Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

ALS as a Multivariable Feedback Instability System: Three-Variable Dynamical Model, Phase-Space Analysis, and Optimal Control

Zhuofan Shen

alexshen2468@gmail.com

UIC Barcelona https://orcid.org/0009-0005-4304-8391

Research Article

Keywords: Amyotrophic lateral sclerosis, ordinary differential equation model, phase-space analysis, optimal control, Pareto front, Michaelis-Menten pharmacodynamics, systems pharmacology

Posted Date: September 22nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11104563/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

ALS as a Multivariable Feedback Instability System: Three-Variable Dynamical Model, Phase-Space Analysis, and Optimal Control

Corresponding Author: Zhuofan Shen
Affiliation: Independent Researcher
Correspondence: alexshen2468@gmail.com
Conflict of interest: The author declares no competing interests.
Data and code availability: All model equations, parameter tables, and numerical workflows are fully reproducible. Source code is available upon reasonable written request.
Disclaimer: This is a computational modelling preprint; it does not provide clinical medical guidance. All model parameters and simulation outputs require independent code verification and experimental dataset calibration prior to clinical translation.

Abstract

Background

Amyotrophic lateral sclerosis (ALS) is a fatal neurodegenerative disease driven by multifactorial crosstalk among energy-metabolic impairment, disrupted proteostasis, and self-amplifying neuroinflammation. Conventional single-target therapeutics have yielded modest clinical benefits, implying ALS should be conceptualised as progressive destabilisation of a multivariable positive-feedback dynamical system rather than a mono-molecular disorder.

Objective

We constructed an analytically tractable three-dimensional ordinary-differential-equation (ODE) model of ALS pathogenesis. Using phase-space geometry and control-theory frameworks, we aimed to quantitatively assess multi-target combination therapy efficacy, optimise intervention timing, and characterise efficacy-versus-dosing-burden trade-offs.

Methods

State variables represented energy-metabolic status (E), proteostasis capacity (P), and neuroinflammatory intensity (I). Three control inputs described metabolic support (u_{meta}), integrated stress response (ISR) inhibition (u_{isr}), and glial-mediated inflammatory clearance (u_{gli}), each governed by Michaelis-Menten-type saturated pharmacodynamic functions. The ODE system was numerically integrated with the LSODA adaptive solver. Analytical workflows included phase-space trajectory analysis, critical-threshold parameter scanning, Pareto multi-objective optimisation, proportional-integral (PI) closed-loop feedback control, and toxicity-constrained piecewise-constant optimal control.

Results

Without therapeutic intervention, model trajectories diverged from early-disease initial conditions and converged onto the pathological attractor $X_D \approx (0{,}0, I_{\max})$ within approximately 4.5 time units. Open-loop triple-agent constant-dose combination therapy redirected system trajectories toward the healthy homeostatic attractor $X_H \approx (1{,}1{,}0)$. Parameter mapping identified a well-defined “survival boundary”; control regimens falling below this threshold could not halt pathological collapse. Pareto optimisation demonstrated that the reference open-loop constant-dose protocol lay within the dominated sub-optimal region, leaving substantial room for dosing-burden reduction. Relative to proportional-only (P) feedback, PI feedback control reduced steady-state error by roughly one order of magnitude, at the expense of higher cumulative drug exposure. Under fixed cumulative-toxicity constraints, constrained optimal control yielded front-loaded piecewise-constant dosing schedules that minimised time-integrated disease burden: higher-magnitude interventions were applied early, followed by low-intensity maintenance.

Conclusions

The performance of multi-modal ALS combination therapy strongly depends on intervention timing and control-strategy design. Closed-loop feedback outperforms static open-loop constant dosing, and constrained optimal control further balances therapeutic efficacy against drug-related toxicity limits. Therapeutic development for ALS should shift from “inhibiting individual molecular targets” toward a systems-oriented goal: steering the overall pathological dynamical trajectory back toward the healthy homeostatic attractor.

