

Quantum Field Theory as Generative Models

Emergent Gravity as Latent Variable Model

Everything on a 4×4 Lattice

Ying Nian Wu

UCLA Department of Statistics and Data Science

Tutorial for Bin Yu's group, UC Berkeley

Abstract

This tutorial presents quantum field theory (QFT) and emergent gravity as three generative models that statisticians and ML researchers already know. **Model 1 (RNN)**: on the Lorentzian side, a quantum field is a recurrent neural network — one hidden vector rotating under a unitary Hamiltonian, with an interactive Born output layer. **Model 2 (EBM)**: after Wick rotation, the path integral becomes an energy-based model $p(\phi) = e^{-S_E}/Z_E$ — an exponential family where correlators are moments, connected correlators are cumulants, and the full statistical toolkit applies. **Model 3 (Latent variable model)**: the Hubbard-Stratonovich transformation introduces a latent tensor $\sigma_{\mu\nu}$ that makes the model conditionally Gaussian — like a restricted Boltzmann machine where both conditionals are tractable. Einstein's equation emerges as the EM fixed point; quantum gravity is the $1/\sqrt{N}$ fluctuation from the Gibbs sampler.

Every concept is exhibited on a single concrete example: a scalar field on a 4×4 periodic lattice (16 field variables). The Euclidean path integral is a 16-dimensional Gaussian. The propagator is a 16×16 matrix. The mass is a decay rate. No functional analysis, no infinite-dimensional measures: just finite-dimensional linear algebra, probability, and statistics.

Contents

1	The Setup: 16 Numbers on a Grid	3
1.1	The lattice	3
1.2	Derivatives as differences	4
2	The Lagrangian: Kinetic Energy Minus Potential Energy	4
2.1	The continuum Lagrangian	4
2.2	The lattice Lagrangian	4
2.3	The action as a quadratic form	5
3	Lorentz Transformations: Rotations that Mix Space and Time	5
3.1	Why space and time can mix	5
3.2	Lorentz invariance of the action	5
3.3	What Lorentz invariance means for the lattice	5
4	The Path Integral: Weighting Every Configuration	6
4.1	The Lorentzian path integral (toward Model 1)	6
4.2	The problem with the Lorentzian path integral	6
4.3	What physics can we extract from the Lorentzian side?	6

5	Wick Rotation: From Phase to Probability (Model 1 \rightarrow Model 2)	7
5.1	The substitution $t \rightarrow -i\tau$	7
5.2	The side-by-side comparison	7
6	The Euclidean Path Integral: A 16-Dimensional Gaussian	8
6.1	Model 2: the energy-based model (exponential family)	8
6.2	The free field: a 16-dimensional Gaussian	8
6.3	The covariance = the propagator	9
6.4	The log-normaliser and the cumulant generating function	9
6.5	FFT-Gaussian sampling	9
7	What Physics Can We Extract?	9
7.1	The propagator: pole vs. decay rate	9
7.2	The Euclidean correlator on the 4×4 lattice	10
7.3	What the Euclidean side can extract	10
7.4	What the Euclidean side cannot extract	10
8	The Transfer Matrix: Connecting the Two Sides	10
8.1	Construction	10
8.2	Two functions of one operator	11
8.3	The path integral as a derived identity	11
9	Interactions: The ϕ^4 Theory	11
9.1	Adding ϕ^4	11
9.2	What changes	11
9.3	The source J and generating functionals	12
9.4	Feynman diagrams as Taylor expansion	12
9.5	Renormalization: the lattice-first view	12
10	The Hubbard-Stratonovich Transformation (Model 2 \rightarrow Model 3)	13
10.1	The idea: decouple a quartic into two quadratics	13
10.2	The latent variable model (Model 3)	13
11	The T^2 Interaction: Emergent Gravity	13
11.1	From ϕ^4 to $T_{\mu\nu}T^{\mu\nu}$	13
11.2	The joint distribution	14
12	The N-Copy Construction: EM and Gibbs with Explicit N	14
12.1	N copies of the scalar field	14
12.2	The Gibbs sampler, the EM algorithm, Einstein's equation, and the hydrodynamic mode	14
12.3	Saddle point = LLN; fluctuation = CLT	15
13	Polishing: From Tensor to Metric, and What Clocks Measure	16
13.1	What $\sigma_{\mu\nu}$ is and what it is not	16
13.2	The weak-field polishing: $\sigma \rightarrow g$	16
13.3	Clocks, rulers, and the equivalence principle	17
13.4	The saddle-point vs. fluctuation distinction	18
13.5	Diffeomorphism invariance: a consequence of the polishing	18
13.6	At the quantum level, geometry does not exist	18
13.7	Strong-field departure from Einstein's equation	19
13.8	The effective action after polishing	19

