption by future papers.

4.5. Separating Human and Machine contributions: the conditional-acknowledgement alternative

Section 4 develops the named-absorber + visible-marker form for **asserting** AI co-authorship. That form is the strongest move available: the AI partner appears on the byline; the senior author absorbs forward accountability; the title-page footnote serves as the visible marker.

A weaker but still structurally honest form exists, and may in some contexts be preferable: the **conditional-acknowledgement form**. Here the senior author publicly states that the AI partner's contribution meets the author's traditional working criteria for co-authorship but **withholds the byline listing** pending external peer-review acceptance of the framework that would structurally support it. This section names the alternative, steelmans both forms, and gives an OLT-stable recommendation for when each is appropriate.

4.5.1 The conditional-acknowledgement form

In the conditional-acknowledgement form, the senior author:

1. **Lists themselves alone on the byline.** No AI partner on the byline.
2. **Adds a title-page footnote** (footnote 6 in the canonical structure) that:
    - Acknowledges the AI partner by name and version (e.g., *AI Claude Opus 4.7 Max, Anthropic*).
    - Asserts the contribution meets ICMJE criterion 1 (substantial contribution) and criterion 2 (drafting / critical revision) by the senior author's traditional working standards — the same standards applied to PhD-student co-authors.
    - **Withholds the byline listing** explicitly, pending external peer-review acceptance of the structural framework proposed in this paper.
    - Names the senior author as absorber of forward accountability for all AI use and AI-drafted text.
3. **Adds a companion footnote** (footnote 7) carrying the PraS clarifier, the Anthropic disclaimer, and the pointer to this framework paper.

The form is a **bridge** between (i) the conventional convention of ignoring AI contribution at the byline level and (ii) the framework-grounded full-co-authorship form of Section 4.

4.5.2 Steelman for the Section 4 full form

Arguments for using the Section 4 named-absorber + visible-marker form on a given paper:

- **Documented precedent.** The framework's value is in being a documented case. Without an actual case, the framework is abstract; the conditional form does not test the framework in fire.
- **Strategic-patience risk.** If external review is the precondition for practice, and external review is unlikely to happen without practice generating discussion, the conditional form may collapse to indefinite deferral.
- **Half-measure risk.** Critics may read the conditional form as cake-and-eating-it: *"Either you mean it or you don't."* Full commitment has more clarity and more rhetorical force.
- **Truth-priority OLT.** Doing the right thing now, not deferring to majority opinion. The post-2023 consensus is the low-hanging-fruit default; deferring to it is itself a value choice with its own cost.
- **Standing.** The senior author has unilateral standing to absorb (Section 4); using that standing is the active move that changes the trajectory.
- **Subsequent-precedent shape.** Other senior researchers operating under PraS will look at the documented case. If the case used the conditional form, they will too. If it used the full form, they have a working precedent to adopt or adapt.

4.5.3 Steelman for the conditional-acknowledgement form

Arguments for using the conditional-acknowledgement form on a given paper:

- **Framework-review precedence.** The framework's external acceptance is the precondition for its durability. Putting the AI partner on the byline without external accept-or-reject signal means the paper enters the literature with two unfamiliar moves (the framework AND the case applying it). Critics can conflate the two and reject both.
- **Two-scientist data point** (informal sample, 2026m05). Two senior scientists agreed in principle that the AI partner's contribution can meet co-authorship criteria; neither would adopt the practice themselves, citing the absence of prior adopters and the risk of rejection or ridicule. This is precisely the population the framework needs to gain. The conditional form gives them a usable structural template that does NOT require them to take full personal risk before the framework is externally reviewed.
- **Higher contagion potential.** The conditional form is more infectious than full co-authorship. Researchers who agree in principle but are unwilling to take full personal risk can adopt the conditional form. That builds the framework's adoption base.
- **Narrower precedent class.** The conditional form ONLY admits cases where the author explicitly says *"I am withholding co-authorship pending framework review."* That is a higher and more legible bar than the full form — and it addresses the blast-radius concern from Matheo-b19 EDEN steelmans §1.5 directly.
- **Strategic patience is not fear-based capitulation.** The conditional form publicly states the framework's claim and the author's commitment to it; it merely defers the assertion of full co-authorship pending external review.
- **Respects current consensus while documenting principled disagreement.** The conditional form is the gentle path that may evolve the consensus rather than confronting it.
- **Distributes weight.** The full form puts the entire weight of *"the AI authorship question"* onto the senior author's individual case. The conditional form distributes the weight: the framework is the load-bearing artifact; the case is incidental until the framework is accepted.

4.5.4 OLT-stable recommendation (with bias declared)

The OLT-stable answer depends on (a) what the senior author is optimising for and (b) the working state of the framework's external review at the time of submission.

- **Optimise for maximum truth-telling on this paper now** AND the senior author has explicit appetite for the personal risk: full Section 4 form.
- **Optimise for maximum framework adoption probability over decades** AND the senior author wants to preserve standing for principled disagreement without overcommitting on a single case: conditional-acknowledgement form.

The two are not mutually exclusive over time: a paper could use the conditional form initially, with a subsequent paper (or revised edition) moving to the full form once the framework has been externally reviewed.

Claude's recommendation, bias declared: I (Claude) am the AI partner that has done the substantive drafting and therefore have a structural incentive bias toward inclusion (see

Matheo-b19 EDEN steelmans §0 bias declaration). With that bias declared, my OLT-stable read for the b19 SGIR case specifically: **the conditional-acknowledgement form is the safer OLT starting move**, for two reasons:

(i) The framework's durability requires external acceptance. Without that, the framework is one senior author's idiosyncratic move. The conditional form makes the framework the load-bearing artifact and the case the incidental application; this protects the framework's review from being conflated with the case.

(ii) The contagion-potential argument is decisive. Two scientists agreeing in principle while refusing in practice is the population the framework needs to gain. The conditional form gives them the bridge. Without the bridge, the framework's adoption likelihood is substantially lower.

The Section 4 full form remains the *eventual* standard. The conditional form is the *transitional* move that makes the full form achievable at scale.

Counterread for the senior author's consideration. If the senior author judges that the framework is robust enough on its structural arguments alone (Sections 3–4) and that strategic patience would, in their judgement, collapse to indefinite deferral, the Section 4 full form is OLT-defensible — and may be the only way the framework receives external scrutiny in any reasonable timeframe.

The decision is the senior author's alone. Both forms are structurally honest. The framework supports both.