Keywords: Amyotrophic lateral sclerosis; ordinary differential equation model; phase-space analysis; optimal control; Pareto front; Michaelis-Menten pharmacodynamics; systems pharmacology

---

1 Introduction

Amyotrophic lateral sclerosis (ALS) causes progressive degeneration of upper and lower motor neurons, with a median post-diagnosis survival of only 3–5 years. Pathogenesis arises from tightly coupled pathological cascades: genetic variants in *SOD1*, *C9orf72*, *TARDBP*, and *FUS* drive protein misfolding and aggregate formation; mitochondrial dysfunction impairs cellular energy supply; and sustained microglial-astrocytic activation establishes

self-propagating neuroinflammation. Rather than acting via simple linear additive effects, these pathological modules amplify each other through interlocking positive-feedback loops. For example, protein aggregates worsen mitochondrial stress; damaged mitochondria release damage-associated molecular patterns (DAMPs) that trigger neuroinflammation; and inflammatory cytokines further repress cellular protein-degradation machineries.

Single-target drug-discovery programmes for ALS have repeatedly faced translational limitations. Riluzole extends median survival by merely 2–3 months; edaravone delivers limited clinical benefit; and the antisense oligonucleotide tofersen achieves target engagement for SOD1-mutant ALS but addresses less than 2 % of the overall ALS patient population. These clinical setbacks motivate a systems-level paradigm shift: if ALS fundamentally represents destabilisation of a multi-variable feedback-regulated biological system, the therapeutic objective ought not to be “blocking a single molecule”, but “re-directing the system’s dynamical trajectory back to its healthy homeostatic attractor”.

Computational systems biology supplies quantitative tools for this paradigm. Lee et al. built a feedback-driven dynamical model for wild-type and SOD1-G93A transgenic mice, showing that untreated ALS-like systems exhibit oscillatory instability from elevated feedback gain, while combinatorial interventions restore near-homeostatic dynamics. Burlando et al. described a multi-stationary loop model highlighting reciprocal coupling between mitochondrial failure and glucose metabolic reprogramming in ALS. De Nardi et al. reported a minimal four-ODE mechanistic model centred on p62-dependent selective autophagy and mTOR signalling, predicting multiple qualitative dynamical behaviours including stable healthy homeostasis, oscillatory stress responses, and stress-driven pathological transitions. Kuznetsov formulated a mathematical description of TDP-43 inclusion-body growth and derived scaling laws for aggregate expansion and neurotoxic accumulation over time.

Nevertheless, most existing ALS dynamical-modelling work focuses on *forward simulation*: forecasting disease trajectories given predefined fixed parameter sets. For clinical translation, the higher-priority open question belongs to *inverse control theory*: given a desired target healthy attractor, what intervention profiles should be deployed? How should timing, magnitude, and duration of multi-agent therapies be selected to jointly optimise clinical outcome and cumulative dosing burden? Addressing these questions requires merging mechanistic dynamical modelling with formal control-theory methodologies.

In the present work, we develop a three-state-variable ODE system defined by energy-metabolic state E, proteostasis capacity P, and neuroinflammatory intensity I, with three distinct classes of therapeutic interventions acting as control inputs: metabolic support u_{meta}, ISR-pathway inhibition u_{isr}, and glial-driven inflammatory clearance u_{gli}. We systematically investigate four core research questions:

1. What are the dynamical signatures and intrinsic timescales of unimpeded pathological system collapse?

2. How are critical “survival boundaries” for multi-agent control inputs distributed across parameter space?

3. What efficacy, steady-state-error, and cumulative-exposure trade-offs separate open-loop static dosing, proportional-only feedback, and proportional-integral closed-loop

feedback strategies?

4. What qualitative temporal architectures characterise toxicity-constrained optimal dosing schedules?

The core contribution of this manuscript is to recast the multifactorial, descriptive pathology of ALS as a quantitatively solvable dynamical-control problem. We deliver a minimal-but-sufficient modelling framework to evaluate multi-target, time-varying combinatorial therapeutic strategies for ALS.

2 Model and Methods

2.1 State variables

The system state vector is defined as

- E: energy-metabolic state, normalised readout for cellular ATP/NAD(H) availability;
- P: proteostasis state, composite metric integrating protein-translational activity and protein-degradation capacity;
- I: pro-inflammatory neurotoxic intensity within the neuronal microenvironment.

