FLASH Radiotherapy as a State Transition: A Memory-Modulated Oxygen Depletion Model

Vaitheeswaran R. *

Posted Date: 5 May 2026

doi: 10.20944/preprints202603.2392.v2

Keywords: FLASH radiotherapy; radiolytic oxygen depletion; biological memory; state transition; dose-rate effects; temporal structure; radiosensitivity; identifiability

 doi:10.20944/preprints202603.2392.v2

Article

FLASH Radiotherapy as a State Transition: A Memory-Modulated Oxygen Depletion Model

Vaitheeswaran R.

Independent Researcher, India; vaithe1985@gmail.com

Abstract

FLASH radiotherapy, characterized by ultra-high dose rates, has been shown to reduce normal tissue toxicity while preserving tumor control, yet its underlying mechanism remains unresolved. Existing models based on radiolytic oxygen depletion (ROD) successfully capture dose-rate dependence but fail to explain key experimental features, including threshold-like onset, saturation of the sparing effect, and sensitivity to temporal delivery structure. Here, we propose a mechanistic framework—Memory-modulated Radiolytic Oxygen Depletion (M-ROD)—that extends classical ROD by incorporating a bounded, history-dependent internal state. The dynamical structure of this state—cooperative activation, bounded feedback, and characteristic decay—is consistent with that of cooperative biological regulatory processes, including gene regulatory networks. In this framework, dose-rate–dependent stress activates a nonlinear biological state that evolves through induction, bounded feedback, and decay, modulating radiosensitivity alongside oxygen effects. We show that the framework reproduces the defining characteristics of FLASH, including sharp threshold-like transitions, plateau behavior, and strong dependence on pulse spacing, duty cycle, and irradiation sequence, while reducing to conventional radiobiology under low dose-rate conditions. The pulse-spacing sensitivity that distinguishes M-ROD from memoryless models requires the state to relax on a characteristic timescale τ_M of approximately 10–100 ms; we show that bioelectric membrane dynamics, treated as a passive RC system using parameter values from standard electrophysiology, naturally produce relaxation in this range without parameter tuning. The model predicts that the magnitude of the FLASH effect is governed by the extent of state activation rather than dose rate alone, providing a mechanistic explanation for variability across experiments. These results support the interpretation of FLASH as an emergent state transition in a dynamical biological system and offer experimentally testable predictions that distinguish it from memoryless models.

Keywords: FLASH radiotherapy; radiolytic oxygen depletion; biological memory; state transition; dose-rate effects; temporal structure; radiosensitivity; identifiability

1. Introduction

FLASH radiotherapy, characterized by the delivery of radiation at ultra-high dose rates (typically >40 Gy/s), has emerged as a promising paradigm in radiation oncology. A growing body of experimental evidence demonstrates that FLASH irradiation can substantially reduce normal tissue toxicity while maintaining tumor control, a phenomenon now widely referred to as the FLASH effect (Favaudon et al., 2014; Montay-Gruel et al., 2017; Vozenin et al., 2019; Bourhis et al., 2019; Simmons et al., 2019). This effect has been observed across multiple tissues, species, and irradiation modalities, suggesting the presence of a general underlying mechanism.

The most widely studied explanation for the FLASH effect is radiolytic oxygen depletion (ROD), in which rapid radiation-induced consumption of oxygen transiently reduces tissue oxygenation and thereby decreases radiosensitivity (Pratx and Kapp, 2019; Labarbe et al., 2020; Spitz et al., 2019). Within this framework, radiosensitivity is directly linked to instantaneous oxygen concentration, providing a physically grounded mechanism connecting dose rate to biological response.

 doi:10.20944/preprints202603.2392.v2

However, despite its conceptual appeal, ROD-based models exhibit important limitations. In particular, they predict a smooth and continuous dependence of radiosensitivity on dose rate, without an intrinsic mechanism for threshold behavior or saturation. In contrast, experimental observations consistently demonstrate a sharp, threshold-like onset of tissue sparing followed by a plateau at high dose rates (Montay-Gruel et al., 2017; Vozenin et al., 2019). Quantitative analyses further suggest that ROD alone is insufficient to account for the magnitude and shape of the observed response under physiologically relevant conditions (Petersson et al., 2020; Wilson et al., 2020; Labarbe et al., 2020).

A more fundamental limitation of ROD lies in its memoryless formulation. By construction, radiosensitivity depends only on instantaneous oxygen concentration and dose rate, and is independent of prior irradiation history. As a result, ROD cannot account for experimentally observed sensitivities to temporal delivery structure, including pulse spacing, duty cycle, and irradiation sequence. From the standpoint of system identifiability (Bellman and Åström, 1970; Cobelli and DiStefano, 1980), a memoryless model class is structurally insufficient to capture phenomena whose outcomes depend on input history at fixed instantaneous inputs. These observations therefore indicate the presence of an internal regulatory process that evolves over time and integrates prior exposure.

Several alternative mechanisms have been proposed to address these limitations, including radical recombination models (Labarbe et al., 2020), reactive oxygen species saturation (Spitz et al., 2019), and vascular or immune-mediated effects (Durante et al., 2018; Venkatesulu et al., 2019). While these approaches provide valuable insights, they remain fundamentally reactive, describing radiosensitivity as a direct function of instantaneous physicochemical variables, and do not introduce an explicit internal state capable of encoding irradiation history.

In parallel, recent advances in systems biology provide a different perspective. Experimental work has demonstrated that even single non-neural cells can encode temporal patterns of stimulation and modify subsequent responses, exhibiting associative learning at the unicellular level (Doan et al., 2026). Complementary theoretical and computational work has shown that gene regulatory networks support associative conditioning and that such conditioning is accompanied by measurable increases in integrative causal emergence (Pigozzi et al., 2025). At the molecular level, biochemical regulatory networks are well known to exhibit bistability, saturation, and switch-like responses through nonlinear feedback interactions.

