Estimating long-time rare-event statistics from short trajectories

Transition models, value functions, and rare-event simulation

Jonathan Weare

ICERM | September 27, 2026

Courant Institute, New York University

Extreme and Rare Events: Modeling, Algorithms, and Applications

Data required for rare-event statistics

Can we learn rare-event statistics with less total data than the return time of the event?

1. EQUATIONS

Short trajectory segments determine long-time statistics through a Bellman equation.

2. PARAMETRIZATION

Parametrize either the transition operator or the desired statistic itself.

3. SAMPLING

Use the estimate to decide where to generate new trajectory data.

Motivating examples

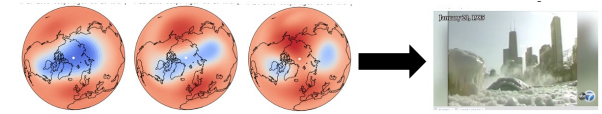

Polar-vortex disruption and its surface impacts.

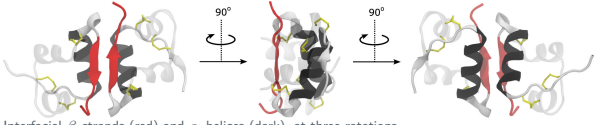

Interfacial β-strands (red) and α-helices (dark), at three rotations.

ATMOSPHERIC EXTREMES

Sudden stratospheric warming: rapid weakening or reversal of the winter polar jet, with cold-air outbreaks at the surface.

MOLECULAR KINETICS

Insulin dimer dissociation: the dimer must come apart before insulin can bind its receptor.

For SSWs, we will estimate neary 1000-year return periods with (many) 46-day forecasts. For insulin, 94 μs of simulation is used (by the Dinner group) to estimate a 100 ms dissociation time.

Rare event statistics

X_t is the (Markov) system state at time t — an atmospheric field, or a molecular configuration.

Initial states are drawn from a **sampling distribution** μ determined by the available data.

We are interested in first escape events from a domain D often with $D = (A \cup B)^{\mathrm{c}}$.

$$T_A = \inf\{t \geq 0 : X_t \in A\}, \quad T_B = \inf\{t \geq 0 : X_t \in B\}, \quad T = \inf\{t \geq 0 : X_t \notin D\}$$

We want to estimate, e.g.

COMMITTOR:

$$q(x) = \mathbf{P}_x(T_B < T_A)$$

Probability of reaching B before A.

MEAN FIRST PASSAGE TIME:

$$m(x) = \mathbf{E}_x[T]$$

Expected time to leave D.

SSW: B = event threshold, A = season end without an event.
Insulin: A = bound dimer, B = dissociated.

Short trajectories and the Bellman equation

Stopped lag-τ trajectories define Bellman equations for committors and stopping-time moments.

Markov state models (MSMs)

Partition state space and estimate short-time transition probabilities.

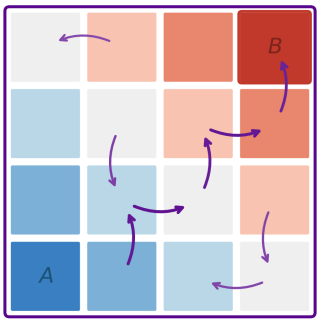

- The short trajectories are typically be much shorter than the event time.
- Long-time statistics for the finite state chain can be computed by linear algebra.
- MSM have a very long history...

Rare event statistics solve the Bellman equation.

For a running reward R and terminal data g, for any $\tau > 0$ the value function

$$u(x) = \mathbf{E}_x\left[g(X_T) + \sum_{t=0}^{T-1} R(X_t)\right] \qquad \text{satisfies} \qquad \boxed{u = S^\tau u + r_\tau, \quad u|_{D^c} = g}$$

where

$$S^\tau u(x) = \mathbf{E}_x[u(X_{\tau\wedge T})]\,, \qquad r_\tau(x) = \mathbf{E}_x\left[\sum_{t=0}^{\tau\wedge T-1} R(X_t)\right].$$

S^τ is the lag-τ transfer operator with trajectories stopped on exit from D. Both terms on the right are defined by segments of length at most τ.

