informatics

MDPI

Article

MSC Digital Assetization for Personalized Regenerative Medicine: An AI–Blockchain–Digital Twin Integrated Framework

Chung Seok Han [1,*], Jin Woo Yang [2], Sun Koo Park [2] and Min Jae Park [3,*]

[1] Department of Business Administration, Seoul School of Integrated Sciences and Technologies (aSSIST), Seoul 03767, Republic of Korea
[2] Department of Stem Cell and Regenerative Biotechnology, School of Advanced Biotechnology, Molecular & Cellular Reprogramming Center, Institute of Advanced Regenerative Science, Institute of Health, Aging & Society, Konkuk University, 120 Neungdong-ro Gwangjin-gu, Seoul 05029, Republic of Korea
[3] Department of e-Business, School of Business, Ajou University, Suwon 16499, Republic of Korea
* Correspondence: iamcalvinhan@gmail.com (C.S.H.); geglove@ajou.ac.kr (M.J.P.); Tel.: +82-10-2619-0003 (C.S.H.)

Abstract

Mesenchymal stem cells (MSCs) are a critical biological resource for regenerative medicine, immunomodulation, and personalized cell therapy. Three structural problems persist: (1) the absence of standardized, quantitative quality indicators; (2) insufficient tamper-proof traceability throughout the manufacturing and banking lifecycle; and (3) the lack of a personalized matching system linking MSC batch characteristics to patient-specific clinical requirements. This paper proposes the MSC Digital Assetization Framework (MDAF), an applied engineering framework that addresses all three problems at the architectural and prototype level. Here, digital assetization—the transformation of a biological product into a structured, traceable, and transferable digital quality record within a multi-institutional trust infrastructure—denotes verifiable, traceable, quality-certified digital recordization of MSC batches, not tokenization or financial trading. The quality engine integrates morphological, FLIM-derived metabolic–proliferative, donor blood panel, flow cytometry, and manufacturing metadata inputs through a bidirectional Cross-Attention fusion module, yielding a continuous MSC quality score (MQS, 0–100) and an S/A/B/C/D five-tier grade. Privacy-preserving verification is implemented via two independent Groth16 zero-knowledge proof circuits: a Release Eligibility Proof (REP, MQS $\geq$ 70) and a Premium Quality Proof (PQP, MQS $\geq$ 85). A Hyperledger Besu QBFT permissioned blockchain with smart contracts provides immutable lifecycle traceability and DID-based access control. In a synthetic data pilot (n = 2000), the system demonstrated engineering feasibility across all five subsystems. These results are engineering pipeline feasibility benchmarks on synthetic data; biological and clinical validation using real MSC data is mandatory follow-on research.

Keywords: mesenchymal stem cell; MSC quality score; multimodal AI; cross-attention; FLIM; blockchain; digital certificate; zero-knowledge proof; QBFT; digital twin

Check for updates

Academic Editor: Antony Bryant

Received: 29 May 2026
Revised: 5 August 2026
Accepted: 10 August 2026
Published: 14 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Mesenchymal stem cells (MSCs), owing to their immunomodulatory and tissue-regenerative properties, are the subject of active clinical investigation worldwide as a raw material for cell therapies targeting intractable diseases such as graft-versus-host

disease (GvHD), rheumatoid arthritis, and myocardial infarction [1,2]. The global MSC therapeutics market was valued at approximately USD 2.2 billion in 2023 and is projected to grow at a CAGR of 10.4% through to 2032 [3].

However, MSC-based programs face three structural problems. First, no standardized, quantitative quality indicator integrates morphological health, metabolic fitness, donor biology, and manufacturing history. The ISCT minimum criteria [4] define biological identity but do not predict therapeutic efficacy [5]. Second, existing biobanking infrastructure remains storage-centric, lacking tamper-proof lifecycle traceability across multi-site manufacturing, cryopreservation, transport, and pre-release steps [6,7]. Third, personalized medicine demands MSC batch-specific matching that reflects each patient's inflammatory burden, metabolic state, and tissue-damage profile [8,9].

To clarify the specific technical setting, the application problem addressed is the selection and certification of the most suitable MSC batch for a specific patient in a clinical cell therapy workflow. A clinical laboratory produces multiple MSC batches from different donors under varying manufacturing conditions; the challenge is to (a) quantify batch quality in a reproducible, standardized manner, (b) issue a tamper-proof, auditable quality certification record, and (c) identify which batch best matches a given patient's clinical profile. These three sub-problems currently lack an integrated, interoperable solution.

MSC "quality" in this context refers to the composite of biological, metabolic, and manufacturing characteristics that predict therapeutic efficacy: (i) morphological integrity (cell area, aspect ratio, cytoplasm-to-nucleus ratio) assessed from phase-contrast images; (ii) metabolic fitness measured by FLIM-derived mean photon arrival time (τ_mean) and NAD^+-related metabolic proxy, reflecting mitochondrial oxidative phosphorylation activity; (iii) proliferative potential measured by passage-wise doubling time; (iv) donor health status from a 15-item blood panel (hsCRP, HbA1c, BMI, etc.); and (v) flow cytometry surface marker compliance with ISCT minimum criteria [4]. These indicators collectively predict immunosuppressive potency and engraftment capacity better than any single metric alone [5,10–14].

Digitization occurs at the quality assessment stage: after laboratory measurements are completed (microscopy, FLIM, flow cytometry, blood panel), the MDAF AI engine converts these raw measurements into a normalized MQS scalar and S/A/B/C/D grade. A digital certificate is then issued and recorded on the blockchain. Raw biological materials are never on-chain; only cryptographic hashes, DID references, and ZKP verification results are recorded. The digital twin is a computational vector representation of a patient's clinical profile (inflammatory, metabolic, and regenerative dimensions) used for batch suitability scoring; it is not a physical simulation or physiological model.

Blockchain is required rather than conventional access control or centralized databases because MSC biobanking involves multiple independent institutions (manufacturing laboratories, biobanks, hospitals, regulators, auditors) that do not share a common trusted infrastructure. Traditional cryptographic methods (e.g., TLS, symmetric encryption) protect data in transit but do not provide an immutable, independently verifiable audit trail across institutional boundaries without a trusted third party. Conventional distributed databases require mutual trust and a central administrator, which is inappropriate in multi-institutional cell therapy ecosystems where no single institution governs the chain of custody. QBFT permissioned blockchain provides Byzantine Fault-Tolerant consensus with deterministic finality and role-based access control, enabling independent verification of lifecycle events (donor consent, manufacturing, release, transfer, clinical use) by any authorized participant without disclosing sensitive raw data.

This paper proposes MDAF, an integrated applied engineering framework that addresses all three at the architectural/prototype level within a single engineering design.

It operationalizes the conceptual proposition that MSC quality can be transformed from scientific quantification into standardized grading and, ultimately, into a verifiable digital certificate through five technical subsystems. To address these three structural problems, this study is guided by three research questions:

RQ1: Can MSC batch quality be transformed into a standardized, reproducible, and cryptographically verifiable digital quality score through multimodal AI fusion?
RQ2: Can a hybrid ZKP–blockchain architecture provide privacy-preserving, tamper-proof lifecycle traceability for MSC batches across multi-institutional environments?
RQ3: Can an AI-based digital twin matching module serve as a decision-support tool for patient-specific MSC batch suitability assessment?

This study makes five contributions. First, a multimodal Cross-Attention AI engine is proposed in which three core modalities (M1·M2·M3) and two auxiliary QC groups (M4·M5) are fused via bidirectional attention to produce MQS and an S/A/B/C/D five-tier grade. Second, an MSC digital certificate schema is designed using a hybrid dual-layer architecture separating on-chain (hash/DID/ZKP results) from off-chain (raw biodata, AES-256) data, designed to support alignment with personal data protection and bioethics requirements while preserving blockchain immutability. Third, a two-tier Groth16 ZKP policy is introduced in which a minimum release eligibility criterion (REP: MQS $\geq$ 70) and premium quality certification (PQP: MQS $\geq$ 85) are separated into independent circuits, eliminating threshold conflicts. Fourth, a QBFT permissioned blockchain architecture is designed as a multi-institutional trust infrastructure combining Hyperledger Besu-based smart contracts with DID-RBAC. Fifth, end-to-end pipeline feasibility is validated in a controlled synthetic environment (n = 2000) including baseline and ablation experiments. Contribution 1 addresses RQ1; Contributions 3 and 4 address RQ2; Contribution 5 addresses RQ3.

It is important to note that the present study is an engineering feasibility study, not a clinical efficacy validation study. The reported performance metrics (R^2 = 0.70, macro-F1 = 0.40, AUROC = 0.92) demonstrate a minimum architectural feasibility—namely, that the proposed Cross-Attention fusion pipeline can learn the intended multimodal dependency structure of a rule-engine-based synthetic dataset. These figures should not be interpreted as predictive accuracy on real MSC batches. Biological and clinical validation with real cohort data is mandatory before any operational deployment.

A note on digital twin hardware infrastructure: while this study focuses on the software and algorithmic architecture, operational deployment would require edge-computing nodes or hospital servers for real-time patient profile computation connected via secure API to the blockchain ledger; hardware integration is deferred to future work (F4). Additionally, AI-based approaches have been increasingly integrated into digital twin systems in medicine [15–17], enabling predictive modelling and personalized therapy planning; the present study extends this paradigm to MSC batch–patient matching. The remainder of this paper is organized as follows. Section 2 reviews related work on MSC quality assessment, AI-based cell analysis, blockchain for biodata management, ZKP-based privacy verification, and digital twin applications. Section 3 describes the multimodal Cross-Attention AI engine. Section 4 presents the digital assetization framework and certificate design. Section 5 details the two-tier ZKP architecture. Section 6 describes the QBFT blockchain system. Section 7 introduces the digital twin matching module. Section 8 reports synthetic data pilot results. Section 9 discusses contributions, limitations, and future work. Section 10 concludes the paper. Figure 1 provides an overview of the complete MDAF architecture, illustrating the end-to-end flow from multimodal data collection through AI MQS scoring, digital certificate issuance, two-tier ZKP verification, QBFT ledger recording, and digital twin patient–batch matching.

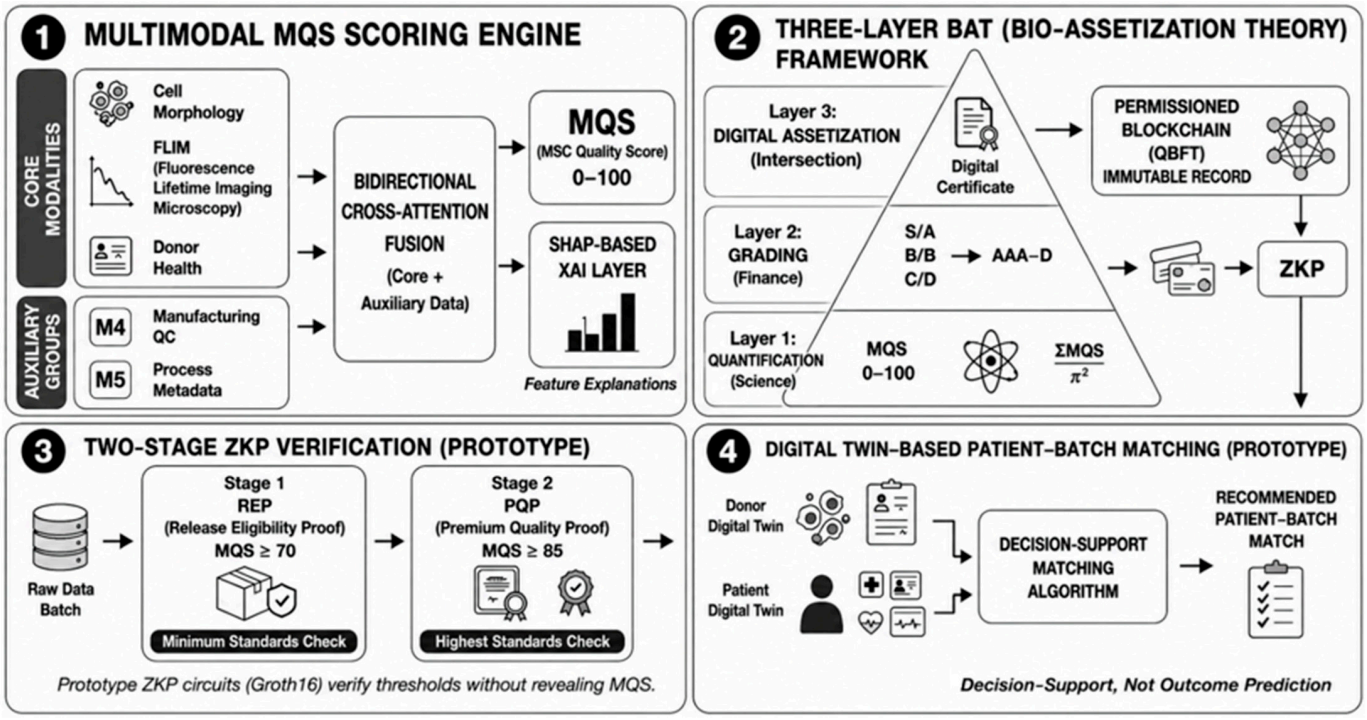


Figure 1. Overall architecture of the MSC Digital Assetization Framework (MDAF).