Two key attractors are defined: the healthy homeostatic attractor $X_H \approx (1,1,0)$ and the irreversible pathological attractor $X_D \approx (0,0,I_{\max})$.

2.2 Governing ordinary differential equations

The coupled dynamical system reads:

The saturated Michaelis-Menten pharmacodynamic operator is

This functional form enforces target-dependent drug action: pharmacological efficacy tends toward zero when the biological target variable $x \to 0$, encoding the biological constraint that therapeutic agents require available molecular substrates to exert measurable effects.

Equation biological interpretation

1. Energy equation (dE/dt): the first term represents logistic self-regenerative capacity of cellular energy metabolism (normalised ceiling = 1). The second and third terms model inhibitory pressure exerted by impaired proteostasis and neuroinflammation upon energy homeostasis. The saturation-form control term for metabolic support scales therapeutic effect according to the magnitude of existing energy deficit $(1 - E)$.

2. Proteostasis equation (dP/dt): the Hill-function term $\frac{I^n}{K_I^n + I^n}$ describes ISR-mediated translational repression triggered by neuroinflammation, exhibiting saturable behaviour above an inflammatory intensity threshold. The ISR-inhibition control term's pharmacological potency diminishes as proteostasis substrate P is depleted.

3. Inflammation equation (dI/dt): $\alpha_I I$ represents positive-feedback self-amplification of neuroinflammation, the core driver of pathological runaway. The coupling term $\beta_I(1 - P)(1 -$

$E)$ captures de-novo inflammatory activation triggered by concurrent failure of both proteostasis and energy metabolism; $\delta_I I$ is intrinsic inflammatory clearance; the saturation-type control term describes glial-modulated inflammatory-cell clearance therapy.

2.3 Baseline parameters and initial-condition specification

Baseline model parameters are summarised in **Table 1**. Initial conditions represent an early-stage pathological disease state offset from healthy homeostasis: $\mathbf{x}_0 = (E_0, P_0, I_0) = (0.6,,0.5,,0.5)$.

Table 1. Baseline model parameter definitions and numerical values

Parameter	Biological description	Value
α_E	Energy-metabolic self-regeneration rate	0.8
γ_E	Inhibitory coefficient: proteostasis failure on energy metabolism	0.3
β_E	Inhibitory coefficient: neuroinflammation on energy metabolism	0.4
α_P	Proteostasis intrinsic self-recovery rate	0.5
β_P	Inhibitory coefficient: inflammation on proteostasis	0.8
K_I	Hill-equation half-saturation constant (inflammatory repression)	0.3
n	Hill coefficient	2
θ_P	Energy-driven supportive coefficient for proteostasis maintenance	0.2
α_I	Self-amplification rate of neuroinflammatory feedback	0.5

β_I	Inflammatory-triggering coefficient for combined proteostasis-energy failure	0.6
δ_I	Intrinsic neuroinflammation clearance rate	0.2
K_m	Half-saturation constant for metabolic-support pharmacodynamics	0.2
K_p	Half-saturation constant for ISR-inhibition pharmacodynamics	0.3
K_g	Half-saturation constant for glial-clearance pharmacodynamics	0.4

2.4 Four intervention control strategies

Four distinct therapeutic-intervention scenarios were evaluated:

Strategy I: No-intervention natural-history baseline

$u_{\mathrm{meta}} = u_{\mathrm{isr}} = u_{\mathrm{gli}} = 0$. Serves as reference for unperturbed disease progression.

Strategy II: Open-loop constant-dose triple combination therapy

Fixed time-invariant control inputs independent of system state feedback:

Strategy III: Closed-loop proportional-integral (PI) feedback control

Target attractor $X_H = (1,1,0)$. Instantaneous state error vector:

Component-wise PI control law:

Proportional gain vector $\mathbf{K}_p = (1.2,;1.2,;1.0)$; integral gain vector $\mathbf{K}_i = (0.3,;0.3,;0.2)$. For pure proportional-only (P) control, $\mathbf{K}_i = (0,,0,,0)$. All control outputs were hard-clipped to the biologically plausible interval $\mathrm{clip}(u,,0,,1.5)$.