We can solve the equation with parametrization and a data set of the form $\{(X_0^\ell, X_1^\ell, \ldots, X_{\tau\wedge T}^\ell)\}_{\ell=1}^M$.

committor: $R = 0$, $g = \mathbf{1}_B$ **MFPT:** $R = 1$, $g = 0$. [J. Strahan et al *JCTC* 2021]

(Stopped) MSMs are a Galerkin approximation

Approximate u by a function **constant on each cell** $C_1, \dots, C_n$ then use the Galerkin orthogonality condition (in the μ weighted inner product).

Then S^τ is represented by an $n \times n$ matrix and the Bellman equation by a finite linear system:

$$N_{ij} = \sum_{\ell=1}^{M} \mathbf{1}\{X_0^\ell \in C_i,\ X_{\tau \wedge T^\ell}^\ell \in C_j\}, \qquad \widehat{P}_{ij} = \frac{N_{ij}}{\sum_k N_{ik}},$$

Committor ($g = \mathbf{1}_B,\ R = 0$)

$$\widehat{q}_i = \sum_j \widehat{P}_{ij} \widehat{q}_j, \quad \widehat{q}|_A = 0,\ \widehat{q}|_B = 1.$$

Mean stopping time ($g = 0,\ R = 1$)

$$\widehat{m}_i = \widehat{r}_i + \sum_j \widehat{P}_{ij} \widehat{m}_j, \quad \widehat{m}|_{D^c} = 0.$$

If we replace $T \wedge \tau$ by τ (i.e. no stopping), the estimator will have a bias that grows with τ from missed escapes.

Finite-state error analysis and applications

Remarkable data efficiency and applications to SSW return periods and insulin dissociation.

A condition-number bound is wildly pessimistic

Errors in the approximate Bellman solution — from finite data, and from the parametrization — are classically bounded by

$$\|\widehat{u} - u\|_D \;\leq\; \left\|(I - S^\tau|_D)^{-1}\right\| \left(\left\|(S^\tau - \widehat{S}^\tau)\,\widehat{u}\right\| + \left\|r_\tau - \widehat{r}_\tau\right\| \right).$$

$S^\tau|_D$ **is nearly singular.** Its smallest eigenvalue satisfies $\lambda_{\min}(I - S^\tau|_D) \leq \tau/\mathbf{E}_\nu[T]$, with ν the quasi-stationary distribution, so $\|(I - S^\tau|_D)^{-1}\| \geq \mathbf{E}_\nu[T]/\tau$ **— extremely large** for a rare event.

The classic perturbation bound predicts that if $\tau \ll \mathbf{E}_\nu[T]$ then error will be huge.

This worst case bound doesn't reflect the errors we make in practice.

A finite-state error bound (nearly) independent of the return period

FOR A STOPPED MSM ON A FINITE STATE SPACE:

For M independent lag-τ segments on an n-state model, the empirical Bellman estimator obeys

$$\sqrt{M}\big(\widehat{u}(i) - u(i)\big) \Longrightarrow \mathcal{N}(0, \sigma_i^2), \qquad \frac{\sigma_i^2}{u(i)^2} \leq C_0\, \tau n^3,$$

where C_0 depends only on the assumption constants (next slide) — **not on** n.

Under the same assumptions both $1/u(i)$ and $\mathbf{E}[T]$ can be exponentially large in n. Our bound is **polylogarithmic in the return period**.

A direct Monte Carlo estimate of $u(i)$ would require **exponentially many, exponentially long** (in n) trajectory samples to achieve the same relative accuracy

[X. Cheng, J. Weare, *NeurIPS* 2024]

Key assumptions and a triple well example

Take the nearest-neighbour chain with $\pi \propto e^{-V/\varepsilon}$ on $D = \{1, \dots, n-1\}$, $A = \{0\}$, $B = \{n\}$:

$$P(i, i \pm 1) = \tfrac{1}{2} \frac{\pi\big((i \pm 1)/n\big)}{\pi\big((i \pm 1)/n\big) + \pi(i/n)}.$$

1. Well-spread μ ($\mu(i) \gtrsim 1/n$): take μ uniform.

2. Connected minorizing graph (Graph with only edge probabilities $\geq c$ is irreducible): $P(i, i \pm 1) \gtrsim e^{-\mathrm{Lip}(V)/(n\varepsilon)}$ — so it holds once $\varepsilon \gtrsim 1/n$.

