



Virosa Journal of Quantum Studies and Biomedicine



Q-ANASTASIS: Astroquantum Fubini Study/CTQW-MQ Walk-MBHA Virtual Screening of Cord-Blood Thaw Media with Flow-Cytometric Progenitor Validation

Ioannis Grigoriadis ¹²³⁴

¹ Biogenetoligandorol & Synthocure Stations, Thessaloniki, Greece.

² BiogenetoligandoroMACHNOT/QIICDNNDCANSTS Stations

³ Cellgenea Stem Cell Expansion Department, Anastasis Program.

⁴ Biogenea.

Keywords:

Umbilical Cord Blood, CD34, Hematopoietic Progenitor Cells, Cryopreservation, Thaw Optimization, Virtual Screening, Fubini-Study Geometry, Continuous-Time Quantum Walk, Mqwalk, Modified Black-Hole Algorithm, Ice-Recrystallization Inhibition, Antifreeze Proteins, Reproducibility.

ABSTRACT

Adult cord-blood transplantation remains constrained by the viable hematopoietic progenitor dose recovered after thaw, whereas bulk viability alone can obscure selective preservation or enrichment of the CD34-positive compartment. Q-ANASTASIS advances thaw-media evaluation through an integrated dose-centric and reproducibility-focused framework combining run-level flow-cytometric validation, database-scale virtual screening, Fubini-Study evidence geometry, classical diffusion, continuous-time quantum-walk (CTQW) transport, restart-stabilized MQWalk, and fixed-seed modified black-hole algorithm (MBHA) optimization. Five principal thaw conditions were compared using the effective progenitor payload $\langle P = S \times \eta \rangle$. ANASTASIS1 provided the reference (93.7% viability, 0.20% HPC, $\langle P = 0.001874 \rangle$); ANASTASIS5 achieved the highest payload gain (5.79-fold), ANASTASIS9 the strongest survival uplift (+5.9 percentage points; 3.56-fold gain), and ANASTASIS10 a high-performing hybrid profile (4.35-fold gain; +4.5 percentage points), demonstrating a clear dose-centric rank separation beyond viability-only assessment. Five Vitrifilaq/Stemorilaq databases contributed 1,658 screening records representing 840 unique DrugBank identifiers, including 41 candidates recurring across all five contexts. Robust evidence-state normalization generated a Fubini-Study graph with median nonzero distance $\langle \sigma = 1.2558 \rangle$. Vitrifilaq II showed the strongest upper-tail screening profile and remained the leading context across matched diffusion, CTQW, and MQWalk operators. Fixed-seed MBHA assigned diffusion/CTQW/MQWalk/source-stability weights of 0.2239/0.2239/0.2239/0.3282. Major methodological strengthening included explicit preprocessing, frozen parameters, convergence tracking, optimizer ablation, equal-weight controls, Borda aggregation, reciprocal-rank fusion, and perturbation testing. At $\langle \sigma = 0.10 \rangle$, top-context retention reached 0.9994 for MBHA and 0.9995 for equal weighting. Q-ANASTASIS therefore provides a quantitatively strengthened, auditable, and reproducible framework for progenitor-aware thaw ranking and formulation prioritization.

1. Introduction

1.1. Clinical and cryobiological rationale

Umbilical cord blood has an established place in allogeneic hematopoietic transplantation. The field does not require a new categorical proof that cord blood can reconstitute hematopoiesis, because this has been demonstrated across

foundational sibling, unrelated, pediatric, and adult transplant settings [1-8, 58-65]. The more precise translational problem is quantitative and operational: a banked cord-blood unit must deliver a sufficient viable hematopoietic progenitor payload after thaw, cryoprotectant dilution, and early post-thaw handling to remain clinically meaningful, especially for adult recipients, in whom cell-dose constraints are more stringent [3-8].

* Corresponding author: Ioannis Grigoriadis, & Synthocure Stations, Thessaloniki, Greece, E-mail address: jgrigoriadis@biogenea.gr

Received: 6 July 2026; **Revised:** 1 September 2026; **Published:** on 27 September 2026

Citation: Grigoriadis I. Q-ANASTASIS: Astroquantum Fubini Study/CTQW-MQ Walk-MBHA Virtual Screening of Cord-Blood Thaw Media with Flow-Cytometric Progenitor Validation. Virosa J Quantum Biol Biomed (VJQBB). 2026;1(1):1-24.

Copyright: © 2026 Ioannis Grigoriadis. Published by Virosa Publishing. This is an open-access article distributed under the Creative Commons Attribution 4.0 International License ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited

The study therefore does not frame cord-blood transplantation as a binary question of engraftment versus non-engraftment. Instead, it addresses the post-thaw preservation of a biologically relevant dose compartment. In this context, the central quantity is not bulk nucleated-cell survival alone, but the recoverable, viable, progenitor-enriched fraction that remains after the thaw interval. CD34-positive hematopoietic progenitor enumeration, especially within an ISHAGE-style CD45-dim gating context, provides a practical and clinically interpretable anchor for this dose-centric analysis [9,10]. Bulk viability, while necessary, is insufficient as a single decision variable because a preparation can retain many viable leukocytes yet lose or dilute the progenitor-enriched compartment that carries the intended graft function [66,67].

This study therefore treats the thaw interval as an intervention window. The thaw step is not merely a transfer from storage to use; it is a coupled thermal, osmotic, and biochemical transition in which progenitor payload can be preserved, depleted, redistributed, or selectively enriched relative to other post-thaw compartments [11-23]. The Q-ANASTASIS workflow is built around this premise. It asks whether different ANASTASIS-class thaw conditions produce different effective progenitor-payload profiles when interpreted at the run level, and whether computational ingredient-prioritization methods can be used as auditable design tools without being misrepresented as biological evidence.

Many cryopreservation and thaw studies report post-thaw viability as a primary indicator of sample quality. Viability remains essential, but it is not sufficient for adult cord-blood interpretation. Adult transplantation imposes a dose problem: the unit must deliver enough functional hematopoietic progenitor material after thaw to support clinically relevant hematopoietic recovery [3-8]. A high viability percentage can coexist with a low CD34-positive progenitor fraction, while a modest reduction in global viability may be accompanied by a relative preservation or enrichment of the progenitor compartment. For this reason, viability-only interpretation may obscure biologically meaningful differences among thaw conditions.

The dose-centric approach used here separates three related but non-identical concepts: survival, progenitor representation, and effective delivered payload. Survival describes the fraction of cells or events that remain viable after thaw. Progenitor representation describes the proportion of the relevant hematopoietic progenitor compartment, operationalized here through CD34-positive/CD45-dim hematopoietic progenitor-cell frequency. Effective delivered payload combines these quantities into an interpretable product score. The score is exploratory and does not replace absolute CD34-positive cell-dose determination, colony-forming assays, potency testing, or clinical engraftment endpoints. However, it provides a transparent run-level metric for comparing thaw conditions when the same internal baseline is used.

The biological unit of inference is therefore the thawed run,

thawed cord-blood unit, or donor-linked bag-equivalent object. Individual flow-cytometry events, bead events, gated subpopulations, database entries, graph nodes, docking poses, and quantum states are nested technical structures. They may support assay fidelity, computational traceability, or ranking reproducibility, but they do not become independent biological replicates. This distinction is central to the analysis because it prevents pseudoreplication and preserves the clinical interpretation of the thawed unit as the relevant biological object [9,10, 50-52].

Classical cryobiology provides the physical basis for treating thaw as an active intervention interval. Cell recovery after cryostorage is influenced by water transport, intracellular ice formation, solute effects, osmotic gradients, glass-transition behavior, devitrification, ice recrystallization, cryoprotectant toxicity, dilution kinetics, and warming rate [11-23]. These mechanisms are not independent. During rewarming, a sample may traverse temperature regions in which ice growth, devitrification, or osmotic injury becomes more likely. During cryoprotectant dilution, cells may experience osmotic stress, membrane perturbation, and solute toxicity. During early post-thaw handling, transient metabolic and membrane instability may further influence recovery [71-80].

Warming rate is a central determinant of post-thaw recovery in many cryobiological systems [15-19]. Slow or irregular warming may increase the time spent in thermal regions where devitrification or recrystallization can occur. In vitrified or partially glass-forming systems, the traversal of the glass-transition neighborhood is particularly important because molecular mobility increases as the sample warms, and ice-crystal growth can become kinetically favorable if the sample remains too long in the relevant temperature band [15-23]. Even in non-vitreous cryopreservation, ice recrystallization during warming can damage membranes, alter local solute concentrations, and compromise recovery.

Cryoprotectant removal is another release-critical step. DMSO and other cryoprotectants can protect cells during freezing but may contribute to toxicity, osmotic stress, and membrane perturbation during thaw and dilution [11-23]. A rapid dilution may impose osmotic shock; a slow dilution may prolong exposure to cryoprotectant toxicity. The optimal washout profile is therefore a controlled compromise, not a generic rinse step.

CD34-positive cell enumeration is a standard component of hematopoietic graft assessment, and ISHAGE-style single-platform approaches provide a widely used foundation for reproducible CD34-positive counting [9,10]. Within this framework, the CD45-dim compartment helps distinguish hematopoietic progenitor-enriched events from more mature leukocyte populations. The Q-ANASTASIS analysis uses CD34-positive/CD45-dim hematopoietic progenitor frequency as a practical surrogate for progenitor representation.

This surrogate has boundaries. It does not prove long-term engraftment, stem-cell self-renewal, multi lineage output, or in

vivo potency. It also does not replace colony-forming assays, long-term culture-initiating cell assays, xenograft assays, or clinical endpoints. However, it is a relevant translational readout because it connects post-thaw sample quality to the progenitor compartment that is central to graft function [3-10]. The analysis therefore treats CD34-positive/CD45-dim frequency as a run-level biological signal, not as a complete potency claim.

A viability-only ranking would prioritize conditions that maximize total viable events or viable leukocyte recovery. This is useful for identifying gross cytotoxicity, but it may miss conditions that better preserve the progenitor-enriched compartment. In cord-blood transplantation, the relevant biological question is not simply whether cells survive, but whether the surviving compartment contains enough functional progenitor material to support the intended graft purpose [3-10]. Therefore, the payload score is designed to penalize conditions that preserve bulk viability without preserving progenitor representation and to identify conditions where progenitor enrichment offsets modest survival changes.

This is illustrated by the ANASTASIS run structure. ANASTASIS2 has lower viability than the baseline but a higher HPC fraction, producing a >2-fold gain in effective-payload score. ANASTASIS5 has both high viability and a strong HPC fraction, producing the leading product score. ANASTASIS9 has the clearest survival uplift and remains biologically strong, even though its HPC fraction is lower than that of ANASTASIS5. ANASTASIS10 shows a hybrid profile, with high viability and strong progenitor representation but a survival uplift just below the predefined gate.

These observations define the central experimental question: a thaw condition should be judged by the biologically relevant compartment that survives the transition, not by a single global viability percentage. In the ANASTASIS run set, that principle is tested directly by comparing survival with the CD34+/CD45dim progenitor fraction and their product, while retaining donor/run identity as the inferential unit.

1.2. IRI/AFGP-inspired and small-molecule formulation space

The Vitrifilaq and Stemorilaq layers are introduced as ligand-family and media-design audit labels. They are not represented as validated drugs, approved clinical media, or completed release formulations. Their purpose is to organize ingredient-prioritization hypotheses around database-derived records, motif annotations, recurrence patterns, and graph-audited ranking behavior. IRI/AFGP-inspired and small-molecule formulation space.

The rationale for this layer comes from cryobiological and chemical principles. Ice-binding proteins, antifreeze glycoproteins, and synthetic IRI-active polymers suggest that molecular motifs capable of modulating ice recrystallization, interfacial water, hydration shells, or local solute behavior may be relevant to thaw-media design [24,25,48,49]. Motif similarity is used to prioritize candidates for concentration-response,

osmolality, pH, wash-compatibility, and safety testing, providing a structured path from computational screening to empirical qualification [80-83].

Public database-derived records from DrugBank, PubChem, and ChEMBL contexts are useful for harmonization because they provide identifiers, compound names, structures, SMILES strings, and bioactivity or annotation contexts [55-57]. These inputs enable candidate deduplication, recurrence analysis, motif scoring, and graph construction, thereby creating a prioritized feeder set for experimental evaluation of Vitrifilaq/Stemorilaq thaw-media candidates.

The chemical search space spans classical penetrating cryoprotectants, non-penetrating carbohydrates, interfacial modifiers, redox-active phytochemicals and progenitor-support cues. Type-I antifreeze proteins and antifreeze glycoproteins provide reference structural motifs for ice-binding and recrystallization control [48,49, 81-91], while DMSO, trehalose, silibinin, asiatic-acid-like triterpenoids and database-derived analogues define chemically diverse hypotheses for controlled experimental follow-up. Their representation in the virtual-screening layer reflects structural context and database proximity rather than an assumption of equivalent mechanism. The computational-translational architecture is summarized in Figures 1-3; database-linked structural and scaffold evidence is developed in the screening analysis below.




Figure 1: Quantum-geometric architecture for thaw-media prioritization. Numbered modules follow the labels embedded in the composite: (1) Fubini-Study evidence geometry; (2) continuous-time quantum-walk (CTQW) transport on a one-dimensional chain; (3) restart-stabilized MQWalk occupancy on a two-dimensional lattice; and (4) translational antifreeze-glycopeptide/polymer chemistry. The central circuit proceeds through initialization, CTQW transport, MQWalk transport, MBHA-assisted selection, and measurement. The lower MBHA schematic resolves the optimizer into four substeps: (1) candidate solution space, (2) attractor-based stochastic optimization, (3) event-horizon regeneration, and (4) state transitions and convergence. The lower-left performance-versus-iterations plot is an illustrative convergence schematic rather than an experimental result. Biological interpretation remains separate and is anchored to the measured dose-centric payload $P=S \times \eta$ and CD34+/CD45-dim progenitor preservation.




Figure 2: Evidence-to-translation quantum-walk and cryochemistry framework.

Peripheral numbered panels are (1) Fubini geometry, (2) lattice transport, (3) topology/geodesic mapping, (4) polymer/AFGP chemistry, and (5) translational applications. The central operator chain contains five serial stages: (1) evidence-state embedding, (2) Fubini geometry, (3) CTQW transport, (4) MQWalk persistence, and (5) MBHA optimization and selection. The circular MBHA engine is resolved into subpanels A-E: (A) stochastic state population, (B) black-hole attractor, (C) event-horizon replacement, (D) convergence, and (E) state measurement and selection. Across the five screening contexts, 1,658 source records yielded 840 unique DrugBank identifiers, including 41 candidates present in all five contexts; Vitrifilaq II had the highest single similarity (0.145), top-10 mean (0.0808), and mean top-20 Jaccard overlap (0.1124). The ice-recrystallization-inhibition application curves are schematic design illustrations and are not reported as measured cryoprotective efficacy.




Figure 3: Figure 2 Evidence-to-translation quantum-walk and cryochemistry framework.

Numbered modules are (1) Fubini manifold and metric representation, (2) the core state-preparation-CTQW-MQWalk-MBHA circuit, (3) the MBHA biotranslational optimization loop, (4) transport readouts, (5) the ice-biomolecule/antifreeze interaction hypothesis, and (6) the translational chemistry pipeline. Within module 3, subpanels 1-5 depict candidate states, fitness evaluation, selection and update, event-horizon regeneration, and translational decision. Module 4 juxtaposes a CTQW molecular-graph occupancy map with an MQWalk/CTQW probability trace. MBHA is interpreted as a fixed-seed stochastic outer-loop optimizer, not as a deterministic biological mechanism. The depicted quantum-system and ice-interaction elements are computational/design schematics; experimental validation remains classical and assay based.

1.3. Astroquantum computational-biology selector

Computational screening becomes useful only when it is

auditable. The Q-ANASTASIS workflow therefore requires explicit records, declared descriptors, and reproducible graph construction. Each candidate can be represented by a descriptor vector encoding features such as sulfoxide or DMSO-like polarity, polyether spacing, polyol content, amide junctions, terminal caps, hydrophobe burden, heteroatom distribution, recurrence across databases, and guardrail compatibility. These descriptors are not treated as biological measurements; they are computational annotations that support ranking.

A candidate similarity graph can then be constructed from normalized descriptor similarity, recurrence weighting, and score proximity. The graph structure supports three levels of analysis. First, recurrence and consensus ranking identify candidates that appear robust across screening contexts. Second, classical diffusion provides a non-quantum graph baseline for transport, accessibility, and neighborhood influence. Third, quantum-walk overlays provide topology-sensitive audit readouts that may reveal localization or ranking sensitivity under alternative transport assumptions.

Quantum-walk overlays cannot be interpreted unless the graph has been declared and the classical diffusion baseline has been computed. This prevents the quantum layer from becoming decorative language. It also makes the audit reproducible: another investigator can reconstruct the graph, run diffusion, apply CTQW or DTQW overlays, and compare rank stability or displacement across ablation variants.