These findings suggest a general principle: biological response is inherently state-dependent and history-dependent, rather than purely reactive to instantaneous inputs. In this view, cellular response is governed by an evolving internal state that encodes prior exposure and modulates sensitivity to subsequent stimuli.

Motivated by this perspective, we propose that the FLASH effect arises from the coupling between radiolytic oxygen depletion and a bounded, history-dependent internal state representing adaptive biological activation. In this framework, oxygen depletion acts as a trigger, while the internal state regulates and constrains radiosensitivity.

We further hypothesize that this interaction gives rise to a nonlinear transition in radiosensitivity when dose-rate–dependent stress exceeds a critical threshold. Under this interpretation, the FLASH effect reflects an emergent transition between distinct biological regimes—a radiosensitive regime and a protected regime—rather than a continuous modulation of response.

To test this hypothesis, we develop a theoretical framework, termed Memory-modulated Radiolytic Oxygen Depletion (M-ROD), which extends classical ROD by incorporating nonlinear state dynamics. The state variable is phenomenological: its dynamical structure is consistent with that of cooperative biological regulatory processes, including gene regulatory networks, but the model does not commit to any specific molecular substrate. We show that this formulation reproduces the defining features of FLASH radiotherapy, including threshold behavior, saturation, and sensitivity to temporal structure, while remaining consistent with conventional radiobiology under low dose-rate conditions.

 doi:10.20944/preprints202603.2392.v2

2. Theory

2.1. Conceptual Basis: State-Dependent Radiobiology

Conventional radiobiological models, including the linear-quadratic (LQ) framework and radiolytic oxygen depletion (ROD), assume that radiosensitivity is determined solely by instantaneous physical and chemical conditions. In these formulations, the biological response to radiation is memoryless, depending only on current dose and oxygen concentration.

However, recent advances in systems biology suggest that cellular responses are inherently state-dependent. Single non-neural cells can encode temporal patterns of stimulation and exhibit associative learning at the unicellular level (Doan et al., 2026), and gene regulatory networks have been shown computationally to support associative conditioning with corresponding increases in integrative causal emergence (Pigozzi et al., 2025). At the molecular level, biochemical regulatory networks display bistability, saturation, and switch-like responses through cooperative feedback interactions.

These observations motivate a shift from purely reactive descriptions to dynamical models in which radiosensitivity depends on an evolving internal state. We therefore extend the ROD framework by introducing a coarse-grained internal variable representing adaptive biological activation. The variable is phenomenological—its dynamical structure is consistent with cooperative regulatory processes broadly, but no specific molecular substrate is committed to.

2.2. Oxygen Dynamics

We retain the standard ROD description of oxygen evolution:

$$\frac{dO}{dt} = -k_D\,\dot{D}(t)\,O + \frac{O_0 - O}{\tau_O}$$

where $O(t)$ is the local oxygen concentration, $\dot{D}(t)$ is the dose rate, k_D is the oxygen depletion coefficient, O_0 is the baseline oxygen level, and τ_O is the oxygen recovery timescale.

The first term represents radiolytic depletion proportional to dose rate and oxygen availability, while the second term captures recovery through diffusion and perfusion. This equation governs the fast physicochemical dynamics of the system.

2.3. Internal State Variable

To incorporate biological memory, we introduce a bounded internal state variable:

$$M(t) \in [0, 1]$$

where M = 0 corresponds to an inactive, radiosensitive state and M = 1 corresponds to a maximally activated, protected state. This variable represents the collective activation of adaptive processes, such as redox buffering, DNA damage response pathways, and transcriptional regulation. Its bounded nature reflects finite biological capacity.

2.4. Dose-Rate–Dependent Activation

Activation of the internal state is assumed to depend on dose rate through a cooperative process. Following standard coarse-grained descriptions of biochemical activation, we define:

$$f(\dot{D}) = \frac{\dot{D}^n}{\dot{D}^n + D_c^n}$$

where D_c is the characteristic activation threshold and n controls the steepness of the transition.

Hill-type responses of this form arise generically in cooperative binding and multi-step activation processes across many regulatory systems, capturing threshold-like activation without requiring explicit molecular detail.

 doi:10.20944/preprints202603.2392.v2

2.5. Internal State Dynamics

We construct the evolution equation for M(t) as a coarse-grained dynamical system. The structure of the equation is consistent with that of cooperative biological regulatory processes, but is presented here as a phenomenological model rather than a derivation from any specific molecular mechanism.

Step 1: Generic regulatory structure

A broad class of biochemical regulatory systems can be expressed as:

$$\frac{dM}{dt} = \text{activation} - \text{decay}$$

representing the balance between induction and degradation of an activated state.

Step 2: External activation

We assume that activation is driven by dose-rate–dependent stress:

$activation = \alpha \cdot f(\dot{D})$

where α sets the rate of induction.

Step 3: Linear decay

In the absence of stimulation, the system relaxes toward baseline:

$decay = \gamma M$

where γ defines the characteristic decay timescale of the internal state.

Step 4: Bounded nonlinear feedback

To incorporate cooperativity and bounded growth, we add a low-order polynomial feedback term that vanishes at both the inactive (M = 0) and saturated (M = 1) limits:

$feedback = \eta \cdot M^2(1 - M)$

This cubic form is the simplest polynomial consistent with positive feedback, saturation, and a bounded fixed-point structure. Functionally similar normal-form reductions arise in the analysis of cooperative regulatory systems near critical transitions, and have been studied in the context of gene regulatory network learning and associative conditioning (Pigozzi et al., 2025). Within the present model, the feedback term is included as a structural ingredient consistent with this broader class of cooperative regulatory dynamics; it is not derived from any specific gene regulatory network.