3. Local tails (If d is the minorizing graph distance, $S^t(i,j) \leq Ce^{-\beta d(i,j)^2}$): the example is nearest-neighbor

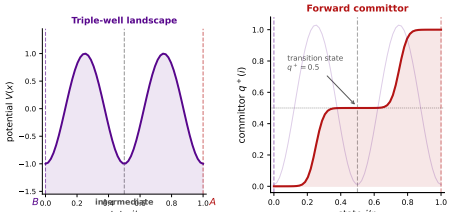

π (left) and q (right) for a triple-well V, $n = 60$.

Yet $\mathbf{E}_\nu[T]$ and $1/q(1)$ are **both exponentially large** in n: the linear system is exponentially poorly conditioned

$\mu = \pi$ **is a poor choice. It fails Assumption 1.**

Two practical lessons from the theory

Lesson 1: keep τ intermediate

Too small: few inter-state transitions; u is sensitive to poorly estimated small entries in S^τ.

Too large: The stopped MSM estimate becomes direct Monte Carlo.

Rule of thumb: use a fraction of the transition-path time.

Lesson 2: μ must cover the pathways

The sampling distribution must cover likely transition pathways.

Stationary sampling is inefficient: π puts too few samples along them.

In practice: Spread initial conditions over collective variables that separate A from B.

SSW return-period estimates from 46-day ECMWF hindcasts

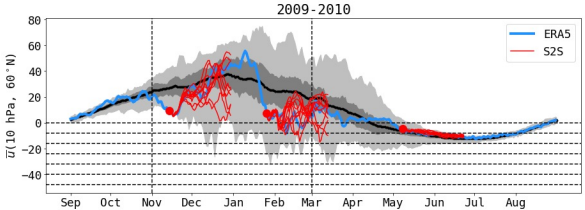

Red: S2S members. Blue: ERA5 reanalysis. Grey: historical range.

Flux counting (a direct estimate for this statistic only) and stopped MSM estimate ($\tau = 1$ day) agree with ERA5 for more common events.

The ensemble supports return-period estimates of nearly **1000 years** from 46-day trajectories.

[**J. Finkel**, E. Gerber, D.S. Abbot, J. Weare, *AGU Advances* 2023]

ECMWF S2S hindcasts: 20 winters (1996–2015), two 10-member 46-day forecasts per week — nearly **900 years** total.

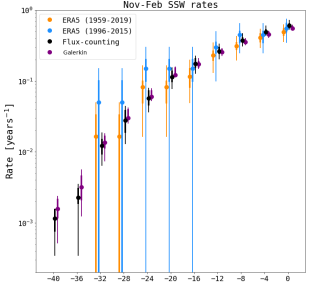

SSW rate vs. zonal-wind threshold, 10 hPa 60°N.

SSW committor and lead time

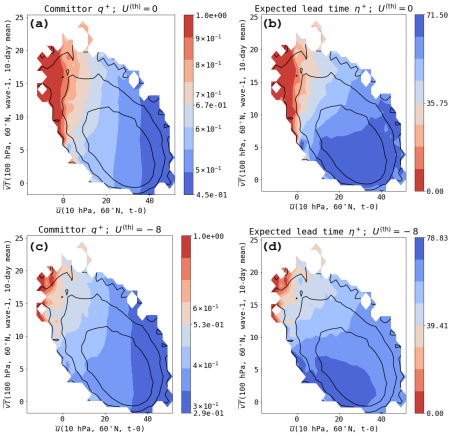

[**J. Finkel**, E. Gerber, D.S. Abbot, J. Weare, *AGU Advances* 2023]

Committor q: primarily a function of the zonal wind alone.

Expected lead time η^+: also depends on the wavenumber-1 heat flux, a measure of upward wave forcing.

Both functions are estimated from the same trajectory sample.