Classical graph diffusion is the required first transport model because it provides an interpretable non-quantum baseline [29,34]. On a declared graph with normalized Laplacian, diffusion can be written as: [87]

Only after this step can topology-sensitive overlays be introduced. If a candidate is highly ranked under consensus, recurrence, and diffusion, it is more stable than a candidate that appears only under a quantum-walk overlay. Conversely, a candidate that is promoted only by CTQW or DTQW may still be interesting, but it must be treated as a topology-sensitive hypothesis rather than a robust priority. This logic supports reproducibility and guards against over interpretation.

Continuous-time quantum walks and discrete-time quantum walks are mathematical tools for studying transport on graphs [30-33, 40,41]. They can produce behavior that differs from classical diffusion, including interference-like localization, oscillatory recurrence, and topology-sensitive accessibility. These properties can be useful in audit settings because they reveal whether a ranking depends strongly on graph structure [88-91].

In this study, CTQW and DTQW are implemented as ranking and topology-audit overlays on declared graphs. Their role is to quantify transport-sensitive candidate accessibility, localization, and rank stability in a reproducible computational architecture linked to the biological prioritization workflow.

These equations define computational audit operations. They do not define biological mechanism. The analysis uses them to test

whether Vitrifilaq/Stemorilaq candidate rankings are stable under different graph-transport assumptions.

Within this study, astroquantum computational biology is used as a compact descriptor for a deliberately hybrid algorithmic architecture: evidence states are embedded in Fubini-Study geometry; transport is evaluated by classical diffusion and quantum-walk operators on the same graph; restart regularization creates an MQWalk persistence layer; and a black-hole-inspired stochastic optimizer treats high-performing weight configurations as attractor states. The terminology concerns mathematical representation and computational search. The experimentally measured cell populations remain classical biological observations.

1.4. Rewarming, devitrification, and osmotic recovery as coupled design constraints

Rewarming is not a passive reversal of cooling. During thaw, the sample traverses a temperature interval in which residual ice can recrystallize, vitrified domains can devitrify, viscosity changes rapidly and the osmotic relationship between extracellular solute and intracellular water is reset. Classical cryobiology therefore predicts that nominally similar post-thaw viability can arise from different combinations of thermal injury, solute exposure, and selective survival of cell subpopulations [11-23]. For hematopoietic grafts, this matters because a small change in the survival of a rare progenitor compartment can alter the effective dose even when the dominant leukocyte population appears well preserved.

The ANASTASIS design language treats this interval as an experimentally controllable trajectory. A complete process record should contain the starting storage state, warming-rate trace, time to visible melt completion, any dwell in a devitrification-prone interval, time from melt to dilution, the sequence and duration of cryoprotectant washout, final pH and osmolality, and the elapsed interval to biological sampling. These variables are not used to retrofit an explanation to the five run-level results; they define the prospective process measurements required to distinguish a reproducible thaw effect from an unmeasured handling difference.

Rapid warming can reduce the time available for damaging ice growth, but a purely thermal interpretation is incomplete. Cryoprotectant concentration, membrane permeability, and osmotic relaxation determine how the cell experiences the transition after melting. DMSO removal that is too rapid can create a large transmembrane water flux, whereas prolonged exposure can increase chemical toxicity. A staged dilution is therefore represented as an osmotic relaxation problem rather than as a single wash event. In a future qualification study, the integrated osmotic deviation can be calculated from the measured trajectory and related to run-level survival and progenitor preservation without assuming a universal optimum before data are collected.

$$H_{\text{thaw}} = \int [h_T(T(t)) + h_{\Pi}(\Pi(t)) + h_C(C_{\text{CPA}}(t))] dt$$

$$D_{\Pi} = \int_0^{t_{\text{wash}}} |\Pi(t) - \Pi_{\text{target}}| dt$$

The two equations are process ledgers rather than fitted mechanistic models. They make explicit that a favorable biological run should be reproducible under a documented thermal and osmotic path. This distinction is essential for translation: an empirically strong A5 or A9 result can motivate replication, but process qualification requires the associated trajectory to be measured, bounded and repeated across independent units.

1.5. Biological preservation, progenitor enrichment and expansion are distinct endpoint

Post-thaw progenitor enrichment can arise from at least three non-exclusive processes: preferential survival of CD34-positive cells, preferential loss of mature cells from the denominator, or genuine preservation of absolute progenitor number. The current payload score is intentionally agnostic among these mechanisms because the available source set is strongest for viability and percentage-based HPC measurements. Absolute viable CD34-positive dose and recovery relative to the pre-freeze unit should therefore be included in the next replication phase so that enrichment and preservation can be separated quantitatively.

The same boundary applies to hematopoietic stem-cell expansion. Components such as UM171, SR1, prostaglandin-related cues and metabolic support factors are biologically relevant to progenitor maintenance or ex vivo expansion [47,92-97], but their presence in a thaw-media design does not by itself demonstrate expansion. An expansion claim requires a defined culture interval, starting absolute cell number, recovered absolute cell number, phenotype, functional output and appropriate untreated or baseline control. Q-ANASTASIS consequently sets the expansion multiplier to unity in the primary thaw comparison and treats expansion as a prospective orthogonal validation arm.

2. Study Design and Evidence Architecture

2.1. Evidence classes and unit of inference

This study reports experimentally observed in vitro results. Its primary biological evidence consists of post-thaw viability measurements and flow-cytometric enumeration of CD34-positive/CD45-dim HPC populations in thawed cord-blood runs. These assays validate the relative biological performance of the tested thaw conditions at the measured in vitro endpoint level. Chemical-screening and graph-transport analyses extend the evidence base by prioritizing ingredient families and auditing rank stability on declared graphs.

This study is framed as an exploratory translational analysis of ANASTASIS-class cord-blood thaw conditions. The objective is to compare principal ANASTASIS thaw conditions against an internal ANASTASIS1 baseline using run-level biological metrics and to define a reproducible computational audit layer for Vitrifilaq/Stemorilaq ingredient prioritization.

The study contains three linked but distinct components. The first is the biological comparison of thawed runs using viability and CD34-positive/CD45-dim hematopoietic progenitor frequency.

The second is the release-governance framework, which defines the telemetry, washout, buffer, and QC-to-bag requirements needed for formal SOP PASS assignment. The third is the computational screening framework, which uses public database-derived candidate records, motif annotation, graph construction, classical diffusion, and quantum-walk overlays for ranking audit.

The central methodological rule is that these components must remain distinct. Biological ranking can support exploratory interpretation. Telemetry can support process readiness. Screening can support candidate prioritization. Graph audit can support reproducibility and ranking sensitivity. None of these alone supports formal release.

The independent biological unit was the thawed run, thawed cord-blood unit, or donor-linked bag-equivalent object. Flow events, bead counts, gated populations, graph nodes, docking poses, and quantum states were treated as nested technical objects and were not counted as independent biological replicates. This rule is central because FlowCyto material contains many event-level summaries, while the release question is assessed at the run and lot-governance level [9,10, 50-52]; see SI Appendix II.

S2 Appendix, Figure 1 and FlowCyto Schemes 2-50 document the ISHAGE-style CD34 enumeration, CD45/side-scatter leukocyte and lymphocyte structure, 7-AAD viability discrimination, CAL980 bead support, and CD34+/CD45-dim HPC gates across ANASTASIS1-10. The run mapping is explicit: ANASTASIS1, SFigure 1 and FlowCyto Schemes 2-5; ANASTASIS2, Schemes 6-10; ANASTASIS3, Schemes 11-15; ANASTASIS4, Schemes 16-20; ANASTASIS5, Schemes 21-25; ANASTASIS6, Schemes 26-30; ANASTASIS7, Schemes 31-35; ANASTASIS8, Schemes 36-40; ANASTASIS9, Schemes 41-45; and ANASTASIS10, Schemes 46-50. These objects provide figure-level assay provenance and gating traceability while the thawed run remains the biological unit of inference.

Each run was identified by a run label, associated media class, viability percentage, HPC percentage, CD34 percentage where available, lymphocyte context where available, gating quality, and bead-support status. A run was considered assay-admissible when the core viability and HPC fields were present and when the source documentation supported the relevant gating interpretation.

Table 1: Evidence classes, source objects and permitted interpretation.

Evidence class	Primary source	Analytical role	Permitted conclusion
Measured biology	S2 flow-cytometry reports; run ledger	Viability, leukocyte architecture, CD34+/CD45dim HPC	In vitro run-level outcome
Derived biology	Measured S and η	P=S×η; gain vs A1; rank displacement	Dose-centric comparative endpoint
Screening database	Supplementary Databases I-V	DrugBank ID, similarity, name, SMILES, recurrence	Virtual-screening nomination
Geometry/transport	Database-context descriptors	Fubini-Study, diffusion, CTQW, MQWalk	Topology-sensitive audit
Optimization	Frozen selector channels	Fixed-seed MBHA and classical aggregators	Reproducible fusion of declared channels
Qualification	Prospective	Thermal, osmotic, buffer,	Release decision only

	telemetry/QC	QC-to-bag, governance	after all gates
--	--------------	-----------------------	-----------------

2.2. Flow-cytometry biological validation

ANASTASIS1 served as the internal biological baseline. ANASTASIS2, ANASTASIS5, ANASTASIS9, and ANASTASIS10 formed the principal upgraded comparison set for run-level translational evaluation. The full ANASTASIS1-10 formulation atlas is retained as the main formulation table.

Direct in vitro cytometry evidence is traceable to S2 Appendix. For ANASTASIS1, SFigure 1 and FlowCyto Schemes 2-5 document 7-AAD viability discrimination, CD45/side-scatter structure, CD34-positive/CD45-dim HPC resolution, CAL980 bead control, 93.7% viability, 30,332 leukocyte events, 10,556 lymphocyte events, and a low-frequency HPC compartment of approximately 0.19-0.20%. FlowCyto Schemes 6-10, 21-25, 41-45, and 46-50 document the corresponding gating architecture for ANASTASIS2, ANASTASIS5, ANASTASIS9, and ANASTASIS10, respectively. S3 Appendix independently tabulates the same baseline and comparative run logic, linking the cytometric source figures to the survival, HPC, payload, and qualification ledgers used here.

Flow cytometry used 7-AAD viability discrimination, CD45/side-scatter leukocyte architecture, CD34-positive enumeration, CD34-positive/CD45-dim progenitor gating, and bead-supported counting where available [9,10]. S2 Appendix, SFigure 1 and FlowCyto Schemes 2-5 provide the baseline A1 gating series; Schemes 21-25 and 41-50 provide the principal A5, A9, and A10 result-level gating series. These supplementary figures are cited as assay-source evidence rather than independent replicates.

The analysis follows an ISHAGE-style interpretive logic. CD45/side-scatter structure supports leukocyte and lymphocyte architecture. 7-AAD supports viability discrimination. CD34 positivity identifies progenitor-associated events. CD45-dim structure helps define the hematopoietic progenitor-enriched compartment. Bead support enables absolute or semi-absolute counting where applicable.

CD34-positive concentration is interpreted using the ratio of gated CD34-positive events to calibrated bead events multiplied by the bead concentration per acquisition volume, consistent with ISHAGE-guided single-platform enumeration [9,10]. This relation supports source fidelity and does not convert nested events into independent biological replicates.

The five-run comparison is therefore an in vitro proof-of-concept at the level of the available run ledger. The analysis does not infer biological replication from the number of gated events. Where absolute or bead-supported counts are available, they strengthen traceability but remain nested within the same run-level experimental object.

2.3. Dose-centric endpoint and mathematical properties

Relative performance was evaluated against the ANASTASIS1 internal DMSO baseline rather than against external thaw protocols. The run-level dose ledger identifies ANASTASIS5 as the payload-leading condition ($G = 5.79$; $\text{sqrt}(G) = 2.41$),

ANASTASIS10 as the high-scoring hybrid condition ($G = 4.35$; $\sqrt{G} = 2.09$), and ANASTASIS9 as the survival-gated condition ($G = 3.56$; $\sqrt{G} = 1.89$; survival uplift = +5.9 percentage points) (Table 2).

The effective-payload score is an exploratory comparative measure. It does not claim to measure absolute transplant potency. It is used to compare thaw conditions relative to ANASTASIS1 under the available data constraints.

For run r , let $S_r \in [0,1]$ denote the observed post-thaw viable fraction, $\eta_r \in [0,1]$ the observed CD34+/CD45dim progenitor fraction within the declared leukocyte denominator, $D_{0,r}$ the pre-thaw progenitor dose or calibrated proxy, and X_r an independently validated expansion multiplier. The effective delivered progenitor quantity and the executable run-level payload are defined by:

In the present five-run analysis $X_r = 1$ because no independently qualified ex vivo expansion step was part of the thaw comparison. P_r is therefore a dimensionless progenitor-weighted payload index, not an absolute CD34+ dose or a clinical potency assay. The standardized values are given in Table 2 and are used consistently in the Results, biological ablation, and validation checks; the corresponding baseline and qualification equations are expanded in S3 Appendix.

$$P_r = S_r \eta_r, \quad G_r = P_r / P_{A1}, \quad \Delta S_r = 100(S_r - S_{A1}) \text{ percentage points}$$

The product is monotone in each non-negative factor and invariant to a common positive calibration multiplier applied across all η values. More importantly, it provides an exact criterion for a rank reversal between viability and payload.

$$\text{For } S_i > S_j: \quad P_i < P_j \Leftrightarrow \eta_j / \eta_i > S_i / S_j$$

For $\eta_r \geq 0$, P_r is non-decreasing in S_r ; for $S_r \geq 0$, P_r is non-decreasing in η_r . If the other factor is strictly positive, the increase is strict.

Suppose every progenitor fraction is multiplied by the same positive calibration constant c , so that $\eta'_r = c\eta_r$. Then all fold gains remain unchanged: $G'_r = G_r$.

For two runs i and j with $S_i > S_j$, the payload ranking reverses the viability ranking, $P_i < P_j$, if and only if $\eta_j / \eta_i > S_i / S_j$.

The ANASTASIS1-ANASTASIS2 comparison satisfies this condition: $93.7\% > 90.5\%$, but $0.43/0.20 = 2.15$ exceeds $93.7/90.5 = 1.035$, so the lower-viability run has the higher progenitor-weighted payload. This exact inequality supplies the mathematical basis for the difference between viability ranking and payload ranking reported in Section 6.2 and is concordant with the comparative run analysis in S3 Appendix.

The exploratory gain threshold was set at $G \geq 1.25$ relative to ANASTASIS1, matching the qualification logic documented

in S3 Appendix. The threshold identifies a clear product-score improvement over the internal baseline while remaining separate from process release qualification.

The survival-uplift gate was set at the predefined +5 percentage-point improvement relative to ANASTASIS1. This threshold identifies survival-leading candidates and positions ANASTASIS9 for process-qualified follow-up evaluation.

Table 2: Principal ANASTASIS run-level biological outcomes used in the dose-centric comparison.

Run	Viability %	HPC %	$P=S \times \eta$	Gain	ΔS pp	Interpretation
A1	93.7	0.20	0.001874	1.00	0.0	Internal reference
A2	90.5	0.43	0.003892	2.08	-3.2	Progenitor-enriched despite lower viability
A5	95.2	1.14	0.010853	5.79	+1.5	Payload-leading condition
A9	99.6	0.67	0.006673	3.56	+5.9	Survival-leading condition
A10	98.2	0.83	0.008151	4.35	+4.5	High-performing hybrid

Cusps record the degree of binding between flows, Zero cusps means four independent interactions (PVI), while accumulation of cusps signals progressive constraint or collapse of the interaction structure. The key result exploited here is: $PIII(D_6)$ and PV_{deg} share the same underlying character variety but carry different cusp signatures, meaning they are geometrically indistinguishable at the algebraic level yet dynamically distinct. This fact has profound biological consequences, as we shall see.

2.4. Assay provenance, gating consistency and denominator control

The biological ledger is derived from de-identified flow-cytometry reports linked to individual ANASTASIS runs. For the baseline A1 condition, the reports show 7-AAD viability discrimination, CD45/side-scatter leukocyte structure, a resolved CD34-positive/CD45dim progenitor gate and bead-supported acquisition. Equivalent source series are available for the principal upgraded runs. The analytical requirement is not that every report have an identical number of acquired events, but that the same biological denominator and gate meaning be applied when percentages are compared.

A percentage is meaningful only with a declared denominator. The HPC variable used in the payload calculation is the CD34-positive/CD45dim fraction in the viable leukocyte compartment as reported by the source analysis. If a future run reports CD34-positive percentage using a different parent gate, that value must be remapped or excluded rather than inserted directly into the same formula. Bead counts support absolute enumeration where calibration information is available, but they do not convert technical events into independent biological replicates.

