Toward a Lagrangian Framework for Biological Digital Twins

Fluid-Driven Tissue Function and the Case for Variational Methods in Computational Physiology

Wayne Eskridge[1]
Decision Sciences, LLC, Boise, Idaho, USA

July 6, 2026

Abstract

Digital twins of biological organs—computational models that mirror the physiology of an individual patient—are an emerging paradigm in precision medicine. Existing approaches rely on finite-element methods, agent-based rules, pharmacokinetic compartments, or Eulerian computational fluid dynamics, none of which employ a variational principle. We propose that biological tissue function is fundamentally *fluid-driven*: the transport of blood, bile, lymph, and interstitial fluid dictates nutrient delivery, waste clearance, mechanical signalling, and cellular fate. Because the Lagrangian (variational) formulation is the natural language of fluid dynamics—encoding symmetries, conservation laws, and multi-scale coupling in a single scalar functional—it is a natural and potentially superior foundation for organ-scale digital twins.

We develop this argument using the liver as the initial target organ, where sinusoidal blood flow governs metabolic zonation, fibrotic remodelling, and portal hypertension. We outline the general form of a tissue Lagrangian, identify the symmetries and conservation laws it implies, and survey the existing literature to establish that no research group has, to our knowledge, proposed or implemented a Lagrangian-based biological digital twin. This paper serves as a conceptual foundation and research programme for pursuing this approach.

Contents

1 Introduction

The concept of a *digital twin*—a computational replica of a physical system that evolves in parallel with its real counterpart— originated in aerospace engineering [**?**], where high-fidelity models of aircraft and turbines are used for predictive maintenance and design optimisation. The extension to biology and medicine promises patient-specific models that can predict disease progression, optimise drug dosing, and simulate surgical outcomes before they are performed on a living patient.

Yet the biological digital twin faces a challenge that aerospace models do not: *the governing equations are not known from first principles.* An aircraft wing obeys the Navier–Stokes equations and linear elasticity; a liver obeys... what, exactly? The standard answer is a patchwork: mass-action kinetics for metabolism, finite-element mechanics for tissue deformation, compartment models for pharmacokinetics, agent-based rules for cell behaviour. These are stitched together with ad hoc coupling terms, each sub-model using different variables, different numerical schemes, and different assumptions about what is conserved.

This paper proposes a different starting point. We argue that:

1. Biological tissue function is *fundamentally fluid-driven*: the transport of blood, lymph, bile, and interstitial fluid is not merely a supporting service but the primary determinant of cellular behaviour, tissue structure, and organ function.
2. Because fluid dynamics has a natural variational (Lagrangian) formulation, organ-scale digital twins can be constructed from a single scalar functional—a *tissue Lagrangian*—from which all governing equations, conservation laws, and coupling terms emerge automatically via the Euler–Lagrange equations and Noether's theorem [**?**].
3. To our knowledge, this approach has not been attempted. Despite extensive literature on biological modelling, no group has proposed or implemented a Lagrangian-based digital twin of any organ.

We develop this argument using the liver as the initial target, because its architecture is dominated by fluid flow at every scale and because liver disease (cirrhosis, fatty liver disease) is a leading cause of morbidity worldwide.

2 Current Approaches to Biological Digital Twins

A survey of the computational physiology literature reveals four dominant modelling paradigms for organ-scale digital twins. None employs a variational principle.

2.1 Finite-Element Methods (FEM)

FEM discretises the organ into small elements and solves partial differential equations (typically elasticity, diffusion, or Navier–Stokes) on the resulting mesh. Major efforts include the *Virtual Physiological Human* (VPH) initiative in Europe, the *Living Heart Project* (Dassault Systèmes), and various cardiac electromechanics models.

Limitation: FEM solves equations that are written down *before* the simulation begins. The equations are not derived from a unifying principle; they are postulated for each sub-domain (mechanics, electrics, perfusion) and coupled via boundary conditions. Conservation laws must be manually enforced.

2.2 Agent-Based Models (ABM)

ABMs represent cells as discrete agents with rule-based behaviours: "if TGF-β exceeds threshold X, activate and produce collagen at rate Y." This approach has been applied to tumour growth, wound healing, and liver fibrosis.

Limitation: The rules are phenomenological, not derived from any principle. There is no guarantee that the model conserves mass, energy, or momentum. Changing one rule can have unpredictable cascading effects on others, making calibration fragile.