Strategy IV: Toxicity-constrained piecewise-constant optimal control

Total simulation time horizon T was partitioned into $N_{\mathrm{seg}} = 10$ equal-duration

piecewise-constant control segments. We minimise the composite objective functional

subject to cumulative-dosage toxicity constraint

with per-segment input bounds $u_{k,i} \in [0,1.5]$. Optimisation used sequential least-squares programming (SLSQP); objective functional evaluation relied on LSODA numerical integration; penalty weight $\lambda = 0.05$.

2.5 Pareto-front multi-objective analysis

Latin-hypercube sampling (LHS) generated 250 distinct control-input tuples sampled over the input domain $[0,0.6] \times [0,1.0] \times [0,1.0]$. Two competing objectives were computed for each regimen:

1. Cumulative total dosage $\sum u_i \cdot T$ (to be minimised, cost metric);
2. Terminal-time energy-metabolic status $E(T)$ (to be maximised, efficacy metric).

Non-dominated sorting was applied to extract the Pareto-optimal front of treatment regimens.

2.6 Numerical implementation

ODE integration was performed via `scipy.integrate.solve_ivp` with the LSODA solver, adaptive switching between Adams (non-stiff) and BDF (stiff) algorithms; relative tolerance 10^{-8}, absolute tolerance 10^{-10}. Jacobian-matrix eigenvalue analysis assessed local stability of system equilibria. One-dimensional critical-threshold searches deployed the Brent-q root-finding algorithm; two-dimensional response surfaces were computed over regular parameter grids. SLSQP optimal-control optimisation convergence tolerance 10^{-6}, maximum iteration count 200. All computational workflows are reproducible within standard Python/SciPy environments.

3 Results

3.1 Un-intervened pathological collapse dynamics

Starting from the early-disease initial state $(E_0, P_0, I_0) = (0.6,0.5,0.5)$, the uncontrolled system exhibits bistable attractor geometry. Initial time-derivative estimates are:

All local vector-field components drive pathological deterioration: E and P decay monotonically, while inflammatory intensity I rises rapidly. Event detection shows energy-metabolic status crosses the collapse threshold $E < 0.05$ at approximately $t \approx 4.5$ time units; at this time point P approaches zero and I climbs to ≈ 1.2. System trajectories fall irreversibly from healthy-region phase space toward the pathological attractor $X_D \approx (0,0, I_{\max})$.

The term $\alpha_I I$ constitutes the core positive-feedback amplifier. Progressive collapse of E and P activates the $(1-P)(1-E)$ inflammatory trigger term; elevated I in turn further accelerates degradation of E and P via $\beta_E IE$ and Hill-type proteostasis repression. This triple interconnected positive-feedback cascade generates self-reinforcing runaway pathology,

rendering late-stage system states intrinsically difficult to reverse.

3.2 Performance of open-loop triple-combination therapy

For the reference open-loop regimen $\mathbf{u} = (0.25,0.40,0.60)$, initial-time derivatives reverse pathological directions:

Energy status E rapidly recovers toward unity after a brief transient dip. Proteostasis P displays smoother restoration kinetics, partially constrained by reduced available biological substrate for the ISR-inhibition control channel. Neuroinflammatory intensity I drops early; as E and P recover, inflammatory positive-feedback driving forces diminish, and I stabilises near zero. By $t \approx 15$ time units, system trajectories enter the neighbourhood of the healthy attractor $X_H \approx (1,1,0)$.

3.3 Survival boundaries and critical intervention thresholds

Fixing $u_{\mathrm{meta}} = 0.25$ and $u_{\mathrm{isr}} = 0.40$, one-dimensional scanning of u_{gli} across $[0,1.0]$ yields a sigmoidal terminal-response curve for $E(30)$. When $u_{\mathrm{gli}} < 0.12$, $E(30) < 0.5$ and the system succumbs to pathological collapse; for $u_{\mathrm{gli}} > 0.35$, $E(30) > 0.9$ and homeostasis is largely restored. The critical threshold $u_{\mathrm{gli}}^{*} \approx 0.18$ marks the dynamical watershed between system rescue versus irreversible collapse.