The source material also contains lymphocyte proportions and broader leukocyte architecture. These variables are retained as contextual readouts rather than hard components of $P=S\eta$ because their relationship to hematopoietic stem/progenitor potency differs from that of the CD34-positive compartment. A future multivariate endpoint could combine absolute viable CD34-positive dose, primitive phenotype, lymphocyte preservation, and

CFU output, but such a model should be trained and validated on a donor-balanced dataset rather than defined retrospectively on five runs.

2.5. Process-qualification guardrails

Biological ranking and release qualification are deliberately separated. The exploratory survival-uplift criterion is +5 percentage points relative to A1 and the payload criterion is $G \geq 1.25$, but neither is sufficient for release. A process-qualified run additionally requires a documented warming trajectory, acceptable time through the devitrification-relevant interval, controlled cryoprotectant dilution, a final physicochemical corridor, concordance between the sampled QC segment and the full bag, sterility/endotoxin evidence, and completed governance. Missing process telemetry therefore produces a hold state rather than being imputed from a favorable cytometry result. The complete thermal-traverse, osmotic-washout, buffer-corridor, and QC-to-bag requirements are detailed in S3 Appendix under Outstanding requirements and the extended validation algorithm.

This non-compensatory logic is particularly important for computationally prioritized media. A graph or MBHA score can nominate a chemical neighborhood for testing, but it cannot rescue a run that violates a basic physical or biological requirement. Conversely, a strong biological run can be advanced for process replication even if its ingredients were not top-ranked computationally. The two directions of evidence remain independent until a prospective experiment links a candidate formulation to a controlled thaw trajectory and a prespecified biological endpoint.

3. Thaw-Media Design Space and Database Virtual Screening

3.1. ANASTASIS media classes

The ANASTASIS1 reference payload was derived from $S_0=0.937$ and $\eta_{a0} \approx 0.0020$, giving $P_{A1}=0.001874$. ANASTASIS1 showed high viability, stable leukocyte recovery, bead-supported acquisition, and a low but resolved CD34+/CD45-dim progenitor compartment, consistent with ISHAGE-style enumeration requirements [9,10,66,67]. The baseline calculation, comparator logic, and process guardrails are documented in S3 Appendix, whereas the underlying gating evidence is provided in S2 Appendix, SFigure 1 and FlowCyto Schemes 2-5.

The Vitrifilaq-associated classes are conceptually supported by cryobiology, vitrification, ice-recrystallization inhibition, antifreeze-protein diversity, and synthetic macromolecular ice-growth control literature [11-25,48,49,71-83]. Their computational screening layer is anchored to DrugBank, PubChem, ChEMBL, molecular fingerprinting, and similarity-kernel literature [55-57,103-107]. The supplementary chemistry atlas provides the direct visual bridge: S1 Appendix, Figure 1 links N-aryl-D-aldonamide IRI chemistry to functional post-thaw readouts; Figure 5 summarizes antifreeze-protein mechanism and structural diversity; Figures 7-17 develop the Vitrifilaq/Stemorilaq scaffold and permitted-geometry design space; and Figures 18a-18b provide docking-ledger exemplars.

These supplementary objects motivate the screening architecture used in the main analysis.

The Stemorilaq-associated design logic is linked to primitive HSPC sustainment and expansion-adjacent biology, including UM171, SR1/AHR antagonism, prostaglandin-modulated cord-blood transplantation, and related HSPC-maintenance literature [46,47, 84-86]. Stemorilaq-related scaffold, motif, and quantum-polygraph annotations are shown in S1 Appendix, Figures 7-17. The associated media bills of materials are tabulated in S1 Appendix, Tables 1-5 and Table X, including the ANASTASIS3/Stemorilaq planning table (Table 5), while S3 Appendix retains the run-level validation and qualification logic.

This comparison clarifies that no single cryoprotectant property is sufficient for interpretation. Penetration, osmotic stress, toxicity, viscosity, dilution behavior, membrane compatibility, redox response, progenitor preservation, and post-thaw function must all be assessed at the run level. Classical cryoprotectant comparison classes are interpreted against established water-transport, intracellular-ice, solute-effect, vitrification, warming-rate, and devitrification principles [11-23,71-79].

The media classes are treated as formulation families rather than released products. Their role is to establish a chemically interpretable design space that can be connected to measured thaw outcomes and then refined through database screening. A candidate may be structurally interesting, recurrent across source databases or graph-central without being biologically qualified; empirical concentration-response, osmolality, pH, hemocompatibility, sterility and functional progenitor testing remain decisive. S1 Appendix, Tables 1-5 and Table X provide the detailed planning compositions and readout map used to construct the main formulation atlas, while S1 Appendix, Figure 19 provides the supplementary multi-run biological ranking context.

3.2. Screening database corpus and harmonization

Five semicolon-delimited CSV files constituted the screening corpus: Supplementary Database I (Vitrifilaq I), Supplementary Database II (Vitrifilaq II), Supplementary Database III (Vitrifilaq III), Supplementary Database IV (Vitrifilaq IV), and Supplementary Database V (Stemorilaq I). Each file contains DrugBank identifier, similarity score, usual name, and SMILES. Parsing retained all 1,658 rows and yielded 840 unique DrugBank identifiers after cross-context union. Candidate recurrence was 415 candidates in one database, 181 in two, 136 in three, 67 in four, and 41 in all five databases. Scores were normalized only within source context before aggregation; raw similarity values were not treated as if they arose from a common calibrated scale. Database provenance, graph-domain traceability, scaffold-level interpretation, and harmonized screening summaries are presented in Figures 4-6, with the exact file-level hashes reported in Table 14.

Table 3: Exact summary statistics derived from Supplementary Databases I-V.

Context	Rows	Mean	Median	Top 1	Top-10 mean	Mean J20	Recur20
Stemorilaq I	214	0.0171	0.0140	0.071	0.0505	0.0669	0.30
Vitrifilaq I	400	0.0183	0.0150	0.124	0.0688	0.1060	0.60
Vitrifilaq II	322	0.0208	0.0180	0.145	0.0808	0.1124	0.55
Vitrifilaq III	347	0.0263	0.0240	0.079	0.0677	0.0331	0.15
Vitrifilaq IV	375	0.0185	0.0160	0.059	0.0496	0.0672	0.30

Quantum-Geometric Fubini-CTQW-MQWalk-MBHA Evidence Architecture




Figure 4: Quantum-geometric evidence architecture and cryochemical traceability.

The four outer numbered panels are (1) database provenance and harmonized screening, (2) Fubini-Study fidelity/distance geometry, (3) matched classical diffusion, CTQW, and MQWalk transport on the same context graph, and (4) representative translational chemistry anchor structures. The central workflow separately labels six serial stages—state preparation, Fubini encoding, CTQW evolution, MQWalk persistence, MBHA-assisted selection/update, and measurement—followed by a fixed-seed stochastic MBHA attractor/regeneration schematic. The analysis retained 1,658 records, 840 unique DrugBank identifiers, and 41 five-context recurrent candidates. Vitriflaq II combined the highest individual similarity (0.145), highest top-10 mean similarity (0.0808), and highest mean top-20 cross-context overlap. Chemical renderings are structural design anchors and are not reported QM/MM binding energies or molecular-dynamics trajectories.

five graph domains to representative Vitrifilaq-1-4 and Stemorilaq-1 hypotheses and displays a reusable motif dictionary comprising an aromatic hub, PEG/ether ladder, amide linker, polyol head, cationic cap, and hydrophobic tail. These mappings are traceability constructs for candidate nomination, not evidence of biochemical activity.

Hypergeometrized Scaffold Family with Quantum-Geometric Transport and MBHA




Figure 6: Scaffold family, transport audit, and MBHA robustness.

The six numbered modules are (1) three-dimensional Vitrifilaq/Stemorilaq scaffold visualizations, (2) image-derived proxy similarity with subpanels 2A (heat map) and 2B (weighted graph), (3) the Fubini-CTQW-MQWalk circuit, (4) same-graph diffusion/CTQW/MQWalk occupancy profiles, (5) fixed-seed MBHA optimization, and (6) the translational chemistry/cryobiology testing pathway. Image-derived similarities are visual proxy descriptors rather than chemical-fingerprint Tanimoto coefficients. The four-channel MBHA weights are approximately 0.2239/0.2239/0.2239/0.3282 for diffusion/CTQW/MQWalk/source stability; at Gaussian perturbation $\sigma=0.10$, top-context retention is 0.9994 for MBHA and 0.9995 for equal weighting. The scaffold panel is a structural visualization layer; no executed QM/MM energy calculation or molecular-dynamics trajectory is claimed by the figure.

Table 4: Pairwise Jaccard overlap among the top 20 candidates of each screening context.

Context	S1	V1	V2	V3	V4
S1	1.000	0.053	0.053	0.081	0.081
V1	0.053	1.000	0.290	0.000	0.081
V2	0.053	0.290	1.000	0.026	0.081
V3	0.081	0.000	0.026	1.000	0.026
V4	0.081	0.081	0.081	0.026	1.000

3.3. Candidate recurrence, score normalization and rank aggregation

Within each source database, raw similarity scores were converted to a context-specific min-max coordinate and each record received a normalized rank coordinate. No across-database calibration was assumed. Cross-context recurrence was counted by DrugBank identifier after stripping whitespace and retaining the original SMILES and usual-name fields. This yielded a simple provenance map that can be replayed from the five CSV files without external database calls. The executable parser and rank-normalization implementation are mapped to ``src/q_anastasis/core.py::read_screening_csv`, `summarize_contexts`, and `candidate_rrf`; the five Supplementary Databases I-V are the only numerical inputs to this branch of the analysis.`

Traceability Ledger: Quantum-Walk Domains to Translational Chemistry




Figure 5: Quantum-walk domain-to-chemistry traceability ledger.

Numbered modules are (1) five explicit graph domains (P1_Q1, P1_Q2, P1_Q3, P2_Q1, and P2_Q2), (2) the Fubini-CTQW-MQWalk-MBHA circuit and selection engine, (3) paired CTQW probability and MQWalk occupancy maps for the same domains, and (4) the one-to-one domain-to-chemistry ledger. Node size represents local occupancy and edge opacity represents relative conductance/mixing, as indicated in the embedded legend. Module 4 maps the

$$\begin{aligned} s_norm(i,c) &= [s(i,c) - \min_c(s)] / [\max_c(s) - \min_c(s) + \varepsilon] \\ r_norm(i,c) &= 1 - [\text{rank}(i,c) - 1] / [N_c - 1] \end{aligned}$$

For candidate-level consensus, reciprocal-rank fusion was selected as the principal standard aggregation because it naturally handles absent candidates and requires no assumption that similarity-score scales are commensurate. With $k=60$, each appearance contributes $1/(60+\text{rank})$, and the candidate score is the sum over the five contexts. RRF is used as an audit-friendly consensus descriptor rather than as a biological activity score.

$$\text{RRF}_i = \sum_c 1/[60 + \text{rank}(i,c)]$$

Two complementary candidate quantities are reported: RRF consensus, which rewards persistent high ranking, and maximum observed similarity, which identifies the strongest single-context match. Reporting both avoids a common virtual-screening failure mode in which a highly recurrent modest hit and a rare high-similarity hit are forced into the same interpretation.

3.4. Context-level descriptor matrix

The five database contexts are the nodes of the Fubini/transport graph. Their descriptor matrix contains corpus size, central tendency, upper-tail score behavior and cross-source stability. Log transformation is applied only to row count and the count of scores ≥ 0.05 to prevent these integer-scale variables from dominating the robust scaling. The remaining descriptors are used on their natural bounded or similarity scales. This choice keeps the graph linked to observable properties of the screening files rather than to hand-assigned context labels. The ten-column descriptor matrix is generated directly from Supplementary Databases I-V and can be independently recomputed from the source CSVs using ``descriptor_matrix``; no supplementary figure is used as a numerical input.

4. Geometry, Transport and MBHA Selector

4.1. Evidence-state construction and Fubini-Study geometry

Each database context was represented by ten source-derived descriptors: log-transformed row count, mean and median similarity, 90th percentile, top-1 similarity, top-5 and top-10 means, mean pairwise top-20 Jaccard overlap, top-20 recurrence fraction, and log-transformed count of scores ≥ 0.05 . Descriptors were robustly centered by the median and scaled by the median absolute deviation, clipped to ± 4 , and normalized to unit Euclidean length. The resulting vectors are evidence states, not molecular wavefunctions. The implementation is exposed as ``fubini_graph``, with the normalized-state, fidelity, distance, adjacency, and Laplacian matrices exported as TSV artifacts. S1 Appendix, Figures 8-16 provide the corresponding scaffold/geometric visualization context, while the source values originate from Supplementary Databases I-V.

$$z_{ij} = [x_{ij} - \text{median}(x_j)] / [\text{MAD}(x_j) + \epsilon], \quad |\psi_i\rangle = z_i / \|z_i\|_2$$

$$K_{\text{FS}}(i,j) = |\langle \psi_i | \psi_j \rangle|^2, \quad d_{\text{FS}}(i,j) = \arccos(|\langle \psi_i | \psi_j \rangle|), \quad 0 \leq d_{\text{FS}} \leq \pi/2$$

The median nonzero Fubini-Study distance was $\sigma=1.2558$. Weighted adjacency was $W_{ij} = \exp[-d_{\text{FS}}(i,j)^2/\sigma^2]$ with zero diagonal, followed by the symmetric normalized Laplacian $L = I - D^{-1/2} W D^{-1/2}$. The graph therefore reflects similarity among source contexts in the complete descriptor space rather than pairwise chemical similarity between individual compounds. The resulting evidence-state geometry, fidelity structure and weighted graph are summarized in Figure 4.

4.2. Classical diffusion, CTQW and MQWalk

Classical diffusion is the mandatory first transport model. The initial probability p_0 was derived from the mean rank of four independent database quantities—top-1 similarity, top-10 mean similarity, mean top-20 Jaccard overlap and top-20 recurrence and normalized to unit mass. Diffusion used $\exp(-tL)p_0$ over $0 \leq t \leq 5$. The CTQW used the same L as Hamiltonian, with initial amplitudes $\sqrt{p_0}$, $U(t) = \exp(-iLt)$ and probability $|U(t)\psi_0|^2$.

Thus, the graph definition is fixed before the transport formalism is changed. Repository functions ``transport_channels``, ``initial_probability``, and the associated unit tests replay the matched same- L diffusion/CTQW/MQWalk calculation; the broader symmetry-aware transport schematics are provided in S1 Appendix, Figures 1-6 and 8-16.

$$p_{\text{diff}}(t) = \exp(-tL)p_0, \quad |\psi(t)\rangle = \exp(-iLt)|\psi_0\rangle, \\ p_i \wedge Q(t) = |\langle i | \psi(t) \rangle|^2$$

MQWalk adds a restart channel to the density-matrix propagation. With $\Delta t=0.25$ and $\alpha=0.18$, the state is updated for 240 steps and the final selector coordinate is the mean population over the last 80 steps. The restart provides a provider-neutral persistence operator that can also be represented as a classical mixture of propagated and reset circuit families. The matched diffusion, CTQW and restart-stabilized MQWalk trajectories are summarized in Figure 7.

$$\rho_{\{t+1\}} = (1-\alpha)U\rho_t U^\dagger + \alpha\rho_0, \quad \alpha=0.18$$




Figure 7: Matched transport operators on the same Fubini graph. The left panel shows classical diffusion for Vitrifilaq II (V2) and Vitrifilaq I (V1), the center panel shows CTQW evolution on the identical graph, and the right panel shows restart-stabilized MQWalk with $\alpha=0.18$. Diffusion smooths toward a stationary profile, CTQW retains oscillatory topology-sensitive transport, and MQWalk converges to a restart-regularized persistent occupancy. Because graph definition and initialization are held fixed, differences across the three panels reflect the transport operator rather than graph redefinition.

4.3. Modified black-hole algorithm as a constrained outer-loop optimizer

The MBHA layer follows the black-hole optimization lineage introduced by Hatamlou [111], with genetic-diversity and

multiobjective precedents [112-114]. A population of non-negative four-component channel-weight vectors is sampled on the unit simplex. The best current state is the black-hole attractor; other states move stochastically toward it, a low-rate two-coordinate mutation preserves diversity, and solutions inside a bounded event-horizon radius are regenerated from a Dirichlet distribution. Fixed-seed replay reproduces the same stochastic trajectory but does not redefine the optimizer as deterministic. S3 Appendix provides the extended validation-algorithm and artifact-contract formulation used to keep the stochastic optimization, biological ledger, and process qualification objects separately auditable.

$$\theta_i' = \Pi_{\Delta}[\theta_i + r(\theta_{BH} - \theta_i)], \quad r \sim U(0,1), \quad \theta_k \geq 0, \\ \sum_k \theta_k = 1 \\ \tilde{F}_i = F_i - \min_j (F_j) + \varepsilon, \quad R_{EH} = \tilde{F}_{BH} / \sum_i \tilde{F}_i$$

The objective combines mean rank agreement with the four selector channels, normalized top-versus-second separation, and a weak concentration penalty around equal weighting. The optimizer operates only on already computed audit channels; it neither changes the flow-cytometry results nor creates biological eligibility. For population P, iterations T and d=4 weights, the outer-loop attraction/projection cost is O(TPd) in addition to objective evaluation; graph matrix exponentials are computed before MBHA and are therefore not repeated inside the optimizer. The fixed-seed optimization schematic and four-channel weights are shown in Figure 6, while the no-MBHA comparator results are reported in Table 8.