2.3 Pharmacokinetic / Compartment Models

These treat the organ as a set of well-mixed compartments connected by flow rates, described by ordinary differential equations (ODEs). They are the workhorse of drug development and pharmacology.

Limitation: Compartment models have no spatial resolution. They cannot capture the gradients, flow patterns, and structural remodelling that define organ pathology. A compartment model of the liver cannot distinguish periportal from pericentral hepatocytes—a distinction that is central to metabolic zonation and disease.

2.4 Eulerian Computational Fluid Dynamics (CFD)

Several groups have applied Navier–Stokes-based CFD to hepatic blood flow, including studies of portal vein haemodynamics and sinusoidal perfusion. These models solve the Eulerian fluid equations directly, often coupled to passive scalar transport for nutrients or drugs.

Limitation: Eulerian CFD solves the equations of motion but does not derive them from a variational principle. As discussed in the companion paper on fluid Lagrangians, the Eulerian Navier–Stokes equations *can* be obtained from a Lagrangian, but standard CFD practice does not exploit this structure. Conservation laws are enforced numerically rather than structurally.

2.5 Summary

Approach	Spatial	Multi-scale	Conservation	Variational	Coupling
FEM	✓	Partial	Manual	No	Ad hoc
ABM	✓	✓	None	No	Rules
Compartment	No	No	Mass	No	ODEs
Eulerian CFD	✓	No	Numerical	No	BCs
Lagrangian (proposed)	✓	✓	**Automatic**	**Yes**	**Structural**

3 The Fluid Nature of Biological Tissue

The central thesis of this paper is that biological tissue function is *fluid-driven* in a deep, not merely supportive, sense. We argue this at three levels.

3.1 Nutrient Transport and Metabolic Zonation

Every cell in the body depends on convective and diffusive transport of oxygen, glucose, hormones, and signalling molecules delivered by blood flow. In the liver, this dependence is extreme: hepatocytes are arranged in lobules around a central vein, and their metabolic programme depends on their position along the *porto-central axis*—a gradient established entirely by sinusoidal blood flow.

Periportal hepatocytes (near the portal triad, receiving oxygen-rich blood) specialise in oxidative metabolism, gluconeogenesis, and urea synthesis. Pericentral hepatocytes (near the central vein, receiving oxygen-depleted blood) specialise in glycolysis, lipogenesis, and drug metabolism via cytochrome P450 enzymes. This *metabolic zonation* is not genetically hard-wired; it is dynamically maintained by the oxygen and morphogen gradients established by flow. Reverse the flow direction experimentally, and the zonation pattern reverses with it.

3.2 Mechanotransduction and Structural Remodelling

Fluid flow exerts shear stress on the cells lining blood vessels and sinusoids. These endothelial cells are not passive pipes; they are mechanosensors that transduce shear stress into biochemical signals:

- **Liver sinusoidal endothelial cells (LSECs)** maintain fenestrations (nanopores) that allow plasma to filter through to hepatocytes. Fenestration is regulated by shear stress and nitric oxide, both flow-dependent quantities.
- **Hepatic stellate cells (HSCs)** reside in the space of Disse and respond to flow-mediated signals (TGF-β, PDGF) by either remaining quiescent or activating into myofibroblasts that produce collagen—the key driver of fibrosis [**?**].
- **Portal pressure** (measured clinically as the hepatic venous pressure gradient, HVPG) is a direct fluid-mechanical quantity: the pressure drop across the hepatic vascular bed. It is the single most important prognostic marker in chronic liver disease.

3.3 Disease as Disrupted Flow

In this view, liver disease is fundamentally a *fluid dynamics pathology*:

- **Fibrosis**: Collagen deposition increases flow resistance → portal pressure rises → sinusoidal shear stress changes → LSEC defenestration and HSC activation → more collagen. This is a positive feedback loop mediated entirely by fluid mechanics.
- **Cirrhosis**: Advanced fibrosis restructures the vascular architecture, creating regenerative nodules with abnormal flow → metabolic dysfunction, ascites, variceal bleeding.
- **Steatosis (fatty liver)**: Fat accumulation in hepatocytes compresses sinusoids → reduced flow → hypoxia → inflammation → further fat accumulation and eventual fibrosis.

In each case, the disease trajectory is governed by the interplay between fluid flow, mechanical forces, and cellular response—exactly the kind of multi-physics coupling that a Lagrangian formulation handles naturally.