Two-dimensional response-surface analysis (u_{meta} versus u_{gli} at fixed $u_{\mathrm{isr}} = 0.40$) demonstrates strong therapeutic synergy. The survival-boundary contour defined by $E(30) = 0.5$ curves convexly toward the origin: moderate dual-agent dosing outperforms high-dose monotherapy. For instance, with $u_{\mathrm{meta}} = 0.1$, survival requires $u_{\mathrm{gli}} \approx 0.55$; increasing u_{meta} to 0.3 lowers the required glial-clearance threshold down to ≈ 0.15.

Variation of u_{isr} produces comparatively modest boundary shifts: raising u_{isr} from 0.0 up to 0.8 displaces the survival boundary leftward by only ~0.1 control-input units. This observation aligns with saturated pharmacodynamic behaviour: ISR-inhibitor efficacy is intrinsically limited when its proteostasis-related biological target P becomes depleted.

3.4 Pareto-front analysis and dosing-economy evaluation

Pareto multi-objective analysis identifies clear efficacy-versus-cumulative-dosage trade-offs. Among 250 Latin-hypercube-sampled regimens, 17 configurations populate the non-dominated Pareto-optimal front. The front exhibits strong convex geometry: at low total cumulative dosage ($\sum u < 3$), each 0.1-unit dosage increment delivers ~0.05 gain in terminal $E(30)$; at high total dosage ($\sum u > 6$), marginal efficacy returns fall below 0.01.

The reference open-loop constant-dose protocol yields $E(30) \approx 0.95$ with cumulative dosage burden of 37.5, and resides well within the dominated (sub-optimal) region of Pareto space. Multiple Pareto-optimal points achieve comparable terminal therapeutic efficacy with only 40 %–60 % of the original cumulative-dosage burden. Optimal trade-off solutions satisfying $E(30) > 0.9$ utilise roughly 55 %–65 % of the reference open-loop total dose. The dominant optimisation pattern consists of reducing u_{isr} (due to its low marginal therapeutic return) while retaining moderate levels of both u_{meta} and u_{gli}.

Without toxicity constraints, static open-loop dosing vastly exceeds the upper toxicity bound $U_{\max} = 6.0$. Under this identical hard cumulative-toxicity ceiling, uniform low-magnitude open-loop interventions cannot prevent pathological collapse; by contrast, piecewise time-varying optimal-control successfully rescues system homeostasis. This highlights that temporal distribution of therapeutic effort carries greater functional importance than raw total cumulative dosage magnitude.

3.5 Comparative assessment of feedback-control performance

Comparing terminal steady-state outputs: proportional-only (P) feedback yields $E \approx 0.92$, $P \approx 0.88$, $I \approx 0.05$; PI closed-loop feedback achieves $E \approx 0.98$, $P \approx 0.96$, $I \approx 0.02$. PI control reduces steady-state error for E from \0.08 down to \0.02 (75 % error reduction). Mechanistically, pure proportional control outputs shrink as state error diminishes, yielding stable off-set equilibria where control forces exactly balance intrinsic system drag. Integral action accumulates historical state error and sustains corrective control inputs until steady-state residual error is eliminated.

Zero residual steady-state error comes at the cost of elevated drug exposure. Under our parameterisation, cumulative drug-exposure area-under-the-curve (AUC) for PI feedback is 1.4–1.6-fold higher than for proportional-only feedback. Translational interpretation: for therapeutics with substantial chronic cumulative toxicity, the small residual steady-state offset under P-only feedback may represent an acceptable clinical compromise; when near-complete physiological normalisation is prioritised, the higher dosing burden of PI control becomes justified.

3.6 Architecture of toxicity-constrained piecewise-constant optimal-control schedules

SLSQP optimisation converged within 50–80 iterations. Optimal piecewise-constant protocols display a characteristic **front-loaded temporal profile**: therapeutic intensity is substantially higher within the first 3–4 time segments ($t \in [0,,9–12]$), then drops to low-amplitude maintenance dosing thereafter. Concretely:

- u_{meta} starts at 0.45–0.55 and decreases to 0.10–0.15;
- u_{gli} initiates at 0.60–0.70, later declining to 0.15–0.25;
- u_{isr} remains comparatively low throughout the whole time horizon (initial 0.20–0.30; late-phase 0.05–0.10), consistent with its low marginal-utility signature identified in Pareto scans.