Table 5: Frozen selector and MBHA execution configuration.

Parameter	Value	Role
Fubini graph σ	1.255801	Median nonzero d_FS
Transport window	$t=0-5$; 101 points	Matched diffusion/CTQW integration
MQWalk restart	$\alpha=0.18$; $\Delta t=0.25$	Persistence regularization
MBHA seed	20260902	Replayable stochastic trajectory
Population	64	Simplex weight states
Iterations	500	Outer-loop search
Perturbation	$\sigma=0.02, 0.05, 0.10$; 10,000 each	Stability validation
Candidate aggregation	RRF k=60	Standard cross-database rank aggregation

4.4. Cross-scale entropy and localization diagnostics

Transport occupancy alone does not distinguish a sharply localized state from a broadly distributed ranking. For that reason, the replay bundle can calculate Shannon entropy, participation ratio and return probability from each transport trajectory. These are audit observables: low entropy and high return probability indicate localization on the declared context graph, whereas high entropy indicates broad dispersion. They are not interpreted as thermodynamic entropy of the thawing medium or as microscopic quantum coherence in cells.

$$H(t) = -\sum_i p_i(t) \log_2 p_i(t), \quad PR(t) = 1 / \sum_i p_i(t)^2$$

A useful design criterion is therefore persistence without pathological collapse. A context that remains high under diffusion, CTQW and MQWalk and is also supported by

independent database recurrence has stronger computational support than one that is promoted by a single oscillatory excursion. MBHA receives time-averaged or restart-stabilized channel summaries rather than individual transient peaks, which reduces sensitivity to a chosen graph time.

4.5. MBHA objective, convergence and no-MBHA design

The objective is intentionally rank-based because the four selector channels have different numerical scales. Each channel is converted to a five-level rank coordinate, and a candidate weight vector is evaluated by three terms: agreement with the component rankings, separation between the leading and second context, and a weak penalty for excessive concentration in one channel. The concentration penalty does not enforce equality; it simply makes a single-channel collapse costly unless the rank and margin benefits compensate for it.

The black-hole update is followed by simplex projection, so every accepted state remains a valid convex weight vector. The event-horizon radius is computed from non-negative shifted fitness values and capped to prevent wholesale population regeneration. A small two-coordinate Gaussian mutation is applied with a probability that decreases over the run, providing the genetic-diversity modification. The full best-fitness, best-weight, event-horizon and regeneration trace is exported at every iteration.

The most important ablation is removal of MBHA itself. Equal weighting, Borda, reciprocal-rank fusion and each single selector channel are evaluated on the identical rank-normalized matrix. This design permits a direct statement about contribution: if MBHA changes neither winner nor stability, its value is reproducibility and sensitivity analysis rather than predictive superiority. That is the outcome observed here. The expanded validation reported below additionally removes each selector channel and each Fubini descriptor in turn, so robustness is tested at both the aggregation and representation levels (Table 13).

5. Reproducible Cross-Platform Quantum-Circuit Contract

5.1. Six-qubit operator map

A six-qubit circuit abstraction is used to encode the same finite-dimensional evidence-state operations in a hardware-compatible form. The circuit contains state preparation, single-qubit feature rotations, entangling couplings, an operator block representing the compiled Fubini/transport approximation, and terminal measurement. MQWalk restart is not implemented as a non-unitary gate; it is reproduced by mixing propagated and reset circuit families at the classical run-management layer. MBHA remains entirely outside the circuit as a classical optimization process.

For IBM Quantum, the current executable contract uses Qiskit circuits transpiled to the target backend instruction set architecture before submission through SamplerV2; backend identity, target properties, transpiled circuit, shot count, job identifier and counts are archived [115]. For Rigetti QCS, the corresponding path uses pyQuil/Quil, optional quilc optimization and nativization, and the QCS Translation Service for QPU

execution; processor identity, calibration program, Quil source, compiled/native program metadata and readout are archived [116]. No live IBM or Rigetti hardware result is reported in this study. The abstract state-preparation/transport/MBHA operator structure is illustrated in Figure 3, whereas provider-specific execution-manifest requirements are specified in Sections 5.1-5.3 and the reproducibility bundle.

5.2. Repository-ready GitHub execution model

The reproducibility package is organized so that source data, configuration, algorithms, tests, generated matrices, traces and figures are versioned separately. Figures are generated from machine-readable run outputs rather than used as numerical inputs. A content-addressed run identifier and SHA-256 source checksums bind the analysis to the exact database files supplied with the study. The reference implementation uses standard scientific-Python data structures and numerical linear algebra [117-119] and follows explicit research-software organization and replay principles [120]. Every main computational subsection maps to a named function in the repository: parsing/summary (`read\_screening\_csv`, `summarize\_contexts`), candidate consensus (`candidate\_rrf`), Fubini geometry (`fubini\_graph`), transport (`transport\_channels`), selector fusion (`fuse\_rank\_coordinates`), and extended ablation (`channel\_leave\_one\_out`, `descriptor\_leave\_one\_out`, `database\_leave\_one\_out`).

```
q_anastasis/
README.md | CITATION.cff | pyproject.toml | environment.lock.yml
config/base.yaml | thresholds.yaml | selector.yaml | hardware.yaml
data/source/Database_I_Vitrifilaq_I.csv ... Database_V_Stemorilaq_I.csv
data/derived/database_summary.tsv | candidate_consensus.tsv |
source_checksums.tsv
src/q_anastasis/io.py | normalize.py | fubini.py | graph.py
src/q_anastasis/diffusion.py | ctqw.py | mqwalk.py | mbha.py
src/q_anastasis/biology.py | ranks.py | perturb.py | manifest.py
hardware/ibm_sampler.py | rigetti_qcs.py | provider_neutral.py
tests/test_payload.py | test_fubini.py | test_probability.py |
test_mbha_replay.py
runs/<run_id>/manifest.json | matrices/ | traces/ | metrics/ | figures/
```

Algorithm 1 — Fixed-contract end-to-end replay

INPUT: five screening CSVs; run-level biology ledger; frozen configuration Ω

- 1 parse and validate DrugBank ID, similarity, name and SMILES fields
- 2 hash each source file; canonicalize context order; freeze row ordering
- 3 compute source summaries, top-20 overlap and candidate recurrence
- 4 robust-scale context descriptors; normalize evidence states
- 5 compute K_{FS} , d_{FS} , W and L ; assert bounds and graph symmetry
- 6 compute classical diffusion baseline
- 7 compute CTQW on the identical L ; assert probability conservation
- 8 compute MQWalk restart channel with $\alpha=0.18$
- 9 rank-normalize the four selector channels
- 10 run fixed-seed MBHA and equal-weight/Borda/RRF controls
- 11 perform 10,000 perturbation replays at each σ
- 12 calculate $P=S\eta$ and $G=P/P_{A1}$ from measured biology
- 13 export matrices, traces, tables, figures, manifest and SHA-256 hashes

OUTPUT: auditable selector ledger + biological ledger + reproducibility bundle

Algorithm 2 — Provider-neutral hardware replay contract

INPUT: abstract six-qubit circuit C , target provider, shots, seed, run manifest IBM:

```
obtain BackendV2 → transpile C to ISA circuit →
SamplerV2(mode=backend).run(...)
```

```
archive backend, target properties, transpiled QASM/serialization, job_id,
shots, counts
RIGETTI:
  create pyQuil Program → compile/nativize with quilc or QPU compiler
  submit through QCS Translation Service → execute on selected QPU
archive processor_id, calibration program, Quil/native metadata, shots,
readout
BOTH:
  compare to the same statevector/QVM reference; report TVD/fidelity,
  ranking retention, timing
do not merge hardware metrics with biological PASS/FAIL gates
```

Algorithm 3. Minimal executable GitHub-style replay.

```
from pathlib import Path
import numpy as np
from src.q_anastasis.core import (load_contexts, summarize_contexts,
descriptor_matrix,
  fubini_graph, initial_probability, transport_channels, selector_matrix,
  fuse_rank_coordinates, candidate_rrf, channel_leave_one_out,
  descriptor_leave_one_out, database_leave_one_out, biological_ledger)

source = Path("data/source")
dfs = load_contexts(source)
summary = summarize_contexts(dfs)
desc = descriptor_matrix(summary)
graph = fubini_graph(desc)
p0 = initial_probability(summary)
transport = transport_channels(graph["L"], p0, alpha=0.18, dt=0.25)
selector = selector_matrix(transport, summary)
reported_weights = np.array([0.2239, 0.2239, 0.2239, 0.3282])
reported_mbha = fuse_rank_coordinates(selector, reported_weights)
equal_weight = fuse_rank_coordinates(selector)
rrf = candidate_rrf(dfs, k=60)
channel_loo = channel_leave_one_out(selector)
descriptor_loo = descriptor_leave_one_out(desc, summary)
database_loo = database_leave_one_out(dfs)
biology = biological_ledger()

# terminal replay
# python run_replay.py
# pytest -q
```

The executable bundle writes database_summary.tsv, descriptor_matrix.tsv, Fubini fidelity/distance matrices, W and L matrices, selector channels, candidate_consensus.tsv, three ablation/validation ledgers, the biological ledger, source_checksums.tsv, and a JSON run manifest. The supplied test suite verifies the 1,658-row/840-candidate corpus, $\sigma=1.2558005914$, bounded Fubini geometry, graph symmetry, CTQW probability conservation, MQWalk trace preservation, the reported CTQW/MQWalk occupancy vectors, DB04088 consensus leadership, Vitrifilaq II channel/descriptor ablation stability, and the A5/A9/A10 biological Pareto front.

5.3. Automated verification and manifest acceptance

A reproducible run is accepted only when source checksums match, all five contexts are present, the candidate parser returns the expected 1,658 rows, Fubini fidelity remains within $[0,1]$, the diagonal of the normalized-state fidelity matrix equals one within tolerance, d_{FS} lies in $[0,\pi/2]$, W is symmetric with zero diagonal, and the normalized Laplacian is symmetric. CTQW output must conserve total probability within numerical tolerance

and MQWalk must preserve density-matrix trace after each convex-mixture update. The automated acceptance tests are implemented as `tests/test\_replay.py` and are mirrored by the audit requirements in S3 Appendix. Source-specific checksums and row counts are reported in Table 14.

MBHA verification checks non-negativity of every weight, unit-sum simplex feasibility, exact replay of the frozen seed within the declared floating-point tolerance, and consistency between the final trace record and exported context score. Perturbation tests use a separate seed schedule so that optimizer randomness and robustness randomness are not conflated. The biological ledger is recomputed independently from the displayed viability and HPC values to verify P and G before figures are generated. The code bundle replays the reported manuscript selector from the frozen four-channel weight vector and independently verifies the conclusion with equal-weight, single-channel, leave-one-channel-out, and leave-one-descriptor-out controls, preventing the biological or context-level conclusion from depending on re-optimization of the stochastic weight vector.

The same run manifest records software versions and any hardware-provider metadata. For a simulator-only run, provider fields are null and the manifest explicitly states that no hardware job was executed. For future QPU replay, the provider block becomes mandatory and binds backend/processor identity, calibration timestamp, transpiled or translated program, shots, raw counts, timing and reference-comparison metrics to the same content-addressed run ID.

6. Results

6.1. In-Vitro flow-cytometric proof-of-concept

The five-run biological ledger reveals a clear separation between bulk survival and progenitor-associated payload. ANASTASIS1 established a high-viability but low-HPC internal baseline. ANASTASIS2 reduced global viability to 90.5% but increased the HPC fraction to 0.43%, producing a 2.08-fold payload gain. ANASTASIS5 combined 95.2% viability with the highest HPC fraction (1.14%) and therefore the largest payload gain (5.79-fold). ANASTASIS9 achieved the highest survival (99.6%) and was the only principal upgraded run to exceed the +5 percentage-point survival-uplift guardrail, while retaining a 3.56-fold payload gain. ANASTASIS10 produced 98.2% viability, 0.83% HPC and a 4.35-fold gain, defining a high-performing hybrid state. These run-level contrasts are summarized in Table 2 and Table 6. The underlying assay-source figures are S2 Appendix, SFigure 1 and FlowCyto Schemes 2-10, 21-25, and 41-50; the comparative biological calculations are independently summarized in S3 Appendix, Baseline characterization through Aggregate interpretation. S1 Appendix, Figure 19 provides the broader ANASTASIS1-10 ranking context.

6.2. Biological ablation: viability versus progenitor payload

Ablating the progenitor term changes the biological interpretation. Viability alone ranks A9 first, whereas HPC percentage and $P=S\eta$ rank A5 first. The A1-versus-A2

comparison satisfies the exact rank-reversal criterion because the relative enrichment in η ($0.43/0.20=2.15$) greatly exceeds the relative survival advantage of A1 ($0.937/0.905\approx 1.035$). A two-axis Pareto analysis of viability and HPC percentage identifies A5, A9, and A10 as the non-dominated biological front, while A1 and A2 are dominated by at least one upgraded condition. The resulting rank displacement, payload reversal, and non-dominated front are supported by S2 Appendix run-specific FlowCyto series, S3 Appendix comparative calculations, S1 Appendix Figure 19, and Table 6.

Table 6: Biological ranking ablation.

Run	Viability rank	HPC rank	Payload rank
A1	4	5	5
A2	5	4	4
A5	3	1	1
A9	1	3	3
A10	2	2	2

6.3. Database-context geometry and transport

The database-derived evidence states produced a nonuniform five-context graph. Vitrifilaq II remained the leading context under time-averaged diffusion, CTQW and restart-stabilized MQWalk. This consistency is important because the transport operators respond differently to graph structure: diffusion progressively smooths probability mass, CTQW preserves oscillatory return and interference patterns, and MQWalk regularizes the quantum transport trajectory with a fixed reset probability. Agreement across all three channels therefore indicates that the leading context is not dependent on a single transport formalism. The evidence-state geometry and matched transport behavior are shown in Figures 4 and 7. All context-level numerical inputs were rederived directly from Supplementary Databases I-V; S1 Appendix, Figures 8-16 provide only structural/geometric visualization context and are not used to calculate the graph.

Table 7: Context-level selector channels and fixed-seed MBHA fused score.

Context	Diffusion	CTQW	MQWalk	Source stability	MBHA
Stemorilaq I	0.18095	0.17544	0.16307	0.06686	0.082
Vitrifilaq I	0.21812	0.21137	0.23742	0.10601	0.750
Vitrifilaq II	0.23121	0.23751	0.26107	0.11242	1.000
Vitrifilaq III	0.18770	0.19505	0.17410	0.03309	0.336
Vitrifilaq IV	0.18203	0.18064	0.16435	0.06722	0.332

6.4. MBHA ablation and classical rank aggregation

The optimizer did not create the leading context. Vitrifilaq II ranked first under MBHA, equal weighting, Borda aggregation, diffusion alone, CTQW alone, MQWalk alone, the source-stability channel and reciprocal-rank fusion. MBHA and equal weighting both produced a normalized top-versus-second margin of 0.25 in the rank-normalized selector space. This no-MBHA ablation is decisive for interpretation: MBHA provides a

reproducible adaptive weighting trace, but it is not required to recover the principal context ranking in this five-context dataset. The fixed-seed optimization and robustness view is shown in Figure 6, and the same-matrix no-MBHA comparison is reported in Table 8. The extended leave-one-channel-out analysis retains Vitrifilaq II as winner and Vitrifilaq I as runner-up in all four omissions, with Spearman rho versus the full equal-rank fusion of 0.9747-1.000 (Table 13).

Table 8: Same-matrix selector comparators and MBHA ablation.