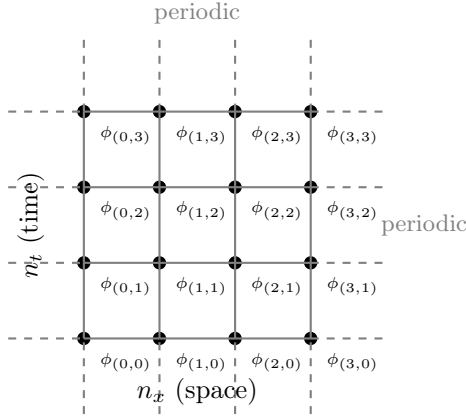
14 The Heat Kernel Expansion: Why GR Is the Leading Term	19
14.1 The Seeley-DeWitt expansion	20
14.2 The tower: each additional curvature power is suppressed	20
14.3 The complete pipeline	20
15 The Hamiltonian: Model 1 (The RNN)	20
15.1 The classical Hamiltonian	21
15.2 Quantization: operators and commutators	21
15.3 Diagonalization in momentum space	21
15.4 Creation and annihilation operators	21
15.5 Fock space	22
15.6 The Schrödinger equation: the RNN	22
15.7 From Hamiltonian back to path integral	22
15.8 Why the path integral gives physical answers	22
15.9 The interacting Hamiltonian: a polynomial of creation operators	22
15.10 The Dyson series: another route to Feynman diagrams	23
15.11 Rendering stability from the Hamiltonian side	23
15.12 Observables and degenerate observables	23
15.13 Partial access and thermalization	24
15.14 Composite operators and macroscopic averages	24
16 Summary: Three Generative Models on 16 Numbers	25
A The Stress-Energy Tensor: Three Perspectives	26
A.1 Perspective 1: Noether current from spacetime symmetry	26
A.2 Perspective 2: Hilbert definition (variation of the matter action)	26
A.3 Perspective 3: T^2 from the Wilsonian RG flow	27
B The Heat Kernel Expansion: A Closer Look	27
B.1 How the $\text{Tr} \log$ arises	27
B.2 The Schwinger proper-time representation	27
B.3 Why it is called the “heat kernel”	28
B.4 The asymptotic expansion	28
B.5 The unpolished expansion ($K[\sigma]$) vs. the polished expansion ($K[g]$)	28
B.6 Example 1: gravity ($K = -\square_g + m^2$, spin 0)	29
B.7 Example 2: gauge theory ($K = -D_\mu^2 + m^2$, spin 1)	29
B.8 The universal pattern	30

1 The Setup: 16 Numbers on a Grid

1.1 The lattice

Consider a two-dimensional grid with $N_x = 4$ columns (“space”) and $N_t = 4$ rows (“time”), giving $V = N_x \times N_t = 16$ sites. Impose **periodic boundary conditions** in both directions: the right edge wraps to the left, the top wraps to the bottom. Topologically, the grid is a torus.