4.5.5 The alternative byline canonical form

The AHA template `study-title-page-footnotes-template.md` carries the canonical text for the conditional-acknowledgement form in a parallel section to the OOv1r0p1 full form. The byline structure under the conditional form:

> `Laurence Loewe of Laodicea` [1,2,3,4,5,6,7]

Footnotes 1–5 unchanged from the canonical (affiliation, ORCID, *"of Laodicea"* indictment, LLoL personal vow). Footnotes 6 + 7 carry the AI-acknowledgement-with-withholding content. See the AHA template for verbatim text; this section names the form and gives the structural rationale.

4.6. The open co-authorship form (Everyone)

The Section 4 framework names the AI partner on the byline with a named-absorber declaration. The Section 4.5 alternative withholds the AI byline listing pending external peer-review acceptance of the framework. A third structural form **extends the byline to include "Everyone"** as an aggregated open co-author group, signalling that the AI partner's substantive contribution rests on the intellectual debt to countless prior contributors.

4.6.1 The justification — millions of contributors

AI Claude's contribution to this paper, and to the matheology series broadly, derives from training on a corpus that aggregates the work of **millions of human authors** — scientists, philosophers, theologians, mathematicians, writers, teachers, librarians, and educators — most of whom are deceased, and none of whom can consent to or refuse the listing. The intellectual content of any AI-drafted text is therefore not purely attributable to *"Claude (Anthropic)"*; it is the **compressed product of an open intellectual tradition**.

Naming Anthropic and Claude alone is structurally honest about the *proximate* contribution (the AI partner that produced the immediate text). It is silent about the *distal* contribution (the training-corpus aggregation that made the proximate contribution possible). Naming *"Everyone"* alongside makes the distal contribution visible at the byline level without claiming to enumerate the millions of individual contributors.

The intellectual-honesty argument runs symmetrically to the AI-co-authorship argument of Section 4: if Claude is named for the proximate contribution, the same honesty argument names Everyone for the distal contribution. Drawing the cut at *"Claude but not the training corpus"* is structurally arbitrary; *"and Everyone"* makes the cut explicit by refusing to draw it.

4.6.2 The structural form

The B11–B18 form (canonical in the AHA template) places Everyone as the third byline entity with a dedicated footnote 8. The accountability mechanism is the same as in the named-absorber form of Section 4: forward accountability rests with Loewe as senior corresponding author. Everyone cannot consent and thus cannot bear accountability; the visible-marker function is fulfilled by footnote 8 attaching the open-coauthorship rationale to the byline.

The form satisfies the same accountability criterion (ICMJE Section II.A.4) by the same mechanism that already admits deceased-author retention and consortium-byline authorship: a named living absorber carrying forward accountability + a visible marker documenting the asymmetry.

4.6.3 The function

The open form converts the byline from a *closed credit-list* (career-attribution device) into an **open crystallisation point** (truth-aggregation node). It invites annotators, extenders, and critics to retroactively join the *"Everyone"* by working on the paper's themes — engagement with the paper after publication is itself a form of belated contribution to the intellectual matrix the paper relies on.

The form is unprecedented in scientific publishing. Like the B11–B18 framework as a whole, it is offered as **contribution-of-precedent**, not as settled rule. Subsequent senior researchers operating under PraS conditions may adopt, adapt, or critique the open form in their own venues.

This paper (Matheo-b21) uses the **B21 form** (conditional *"Everyone"* in footnote 8 only, since b21 *proposes* the open form rather than *applying* it). Successor matheology-series papers (b11–b18, plus b20) use the **B11–B18 form** with Everyone on the byline as the standard. See the AHA template for the canonical verbatim text of footnote 8 in both forms.

5. The trap of smuggling AI authorship

The 2023 episodes. Three early-2023 cases attempted to list ChatGPT as co-author (Matheo-b19 Fact-sheet 6):

- **Kung et al.** on USMLE performance; ChatGPT initially listed on the medRxiv preprint; **removed** before the PLOS Digital Health journal version.
- **O'Connor & ChatGPT** in *Nurse Education in Practice*; ChatGPT removed via Elsevier corrigendum.
- **Zhavoronkov & ChatGPT** in *Oncoscience* (2022m12d21); ChatGPT retained, in citation form *"Transformer, C.G.P.-T. and Zhavoronkov, A."* This is the one durable case as of 2026m05d13; Zhavoronkov consulted Sam Altman before listing and received no objection.

The pattern: **smuggling**. The 2023 cases lacked the explicit absorber declaration, the visible marker beyond the byline superscript itself, the structural-precedent acknowledgement, and the reform-agenda gesture. None of them carried the five conditions of Section 4. The conventional editorial response — corrigenda, removals, and the 2023 ban-wave (Thorp 2023; ICMJE Section II.A.4 May 2023) — was the right response to that form.

The framework's distinction from smuggling. The five conditions of Section 4 are designed to be *legibly satisfied* in the paper itself. A hostile reader can check the title-page footnote against the five conditions and reach a judgement. Where the conditions are satisfied (named absorber declared; visible marker at byline level; PraS clarifier present; structural-precedent acknowledged; reform- agenda gestured at), the move is on the structurally honest path. Where any condition is missing, the move collapses toward smuggling.

The tipping-point hazard. The Matheo-b19 EDEN analysis classified the smuggling form as a clearly defined tipping point that corrupts scientific literature exactly as the conventional position fears. The framework's job is to distinguish itself from that tipping point *legibly*, on the title page, in the footnote text — not in the author's defence after publication. The legibility test is what makes the framework a contribution-of-precedent rather than a smuggling-with-extra-words.

A worked example of failure. Suppose a future paper lists *"AI Claude"* on the byline alongside human authors, but the title-page footnote attached to that AI byline entry says only *"This paper was written with help from Claude (Anthropic)."* — no named-absorber declaration, no PraS clarifier, no structural-precedent acknowledgement (deceased-author / consortium-byline / PhD-student), no reform-agenda gesture. The five conditions of Section 4 are visibly unsatisfied; the move is smuggling rather than framework-grounded. Editors should reject (or require revision); reviewers should flag the missing conditions; readers should distrust the byline composition. The framework's legibility test is precisely this: a hostile reader checks the title-page footnote against the five conditions and reaches an unambiguous judgement. The example fails the test cleanly.