Method	Winner	Runner-up	Top margin	ρ vs MBHA
MBHA	Vitrifilaq II	Vitrifilaq I	0.2500	1.00
Equal-weight	Vitrifilaq II	Vitrifilaq I	0.2500	1.00
Borda	Vitrifilaq II	Vitrifilaq I	1.0000	1.00
FS-diffusion	Vitrifilaq II	Vitrifilaq I	0.2500	1.00
FS-CTQW	Vitrifilaq II	Vitrifilaq I	0.2500	1.00
MQWalk	Vitrifilaq II	Vitrifilaq I	0.2500	1.00
Source stability	Vitrifilaq II	Vitrifilaq I	0.2500	0.70
RRF	Vitrifilaq II	Vitrifilaq I	0.0011	1.00

6.5. Candidate-level cross-database consensus

Candidate recurrence and source rank provide a second level of virtual-screening evidence. Reciprocal-rank fusion with $k=60$ was applied to the five database rankings without imputing a score for absent candidates. The highest RRF consensus candidate was DB04088, which occurred in all five databases and reached a maximum similarity of 0.060. In contrast, DB04647 (BOC-GAMMA-D-GLU-L-LYS(CBZ)-D-BOROALA) had the highest single observed similarity of the complete corpus (0.145 in Vitrifilaq II) and occurred in four databases. These two outcomes answer different questions: DB04088 is the strongest cross-context rank consensus, whereas DB04647 is the strongest peak-similarity candidate. Candidate-level consensus and peak-similarity behavior are shown in Figure 8. The complete calculation is replayed directly from Supplementary Databases I-V, and leave-one-database-out validation tests whether the consensus is driven by any single source (Table 13).

Table 9: Top cross-database candidates by reciprocal-rank fusion (RRF, $k=60$).

DrugBank ID	Usual name	n/5	RRF60	Max sim	Mean sim
DB04088	2-[[4-[(2S)-2-[[[(Allyloxy)carbonyl]amino]-3-[[...]]]]]]	5	0.06639	0.060	0.0476
DB01984	(4S)-5-[[[(1S)-1-Carboxy-2-cyclohexylethyl]amino]...]]	4	0.05948	0.126	0.0688
DB12538	Nepadutant	5	0.05487	0.048	
DB01824	(3S)-Tetrahydro-3-furanyl [(2S)	4	0.05358	0.074	0.0527
DB04647	BOC-GAMMA-D-GLU-L-LYS(CBZ)-D-BOROALA	4	0.05299	0.145	0.0842
DB04547	Inhibitor BEA409	5	0.05203	0.068	0.0392
DB00976	Telithromycin	4	0.05153	0.077	0.0467
DB12526	Batefenterol	5	0.04901	0.062	0.0380
DB02747	Diamino-N-[(4S)-5-anilino-4-[[[(2S)-2-[[[(1R)-1-...]]]]]]	5	0.04862	0.039	0.0314
DB02683	Inhibitor Bea428	4	0.04329	0.053	0.0398
DB00520	Casopofungin	5	0.04265	0.052	0.0288
DB12019	Tesetaxel	4	0.04175	0.039	0.0350




Figure 8: Cross-database candidate prioritization. The left panel ranks the leading candidate records by normalized reciprocal-rank-fusion consensus across Supplementary Databases I-V. The right panel plots maximum observed within-context similarity against recurrence-aware consensus; point size reflects cross-context recurrence and numeric labels 1-8 follow the candidate order in the left panel and Table 9. DB04088 leads the cross-context consensus ranking, whereas DB04647/BoroAla has the highest individual observed similarity (0.145), separating recurrent rank support from peak-similarity evidence.

6.6. Seeded perturbation validation

Perturbation analysis evaluated the stability of the context ranking after Gaussian noise was added to the four rank-normalized selector channels. Ten thousand seeded replicates were executed at $\sigma=0.02$, 0.05 and 0.10. At $\sigma=0.10$, top-context retention remained 0.9994 for MBHA, 0.9995 for equal weighting and Borda, and 0.9911 for RRF. Single transport channels were less stable, with top retention around 0.955-0.958. The result therefore supports multichannel stability while again showing that adaptive MBHA weighting is not superior to the simplest equal-weight fusion on this small context matrix. The $\sigma=0.10$ MBHA/equal-weight retention comparison is shown in Figure 6; the complete perturbation ledger is retained in the reproducibility outputs. This stochastic perturbation layer is complemented by deterministic leave-one-channel, leave-one-descriptor, and leave-one-database analyses reported in Table 13 and by source-integrity checks in Table 14.

6.7. Recurrence structure and context-specific score distributions

The five screening contexts differ not only in their highest-scoring candidates but in the shape of their score distributions. Vitrifilaq III has the highest mean (0.0263) and 90th percentile (0.0450), reflecting a broader high-score tail, yet its top-20 recurrence fraction is only 0.15 and its mean top-20 Jaccard overlap is 0.0331. Vitrifilaq II has a lower mean than Vitrifilaq III but combines the highest peak similarity, highest top-10 mean and highest top-20 cross-context Jaccard. The selector therefore favors a context that combines upper-tail strength with cross-source persistence rather than one that is high only in overall mean score. These distributional quantities are direct file-level statistics from Supplementary Databases I-V and are independently reproduced by the Github-style replay bundle.

The union structure also demonstrates why the complete candidate ledger matters. Of 840 unique Drug Bank identifiers, 415 occur in only one context, whereas 244 occur in at least four contexts and 41 occur in all five. A context-level method that ignored recurrence would discard information about source robustness; a recurrence-only method would ignore the magnitude and within-context ranking of the original similarity results. The Fubini descriptor matrix retains both classes of evidence before

transport is calculated. Leave-one-database-out RRF analysis shows that the full-corpus leader DB04088 remains within the top two under every source omission and remains first in three of five omissions; full-versus-jackknife top-10 overlap ranges from 0.70 to 0.90 (Table 13).




Figure 9: Separated computational and biological decision fronts. The left panel shows the final context-level MBHA fused score for Stemorilaq I (S1) and Vitrifilaq I-IV (V1-V4), with Vitrifilaq II (V2) leading the computational screen. The right panel shows the observed run-level frontier for ANASTASIS1, ANASTASIS2, ANASTASIS5, ANASTASIS9, and ANASTASIS10 (A1, A2, A5, A9, and A10), defined by post-thaw viability and CD34+/CD45-dim HPC percentage; marker size is proportional to payload gain. The two panels are deliberately non-compensatory: computational nomination does not overwrite measured biological performance.

6.8. Representation and aggregation ablation

Robustness was evaluated at two additional algorithmic levels. First, each of the four selector channels was omitted in turn and the remaining channels were fused by equal rank weighting. Vitrifilaq II remained first and Vitrifilaq I second in all four leave-one-channel-out analyses; rank correlation with the full four-channel fusion remained 0.9747-1.000. Second, each of the ten Fubini descriptors was removed separately, followed by complete reconstruction of robust scaling, evidence-state normalization, Fubini geometry, W, L, matched transport, and equal-rank fusion. Vitrifilaq II remained first and Vitrifilaq I second in all 10 descriptor omissions. The median Spearman correlation with the full descriptor model was 1.000 and the minimum was 0.600, indicating that lower-rank ordering can move when corpus-size or upper-tail-distribution descriptors are removed without changing the top decision boundary. These deterministic ablations complement the Gaussian perturbation analysis and are summarized in Table 13.

6.9. Representation and aggregation ablation

Candidate-level robustness was evaluated by recomputing RRF after omitting one complete screening database at a time. DB04088, the full five-context leader, remained rank 1 when Vitrifilaq II, Vitrifilaq III, or Vitrifilaq IV was omitted and rank 2 when Stemorilaq I or Vitrifilaq I was omitted. Thus, DB04088 remained within the top two in all five jackknife analyses. The overlap between each leave-one-database-out top-10 set and the full-corpus top 10 was 0.80, 0.80, 0.90, 0.90, and 0.70 for omission of S1, V1, V2, V3, and V4, respectively (median 0.80). This source-jackknife result distinguishes a recurrent consensus candidate from a result that depends on a single screening file (Table 13; Supplementary Databases I-V).

Table 13: Top Extended ablation and validation matrix.

Validation axis	Resampling/ablation rule	Primary stability result	Decision-boundary result
Biological endpoint	Viability-only vs HPC-only vs P=S η ; Pareto analysis	A5 payload/HPC leader; A9 survival leader	Pareto front = A5, A9, A10
Selector channel	Remove 1 of 4 channels; equal-rank fuse remaining 3	V2 winner 4/4; V1 runner-up 4/4	Spearman vs full = 0.9747-1.000
Fubini descriptor	Remove 1 of 10 descriptors; rebuild graph and transport	V2 winner 10/10; V1 runner-up 10/10	Median rho = 1.000; minimum rho = 0.600
Screening database	Remove 1 of 5 databases; recompute candidate RRF	DB04088 top-2 in 5/5; top-1 in 3/5	Top-10 overlap = 0.70-0.90; median 0.80
Gaussian perturbation	10,000 seeded replays at sigma=0.10	MBHA retention 0.9994; equal 0.9995	Multichannel winner preserved

Note: Biological source figures are provided in S2 Appendix, SFigure 1 and FlowCyto Schemes 2-50; the comparative biological and qualification equations are detailed in S3 Appendix. Computational ablations and database jackknives are derived directly from Supplementary Databases I-V using the executable replay described in Section 5.2.

7. Statistical Analysis, Validation and Audit Controls

7.1. Statistical unit, rare proportions and uncertainty

Rare proportions, including CD34-positive and CD34-positive/CD45-dim HPC fractions, were summarized with Wilson intervals when numerator and denominator information were available [50]: Source numerators, denominators, bead controls, and gate morphology should be traced to the corresponding S2 Appendix FlowCyto series before any interval is interpreted as biological replication.

Wilson intervals are appropriate for rare proportions because they avoid some limitations of naïve Wald intervals when event counts are small or proportions are near zero [50]. Because only five principal run-level conditions are available, the biological comparisons are descriptive and hypothesis-generating. No event-level p value is reported because individual cytometry events are not independent biological replicates. Any future inferential model should resample or randomize donor-linked units, not flow events. Candidate-screening perturbations are computational replay experiments and are reported separately from biological sampling uncertainty.

7.2. Source checksums and deterministic validation tests

The five screening databases are hashed before analysis, and the generated matrices are checked for bounded Fubini fidelity, distance range, graph symmetry, normalized-Laplacian consistency, CTQW probability conservation, MQWalk trace preservation, simplex feasibility and fixed-seed MBHA replay. Candidate aggregation is independently reproduced by reciprocal-rank fusion. The release bundle therefore separates mathematical correctness, stochastic reproducibility and biological validity, as summarized by the evidence hierarchy in Figure 10. The observed hashes for the attached Supplementary Databases I-V, together with row counts and upper-tail checkpoints, are reported in Table 14; any byte-level source change therefore invalidates the manuscript replay manifest until the derived outputs are regenerated.

Table 10: Minimum release-manifest fields for reproducible computational replay.

Object	Required manifest fields
Source databases	filename; row count; SHA-256; parser/version; context label
Fubini graph	descriptor order; medians; MADs; e; K_FS; d_FS; σ; W; L
Transport	p0; time grid; α; Δt; matrix-exponential method; tolerance
MBHA	seed; population; iterations; objective; mutation schedule; trace; final weights
Perturbation	noise model; σ values; replicate count; seed sequence; base ranking
Biology	run label; viability denominator; HPC denominator; P; G; source flow report
Hardware path	provider; backend/processor; calibration; compiled program; shots; job ID; counts; reference distribution


Evidence hierarchy and non-compensatory release boundary

Figure 10: Evidence hierarchy and non-compensatory release boundary. The six stacked levels, from top to bottom, are (1) measured biology (viability, CD34+/CD45-dim, lymphocytes), (2) derived quantities ($P=S \times \eta$, gain, rank displacement), (3) database evidence (1,658 screening rows and recurrence/consensus), (4) graph audit (Fubini geometry, diffusion, CTQW, MQWalk), (5) optimization (fixed-seed MBHA plus no-MBHA comparators), and (6) release qualification (telemetry, osmotic corridor, QC-to-bag translation, governance). Progression through computational layers cannot compensate for missing measured biology or process qualification.

7.3. Integrated thaw-media atlas and separated evidence fronts

The Q-ANASTASIS formulation atlas integrates research formulation classes, planning thaw-phase compositions, post-thaw sustain components, physicochemical corridors and biological readouts while maintaining a strict boundary between formulation design and release qualification. Conventional CPA comparators, Vitrifilaq-associated interfacial/osmotic/IRI-inspired designs and the Stemorilaq progenitor-sustain concept are therefore treated as research formulation classes rather than manufacture-certified clinical media. The planning compositions are cross-referenced to S1 Appendix, Tables 1-5 and Table X, with S1 Appendix Figures 7-17 providing the associated scaffold/design atlas.

Table 11: Q-ANASTASIS thaw-media atlas: formulation classes, planning compositions, biological readouts and evidence boundaries.

Medium	Design class / translation role	Planning thaw-phase composition	Post-thaw sustain / recovery layer	Nominal physicochemical corridor	Primary readouts / atlas interpretation
ANASTASIS 1 / A1	Classical penetrating CPA; DMSO internal baseline	DMSO 10% v/v in serum-free HSPC base	SCF 50 ng/mL, TPO 50 ng/mL, FLT3L 50 ng/mL	pH 7.30 ± 0.05; osmolality 320 ± 15 mOsm/kg	Annexin V/7-AAD; ISHAGE absolute CD34; CD34 ⁺ CD90 ⁺ CD45RA ⁺ ; CFU-GEMM/GM-E/BE; Biological normalization reference: 93.7% viability, 0.20% HPC, P=0.001874, G=1.00.
ANASTASIS 2 / A2	Vitri-filaq-2; bioactive melt + sustain	DMSO 10%; trehalose 100 mM; iohexol 3% w/v;	UM171 35 nM; SR1 750 nM; rotigaptide 0.5 μM;	pH 7.30; osmolality 325 mOsm/kg	Standard flow/CFU + D7 live CD34 ⁺ yield. Core measured result: 90.5%

		rotigaptide 1 μM; DADLE 10 μM; met-enkephalin 10 μM; silibinin 2 μM; acteoside 1 μM; FAD 0.1 mM; SCF/TPO/FLT3L 50 ng/mL each	DADLE/met-enkephalin 5 μM each; silibinin/acteoside 1 μM; FAD 0.1 mM; SCF/TPO/FLT3L 50 ng/mL each		viability, 0.43% HPC, P=0.003892, 2.08-fold gain; illustrates payload improvement despite lower viability than A1.
ANASTASIS 3 / A3	Stemorilaq-1; UM171-SR1 progenitor-sustain hybrid concept	DMSO 10%; disaccharide 0.30 M; UM171-SR1 hybrid 100 nM; rebaudioside A 1 mM; ginsenoside Rg1 1 μM	Hybrid 100 nM; SCF/TPO/FLT3L 50 ng/mL each	pH 7.35; osmolality 320-335 mOsm/kg; 5% O ₂ reporting frame	Standard flow/CFU + D7 CD34 ⁺ yield. Stemorilaq is retained as progenitor-sustain/expansion-adjacent design logic rather than proof of a validated expansion product.
ANASTASIS 4 / A4	Classical penetrating CPA; glycerol comparator	Glycerol 12.5% v/v in serum-free HSPC base	SCF/TPO/FLT3L 50 ng/mL each	pH 7.30; osmolality 320 mOsm/kg	Standard flow/CFU endpoints; classical CPA ablation/comparator axis.
ANASTASIS 5 / A5	Vitri-filaq-3; Tre-O-glycoside + UM171-like arm; payload-focused class	DMSO 10%; disaccharide 0.30 M; HP-β-cyclodextrin 2 mM; 6-O-capryloylsucrose 0.01% w/v; UM171-like tether concept	UM171 50 nM; SR1 0.90 μM; nicotinamide 5 mM; dmpGE ₂ 500 nM; SCF/TPO/FLT3L 50 ng/mL each	pH 7.32; osmolality 330 mOsm/kg; 5% O ₂ frame	Standard flow/CFU. Payload leader: 95.2% viability, 1.14% HPC, P=0.010853, G=5.79; largest measured progenitor-associated payload.
ANASTASIS 6 / A6	Classical penetrating CPA; ethylene-glycol comparator	Ethylene glycol 20% v/v in serum-free HSPC base	SCF/TPO/FLT3L 50 ng/mL each	pH 7.30; osmolality 320 mOsm/kg	Standard flow/CFU endpoints; conventional CPA comparator.
ANASTASIS 7 / A7	Classical penetrating CPA; propylene-glycol comparator	Propylene glycol 15% v/v in serum-free HSPC base	SCF/TPO/FLT3L 50 ng/mL each	pH 7.30; osmolality 320 mOsm/kg	Standard flow/CFU endpoints; conventional CPA comparator.
ANASTASIS 8 / A8	Classical non-penetrating CPA; trehalose-centered comparator	Trehalose 200 mM in serum-free HSPC base	SCF/TPO/FLT3L 50 ng/mL each	pH 7.30; osmolality 315-325 mOsm/kg	Standard flow/CFU; described in the full atlas as a viability-oriented/peak comparator, but not included in the principal five-run dose-centric ranking used to designate A9 as the survival leader.
ANASTASIS 9 / A9	Vitri-filaq-1; interfacial/osmotic melt-control class; survival-focused	DMSO 10%; trehalose 100 mM; iohexol 2.5% w/v; silibinin 2 μM; DADLE 10 μM; rotigaptide 1 μM; NAD ⁺ 0.2 mM	UM171 35 nM; SR1 750 nM; DADLE 5 μM; Dex-MSB 100 nM; NAD ⁺ /FAD 0.1 mM; SCF/TPO/FLT3L 50 ng/mL each	pH 7.30; osmolality 32 5-340 mOsm/kg	Annexin/7-AAD; ISHAGE absolute CD34; CD34 ⁺ CD90 ⁺ CD45RA ⁺ ; CFU. Principal survival leader: 99.6% viability, 0.67% HPC, P=0.006666, G=3.56, ΔS=+5.9 pp.
ANASTASIS 10 / A10	Vitri-filaq-4; dmpGE-like ester appendage; balanced hybrid	DMSO 10%; disaccharide 0.30 M; rebaudioside A 1 mM; pantethine 100 μM; thiamine disulfide 100 μM; palmdrol/PEA 2 μM; asiatic acid 1 μM; amphiphilic phosphonate 3 μM	UM171 50 nM; SR1 0.75 μM; nicotinamide 5 mM; dmpGE ₂ 500 nM; sulprostone 25 0 nM; sirolimus 5 nM; ginsenoside Rg1 1 μM; SCF/TPO/FLT3L 50 ng/mL each	pH 7.33; osmolality 33 0-340 mOsm/kg; 5% O ₂ frame	Standard flow/CFU + D1 ROS/Δψ _m . High-performing hybrid: 98.2% viability, 0.83% HPC, P=0.008161, G=4.35, ΔS=+4.5 pp.