Insulin dimer dissociation (Dinner group): 100 ms from 94 μs of data

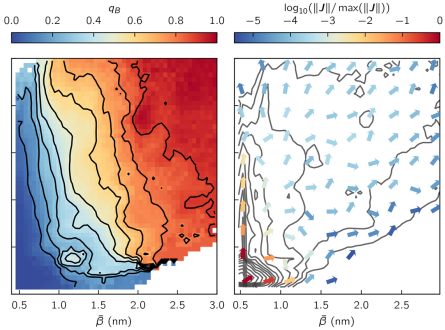

Committor q_B (left) and reactive current $\mathbf{j}^{AB}$ (right) over the mean interfacial β-strand ($\bar{\beta}$) and α-helix ($\bar{\alpha}$) separations. A: dimer, B: dissociated.

Data: 18,816 trajectories of 5 ns; lag $\tau = 2$ ns; **94** μ**s** total. Restrained MD generates initial states spanning the $\bar{\alpha}$ and $\bar{\beta}$ interface-separation coordinates.

Result: Estimated mean dissociation time $\approx$**100 ms** — roughly **1000**$\times$ the total data length.

The reactive current branches among multiple dissociation pathways; no single pathway dominates.

[K. Jeong et al., *J. Phys. Chem. B* **128**, 12728 (2024)]

Nonlinear parametrizations

Nonlinear methods parametrize either the transition operator or the value function.

Two nonlinear parametrizations of the Bellman equation

An MSM is both a parametrization of the dynamics and a parametrization of u.

For non-linear parametrizations we need to choose.

Transition model S_θ^τ (or T_θ^τ)

Parametrize a short-time transition law, $S_\theta^\tau \approx S^\tau$ (or $T_\theta^\tau \approx T^\tau$).

Generate simulated trajectories (roll-outs).

One parametrization allows estimation of multiple statistics.

Value function u_θ

Parametrize the desired value function, $u_\theta \approx u$.

Enforce Bellman consistency directly on short trajectory pairs — no rollouts, and no model of the full transition law.

One fit per statistic.

Minimizing the Bellman residual requires two endpoints per state

Since $u = S^\tau u + r_\tau$ on D, the natural least-squares loss is

$$J(\theta) = \mathbf{E}_\mu\left[\left(S^\tau u_\theta + r_\tau - u_\theta\right)^2\right], \qquad J(\theta) = 0 \iff u_\theta \text{ solves the Bellman equation.}$$

Its gradient is a product of two conditional expectations.

$$\tfrac{1}{2}\nabla_\theta J = \mathbf{E}_\mu\left[\left(S^\tau u_\theta + r_\tau - u_\theta\right)\left(S^\tau \nabla_\theta u_\theta - \nabla_\theta u_\theta\right)\right]$$

An unbiased product requires **two conditionally independent** copies of $X_{\tau\wedge T}$ for each X_0.

We can either:

- Generate more than one conditionally independent $X_{\tau\wedge T}$ for each X_0 in our dataset [Strahan et. al. *J. Comp. Phys.* 2023], or
- Use a fixed point iteration instead of minimization.

Fixed-point iteration needs only one endpoint per state

For any $\varepsilon \in (0, 1]$ the damped (Richardson) iteration

$$\tilde{u}_{k+1} = (1 - \varepsilon)\, u_{\theta_k} + \varepsilon\big(S^\tau u_{\theta_k} + r_\tau\big) \quad \text{on } D, \qquad \tilde{u}_{k+1} = g \quad \text{on } A \cup B,$$

converges to u in the exact function space.

A parametrized implementation fits $\theta_{k+1} \approx \arg\min_\theta d(\tilde{u}_{k+1}, u_\theta)$.

For a sample pair $(X_0^\ell, X_{\tau\wedge T}^\ell)$ and corresponding accumulated reward r^ℓ,

$$Z_\ell = (1 - \varepsilon) u_{\theta_k}(X_0^\ell) + \varepsilon\big(u_{\theta_k}(X_{\tau\wedge T}^\ell) + r^\ell\big)$$

satisfies $\mathbf{E}[Z_\ell \mid X_0^\ell] = \tilde{u}_{k+1}(X_0^\ell)$.