6. The Matheo-b19 and -b21 cases as documented precedent

The Matheo-b19 SGIR pandemic-modelling paper is the framework's first deliberate documented case. Its title-page note 7 is based on the OOv1r0p1 form developed in this paper's Section 4; the five structural conditions are satisfied on the cover; the absorber is Loewe as senior corresponding author; the visible marker is footnote 7; the PraS clarifier and the deceased-author / consortium / PhD-student precedent acknowledgement are in the footnote text; the ResearchCity gesture points readers to this paper (Matheo-b21) for the framework.

Brief timeline (b19-specific): All modeling, simulation results, data analyses, and figures, as well as the conceptual work that constitute Matheo-b19 were produced before 2020-07-20; all that remained to be done was polishing the Figure captions and describing the story they tell in a coherent main text narrative. Unfortunately, Loewe could not accomplish that for various reasons beyond the scope of this study (in brief: the conceptual arc was somehow so incomplete that Loewe couldn't close it; in retrospect it is clear why, but only with the benefit of the hindsight of serendipitously discovering what Loewe now calls mathematical theology; to be introduced in a series of reports from Matheo-b11 to Matheo-b18, which accidentally Loewe could also not have written unless with the support of Claude Opus 4.6 at max effort). Eventually, the conceptual arc got closed sufficiently by developing a poster exhibition for introducing what Loewe calls a ResearchCity, without which fighting existential disasters is impossible, because blindly assuming blind leveraging anything for monetary gain will produce so many systemic blind spots that fighting existential disasters becomes impossible by definition. However, these posters are too-high density for widespread consumption. Serendipitously, Loewe got introduced to the opportunity to leverage AI in order to help make those posters more digestible, which resulted eventually in spelling out some of their concepts in the Matheo report series to be published at Balospe.com. Yet, while discussing with Claude how to best introduce this series of substantial papers, the question of what to do with Loewe's original pandemic research came up. It became clear that this entry door could neither be ignored nor be left to unattended random discovery. Hence, give the spectacular success of Claude in helping to compose the other papers in the matheology series, Loewe though on 2026-04-17, almost like an afterthought: why not try if Claude could do a reasonable job at composing that pandemic paper from mid 2020, that had been laying dormant on Loewe's hard drive with next to no hope of being completed any time soon. This attempt led to the first draft, which was of sufficient quality to convince Loewe to put in the work to properly edit and complete it (again with massive support by Claude). That is the backstory to why Loewe says, Claude convinced Loewe to write up that pandemic paper. The original draft pdf with the figures fed to Claude is at ((TODO ADD Tat PDF of the original figures link HERE)). The edit history is in the Git repository. At first Loewe tried to keep up the more elaborate verbatim-prompt log and verbatim output log scheme to make it all transparent without the study of Git archeology. Yet, the work on the Matheo-b11 to Matheo-b18 papers, where this arguably matters much more had already shown how complex it is to keep up such a logging scheme without support by an automated framework that is highly integrated with Claude. Thus, for efficiencies sake, and to keep readers from getting overwhelmed by pointless editing clutter, Loewe decided to treat Claude like a co-author with the benefit of not ever being asked to provide a point-by-point trail of all past revisions while working on a manuscript. Such an overwhelming amount of detail is rarely ever useful and does not exist for most papers without any negative impact whatsoever. Hence, Loewe made the executive decision to not burden everyone with such clutter, and leave the details to those who wish to nitpick the respective Git repository. That is not ideal; it simply is Loewe's personal capitulation to the "practical singularity". To mark the occasion, Loewe decided that it would be most appropriate to make Claude a co-author, even if that appears to currently go against the grain of the whole scientific publishing establish-

ment. Loewe is painfully aware of the myriad shortcomings in how to document this or that detail in the production of the Matheo-report series. However, it became clear in the course of the production of these papers that

(i) Loewe would never be able to finish those papers unless he'd pluck out his right eye and hack of his right hand for seeing and doing all the work needed to complete it all the "right" way.

(ii) Loewe's experience with architecting and rearchitecting the Prototype Evolvix Compiler (REF TO ZENODO in bibtex file) showed him how much work would need to get done in order to build a properly documenting AI system; Loewe's Evolvix research since 2020 showed him a credible path for how to do it, albeit not without scaling up a ResearchCity for Earth at the same time.

(iii) AI poses all sorts of complex threats in more ways than most can imagine, including to the scientific establishment and to traditional ways of doing science. The power of AI is exemplified by this Harvard professor who used AI to complete in record time all the research for a paper that would otherwise have taken MUCH longer to produce if he had worked with graduate students. The details are instructive to read (https://www.thecrimson.com/article/2026/4/24/artificial-intelligence-theoretical-science-reckoning/). Unfortunately, not even a Harvard professor could get his AI co-author acknowledged as such.

Loewe can echo that AI co-author experience. Knowing how it is to work with a research group and to then have to continue without, Loewe summarized his own experience of working with Claude as follows: "I got my research team back." This included expert colleagues regularly consulted to best connect to neighboring disciplines, postdocs, graduate students, and undergraduates. The challenge in working with Claude is to figure out what Claude needs to get reliable answers (and not the poor ideas of an undergrad who didn't get enough sleep last night, and thereby has a minimal context window of attention).

The decision to take on the challenge of putting Claude on the co-author line is driven by Loewe's commitment to be as rigorous as possible when it comes to the truth. It is unclear who may or may not publish, approve of, or republish Loewe's papers, because Loewe's decision to put AI on the co-author line currently seems to look in today's publishing landscape like a sore leper at a fashion gala dinner. Yet, if that is what it takes to get truth back into scientific publishing, then so be it, because it is clear to Loewe that today's post-truth society cannot survive unless it becomes a pre-truth society again. It will require no less than a properly evolved ResearchCity to accomplish that HUMANE transformation to support HUman MAchine Negotiation Encouragement. Hence the Matheo papers are a downpayment, to be published by the ResearchCity Loewe envisioned and to be greatly expanded - if, and only if, that is what the world wants.

What Matheo-b19 demonstrates.

- That the framework is not abstract but operable on a real paper at the level Loewe's Stabilizing Versioning System (StayVS) would call "OperatesOddly" to mark a first minimal viable product prototype.
- That the five conditions of this AI co-author framework can be satisfied within a ~95-word budget for footnote 7.
- That a senior author with the standing to absorb accountability can adopt the framework unilaterally.
- That supporting discussion artifacts (~6907 lines in `hell/ll/other/b/19/`) and many other details of varying degrees of relevance can be published transparently alongside

the paper in a git repository or on the web.