4 Why the Lagrangian Formalism Is Appropriate

Given that tissue function is fluid-driven, and that fluid dynamics has a natural variational formulation (as developed in the companion paper), the application of Lagrangian methods to biological tissue becomes not just possible but arguably *natural*. The specific advantages are:

4.1 Automatic Conservation Laws

A Lagrangian formulation guarantees, via Noether's theorem [**?**], that the model automatically conserves the quantities mandated by its symmetries. For a tissue model, these include:

Symmetry	Conservation Law	Biological Meaning
Time translation	Energy	Metabolic energy balance
Space translation	Momentum	Force balance in tissue
Particle relabeling	Circulation	Vascular flow topology
Mass-phase symmetry	Mass	Cell & fluid mass conservation

In current ABM and FEM approaches, these conservation laws must be enforced explicitly and may be violated at insufficient resolution or by modelling inconsistency.

4.2 Structural Coupling

In a Lagrangian formulation, different physical processes are not coupled via ad hoc terms; they are coupled *structurally* through the action principle. Adding a new physics (e.g., bile transport, immune cell migration) means adding terms to the Lagrangian, and all cross-couplings emerge automatically from the variational principle. This eliminates the "integration tax" that plagues multi-physics models.

4.3 Multi-Scale Consistency

The Lagrangian formalism scales naturally from microscopic to macroscopic. A coarse-grained Lagrangian can be derived from a fine-grained one by integrating out fast degrees of freedom—a procedure well-established in condensed matter physics (the Wilsonian renormalisation group) [**?**]. This provides a principled way to connect molecular-scale processes (receptor binding, enzyme kinetics) to tissue-scale observables (portal pressure, fibrosis score) without ad hoc scale-bridging assumptions.

4.4 Natural Treatment of Dissipation

Biological systems are dissipative: energy is continuously consumed and entropy produced. The Lagrangian framework handles dissipation through the Rayleigh dissipation function or, more fundamentally, through the inclusion of irreversible entropy production terms. The resulting equations are thermodynamically consistent by construction—a property difficult to guarantee in ad hoc models.

4.5 Approximation Preserves Structure

Following Salmon's principle [?] (discussed in the companion paper), approximating the Lagrangian *before* deriving equations of motion ensures that the approximate model inherits correct (approximate) conservation laws. This is critical for long-time simulations of disease progression, where accumulated conservation violations in conventional models can produce unphysical drift.

5 General Form of a Tissue Lagrangian

We now sketch the general form of a Lagrangian density for a fluid-perfused biological tissue. This is presented as a research framework, not a final model; the specific functional forms must be determined by calibration against experimental data.

5.1 Degrees of Freedom

The dynamical fields of a perfused tissue include:

- $\mathbf{v}(\mathbf{x}, t)$: blood (sinusoidal) velocity field
- $\rho(\mathbf{x}, t)$: blood density (effectively incompressible)
- $c_i(\mathbf{x}, t)$: concentrations of key solutes (oxygen, glucose, TGF-β, etc.)
- $f(\mathbf{x}, t)$: fibrosis (collagen) fraction
- $\phi(\mathbf{x}, t)$: porosity / fenestration density
- $s(\mathbf{x}, t)$: entropy density (thermodynamic state)

5.2 Proposed Structure

The tissue Lagrangian density takes the general form:

$$\boxed{\mathcal{L}_{\text{tissue}} = \underbrace{\mathcal{L}_{\text{flow}}}_{\text{fluid kinetics}} + \underbrace{\mathcal{L}_{\text{transport}}}_{\text{solute dynamics}} + \underbrace{\mathcal{L}_{\text{matrix}}}_{\text{structural remodelling}} + \underbrace{\mathcal{L}_{\text{cell}}}_{\text{cellular energetics}} + \underbrace{\mathcal{L}_{\text{dissipation}}}_{\text{entropy production}}} \tag{1}$$

Each term has a clear physical interpretation:

5.2.1 Flow term $\mathcal{L}_{flow}$

This is the fluid Lagrangian of sinusoidal blood flow, modified by the tissue microstructure:

$$\mathcal{L}_{\text{flow}} = \frac{1}{2}\rho\,\phi\,|\mathbf{v}|^2 - p_{\text{eff}}(\rho, f, \phi) - \rho\,\Phi_{\text{grav}}, \tag{2}$$

where ϕ (porosity) modulates the effective flow volume and p_{eff} is an effective pressure that depends on both fluid density and tissue structure (fibrosis narrows sinusoids, increasing flow resistance).