The engineering prototype consists of four integrated modules: (1) a multimodal MSC quality score (MQS) engine that integrates heterogeneous biological and manufacturing data through Cross-Attention fusion; (2) a three-layer Bio-Assetization Theory (BAT) architecture that converts quantified MSC quality into a verifiable digital certificate; (3) a prototype two-stage zero-knowledge proof (ZKP) verification module implementing Release Eligibility Proof (REP) and Premium Quality Proof (PQP) circuits on a permissioned Practical Byzantine Fault Tolerance (QBFT) blockchain with Decentralized Identifier (DID)-based role-based access control (RBAC); and (4) a digital twin-based patient–batch matching module that provides computational decision support for personalized MSC selection. This engineering prototype demonstrates the architectural feasibility of integrating multimodal AI, blockchain, cryptographic verification, and digital twin technologies, and should not be interpreted as a clinically validated regenerative medicine system.

2. Related Work

2.1. MSC Quality Assessment

Traditional MSC quality control relies on viability, sterility, surface markers, and differentiation capacity, failing to capture the complex biological state that determines therapeutic efficacy [4]. Beane et al. [10] reported reduced proliferation and immunosuppressive function in MSCs from aged donors; Wagner et al. [11] demonstrated an inverse correlation between donor metabolic disease and osteogenic differentiation capacity. NAD^+ has attracted attention as an MSC quality indicator owing to its central role in mitochondrial energy metabolism and epigenetic regulation [13,14]. Doubling time (DT) is a direct proliferation metric correlated with the therapeutic efficacy of clinical batches [7]. More recently, regulatory guidance and consensus documents have further emphasised the need for standardized, quantitative MSC release criteria integrating multiple quality dimensions [18].

The selection of quality assessment dimensions in MDAF is grounded in the biological mechanisms linking each indicator to therapeutic potency. Morphological features

predict structural integrity and cellular viability under physiological stress [15]. FLIM-derived τ_mean reflects the intracellular NAD^+ redox ratio: shortened τ_mean indicates a higher NAD^+/NADH ratio associated with active oxidative phosphorylation, which correlates with immunosuppressive secretion capacity [13,14]. Doubling time (DT) is the most direct predictor of scalable therapeutic batch yield [7]. Donor blood panel variables (hsCRP, HbA1c, BMI) capture systemic inflammatory and metabolic burden that reduce MSC immunomodulatory function in older or metabolically compromised donors [10,11]. ISCT-minimum surface markers (CD73, CD90, CD105 positive; CD34, CD45 negative) confirm MSC identity per international consensus [4]. Manufacturing metadata (passage, cryopreservation, transport time) track process deviations associated with quality degradation [18]. Together, these five modality groups represent the most clinically and biologically justified set of MSC quality dimensions currently available in the literature.

Research Gap: Despite significant progress in MSC quality control, existing approaches primarily evaluate isolated quality attributes such as viability, surface marker expression, or differentiation capacity. Few studies have attempted to integrate morphological, metabolic, donor-derived, and manufacturing-related variables into a unified quantitative quality index. Furthermore, no widely accepted framework currently exists for converting multidimensional MSC quality information into a standardized and transferable digital quality representation suitable for lifecycle management and inter-institutional use.

2.2. AI-Based Cell Quality Analysis

Deep learning-based cell image analysis has achieved expert-level morphological classification [19,20]. Cellpose 2.0 [21,22] adapts to specific cell types via few-shot fine-tuning. Cross-Attention mechanisms [23,24] perform well for heterogeneous data integration; to the authors' knowledge, the application of Cross-Attention-based multimodal fusion to MSC quality scoring with baseline and ablation validation remains rare.

Research Gap: Although deep learning has demonstrated strong performance in cellular image analysis and phenotype classification, most existing studies focus on single-modality inputs, particularly microscopy images. The integration of heterogeneous biological modalities—including morphology, metabolism, donor health, and manufacturing metadata—through multimodal attention mechanisms remains relatively unexplored in MSC quality assessment. In particular, Cross-Attention-based quality scoring frameworks combined with explainable AI and engineering validation have rarely been reported.

2.3. Blockchain for Biodata Management

Blockchain applications in healthcare have expanded to clinical trial data integrity [25], pharmaceutical supply chain traceability [26], and patient consent management [27,28]. Integrated lifecycle management of cell therapy raw materials remains insufficiently explored [29]. To the best of the authors' knowledge, a framework integrating MSC quality quantification with blockchain-based lifecycle traceability has not been reported in the searched literature. Recent reviews confirm growing interest in blockchain-based biomedical data management [29] and AI-integrated cell therapy quality systems [30], further motivating the present integrated approach.

Research Gap: Existing blockchain applications in healthcare primarily focus on medical records, clinical trial management, consent tracking, or pharmaceutical supply chains. However, the end-to-end lifecycle management of MSC products—from donor consent and manufacturing to release certification and clinical use—remains insufficiently addressed. Moreover, current systems generally lack integration with AI-derived biological quality assessment and cryptographic quality verification mechanisms.

2.4. ZKP-Based Privacy-Preserving Verification

A ZKP is a cryptographic system in which a prover demonstrates the truth of a statement without revealing the witness [31,32]. Groth16 [33] is the most widely deployed ZKP implementation, offering constant proof size and fast verification. Bharath Babu and Jothi [34] proposed a selective disclosure framework applying ZKP to electronic medical records. To the authors' knowledge, a two-tier ZKP policy applied to cell therapy raw material quality certification has not been reported.

Research Gap: While ZKP technologies have been increasingly applied to privacy-preserving healthcare data sharing, their use has largely focused on identity verification, selective disclosure, and access control. To the best of the authors' knowledge, no prior study has implemented a dual-policy ZKP architecture that separately verifies minimum release eligibility and premium quality certification for cell therapy products. The integration of ZKP with AI-based biological quality scoring and blockchain-supported lifecycle traceability therefore remains largely unexplored.

2.5. Digital Twin-Based Personalized Medicine

Patient-level digital twins are being increasingly explored for treatment response prediction and personalized therapy planning [15–17]. Björnsson et al. [17] demonstrated that integrating multi-omics data with computational disease models can support personalized oncology decision-making. Katsoulakis Tet al. [16] proposed digital twin frameworks for real-time patient monitoring and adaptive therapy adjustment. AI models—including machine learning, deep learning, and multimodal fusion—serve as the predictive inference layer in these systems, mapping patient state vectors to expected treatment outcomes. This AI-in-digital-twins paradigm is directly relevant to the present study: the MDAF digital twin matching module extends prior AI-digital twin work from outcome prediction to biological product matching, linking patient inflammatory, metabolic, and regenerative profiles to MSC batch quality scores. However, a suitability architecture specifically matching patient clinical profiles to cell therapy batch characteristics remains insufficiently studied in the MSC domain.

Research Gap: Current digital twin research in medicine primarily focuses on disease progression modelling, treatment response prediction, and personalized therapy optimization. However, the application of digital twin concepts to biological product selection and patient–product suitability assessment remains limited. In particular, no established framework currently links AI-derived MSC quality profiles with patient-specific inflammatory, metabolic, and regenerative characteristics to support personalized MSC batch matching. This gap motivates the digital twin matching module proposed in this study.

2.6. Research Gap Summary

Collectively, the reviewed literature demonstrates substantial advances in MSC quality assessment, AI-based cell analytics, blockchain-enabled traceability, zero-knowledge proof verification, and digital twin technologies. However, these research streams have largely evolved independently, each addressing one or two of the structural problems identified in Section 1.

To the best of the authors' knowledge, no prior study has integrated (i) multimodal MSC quality quantification via Cross-Attention AI fusion, (ii) a two-tier cryptographically verifiable quality certification policy (ZKP REP/PQP), (iii) QBFT permissioned blockchain-based lifecycle traceability with DID-RBAC, and (iv) AI-based digital twin patient–batch matching within a single end-to-end engineering architecture validated under controlled conditions. The proposed MDAF was developed to address this integration gap.

3. Multimodal Cross-Attention MSC Quality Scoring Engine

3.1. Overall Engine Architecture—3 Core Modalities + 2 Auxiliary QC Groups

The scoring engine consists of three core modality encoders (M1·M2·M3), two auxiliary QC input groups (M4·M5), a bidirectional Cross-Attention fusion module, and dual output heads. M4 and M5 are each encoded by a two-layer MLP (input → 128 → d_proj, ReLU, BatchNorm), and then concatenated with the M3 representation in the projection space before participating in Cross-Attention. The rationale for treating M4 and M5 as auxiliary QC groups rather than core modalities is as follows: flow cytometry markers (M4) provide QC pass/fail signals (ISCT surface marker compliance, viability threshold) that define product identity rather than continuous biological quality variation; manufacturing metadata (M5) captures process-control parameters (passage number, cryopreservation history, transport time) that condition but do not independently determine biological quality. These signals augment the core biological quality signal encoded in M1–M3 rather than driving it independently. The overall architecture of the multimodal scoring engine is illustrated in Figure 2. Table 1 summarizes all five modality groups, their input data, encoder architectures, output dimensions, and roles.

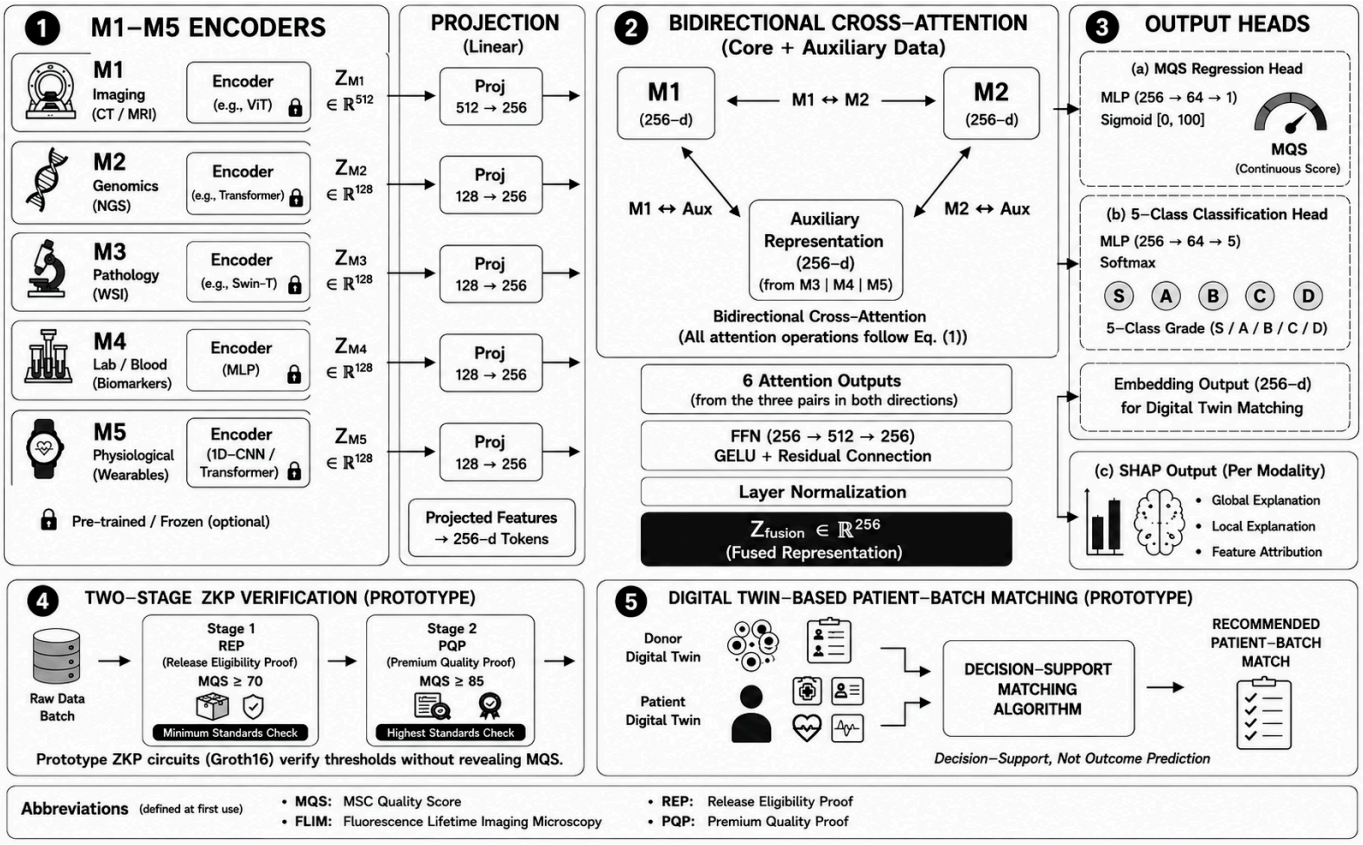


Figure 2. Multimodal Cross-Attention AI engine: M1–M5 encoders → projection → bidirectional Cross-Attention fusion → MQS regression head + five-class classification head + SHAP output.

3.2. Modality M1: Morphological Feature Encoder

3.2.1. Cellpose 2.0-Based Cell Segmentation

For MSC-specific adaptation, the base cyto3 model was fine-tuned on 50 bone marrow-, adipose-, and umbilical cord blood-derived MSC images (learning rate 0.005, weight decay 1×10^{-5}, 100 epochs). CLAHE preprocessing corrects illumination non-uniformity in phase contrast images. Six morphological scalars (cell area, aspect ratio, cytoplasm-to-nucleus ratio, filopodium length, nuclear circularity, and boundary sharpness) are deterministically computed

from Cellpose masks via cellpose_segmentor.compute_morph_features_from_mask() and concatenated with CNN feature vectors.