Cumulative dosage for the optimised schedule totals \5.1, staying safely beneath $U_\mathrm{max}=6.0$, and yields terminal outcomes $E(30)\approx0.93$, $P(30)\approx0.89$, $I(30)\approx0.06$. Relative to the reference open-loop regimen (cumulative dose 37.5), optimal control attains an \85 % reduction in total drug burden with only minor (~0.03) loss of terminal-state E. Under the identical $U_{\max} = 6.0$ toxicity constraint, flat uniform open-loop dosing fails to rescue the system.

These simulation results expose a core dynamical principle: early high-intensity intervention

acts as a “dynamical investment”. When E and P are still partially preserved, corrective interventions incur lower effective biological cost and enable long-term stabilisation using only minimal later-phase maintenance dosing. Once E and P have collapsed and inflammatory positive feedback dominates system dynamics, overcoming the accumulated pathological momentum demands drastically larger total cumulative therapeutic input.

4 Discussion

4.1 Core theoretical contributions

This work translates the descriptive multifactorial pathology of ALS into a formally treatable dynamical-control problem. Our three-dimensional state-variable selection (E, P, I) is grounded in transcriptomic meta-analyses identifying energy-metabolic failure, proteostasis breakdown, and neuroinflammation as shared pan-neurodegeneration hallmarks. The coupling topology — inflammation suppressing metabolism and proteostasis; proteostasis dysfunction fuelling inflammation; energy deficit exacerbating both modules — directly recapitulates experimentally documented positive-feedback interconnections.

Relative to high-dimensional transgenic-mouse dynamical models reported in prior literature, our minimal-sufficient three-variable formulation retains essential feedback architecture while drastically lowering parameter-identification complexity and computational overhead. Low-dimensional state space permits phase-plane visualisation, two-dimensional parameter-boundary mapping, and fast-converging optimal-control workflows, capabilities largely inaccessible for models containing dozens of state variables. While multi-stationary loop models have advanced understanding of mitochondrial-glucose metabolic crosstalk in ALS, their higher dimensionality limits direct deployment for optimal-control synthesis. Our framework balances mechanistic plausibility with analytical and computational tractability.

The Michaelis-Menten saturated pharmacodynamic operator represents another key modelling advance. Unlike linear pharmacology formulations, saturation enforces strict target-dependence: as a biological-target state variable approaches zero, corresponding control-channel potency spontaneously declines. This structural feature reshapes optimal-control behaviour: linear models tend to push extreme therapeutic effort into single channels; saturable pharmacodynamics naturally distribute intervention burden across multiple agents due to diminishing marginal returns. The comparatively low priority assigned to u_{isr} in Pareto optimisations directly originates from this saturation-driven mechanism.

4.2 Context within biomedical control-theory applications

Closed-loop feedback control has well-established precedents in biomedicine, for example PID-regulated artificial-pancreas glucose control or Parkinson’s-disease levodopa delivery optimisation. Those canonical examples are predominantly single-input-single-output (SISO) scalar-target applications. By contrast, our ALS system constitutes a multi-input-multi-output (MIMO) coupled dynamical system with strong cross-channel crosstalk: for instance, metabolic-support input u_{meta} indirectly promotes recovery of P and suppresses I via restoring E. In this proof-of-concept study we adopted independent per-channel gain tuning

as a practical simplification for MIMO feedback; future iterations may deploy more formal MIMO controller-synthesis pipelines.

In optimal-control methodology, constrained drug-scheduling problems for cancer chemotherapy share important conceptual parallels with our formulation. Oncological optimal-control frequently minimises tumour burden subject to toxicity limits, often yielding bang-bang-type solutions. Our piecewise-constant discretisation solved via SLSQP implements a “first-discretise-then-optimise” workflow. Compared with full direct transcription approaches, this discretisation strategy uses coarser state-space partitioning and has reduced dependence on analytical gradients, rendering it robust for non-convex objective landscapes characteristic of biological systems.