5.2.2 Transport term $\mathcal{L}_{transport}$

Solute transport along the sinusoid is described by advection-diffusion-reaction terms:

$$\mathcal{L}_{\text{transport}} = \sum_i \left[\frac{1}{2} D_i \left| \nabla c_i \right|^2 + c_i \left(\mathbf{v} \cdot \nabla \mu_i \right) - U_i(c_i, f, \phi) \right], \tag{3}$$

where D_i are diffusion coefficients, μ_i are chemical potentials, and U_i encodes reaction terms (oxygen consumption, growth factor binding, etc.).

5.2.3 Matrix term $\mathcal{L}_{matrix}$

The extracellular matrix (collagen) has both elastic energy and a remodelling potential:

$$\mathcal{L}_{\text{matrix}} = \frac{1}{2} \kappa \left| \nabla f \right|^2 - V_{\text{matrix}}(f, c_{\text{TGF}}, c_{\text{MMP}}), \tag{4}$$

where the gradient term penalises sharp fibrosis boundaries (surface tension of fibrotic fronts), and V_{matrix} encodes the balance between collagen deposition (driven by TGF-β) and degradation (driven by matrix metalloproteinases, MMPs).

5.2.4 Cellular energetics $\mathcal{L}_{cell}$

Hepatocyte and stellate cell behaviour is encoded as an effective potential landscape:

$$\mathcal{L}_{\text{cell}} = -W(\phi, f, c_{\text{O}_2}, c_{\text{TGF}}), \tag{5}$$

where W captures stellate cell activation, hepatocyte viability, fenestration regulation, and the feedback loops that drive disease progression. The specific form of W is where biological knowledge enters the model.

5.2.5 Dissipation $\mathcal{L}_{dissipation}$

Viscous dissipation and irreversible processes are included via a Rayleigh-type dissipation function:

$$\mathcal{R} = \frac{1}{2} \eta(\phi, f) \left| \nabla \mathbf{v} \right|^2 + \sum_i \frac{1}{2} \gamma_i \, \dot{c}_i^2, \tag{6}$$

where η is the effective viscosity (which increases with fibrosis as sinusoids narrow) and γ_i are relaxation coefficients for solute dynamics. The dissipation function modifies the Euler–Lagrange equations via $\frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = -\frac{\partial \mathcal{R}}{\partial \dot{q}}$.

6 The Liver as Initial Target

The liver is an ideal first target for a Lagrangian digital twin for several reasons:

6.1 Flow-Dominated Architecture

The liver receives approximately 25% of cardiac output—about 1.5 litres of blood per minute through both the portal vein (75%) and hepatic artery (25%). This dual blood supply perfuses a vast sinusoidal network with a total endothelial surface area estimated at 100–200 m^2 in the adult human. No other solid organ has such an intimate coupling between blood flow and parenchymal cell function.

6.2 Measurable Fluid-Mechanical Endpoints

Liver disease progression has well-validated fluid-mechanical biomarkers:

Biomarker	Measurement	Clinical Threshold
HVPG	Catheterisation (mmHg)	≥ 10: clinically significant
Liver stiffness	Elastography (kPa)	≥ 12.5: cirrhosis
Flow velocity	Doppler ultrasound (cm/s)	Reversal = decompensation

These are direct outputs of a fluid-mechanical model, providing natural validation targets for a Lagrangian digital twin.

6.3 Well-Characterised Disease Progression

Chronic liver disease follows a stereotyped trajectory (F0 → F1 → F2 → F3 → F4 → decompensation) with well-studied intermediate biomarkers. The METAVIR fibrosis scoring system, HVPG thresholds, and Child–Pugh/MELD scores provide a rich validation framework.

6.4 Clinical Urgency

Chronic liver disease affects over 1.5 billion people worldwide. Non-alcoholic fatty liver disease (now termed metabolic dysfunction- associated steatotic liver disease, MASLD) alone affects approximately 25% of the global population. A predictive digital twin could guide drug development (e.g., resmetirom, semaglutide) and personalise treatment decisions.

7 Literature Survey: Novelty of the Approach

To substantiate the claim that a Lagrangian-based biological digital twin is novel, we surveyed the literature across several relevant domains. Our conclusion is that, **to our knowledge, no research group has proposed or implemented a digital twin of any biological organ based on a variational (Lagrangian) principle**.