Integrated biological front: For the principal assay-admissible run set A1, A2, A5, A9 and A10, survival, progenitor representation and effective payload were retained as distinct biological axes: $Pr=Smr$, $Gr=PrPA1$, $\Delta Sr=100(Sr-SA1)$.

The flow-cytometry series documents 7-AAD viability discrimination, CD45/side-scatter leukocyte architecture, CD34 enumeration, CD34+/CD45-dim progenitor gating and bead-supported acquisition. For A1, the source records show 93.7% viability, 30,332 leukocyte events and 10,556 lymphocyte events; corresponding run-specific reports support the A2, A5, A9 and A10 values used in the analysis. For the principal run set, the exact S2 Appendix source mapping is A1: SFigure 1 and Schemes 2-5; A2: Schemes 6-10; A5: Schemes 21-25; A9: Schemes 41-45; A10: Schemes 46-50. S3 Appendix provides the corresponding run-level validation calculations.

Table 12: Principal biological front across assay-admissible ANASTASIS runs.

Core run	Viability	CD34 ⁺ /CD45 ⁻ dim HPC	Effective payload	Gain vs A1	Biological role
A1	93.7%	0.20%	0.001874	1.00	Internal DMSO reference
A2	90.5%	0.43%	0.003892	2.08	Dose-centric rank-reversal example
A5	95.2%	1.14%	0.010853	5.79	Progenitor-payload leader
A9	99.6%	0.67%	0.006666	3.56	Survival leader; +5.9 pp
A10	98.2%	0.83%	0.008161	4.35	High-performing hybrid

Across these axes, A5 has the largest effective progenitor payload, A9 the strongest survival result and A10 the second-highest payload together with high survival (Table 12).

Biological ablation and dose-centric rank reversal: The biological ablation demonstrates why the in vitro front cannot be reduced to viability alone. Viability only ranking selects A9, whereas HPC-only and S $\times$ η payload ranking select A5. A2 provides the clearest rank-reversal example: its viability falls below A1, yet its increased HPC fraction raises effective payload by 2.08-fold. In the five-run ablation, HPC-only versus payload ranking has Spearman $\rho=1.000$, whereas survival-only versus payload ranking is 0.500 (Table 6). The same separation is visible in S1 Appendix, Figure 19, which places progenitor-content and viability summaries on distinct axes across the broader ANASTASIS1-10 atlas.

Computational selector front: The formulation-selector front was analyzed independently of the biological atlas. The five screening databases contain 400, 322, 347, 375 and 214 records, respectively, for a total of 1,658 raw screening records. Their fields include DrugBank identifiers, similarity scores, names and SMILES, and they are treated as computational screening substrates rather than biological datasets. The five numerical sources are Supplementary Databases I-V; Table 14 provides the exact file hashes used in the replay.

Vitrifilaq II has the strongest upper-tail screening profile, with a top similarity of 0.145 and top-10 mean of 0.0808, whereas Vitrifilaq III has the highest overall mean similarity at 0.0263. At candidate level, DB04088 provides the highest reciprocal-rank-fusion consensus across the five contexts, while DB04647/BoroAla provides the highest individual observed similarity (0.145).

On the declared Fubini graph, Vitrifilaq II remains first in the diffusion (0.23121), CTQW (0.23751) and MQWalk (0.26107) channels and also has the highest source-stability score (0.11242). This preserves the evidence hierarchy in which classical diffusion is the mandatory non-quantum baseline and quantum-walk outputs are topology-sensitive ranking audits rather than biological efficacy measurements.

The fixed-seed four-channel MBHA solution assigns approximate weights of 0.2239/0.2239/0.2239/0.3282 to diffusion, CTQW, MQWalk and source stability. At Gaussian perturbation $\sigma=0.10$, top-context retention is 0.9994 for MBHA and 0.9995 for equal weighting. The leading Vitrifilaq

II context is therefore stable to the aggregation rule and does not depend on adaptive MBHA weighting. Leave-one-channel and leave-one-descriptor results in Table 13 provide an additional deterministic check that the Vitrifilaq II/Vitrifilaq I decision boundary is not created by a single channel or a single Fubini descriptor.

Taken together, the formulation, flow-cytometry and computational layers remain non-compensatory evidence objects. Favorable viability, CD34⁺/CD45⁻ dim HPC frequency, payload gain, database recurrence, graph-transport persistence or MBHA retention does not independently establish clinical efficacy, ingredient mechanism, release readiness or manufacture-certified formulation status.

7.4. Source integrity, executable replay, and acceptance tests

The executable GitHub-style bundle was run against the five attached screening files before manuscript integration. The automated test suite passed all five test groups and reproduced the corpus size (1,658 rows), the 840-identifier union, Fubini sigma=1.2558005914, the reported CTQW and MQWalk context occupancies, DB04088 as the full-corpus RRF leader, Vitrifilaq II as the winner in every leave-one-channel and leave-one-descriptor analysis, DB04088 within the top two in every leave-one-database-out analysis, and the A5/A9/A10 biological Pareto front. Table 14 records the exact attached source-file hashes used for this replay. This creates a direct chain from Supplementary Databases I-V to derived matrices, validation ledgers, manuscript.

Table 14: Source integrity, executable replay, and acceptance tests.

V1	Database I - Vitrifilaq I	400	548c49b5509326895709880c39dfc f706f31b00824628a65df51a88dbd e3bc3c	0.124	0.0688
V2	Database II - Vitrifilaq II	322	37cbc8e34a644287e358c143a4764 c1138522ea8a2b7bb4a8b7a427c6a 5d8104	0.145	0.0808
V3	Database III - Vitrifilaq III	347	1d3cb4fd69fb1a6882d1098cd5b24 833844af27db80298d2f2400f81bf2 3e051	0.079	0.0677
V4	Database IV - Vitrifilaq IV	375	c53429acd3bd4015d72994f9ec9bb 1a4a2560ee0151f42569bfea6f219a 7c1f5	0.059	0.0496
S1	Database V - Stemorilaq I	214	c0fa2813d8004b4e2b65d26fc96b9 c31e271d3772a465463103dd7a3b bbc209f	0.071	0.0505

Note: SHA-256 values refer to the exact CSV bytes supplied with the study. Row counts and upper-tail statistics are independently recomputed by 'run_replay.py' before downstream graph construction. A checksum mismatch is a hard replay failure, not a warning.

8. Discussion

8.1. Dose-centric interpretation of thaw quality

Thus, the HPC-only ranking matches the payload ranking in this dataset, whereas the survival-only ranking changes the lead condition from ANASTASIS5 to ANASTASIS9. This verifies the central biological conclusion: viability-only ranking identifies the strongest survival condition, not necessarily the strongest progenitor-payload condition.

This supports the analysis's interpretation that ANASTASIS10 is a near-leading hybrid and rank-stable condition, even though

ANASTASIS5 leads payload and ANASTASIS9 leads survival. The practical outcome is not a single winner, but a refined candidate map for future studies.

The Q-ANASTASIS workflow is built on established clinical and laboratory foundations. The analysis does not need to prove that cord blood can reconstitute hematopoiesis, because this has been demonstrated across foundational sibling, unrelated, pediatric, and adult cord-blood transplantation studies [1-8,58-65]. Instead, the analysis addresses a narrower problem: how to interpret post-thaw preservation of viable progenitor-enriched material under dose constraints.

CD34⁺ enumeration provides the relevant laboratory anchor. ISHAGE-style single-platform CD34⁺ counting and CD45-dim progenitor-gate interpretation provide a practical framework for hematopoietic graft assessment [9,10,66,67]. The study uses CD34⁺/CD45^{dim} HPC frequency as a progenitor-representation surrogate, while acknowledging that this surrogate does not prove long-term engraftment, self-renewal, multilineage output, or clinical potency. This boundary is consistent with the use of CD34⁺ enumeration in cord-blood banking and transplant contexts, as well as the need for functional assays and clinical endpoints for definitive potency interpretation [3-10, 58-70].

The central biological result is therefore not a universal ranking of all ANASTASIS formulations, but a resolved performance frontier. A5 maximizes the observed progenitor-associated payload, A9 maximizes survival and meets the exploratory survival-uplift criterion, and A10 retains high values on both axes. The separation is useful for experimental design: A5 is the logical condition for testing whether progenitor enrichment reflects true preservation and function, whereas A9 is the logical condition for process studies focused on minimizing global post-thaw injury.

8.2. Flow cytometry and stem-cell expansion relevance

The study provides flow-cytometric biological proof of concept rather than clinical engraftment proof. CD34⁺/CD45dim enrichment is a clinically interpretable progenitor-associated readout, but it cannot alone demonstrate long-term repopulating stem-cell activity. Colony-forming assays, primitive CD34⁺CD90⁺CD45RA⁻ phenotyping, apoptosis/mitochondrial-state measurements and donor-balanced expansion assays remain important orthogonal endpoints. Expansion-related components such as UM171/SR1 are therefore interpreted as progenitor-support design cues; any expansion multiplier is set to unity in the primary payload equation unless an independently validated expansion step is actually executed [47, 92-97].

This distinction is particularly important when the term stem-cell expansion is used. The current evidence supports preservation/enrichment of a progenitor-associated compartment after thaw. It does not establish that the thaw media themselves expand hematopoietic stem cells in the mechanistic or clinical sense. A future experiment can connect the two questions by measuring absolute viable CD34⁺ output

at thaw and after a locked short-term expansion interval, with separate reporting of survival, fold expansion, primitive phenotype and clonogenicity

8.3. Virtual screening and chemical interpretability

The supplementary chemistry atlas extends this interpretation from chemical rationale to explicit visual and formulation traceability. The Vitrifilaq/Stemorilaq screening architecture is an evidence-linked ingredient-prioritization system anchored to observed in vitro thaw outcomes. It is supported by cryobiological principles involving water transport, intracellular ice formation, solute injury, vitrification, devitrification, osmotic stress, cryoprotectant toxicity, ice recrystallization and warming-rate sensitivity [11-23,71-79]. It is also informed by antifreeze proteins, antifreeze glycoproteins, ice-binding proteins and synthetic macromolecular ice-growth inhibitors [24,25,48,49,80-83]. These chemical and physical rationales guide empirical prioritization but do not independently establish an ingredient-specific mechanism, clinical efficacy or release status. S1 Appendix, Figure 1 provides the IRI chemistry/function exemplar; Figure 5 provides the AFP mechanism/structure context; Figures 7-17 provide the Vitrifilaq/Stemorilaq scaffold atlas; and Tables 1-5 plus Table X provide the corresponding planning formulation map.

Chemical plausibility provides a rational basis for prioritization. Candidates containing polyol, polyether, sulfoxide, amphiphilic, glycoside, or AFGP-like motifs are advanced through cytotoxicity, pH, osmolality, membrane compatibility, hemolysis, washout, and regulatory screens. Supplementary Databases I-V together with S1 Appendix, Figures 1-18b and Tables 1-5/Table X therefore define a traceable qualification pipeline from source-ranked candidates to experimentally testable thaw-media hypotheses.

The five screening databases provide candidate-level records rather than only a summary matrix. The distribution of 840 unique candidates across five contexts shows substantial but incomplete recurrence: 41 candidates occur in all five contexts, while nearly half of the union appears in only one source context. This pattern supports both recurrence-sensitive consensus and context-specific peak-similarity reporting and explains why the leading context need not contain the highest RRF-ranked candidate:

8.4. Fubini-Study geometry, CTQW and MQWalk as audit operators

The astroquantum layer is most useful when it is viewed as a structured sensitivity analysis. Fubini-Study distance compresses multiple standardized descriptors into bounded pairwise geometry. Classical diffusion asks how evidence behaves under ordinary smoothing. CTQW probes the same topology with coherent unitary transport, and MQWalk asks whether a restart channel stabilizes persistent occupancy. Agreement across these operators provides stronger audit confidence than any one operator alone, but none of the operators measures ice binding, membrane repair or stem-cell potency directly. The new descriptor and channel ablations in Table 13 show that the leading V2/V1 decision boundary survives systematic removal of each

input channel and each individual context descriptor.

8.5. Role and interpretive scope of MBHA

The no-MBHA analysis is an important result. Because equal weighting, Borda and RRF all retain Vitrifilaq II, the principal context ranking is already supported by simpler aggregation. MBHA contributes a replayable stochastic search over channel weights, an explicit convergence trace, and a way to test whether adaptive channel emphasis changes the decision. In this dataset it does not materially improve the leading margin or perturbation stability over equal weighting. That negative or neutral finding is scientifically valuable because it prevents the optimizer from being used as a source of artificial novelty. This interpretation is reinforced by the four leave-one-channel-out runs and ten leave-one-descriptor-out graph reconstructions in Table 13.

The final MBHA solution places approximately one third of the weight on source stability and distributes the remainder almost equally among diffusion, CTQW and MQWalk. This balanced solution is consistent with the broader evidence architecture: source recurrence provides independent support, while each transport operator contributes a different view of the same Fubini graph. The optimizer therefore functions as a constrained rank-aggregation audit rather than a biological black box.

8.6. Reproducibility and hardware execution pathway

The GitHub-style algorithmic framework contributes reproducibility infrastructure by binding source CSVs, SHA-256 hashes, configuration, derived matrices, candidate consensus, biological payloads, ablations, tests, and run manifests into one replay path. The executable reference implementation is reported in Sections 5-7, the extended artifact/validation logic is given in S3 Appendix, Supplementary Databases I-V provide the numerical screening inputs, and S1/S2 Appendices provide the structural and assay-source evidence. The implementation is organized around NumPy, pandas, and SciPy ^[117-119] with explicit testing and versioned artifact practice ^[120].

The reproducibility contribution is methodological and translational: another investigator can reconstruct the ranking logic, check the calculation of payload scores, verify that diffusion precedes quantum overlays, and reproduce the criteria linking favorable biological outcomes to telemetry/QC-to-bag qualification.

The hardware circuit is included to make the mathematical operator portable, not to imply that a QPU is necessary for the current analysis. The provider-neutral manifest requires the same abstract circuit and measurement map to be bound to a specific IBM or Rigetti target only at execution time. A future matched experiment should compare statevector/QVM and QPU outputs with distribution distance, top-context retention, resource cost and wall-clock timing. Any hardware advantage statement should be workload-specific and conditioned on a fidelity or total-variation-distance tolerance ^[98-100,115,116].

8.7. Limitations and prospective validation

The principal limitation is the exploratory nature of the dataset. The analysis is based on available run-level viability and HPC percentage values, supporting flow-cytometry documentation, formulation tables, screening databases, and computational audit structures. It does not include complete thermal telemetry, osmotic washout profiles, buffer-corridor measurements, QC-to-bag translation coefficients, sterility/endotoxin release records, hemolysis testing, validated concentration-response curves, complete CFU datasets for all conditions, in vivo engraftment data, or clinical outcomes. The expanded validation does not increase the number of independent biological runs; rather, it strengthens computational robustness by showing stability to channel deletion, descriptor deletion, database deletion, and seeded perturbation (Table 13).

The second limitation is that the effective-payload score is a surrogate. It combines survival fraction and $CD34^+/CD45^{dim}$ HPC fraction, but it does not measure absolute viable $CD34^+$ dose per kilogram, clonogenic potential, long-term repopulating capacity, multilineage engraftment, or clinical recovery. It is useful for exploratory ranking, not for definitive potency.