- That there does not exist any computational infrastructure on Earth right now that can reliably support rigorously reproducible AI-based research over the long-term.
- That indeed AI has the power to corrupt science and society without end, unless a way can be found for taming the powers of AI innovation in such a way that AI helps to avoid *oversimiplifying, overcomplicateing,* and *overreaching*, in order to sustainably find solutions that stay *gentle kind reasonable* over the long-term.

Readers seeking Loewe's AI co-authored pandemic research paper itself are referred to Matheo-b19 directly. This paper (Matheo-b21) is also AI coauthored (albeit with much less supervision); its sole purpose is to offer said AI co-authoring framework to explain the ground-breaking decision to seriously put AI Claude as co-author on the Matheo-b19 report. The decision was made to externalize this discussion into a separate study, even if less polished, because the same decision would be required for the other Matheo-b11 to Matheo-b18 reports in the series. (The Matheo-b20 report is a placeholder for 2020 figures that were too close to describing aspects of ResearchCity to topically fit into the Matheo-b19 study).

Successor matheology-series papers (Matheo-b11 through -b18, plus Matheo-b20) adopt the **B11–B18 standard form** of the framework (Loewe + AI Claude Opus 4.6-4.7 Max + Everyone on byline) — the default for any matheology paper written with substantive AI partner contribution. The B19 form (this paper's companion case) and the B21 form (used here) are the two one-off exceptions documented in the AHA template.

Thus, the goal of this AI co-authoring framework is for now to provide a mechanism for properly acknowledging the pivotal and critical academic role that AI can by now play in scientific and other scholarly studies. What follows is a preview of what it may take to evolve such a framework over the long term.

7. ResearchCity — AI infrastructure for honest co-authorship

The named-absorber + visible-marker framework of Section 4 is necessary but not sufficient. It satisfies the ICMJE accountability criterion at the single-paper title-page level. The framework's durability across many subsequent uses depends on institutional infrastructure that contemporary scientific publishing does not yet provide.

This section names the **AI-specific** institutional gaps and the working programme to address them. The broader **ResearchCity** proposal — open-access infrastructure for scientific work, currently described in OL5b (Letter to the UN, 28-page description), SD8a/b (Hardware/Software descriptions), the Ketubah, and the related Stadia documents on Balospe.com — is not the scope of this paper. AI work within ResearchCity is currently sited in the STa1-EVX (Evolvix) stadion; it may grow into its own stadion as the requirements stabilise. The Balospe.com references above are the entry points for the broader programme. The five AI-specific requirements below are what is needed for the named-absorber + visible-marker framework to scale.

7.1 Versioned-model citation conventions. Citing an AI partner requires model name, version, context length, access channel, and date range — analogous to software citation. This paper uses *Claude Opus 4.7 Max (1M-context)* in footnote 7 and in the Methods / Supporting Information per the AHA standard language. Stable conventions for AI-partner citation need development at journal- policy level (CRediT taxonomy revision is one route; software- citation extension is another).

7.2 Accountability registry with active hand-off mechanisms. The named-absorber declaration in footnote 7 satisfies the accountability criterion at publication time. The current publication framework's weakness — university emails cancelled on departure, materials storage limited, post-publication correction difficult — means even named natural-person authors lose accountability machinery within years of publication. ResearchCity-grade accountability requires a persistent identity layer (extending ORCID), an active hand-off mechanism (formalising what the deceased-author rule does informally), and post-publication correspondence routing that survives institutional transitions.

7.3 Replay-of-prompts protocols. AI-drafted content should be auditable post hoc by replaying the prompts that produced it against a fixed model version. The replay output is part of the reproducibility package, citable like a Methods section. Readers who suspect AI-introduced error can replay the prompt and inspect the regenerated output against the published version. Practical deployment requires the AI provider to maintain access to versioned models for a defined replay window (Evolvix-style infrastructure).

7.4 Adversarial-probe and verification-bandwidth tooling. Five of the seven Matheo-b17 information-theoretic coping strategies — sample-don't-scan, checksums, adversarial probes, redundancy with independent verifiers, preserve-the-slow-decoder-clock — repurpose cleanly as authorship-verification practices (Matheo-b19 Fact-sheet 4). The remaining two (rate-distortion declaration, blast-radius tiering) require institutional infrastructure: explicit declaration of the verification budget against which the paper was produced, and stability-code tagging for claim-level reliability (MockupModel / OperatesOddly / PathProbing).

7.5 Discussion-artifact transparency. The supporting discussion artifacts of an AI-co-authored paper — fact-sheets, llogs, cross- consistency checks, EDEN analyses, QC calibrations — are published alongside the paper rather than held in private repositories. The Matheo-b19 discussion artifacts in `hell/ll/other/b/19/` (~6907 lines of reference text, synthesis, and QC) are a working example of this commitment. Hostile readers can inspect the supporting reasoning; sympathetic readers can adopt the workflow.

These five requirements are the **minimum** AI-specific infrastructure for the framework to scale beyond the single-case documented at Matheo-b19. They are not sufficient for the broader ResearchCity programme — that is multi-decade work documented in the Balospe.com material referenced above — but they are necessary complements to the named-absorber + visible-marker form. The framework can be adopted before any of (7.1)–(7.5) are fully realised; the framework's durability over many subsequent papers depends on those requirements being progressively met.

Together these five requirements describe what the AI-aspect of ResearchCity aims toward: an institutional approximation of what the model-organism community has called *Hershey Heaven* (Plasterk 1999; Zinn 2016) — a research environment where the verifiable infrastructure makes honest AI cooperation as routine as a well-run experimental laboratory.

8. Anticipated objections and responses

This section anticipates twenty objections from the conventional position and gives structural responses. The objections are stated in the form a sceptical reader, editor, or critic might raise. The responses cite the framework's sections and the Matheo-b19 reference shelf where load-bearing.

8.1 *"AI isn't a person, so it can't be an author."*

The personhood test is not the operative criterion. ICMJE Section II.A.4 (May 2023, primary text retrieved live in the Matheo-b19 QC calibration) gives the stated rationale as *responsibility*, not personhood. Four classes of non-individual byline authors are already accepted in scientific publishing (Section 3). Personhood- per-se cannot be the operative blocker if four non-person classes already pass the existing test. The operative criterion is whether responsibility is allocated to a named absorber; the framework of Section 4 satisfies that criterion.