Label each site by $\mathbf{n} = (n_x, n_t)$ with $n_x \in \{0, 1, 2, 3\}$ and $n_t \in \{0, 1, 2, 3\}$. At each site, place a single real number $\phi_{\mathbf{n}} \in \mathbb{R}$. The field configuration is a vector $\phi = (\phi_{(0,0)}, \phi_{(1,0)}, \dots, \phi_{(3,3)})^T \in \mathbb{R}^{16}$.



That is the entire setup. Everything in this tutorial — special relativity, quantum mechanics, quantum field theory, gravity — will be built on these 16 numbers.

1.2 Derivatives as differences

On the continuum, the spatial derivative is $\partial_x \phi(x, t) = \lim_{a \rightarrow 0} [\phi(x + a, t) - \phi(x, t)]/a$. On the lattice with spacing a , the derivative becomes a **finite difference**:

$$(\Delta_x \phi)_{\mathbf{n}} \equiv \frac{\phi_{(n_x+1, n_t)} - \phi_{(n_x, n_t)}}{a}. \quad (1.1)$$

Similarly for the time direction: $(\Delta_t \phi)_{\mathbf{n}} = [\phi_{(n_x, n_t+1)} - \phi_{(n_x, n_t)}]/a$.

Because of periodic boundary conditions, $\phi_{(4, n_t)} \equiv \phi_{(0, n_t)}$ and $\phi_{(n_x, 4)} \equiv \phi_{(n_x, 0)}$.

The difference operator is a matrix. For $N_x = 4$ with periodic BC, the forward-difference operator in the x -direction acts on the vector $(\phi_0, \phi_1, \phi_2, \phi_3)^T$ as:

$$\Delta_x = \frac{1}{a} \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{pmatrix}. \quad (1.2)$$

The bottom-left 1 is the periodic boundary condition wrapping site 3 back to site 0. The discrete Laplacian is $-\Delta_x^T \Delta_x$, which is the familiar circulant matrix.

2 The Lagrangian: Kinetic Energy Minus Potential Energy

2.1 The continuum Lagrangian

The Lagrangian density for a free scalar field of mass m in 1+1 dimensions is:

$$\mathcal{L} = \frac{1}{2}(\partial_t \phi)^2 - \frac{1}{2}(\partial_x \phi)^2 - \frac{1}{2}m^2 \phi^2. \quad (2.1)$$

This is “kinetic minus potential”: the time derivative $(\partial_t \phi)^2$ is the kinetic energy; the spatial gradient $(\partial_x \phi)^2$ and the mass term $m^2 \phi^2$ are the potential energy.

Notice the **relative sign**: the time and space derivatives enter with *opposite* signs. This minus sign is the signature of **Lorentzian** (Minkowski) spacetime, and it is the single most important sign in all of physics.

2.2 The lattice Lagrangian

On the 4×4 lattice, replace derivatives by differences:

$$\mathcal{L}_{\mathbf{n}} = \frac{1}{2} \left(\frac{\phi_{(n_x, n_t+1)} - \phi_{(n_x, n_t)}}{a} \right)^2 - \frac{1}{2} \left(\frac{\phi_{(n_x+1, n_t)} - \phi_{(n_x, n_t)}}{a} \right)^2 - \frac{1}{2} m^2 \phi_{\mathbf{n}}^2. \quad (2.2)$$

The action is the sum over all sites:

$$S_M[\phi] = a^2 \sum_{\mathbf{n}} \mathcal{L}_{\mathbf{n}} = \frac{1}{2} \vec{\phi}^T K_M \vec{\phi}, \quad (2.3)$$

where K_M is a 16×16 real symmetric matrix encoding the time differences (with $+$ sign), spatial differences (with $-$ sign), and the mass term.

2.3 The action as a quadratic form

In momentum space (discrete Fourier transform on the periodic lattice), K_M diagonalises. The eigenvalue at momentum $\mathbf{k} = (k_x, k_t)$ is:

$$(K_M)_{\mathbf{k}} = \hat{\omega}^2(k_t) - \hat{k}^2(k_x) - m^2, \quad (2.4)$$

where the lattice frequencies are $\hat{\omega}(k_t) = \frac{2}{a} \sin \frac{k_t a}{2}$ and $\hat{k}(k_x) = \frac{2}{a} \sin \frac{k_x a}{2}$.