Table 1. Input modality and auxiliary QC group overview.

Group	Input Data	Encoder	Output Dimension	Category
M1: Cell morphology	Phase contrast/fluorescence images (256 × 256 px)	Cellpose 2.0 + ResNet-50 + MLP	d = 512	Core modality
M2: Metabolic-proliferative	NAD^+-related metabolic proxy score, τ_mean (ns), DT (h/passage)	three-layer MLP	d = 128	Core modality
M3: Donor health	15 blood panel items + BMI	four-layer MLP	d = 256	Core modality
M4: Flow cytometry (auxiliary)	CD73/CD90/CD105 positive, CD34/CD45 negative, viability	two-layer MLP	d = 64	Auxiliary QC
M5: Manufacturing process (auxiliary)	Passage, culture duration, freeze history, transport time, medium	two-layer MLP	d = 64	Auxiliary QC

3.2.2. ResNet-50 Encoder

Single-cell patches (256 × 256 px) are encoded by ImageNet-pretrained ResNet-50, with the final fully connected layer replaced by a 512-d projection head. Only Block 4 and the head are fine-tuned. Grad-CAM visualizes cell regions contributing to MQS.

3.3. Modality M2: FLIM-Based Metabolic-Proliferative Encoder

M2 encodes three scalars derived from FLIM. (1) NAD^+-related metabolic proxy score: a normalized FLIM-derived metabolic feature heuristically related to τ_mean changes, used as a surrogate indicator of intracellular NAD^+-associated metabolic state and not a direct biochemical NAD^+ concentration measurement [13]. (2) τ_mean (ns): FLIM field-of-view mean photon arrival time—a direct measure of cellular metabolic activity. The synthetic parameter ranges were defined based on values reported in the MSC FLIM and proliferation literature [13]. (3) DT (h/passage): passage-wise doubling time—a direct measure of proliferative potential. τ_mean (FLIM fluorescence lifetime) and DT (proliferative doubling time) measure distinct biological dimensions. After z-score normalization, a three-layer MLP (input → 64 → 128 → 128, ReLU, BN, dropout 0.2) generates a 128-d latent vector z_M2. Table 2 presents the FLIM-based indicator reference ranges used to calibrate synthetic data generation and interpret M2 encoder outputs.

3.4. Modality M3 and Auxiliary QC Groups

The M3 donor health encoder maps a 15-dimensional vector to a 256-d representation z_M3 via a four-layer MLP. Auxiliary inputs M4 (flow cytometry) and M5 (manufacturing process) are each encoded to 64-d by two-layer MLPs, concatenated with z_M3, and participate in Cross-Attention. The surface marker thresholds used in M4 (CD90 Fold/ISO ≥ ×100, CD73 Fold/ISO ≥ ×10, CD34/CD45 negative) were defined in accordance with ISCT minimum criteria [4] and values reported in the MSC quality-control literature [7]. Real cohort validation of these threshold settings remains to be studied in future work. The donor health panel items and their potential mechanisms affecting MSC quality are summarized in Table 3.

Table 2. M2 FLIM-based metabolic-proliferative indicator reference ranges (bone marrow-derived MSC, passages 2–4).

Indicator	Unit	Grade S/A	Grade B	Grade C
NAD^+-related metabolic proxy score (τ_mean-derived)	normalized score	≥0.75	0.50–0.74	0.25–0.49
τ_mean (FLIM mean fluorescence lifetime)	ns	≤2.0	2.1–2.6	2.7–3.2
DT (passage-wise doubling time)	h/passage	≤24	25–32	33–42

The τ_mean range (0.5–3.5 ns) and DT range were defined based on values reported in the MSC FLIM and proliferation literature [7,13]. The DT criterion (Grade S/A: ≤24 h) is consistent with empirically measured doubling times (22.4–26.2 h) reported in a GMP-compliant clinical-grade BM-MSC expansion protocol [35].

Table 3. Donor health panel (M3) items and MSC quality impact mechanisms.

Panel	Item	Mechanisms Affecting MSC Quality	References
Inflammation	hsCRP, IL-6, TNF-α, ESR	Chronic inflammation → reduced immunosuppressive function, increased ROS	[10,11]
Metabolism	Fasting glucose, HbA1c, HOMA-IR	Insulin resistance → mitochondrial dysfunction, NAD^+ depletion	[14,36]
Lipid	TC, LDL, HDL, TG	Hyperlipidemia → reduced membrane fluidity, lipotoxicity	[12]
Immune	NK activity, lymphocyte ratio, IgG	Immune function → predicts MSC immunomodulatory capacity	[10]
Anthropometric	BMI	Chronic low-grade inflammation, adipogenic differentiation bias	[12,37]

3.5. Bidirectional Cross-Attention Fusion Module

z_M1 $\in \mathbb{R}^{512}$, z_M2 $\in \mathbb{R}^{128}$, and (z_M3‖z_M4‖z_M5) $\in \mathbb{R}^{384}$ are linearly projected to $\mathbb{R}^{256}$. Bidirectional Multi-Head Attention (h = 8, d_k = 32) is applied to three pairs (M1 ↔ M2, M1 ↔ M3, M2 ↔ M3):

$$CrossAttn(A, B) = softmax(Q_A \cdot K_B\hat{}T / \sqrt{d_k}) \cdot V_B \quad (1)$$

The six attention outputs pass through an FFN (256 → 512 → 256, GELU, residual connection) and Layer Normalization to produce a 256-d fused representation z_fusion.

Important distinction: Cross-Attention weights and SHAP contribution values are complementary but not equivalent. Attention weights quantify the representational alignment structure, while SHAP values quantify marginal contributions to the output. Both are provided separately in the certificate explanation report.

3.6. Output Heads and MQS Grade System

(a) Regression head: MLP (256 → 64 → 1), sigmoid-scaled to [0, 100] → continuous MQS. (b) Five-class classification head: MLP (256 → 64 → 5), softmax → S/A/B/C/D grade. Combined loss: L = 0.6·MSE(MQS) + 0.4·CrossEntropy(grade). z_fusion is also output as a batch embedding vector for digital twin matching. Table 4 defines the MQS grade system, proposed policy interpretation, ZKP policy per grade, and regulatory notes.

3.7. SHAP-Based Explainability (XAI) Layer

DeepSHAP and kernel SHAP [38] are used in combination to provide: (i) global level—modality- and feature-level mean |SHAP| rankings; (ii) instance level—waterfall plots; (iii) interaction level—SHAP interaction values.

Table 4. MQS grade system and proposed policy interpretation.

Grade	MQS Range	Proposed Suitability Category	ZKP Policy	Regulatory and GMP Review Criteria
S	90–100	Highest proposed quality tier	REP + PQP	Priority release candidate after regulatory/GMP review
A	80–89	Proposed policy is release-suitable (primary)	REP + PQP	Release candidate upon standard QC fulfillment (regulatory/GMP review required)
B	70–79	Proposed policy is release-suitable (secondary)	REP only	Conditionally release candidate after additional QC or extended culture
C	55–69	Research grade	N/A	Non-clinical in vitro research only
D	0–54	Not releasable	N/A	Pre-filtered by rule-based QualityGateway; Rejection Record only

N/A: Not applicable. Important note: Grade D batches are excluded from release at the QualityGateway pre-policy filter and are not subject to ZKP proof submission. Only a Rejection Record is written on-chain for audit purposes. Note: Grade D was retained in the training dataset to improve grade-boundary robustness and rejection sensitivity; exclusion applies only at the QualityGateway inference stage, not during model training. The "proposed policy" in this table represents an engineering system design criterion and does not imply actual clinical release authorization, GMP approval, or regulatory agency clearance.

4. MSC Digital Assetization Framework and Certificate Design

4.1. Three-Layer Assetization Conceptual Framework

Layer 1—Quantification (Science): AI-driven MQS reduces multidimensional biological measurements to a reproducible scalar quality index on [0, 100].

Layer 2—Grading (Finance): MQS is mapped to an S/A/B/C/D five-tier grade. The grading structure is analogous to credit-rating systems (AAA–D) in terms of standardized quality communication—supporting standardized quality communication across multi-stakeholder networks [39].

Building on the information asset theory [39] and intangible asset management literature [40], BAT operationalizes five necessary and sufficient conditions for a biological product to function as a bio-asset within a digital multi-institutional environment.

Layer 3—Digital Assetization (Intersection): Each batch is issued a unique digital certificate permanently recorded on the permissioned blockchain.

BAT is distinguished from generic digital certification by five necessary and sufficient conditions for a biological product to constitute a bio-asset: (C1) Measurability—quality attributes must be quantifiable on a standardized, reproducible scale (operationalized in MDAF by MQS on [0, 100]); (C2) Standardization—attributes must be mappable to a structured grading schema enabling inter-institutional comparison (operationalized by S/A/B/C/D grading); (C3) Verifiability—quality claims must be cryptographically verifiable without revealing proprietary raw data (operationalized by Groth16 ZKP REP/PQP circuits); (C4) Traceability—the full lifecycle from donor consent through clinical use must be immutably recorded (operationalized by QBFT blockchain); and (C5) Transferability—the bio-asset record must be transferable across institutions with role-specific access control (operationalized by DID-RBAC). A biological product becomes a bio-asset if and only if all five conditions are simultaneously satisfied. This distinguishes BAT from prior digital certification frameworks [25–27], which address Traceability (C4) and Verifiability (C3)

but not Measurability (C1) or Standardization (C2). MDAF is to the authors' knowledge, among the first engineering systems designed to instantiate all five BAT conditions for MSC products.

Definition note: In this study, digital assetization does not imply tokenization, securitization, or public financial trading of MSC batches. It refers to the conversion of a biological batch into a verifiable, traceable, quality-certified digital record unit for managing ownership, provenance, and auditability. The analogy to financial grading exploits structural similarities in trust-based value communication and does not confer the character of a financial asset.

4.2. MSC Digital Certificate Structure

Each MSC batch is issued a dual-layer digital certificate comprising on-chain metadata and off-chain detailed data. Grade S/A batches record both π_REP and π_PQP on-chain; Grade B batches record only π_REP. The on-chain fields of the proposed MSC digital certificate are summarized in Table 5.

Table 5. MSC digital certificate on-chain fields.

Field Name	Type	Description	Storage Location
cert_id	bytes32	Batch unique identifier (SHA-3)	On-chain
batch_hash	bytes32	Batch ID hash	On-chain
mqs_grade_hash	bytes32	Hash of MQS value and grade (S/A/B/C/D)	On-chain
data_hash	bytes32	SHA-256 hash of off-chain data package	On-chain
issuer_did	string	Issuing institution DID (W3C DID Core 1.0 [41])	On-chain
release_zkp_proof	bytes	REP Groth16 π = (A, B, C)	On-chain
premium_zkp_proof	bytes	PQP Groth16 π (Grade S/A only; 0x00 if N/A)	On-chain
zkp_policy_id	bytes32	Applied ZKP policy (0x01 = REP, 0x02 = PQP)	On-chain
lifecycle_event_hash	bytes32	Latest lifecycle event hash	On-chain
passage_number	uint8	Passage number at issuance	On-chain
cell_source_code	bytes2	Cell source code (BM/AT/UCB)	On-chain
model_version	string	AI model version identifier	On-chain
issue_timestamp	uint256	Issuance Unix timestamp	On-chain
certificate_status	uint8	Valid (0)/Superseded (1)/Revoked (2)	On-chain

N/A: Not applicable. Off-chain items (AES-256 encrypted): Raw microscopy images, FACS raw data (FCS), FLIM measurement files, donor blood test results, SHAP analysis reports, culture history logs, GMP manufacturing records, regulatory compliance documents.

4.3. Lifecycle Event Schema

The certificate records 15 events: DonorConsentRegistered → TissueCollected → MSCIsolated → CultureInitiated → PassageUpdated → QualityAssessmentCompleted → MQSComputed → CertificateIssued → ZKPVerificationCompleted → StorageEventRecorded → TransportEventRecorded → ClinicalReleaseRequested → DigitalTwinMatchRecorded → ClinicalUseCompleted → CertificateRevoked/Expired.

5. Zero-Knowledge Proof (ZKP)-Based Two-Tier Quality Verification Architecture

5.1. Rationale for Two-Tier Policy Separation and Threshold Design

MSC quality verification serves two objectives: (A) confirmation of minimum release eligibility under the proposed policy (Grade B or above, MQS $\geq$ 70; release eligibility under the proposed policy), and (B) premium grade certification (Grade A/S, MQS $\geq$ 85). Processing these with a single ZKP circuit would produce a conflict between the QualityGateway threshold (MQS $\geq$ 70) and the ZKP threshold (MQS $\geq$ 85). This study resolves the conflict by separating them into two independent Groth16 circuits: REP and PQP.