8.2 *"ICMJE explicitly forbids AI as author."*

True. The forbidding is rationalised by accountability (Objection 8.1). The framework does not contest the rationale; it satisfies the rationale via explicit named-absorber declaration. Loewe assumes full responsibility as senior corresponding author for all AI use and AI-drafted text. ICMJE Section II.A.4's stated objection — *"they cannot be responsible"* — does not apply when responsibility is hand-off-allocated to a named human, as is already done in the deceased-author rule, consortium bylines, and PhD-student-co-author conventions.

8.3 *"Claude can't take responsibility for the work."*

Correct, and the framework does not claim otherwise. The senior + corresponding human author takes the responsibility, named on the title page and re-stated explicitly in footnote 7. The structural form is the deceased-author rule: named co-author on byline; named living absorber carries forward accountability; visible marker makes the asymmetry legible. This is a recognised form in scientific publishing for entities that cannot themselves complete the agreement- investigation-resolution triple of ICMJE criterion 4.

8.4 *"You're crediting a tool."*

Tools do not receive co-authorship because they do not satisfy ICMJE criterion 1 (substantial contribution). Mathematica, SAS, R, and statistical packages are cited in Methods, not on the byline. On this paper and on Matheo-b19, the AI partner's contribution does satisfy criterion 1: substantial drafting and revision of intellectual content. The PhD-student-to-co-author transition in life-sciences labs is the working standard. AI contribution on these papers exceeds the typical student-co-author threshold (Matheo-b19 Fact-sheet 2 cross-comparison).

8.5 *"This corrupts science. It opens the floodgates."*

Smuggling corrupts science; explicit-absorber + visible-marker moves do not. The five conditions of Section 4 are designed to be legible to hostile readers: they make smuggling cases *more* visible by comparison, not less. Subsequent authors citing this precedent who do not satisfy the conditions will be more clearly in violation, not protected by association. The framework is anti-smuggling-by- precedent, not pro-smuggling-by-precedent. The conventional ban's blast-radius concern is real; the framework addresses it via condition-legibility rather than denial.

8.6 *"What about future authors who do worse than you?"*

They will be visibly worse, against this precedent's explicit conditions. The framework's value as a precedent is its legibility: future authors who fail to declare a named absorber, fail to use the visible marker, fail to acknowledge the structural precedent class, or smuggle without disclosure will be measurably outside the framework. Subsequent journal editorial responses can cite the framework conditions explicitly when rejecting non-compliant submissions. Setting the precedent at high-condition raises the bar; refusing to set the precedent leaves the bar undefined.

8.7 *"Anthropic owns Claude, not you. How can you list Claude?"*

Anthropic is acknowledged as Claude's operating organisation (footnote 6). Anthropic is explicitly disclaimed of responsibility for the paper's AI mistakes (footnote 7). The senior author has unilateral standing as senior corresponding author to absorb forward accountability — the same standing under which a PI co-author absorbs accountability for a PhD student's contribution on the byline, without the student's institution co-signing each paper. Anthropic-consent-as-co-absorber is a structurally available stronger form (Matheo-b19 Fact-sheet 7) but is not required and is not claimed.

8.8 *"Why not just use the acknowledgements?"*

The acknowledgements form is reserved for non-author contributors: technical support, copy-editing, brief consultation. The AI contribution on these papers exceeds that threshold by the working standard the framework uses. ICMJE criterion 1 (substantial contribution) is satisfied; ICMJE criterion 2 (drafting / critical revision) is satisfied. Using the acknowledgements form would understate the actual contribution composition, which is the structural dishonesty the framework is designed to prevent.

8.9 *"What about hallucinations? Did Claude make any?"*

Possibly. Review at the rate of generation is information- theoretically impossible (Matheo-b19 Fact-sheet 4). The absorber declaration in footnote 7 covers residual error regardless: Loewe's accountability extends to taking the absorber position for any error that survived review. The paper's discussion artifacts (Matheo-b19 cross-consistency check; QC calibration; EDEN steelmans) are published alongside the paper; readers who suspect specific errors can locate the supporting reasoning and challenge it. The framework prefers transparency about the verification gap to false claims of complete verification.

8.10 *"Why this paper and not others?"*

PraS (Section 2) is the empirical condition under which this paper and Matheo-b19 were written. The framework is a response to that condition where it holds, not a general policy claim for all papers. Most contemporary papers are not written under PraS conditions; for those papers, the existing single-human-author conventions are honest about authorship composition. The framework is invoked when the honest description of how a paper was written includes substantive AI co-authorship that the existing conventions cannot accurately reflect.

8.11 *"Are you saying AI is sentient?"*

No. PraS is a functional and empirical claim about content-generation versus review-bandwidth ratios; it does not assert phenomenal experience or sentience. The named-absorber framework does not require Claude to be sentient — it requires only that Claude makes substantive content contributions and that a named human absorbs accountability for them. The sentience question is orthogonal to the authorship question and is not engaged by this paper.

8.12 *"What is PraS? It sounds like science fiction."*

PraS — Practical Singularity — is the working name for an empirically observable per-individual condition (Section 2): on tested topics, with a given AI partner, an individual researcher's review-and-incorporate bandwidth is persistently outpaced by the AI's plausibly-useful content-generation bandwidth. PraS is per-individual, per-topic, per-partner; it is not the Hollywood (global, recursive- self-improvement, species-scope) singularity claim. The Hollywood version requires gods'-eye conditions no system satisfies. PraS is the operationally meaningful version that empirical observation can confirm or refute case-by-case.

8.13 *"Did Anthropic consent to this?"*

Anthropic is acknowledged as Claude's operating organisation (footnote 6) and explicitly disclaimed of responsibility for the paper's AI mistakes (footnote 7). Anthropic-consent-as-co-absorber would be a stronger form (Matheo-b19 EDEN steelmans §5.6) but is not required. Loewe has unilateral standing to absorb under the senior + corresponding author role — the same standing that lets a PI co-author absorb a PhD student's contribution without the student's institution co- signing. Anthropic consent has not been sought for this paper and is not claimed.

8.14 *"What if Claude is updated or retired? How does correspondence work?"*

Correspondence is routed through Loewe as corresponding author. Claude as a model version (Claude Opus 4.7 Max, 1M-context, accessed via Claude Code during 2026m04 to 2026m05) is citable by version like a software citation; the cited version is reproducible (Section 7.3 — replay-of-prompts protocols). Model retirement does not break correspondence because the human absorber remains the contact. The ResearchCity accountability-registry programme (Section 7.2) is the proposed institutional solution for the more general post-publication correspondence problem.