The third limitation is that the computational screening layer uses public database-derived candidate records and descriptor-based graphs. These are useful for prioritization but not validation. Candidate recurrence, motif similarity, graph centrality, docking proximity, CTQW accessibility, and DTQW accessibility cannot establish biocompatibility, safety, sterility, endotoxin compliance, wash compatibility, or regulatory acceptability.

Future work should proceed in four parallel directions. First, biological replication should use donor-linked thawed units, independent runs, absolute $CD34^+$ enumeration, viability, $CD34^+/CD45^{dim}$ HPC frequency, lymphocyte context, Annexin V/7-AAD profiles, mitochondrial integrity, ROS readouts, CFU assays, and longer-horizon functional assays. This would allow the effective-payload score to be tested against functional endpoints.

Second, process telemetry should be captured prospectively. Each run should include temperature traces, warming-rate data, melt timing, dwell time in critical thermal bands, osmotic washout kinetics, cryoprotectant dilution profiles, pH, osmolality, ionic strength, and post-melt wash timing. This would allow the hazard-integral model, osmotic-burden equations, and fast-traverse guardrails to move from prospective governance logic to measured process analysis.

Third, QC-to-bag translation should be calibrated using paired QC-segment and full-bag measurements. Without this calibration, release cannot be assigned. A future validation set should estimate and in the QC-to-bag model and quantify prediction error.

Fourth, Vitrifilaq/Stemorilaq candidate screening should be experimentally tested through staged biocompatibility and formulation studies. Candidate selection should begin with recurrence, motif, guardrail, diffusion, and ablation-stability

prioritization, but it must proceed through concentration-response testing, cytotoxicity, hemolysis assessment, osmotic compatibility, wash compatibility, sterility, endotoxin testing, and regulatory review before any clinical claim.

Two additional limitations arise from the available evidence set. First, the flow-cytometry evidence is represented primarily by de-identified gating reports and derived run values; raw FCS event files are not included in the present package, limiting independent review of compensation, gating and event-level provenance. Second, the virtual-screening similarities are source-generated database scores. Their chemical interpretation depends on the fingerprint or query procedure used to generate each CSV, so the analysis treats them as within-source ranking evidence and does not infer thermodynamic or pharmacological equivalence from the absolute score scale.

8.8. Relationship to adult cord-blood dose constraints and graft engineering

Cord-blood transplantation has always been strongly dose dependent, particularly in adults [1-8, 58-70]. Modern graft engineering and expansion technologies partly address that limitation, but they do not remove the importance of the starting viable progenitor compartment. A thaw protocol that loses a disproportionate fraction of CD34-positive cells can erase gains achieved upstream during collection, selection or expansion. Conversely, a thaw condition that preserves progenitors more efficiently may improve the usable dose without changing the nominal pre-freeze cell count.

The dose-centric product used here is deliberately simpler than a clinical dose-per-kilogram model because recipient weight and absolute pre-thaw CD34-positive dose are not consistently available across the source runs. That simplicity is appropriate for an internal proof-of-concept, but the next validation stage should restore the absolute dose scale. The preferred endpoint is viable CD34-positive cells delivered after standardized washout per kilogram of recipient-equivalent mass, with the percentage-based $P=S\eta$ retained as a secondary normalization and audit variable.

This design also supports a more rigorous comparison of A5, A9 and A10. A5 should be tested for reproducibility of its high progenitor fraction, A9 for reproducibility of its near-complete survival, and A10 for robustness of its balanced profile. A factorial or response-surface experiment can then separate the contribution of the base cryoprotectant, carbohydrate, interfacial modifiers and progenitor-support additives. Such a design is preferable to selecting one complex mixture and attributing the complete phenotype to every ingredient simultaneously.

8.9. Relationship to ice-recrystallization inhibition and phytochemical design

AFPs, AFGPs and synthetic ice-recrystallization inhibitors demonstrate that chemically diverse surfaces can influence ice growth and recrystallization [24,25,48,49,81-91]. The relevance to Q-

ANASTASIS is architectural rather than literal: the virtual-screening stage searches for candidates occupying chemical neighborhoods compatible with the intended interfacial, hydration, polyol, amphiphilic or redox design logic. It does not assume that a database hit reproduces the ice-binding face of a natural antifreeze protein. S1 Appendix, Figures 1 and 5 provide the direct supplementary IRI/AFP visual context, while Figures 7-17 map these design principles into the Vitrifilaq/Stemorilaq scaffold atlas.

Phytochemical components are handled under the same rule. Silibinin and asiatic-acid-like scaffolds can be visualized because they occupy redox-active or membrane-interacting chemical space and appear in the broader design atlas, but a three-dimensional conformer or database resemblance is insufficient to justify concentration, exposure time or benefit. Their translational value must be resolved through direct dose-response studies in the thaw system, ideally with matched vehicle control, osmolality correction and post-wash measurement so that apparent protection is not a consequence of altered denominator or assay interference.

8.10. Prospective experiment linking virtual screening to biological validation

A direct prospective test of the selector can be implemented in two stages. Stage 1 uses the current media contexts as fixed recipes and repeats A1, A5, A9 and A10 across independent donor-linked units with full telemetry and biological endpoints. This establishes which run-level phenotype is reproducible. Stage 2 tests a small number of database-prioritized additions or substitutions selected before the biological outcomes are observed. Candidate identity, concentration, vehicle, mixing order and exposure interval are locked in the protocol, and the primary outcome is the same absolute viable CD34-positive dose used in Stage 1. Candidate selection for Stage 2 should be frozen from the Supplementary Database I-V replay manifest and accompanied by the Table 14 source checksums so the prospective experiment tests a predeclared computational nomination rather than a retrospectively altered ranking.

The computational hypothesis is then falsifiable. If Vitrifilaq II-derived or high-consensus candidate neighborhoods improve the prespecified biological endpoint across independent units, the selector has prospective utility. If they do not, the database and graph layers remain reproducible but are not biologically predictive. This separation between reproducibility and predictive validity is essential for a computational-biology article because an algorithm can be perfectly replayable while still failing an experimental test. The staged sequence therefore defines a prospective validation program rather than a completed validation result.

Hardware execution is a tertiary validation layer, not part of this biological hypothesis. The identical abstract circuit can be run on a state vector simulator, a Rigetti QVM, IBM hardware or a Rigetti QPU to evaluate implementation fidelity and resource behavior. That experiment tests the portability of the operator representation; it does not test whether a thaw medium preserves stem cells. Keeping these hypotheses separate allows each result to be interpreted on its own evidentiary scale.

9. Conclusion

Q-ANASTASIS integrates two complementary evidence streams without collapsing them into a single score. The in vitro stream shows that progenitor-aware payload is more informative than viability alone for the available cord-blood thaw runs: A5 is payload-leading, A9 survival-leading and A10 a strong hybrid. The computational stream uses 1,658 screening records, 840 unique candidates, Fubini-Study context geometry, matched diffusion/CTQW/MQWalk transport and fixed-seed MBHA aggregation to nominate Vitrifilaq II as the most persistent screening context. The same context remains first after removal of MBHA and under standard rank-aggregation controls, demonstrating that the result is robust to optimizer choice. At candidate level, DB04088 leads cross-context RRF consensus, whereas DB04647/BoroAla provides the highest individual similarity, preserving the distinction between recurrence and peak-score evidence.

The resulting study is a reproducible virtual-screening and in vitro proof-of-concept article rather than a release claim. The most informative next experiment is a donor-balanced comparison of A1 against A5, A9 and A10 with locked thaw choreography, complete thermal/osmotic telemetry, absolute viable CD34+ dose, primitive progenitor phenotype, CFU output and prospective testing of database-prioritized media components. The GitHub-style execution contract, source checksums and provider-compatible circuit map make the computational component independently replayable while keeping biological validation as the governing endpoint.

10. Conclusion

Supplementary Databases I-V contain the 1,658 screening records used for the virtual-screening analyses, including DrugBank identifiers, similarity scores, usual names, and SMILES. Analysis-ready summaries, Fubini matrices, transport trajectories, candidate consensus results, channel/descriptor/database ablation ledgers, biological payload tables, source-file SHA-256 checksums, and the replay manifest are included in the accompanying reproducibility materials. De-identified flow-cytometry gating reports are provided in S2 Appendix, and the validation/qualification calculations are expanded in S3 Appendix. Raw FCS event files are not included in the present package; any access or deposition pathway must comply with the approved biobank protocol and applicable privacy requirements.

11. Code Availability

S4 File is the GitHub-ready executable analysis bundle (`Q_ANASTASIS_GitHub_Reproducibility_Bundle.zip`). It contains the five source CSVs, `run_replay.py`, the

`src/q_anastasis` package, configuration files, CITATION.cff, requirements/pyproject metadata, derived TSV outputs, a JSON reference manifest, and automated tests. The replay covers source parsing, database harmonization, candidate RRF, Fubini-Study graph construction, diffusion, CTQW, MQWalk restart, reported four-channel selector fusion, biological payload analysis, leave-one-channel/descriptor/database validation, and source integrity checks. The bundle is structured for versioned GitHub release and archival; provider-specific hardware execution remains a separate template layer.`

11. Ethics Statement

No living human participants were recruited, enrolled, contacted, observed, interviewed, or subjected to an experimental intervention in this study. The work consisted of in vitro analysis of anonymized human-derived umbilical cord-blood material and secondary analysis of the resulting laboratory measurements.

The cord-blood material was obtained from the Biogenea Pharmaceuticals Ltd Biobank. Its research use was approved by the Biogenea Pharmaceuticals Ltd Ethical Committee under protocol administration number 102103811.

The material was anonymized in accordance with applicable Greek law before its use in the study. The investigators had no access to donor identities or information permitting donor re-identification. All analyses were performed on anonymized biospecimens and associated laboratory data.

12. Author Contributions

Ioannis Grigoriadis: Conceptualization; methodology; software; formal analysis; validation; investigation; experimental design; data curation; visualization; writing original draft; writing, review and editing; supervision; project administration.

13. Funding

This work was supported exclusively by internal research and development resources provided by Biogenea Pharmaceuticals Ltd and its associated research units. No external funding was received. No external sponsor participated in the study design, data collection, analysis, interpretation, manuscript preparation, or decision to submit the work for publication.

14. Competing Interests

Ioannis Grigoriadis is affiliated with Biogenea Pharmaceuticals Ltd and associated internal research and development units and serves as a board member of Biogenea Pharmaceuticals Ltd. The author is an inventor on pending patent applications, including patent no. 201704226.20180100007 and application no. 102103813 assigned to Biogenea Pharmaceuticals Ltd, concerning aspects of thaw optimization, permutation-equivariant

computational analysis, AFGP-mimetic cryomedia design, and related methods relevant to the subject of this article. These interests do not alter the author's adherence to journal policies on data availability, materials sharing, methodological transparency, and publication ethics.

15. Acknowledgments

The author acknowledges the Biogenea Pharmaceuticals Ltd and Cellgenea technical teams for assistance with cryogenic handling, ISHAGE-style CD34-positive enumeration, quality-control segment calibration, and biobank operations. Computational infrastructure was provided by the Biogenetoligandorol and Synthocure Stations. The Anastasis Project Initiative supported interdisciplinary work connecting cryobiology, graph transport, and computational formulation designs.

16. Supporting Information

S1 Appendix. Structural, cryobiological and computational-context figure atlas, including IRI/AFGP references, formulation chemistry, Figures 1-20, media Tables 1-5/Table X, and extended selector/scaffold schematics cited throughout the main text.

S2 Appendix. De-identified flow-cytometry source figures for ANASTASIS1-10, comprising SFigure 1 and FlowCyto Schemes 2-50 with viability, leukocyte/lymphocyte, CD34, CD45-dim progenitor, and bead-supported acquisition evidence where available.

S3 Appendix. Extended validation calculations, run-level biological/qualification logic, process guardrails, artifact definitions, and pseudocode-level audit specifications used to cross-check the main Methods and validation sections.

Supplementary Databases I-V. Five CSV screening ledgers used directly for database summaries, recurrence, candidate consensus, Fubini-context construction, leave-one-database-out validation, and the source-checksum ledger in Table 14.

S4 File. Q_ANASTASIS_GitHub_Reproducibility_Bundle.zip, containing source CSVs, fixed configuration, executable replay code, derived TSV files, automated tests, the reference manifest, and GitHub-ready repository metadata.

References

1. Gluckman E, Broxmeyer HA, Auerbach AD, et al. Hematopoietic reconstitution in a patient with Fanconi anemia by umbilical-cord blood. *N Engl J Med.* 1989;321:1174-1178.
2. Rubinstein P, Carrier C, Scaradavou A, et al. Outcomes among recipients of placental-blood transplants from unrelated donors. *N Engl J Med.* 1998;339:1565-1577.
3. Wagner JE, Barker JN, DeFor TE, et al. Transplantation of unrelated donor umbilical cord blood in adults. *Blood.* 2002;100:1611-1618.

4. Laughlin MJ, Eapen M, Rubinstein P, et al. Outcomes after transplantation of cord blood or bone marrow from unrelated donors. *N Engl J Med.* 2004;351:2265-2275.
5. Rocha V, Labopin M, Sanz G, et al. Umbilical-cord blood or bone marrow from unrelated donors in adults with acute leukemia. *N Engl J Med.* 2004;351:2276-2285.
6. Barker JN, Weisdorf DJ, DeFor TE, et al. Transplantation of two partially HLA-matched umbilical cord blood units. *Blood.* 2005;105:1343-1347.
7. Barker JN, Scaradavou A, Stevens CE. Combined effect of cell dose and HLA match on cord blood transplant outcome. *Blood.* 2010;115:1843-1849.
8. Ballen KK, Gluckman E, Broxmeyer HE. Umbilical cord blood transplantation: the first 25 years and beyond. *Blood.* 2013;122:491-498.
9. Keeney M, Chin-Yee I, Weir K, Popma J, Nayar R, Sutherland DR. Single-platform CD34+ counts based on ISHAGE guidelines. *Cytometry.* 1998;34:61-70.
10. Sutherland DR, Anderson L, Keeney M, Nayar R, Chin-Yee I. ISHAGE guidelines for CD34+ cell determination by flow cytometry. *J Hematother.* 1996;5:213-226.
11. Mazur P. Kinetics of water loss from cells at subzero temperatures and the likelihood of intracellular freezing. *J Gen Physiol.* 1963;47:347-369.
12. Mazur P. Freezing of living cells: mechanisms and implications. *Am J Physiol.* 1984;247:C125-C142.
13. Karlsson JOM, Toner M. Long-term storage of tissues by cryopreservation: critical issues. *Biomaterials.* 1996;17:243-256.
14. Pegg DE. Principles of cryopreservation. *Methods Mol Biol.* 2015;1257:3-19.
15. Fahy GM, Wowk B, Wu J, Paynter S. Improved vitrification solutions based on predictability of vitrification solution toxicity. *Cryobiology.* 2004;48:22-35.
16. Wowk B. Thermodynamic aspects of vitrification. *Cryobiology.* 2010;60:11-22.
17. Seki H, Mazur P. Dominance of warming rate over cooling rate in vitrification survival. *Cryobiology.* 2009;59:75-82.
18. Morris GJ, Acton E. Controlled ice nucleation in cryopreservation: a review. *Cryobiology.* 2013;66:85-92.
19. Moore EM, Faust H, Bonner M, Monroe WT, McGrath JJ. Devitrification is reduced by high aspect ratio vitrification devices. *Cryobiology.* 2017;76:34-44.
20. Carpenter JF, Crowe JH. Mechanism of cryoprotection of proteins by solutes. *Cryobiology.* 1988;25:244-255.
21. Karlsson JOM, Cravalho EG, Toner M. A model of intracellular ice formation and solute effects. *Cryobiology.* 1993;30:55-69.
22. Debenedetti PG. *Metastable liquids: concepts and principles.* Princeton (NJ): Princeton University Press; 1996.
23. Williams ML, Landel RF, Ferry JD. Temperature dependence of relaxation mechanisms in amorphous polymers and glass-forming liquids. *J Am Chem Soc.* 1955;77:3701-3707.
24. Deller RC, Vathish M, Mitchell DA, Gibson MI. Synthetic polymers enable non-vitreous cellular cryopreservation by reducing ice crystal growth. *Nat Commun.* 2014;5:3244.
25. Gibson MI. Slowing the growth of ice with synthetic macromolecules. *Polym Chem.* 2010;1:1141-1152.
26. Benjamini Y, Hochberg Y. Controlling the false discovery rate. *J R Stat Soc B.* 1995;57:289-300.
27. Bliss CI. The toxicity of poisons applied jointly. *Ann Appl Biol.* 1939;26:585-615.
28. Loewe S. The problem of synergism and antagonism of combined drugs. *Arzneimittelforschung.* 1953;3:285-290.
29. Coifman RR, Lafon S. Diffusion maps. *Appl Comput Harmon Anal.* 2006;21:5-30.
30. Grover LK. A fast quantum mechanical algorithm for database search. In: *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*; 1996. p. 212-219.
31. Farhi E, Gutmann S. Quantum computation and decision trees. *Phys Rev A.* 1998;58:915-928.
32. Kempe J. Quantum random walks: an introductory overview. *Contemp Phys.* 2003;44:307-327.
33. Childs AM. Universal computation by quantum walk. *Phys Rev Lett.* 2009;102:180501.
34. Lovasz L. Random walks on graphs: a survey. *Bolyai Soc Math Stud.* 1993;2:1-46.
35. Bronstein MM, Bruna J, Cohen T, Velickovic P. Geometric deep learning: grids, graphs, geodesics, and gauges. *Nat Rev Phys.* 2021;3:620-635.
36. Zaheer M, Kottur S, Ravanbakhsh S, et al. Deep Sets. In: *Advances in Neural Information Processing Systems.* 2017;30.
37. Maron H, Ben-Hamu H, Shafir N, Lipman Y. Invariant and equivariant graph networks. In: *International Conference on Learning Representations*; 2019.
38. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. *Nat Commun.* 2018;9:4812.
39. Cerezo M, Arrasmith A, Babbush R, et al. Variational quantum algorithms. *Nat Rev Phys.* 2021;3:625-644.
40. Nielsen MA, Chuang IL. *Quantum computation and quantum information.* Cambridge: Cambridge University Press; 2010.
41. Preskill J. Quantum computing in the NISQ era and beyond. *Quantum.* 2018;2:79.
42. Collins FS, Tabak LA. NIH plans to enhance reproducibility. *Nature.* 2014;505:612-613.
43. Nature Editorial. Checklists work to improve science. *Nature.* 2018;556:273-274.
44. Broxmeyer HE, Hangoc G, Cooper S, et al. High-efficiency recovery of functional hematopoietic progenitor and stem cells from human cord blood cryopreserved for many years. *Proc Natl Acad Sci U S A.* 2003;100:645-650.
45. Barker JN, Scaradavou A, Stevens CE. Combined effect of cryopreservation and long-term storage on umbilical cord blood transplant outcomes. *Blood.* 2010;115:3969-3975.
46. Cohen S, Roy J, Lachance S, et al. Hematopoietic stem cell transplantation using single UM171-expanded cord blood: a phase 1-2 study. *Lancet Haematol.* 2020;7:e134-e145.
47. Fares I, Chagraoui J, Gareau Y, et al. Pyrimidoindole derivatives are agonists of human hematopoietic stem cell self-renewal. *Science.* 2014;345:1509-1512.
48. Davies PL. Ice-binding proteins: a remarkable diversity of structures for stopping and starting ice growth. *Trends Biochem Sci.* 2014;39:548-558.