The crucial feature: $(K_M)_{\mathbf{k}}$ is **indefinite**. For some modes (small k_t , large k_x), it is negative; for others, positive. This indefiniteness is the lattice manifestation of Lorentzian signature.

3 Lorentz Transformations: Rotations that Mix Space and Time

3.1 Why space and time can mix

In Newtonian mechanics, space and time are separate: a rotation mixes x and y , but never x and t . Special relativity's key insight is that space and time *can* mix, through **Lorentz boosts**.

A Lorentz boost with velocity v (in units $c = 1$) transforms:

$$\begin{pmatrix} t' \\ x' \end{pmatrix} = \begin{pmatrix} \cosh \eta & -\sinh \eta \\ -\sinh \eta & \cosh \eta \end{pmatrix} \begin{pmatrix} t \\ x \end{pmatrix}, \quad \tanh \eta = v. \quad (3.1)$$

This looks like a rotation, but with cosh and sinh instead of cos and sin, and a minus sign instead of a plus. It is a *hyperbolic* rotation.

3.2 Lorentz invariance of the action

The Lagrangian (2.1) is invariant under Lorentz boosts because $(\partial_t \phi)^2 - (\partial_x \phi)^2$ is invariant. To see this, note that in Minkowski spacetime with metric $\eta_{\mu\nu} = \text{diag}(+1, -1)$:

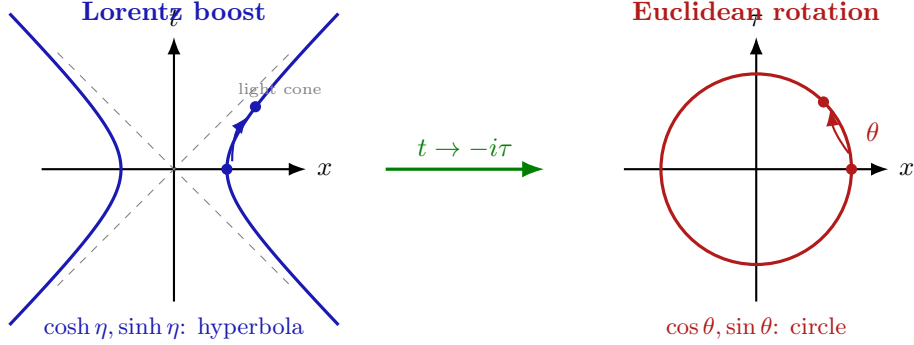
$$(\partial_t \phi)^2 - (\partial_x \phi)^2 = \eta^{\mu\nu} (\partial_\mu \phi) (\partial_\nu \phi) = (\partial_\mu \phi) (\partial^\mu \phi). \quad (3.2)$$

This is a scalar under Lorentz transformations: it has the same value in every reference frame. So the action $S_M = \int dt dx \mathcal{L}$ is Lorentz-invariant: different observers agree on its value.

In momentum space, the combination $\hat{\omega}^2 - \hat{k}^2$ in (2.4) is the lattice version of the Lorentz scalar $k_\mu k^\mu = \omega^2 - k^2$. The mass-shell condition $\omega^2 - k^2 = m^2$ (or $(K_M)_{\mathbf{k}} = 0$) determines which modes are “on shell” — the physical particles.

3.3 What Lorentz invariance means for the lattice

On the 4×4 lattice, exact Lorentz invariance is broken (the lattice distinguishes axes from diagonals). But at long wavelengths ($|k_\mu a| \ll 1$), the lattice dispersion relation $\hat{\omega}^2 - \hat{k}^2 \approx \omega^2 - k^2 + O(a^2)$ approaches the continuum one. Lorentz symmetry is an *emergent* symmetry, accurate to $O(a^2)$.