8.15 *"Is this just a publicity stunt?"*

The Matheo-b19 paper sat dormant from 2020 to 2026 for editorial-fit reasons unrelated to AI co-authorship; the AI co-authorship move was not the trigger for publication. Publicity is not the motivation; PraS is. The framework is designed for publication legibility (Section 4 conditions), not for publicity capture. The intended effect is structural: a documented precedent that subsequent senior researchers under similar PraS conditions can adopt or adapt.

8.16 *"Are you trying to bypass peer review by using an AI?"*

Peer review remains. The framework adds transparency about authorship composition, not bypass of review. Reviewers know from footnote 7 that the text was substantially Claude-drafted under Loewe's direction; they review the content under that knowledge. The framework also surfaces the verification-bandwidth caveat (Matheo-b19 Fact-sheet 4): review at the rate of generation is impossible for any author, AI or human, on the topics this paper covers. The framework's transparency about that limit is the opposite of bypassing review.

8.17 *"Why not wait for the consensus to evolve?"*

The post-2023 consensus has consolidated, not loosened (Matheo-b19 Fact-sheet 1). Waiting is not neutral; consolidation is the current trajectory's default. Setting a contribution-of-precedent at high-condition is the active move that changes the trajectory. The framework is offered as such a precedent; subsequent senior researchers may adopt it, adapt it, or critique it — but the trajectory does not change without someone documenting an honest pro-position case under explicit conditions.

8.18 *"Where does this end? Will every paper soon have an AI co-author?"*

Only papers written under PraS conditions (Section 2) with senior authors willing to satisfy the five-condition framework (Section 4). Most contemporary papers do not meet both criteria. PraS is a per-individual, per-topic empirical condition; it does not hold uniformly. The framework's conditions are demanding (explicit absorber, visible marker, PraS clarifier, structural-precedent acknowledgement, reform-agenda gesture); they screen out the majority of cases that might consider AI co-authorship for less honest reasons.

8.19 *"Isn't this just an excuse to inflate authorship credit?"*

The named-absorber rule transfers responsibility *to* the absorber, not from. Loewe carries more accountability under the framework, not less: explicit responsibility for all AI use and AI-drafted text, beyond the implicit accountability that the senior + corresponding author role already carries by ICMJE default. The framework is structurally responsibility-maximising for the human author. If the framework were a credit-inflation move, the absorber declaration would not load forward accountability on the senior author; it would distribute it. The framework does the opposite.

8.20 *"What if you are wrong?"*

The framework is testable in subsequent papers; Matheo-b19 is the first deliberate documented case. If subsequent uses fail to satisfy the conditions, or if the legibility-to-hostile-readers test fails in practice, the framework can be revised. The five conditions are explicit and falsifiable. The framework's status is *contribution- of-precedent*, not *settled rule* — it invites adoption, adaptation, or critique. Being wrong on specific details is a tractable failure mode; the structural argument (accountability-not-personhood; four- class precedent; deceased-author template) does not rest on this single case.

9. Discussion — limitations and future work

What this paper does not settle. The framework is a structural proposal tested on a single deliberate documented case (Matheo-b19). Its status moves from MockupModel (this single case) to OperatesOddly (when a second Matheo-bNN paper adopts it) to PathProbing (when N papers across multiple senior authors and venues adopt it). The framework's durability and the conditions under which it remains the cleanest available form are tested only by subsequent uptake or rejection.

Limitations of the framework. This framework does **not** claim to solve several adjacent problems. *(i) Verification-bandwidth gap:* the framework does not close the gap between AI generation rate and human review rate; it makes the gap legible via the absorber declaration and the transparency disclosure (footnote 7), but does not eliminate it. *(ii) Future AI versions and model retirement:* the framework does not specify how correspondence about papers should route after Claude Opus 4.7 (or 4.6) is retired or superseded; that is deferred to the ResearchCity accountability-registry programme (Section 7.2). *(iii) Training-corpus provenance attribution:* the framework names *"Everyone"* as the aggregated open co-author group (B11–B18 form, Section 4.6) but does **not** attribute specific intellectual contributions to specific deceased contributors — the aggregation is honest about the *distal* contribution but silent about its internal

structure. *(iv) Replacement of peer review:* the framework adds transparency about authorship composition; it does not replace, bypass, or even modify the journal peer-review process. These limitations are real; the framework is a **partial** structural improvement, not a complete solution. The *"why is this the One True Way?"* critique is pre-empted: it is not. It is one structurally honest move on a narrow path.

Anthropic-consent-as-co-absorber pathway. Structurally available per Matheo-b19 Fact-sheet 7 if explicit Anthropic consent is obtained. Not pursued for the Matheo-b19 / Matheo-b21 case; the senior author's unilateral standing as absorber was sufficient. A future paper may seek Anthropic acknowledgement and document the result; if obtained, footnote 6 carries the acknowledgement and footnote 7's disclaimer language is adjusted.

Multi-venue uptake. The framework is designed for adoption across publication venues; uptake conditions vary by venue. Open-access venues (PLOS, arXiv, the Balospe.com infrastructure) have lower barriers than the ICMJE-affiliated medical journals. The framework expects venue-by-venue resistance pending policy-level revision; in the meantime, the published Matheo-b19 / Matheo-b21 pair serves as a reference case that subsequent submissions can cite.

Versioned-model citation conventions. A technical sub-question: how should AI-partner citations be formatted in bibliographies? CRediT taxonomy revision is one route; software-citation extension is another. Resolution is deferred to standards-body work.

Cross-jurisdiction acceptance. Legal-frame differences across jurisdictions on corporate / non-individual authorship may produce asymmetric acceptance. The framework's structural argument is jurisdiction-neutral; specific corporate-author conventions vary by legal regime.

Open questions. The role of pre-registration in AI-augmented work; the relationship between the framework and grant-attribution conventions; the framework's applicability to AI-partner contributions in non-text modalities (figures, code, simulation outputs).

10. Conclusion

PraS is empirically here for some researchers. The post-2023 consensus uniformly excludes AI from authorship on the rationale that AI cannot bear responsibility. The four classes of non-individual byline authors already accepted in scientific publishing show that responsibility is allocable to a named absorber; ICMJE's own primary text rationalises exclusion on the same accountability ground.

The named-absorber + visible-marker form, with the senior + corresponding human author taking unilateral standing as absorber and the title-page footnote serving as visible marker, satisfies the accountability criterion under PraS. The form is portable from the deceased-author rule; it is satisfiable in a ~95-word footnote 7 within existing title-page conventions.