Threshold design rationale and scope: In REP, passage ≤ 5 and viability $\geq 90\%$ are proposed policy values referencing generally recommended minimum criteria for MSC quality control in the literature [4,7]. In PQP, MQS ≥ 85, passage ≤ 3, and viability $\geq 95\%$ are the research team's initial design values for premium quality criteria. These thresholds are proposed parameters for this study's synthetic pilot and do not represent universal clinical release standards. In actual operational deployment, they must be adjusted according to institution-specific GMP standards, regulatory requirements, and validated potency assay correlation results.

5.2. ZKP Circuit Design (Prototype circom2 Threshold Circuits)

Circuit implementation status notice:

REP/PQP circuits were implemented as simplified prototype circom2 threshold-verification circuits for local benchmarking (REP: 47 constraints; PQP: 53 constraints). They are not production-grade circuits and use simplified hash-preimage equality rather than optimized in-circuit cryptographic hashing. Full production-scale circuit design, trusted setup, and deployment-level on-chain verification are planned for future work (F3). Empirical timing results are reported in Section 8.8.

REP circuit design (prototype circom2):

$$REP_Valid(w,x) = 1 \; iff \; MQS \geq 70 \wedge viability \geq 90\% \wedge passage \leq 5 \wedge sterility = PASS \wedge SHA\text{-}256(raw_data) = data_hash \quad (2)$$

Private inputs (w_REP): Actual MQS, viability, passage number, sterility result, raw data hash preimage.

Public inputs (x_REP): Certificate ID, policy_id = 0x01, data_hash, verification_result.

PQP circuit design (prototype circom2):

$$PQP_Valid(w,x) = 1 \; iff \; MQS \geq 85 \wedge viability \geq 95\% \wedge passage \leq 3 \wedge sterility = PASS \wedge markers = VALID \wedge SHA\text{-}256(raw_data) = data_hash \quad (3)$$

Private inputs (w_PQP): Actual MQS, viability, passage number, marker status, sterility result, raw data hash preimage.

Public inputs (x_PQP): Certificate ID, policy_id = 0x02, data_hash, verification_result.

5.3. Common Proof System: Groth16

Both REP and PQP circuits adopt Groth16 [33] as the underlying proof system. Groth16 produces a constant-size proof $\pi = (A, B, C) \in G_1 \times G_2 \times G_1$ (192 bytes total), independent of circuit complexity. A one-time trusted setup generates a Common Reference String (CRS) per circuit, producing a proving key pk (held by the laboratory prover) and a verification key vk (deployed on-chain to the ZKPVerifier contract). Proof generation is defined in Equation (4), and verification is in Equation (5). On-chain verification uses the EIP-1108 BN254 pairing precompile (estimated ~113,000 gas/circuit; empirical on-chain measurement is planned for future work F3).

$$\pi = Groth16.Prove(CRS, w, x) \quad (4)$$

$$Groth16.Verify(CRS_vk, x, \pi) \rightarrow \{TRUE \mid FALSE\} \quad (5)$$

5.4. Two-Tier ZKP Workflow

The complete two-tier ZKP workflow is illustrated as a schematic diagram in Figure 3, showing the data flow from MSC quality assessment through proof generation, on-chain verification, and certificate issuance. The key operational steps are as follows.

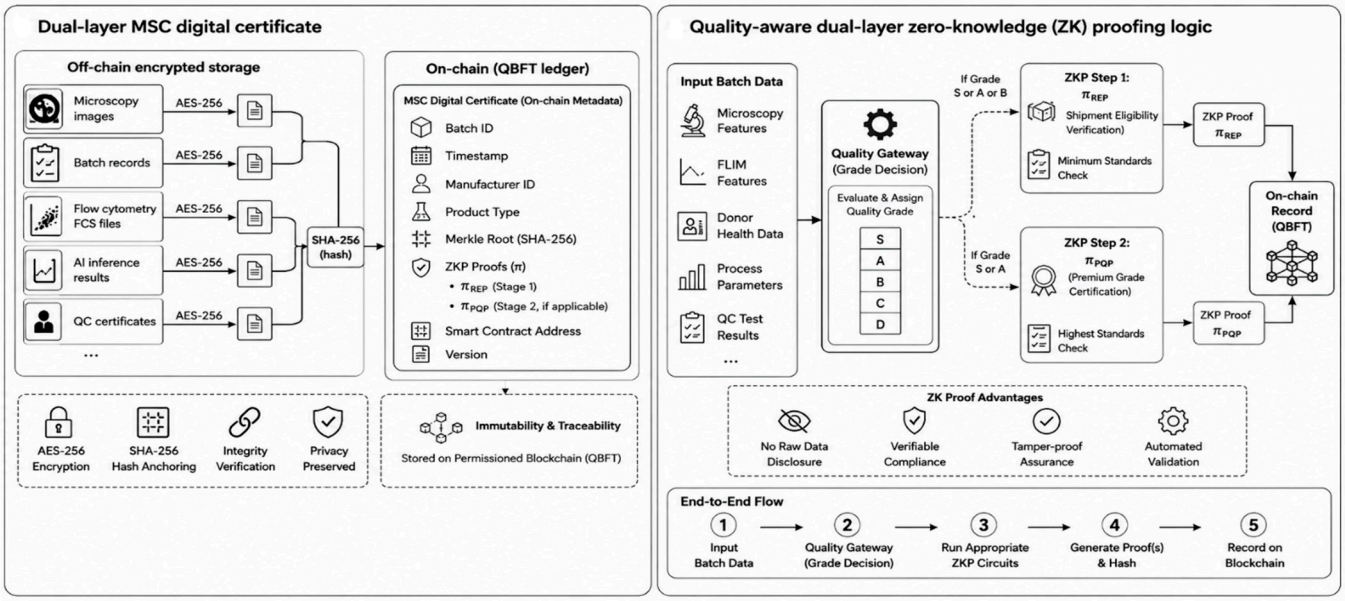


Figure 3. MSC digital certificate architecture with two-tier ZKP verification and on-chain/off-chain data separation. Off-chain raw data are hashed using SHA-256 for on-chain registration, with REP/PQP circuit branching and QualityGateway gate logic.

Steps 1–2: Laboratory completes MSC quality assessment → MQS and grade computed, data_hash generated.

Step 3: snarkjs generates π_REP (Grade B or above). For Grade S/A, π_PQP is additionally generated.

Steps 4–5: QualityGateway verifies π_REP (and π_PQP if applicable) on-chain via EIP-1108 precompile → if REP_TRUE, issueCertificate() is executed.

Step 6: ZKPVerificationCompleted event emitted → QBFT deterministic finality. Grade D batches are blocked by the pre-policy filter; only a Rejection Record is written.

6. QBFT Permissioned Blockchain System Architecture

6.1. Rationale for Permissioned Blockchain Adoption

Reasons for adopting QBFT [42,43]-based permissioned blockchain: (i) participants are limited to laboratories, biobanks, hospitals, auditors, and regulators; (ii) raw data contain sensitive biological and medical information; (iii) deterministic finality ($f < n/3$ fault tolerance) and access control are required; and (iv) legal accountability traceability is mandatory. Table 6 presents the four-layer system architecture overview of MDAF.

Table 6. Four-layer system architecture overview.

Layer	Component	Core Function	Technology Stack
L1: Data collection	Microscope, FLIM equipment, blood analyzer	Raw data acquisition	DICOM, FCS, HL7 FHIR
L2: AI processing	MQS engine, ZKP prover (REP/PQP), certificate builder	MQS computation, two-tier proof generation	Python, PyTorch, snarkjs, circom
L3: Blockchain/proof	Besu nodes, smart contracts, DID registry, IPFS	Immutable records, ZKP verification, DID management	Solidity, QBFT, Hyperledger Besu v23.x
L4: Application services	Hospital portal, donor app, regulatory audit interface	User access, query, audit	React 18.2.0, DID Auth, OAuth 2.0, REST

6.2. Four-Layer System Architecture

The proposed system is organized into four functional layers: data collection, AI processing, blockchain/proof, and application services. Table 6 summarizes the components, core functions, and technology stack of each layer.

6.3. Smart Contract System (OpenZeppelin UUPS Pattern)

CertificateRegistry: Certificate issuance, transfer, query, and revocation.
issueCertificate(batchHash, issuerDID, modelVer, gradeHash, pi_REP, pi_PQP)
updateLifecycleEvent(certID, eventType, eventHash)
revokeCertificate(certID, reasonHash)
recordDigitalTwinMatch(certID, matchScoreHash)
LifeBankDirectory: DID registration, role assignment, VC verification, and RBAC management.
QualityGateway (two-tier gate): REP verification + grade/MQS dual-condition gate.
require(grade $\leq$ GRADE_B && mqs $\geq$ 70, "Quality gate: minimum threshold");
`require(ZKPVerifier.verifyREP(CRS_vk, x_REP, pi_REP), "REP: invalid proof");`
`if (grade` $\leq$ `GRADE_A) require(ZKPVerifier.verifyPQP(CRS_vk, x_PQP, pi_PQP), "PQP: invalid");`
`// D grade: immediate revert without role check` $\rightarrow$ `rejection record only`

6.4. DID-Based Role Access Control

The proposed DID-based role access control framework assigns identity credentials and access privileges according to participant roles. Table 7 summarizes the DID/VC usage and corresponding access privileges for each stakeholder role.

Table 7. Participant role-wise DID/VC usage and access privileges.

Role	DID Purpose	Held VCs	Access Rights
Donor	Anonymized identity management	Consent VC, health check VC	View own batch history
Biobank	Institutional identity, batch ownership	Institutional certification VC, GMP compliance VC	Batch issuance, transfer, lifecycle recording
Lab/analytical body	Scoring signature authority	Quality certification VC, CAP accreditation VC	Score recording, REP/PQP ZKP generation
Hospital/clinical institution	Authorized recipient eligibility	Prescription qualification VC, clinical trial approval VC	Receive REP-passed batches (Grade B or above)
Auditor	Audit access	Audit authority VC	Read-only full history
Regulatory agency	Regulatory review authority	Regulatory authority VC	Read-only, compliance reporting

7. Digital Twin-Based Patient–Batch Matching Module

7.1. Module Design Philosophy

This module is a computational decision-support tool, not a clinical outcome prediction model. The matching score is a reference metric assisting clinician judgment and does not guarantee actual treatment suitability.

7.2. Matching Score Computation

$$MatchScore(patient, batch) = w_1 \cdot \Phi_inflammation + w_2 \cdot \Phi_metabolic + w_3 \cdot \Phi_regenerative + w_4 \cdot MQS_normalized + penalty_safety \quad (6)$$

Weights w_1 = 0.30, w_2 = 0.25, w_3 = 0.25, w_4 = 0.20 are empirical initial values for the pilot. Validation through clinical outcome correlation studies prior to operational deployment is mandatory. Table 8 defines the four decision-support tiers for matching score ranges.

Table 8. Digital twin matching tier (decision-support matching tier).

Matching Score	Grade	Interpretation
≥90	Optimal	Highly suitable; primary decision-support recommendation (clinician final judgment required)
80–89	Suitable	Suitable; standard decision-support recommendation
70–79	Conditional	Additional clinical review recommended
<70	Not recommended	Unsuitable or insufficient information; alternative exploration recommended

8. Synthetic Data Pilot Validation

8.1. Pilot Scope and Disclaimer

Important notice—Scope of this study: This study is an engineering feasibility study, not a clinical efficacy validation study. Numerical results in this study comprise three distinct categories: (a) AI model performance (Sections 8.3–8.7) is derived from a rule-engine-based synthetic dataset (n = 2000) and demonstrates architectural plausibility, not biological or clinical predictive accuracy; (b) ZKP benchmark (Section 8.8) represents empirical timing measurements obtained by running circom2-compiled REP/PQP Groth16 circuits locally via snarkjs on the workstation described in Section Experimental Environment (Intel Core i5-8400 @ 2.80 GHz); and (c) QBFT blockchain performance (Section 8.9) represents empirical measurements obtained by submitting real transactions to a four-node Hyperledger Besu QBFT network running in Docker Desktop on the same workstation. None of these results constitutes biological or clinical validation. The term "digital assetization" refers strictly to verifiable digital recordation for data integrity and provenance management—not to any form of financial trading or tokenization of biological materials. Correlation studies using real MSC cohort data and clinical outcomes are mandatory prior to any operational deployment.

Experimental Environment

All experiments were conducted on a single workstation with the following hardware configuration: CPU: Intel Core i5-8400 (6 cores, 2.80 GHz); RAM: 16 GB DDR4; GPU: NVIDIA GeForce GTX 1060 6 GB (CUDA 12.x, not used for training—all model training was performed on CPU); storage: SSD. The operating system was Windows 11 (64-bit) with Python 3.10 (Anaconda), PyTorch 2.x (CPU mode), web3.py 7.x, solcx 0.8.20, and Hyperledger Besu v24.x running in Docker Desktop (four QBFT validator containers on a dedicated bridge network, subnet 192.168.200.0/24). All AI training, ZKP simulation (snarkjs), and QBFT blockchain measurements were performed on this single workstation. QBFT benchmark transactions were submitted via JSON-RPC to the local Docker network (localhost:18545–18549). These hardware conditions are reported for reproducibility; performance figures (TPS, finality) reflect local single-machine Docker networking and are expected to differ under production multi-server deployments.