Requirement: $\nabla_\theta d$ must be affine in $\tilde{u}_{k+1}$
Then $\nabla_\theta d(Z_\ell, u_\theta)$ averages to $\nabla_\theta d(\tilde{u}_{k+1}, u_\theta)$.

Descend $\frac{1}{|\mathcal{B}|}\sum_{\ell\in\mathcal{B}} d(Z_\ell, u_\theta)$ with θ_k frozen then increment k and repeat. [Strahan et. al. *J. Chem. Phys.* 2023]

Two examples of d, and the $\tau \to \infty$ limit

Squared $L^2(\mu)$ error

$$d = \mathbf{E}_\mu\left[(\tilde{u}_{k+1} - u_\theta)^2\right]$$

$$\nabla_\theta d = -2\,\mathbf{E}_\mu\left[(\tilde{u}_{k+1} - u_\theta)\,\nabla_\theta u_\theta\right]$$

Binary cross-entropy (committor)

$$d = -\mathbf{E}_\mu\left[\tilde{u}_{k+1}\log u_\theta + (1-\tilde{u}_{k+1})\log(1-u_\theta)\right]$$

$$\nabla_\theta d = -\,\mathbf{E}_\mu\left[\frac{\tilde{u}_{k+1} - u_\theta}{u_\theta(1-u_\theta)}\,\nabla_\theta u_\theta\right]$$

These are nonlinear in u_θ but affine in $\tilde{u}_{k+1}$.

If we take $\tau \to \infty$, $S^\tau|_D \to 0$ and the Bellman equation becomes the explicit formula for u.
With $\varepsilon = 1$, as $\tau \to \infty$ one fixed-point step becomes regression on the labels $g(X_T^\ell) + r^\ell$

[Strahan et al. *J. Chem. Phys.* 2023]

AIB_9: a neural committor from short trajectory pairs

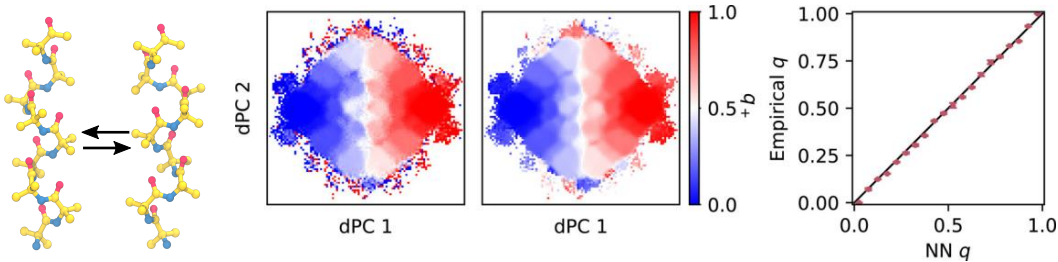

Reference and network committor on dPC 1, dPC 2; network vs. empirical at right.

System and data. AIB_9, a nine-residue α-aminoisobutyric acid peptide, switches between left- and right-handed helices on ≈ 60 ns. Data: 6,910 trajectories of 20 ns; lag $\tau = 400$ ps.

[Strahan et al. *J. Chem. Phys.* 2023]

Transition models in practice: AI weather emulators

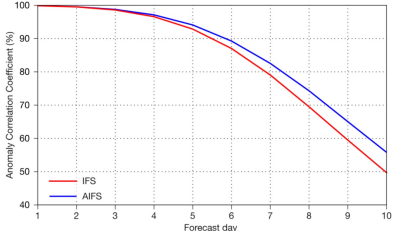

DeepMind ECMWF Google HUAWEI NVIDIA

An emulator approximates a one-step transition map F_θ or conditional law P_θ, fitted by a regression or likelihood loss:

$$L(\theta) = \sum_{\ell} \ell_{\text{forecast}}\left(F_\theta(X_0^\ell), X_\tau^\ell\right).$$

defining an approximate lag-τ transfer operator.

Rollouts estimate multiple statistics.

Rare-event simulation with incomplete transition data

Both parametrizations require training data that cover the relevant transition paths.