Final internal state equation

Combining these terms yields:

$$\frac{dM}{dt} = \alpha f(\dot{D}) - \gamma M + \eta M^2(1 - M)$$

Interpretation:

- $\alpha \cdot f(\dot{D})$: external activation driven by dose rate
- $-\gamma M$: decay of the internal state toward baseline
- $\eta \cdot M^2(1 - M)$: bounded nonlinear reinforcement and saturation

This equation represents a minimal dynamical system capable of producing threshold behavior, saturation, and history dependence. The four ingredients—cooperative activation, linear decay, bounded feedback, and the bounded support of M itself—together define the model's architecture. Quantitatively, threshold behaviour is set primarily by the Hill function $f(\dot{D})$, saturation by the bounded support and the $(1 - M)$ factor in the feedback term, and history dependence by the decay timescale $1/\gamma$. The cubic feedback term $M^2(1 - M)$ makes a smaller direct contribution to threshold and saturation but completes the model's structural architecture: it provides positive feedback within the activated regime and ensures the model belongs to the family of cooperative regulatory systems whose qualitative behaviour the framework intends to capture.

2.6. State-Dependent Radiosensitivity

We define radiosensitivity as:

$S(t) = S_0 \cdot g(O(t)) \cdot (1 - \beta M(t))$

 doi:10.20944/preprints202603.2392.v2

where S_0 is baseline radiosensitivity, g(O) represents oxygen-dependent modulation, and $\beta \in (0, 1)$ quantifies the strength of biological protection.

The oxygen term captures fast physicochemical effects, while the factor $(1 - \beta M)$ represents slower biological modulation. This linear dependence is the lowest-order approximation consistent with bounded protection.

2.7. Damage Accumulation

Biological damage evolves according to:

$$\frac{dD_{\text{bio}}}{dt} = \dot{D}(t) \cdot S(t)$$

and integrates to:

$$D_{\text{bio}} = \int \dot{D}(t)\, S(t)\, dt$$

Because $S(t)$depends on $M(t)$, accumulated damage depends on the full irradiation history.

2.8. Mapping to Survival

We map biological damage to survival fraction using:

$$SF = \exp(-\kappa D_{\text{bio}})$$

where $\kappa$is a scaling parameter. This ensures consistency with standard radiobiological models in the appropriate limit.

2.9. Limiting Behavior

The model reduces correctly under relevant limits:

- $\dot{D} \ll Dc \Rightarrow M \to 0 \to$ ROD
- $\alpha \to 0 \Rightarrow M \to 0 \to$ ROD
- $O \approx O_0 \to$ LQ-like behavior

Thus:
$LQ \subset ROD \subset M\text{-}ROD$

2.10. Emergent Properties

The coupled system naturally produces:

- threshold behavior, via $f(\dot{D})$
- saturation, via bounded M and the $(1 - M)$ factor in the feedback term
- history dependence, via the time-evolution of M(t)
- sensitivity to temporal structure, via the finite decay timescale $1/\gamma$

These properties emerge directly from the governing equations and are consistent with the qualitative behavior of cooperative biological regulatory systems.

Summary

The M-ROD framework provides an extension of radiobiology by embedding oxygen dynamics within a nonlinear, state-dependent system. By incorporating an internal state variable whose dynamics are consistent with cooperative biological regulatory processes, the model represents the FLASH effect as an emergent transition between biological regimes rather than a purely physicochemical phenomenon.

 doi:10.20944/preprints202603.2392.v2

3. Results

3.1. Model Parameters and Simulation Setup

The simulations presented in this work are based on a set of parameters governing oxygen dynamics, internal state evolution, and radiosensitivity modulation. These parameters were chosen to reproduce the qualitative features of the FLASH effect, including threshold behavior, saturation, and sensitivity to temporal delivery structure, rather than to fit a specific experimental dataset. Qualitative trends were robust to moderate variation in parameter values. The oxygen dynamics are controlled by the depletion coefficient k_D, baseline oxygen level O_0, and recovery timescale τ_O, which together determine the balance between radiolytic consumption and physiological replenishment. The internal state dynamics are governed by the activation rate α, decay rate γ, and nonlinear feedback strength η, which collectively define the timescale and extent of biological activation. The dose-rate–dependent activation function is characterized by the threshold D_c and Hill coefficient n, which control the onset and sharpness of the transition. Radiosensitivity modulation is determined by the protection factor β, while the mapping from accumulated damage to survival is set by the scaling constant κ. All simulations were performed using a fixed time step and total irradiation duration sufficient to capture the relevant dynamical behavior of both oxygen and internal state variables. The full set of parameters used in the simulations is summarized in Table 1.

Table 1. Model parameters used in simulations.

Parameter	Symbol	Value	Description
Oxygen depletion coefficient	k_D	0.05	Rate of radiolytic oxygen consumption
Baseline oxygen level	O_0	1.0	Normalized initial oxygen concentration
Oxygen recovery timescale	τ_O	0.5	Characteristic time for oxygen replenishment
Activation threshold dose rate	D_c	55 Gy/s	Dose-rate scale for internal state activation
Hill coefficient	n	10	Controls sharpness of activation threshold
Activation rate	α	3.0	Strength of dose-rate–induced activation
Decay rate	γ	3.5	Relaxation rate of internal state (memory loss)
Feedback strength	η	3.0	Nonlinear self-reinforcement of internal state
Protection factor	β	0.6	Fractional reduction in radiosensitivity due to activation
Survival scaling constant	κ	2.5	Maps accumulated damage to survival fraction
Simulation duration	T	1.0 s	Total irradiation time
Time step	Δt	5×10^{-4} s	Numerical integration step size

3.2. Consistency with Conventional Radiobiology

We first verify that the M-ROD framework reproduces standard radiobiological behavior under conventional (low) dose-rate conditions. In this regime, the activation function satisfies $f(\dot{D}) \approx 0$, and the internal state relaxes to $M \to 0$. As a result, radiosensitivity depends only on oxygen concentration, recovering the classical ROD description.

Figure 1 shows that survival curves predicted by M-ROD are indistinguishable from those obtained using ROD and the LQ model in conventional dose rate regimes. All three models exhibit

 doi:10.20944/preprints202603.2392.v2

the expected monotonic decrease in survival with dose, with no evidence of threshold or saturation effects. A marked difference arises only under FLASH dose rates.