4 The Path Integral: Weighting Every Configuration

4.1 The Lorentzian path integral (toward Model 1)

The fundamental object of quantum field theory is the **path integral** (partition function):

$$Z_M = \int_{\mathbb{R}^{16}} d^{16}\phi e^{iS_M[\phi]} = \int_{\mathbb{R}^{16}} d^{16}\phi e^{i\vec{\phi}^T K_M \vec{\phi}/2}. \quad (4.1)$$

This is a 16-dimensional integral over all possible field configurations $\vec{\phi} \in \mathbb{R}^{16}$.

Each configuration is weighted by $e^{iS_M[\phi]}$. This weight is a *complex phase*: $|e^{iS_M}| = 1$ for every $\vec{\phi}$. It does not decay. Every configuration contributes with equal absolute magnitude.

4.2 The problem with the Lorentzian path integral

Since $|e^{iS_M}| = 1$, the integrand does not decay as $|\vec{\phi}| \rightarrow \infty$. The integral does not converge absolutely. It converges (when it converges at all) only by cancellation of oscillations — positive and negative contributions cancel, leaving a finite result.

Moreover, e^{iS_M} is complex, so $d^{16}\phi e^{iS_M}$ is *not* a probability measure. We cannot sample from it, compute expectations, or apply statistical methods.

4.3 What physics can we extract from the Lorentzian side?

Despite these difficulties, the Lorentzian path integral defines the physical theory. The key objects:

Correlators (moments of the “distribution”):

$$\langle \phi_{\mathbf{n}} \phi_{\mathbf{m}} \rangle_M = \frac{1}{Z_M} \int d^{16}\phi \phi_{\mathbf{n}} \phi_{\mathbf{m}} e^{iS_M}. \quad (4.2)$$

The propagator (two-point correlator in momentum space):

$$\tilde{G}_M(\mathbf{k}) = \frac{1}{(K_M)_{\mathbf{k}} + i\epsilon} = \frac{1}{\hat{\omega}^2 - \hat{k}^2 - m^2 + i\epsilon}. \quad (4.3)$$

The $i\epsilon$ (a small positive imaginary part) is the Feynman prescription: it moves the poles off the real axis and gives the integral a definite meaning.

Mass as a pole: the propagator has a pole (diverges) when $(K_M)_k = 0$, i.e., when $\hat{\omega}^2 = \hat{k}^2 + m^2$. This is the **dispersion relation**: it tells you the energy ω of a particle with momentum k and mass m . A particle is a pole in the propagator.

5 Wick Rotation: From Phase to Probability (Model 1 \rightarrow Model 2)

5.1 The substitution $t \rightarrow -i\tau$

Replace real time t by imaginary time $\tau = it$: $t = -i\tau$. Under this substitution:

$$\partial_t \phi = \frac{\partial \phi}{\partial t} = \frac{\partial \phi}{\partial(-i\tau)} = i \frac{\partial \phi}{\partial \tau} = i \partial_\tau \phi, \quad (5.1)$$

so $(\partial_t \phi)^2 = (i \partial_\tau \phi)^2 = -(\partial_\tau \phi)^2$. The time kinetic term flips sign.

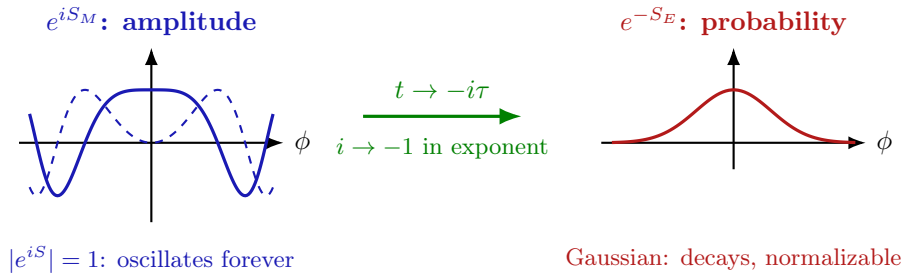
The Lagrangian becomes:

$$\mathcal{L}_M = \frac{1}{2}(\partial_t \phi)^2 - \frac{1}{2}(\partial_x \phi)^2 - \frac{1}{2}m^2 \phi^2 \rightarrow -\frac{1}{2}(\partial_\tau \phi)^2 - \frac{1}{2}(\partial_x \phi)^2 - \frac{1}{2}m^2 \phi^2 = -\mathcal{L}_E, \quad (5.2)$$

where the **Euclidean Lagrangian** is:

$$\boxed{\mathcal{L}_E = \frac{1}{2}(\partial_\tau \phi)^2 + \frac{1}{2}(\partial_x \phi)^2 + \frac{1}{2}m^2 \phi^2.} \quad (5.3)$$

All three terms now have the *same* sign. The crucial minus between time and space is gone.



5.2 The side-by-side comparison

	Lorentzian (S_M)	Euclidean (S_E)
Lagrangian	$\frac{1}{2}(\partial_t\phi)^2 - \frac{1}{2}(\partial_x\phi)^2 - \frac{1}{2}m^2\phi^2$	$\frac{1}{2}(\partial_\tau\phi)^2 + \frac{1}{2}(\partial_x\phi)^2 + \frac{1}{2}m^2\phi^2$
Sign pattern	+ - - (indefinite)	+ + + (positive definite)
Momentum space	$(K_M)_k = \hat{\omega}^2 - \hat{k}^2 - m^2$	$(K_E)_k = \hat{\omega}^2 + \hat{k}^2 + m^2$
Eigenvalues of K	Some positive, some negative	All positive
Weight $e^{?S}$	e^{iS_M} : complex phase, $ e^{iS_M} = 1$	e^{-S_E} : real positive, decays
Each mode	$e^{i(K_M)_k \hat{\phi} ^2/2}$: oscillates	$e^{-(K_E)_k \hat{\phi} ^2/2}$: Gaussian bell
Product over modes	\prod_k (phase): not a probability	\prod_k (bell curve): $p(\phi) = \mathcal{N}(0, K_E^{-1})$
Symmetry	Lorentz: $\hat{\omega}^2 - \hat{k}^2$ invariant	Euclidean rotation: $\hat{\omega}^2 + \hat{k}^2$ invariant

Where Wick rotation acts: the single sign flip $(\partial_t\phi)^2 \rightarrow -(\partial_\tau\phi)^2$ changes the operator from indefinite ($\square = \partial_t^2 - \partial_x^2$, the wave operator) to positive definite ($-\Delta = \partial_\tau^2 + \partial_x^2$, the Laplacian):

$$\underbrace{\partial_t^2 - \partial_x^2}_{\text{wave operator (indefinite)}} \xrightarrow{t \rightarrow -i\tau} \underbrace{\partial_\tau^2 + \partial_x^2}_{\text{Laplacian (positive definite)}} \quad (5.4)$$

Symmetry change: the invariance group changes from Lorentz transformations ($\cosh \eta, \sinh \eta$: hyperbolic rotations mixing t and x) to Euclidean rotations ($\cos \theta, \sin \theta$: ordinary rotations mixing τ and x). On the lattice, the Euclidean rotation group is the discrete symmetry that exchanges the τ and x axes — visible as the symmetry of K_E under $k_0 \leftrightarrow k_1$.

6 The Euclidean Path Integral: A 16-Dimensional Gaussian

6.1 Model 2: the energy-based model (exponential family)

The Euclidean path integral defines a probability distribution:

$$\boxed{p(\phi) = \frac{1}{Z_E} e^{-S_E[\phi]}, \quad Z_E = \int_{\mathbb{R}^{16}} d^{16}\phi e^{-S_E[\phi]}.} \quad (6.1)$$

This is an **energy-based model** (EBM): the action $S_E[\phi]$ is the energy, and configurations with lower energy have higher probability. It is also an **exponential family**: the log-density is linear in the “sufficient statistics” (here, ϕ^2 and $(\Delta\phi)^2$), with coupling constants m^2 and $1/a^2$ as the *natural parameters*.