7.1 Liver-Specific Digital Twins

The most advanced liver modelling efforts include:

- The **Virtual Liver Network** [?] (Germany, 2010–2016): Multi-scale modelling of liver metabolism using coupled ODEs and agent-based models. No variational formulation.
- **HepatoNet1** [?]: A genome-scale metabolic network model solved by flux balance analysis. Algebraic constraint-based, not dynamical.
- **Ricken et al.** [?]: FEM models of liver perfusion and lobule mechanics using porous media theory (Biot equations). Continuum mechanics, not variational.
- **Debbaut et al.** [?]: CFD modelling of hepatic blood flow using patient-specific vascular geometries from CT scans. Eulerian Navier–Stokes, not variational.
- **Ashworth et al.** [?]: Computational model of zonated hepatic energy metabolism, used to study steatosis in NAFLD. ODE-based compartmental approach—no variational formulation.

7.2 Cardiac Digital Twins

The most mature organ digital twin field is cardiac modelling:

- **Living Heart Project** (Dassault, 2014–present): FEM-based electromechanical model. Uses constitutive laws, not a Lagrangian.

- **Niederer et al.** [?]: Review of cardiac digital twins. Discusses FEM, lumped-parameter, and electrophysiology models. No variational formulation mentioned.
- **Corral-Acero et al.** [?]: "The 'Digital Twin' to enable the vision of precision cardiology." Surveys imaging-based models and machine learning. No Lagrangian framework.

7.3 Variational Methods in Biology

Variational principles have been applied to specific biological sub-problems, but never as a unifying framework for organ-scale digital twins:

- **Helfrich free energy** [?]: A variational formulation for cell membrane shape. Widely used in biophysics but describes equilibrium membrane geometry, not organ-scale dynamics.
- **Murray's law** [?]: Minimisation of total power dissipation to predict vascular branching ratios. A static optimality principle, not a dynamical Lagrangian.
- **Prigogine's minimum entropy production** [?]: A variational principle for near-equilibrium thermodynamics applied to biological systems. Controversial and limited to linear irreversible thermodynamics; not used for spatially-resolved tissue models.
- **Mori et al.** [?]: Variational formulations for cell motility using phase-field models. Single-cell scale, not organ-scale.
- **Onsager variational principle** [?]: Application to soft matter dynamics including gels and polymers. Relevant to extracellular matrix mechanics but not applied to organ-scale modelling.

7.4 Assessment

The gap is clear: variational principles are established in fluid mechanics, well-explored at the molecular and cellular scale in biophysics, and extensively theorised in non-equilibrium thermodynamics—but **no group has assembled these ingredients into a Lagrangian-based digital twin of a biological organ**. The approach proposed here is, to our knowledge, **without precedent** in the literature.

8 Advantages Over Existing Approaches

1. **Principled multi-physics coupling.** Adding a new biological process (e.g., immune cell infiltration, bile transport, lymphatic drainage) requires only adding a term to $\mathcal{L}_{\text{tissue}}$. All cross-couplings emerge from the Euler–Lagrange equations. In FEM or ABM, each new process requires hand-crafted coupling code.

2. **Conservation by construction.** Mass, momentum, energy, and topological invariants (vascular connectivity) are automatically conserved. In multi-year disease simulations, this prevents unphysical drift.

3. **Structural testability.** The tissue Lagrangian makes explicit predictions about conservation laws and symmetries. If the model predicts that a certain quantity is conserved and experiment shows it is not, the Lagrangian must be modified—providing a clear falsification pathway.

4. **Approximation hierarchy.** Coarse-grained models (lobule-scale → organ-scale → whole-body) can be derived systematically by averaging the Lagrangian over fast/small-scale degrees of freedom.

5. **Natural interface with control theory.** Drug interventions and clinical decisions can be formulated as optimal control problems on the action functional, leveraging well-developed mathematical tools (Pontryagin's maximum principle, Hamilton–Jacobi–Bellman equations).

6. **Connection to fundamental physics.** Because the tissue Lagrangian has the same mathematical structure as field-theoretic Lagrangians in physics [**??**], the entire apparatus of modern theoretical physics—symmetry analysis, perturbation theory, renormalisation, topological methods—becomes available for biological modelling.