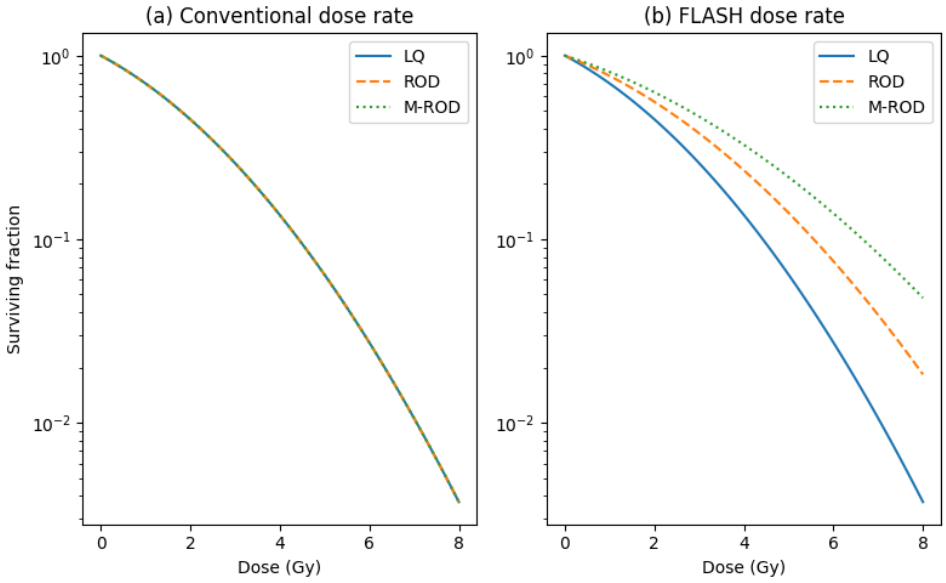

Figure 1. Survival as a function of dose under (a) conventional low dose-rate conditions (~0.01–0.1 Gy/s) and (b) FLASH-like conditions (~40–100 Gy/s). M-ROD reproduces standard radiobiological behavior and overlaps with ROD and LQ predictions in conventional dose rate regimes.

This result confirms that the introduction of the internal state variable does not alter baseline radiobiological behavior and that the model reduces correctly in the appropriate limit.

3.3. Emergence of Threshold and Saturation in Dose-Rate Response

We next examine the dependence of biological response on dose rate under ultra-high dose-rate conditions. Survival is normalized to its low dose-rate value to allow direct comparison across models.

As shown in Figure 2, the ROD model predicts a smooth and continuous reduction in normal tissue response with increasing dose rate, reflecting gradual oxygen depletion. In contrast, the M-ROD model exhibits a sharply nonlinear response characterized by a threshold-like onset of sparing followed by saturation at high dose rates.

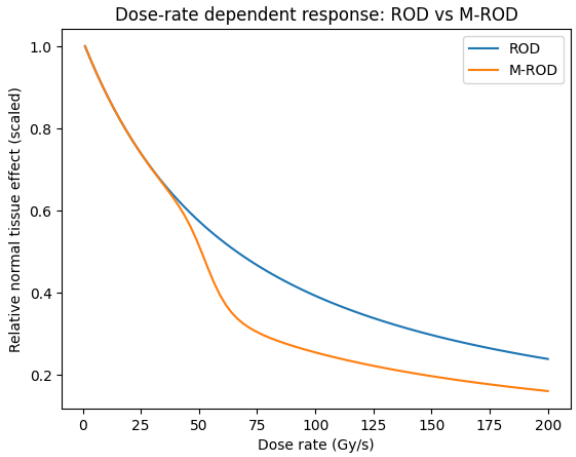

Figure 2. Relative normal tissue response as a function of dose rate. The M-ROD model exhibits a threshold-like onset and saturation of the sparing effect, while the ROD model predicts a smooth, continuous response.

 doi:10.20944/preprints202603.2392.v2

This behavior arises directly from the internal state dynamics. The Hill function f(Ḋ) introduces a threshold for activation, while the bounded support of M(t) and the (1 − M) factor in the feedback term enforce saturation. As dose rate exceeds the characteristic threshold Dc, the system transitions sharply between two dynamical regimes—an inactive regime in which M relaxes to baseline, and an activated regime in which M settles at a partially saturated value sustained by the cooperative dynamics. This regime transition produces the rapid reduction in effective radiosensitivity. The resulting dose-rate response reproduces the key qualitative features of the FLASH effect, including abrupt onset and plateauing, which are not captured by ROD.

3.4. Pulse Spacing Dependence

To investigate the role of temporal structure, we examine the effect of pulse spacing while keeping peak dose rate and pulse width fixed.

Figure 3 shows that the M-ROD model predicts a strong dependence of survival on inter-pulse spacing. When pulses are closely spaced, the internal state accumulates across successive exposures, leading to enhanced protection. As the spacing increases, the decay term −γM reduces the internal state between pulses, resulting in diminished sparing.

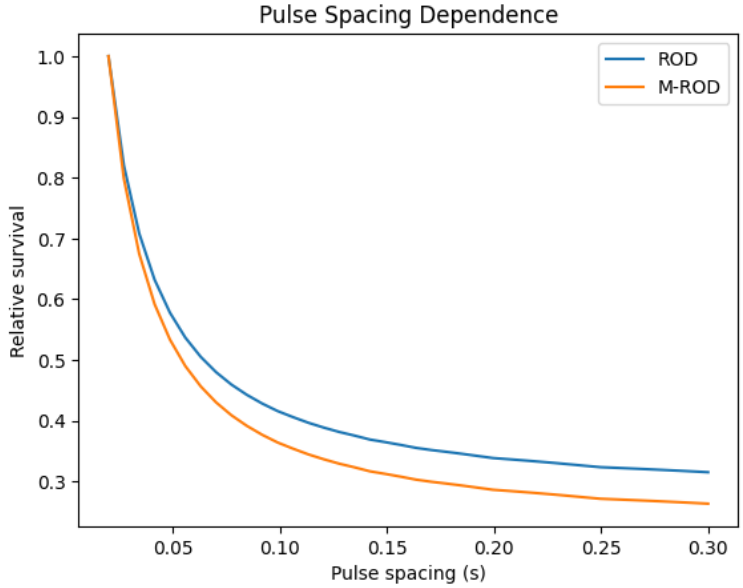

Figure 3. Survival as a function of pulse spacing. M-ROD shows strong sensitivity to inter-pulse interval due to decay of the internal state, whereas ROD shows weak dependence.