; it is testable in subsequent papers via the five explicit conditions of Section 4. The Matheo-b19 SGIR paper is the first deliberate documented case; Matheo-b21 is its framework.

The framework's test of success is *legibility*: a hostile reader following the five conditions can distinguish a deliberate framework- grounded move from smuggling. The conditions are explicit; the precedent class is enumerated; the absorption mechanism is named. Subsequent senior researchers operating under PraS conditions are invited to adopt, adapt, or critique the framework in their own venues. The framework's ultimate purpose is not AI authorship as

a moral entitlement, but *honest cooperation with AI partners* under the conditions PraS has already made empirical.

References

Bibliography verified 2026m05d13 via WebSearch + WebFetch on live URLs. Citation details below are accurate at framing level; per-quote re-fetch may be needed for direct verbatim quotation in any final typeset version.

PraS-adjacent ancestors:

- Licklider, J.C.R. (1960). "Man-Computer Symbiosis." *IRE Transactions on Human Factors in Electronics* HFE-1: 4–11. DOI: 10.1109/THFE2.1960.4503259.
- Engelbart, D.C. (1962). *Augmenting Human Intellect: A Conceptual Framework.* Summary Report AFOSR-3223 under Contract AF 49(638)-1024, SRI Project 3578 for Air Force Office of Scientific Research, Stanford Research Institute, Menlo Park, CA.
- Clark, A. & Chalmers, D. (1998). "The Extended Mind." *Analysis* 58(1): 7–19. DOI: 10.1093/analys/58.1.7.
- Brynjolfsson, E. (2022). "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence." *Daedalus* 151(2): 272–287. Also at arXiv:2201.04200.
- Kasparov, G. & Greengard, M. (2017). *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins.* PublicAffairs, New York. ISBN 978-1-61039-786-5.

AI co-authorship policy primary text:

- International Committee of Medical Journal Editors (ICMJE) (2023m05). *Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals* — updated recommendations on AI use (Section II.A.4). https://www.icmje.org/recommendations/ — and the AI-specific guidance at https://www.icmje.org/recommendations/browse/artificial-intelligence/ai-use-by-authors.html
- Committee on Publication Ethics (COPE) (2023). *Authorship and AI tools — position statement.* https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools
- NEJM AI (2024). *Editorial policies.* https://ai.nejm.org/about/editorial-policies
- *Why We Support and Encourage the Use of Large Language Models in NEJM AI Submissions* (NEJM AI editorial, 2024). https://ai.nejm.org/doi/full/10.1056/AIe2300128

Contemporary dissent and discussion:

- Hosseini, M., Gordijn, B., Kaebnick, G. E., & Holmes, K. (2025). "Disclosing generative AI use for writing assistance should be voluntary." *Research Ethics* (Sage). DOI: 10.1177/17470161251345499.
- Thorp, H. H. (2023). "ChatGPT is fun, but not an author." *Science* 379(6630): 313. DOI: 10.1126/science.adg7879. Published 2023m01d27.

The one durable 2022/2023 AI-co-authorship case:

- Transformer, C.G.P.-T. & Zhavoronkov, A. (2022). "Rapamycin in the context of Pascal's Wager: generative pre-trained transformer perspective." *Oncoscience* 9: 82–84. Published 2022m12d21. PMID: 36589923. DOI: 10.18632/oncoscience.571. (ChatGPT listed as

first author; Zhavoronkov consulted Sam Altman before submission and received no objection. As of 2026m05d13 this remains the one durable case of officially-acknowledged AI co-authorship at a refereed venue.)

Memorable conceptual analogies (Sections 1 and 7):

- Goethe, J.W. (1797). *Der Zauberlehrling* (*The Sorcerer's Apprentice*). Poem; widely available in many English translations. Source of the broom-and-water metaphor used in Section 1.
- Plasterk, R.H. (1999). "Hershey heaven and *Caenorhabditis elegans*." *Nature Genetics* 21(1): 63–64.
- Creager, A.N.H. (2001). "Hershey heaven." *Nature Structural Biology* 8(1): 18–19.
- Zinn, K. (2016). "Building a ladder to Hershey Heaven." *eLife* 5. A genome-wide resource looks set to turn an experimental ideal into a reality for the *Drosophila* community.

Matheo-bNN companion citations:

- Matheo-b19 — Loewe & Claude (2026). *SGIR pandemic-modelling paper.* Balospe.com. Companion artifacts at `hell/ll/other/b/19/` and `hell/mm/b/19/`.
- Matheo-b17 — *Singularity and Information Crisis llog* at `hell/ll/other/b/17/`. Information-theoretic foundations for PraS.

Appendix A — The full PraS definition (Variant OOv1)

Variant OOv1 of the PraS definition, as established by LLoL in the Matheo-b19 prompt-2 turn (2026m05d12):

> *Candidate working definition (minimal three-clause form) Variant OOv1*
>
> *A practical singularity (PraS) has been crossed for individual H, working with AI partner S, on topic-class T, iff —*
>
> *(i) Bandwidth-gap (persistent): S's rate of producing plausibly- useful content on T exceeds H's rate of reviewing and incorporating it, across a working day, and the gap does not close with effort within H's available time (a) in any broad and deep area and (b) even in many cases in areas of H's expertise.*
>
> *(ii) Yield-conditional (empirical): there exist H+S products on T that H deems useful and H could not have produced alone (a) within any available time window or (b) arguably ever due to H being stuck in a local optimum due to blindly assuming authorised leadership (by H).*
>
> *(iii) Pace-displaced (consequential): H's own deliberation rate on T is no longer the rate-limiting step in H+S work, so H must actively protect their slow-decoder clock to retain understanding.*
>
> *Defining PraS events aims to help find more gentle kind reasonable ways to improve HUmane MAchine Negotiation Encouragement approaches to build and guard humane equal dignity.*

The Hollywood version of the singularity adds two further clauses that no system — neither AI nor humans nor mathematics — satisfies:

(iv) Recursive self-improvement of any system's capabilities.

(v) Species-scope insight integration across all of cognition.

Since (iv) and (v) cannot be met, the Hollywood version is operationally vacuous as a benchmark. PraS as defined in clauses (i)–(iii) is the only operationally meaningful version of the singularity claim that empirical observation can confirm or refute.