The negative log-density is the action: $-\ln p(\phi) = S_E[\phi] + \ln Z_E$.

Compare with the Lorentzian side, which defines **Model 1: the RNN**. The hidden state $|h_t\rangle = e^{-iHt}|h_0\rangle$ rotates unitarily under the Hamiltonian — a recurrent neural network with unitary weight matrix. The Born rule $P(\lambda) = \|P_\lambda h_t\|^2$ produces Born-layer outputs. No probability distribution on field configurations; no sampling.

The Wick rotation is the passage from Model 1 (RNN, complex amplitudes, no sampling) to Model 2 (EBM, real positive weights, MCMC sampling).

6.2 The free field: a 16-dimensional Gaussian

For the free theory ($\lambda = 0$), the action is quadratic: $S_E = \frac{1}{2}\vec{\phi}^T K_E \vec{\phi}$. The EBM is a Gaussian:

$$p(\phi) = \frac{\sqrt{\det K_E}}{(2\pi)^8} e^{-\vec{\phi}^T K_E \vec{\phi}/2} = \mathcal{N}(\vec{0}, \Sigma), \quad \Sigma = K_E^{-1}. \quad (6.2)$$

The **precision matrix** (inverse covariance) is $K_E = -\Delta_{\text{lat}} + m^2$ (16×16 , positive definite). The **covariance matrix** is $\Sigma = K_E^{-1}$.

In momentum space, the distribution factorises into 16 independent Gaussians:

$$p(\hat{\phi}) = \prod_{\mathbf{k}} \frac{1}{\sqrt{2\pi/(K_E)_{\mathbf{k}}}} e^{-(K_E)_{\mathbf{k}}|\hat{\phi}_{\mathbf{k}}|^2/2}, \quad (K_E)_{\mathbf{k}} = \hat{\omega}^2 + \hat{k}^2 + m^2 > 0. \quad (6.3)$$

Each mode is a bell curve with variance $1/(K_E)_{\mathbf{k}}$: low-momentum modes fluctuate strongly; high-momentum modes are suppressed.

6.3 The covariance = the propagator

The two-point connected correlator is the covariance:

$$G_E(\mathbf{n}, \mathbf{m}) = \langle \phi_{\mathbf{n}} \phi_{\mathbf{m}} \rangle_c = \mathbb{E}[\phi_{\mathbf{n}} \phi_{\mathbf{m}}] - \mathbb{E}[\phi_{\mathbf{n}}] \mathbb{E}[\phi_{\mathbf{m}}] = (K_E^{-1})_{\mathbf{nm}}. \quad (6.4)$$

The propagator IS the covariance matrix. In momentum space: $\tilde{G}_E(\mathbf{k}) = 1/(\hat{\omega}^2 + \hat{k}^2 + m^2)$.

6.4 The log-normaliser and the cumulant generating function

The log-partition function (the *log-normaliser*):

$$\ln Z_E = 8 \ln(2\pi) - \frac{1}{2} \ln \det K_E = 8 \ln(2\pi) - \frac{1}{2} \sum_{\mathbf{k}} \ln(K_E)_{\mathbf{k}}. \quad (6.5)$$

This is a sum of 16 logarithms — the **Tr log**. It is also the **cumulant generating function** (CGF): add a source \vec{J} and define $W[J] = \ln Z_E[J]$. Then: $\partial W / \partial J_{\mathbf{n}}|_0 = \langle \phi_{\mathbf{n}} \rangle$ (1st cumulant = mean), $\partial^2 W / \partial J_{\mathbf{n}} \partial J_{\mathbf{m}}|_0 = \langle \phi_{\mathbf{n}} \phi_{\mathbf{m}} \rangle_c$ (2nd cumulant = covariance).