In contrast, the ROD model shows only weak dependence on pulse spacing, as oxygen dynamics alone do not retain memory of prior irradiation. This difference highlights the role of the internal state variable in encoding temporal information. The M-ROD model predicts that biological response depends strongly on pulse spacing even when total dose and average dose rate are held constant—a behavior not captured by ROD-based descriptions.

3.5. Time-Ordering Dependence

A key feature of state-dependent systems is sensitivity to the order of inputs. To test this, we compare two irradiation sequences delivering identical total dose: a low-to-high dose-rate sequence and a high-to-low sequence.

As shown in Figure 4, the M-ROD model predicts significantly different outcomes for these two cases. When high dose-rate irradiation is delivered first, the internal state is activated early, reducing radiosensitivity during subsequent exposure. In contrast, when low dose-rate irradiation precedes high dose-rate delivery, activation occurs later, resulting in less cumulative protection. The ROD

 doi:10.20944/preprints202603.2392.v2

model predicts identical outcomes for both sequences, as it depends only on instantaneous oxygen concentration.

Figure 4. Time-ordering dependence. Identical total dose delivered in different temporal sequences produces different outcomes in M-ROD but identical outcomes in ROD.

This contrast demonstrates that the M-ROD framework captures history-dependent effects that are absent in conventional models.

3.6. Duty Cycle Dependence

We next examine the effect of pulse duty cycle by varying pulse width while keeping the period and peak dose rate fixed.

Figure 5 shows that the M-ROD model exhibits a strong dependence on pulse width, with increased sparing observed for longer pulses. This behavior arises because extended high-dose-rate exposure sustains activation of the internal state, allowing nonlinear feedback to amplify the response. In contrast, the ROD model shows a weaker and more gradual dependence, reflecting its sensitivity primarily to average dose rather than instantaneous delivery structure.

Figure 5. Duty cycle dependence. M-ROD shows strong sensitivity to pulse width, while ROD exhibits weaker dependence.

 doi:10.20944/preprints202603.2392.v2

3.7. Partial Activation and Saturation Limits

Finally, we examine the extent to which the internal state approaches full activation. The model imposes a theoretical upper bound on protection through the factor $(1 - \beta M)$, with maximum sparing achieved when $M \to 1$.

However, Figure 6 shows that under realistic conditions, the internal state typically saturates at intermediate values ($M < 1$) due to competition between activation and decay. As a result, the observed plateau in survival remains below the theoretical maximum. This provides a mechanistic explanation for variability in the FLASH effect: the magnitude of sparing depends not only on dose rate but also on the degree of activation achieved, which is governed by irradiation time, delivery structure, and system parameters.

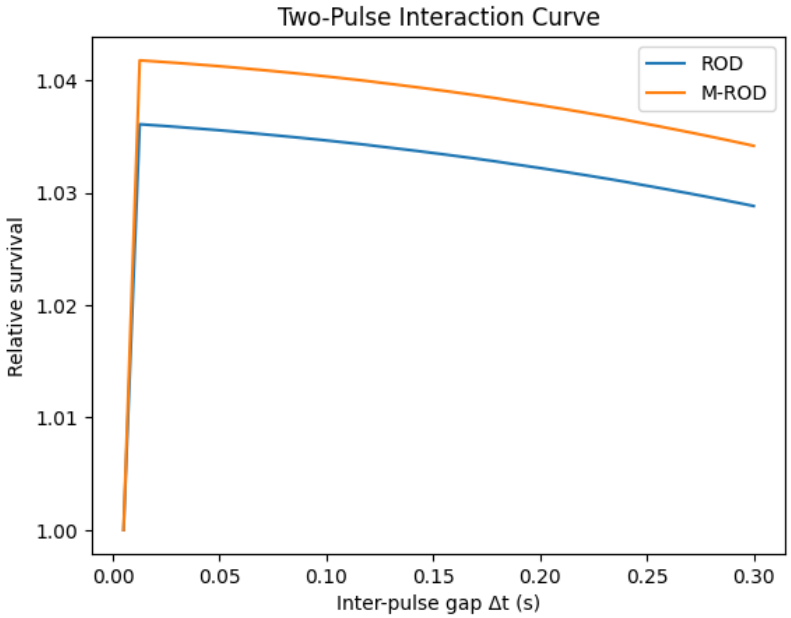

Figure 6. Partial activation of the internal state. The system does not reach full activation under realistic conditions, leading to a plateau below the theoretical maximum.

3.8. A Candidate Substrate Consistent with the Required Timescale

The pulse-spacing sensitivity that distinguishes M-ROD from memoryless models depends on the relationship between the state's decay timescale $\tau_M = 1/\gamma$ and the inter-pulse interval Δt of the FLASH delivery. If τ_M is much shorter than Δt, M decays completely between pulses and each pulse starts from baseline; pulse spacing becomes irrelevant and the FLASH effect loses its temporal-structure dependence. If τ_M is much longer than the total delivery time, M integrates across all pulses regardless of spacing; pulse spacing again becomes irrelevant. Pulse-spacing sensitivity is observable only when τ_M is comparable to Δt itself.