Three AI weather emulator training sets

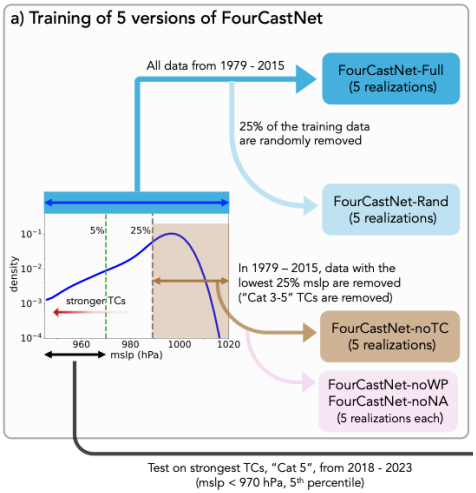

[Y.Q. Sun et al., *PNAS* 2025]

ERA5 training period: 1979–2015.
Full — every sample.

Rand — random removal matched to noTC in sample count and seasonal distribution; Category 3–5 samples retained.

noTC — the 25% with the lowest tropical sea-level pressure removed: Category 3–5 cyclones.

Test set: twenty Category 5 storms from 2018–2023; noTC contains no Category 3–5 training samples.

FourCastNet underestimates excluded Category 5 cyclones

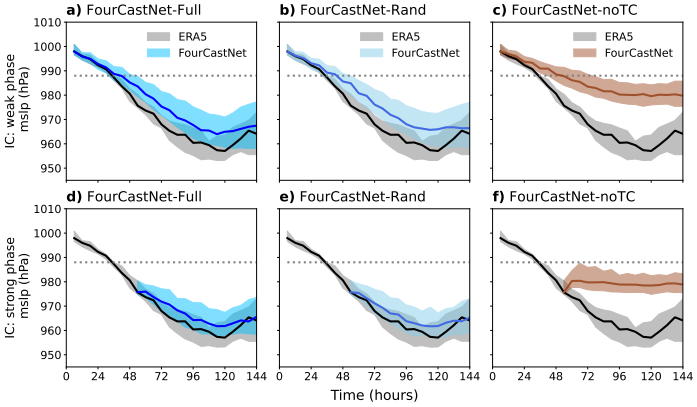

[Y.Q. Sun et al., *PNAS* 2025]

FourCastNet-noTC, trained without samples containing Category 3–5 cyclones, leaves central pressure too high and misses Category 5 intensification.

Top and bottom rows: initialized one day before and one day after crossing 988 hPa.

Nonetheless, emulators can be useful for generating extreme training examples.

FourCastNet underestimates excluded Category 5 cyclones

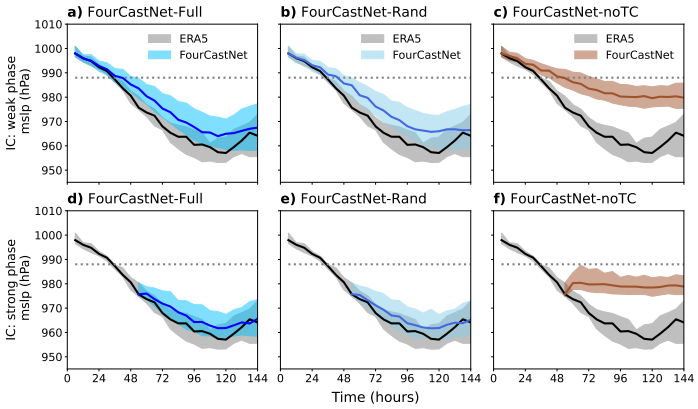

[Y.Q. Sun et al., *PNAS* 2025]

FourCastNet-noTC, trained without samples containing Category 3–5 cyclones, leaves central pressure too high and misses Category 5 intensification.

Top and bottom rows: initialized one day before and one day after crossing 988 hPa.

Nonetheless, emulators can be useful for generating extreme training examples.

Rare Event Simulation reallocates weighted trajectories

A **walker** is one physical-model state with statistical weight w_i; the weights satisfy $\sum_i w_i = 1$.

1. PROPAGATE

Advance an ensemble of weighted walkers using the **physical model**.

2. RANK

Score the walker at position x by a prediction $\theta(x)$ of the rare event probability from initial condition x.

3. RESAMPLE

Branch high-scoring walkers and prune low-scoring walkers while correcting the bias with weights.