4.3 Biological interpretation and translational therapeutic implications

Survival boundaries and threshold phenomena: The computed survival-boundary surfaces demonstrate clear threshold behaviour for ALS combination therapy. When intervention intensity falls beneath critical thresholds, even large-magnitude late-stage efforts cannot reverse system collapse. Clinically, precipitous late-period progression observed in many ALS patients may correspond to crossing this “point-of-no-return” dynamical threshold, where monotherapies fail and adequately-dosed multi-target combinations become mandatory. Our numerically derived critical thresholds offer quantitative benchmarks for minimum-effective combinatorial intervention magnitude, conditional upon the present model parameterisation.

Therapeutic time-window hypothesis: The front-loaded optimal-dosing architecture provides quantitative dynamical grounding for the clinical concept of an early therapeutic window. After state variables pass critical thresholds, motor-neuron loss becomes largely irreversible. Early neuroprotective preservation of E and P achieves desired clinical endpoints with far lower cumulative drug-toxicity burden than attempting rescue of an already severely deteriorated system.

Combination-therapy synergy: Convex survival-boundary contours in the $(u_{\mathrm{meta}}, u_{\mathrm{gli}})$ input plane quantitatively illustrate supra-additive $(1 + 1 > 2)$ therapeutic synergy. This supports prior simulation-based observations that combinatorial regimens restore near-wild-type dynamical profiles in ALS model systems. Our simulations further show that synergy is strongest at moderate dosing magnitudes; once any single therapeutic channel enters saturation, marginal therapeutic gains from further dose escalation diminish sharply.

4.4 Model limitations

1. **Parameter uncertainty**: Current baseline-parameter values represent biologically plausible estimates tuned to reproduce qualitative ALS dynamical behaviours, rather than experimentally fitted against longitudinal in-vivo time-series datasets. Key coupling coefficients such as β_I and δ_I directly shift survival-boundary locations and optimal-control profiles. Future work should perform model calibration against SOD1-G93A mouse longitudinal data, complemented by global-sensitivity analysis (e.g. Sobol indices).

2. **Spatial homogeneity assumption**: The model assumes perfectly uniform spatial distribution of E, P, and I across neuronal-glial populations; it ignores spinal-cord spatial heterogeneity and neuron-glia spatial proximity effects. Extending the framework toward reaction-diffusion partial-differential-equation (PDE) formulations constitutes a natural next step, albeit with substantially increased computational-control complexity.

3. **Abstracted therapeutic control channels**: Each control input aggregates an entire pharmacological class rather than representing individual small-molecule or oligonucleotide agents. For example, u_{meta} lumps together NAD+ boosters and mitochondrial-modulating compounds; u_{gli} aggregates multiple anti-neuroinflammatory modalities. Intra-class pharmacokinetic variability and drug-drug interaction terms are omitted. While acceptable for proof-of-principle dynamical exploration, these simplifications must be refined before generating actionable clinical-trial candidate regimens.

4. **Deterministic ODE framework**: The model is fully deterministic; gene-expression stochastic noise, probabilistic cell-fate decisions, and inter-patient biological heterogeneity are absent. Follow-up work will adapt this structure to stochastic-differential-equation (SDE) or population-simulation workflows to quantify control-strategy robustness under biological variability.

5 Conclusions

We constructed and analysed a minimal three-variable coupled ODE dynamical model capturing core positive-feedback pathological cascades in ALS, integrating phase-space geometry, closed-loop feedback control, and toxicity-constrained optimal-control within a unified evaluation pipeline. Principal findings are summarised below:

1. Un-intervened ALS-like system collapse is driven by triple interlocking positive feedback among energy metabolism E, proteostasis P, and neuroinflammation I. Trajectories converge to the pathological attractor $X_D \approx (0,0,I_{\max})$ on an intrinsic timescale of approximately 4.5 time units.

2. Multi-target combinatorial therapy exhibits convex survival-boundary geometry across control-input parameter space, demonstrating strong therapeutic synergy across intervention channels. ISR-inhibition displays comparatively low marginal therapeutic utility, whereas metabolic-support and glial-inflammatory-clearance represent high-leverage control axes.