Modern FLASH accelerators deliver pulses of microsecond to millisecond duration, with inter-pulse intervals ranging from approximately 2 ms to 500 ms depending on modality and pulse structure. Numerical simulation of the M-ROD dynamics across this range, with parameters from Table 1, shows pulse-spacing sensitivity peaking at $\tau_M \approx 80$ ms and remaining substantial across $\tau_M \approx 10$–200 ms. Below approximately 5 ms, sensitivity falls because M cannot accumulate across pulses. Above approximately 1 s, sensitivity falls because M approaches saturation regardless of spacing. The framework therefore predicts:

$$\tau_M \approx 10\text{–}100 \text{ ms}$$

 doi:10.20944/preprints202603.2392.v2

This prediction is independent of the Hill function parameters that set the dose-rate threshold (Dc, n) and independent of the activation rate α. It comes specifically from the requirement that pulse-spacing sensitivity be observable in the regime where FLASH accelerators operate.

It is therefore worth asking whether any plausible biological substrate naturally produces relaxation in this range without parameter tuning. The simplest candidate is the bioelectric membrane response. To first approximation, a cell membrane behaves as a passive RC circuit, with capacitance C_m provided by the lipid bilayer and resistance R_m provided by ion channel leak conductance. The relaxation timescale is:

$$\tau = R\_m \cdot C\_m$$

The specific capacitance of cell membranes is constrained by the physical properties of the lipid bilayer and is essentially constant across cell types: C_m ≈ 1 μF/cm² (Hille, 2001). Membrane resistance varies substantially with cell type and physiological state. Electrically active tissues such as cardiac muscle and excitable neurons have R_m on the order of 10^3 Ω·cm², while more quiescent tissues have R_m on the order of 10^5 Ω·cm² (Hille, 2001; Koch, 1999). Substituting:

- R_m = 10^3 Ω·cm² → τ = 1 ms
- R_m = 10^5 Ω·cm² → τ = 100 ms

The membrane RC timescale therefore spans approximately 1–100 ms across the physiologically relevant range of tissue membrane resistance. This range is not chosen to fit the FLASH observation; it follows from a single textbook relation and two parameter values independently established by decades of electrophysiology.

Two independent calculations therefore give overlapping ranges:

- From FLASH pulse-spacing phenomenology: τ_M ≈ 10–100 ms
- From membrane RC biophysics: τ = R_m · C_m ≈ 1–100 ms

The overlap is approximately 10–100 ms. This is not a precise numerical match—both ranges are order-of-magnitude estimates—and it is not a uniqueness result, since other biological substrates with fast dynamics (voltage-gated calcium channels, some rapid redox processes) can also produce sub-second relaxation. What the convergence establishes is more modest: that at least one biological substrate, namely passive bioelectric membrane dynamics, naturally produces a decay timescale in the range that the framework's FLASH prediction requires, using parameters that the electrophysiology literature established independently of any radiotherapy context.

Seen this way, the convergence is a sanity check rather than a proof. A framework that requires τ_M ≈ 10–100 ms would be in difficulty if no known biological process produced relaxation in that range, or if the only candidates required fine-tuning to reach it. The check shows neither obtains. The agreement is between two numbers that were neither chosen nor tuned to match. The check does not, however, demonstrate that bioelectric membrane dynamics are in fact the substrate of M(t) in clinical tumours; it shows only that bioelectric dynamics are a candidate substrate consistent with the framework's timescale requirement. Direct experimental discrimination among candidate substrates is the subject of pulse-spacing experiments designed to fit the decay constant from observed survival data.

4. Discussion

The present work introduces a state-dependent extension of radiolytic oxygen depletion (ROD), motivated by recent advances in systems biology and cellular regulation. In contrast to conventional models, which treat radiosensitivity as an instantaneous function of oxygen and dose rate, the M-ROD framework incorporates a dynamically evolving internal state that encodes prior irradiation history. This shift enables a unified explanation of key experimental features of FLASH radiotherapy that are difficult to reconcile within memoryless formulations. The model occupies an intermediate level of complexity, introducing a minimal set of additional parameters sufficient to capture dynamical behaviors such as thresholding and temporal integration, without resorting to high-dimensional parameterizations required in detailed radiochemical models.

 doi:10.20944/preprints202603.2392.v2

FLASH irradiation operates across multiple timescales, with ultrafast physicochemical processes acting as triggers for slower biological responses. While radiochemical reactions and oxygen depletion occur on microsecond to millisecond timescales (Pratx and Kapp, 2019; Spitz et al., 2019), cellular regulatory and adaptive processes operate over milliseconds to seconds and beyond, with associative-conditioning behavior demonstrated in single non-neural cells (Doan et al., 2026) and in computational gene regulatory network models (Pigozzi et al., 2025). The present framework captures this separation by coupling rapid oxygen dynamics to a slower internal state, enabling memory, threshold behavior, and sensitivity to temporal delivery structure.

A central result of this work is that the defining characteristics of the FLASH effect—namely threshold-like onset, saturation of the sparing effect, and sensitivity to temporal delivery structure—emerge naturally from nonlinear state dynamics. In the proposed framework, these features are not imposed phenomenologically but arise directly from the interaction between dose-rate–dependent activation and bounded feedback within the internal state variable. The activation function introduces a threshold for state transition, while the bounded nature of M and the (1 − M) factor in the feedback term enforce saturation and persistence. This interpretation provides a conceptual reframing of the FLASH effect: rather than a continuous modulation of radiosensitivity driven solely by physicochemical processes, the results suggest that FLASH reflects a transition between distinct biological regimes—a baseline radiosensitive regime and an activated, protected regime. Within this picture, oxygen depletion plays an important but incomplete role, acting primarily as a trigger that initiates state activation rather than as the sole determinant of response.