9 Research Programme

We propose a phased research programme to develop and validate the Lagrangian digital twin approach:

9.1 Phase 1: Proof of Concept

- Implement a minimal tissue Lagrangian for a single liver sinusoid, coupling blood flow, oxygen transport, and collagen dynamics.
- Validate against known sinusoidal haemodynamics and fibrosis progression rates.
- Demonstrate that conservation laws (mass, momentum) are maintained over long simulation times without explicit enforcement.

9.2 Phase 2: Multi-Scale Extension

- Extend from single sinusoid to lobule scale using coarse-graining of the Lagrangian.
- Incorporate hepatocyte metabolic zonation as a flow-dependent field.
- Validate against HVPG measurements and elastography data at different fibrosis stages.

9.3 Phase 3: Disease Trajectory Prediction

- Simulate full F0 $\rightarrow$ F4 disease progression for alcohol-associated liver disease (ALD) and MASLD.
- Introduce drug intervention as perturbations to the Lagrangian and predict treatment response.
- Compare with clinical trial data (e.g., MAESTRO-NASH for resmetirom).

9.4 Phase 4: Generalisation

- Adapt the framework to other flow-dominated organs: kidney (glomerular filtration), lung (alveolar ventilation/ perfusion), brain (cerebrovascular coupling).
- Develop a general-purpose "tissue Lagrangian toolkit" for the computational physiology community.

10 Discussion: Challenges and Limitations

We acknowledge several open challenges:

1. **Choosing the right Lagrangian.** Unlike fundamental physics, where symmetries strongly constrain the Lagrangian, biology has many possible terms and the "correct" Lagrangian is not known a priori. It must be determined by a combination of physical reasoning and empirical calibration.

2. **Dissipation.** Biological tissues are far from equilibrium. While the Rayleigh dissipation function and GENERIC (General Equation for Non-Equilibrium Reversible- Irreversible Coupling) framework provide tools, the treatment of strong dissipation within a variational framework remains an active research area.

3. **Biological complexity.** The liver contains dozens of cell types, hundreds of signalling molecules, and feedback loops at every scale. Any tractable Lagrangian must represent a radical simplification, retaining only the fields that dominate the dynamics. The claim is not that a Lagrangian captures everything, but that it captures the essential fluid-mechanical backbone on which the rest is scaffolded.

4. **Computational cost.** Solving the full Euler–Lagrange equations on a 3D tissue domain with multiple coupled fields is computationally demanding. However, GPU acceleration and the inherent parallelism of lattice-based Lagrangian methods make this increasingly tractable.

11 Conclusion

We have argued that:

1. Biological tissue function is fundamentally fluid-driven, with blood flow, shear stress, and pressure gradients governing cellular behaviour, tissue structure, and organ function.

2. The Lagrangian (variational) formulation of fluid dynamics—which encodes conservation laws, symmetries, and multi-scale structure in a single scalar functional—is a natural and potentially superior foundation for biological digital twins.

3. The liver is an ideal initial target due to its flow-dominated architecture, measurable fluid-mechanical biomarkers, well-characterised disease progression, and enormous clinical need.

4. No research group has, to our knowledge, proposed or implemented a Lagrangian-based digital twin of any biological organ. This approach is novel and fills a structural gap in the computational physiology landscape.

The tissue Lagrangian proposed here is a *framework*, not a finished model. Its value lies in providing a principled starting point from which all governing equations, conservation laws, and coupling terms emerge from a single variational principle—replacing the ad hoc integration of disparate sub-models that characterises current practice.

We believe this approach has the potential to transform biological digital twins from primarily empirical constructions toward more theoretically grounded physical models, bringing to biology the same clarity that the Lagrangian formulation brought to physics over the past two and a half centuries.

> *"If the Lagrangian can describe the Standard Model, general relativity, and electromagnetism, why not a liver? The blood already knows the variational principle—it follows the path of least resistance through every sinusoid, every capillary, every day of our lives. We need only write down what it already knows."*

Companion Papers

This paper is the third in a background series:

1. *The Lagrangian Approach to Physics: A Primer on Why Modern Physics Speaks the Language of Action* — introduces the variational framework for particles and fields.
2. *The Lagrangian Approach to Fluid Dynamics: A Primer on Variational Methods for Continuous Media* — extends the framework to fluids, including rotating and quantum fluids.
3. *This paper* — proposes the application to biological tissue and organ-scale digital twins.

Use of AI Tools

Portions of this manuscript were drafted with the assistance of large language models (Google Gemini, Anthropic Claude). The author reviewed, verified, and takes full responsibility for all content, including all mathematical derivations, citations, and factual claims.