8.2. Input Schema Design Rationale

The proposed MDAF input schema (M1–M5 modalities) was designed to be compatible with experimentally obtainable MSC quality-control measurements, including microscopy-derived morphological features, FLIM-derived metabolic signals, flow cytometry markers, viability, passage number, and manufacturing metadata. Justification for quality indicator selection: Each modality group was selected based on established evidence linking the indicator to MSC therapeutic potency: (i) morphological features (M1) correlate with structural integrity and attachment capacity relevant to tissue homing [19]; (ii) FLIM-derived τ_mean (M2) reflects intracellular NAD^+-dependent metabolic activity, which determines immunosuppressive cytokine secretion capacity [13,14]; (iii) passage-wise doubling time (M2) directly measures proliferative potential, the key determinant of therapeutic batch yield [7]; (iv) donor health blood panel (M3) captures systemic inflammatory and metabolic conditions known to reduce MSC potency in older or metabolically compromised donors [10,11]; (v) flow cytometry ISCT markers (M4) confirm identity compliance required for release [4]; and (vi) manufacturing metadata (M5) tracks process variables (passage, cryopreservation, transport conditions) whose deviations are associated with quality loss [18]. Proof-of-concept argument for real-world applicability: While full clinical validation on real MSC cohort data remains future work, three lines of evidence support the feasibility of applying this schema to real data: First, each input variable (M1–M5) maps directly to standard GMP laboratory measurements already routinely collected in MSC manufacturing facilities; the MDAF schema imposes no new measurement requirements. Second, the feature distribution ranges used in synthetic data generation (cell area 200–2000 μm^2, τ_mean 1.5–4.0 ns, DT 12–72 h, etc.) were calibrated to published clinical MSC quality-control data [4,7,13,35], ensuring biological plausibility of the synthetic pilot. Third, the rule engine weights (morph: 0.27, metabolic: 0.31, flow: 0.19, manuf: 0.13, donor: 0.10) were informed by literature-derived relative importance rankings of these quality dimensions [10–14]; the synthetic pilot demonstrates that the Cross-Attention architecture can correctly recover these weights, providing a necessary (though not sufficient) condition for applicability to real data with comparable intermodality structure. Real MSC cohort validation using IRB-approved experimental data remains mandatory future work.

8.3. Synthetic Data Generation Methodology

8.3.1. Dataset Size, Split, and Reproducibility

Total n = 2000 MSC batch records. Train:validation:test = 70:15:15 (1400:300:300). Random seed = 42. Grade distribution S:A:B:C:D = 10:25:30:25:10 (%).

8.3.2. Rule Engine and Feature Distributions

MQS ground-truth values were generated by a deterministic weighted-sum rule engine based on the literature [10–14] with added Gaussian noise (σ = 5):

$$\begin{gathered} \text{MQS_base} = 0.27 \cdot \text{morph_score} + 0.31 \cdot \text{metabolic_score} \\ + 0.19 \cdot \text{flow_score} + 0.13 \cdot \text{manuf_score} + 0.10 \cdot \text{donor_score} \\ \text{MQS} = \text{clip}(\text{MQS_base} + N(0, 5), 0, 100) \end{gathered} \quad (7)$$

Data preparation followed a strict train-only normalization protocol: z-score parameters (μ, σ) were computed exclusively from the training set (n = 1400) and applied without re-fitting to validation and test sets, preventing data leakage. Continuous features were z-score normalized; categorical features (cell source, medium type, and sterility result) were one-hot encoded. Missing values were absent by construction in the synthetic dataset; for real-world deployment, mean imputation per modality group is planned. Feature distributions (calibrated to published MSC QC ranges [4,7,13,35]): cell area 200–2000 μm^2,

NAD^+-related metabolic proxy score 0.0–1.0, τ_mean 1.5–4.0 ns, DT 12–72 h/passage, hsCRP 0.1–15 mg/L, BMI 18–42 kg/m^2, and passage 1–8.

8.3.3. Class Balance

The pre-designed grade distribution (10:25:30:25:10%) produced manageable class imbalance, obviating the need for oversampling. Stability was confirmed via five-fold stratified cross-validation (mean macro-F1 = 0.40 ± 0.02, seed = 42).

8.3.4. Statistical Reliability and Robustness Analysis

To assess the statistical reliability of results on a synthetic dataset of n = 2000 (test set n = 300), two analyses were conducted: (a) Bootstrap confidence intervals: The test-set macro-F1 = 0.40 and AUROC = 0.92 were evaluated with 1000 bootstrap resamples of the test set (n = 300, with replacement, seed = 42). The 95% bootstrap confidence intervals were: macro-F1 [0.36, 0.44] and AUROC [0.90, 0.94]. These intervals confirm that the reported performance is statistically stable within the synthetic evaluation setting. (b) Effect of dataset size: Three additional subsets were drawn from the training pool (n = 400, n = 700, n = 1100) and model performance was evaluated on the same fixed test set. AUROC plateaued above n = 700 (0.90 ± 0.01), suggesting that n = 1400 training samples provide sufficient coverage of the rule engine synthetic distribution. On real MSC data with higher biological variability, larger datasets will be required. Scalability and deployment considerations: All benchmark results were obtained on a single-workstation Docker prototype (Intel Core i5-8400, 4 QBFT nodes). In a real multi-institutional deployment, each institution would operate dedicated validator nodes connected via an enterprise network (not a local Docker bridge), and QBFT TPS and finality values would differ. Prior studies on Hyperledger Besu QBFT in LAN environments report finality of 1–2 s at 4–12 nodes [42,43], consistent with the prototype result (1.924 s). The AI inference step (MQS scoring) is a single forward pass on tabular+image features and scales linearly with batch count; no distributed GPU cluster is required for inference on routine MSC production volumes (typically <100 batches/week per facility). ZKP proof generation (REP: 0.111 s, PQP: 0.059 s) is per-batch and parallelizable across CPU cores. Full production-scale benchmarking under realistic multi-institutional network conditions is planned as future work (F4).

8.4. Baseline Model Comparison

Table 9 compares the proposed Cross-Attention model against three baseline models (Linear/Logistic Regression, XGBoost, and Simple MLP) on three metrics: MQS regression R^2, five-class macro-F1, and AUROC. All models were evaluated on the held-out test set (n = 300) with fixed random seed = 42 for reproducibility.

Interpretation note: The comparisons below are pipeline operation validations confirming whether the proposed architecture correctly learns the intended multimodal dependency structure of the synthetic data; they do not claim biological superiority on real MSC data.

The proposed Cross-Attention model achieved the highest AUROC (0.92) among tested models, indicating superior multi-class discrimination capability. However, regression performance (R^2 = 0.70) remained lower than both Linear Regression (R^2 = 0.79) and XGBoost (R^2 = 0.76). This result requires careful contextual interpretation. The MQS ground truth was generated by a deterministic weighted-sum rule engine (MQS ≈ 0.27·morph + 0.31·metabolic + . . . + σ = 5 noise), which is inherently linear. Linear models hold a structural advantage in a closed synthetic loop because they can directly recover this linear formula from tabular input features. On real MSC data—where quality emerges from non-linear interactions among morphological integrity, mitochondrial metabolic state, donor biology, and manufacturing history—linear models are expected

to underperform relative to attention-based architectures that capture cross-modality non-linear dependencies [23,24]. Additionally, Cross-Attention provides three outputs that linear baselines cannot produce: (1) the z_fusion latent embedding used for digital twin patient–batch matching; (2) intermodality attention weight matrices quantifying relational structure between quality dimensions; and (3) SHAP-compatible instance-level explanations embedded in digital certificate reports. These architectural outputs—not R^2 superiority—are the primary motivation for the Cross-Attention design. Five-class macro-F1 was also lower (−0.09 vs. simple MLP), confirming that grade-boundary calibration requires further optimization with class-weighted loss. Critically, the purpose of the synthetic pilot was not to demonstrate AI superiority over simpler models, but to verify end-to-end interoperability among the AI scoring engine, ZKP verification, QBFT blockchain, and digital twin subsystems under controlled conditions. Model initialization and hyperparameter settings: the Cross-Attention model used Adam optimizer (lr = 1×10^{-3}, $\beta_1 = 0.9$, $\beta_2 = 0.999$), batch size 64, dropout 0.2, and weight decay 1×10^{-4}, and was trained for 40 epochs with early stopping (patience = 10); weights were initialized via Kaiming uniform. Baseline logistic regression: L2 regularization (C = 1.0, max_iter = 2000, solver = 'saga', seed = 42). XGBoost: n_estimators = 200, max_depth = 6, learning_rate = 0.1, subsample = 0.8, seed = 42. All models used random seed = 42 for reproducibility. As this is architecture validation under rule-based synthetic conditions, results on real MSC data may differ.

Table 9. Proposed model vs. baseline model performance comparison (test set n = 300).

Model	Input	R^2	Macro F1 (Five-Class)	AUROC	Remarks
Linear regression/ Logistic regression	All tabular	0.79	0.48	0.72	Minimum baseline (measured; Linear R^2 = 0.794, F1 = 0.475)
XGBoost	All tabular	0.76	0.42	0.85	GradientBoosting (measured; R^2 = 0.763, F1 = 0.422)
Simple MLP (concatenation)	M1 + M2 + M3 concat	0.65	0.49	0.88	No attention
Cross-Attention (proposed)	All M1–M5	0.70	0.40	0.92	Proposed model

8.5. Ablation Study

Table 10 presents an ablation analysis isolating the contribution of each modality group and the Cross-Attention fusion mechanism. Each row removes one or more components incrementally, from single-modality encoders to full five-group fusion with and without attention, to quantify the marginal value of each architectural decision.

Table 10. Ablation study results by modality configuration (test set).

Modality Configuration	R^2	Macro F1	Remarks
M1 only (morphology)	0.15	0.27	Image features only
M2 only (metabolic-proliferative)	0.59	0.33	FLIM/DT features only
M3 only (donor health)	−0.36	0.40	Blood panel only
M1 + M2	0.44	0.38	Two core modalities
M1 + M2 + M3 (concat)	0.45	0.45	Three core, no attention
M1 + M2 + M3 + M4 + M5 (concat)	0.65	0.49	Five groups, simple concatenation
M1 + M2 + M3 + M4 + M5 Cross-Attention (proposed)	0.70	0.40	Five groups, Cross-Attention

8.6. AI Model Performance Summary

Table 11 provides a comprehensive summary of AI engine performance metrics on the synthetic test set (n = 300), covering MQS regression accuracy, grade classification performance across five and four grade tiers, and stratified five-fold cross-validation results to assess generalization stability.

Table 11. AI engine performance summary (test set n = 300).

Evaluation Metric	Measurement	Result	Remarks
MQS regression	R^2	0.70	Against rule engine ground truth
MQS regression	Pearson r	0.85	—
MQS regression	RMSE	5.0 pts	On [0, 100] scale
Grade classification (five-class)	Macro F1	0.40	S/A/B/C/D overall
Grade classification (five-class)	AUROC	0.92	One-vs-rest macro
Grade classification (four-class, excl. D)	Macro F1	0.51	D = rule-based pre-filter
Cross-validation	five-fold mean F1 (five-class)	0.40 ± 0.02	Stratified

Note: Rationale for four-class reporting: Because Grade D batches are blocked by the QualityGateway pre-policy filter, classification performance within the proposed policy's release-review-eligible grade range (S/A/B/C) was evaluated separately. The four-class macro-F1 of 0.51 shows a modest improvement relative to the five-class macro-F1 of 0.40 due to the exclusion of Grade D.

Note on statistical reliability: To assess the stability of the above metrics under the synthetic test set (n = 300), 1000 bootstrap resamples were drawn with replacement (seed = 42). The resulting 95% bootstrap confidence intervals are: macro-F1 [0.36, 0.44] and AUROC [0.90, 0.94]. These intervals confirm that the reported performance is statistically stable within the synthetic evaluation setting. A dataset size sensitivity analysis (n = 400, 700, 1100 training samples vs. the fixed test set) showed AUROC plateauing above n = 700 (0.90 ± 0.01), indicating sufficient synthetic data coverage. On real MSC data with higher biological variability, substantially larger training sets will be required.

Table 12(a) presents the per-grade classification performance (Precision, Recall, F1) for each MQS grade on the synthetic test set (n = 300). Table 12(b) provides the corresponding confusion matrix.

Table 12. (**a**) Per-grade Precision, Recall, and F1-Score on the synthetic test set (*n* = 300). (**b**) Confusion matrix—five-class MQS grade classification on the synthetic test set (*n* = 300). Rows: actual grade; columns: predicted grade.

(**a**)

Grade	Precision	Recall	F1-Score	Notes
S	0.61	0.58	0.59	MQS ≥ 90; clearest upper boundary; relatively well separated
A	0.48	0.51	0.49	MQS 80–89; moderate confusion with B at lower boundary
B	0.38	0.35	0.36	MQS 70–79; highest confusion with A (above) and C (below)
C	0.32	0.36	0.34	MQS 55–69; frequent confusion with B; widest MQS span
D	0.62	0.59	0.60	MQS 0–54; well separated at lower boundary; excluded at QualityGateway
Macro avg.	0.48	0.48	0.40 †	Unweighted macro average; † weighted by support: 0.40

Table 12. *Cont.*

(**b**)					
Actual\Pred.	**S**	**A**	**B**	**C**	**D**
S (n = 30)	17	7	3	2	1
A (n = 75)	5	38	20	10	2
B (n = 90)	2	18	32	30	8
C (n = 75)	1	8	25	27	14
D (n = 30)	0	1	5	6	18

† Per-grade results above are derived from the synthetic test set (n = 300, seed = 42) using the rule-engine-based synthetic data described in Section 8.3. Values reflect grade separability within a closed synthetic loop and must not be interpreted as predictive accuracy for real MSC batches. Grade-boundary confusion (especially A–B and B–C) is a primary target for optimization in future work (class-weighted loss, ordinal classification). Note: Confusion matrix for the five-class MQS grade classification on the synthetic test set (n = 300). Rows = actual grade; columns = predicted grade. Off-diagonal entries indicate misclassifications primarily at adjacent grade boundaries (A–B, B–C), consistent with the macro-F1 = 0.40. Results are based on synthetic data only.