ROD provides a plausible mechanism for dose-rate–dependent modulation of radiosensitivity but does not inherently account for threshold-like transitions, saturation, or sensitivity to temporal delivery structure. The limitations of purely physicochemical models become particularly evident when considering temporal delivery structure. ROD-based models, which depend only on instantaneous oxygen concentration, predict identical outcomes for irradiation schemes with the same total dose and similar average dose rates. In contrast, the M-ROD framework predicts strong sensitivity to pulse spacing, duty cycle, and irradiation sequence, arising from the finite relaxation timescale of the internal state. Existing models incorporating oxygen depletion or radiochemical dynamics (Pratx and Kapp, 2019; Spitz et al., 2019; Labarbe et al., 2020; Petersson et al., 2020) may exhibit limited sensitivity to pulse structure through recovery processes, but do not predict strong dependence on pulse spacing at fixed total dose and average dose rate. A direct experimental test of the proposed mechanism can be performed by delivering identical total dose and average dose rate using pulsed irradiation while systematically varying pulse spacing. The M-ROD framework predicts a strong dependence of biological response on pulse separation, whereas ROD-based models predict only weak sensitivity arising from oxygen recovery. Such experiments would provide a direct means of distinguishing between memoryless and state-dependent mechanisms.

The argument that memoryless model classes are structurally insufficient to capture history-dependent phenomena has a precise statement in identifiability theory (Bellman and Åström, 1970; Cobelli and DiStefano, 1980). A model whose predictions depend only on instantaneous local inputs cannot, regardless of parameter choice, reproduce input-output behavior in which outcomes depend on input history at fixed instantaneous inputs. The reported dependence of FLASH response on pulse spacing, duty cycle, and irradiation sequence places it in this latter regime. Closing the gap therefore requires the introduction of an internal state variable. M-ROD provides a minimal such extension. It does not claim to be the unique resolution: alternative state-carrying formulations, including different functional forms for the feedback term and different identifications of the substrate, could in principle close the same gap. The model proposed here is offered as a dynamical structure consistent with the observed phenomenology and with the qualitative behavior of cooperative biological regulatory systems.

An important aspect of the present formulation is the relationship between the internal state dynamics and the broader class of cooperative biological regulatory processes. The cubic feedback term $M^2(1 - M)$ is the simplest polynomial consistent with positive feedback, saturation, and a

 doi:10.20944/preprints202603.2392.v2

bounded fixed-point structure, and functionally similar normal-form reductions arise in the analysis of cooperative regulatory systems near critical transitions. Recent work has shown that gene regulatory networks support associative conditioning, with conditioning accompanied by measurable increases in integrative causal emergence (Pigozzi et al., 2025), and that single non-neural cells exhibit associative learning under temporal stimulation patterns (Doan et al., 2026). These findings establish that the kind of behavior captured by M-ROD—bounded, history-dependent, cooperatively regulated state activation—exists in biological systems that share the formal structure of the present model. M-ROD does not derive from any specific molecular pathway. Rather, its dynamical structure is consistent with this broader class of regulatory dynamics, and the model is offered as a coarse-grained phenomenological description whose qualitative predictions are insensitive to the specific molecular substrate that ultimately implements the state.

An important implication of this framework is that the magnitude of the FLASH effect is governed not solely by dose rate but by the extent of internal state activation. Because activation depends on both dose-rate intensity and temporal delivery structure, identical total doses delivered under different conditions can lead to different biological outcomes. This provides a mechanistic explanation for variability observed across experimental studies and suggests that optimization of FLASH protocols may require control of both dose rate and temporal structure.

Despite these strengths, several limitations should be noted. First, the internal state variable is phenomenological. While its dynamical structure is consistent with cooperative regulatory processes, and while the timescale calculation in Section 3.8 identifies bioelectric membrane dynamics as a plausible candidate substrate, direct experimental confirmation of the substrate awaits dedicated measurement. Second, the model assumes spatial homogeneity, neglecting heterogeneity in oxygen distribution and cellular response that may be important in vivo. Third, parameter values were chosen to reproduce qualitative features rather than fitted to a specific dataset; quantitative predictions therefore await calibration against carefully designed experimental measurements. Future work should focus on experimental validation of the model's distinguishing predictions, particularly those related to temporal structure, such as pulse spacing and time-ordering effects. Measurement of the characteristic relaxation timescale would provide direct evidence for or against the presence of a memory-like internal state, and would discriminate among candidate substrates with different intrinsic timescales. In parallel, integrating spatial effects and identifying additional candidate molecular pathways would further strengthen the biological grounding of the model.

Essentially, the M-ROD framework provides an extension of conventional radiobiology that incorporates state-dependent dynamics consistent with cooperative biological regulatory processes. By coupling oxygen depletion with nonlinear internal state evolution, the model reproduces key features of the FLASH effect and generates experimentally testable predictions. These results support the interpretation of FLASH as an emergent state transition in biological response, rather than a purely physicochemical phenomenon. The key conceptual contribution of the present work is the identification of temporal structure, specifically pulse spacing, as an independent determinant of radiobiological response, beyond total dose and dose rate.

Existing models incorporating oxygen depletion or radiochemical dynamics may exhibit limited sensitivity to temporal delivery through recovery processes, but remain fundamentally governed by instantaneous physicochemical conditions (Pratx and Kapp, 2019; Spitz et al., 2019; Labarbe et al., 2020; Petersson et al., 2020). In contrast, the M-ROD framework predicts that biological response depends strongly on pulse spacing even when total dose and average dose rate are held constant, reflecting the presence of an internal state that integrates prior exposure. This establishes temporal structure as an independent control parameter in radiobiology, consistent with broader evidence that biological systems exhibit memory, thresholding, and time-integrated responses (Doan et al., 2026; Pigozzi et al., 2025).

 doi:10.20944/preprints202603.2392.v2

5. Conclusions

The present work introduces a state-dependent extension of radiolytic oxygen depletion (ROD) to explain the FLASH effect. By incorporating a bounded internal state whose dynamics are consistent with cooperative biological regulatory processes, the M-ROD framework reproduces key features of FLASH, including threshold-like onset, saturation, and sensitivity to temporal delivery structure, while remaining consistent with conventional radiobiology at low dose rates. The pulse-spacing sensitivity that distinguishes the model from memoryless formulations requires a state relaxation timescale of approximately 10–100 ms, a range that bioelectric membrane RC dynamics produce naturally without parameter tuning. These results support the interpretation of FLASH as an emergent state transition in biological response, rather than a purely physicochemical effect. Importantly, the model generates experimentally testable predictions—particularly regarding pulse structure, irradiation sequence, and the fitted relaxation timescale—that provide a pathway for distinguishing state-dependent mechanisms from memoryless models, and for discriminating among candidate biological substrates.

