Version History

Structural Crystallization: A Unified Computational Framework for Memory Formation, Persistence, and Modification

Po-Ting Lin

May 2026

Version 3 (May 2026)

This version corrects the implementation of one baseline model used for comparison. **The ODE system, calibrated parameter values, and all out-of-sample predictions of the Structural Crystallization framework are unchanged.** The correction affects only the default-parameter variant of the Gershman et al. (2017) latent cause model; the fitted variant used in the main out-of-sample comparison is unaffected.

Correction

1. **Latent cause model default parameters** (Method, Latent Cause Model section; Results, Latent Cause Model section and Table 1). Three default parameter values in the Python port of the Gershman et al. (2017) model diverged from the official MATLAB implementation (`imm_localmap.m`): the learning rate η ($0.3 \to 0.2$), the response threshold θ ($0.02 \to 0.03$), and the response gain λ ($0.01 \to 0.005$). The Method section now lists the official default values explicitly. The default-parameter variant was rerun with the corrected values:
    - The Schiller-style blocking ordering is preserved as an ordinal floor effect (standard $\approx 1 \times 10^{-4}$ vs. retrieval-extinction $\approx 1 \times 10^{-5}$).
    - The Monfils-style interval gradient is preserved as a discrete step (≈ 0 for ITI $\in \{1, 3, 10, 30\}$ and ≈ 0.49 for ITI $= 100$).
    - The within-subject CSa $<$ CSb dissociation is **not** reproduced at default parameters (both saturate at ≈ 1.0); the corresponding qualitative claim has been retracted for the default-parameter variant.

What Is Unchanged in Version 3

- The fitted Gershman variant (RMSE $= 0.050$; Schiller 0.35 vs. 0.42; inverted Monfils gradient): its fitting procedure optimizes over η, θ, λ within bounds spanning both the previous and the official defaults, so fitted parameters, in-sample RMSE, and out-of-sample predictions are identical.
- All Version 2 corrections (Items 1–6 below).
- Everything listed under "What Is Unchanged" for Version 2.

Version 2 (April 2026)

This version incorporates minor corrections identified during post-submission self-audit. **The ODE system, calibrated parameter values, and all out-of-sample predictions are unchanged.** The corrections affect descriptive text, one diagnostic label, and one figure data point.

Corrections

1. **Fear-pathway half-life** (Results, spontaneous recovery section). The fear-pathway half-life was originally stated as approximately 34 days. This value was a diagnostic field computed at a reference arousal of $M = 1$ rather than the fear pathway's $M_i = 5.58$. Under $M_i = 5.58$, the operating half-life is $\sim 5.6 \times 10^5$ days, which is consistent with the near-flat D_{fear} trajectory shown in Figure 1(a). The safety-pathway half-life of $\sim$ 37 days at $M_i = 1$ is now stated explicitly.

2. **SEFL half-life span** (Results, Discussion, and Figure 3 caption). The range was originally described as spanning "four orders of magnitude" (factors of 10^4). The correct ratio is $3{,}850{,}429/4.7 \approx 8 \times 10^5$, spanning nearly six orders of magnitude. The individual half-life values are unchanged.

3. **Rescorla–Wagner day-1 Schiller values** (Results, Competing Model Comparison). The day-1 recovery values were stated with the standard-extinction and retrieval-extinction labels reversed. The correct values are 0.43 for standard extinction and 0.42 for retrieval-extinction. The qualitative conclusion is unchanged.

4. **SEFL 1-shock day-365 value** (Results, SEFL section). The text stated the day-365 value for the 1-shock group as 0.11, corresponding to a linear extrapolation from the initial half-life. The actual simulation output (and the value shown in Figure 3(b), which was rendered directly from the simulation JSON) is $\sim 10^{-5}$, reflecting the nonlinear acceleration of dissolution as D decreases ($\beta = \beta_0 \cdot \exp(-D \cdot M_i/D_c)$ grows as D drops). This version aligns the text with the figure.

5. **Schiller CSa day-30 rounding** (Results, within-subject analysis). The CSa day-30 value of 0.0148 was rounded to 0.02; the correct two-decimal rounding is 0.01.

6. **Calibration worker count** (Method). The calibration was described as parallelized across 16 workers. The code uses `scipy`'s `workers=-1`, which delegates to all available CPU cores; the manuscript now reflects this.

What Is Unchanged

- The four governing ordinary differential equations
- The 10 fitted parameter values calibrated to Quirk and Monfils
- The calibration loss and RMSE
- All Schiller, Monfils, SEFL, and appetitive predictions
- All figures showing model curves (including Figure 3(b), which was already correct in Version 1)
- The theoretical claims and conclusions of the paper

Trace

Versions 1 and 2 of this preprint remain accessible via the Research Square version history.