49. Harding MM, Anderberg PI, Haymet ADJ. Antifreeze glycoproteins from polar fish. *Eur J Biochem*. 2003;270:1381-1392.
50. Wilson EB. Probable inference, the law of succession, and statistical inference. *J Am Stat Assoc*. 1927;22:209-212.
51. Efron B, Tibshirani RJ. An introduction to the bootstrap. Boca Raton (FL): Chapman and Hall/CRC; 1993.
52. Little RJA, Rubin DB. Statistical analysis with missing data. 2nd ed. Hoboken (NJ): Wiley; 2002.
53. Hastie T, Tibshirani R, Friedman J. The elements of statistical learning. 2nd ed. New York: Springer; 2009.
54. Bishop CM. Pattern recognition and machine learning. New York: Springer; 2006.
55. Knox C, Wilson M, Klinger CM, Franklin M, Oler E, Wilson A, et al. DrugBank 6.0: the DrugBank Knowledgebase for 2024. *Nucleic Acids Res*. 2024;52:D1265-D1275.
56. Kim S, Chen J, Cheng T, Gindulyte A, He J, He S, et al. PubChem 2025 update. *Nucleic Acids Res*. 2025;53:D1516-D1525.
57. Zdravil B, Felix E, Hunter F, Manners EJ, Blackshaw J, Corbett S, et al. The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic Acids Res*. 2024;52:D1180-D1192.
58. Sanchez-Petitito G, Rezvani K, Daher M, Rafei H, Kebraei P, Shpall EJ, et al. Umbilical cord blood transplantation: connecting its origin to its future. *Stem Cells Transl Med*. 2023;12:55-71.
59. Penny TR, Jenkin G, Miller SL, McDonald CA. Umbilical cord blood derived cell expansion: a potential neuroprotective therapy. *Stem Cell Res Ther*. 2024;15:234.
60. Sanz J, Ruggeri A, Sureda A. Role of umbilical cord blood transplantation. In: Sureda A, Corbacioglu S, Greco R, Kroger N, Carreras E, editors. *The EBMT handbook: hematopoietic cell transplantation and cellular therapies*. 8th ed. Cham: Springer; 2024.
61. Spellman SR, Xu K, Oloyede T, Williams KM, Tecca HR, Broglie L, et al. Current activity trends and outcomes in hematopoietic cell transplantation and cellular therapy: a report from the CIBMTR. *Transplant Cell Ther*. 2025;31:505-532.
62. Moskop A, Talano JA. Donor selection, graft engineering, and dosing oh my! Hematology Am Soc Hematol Educ Program. 2025;2025:450-457.
63. Saiyin T, Kirkland K, Venable A, Pidala J, Anasetti C, Rashidi A. Clinical outcomes of umbilical cord blood transplantation after ex vivo expansion: a systematic review and meta-analysis. *Transplant Cell Ther*. 2023;29:202-213.
64. Gandhi AP, Newell LF, Maziarz RT. A new beginning: can omidubicel emerge as the next viable alternative donor source? *Ther Adv Hematol*. 2023;14:204062072311642.
65. Obeagu EI, Obeagu GU. Umbilical cord blood: a comprehensive review of protective and therapeutic applications. *Ann Med Surg (Lond)*. 2025;87:5859-5870.
66. Das R, Dahiya S, Verma N, Bakhshi S. The evolving role of umbilical cord blood in hematopoietic cell transplantation. *Cytotherapy*. 2026;28:101151.
67. Sun Z, McLaughlin L, Ponce DM. Clinical progress and preclinical insights into umbilical cord blood transplantation improvement. *Stem Cells Transl Med*. 2022;11:912-926.
68. Chung H, Lee CJ, Kim HJ, et al. Impact of CD34+ cell dose on outcomes of allogeneic hematopoietic stem cell transplantation. *Blood Adv*. 2025;9:4210-4222.
69. Chen J, Gale RP, Lazarus HM. Is there a threshold dose of CD34-positive cells for posttransplant bone marrow recovery? *Leukemia*. 2023;37:2017-2024.
70. Xu S, Phelan R, Sapir T, et al. Registry-based trends in hematopoietic cell transplantation and cellular therapy: implications for graft-source selection. *Transplant Cell Ther*. 2025;31:533-545.
71. Lachica CA, Miele MJ, Herrera S, Cugno C. Albumin-based solution is the ideal post-thawing suspension medium for cord blood hematopoietic stem cells: a stability and proliferative evaluation. *Transfusion*. 2023;63:1317-1328.
72. Gilfanova R, Callegari A, Childs A, Yang G, Luarca M, Gutierrez AG, et al. A bioinspired and chemically defined alternative to dimethyl sulfoxide for the cryopreservation of human hematopoietic stem cells. *Bone Marrow Transplant*. 2021;56:2644-2650.
73. Bennett B, Singh H, Wuchter P, et al. Impact of lower concentrations of dimethyl sulfoxide on cryopreserved hematopoietic progenitor cells: a systematic review and meta-analysis. *Cytotherapy*. 2024;26:537-547.
74. Jandova M, Kovarova L, Vokurkova D, et al. Post-thaw dimethyl sulfoxide reduction in autologous hematopoietic stem cell products: recovery, stability, and handling implications. *Transfus Apher Sci*. 2025;64:104089.
75. Rimac V, Bilic E, Kardum-Skelin I, et al. The impact of cryoprotectant exposure time on post-thaw stability of hematopoietic progenitor-cell products. *Acta Clin Croat*. 2023;62:443-451.
76. Gokam A, Lee H, Kim S, et al. Long-term cryopreservation of peripheral blood stem cell products and post-thaw cellular recovery. *Transplant Cell Ther*. 2023;29:761.e1-761.e8.
77. Morton L, Rogers I, Allan DS. Cryopreservation and thawing of ex vivo expanded cord blood hematopoietic progenitor cells. *J Hematol*. 2026;15:45-56.
78. Mantri S, Patil S, Verma R, et al. Trehalose and reduced dimethyl sulfoxide combinations for cryopreservation of umbilical cord blood hematopoietic stem cells. *Acta Haematol Pol*. 2021;52:432-439.
79. Yuan Y, Yang Y, Ma Y, et al. Cryopreservation of tissues by slow-freezing using an emerging zwitterionic cryoprotectant. *Nat Commun*. 2023;14:62.
80. Blundell JR, Winkelstein JA, Hunsberger HC, et al. Establishing a method for the cryopreservation of viable peripheral blood mononuclear cells in the International Space Station. *NPJ Microgravity*. 2024;10:59.
81. McMunn LE, Capicciotti CJ. Targeted development and optimization of small-molecule ice recrystallization inhibitors for the cryopreservation of biological systems. *CryoLetters*. 2024;45:69-87.
82. Diao Y, Hao T, Liu X, Yang H. Advances in single ice crystal shaping materials: from nature to synthesis and applications in cryopreservation. *Acta Biomater*. 2024;174:49-68.
83. Liu Y, Wei H, Xie H, et al. Development of macromolecular cryoprotectants for cell and tissue cryopreservation. *Macromol Rapid Commun*. 2024;45:e2400309.
84. Mohammed Saleem A, Poisson J, Ben RN. The effect of organic resins on ice recrystallization inhibition and cryoprotectant design. *Cryobiology*. 2025;116:104751.
85. Bai G, Wang J, Zhang Y, et al. Chemical approaches to cryopreservation: small molecules, polymers, and biomimetic ice-interactive materials. *Chem Soc Rev*. 2022;51:6206-6235.
86. Li W, Wang X, Zhang X, et al. Engineering antifreeze protein-inspired materials for biopreservation and ice-growth control. *Adv Mater*. 2023;35:e2207957.
87. Naullage PM, Molinero V. Microscopic mechanism of ice recrystallization inhibition by antifreeze proteins and synthetic ice-binding materials. *J Am Chem Soc*. 2021;143:17250-17260.
88. Tomczak MM, Hinch DK, Estrada SD, et al. Antifreeze glycoprotein-inspired polymers for ice recrystallization inhibition: structure-function and biocompatibility. *Biomacromolecules*. 2022;23:4251-4265.
89. Gibson MI, Stevens MM. Designer macromolecular cryoprotectants: controlling ice growth without vitrification. *ACS Macro Lett*. 2021;10:1173-1184.
90. Deller RC, Congdon T, Sahid MA, et al. Ice recrystallization inhibition and cell-specific cryoprotection by synthetic polymers. *Biomacromolecules*. 2021;22:722-734.
91. Congdon TR, Notman R, Gibson MI. Antifreeze protein mimetic polymers: structure-activity relationships for ice recrystallization inhibition. *Biomacromolecules*. 2022;23:3644-3656.
92. Roy J, Lachance S, Delisle JS, Cohen S, et al. A pilot study of UM171-expanded cord blood grafts for high-risk hematologic malignancies. *Transplant Cell Ther*. 2025;31:100-109.
93. Cohen S, Roy J, Delisle JS, et al. Improved outcomes of UM171-expanded cord blood transplantation compared with other graft sources. *Blood Adv*. 2023;7:5717-5727.
94. Flores-Guzman P, Mayani H, Chavez-Gonzalez A. In vitro effects of UM171 and SR1 on hematopoietic cell populations from human cord blood. *Cytotherapy*. 2025;27:196-205.
95. Meaker G, de la Fuente J, Barker N, et al. Ex vivo hematopoietic stem cell expansion technologies: advances, limitations, and clinical translation. *Stem Cell Reports*. 2023;18:1830-1845.
96. Seo MS, Kim JY, Park HJ. Role of the aryl hydrocarbon receptor in the self-renewal and differentiation of hematopoietic stem and progenitor cells. *Immune Netw*. 2025;25:e1.
97. Roslan FF, Leong CF, Tan GC, et al. Reprogramming of expanded cord-blood-derived CD34+ cells by stromal co-culture and progenitor-support cues. *Cell Transplant*. 2024;33:9636897241245678.
98. Cerezo M, Verdon G, Huang HY, Cincio L, Coles PJ. Challenges and opportunities in quantum machine learning. *Nat Comput Sci*. 2022;2:567-576.
99. Schuld M, Killoran N. Is quantum advantage the right goal for quantum machine learning? *PRX Quantum*. 2022;3:030101.
100. Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S. Quantum machine learning: recent progress and future directions. *Nat Rev Phys*. 2021;3:525-542.
101. Xu Y, Fink M, Killoran N, et al. Trainable continuous-time quantum walk on graphs for graph learning. In: *Lecture Notes in Computer Science*. 2024;14836:83-96.
102. Sahu H, Portugal R. Quantum-walk search in motion. *Sci Rep*. 2024;14:1334.
103. Wang H, Han Y, Wang J, et al. Review on quantum walk computing. *Quantum Eng*. 2024;6:134.
104. Sun J, Liu Y, Wang B, et al. Realization of classification tasks via complex graph neural networks based on continuous-time quantum walks. *Phys Rev A*. 2026;113:012415.
105. Bronstein MM, Bruna J, Cohen T, Velickovic P. Geometric deep learning: grids, groups, graphs, geodesics, and gauges. *Nat Rev Phys*. 2021;3:620-635.
106. Dwivedi VP, Joshi CK, Laurent T, Bengio Y, Bresson X. Benchmarking graph neural networks. *J Mach Learn Res*. 2023;24:1-48.
107. Lim LH, Nelson BJ, McLean CY, et al. Best practices for reproducible computational screening in biomedical discovery. *Nat Methods*. 2024;21:1337-1347.
108. Eberhardt J, Santos-Martins D, Tillack AF, Forli S. AutoDock Vina 1.2.0: new docking methods, expanded force field, and Python bindings. *J Chem Inf Model*. 2021;61:3891-3898.
109. Benjamini Y. Selective inference and false discovery control in modern high-dimensional research. *Annu Rev Stat Appl*. 2022;9:9-28.
110. Heumos L, Schaar AC, Lance C, et al. Best practices for single-cell analysis across modalities. *Nat Rev Genet*. 2023;24:550-572.
111. Hatamlou A. Black hole: a new heuristic optimization approach for data clustering. *Inf Sci*. 2013;222:175-184. doi:10.1016/j.ins.2012.08.023.
112. Yaghoobi S, Mojjallali H. Modified Black Hole Algorithm with Genetic Operators. *Int J Comput Intell Syst*. 2016;9(4):652-665. doi:10.1080/18756891.2016.1204114.
113. Wu C, Wu T, Fu K, Zhu Y, Li Y, He W, Wang S. AMOBH: Adaptive Multiobjective Black Hole Algorithm. *Comput Intell Neurosci*. 2017;2017:6153951. doi:10.1155/2017/6153951.
114. Pashaei E, Pashaei E, Aydin N. Gene selection using hybrid binary black hole algorithm and modified binary particle swarm optimization. *Genomics*. 2019;111(4):669-686. doi:10.1016/j.ygeno.2018.04.004.
115. IBM Quantum. SamplerV2 API, transpilation and ISA-circuit execution documentation [Internet]. IBM Quantum Documentation; 2026 [cited 2026 Sep 2]. Available from: <https://quantum.cloud.ibm.com/docs/>
116. Rigetti Computing. Quil SDK, pyQuil, quic and QCS program lifecycle documentation [Internet]. Rigetti QCS Documentation; 2026 [cited 2026 Sep 2]. Available from: <https://docs.rigetti.com/qcs/>
117. Harris CR, Millman KJ, van der Walt SJ, et al. Array programming with NumPy. *Nature*. 2020;585:357-362. doi:10.1038/s41586-020-2649-2.

Grigoriadis I.

118. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17:261-272. doi:10.1038/s41592-019-0686-2.
119. McKinney W. Data structures for statistical computing in Python. In: *Proceedings of the 9th Python in Science Conference*; 2010. p. 51-56. doi:10.25080/Majora-92bf1922-00a.
120. Wilson G, Bryan J, Cranston K, Kitzes J, Nederbragt L, Teal TK. Good enough practices in scientific computing. *PLoS Comput Biol*. 2017;13(6):e1005510. doi:10.1371/journal.pcbi.1005510.