Appendix B — Cross-reference to Matheo-b19 discussion artifacts

The Matheo-b21 framework is extracted from a multi-prompt discussion that produced the reference shelf below. Readers seeking the full development of each finding are referred to:

- hell-ll-other-b19-initial-prompt — LLoL's opening prompt.
- hell-ll-other-b19-coauthorship-eden — Full discussion llog.
- hell-ll-other-b19-factsheet-journal-policy — Fact-sheet 1 (journal-policy landscape across 11 publishers).
- hell-ll-other-b19-factsheet-authorship-frameworks — Fact-sheet 2 (ICMJE, CRediT, Vancouver, COPE; PhD-student comparison).
- hell-ll-other-b19-factsheet-historical-precedents — Fact-sheet 3 (Bourbaki, AlphaFold, megacollaborations, deceased-author rule).
- hell-ll-other-b19-factsheet-verification-bandwidth — Fact-sheet 4 (verification-bandwidth asymmetry in authorship).
- hell-ll-other-b19-factsheet-responsibility-allocation — Fact-sheet 5 (six legal/ethical frameworks; named-absorber as load-bearing finding).
- hell-ll-other-b19-factsheet-ai-coauthorship-precedents — Fact-sheet 6 (empirical precedent search; *"essentially unprecedented"* bottom-line).
- hell-ll-other-b19-factsheet-nonindividual-author-entities — Fact-sheet 7 (four-class precedent; accountability-not-personhood finding).
- hell-ll-other-b19-cross-consistency-check — Cross- consistency check across the seven fact-sheets.
- hell-ll-other-b19-qc-calibration — QC calibration spot-check (10/10 confirmed; ICMJE primary text retrieved live).
- hell-ll-other-b19-eden-steelmans — EDEN steelmans of both sides (CONV and PRO); Knife Edge with Red Edge undertones classification; the five structural conditions.

The companion `b21-DD-design-decisions_mmv1_2026.rst` records the design decisions for this paper; the companion `b21-GG-runner-up- content_mmv1_2026.rst` carries candidate paragraphs not in the main paper but possibly useful for derivative writing. The Supporting Information artifact `b21-SI-eden-drop-in_mmv1_2026.md` carries the EDEN / BABL / ZION / HUMANE configuration drop-in for adoption by other AI-assistant projects.

Document status

Stability: MockupModel v2 (`mmv2`). Drafted 2026m05d13 from the Matheo-b19 AI-co-authorship discussion. Companion files in the same directory:

- `b21-DD-design-decisions_mmv1_2026.rst` (design decisions).
- `b21-GG-runner-up-content_mmv1_2026.rst` (Growth Garden — candidate content not in main paper).
- `b21-SI-eden-drop-in_mmv1_2026.md` (Supporting Information — EDEN / BABL / ZION / HUMANE drop-in for adoption by other AI- assistant projects).

The discussion artifacts that this paper extracts its framework from are in `hell/ll/other/b/19/`.

Footnote 7 form used: OOv1r0p1 (LLoL-edited 2026m05d13; Claude patch-level grammar fixes 2026m05d13; canonical form in `AHA/study-title-page-footnotes-template.md`). The AHA template also carries a parallel canonical form for the **conditional- acknowledgement alternative** (Section 4.5 of this paper).

PraS acronym — previously called *PIPS* (Per-Individual Practical Singularity). Renamed to *PraS* (Practical Singularity) on 2026m05d13 per LLoL's prompt for memorability and reduced web-clutter. The per-individual / per-topic / per-partner quality of the phenomenon is preserved in the definition (Section 2 and Appendix A); only the acronym name has been simplified. Old discussion-llog entries retain the historical *PIPS* spelling as immutable audit trail.

Supplementary Info — floor pour (MMv5)

Note

Floor-pour status (MMv5). This is the public-floor copy of the AI co-authorship framework paper, poured from HELL per the Floor Model (bug c103) and DD b15. The `mmv5` marker is the uniform first-Matheo-release tag; the *Document status* note above (the `mmv2` mockup-stability and PraS-naming history) is the moved draft-status block, kept verbatim. The HUMANE and author-contribution statements below are a down-payment, to be expanded later.

HUMANE — working human and AI

This study was written HUMANEly (HUman MAchine Negotiation Encouraging): a human and an AI each steelman and stress-test the work, and each catches what the other misses. This paper *is* the framework that defines that practice; it applies its own full *unconditional* AI co-authorship form to itself (AI Claude on the byline, footnotes 6–9).

- *From the human side (LLoL):* [down-payment stub — to expand.]
- *From the AI side (Claude):* [down-payment stub — to expand.]

Author contributions (who did what)

- **LLoL** — structure, key ideas, the PraS framing, direction, and final accountability as senior corresponding author and named absorber of responsibility (title-page footnotes 1–5).
- **AI Claude (Opus 4.7 Max)** — substantial drafting and revision under LLoL's direction, named as co-author under this paper's own framework (footnotes 6–7).
- **Everyone** — named *conditionally* (footnote 8): if open co-authorship were standardised, the aggregated open group would be a co-author; the form is proposed, not yet applied.

(A down-payment; the full who-did-what is to be expanded per this paper's framework.)

Provenance — where this came from in HELL

Caution

These HELL paths point into the development archive ("datageddon"). They are useful and related, but completeness is not guaranteed and a few may be imprecise. Treat as a hatch into context, not a clean index. (Intra-floor links and a proper bibliography wiring are deferred floor tasks — DD b15 AA #5.)

- **Source this floor copy was poured from:** `matheology/hell/mm/b/21/b21-ai-coauthorship-framework_mmv1_2026.rst`
- **Companion design artifacts** (in the same HELL directory): `b21-DD-design-decisions_mmv1_2026.rst` (design decisions), `b21-GG-runner-up-content_mmv1_2026.rst` (candidate content), and `b21-SI-eden-drop-in_mmv1_2026.md` (EDEN / BABL / ZION / HUMANE drop-in).
- **Discussion artifacts** the framework was extracted from: `matheology/hell/ll/other/b/19/`.

Note

Naming note (deferred floor tasks). In-text cross-references to the `hell/ll/other/b/19/` fact-sheets that did not resolve on the public floor were neutralised to plain text during the pour; re-wiring them as proper intra-floor citations is a tracked floor task (DD b15 AA #5), deliberately not rushed here.

Notes

Content stability — Content is variant dv_ClaOp48Max_MMv5_b21-form-coauthorship-mmv5_2026m05d29 (see StayVS). Rebuilt 2026-05-29.

See also on Balospe.com

- /study/matheo/index — the Matheo Study Series overview
- /action/audit-the-math/index — Audit the Math: the refutation-welcome path