Data Availability Statement: Simulation code and data are available upon request.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Bellman, R., & Åström, K. J. (1970). On structural identifiability. Mathematical Biosciences, 7(3–4), 329–339.
2. Bourhis, J., Sozzi, W. J., Jorge, P. G., Gaide, O., Bailat, C., Duclos, F., Patin, D., Ozsahin, M., Bochud, F., Germond, J. F., Moeckli, R., & Vozenin, M. C. (2019). Treatment of a first patient with FLASH-radiotherapy. Radiotherapy and Oncology, 139, 18–22.
3. Cobelli, C., & DiStefano, J. J. (1980). Parameter and structural identifiability concepts and ambiguities: A critical review and analysis. American Journal of Physiology—Regulatory, Integrative and Comparative Physiology, 239(1), R7–R24.
4. Doan, N., Theroux, A., Ramdas, T., & Gershman, S. J. (2026). Associative learning in the protozoan Stentor coeruleus. bioRxiv, 2026-02.
5. Durante, M., Bräuer-Krisch, E., & Hill, M. (2018). Faster and safer? FLASH ultra-high dose rate in radiotherapy. British Journal of Radiology, 91(1082), 20170628.
6. Favaudon, V., Caplier, L., Monceau, V., Pouzoulet, F., Sayarath, M., Fouillade, C., Poupon, M. F., Brito, I., Hupé, P., Bourhis, J., Hall, J., Fontaine, J. J., & Vozenin, M. C. (2014). Ultrahigh dose-rate FLASH irradiation increases the differential response between normal and tumor tissue in mice. Science Translational Medicine, 6(245), 245ra93.
7. Hille, B. (2001). Ion Channels of Excitable Membranes (3rd ed.). Sinauer Associates.
8. Koch, C. (1999). Biophysics of Computation: Information Processing in Single Neurons. Oxford University Press.
9. Labarbe, R., Hotoiu, L., Barbier, J., & Favaudon, V. (2020). A physicochemical model of reaction kinetics supports peroxyl radical recombination as the main determinant of the FLASH effect. Radiotherapy and Oncology, 153, 303–310.
10. Montay-Gruel, P., Petersson, K., Jaccard, M., Boivin, G., Germond, J. F., Petit, B., Doenlen, R., Favaudon, V., Bochud, F., Bailat, C., Bourhis, J., & Vozenin, M. C. (2017). Irradiation in a flash: Unique sparing of memory in mice after whole brain irradiation with dose rates above 100 Gy/s. Radiotherapy and Oncology, 124(3), 365–369.
11. Petersson, K., Adrian, G., Butterworth, K., & McMahon, S. J. (2020). A quantitative analysis of the role of oxygen tension in FLASH radiation therapy. International Journal of Radiation Oncology, Biology, Physics, 107(3), 539–547.
12. Pigozzi, F., Goldstein, A., & Levin, M. (2025). Associative conditioning in gene regulatory network models increases integrative causal emergence. Communications Biology, 8(1), 1027.

 doi:10.20944/preprints202603.2392.v2

13. Pratx, G., & Kapp, D. S. (2019). A computational model of radiolytic oxygen depletion during FLASH irradiation could explain the antitumor effect. Radiotherapy and Oncology, 139, 23–27.
14. Simmons, D. A., Lartey, F. M., Schüler, E., Rafat, M., King, G., Kim, A., Ko, R., Semaan, S., Gonzalez, S., Jenkins, M., Pradhan, P., Shih, Z., Wang, J., von Eyben, R., Graves, E. E., Maxim, P. G., Longo, F. M., & Loo, B. W. (2019). Reduced cognitive deficits after FLASH irradiation of whole mouse brain are associated with less hippocampal dendritic spine loss and neuroinflammation. Radiotherapy and Oncology, 139, 4–10.
15. Spitz, D. R., Buettner, G. R., Petronek, M. S., St-Aubin, J. J., Flynn, R. T., Waldron, T. J., & Limoli, C. L. (2019). An integrated physico-chemical approach for explaining the differential impact of FLASH versus conventional dose rate irradiation on cancer and normal tissue responses. Radiotherapy and Oncology, 139, 23–27.
16. Venkatesulu, B. P., Sharma, A., Pollard-Larkin, J. M., Sadagopan, R., Symons, J., Neri, S., Singh, P. K., Tailor, R., Lin, S. H., & Krishnan, S. (2019). Ultra high dose rate (35 Gy/sec) radiation does not spare the normal tissue in cardiac and splenic models of lymphopenia and gastrointestinal syndrome. Scientific Reports, 9(1), 17180.
17. Vozenin, M. C., De Fornel, P., Petersson, K., Favaudon, V., Jaccard, M., Germond, J. F., Petit, B., Burki, M., Ferrand, G., Patin, D., Bouchaab, H., Ozsahin, M., Bochud, F., Bailat, C., Devauchelle, P., & Bourhis, J. (2019). The advantage of FLASH radiotherapy confirmed in mini-pig and cat-cancer patients. Clinical Cancer Research, 25(1), 35–42.
18. Wilson, J. D., Hammond, E. M., Higgins, G. S., & Petersson, K. (2020). Ultra-high dose rate (FLASH) radiotherapy: Silver bullet or fool's gold? Frontiers in Oncology, 9, 1563.