RES reallocates physical-model trajectories; weights correct the resampling bias. The score function can come from transition-model forecasts or a directly fitted value function.

AI+RES estimates the ranking coordinate with an emulator

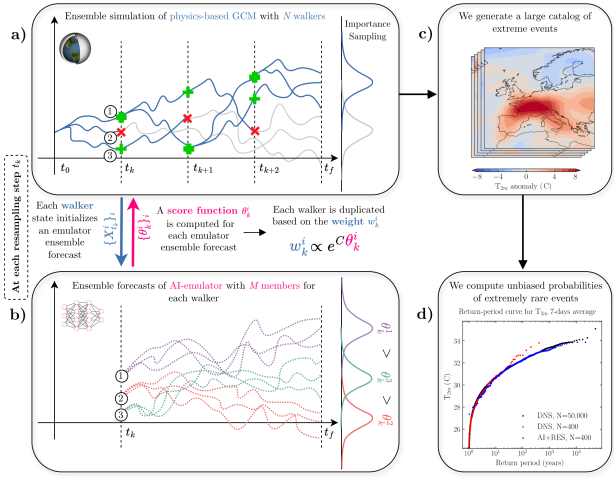

$A_7(t_f)$: 7-day mean 2 m temperature over a 3×3 PlaSim grid-point region centered on France or Chicago; t_f = August 1.

For each physical walker x_i, run $M = 100$ emulator forecasts and use the mean final heatwave index to rank walkers:

$$\widehat{\theta}(x_i) = M^{-1} \sum_{m=1}^{M} \widehat{A}_7^{(m)}(t_f \mid x_i).$$

$A_7^{(m)}(t_f \mid x_i)$ is the average temperature over France (or Chicago) averaged over a week in August .

For the physics model we use PlaSim, an intermediate complexity global circulation model.

[A. Lancelin, A. Wikner et al., PRL 2026]

AI+RES matches the 50,000-member PlaSim return-period curve

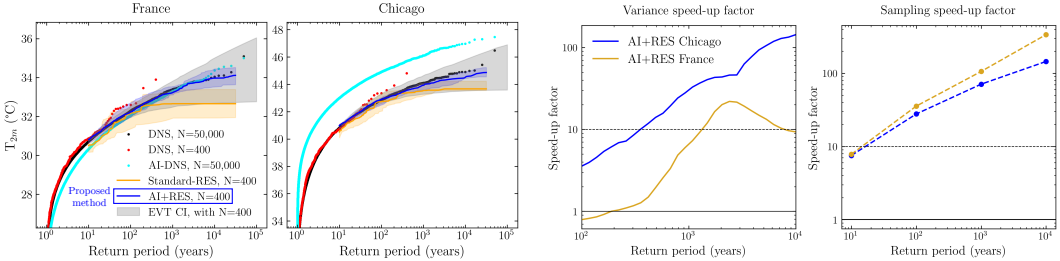

Left: 7-day mean temperature vs. return period; EVT CL: extreme-value confidence limits. Right: PlaSim walker-budget ratio of DNS to AI+RES at matched variance or expected exceedance count.

AI+RES with **400 walkers** matches the 50,000-member reference. At the same PlaSim budget, direct simulation reaches $\sim$50 years; emulator-only estimates are biased at decadal return periods.

For **Chicago**, both speed-up factors exceed **100**$\times$ near 10^4-year return periods. The variance gain is smaller for France.

[A. Lancelin, A. Wikner et al., PRL 2026]

Thanks!

Collaborators

D.S. Abbot, S. Allaw, F. Bouchet, A. Chattopadhyay, X. Cheng, A.R. Dinner, L. Dubus, J. Finkel, E.P. Gerber, S.C. Guo, P. Hassanzadeh, K. Jeong, A. Lancelin, C. Le Priol, C. Lorpaiboon, E. Persson, G. Stadler, J. Strahan, Y.Q. Sun, A. Wikner, M. Zand

Funding

Molecular Dynamics review article now available

J. Weare, A.R. Dinner, *Machine learning kinetics from molecular dynamics data*, `https://arxiv.org/abs/2609.17736`