3. Reference static open-loop constant-dose regimens are sub-optimal in Pareto-objective space; multi-objective optimisation permits 40 %–60 % reduction in cumulative therapeutic burden without meaningful loss of terminal-state efficacy.

4. Relative to proportional-only feedback, PI closed-loop feedback reduces steady-state error by ~75 % at the cost of roughly 50 % higher cumulative drug exposure. This efficacy-versus-burden trade-off should be matched against the known toxicological profile of candidate therapeutic agents.

5. Toxicity-constrained optimal-control solutions adopt front-loaded piecewise-constant dosing patterns. Early high-intensity therapeutic effort functions as a dynamical “investment”,

enabling sustained homeostatic stabilisation using only low-magnitude maintenance doses in later disease phases.

This study supplies a quantitative dynamical-systems foundation for designing combinatorial ALS therapeutic regimens, illustrating the translational potential of control-theory approaches for neurodegenerative-disease drug development. Upcoming research directions include experimental parameter calibration, spatially extended model variants, and robustness-oriented robust-control formulations under biological stochasticity.

In conceptual summary: ALS should be viewed not merely as a single-molecule disease, but as a system-level feedback-control disorder. Therapeutic priorities should shift away from isolated molecular-target inhibition toward time-aware multi-modal interventions that steer the overall pathological dynamical trajectory back toward healthy homeostasis.

References

1. Lee AJB, Bi S, Ridgeway E, et al. Restoring Homeostasis: Treating Amyotrophic Lateral Sclerosis by Resolving Dynamic Regulatory Instability. *Int J Mol Sci*. 2025;26(3):872.

2. Burlando B, Milanese M, Giordano G, et al. A multistationary loop model of ALS unveils critical molecular interactions involving mitochondria and glucose metabolism. *PLoS One*. 2020;15(12):e0244234.

3. De Nardi A, Paris A, Lauria M. A minimal mechanistic QSP-ready model to bridge aggregate and mTOR signaling in ALS. PAGE Meeting Abstract IV-009; 2026.

4. Kuznetsov AV. Modeling the growth of cytosolic TDP-43 inclusion bodies and accumulated neurotoxicity of misfolded oligomers in neurons. bioRxiv; 2025. doi:10.1101/2023.11.28.569118

5. Meta-analytic human transcriptomics identifies neuroinflammation, deficient energy metabolism, and proteostasis failure across neurodegeneration. *Neurobiol Dis*. 2020;149.

6. Mathematical Modelling of Neuroinflammation in Neurodegenerative Diseases. *CPT Pharmacometrics Syst Pharmacol*. 2025.

7. Constrained optimal control drug scheduling models with different toxicity metabolism in cancer chemotherapy. *Comput Chem Eng*. 2025.

8. PID-controller enhanced artificial β-cells. *PLoS One*. 2026.

9. Nocedal J, Wright SJ. *Numerical Optimization*. 2nd ed. Springer; 2006.

10. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17:261-272.

11. Hindmarsh AC, Brown PN, Grant KE, et al. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans Math Softw*. 2005.

12. Deb K, Pratap A, Agarwal S. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput*. 2002.

Appendix A: Nomenclature

Symbol	Definition
E	Energy-metabolic state
P	Proteostasis / translational-activity state
I	Neuroinflammatory intensity
u_{meta}	Metabolic-support control input
u_{isr}	ISR-inhibition control input
u_{gli}	Glial inflammatory-clearance control input
X_H	Healthy homeostatic attractor
X_D	Pathological attractor
$\mathrm{sat}(\cdot)$	Michaelis-Menten saturated pharmacodynamic function
$U_{\max}$	Cumulative-dosage toxicity upper bound

Appendix B: Reproducibility statement

The dynamical system is a three-dimensional set of ordinary differential equations. All numerical experiments are reproducible within standard Python-SciPy computational ecosystems using the following workflow:

1. Solve the coupled ODE system via `solve_ivp` with LSODA solver;
2. Execute one-dimensional critical-threshold searches using the Brent-q root-finding method;
3. Generate 250 control-input samples by Latin-hypercube sampling;
4. Extract Pareto-optimal front via non-dominated sorting;
5. Solve constrained optimal-control problems with `scipy.optimize.minimize` (SLSQP method).