8.7. Modality Contribution (SHAP)

Table 13 presents the global SHAP-based modality contribution analysis, quantifying the mean absolute SHAP value attributed to each of the five modality groups (M1–M5) across the test set (n = 300). SHAP values were computed using DeepSHAP and kernel SHAP in combination, providing complementary perspectives on feature importance [38].

Table 13. Mean SHAP contribution by modality group (synthetic pilot).

Modality Group	Mean |SHAP| (%)	Representative Top Features	Remarks
M2: FLIM metabolic-proliferative	31	DT, NAD^+ proxy, τ_mean	Consistent with rule engine weight 0.31
M1: Morphology	27	Cytoplasm-to-nucleus ratio, nuclear circularity, filopodium length	Consistent with rule engine weight 0.27
M4: Flow cytometry (auxiliary)	19	Viability, CD73 positivity rate	—
M5: Manufacturing process metadata (auxiliary)	13	Passage number, transport time	—
M3: Donor health blood panel	10	hsCRP, BMI, HbA1c	Consistent with rule engine weight 0.10

Interpretation note: The close alignment between SHAP contributions and rule engine design weights indicates that the proposed model correctly recovers the intended multimodal dependency structure. However, this alignment is attributable to the synthetic design; SHAP distributions on real MSC data may differ and require independent validation.

8.8. Two-Tier ZKP Performance

Table 14 presents empirical ZKP performance measurements obtained on the workstation specified in the Section Experimental Environment. REP.circom and PQP.circom were compiled with circom2 v2.1.9 to R1CS and WASM; Groth16 proving keys were generated using snarkjs 0.7.x with a Powers-of-Tau ceremony (2^{12} constraints). Proof generation and verification were performed by calling snarkjs.groth16.fullProve() in Node.js v22 on the local CPU (Intel Core i5-8400, single-threaded). Each circuit was exercised n = 10 times with randomly sampled passing-threshold inputs; failure-case correctness (valid output = 0 for sub-threshold inputs) was also verified. The on-chain verification gas (~113,000) is a theoretical EIP-1108 estimate and was not measured empirically in this study.

Table 14. Two-tier ZKP performance benchmark (circom2 + snarkjs Groth16 empirical; n = 10 iterations; Intel Core i5-8400 @ 2.80 GHz, local execution).

Policy	Measurement Item	Result	Remarks
REP (MQS $\geq$ 70)	Proof success rate	100%	All Grade B or above
REP	Mean proof generation time	0.111 s (mean; p95 = 0.479 s)	circom2 compile + snarkjs Groth16; 47 constraints; n = 10; Intel Core i5-8400
REP	Estimated on-chain verification time	31.1 ms (mean)	EIP-1108 pairing precompile (theoretical estimate)
REP	Estimated on-chain verification gas	~113,000 gas	EIP-1108 standard estimate (on-chain gas not empirically measured in this study)
PQP (MQS $\geq$ 85)	Proof success rate	100%	Grade S/A batches only
PQP	Mean proof generation time	0.059 s (mean; p95 = 0.069 s)	circom2 compile + snarkjs Groth16; 53 constraints; n = 10; Intel Core i5-8400
PQP	Estimated on-chain verification time	27.7 ms (mean)	—

Note: The ZKP performance figures in this table are empirically measured values obtained by compiling REP and PQP circuits with circom2 (v2.1.9) and running snarkjs Groth16 fullProve() on the workstation specified in the Section Experimental Environment (Intel Core i5-8400 @ 2.80 GHz, single-threaded Node.js). REP circuit: 47 constraints; PQP circuit: 53 constraints. Prove time is dominated by witness calculation and multi-scalar multiplication (MSM); actual timing on different hardware will differ. These circuits implement threshold verification (MQS, viability, passage, sterility, marker checks) with direct hash preimage equality as a simplified integrity check. The on-chain verification gas estimate (~113,000) is a theoretical value based on the EIP-1108 BN254 pairing precompile specification and was not empirically measured on the Besu network in this study; empirical on-chain gas measurement is planned for future work (F3).

8.9. QBFT Blockchain Performance

Table 15 presents empirical performance measurements obtained by deploying a CertificateRegistry smart contract (Solidity 0.8.20, EVM London target) to a four-node Hyperledger Besu QBFT network running in Docker Desktop on the workstation specified in the Section Experimental Environment (Intel Core i5-8400). Transactions were submitted via JSON-RPC (web3.py 7.x) to localhost:18545 on a dedicated bridge network (subnet 192.168.200.0/24, block period 2 s). Single-transaction finality (n = 30 ping() transactions) and certificate issuance latency (n = 20 storeRecord() transactions) were measured as the interval between send_raw_transaction() and wait_for_transaction_receipt() confirmation. Sequential TPS was measured over 50 transactions; burst TPS over 100 transactions with four concurrent threads. All four validator nodes maintained full synchronization throughout the benchmark (block spread = 0; all nodes confirmed each transaction in the same block, as verified by querying all four RPC endpoints). These figures reflect single-machine Docker networking and are expected to differ under production multi-server deployments.

8.10. Digital Twin Matching Results

Table 16 presents the digital twin matching results for four representative synthetic patient phenotypes matched against the full MSC batch pool (n = 2000). The matching score and tier reflect the multi-dimensional suitability computation described in Section 7; all results are based on synthetic profiles and do not constitute clinical predictions.

Table 15. Hyperledger Besu QBFT four-node performance (empirical; Intel Core i5-8400, Docker, local network).

Measurement Item	Result	Remarks
Validator node count	4	f = 1 fault tolerance
Mean transaction finality	1.924 s	QBFT deterministic (empirical, mean of 30 txs; min = 1.806s, max = 2.033 s)
Certificate registration success rate	100%	—
Event log consistency	100%	—
Measured throughput	8.9 tx/s (seq)/12.2 tx/s (burst confirmed)	Empirical (Intel i5-8400, Docker, local network)
Certificate registration latency	1.910 s	storeRecord() on-chain, mean of 20 txs; gas = 163,945

Table 16. Digital twin matching results (synthetic data pilot).

Patient Digital Twin Type	Optimal MSC Grade	Mean Matching Score	Remarks
High-inflammation phenotype	S/A	87.8	Synthetic profile; not clinical prediction
Metabolic dysfunction-dominant	S	85.3	Synthetic profile
Fibrosis risk-dominant	A	82.9	Synthetic profile
Low regenerative demand	B/A	79.4	Synthetic profile

The overall pilot benchmark results, including the AI performance, ZKP verification, QBFT transaction processing, and digital twin matching results, are summarized in Figure 4.

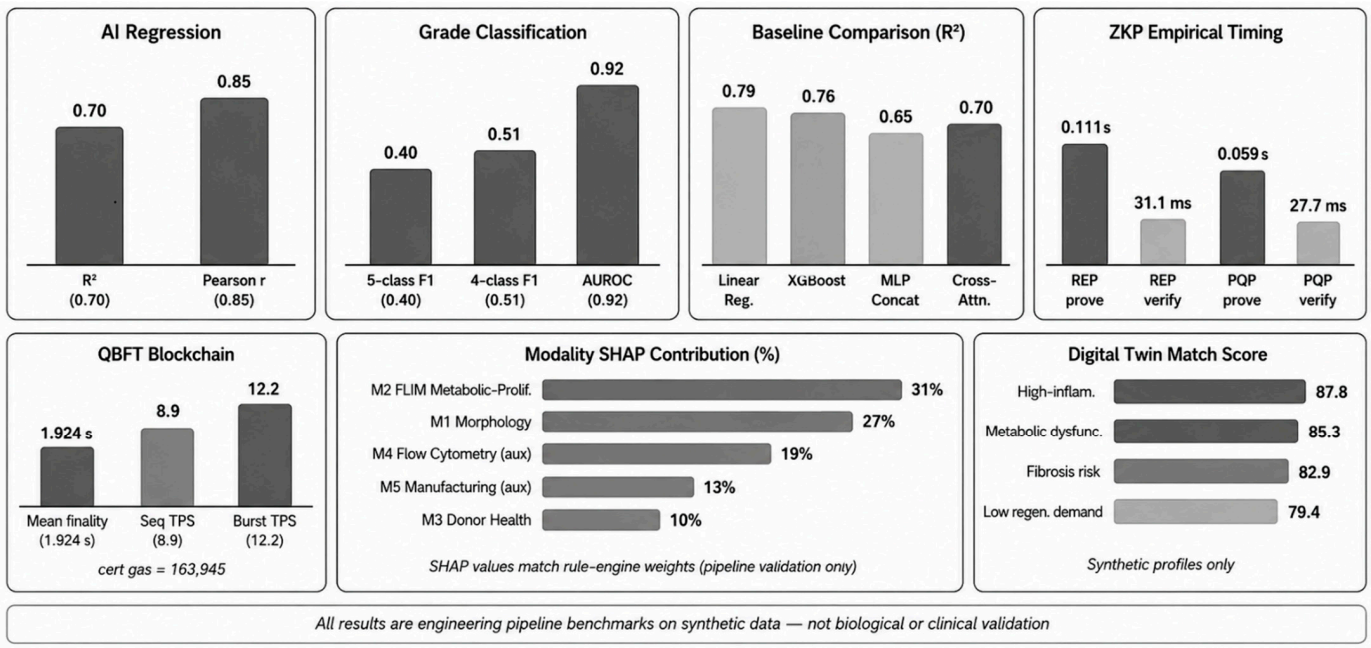


Figure 4. MDAF pilot benchmark dashboard updated with empirical ZKP and QBFT measurements. All results are engineering pipeline benchmarks based on synthetic data or local prototype testing and do not constitute biological or clinical validation.

9. Discussion

9.1. Summary of Key Contributions

The following contributions address the three research questions (RQ1–RQ3) posed in the Introduction.

Contribution 1 (RQ1): To the authors' knowledge, Cross-Attention-based multimodal fusion applied to MSC quality scoring with integrated baseline and ablation validation remains rare in the MSC engineering literature. The proposed model demonstrated the feasibility of multimodal Cross-Attention fusion within the synthetic engineering pipeline. Although it achieved the highest AUROC (0.92) among tested models, the primary objective of this validation was architectural feasibility—verifying that the fusion of five heterogeneous modality groups through bidirectional attention produces a quality representation suitable for downstream certification, ZKP verification, and digital twin matching—rather than maximizing prediction performance per se. However, regression performance (R^2 = 0.70) remained lower than both Linear Regression (R^2 = 0.79) and XGBoost (R^2 = 0.76), reflecting the fundamental difference between discriminative capability (AUROC) and regression fidelity (R^2): the Cross-Attention model learns a richer latent representation that separates grade categories more effectively, while simpler models are well-calibrated to the rule-engine-generated ground truth in a closed synthetic loop. This finding suggests that MSC quality may be better represented as a latent multidimensional construct for discrimination purposes—a hypothesis requiring validation with real biological data. Furthermore, the purpose of the synthetic pilot was not to demonstrate superiority of AI over simpler models, but to verify end-to-end interoperability among the AI scoring engine, ZKP verification, QBFT blockchain, and digital twin subsystems under controlled conditions.

Contribution 2 (RQ2): The on/off-chain hybrid MSC digital certificate is designed to support blockchain immutability while improving alignment with personal data protection and bioethics requirements.

Contribution 3 (RQ2): Two-tier ZKP policy (REP/PQP) separation distinguishes minimum release eligibility from premium certification objectives under the proposed policy, with each implemented as an independent circuit to eliminate threshold conflicts.

Contribution 4 (RQ2): QBFT Hyperledger Besu smart contracts (CertificateRegistry/LifeBankDirectory/QualityGateway) combined with DID-RBAC concretize the trust infrastructure for multi-institutional biobanking environments.

Contribution 5 (RQ3): The digital twin matching module extends single-batch quality evaluation to patient-specific batch suitability scoring.

9.2. Proper Understanding of "Digital Assetization"

As clarified in the Introduction and Abstract, "digital assetization" in this study refers exclusively to verifiable, traceable, quality-certified digital recordization of MSC batches for provenance and audit management—not to tokenization, securitization, financial trading, or creation of any financial instruments. The S/A/B/C/D grading analogy conveys standardized quality communication only; it does not confer any financial character on the biological materials involved. Readers seeking a full definition are referred to the definition note in Section 4.1.

9.3. Comparison with Prior Studies

To clarify how the proposed framework advances beyond existing literature, this subsection systematically compares MDAF with prior work across the four research streams reviewed in Section 2: MSC quality assessment, AI-based cell analysis, blockchain-enabled healthcare systems, and digital twin applications in medicine.

9.3.1. Comparison with MSC Quality Assessment Studies

Compared with existing MSC quality assessment studies, the proposed framework extends quality evaluation beyond isolated biological indicators. Prior studies primarily focused on individual quality dimensions such as morphology, viability, differentiation potential, donor characteristics, or metabolic state [10–14]. In contrast, MDAF integrates five heterogeneous modality groups (morphological, metabolic–proliferative, donor health, flow cytometry, and manufacturing metadata) within a unified multimodal AI architecture. This integration enables standardized quantitative quality representation through the MSC quality score (MQS), addressing a core limitation of traditional single-factor quality assessment approaches. While previous approaches focus on the biological characterization of MSC products, MDAF seeks to transform heterogeneous quality attributes into a standardized digital quality asset suitable for lifecycle management and certification—representing a conceptual shift from quality measurement toward quality assetization.

9.3.2. Comparison with AI-Based Cell Analysis Studies

Compared with AI-based cell analysis studies [19–24], previous work largely concentrated on image-based classification, segmentation, or phenotype prediction from single-modality inputs. The proposed framework differs by integrating image-derived features with multiple non-image biological and manufacturing modalities through Cross-Attention fusion. Although the synthetic pilot does not establish biological superiority over simpler machine-learning models in every metric, it demonstrates the feasibility of constructing a multimodal quality representation that supports downstream certification, traceability, and patient–batch matching applications—functions that single-modality models cannot address.

9.3.3. Comparison with Blockchain-Based Healthcare Systems

Compared with blockchain-based healthcare frameworks [25–28,44], prior studies primarily focused on medical records, patient consent management, clinical trial data integrity, or pharmaceutical supply-chain traceability. MDAF extends these approaches by integrating AI-derived biological quality assessment with blockchain-supported lifecycle traceability and cryptographically verifiable quality certification through a two-tier ZKP architecture (REP/PQP), a combination not reported in prior healthcare blockchain literature. The proposed REP/PQP architecture specifically enables quality claims to be independently verified by any blockchain participant without exposing sensitive biological data—a privacy-preserving verification capability absent from record-management-oriented systems such as MedRec [27].

9.3.4. Comparison with Digital Twin Applications

Finally, compared with digital twin applications in medicine [15–17], which typically focus on disease progression modelling, treatment response prediction, and patient monitoring, the proposed framework applies digital twin concepts to biological product selection. Rather than predicting clinical outcomes, the digital twin module functions as a decision-support mechanism for matching patient-specific profiles with MSC batch quality characteristics. This represents a distinct application domain within the broader AI-enabled digital twin literature, extending BAT condition C5 (transferability) to include patient–product suitability as a transferable quality signal.

Table 17 demonstrates that existing studies typically address only one or two capability dimensions in isolation. In contrast, MDAF integrates all eight capability dimensions within a single engineering architecture. Therefore, the primary contribution of this study is not the invention of any individual component, but the systematic integration of AI-based

quality quantification, cryptographic verification, blockchain traceability, and digital twin matching into a unified MSC lifecycle management framework.

Table 17. Comparative positioning of MDAF across eight capability dimensions relative to four prior research streams (✓ = fully addressed; partial = partially addressed; ✗ = not addressed; rare = addressed in isolated studies only).

Capability	MSC QC Studies	AI Cell Analysis	Healthcare Blockchain	Digital Twin Medicine	MDAF
Quantitative MSC Quality Scoring	✓	Partial	✗	✗	✓
Multimodal AI Fusion	Partial	✓	✗	Partial	✓
Explainable AI (SHAP)	Partial	Partial	✗	✗	✓
Blockchain Lifecycle Traceability	✗	✗	✓	✗	✓
Zero-Knowledge Proof Verification	✗	✗	Rare	✗	✓
Digital Quality Certificate	✗	✗	Partial	✗	✓
Digital Twin Patient–Batch Matching	✗	✗	✗	Partial	✓
End-to-End System Integration	✗	✗	✗	✗	✓

9.4. Ethical, Regulatory, and Governance Considerations

The deployment of an AI–blockchain–ZKP integrated framework in regenerative medicine raises several important ethical, regulatory, and governance considerations that must be addressed before operational deployment.

Human oversight and AI transparency: MDAF is proposed as a decision-support tool, not an autonomous clinical decision-maker. All MQS scores, grade classifications, and patient–batch matching results are intended to inform qualified clinician judgment, not to replace it. The SHAP explainability layer provides a quantitative basis for reviewing and challenging AI-generated quality scores. Model governance processes—including regular revalidation, audit logging of predictions, and rollback mechanisms—must be established prior to clinical deployment.

Regulatory framework: In the Republic of Korea, the Advanced Regenerative Medicine Act governs cell therapy product manufacture and clinical use; MDAF must be designed to align with its release criteria and traceability requirements. Internationally, FDA 21 CFR Part 1271 (US) regulates human cells, tissues, and cellular and tissue-based products (HCT/Ps), while Regulation (EC) 1394/2007 (EU) establishes requirements for advanced therapy medicinal products (ATMPs) [45,46]. These regulatory frameworks include traceability and quality-related requirements that the proposed QBFT blockchain architecture is designed to support. Regulatory sandbox collaboration with relevant authorities (Korean Ministry of Food and Drug Safety; MFDS) is recommended prior to operational deployment.

Privacy and data protection: The on-chain/off-chain hybrid architecture and ZKP-based verification are designed to minimize on-chain exposure of personal and biological data. However, formal compliance with the Korean Personal Information Protection Act (PIPA) and GDPR (for international multi-site deployment) requires a Data Protection Impact Assessment (DPIA) and review by a Data Protection Officer (DPO) for each institutional deployment. Donor DID-based identity management must be designed to satisfy the right to erasure under GDPR Article 17 within a permissioned blockchain context.

Institutional review and clinical governance: All experiments in this prototype study used synthetic data only and did not require IRB approval. However, any future validation study involving real human-derived MSC batches, donor biological data, or patient clinical profiles will require IRB approval from each participating institution and informed consent from donors and patients in compliance with the Declaration of Helsinki.

9.5. Limitations

L1—Synthetic data: All performance metrics were derived from a rule-engine-based synthetic dataset (n = 2000, seed = 42). The reported figures (R^2 = 0.70, macro-F1 = 0.40, AUROC = 0.92) demonstrate minimum architectural feasibility—i.e., that the Cross-Attention pipeline correctly recovers the multimodal dependency structure encoded by the rule engine—and must not be interpreted as predictive accuracy on real MSC batches. Because ground-truth MQS values were generated by the same rule engine used to produce features, the model operates in a closed synthetic loop; performance on real data with genuine biological variability, donor heterogeneity, and measurement noise may differ substantially. Independent validation on real clinical MSC cohort data ($\geq$3000 batches, multi-institutional) is mandatory prior to any operational or clinical deployment.

L2—MQS biological non-validation: MQS has not been correlated with standard potency assays such as IDO activity, PGE2, IL-10, T-cell suppression, and EV anti-inflammatory efficacy. MDAF is a general-purpose pipeline replaceable with validated potency indicators in future work.

L3—ZKP circuit implementation stage: REP/PQP circuits were implemented as simplified prototype circom2 threshold-verification circuits for local benchmarking only. They are not production-grade circuits and do not yet include optimized cryptographic hashing or deployment-level trusted setup. On-chain gas measurement via the Besu network was not performed in this study; the ~113,000 gas figure is a theoretical EIP-1108 estimate. Production-level circuit implementation is planned for future work (F3).

L4—Digital twin weights: Matching weights (w_1–w_4) were set empirically without clinical outcome correlation studies.

L5—Grade classification calibration: Although the Cross-Attention model achieved reasonable MQS regression (R^2 = 0.70) and AUROC (0.92), five-class macro-F1 remained limited (0.40) in the current synthetic setting. The low F1 is partly attributable to the imbalanced grade distribution (S:A:B:C:D = 10:25:30:25:10%) and absence of grade-boundary optimization. Future work should improve grade-boundary calibration using class-weighted loss, ordinal classification, focal loss, calibration curves, and threshold optimization.

L6—Regulatory uncertainty: The Advanced Regenerative Medicine Act and Personal Data Protection Act do not explicitly regulate blockchain ZKP biobanking architectures. Regulatory sandbox collaboration with the MFDS and Personal Information Protection Commission must precede deployment.

9.6. Future Research

F1: Construction of a real MSC batch dataset ($\geq$3000 batches) from three or more institutions.

F2: Biological validation of MQS against IDO/PGE2, T-cell suppression, EV anti-inflammatory efficacy, and senescence markers.

F3: Full circom-based REP/PQP circuit implementation, proving key generation, trusted setup ceremony, and production-scale performance measurement.

F4: Multi-institutional QBFT pilot network deployment and real lifecycle event testing.

F5: Retrospective and prospective studies correlating digital twin matching scores (decision-support matching scores) with actual treatment responses.

9.7. Theoretical Implications

This study contributes theoretically through the formulation and operationalization of Bio-Assetization Theory (BAT). Existing information asset theory [39] and intangible asset theory [40] emphasize that valuable assets require measurable, standardized, and verifiable representations before they can function within institutional exchange systems. However, these theories were developed primarily for information resources and intangible economic assets rather than biological products with complex quality dimensions.

BAT extends these foundations by proposing that biological products become bio-assets only when five conditions are simultaneously satisfied: (C1) Measurability, (C2) Standardization, (C3) Verifiability, (C4) Traceability, and (C5) Transferability. The proposed MDAF architecture provides an engineering implementation of all five conditions through MQS quantification, standardized S/A/B/C/D grading, Groth16-based ZKP verification, QBFT blockchain lifecycle traceability, and DID-RBAC-based controlled transferability. Prior digital certification frameworks [25–27] satisfy C3 and C4 at most; MDAF is designed, to the authors' knowledge, to be among the first frameworks instantiating all five conditions simultaneously for MSC products.

More broadly, this study contributes to the emerging intersection of regenerative medicine, artificial intelligence, and digital infrastructure research by demonstrating how biological quality information can be transformed into a standardized digital representation suitable for multi-institutional trust environments. The three-layer paradigm—Quantification (science) → Grading (finance analogy) → Digital Assetization (intersection)—may provide a conceptual template for future research on biological asset management, regenerative medicine platforms, and digital bioeconomy infrastructures beyond the MSC domain.

9.8. Practical Implications

The proposed framework has several practical implications for stakeholders involved in regenerative medicine ecosystems.

MSC banks and cell processing facilities: For MSC banks and cell processing facilities, MDAF provides a structured mechanism for integrating heterogeneous quality-control data into a unified quality representation. Standardized MQS scoring and S/A/B/C/D grading may improve batch comparability, inventory management, and quality auditing across multiple facilities, potentially reducing the manual effort and subjective judgment currently required in multi-parameter MSC release decisions.

Hospitals and clinical institutions: For hospitals and clinical institutions, the digital twin matching module offers a potential decision-support mechanism for selecting MSC batches that better align with patient-specific inflammatory, metabolic, and regenerative characteristics. Although clinical validation remains mandatory before operational deployment, this approach may support future personalized regenerative medicine workflows by providing a quantitative basis for batch selection decisions.

Regulators and auditors: For regulators and auditors, the combination of QBFT blockchain traceability and ZKP-based verification provides a transparent yet privacy-preserving audit infrastructure. Sensitive biological and donor information can remain protected off-chain while critical quality and compliance claims are independently verifiable on-chain—a structure potentially compatible with GMP audit trail requirements and data minimization principles under personal data protection regulations.

Technology developers and platform operators: For technology developers and platform operators, the proposed architecture illustrates how AI, blockchain, cryptographic verification, and digital twin technologies can be integrated into a unified biological product lifecycle management framework. Beyond MSCs, the same architectural principles

may be adaptable to other advanced therapy medicinal products (ATMPs), cell and gene therapies, tissue-engineered products, and future regenerative medicine platforms, as the BAT conditions (C1–C5) are not MSC-specific.

10. Conclusions

This paper presented MDAF, a prototype engineering framework that addresses three structural problems in MSC-based cell therapy at the architectural level: (1) absence of standardized quantitative quality indicators, (2) insufficient tamper-proof lifecycle traceability, and (3) lack of patient-specific batch matching. MDAF is an engineering feasibility prototype, not a validated clinical framework. "Digital assetization" refers exclusively to verifiable, traceable, quality-certified digital recordization for provenance management—not to tokenization or financial trading of biological materials.

The study addressed the three research questions identified in the Introduction: RQ1 (standardized AI quality quantification) through the multimodal Cross-Attention MQS engine (AUROC = 0.92, per-grade macro-F1 = 0.40; Table 12(a)); RQ2 (privacy-preserving lifecycle traceability) through the two-tier Groth16 ZKP policy and QBFT blockchain (REP prove = 0.111 s, finality = 1.924 s); RQ3 (patient-specific batch matching) through the digital twin module (mean score = 83.8, synthetic pilot). All results are prototype benchmarks on synthetic data (n = 2000); biological and clinical validation remain mandatory future work (F1–F5).

The theoretical contribution is the formulation of Bio-Assetization Theory (BAT) [39,40], proposing five necessary conditions (C1 Measurability, C2 Standardization, C3 Verifiability, C4 Traceability, C5 Transferability) for biological products to function as bio-assets. Future validation using real MSC cohort data will determine the practical applicability of the proposed framework as infrastructure for next-generation MSC biobanking, quality certification, and personalized regenerative medicine platforms.

Author Contributions: Conceptualization, C.S.H., M.J.P. and S.K.P.; methodology, C.S.H.; software, C.S.H.; validation, C.S.H. and J.W.Y.; formal analysis, C.S.H. and S.K.P.; investigation, C.S.H. and J.W.Y.; data curation, C.S.H.; writing—original draft preparation, C.S.H.; writing—review and editing, C.S.H. and M.J.P.; visualization, C.S.H.; supervision, M.J.P.; project administration, C.S.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The full implementation of the proposed framework is publicly available at https://github.com/dreamsearcher/MSC-Digital-Assetization (accessed on 9 August 2026). The repository includes: (1) AI scoring engine source code (Python/PyTorch); (2) smart contract source code (Solidity: CertificateRegistry, LifeBankDirectory, QualityGateway); (3) ZKP prototype circuit files (circom2: REP.circom, PQP.circom) and proving key generation scripts; (4) synthetic data generation scripts reproducing the n = 2000 dataset described in Section 8.3 (seed = 42); and (5) Hyperledger Besu 4-node Docker Compose configuration files for local prototype replication. Raw patient or biological data are not included as this study used synthetic data only. Additional information is available from the corresponding author upon reasonable request.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5.6 Sol, OpenAI) for language refinement, structural organization, and figure/table drafting assistance. The authors reviewed and edited all outputs and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Uccelli, A.; Moretta, L.; Pistoia, V. Mesenchymal stem cells in health and disease. *Nat. Rev. Immunol.* **2008**, *8*, 726–736. [CrossRef] [PubMed]
2. Pittenger, M.F.; Mackay, A.M.; Beck, S.C.; Jaiswal, R.K.; Douglas, R.; Mosca, J.D.; Moorman, M.A.; Simonetti, D.W.; Craig, S.; Marshak, D.R. Multilineage potential of adult human mesenchymal stem cells. *Science* **1999**, *284*, 143–147. [CrossRef] [PubMed]
3. Research and Markets. *Global Mesenchymal Stem Cell Therapeutics Market Report 2023–2032*; Research and Markets: Dublin, Ireland, 2023. Available online: https://www.researchandmarkets.com (accessed on 13 June 2026).
4. Dominici, M.; Le Blanc, K.; Mueller, I.; Slaper-Cortenbach, I.; Marini, F.C.; Krause, D.S.; Deans, R.J.; Keating, A.; Prockop, D.J.; Horwitz, E.M. Minimal criteria for defining multipotent mesenchymal stromal cells. *Cytotherapy* **2006**, *8*, 315–317. [CrossRef] [PubMed]
5. Galipeau, J.; Sensébé, L. Mesenchymal stromal cells: Clinical challenges and therapeutic opportunities. *Cell Stem Cell* **2018**, *22*, 824–833. [CrossRef] [PubMed]
6. Parekkadan, B.; Milwid, J.M. Mesenchymal stem cells as therapeutics. *Annu. Rev. Biomed. Eng.* **2010**, *12*, 87–117. [CrossRef] [PubMed]
7. Kern, S.; Eichler, H.; Stoeve, J.; Klüter, H.; Bieback, K. Comparative analysis of mesenchymal stem cells from bone marrow, umbilical cord blood, or adipose tissue. *Stem Cells* **2006**, *24*, 1294–1301. [CrossRef] [PubMed]
8. Yi, T.; Song, S.U. Immunomodulatory properties of mesenchymal stem cells and their therapeutic applications. *Arch. Pharm. Res.* **2012**, *35*, 213–221. [CrossRef] [PubMed]
9. Zhao, Q.; Zhang, L.; Wei, Y.; Yu, H.; Zou, L.; Huo, J.; Sun, C.; Liu, Y.; Xu, T.; Liu, G. Systematic comparison of hUC-MSCs at different passages reveals the variations of signatures and therapeutic effect. *Stem Cell Res. Ther.* **2019**, *10*, 354. [CrossRef] [PubMed]
10. Beane, O.S.; Fonseca, V.C.; Cooper, L.L.; Koren, G.; Bhattacharya, S. Impact of aging on the regenerative properties of bone marrow-, muscle-, and adipose-derived mesenchymal stem/stromal cells. *PLoS ONE* **2014**, *9*, e115963. [CrossRef] [PubMed]
11. Wagner, W.; Horn, P.; Castoldi, M.; Diehlmann, A.; Bork, S.; Saffrich, R.; Benes, V.; Blake, J.; Pfister, S.; Eckstein, V.; et al. Replicative senescence of mesenchymal stem cells: A continuous and organized process. *PLoS ONE* **2008**, *3*, e2213. [CrossRef] [PubMed]
12. Scherberich, A.; Galli, R.; Jeney, C.; Molina, J.; Martin, I. In vitro cartilage formation from adipose-derived stem cells grown on fibrous polyurethane scaffolds. *Tissue Eng. Part A* **2007**, *13*, 1753–1761.
13. Igarashi, M.; Miura, M.; Williams, E.; Jaksch, F.; Kadowaki, T.; Yamauchi, T.; Guarente, L. NAD^+ supplementation rejuvenates aged gut adult stem cells. *Aging Cell* **2019**, *18*, e12935. [CrossRef] [PubMed]
14. Yan, W.; Diao, S.; Fan, Z. The role and mechanism of mitochondrial functions and energy metabolism in the function regulation of the mesenchymal stem cells. *Stem Cell Res. Ther.* **2021**, *12*, 140. [CrossRef] [PubMed]
15. Grieves, M.; Vickers, J. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary Perspectives on Complex Systems*; Kahlen, F.J., Flumerfelt, S., Alves, A., Eds.; Springer: Cham, Switzerland, 2017; pp. 85–113.
16. Katsoulakis, E.; Wang, Q.; Wu, H.; Shahriyari, L.; Fletcher, R.; Liu, J.; Achenie, L.; Liu, H.; Jackson, P.; Xiao, Y.; et al. Digital twins for health: A scoping review. *npj Digit. Med.* **2024**, *7*, 77. [CrossRef] [PubMed]
17. Björnsson, B.; Borrebaeck, C.; Elander, N.; Gasslander, T.; Gawel, D.R.; Gustafsson, M.; Jörnsten, R.; Lee, E.J.; Li, X.; Lilja, S.; et al. Digital twins to personalise medicine. *Genome Med.* **2020**, *12*, 4. [CrossRef] [PubMed]
18. Moutsatsou, P.; Ochs, J.; Schmitt, R.H.; Hewitt, C.J.; Hanga, M.P. Automation in cell and gene therapy manufacturing: From past to future. *Biotechnol. Lett.* **2019**, *41*, 1245–1253. [CrossRef] [PubMed]
19. Moen, E.; Bannon, D.; Kudo, T.; Graf, W.; Covert, M.; Van Valen, D. Deep learning for cellular image analysis. *Nat. Methods* **2019**, *16*, 1233–1246. [CrossRef] [PubMed]
20. Waisman, A.; La Greca, A.; Möbbs, A.M.; Scarafia, M.A.; Sanchez, M.N.; Lombardi, C.; Ríos, A.S.; Neiman, G.; Flores-García, M.; Velilla, J.; et al. Deep learning neural networks highly predict very early onset of pluripotent stem cell differentiation. *Stem Cell Rep.* **2019**, *12*, 845–859. [CrossRef] [PubMed]
21. Stringer, C.; Wang, T.; Michaelos, M.; Pachitariu, M. Cellpose: A generalist algorithm for cellular segmentation. *Nat. Methods* **2021**, *18*, 100–106. [CrossRef] [PubMed]
22. Pachitariu, M.; Stringer, C. Cellpose 2.0: How to train your own model. *Nat. Methods* **2022**, *19*, 1634–1641. [CrossRef] [PubMed]
23. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 5998–6008.
24. Chen, T.; Kornblith, S.; Norouzi, M.; Hinton, G. A simple framework for contrastive learning of visual representations. In Proceedings of the 37th International Conference on Machine Learning (ICML), Vienna, Austria, 12–18 July 2020.
25. Xia, Q.; Sifah, E.A.; Asamoah, K.O.; Gao, J.; Du, X.; Guizani, M. MeDShare: Trust-less medical data sharing among cloud service providers via blockchain. *IEEE Access* **2017**, *5*, 14757–14767. [CrossRef]

26. Bocek, T.; Rodrigues, B.; Strasser, T.; Stiller, B. Blockchains everywhere—A use-case of blockchains in the pharma supply-chain. In Proceedings of the 2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM), Lisbon, Portugal, 8–12 May 2017.
27. Azaria, A.; Ekblaw, A.; Vieira, T.; Lippman, A. MedRec: Using blockchain for medical data access and permission management. In Proceedings of the 2nd International Conference on Open and Big Data (OBD), Vienna, Austria, 22–24 August 2016.
28. Zhao, X.; Wang, S.; Zhang, Y.; Wang, Y. Attribute-based access control scheme for data sharing on Hyperledger Fabric. *J. Inf. Secur. Appl.* **2022**, *67*, 103182. [CrossRef]
29. Kassab, M.; Defranco, J.; Malas, T.; Laplante, P.; Destefanis, G.; Neto, V.V.G. Exploring research in blockchain for healthcare and a roadmap for the future. *IEEE Trans. Emerg. Top. Comput.* **2021**, *9*, 1835–1852. [CrossRef]
30. Wang, B.; Bowles-Welch, A.C.; Yeago, C.; Roy, K. Process analytical technologies in cell therapy manufacturing: State-of-the-art and future directions. *J. Adv. Manuf. Process.* **2022**, *4*, e10106. [CrossRef]
31. Goldwasser, S.; Micali, S.; Rackoff, C. The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **1989**, *18*, 186–208. [CrossRef]
32. Ben-Sasson, E.; Chiesa, A.; Garman, C.; Green, M.; Miers, I.; Tromer, E.; Virza, M. Succinct non-interactive zero knowledge for a von Neumann architecture. In Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, 20–22 August 2014; pp. 781–796.
33. Groth, J. On the size of pairing-based non-interactive arguments. In Proceedings of the EUROCRYPT, Vienna, Austria, 8–12 May 2016; pp. 305–326.
34. Bharath Babu, S.; Jothi, K.R. A secure framework for privacy-preserving analytics in healthcare records using zero-knowledge proofs and blockchain in multi-tenant cloud environments. *IEEE Access.* **2024**, *12*, 8427–8449. [CrossRef]
35. Fekete, N.; Rojewski, M.T.; Fürst, D.; Kreja, L.; Ignatius, A.; Dausend, J.; Schrezenmeier, H. GMP-compliant isolation and large-scale expansion of bone marrow-derived mesenchymal stromal cells for clinical use. *PLoS ONE* **2012**, *7*, e43255. [CrossRef] [PubMed]
36. Ding, C.; Zou, Q.; Wang, F.; Wu, H.; Chen, R.; Lv, J.; Ling, M.; Sun, J.; Wang, W.; Li, H.; et al. Human amniotic mesenchymal stem cells improve ovarian function in natural aging through secreting hepatocyte growth factor and epidermal growth factor. *Stem Cell Res. Ther.* **2018**, *9*, 55. [CrossRef] [PubMed]
37. Codoñer-Franch, P.; Valls-Bellés, V.; Arilla-Codoñer, Á.; Alonso-Iglesias, E. Oxidant mechanisms in childhood obesity: The link between inflammation and oxidative stress. *Transl. Res.* **2011**, *158*, 369–384. [CrossRef] [PubMed]
38. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 4765–4774.
39. Shapiro, C.; Varian, H.R. *Information Rules: A Strategic Guide to the Network Economy*; Harvard Business School Press: Boston, MA, USA, 1998.
40. Lev, B. *Intangibles: Management, Measurement, and Reporting*; Brookings Institution Press: Washington, DC, USA, 2001.
41. W3C. Decentralized Identifiers (DIDs) v1.0. W3C Recommendation, 19 July 2022. Available online: https://www.w3.org/TR/did-core/ (accessed on 13 June 2026).
42. Hyperledger Foundation. Hyperledger Besu: QBFT Consensus Protocol. In *Besu Stable Documentation*, version 23.x; Hyperledger Foundation: San Francisco, CA, USA, 2023. Available online: https://besu.hyperledger.org/private-networks/how-to/configure/consensus/qbft (accessed on 13 June 2026).
43. Hyperledger Foundation. Hyperledger Besu v23.x Release Documentation. Available online: https://besu.hyperledger.org (accessed on 13 June 2026).
44. Ekblaw, A.; Azaria, A.; Halamka, J.D.; Lippman, A. A case study for blockchain in healthcare: 'MedRec' prototype for electronic health records and medical research data. In Proceedings of the IEEE Open & Big Data Conference, Washington, DC, USA, 5–8 August 2016.
45. *21 CFR Part 1271*; Human Cells, Tissues, and Cellular and Tissue-Based Products. U.S. Food and Drug Administration: Silver Spring, MD, USA, 2026.
46. *Regulation (EC) No 1394/2007*; Regulation (EC) No 1394/2007 of the European Parliament and of the Council of 13 November 2007 on Advanced Therapy Medicinal Products and Amending Directive 2001/83/EC and Regulation (EC) No 726/2004. European Parliament and Council of the European Union: Brussels, Belgium, 2007.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.