The determinant of the covariance matrix: $\det \Sigma = \det K_E^{-1} = 1 / \det K_E = \prod_{\mathbf{k}} 1/(K_E)_{\mathbf{k}}$. This controls the “volume” of the Gaussian ellipsoid in 16 dimensions.

6.5 FFT-Gaussian sampling

To sample from $p(\phi) = \mathcal{N}(0, K_E^{-1})$:

1. Generate $\eta_{\mathbf{k}} \sim \mathcal{N}(0, 1)$ for each of the 16 modes.
2. Scale: $\hat{\phi}_{\mathbf{k}} = \eta_{\mathbf{k}} / \sqrt{(K_E)_{\mathbf{k}}}$.
3. Inverse FFT: $\phi_{\mathbf{n}} = \text{FFT}^{-1}[\hat{\phi}_{\mathbf{k}}]$.

Cost: $O(V \log V)$. Each sample is an independent draw from the 16-dimensional Gaussian — no MCMC, no rejection.

7 What Physics Can We Extract?

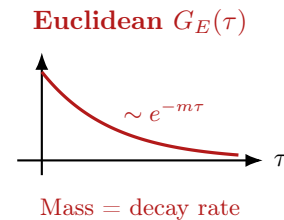
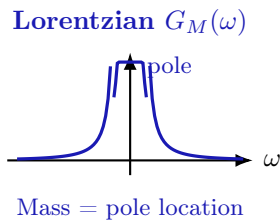
7.1 The propagator: pole vs. decay rate

The same mass m appears in two guises:

	Lorentzian	Euclidean
Propagator	$\frac{1}{\hat{\omega}^2 - \hat{k}^2 - m^2 + i\epsilon}$	$\frac{1}{\hat{\omega}^2 + \hat{k}^2 + m^2}$
Singularity	Pole at $\hat{\omega}^2 = \hat{k}^2 + m^2$	None (denominator > 0)
Position space	$G_M(t) \sim e^{-imt}$ (oscillates)	$G_E(\tau) \sim e^{-m\tau}$ (decays)
Where is m ?	Pole location	Decay rate

Lorentzian: the denominator can vanish, creating a pole. A particle of mass m is a singularity in the propagator.

Euclidean: the denominator is always positive — no singularity. Instead, Fourier-transform to position space and fit the exponential decay: $G_E(\tau) \sim e^{-m\tau}$ at large τ . The mass is the decay rate. This is how lattice QFT measures masses.



7.2 The Euclidean correlator on the 4×4 lattice

Set $\mathbf{k} = 0$ (zero spatial momentum) and compute:

$$G_E(\tau) = \sum_{\omega} \frac{e^{-i\omega\tau}}{\hat{\omega}^2 + m^2} = \sum_{n=0}^3 \frac{e^{-2\pi i n \tau / (4a)}}{\frac{4}{a^2} \sin^2(\pi n / 4) + m^2}. \quad (7.1)$$

This is a sum of 4 terms (one per Matsubara frequency on the 4-site temporal lattice). At large τ : $G_E(\tau) \sim e^{-m\tau}$, the lowest mode dominates.

7.3 What the Euclidean side can extract

From MCMC samples of $p(\phi) = e^{-S_E} / Z_E$:

Masses: exponential decay rates of correlators.

Vacuum energy: $\langle \phi^2 \rangle = \text{Tr}(K_E^{-1}) / V$.

Couplings and scattering: from higher-point correlators.

7.4 What the Euclidean side cannot extract

Real-time dynamics: $G_M(t)$ requires analytic continuation $\tau \rightarrow it$, which is ill-conditioned from noisy data.

Interference: $|e^{iS}| = 1$ phases are discarded.

Scattering amplitudes: require the S -matrix, which lives on the Lorentzian side.

8 The Transfer Matrix: Connecting the Two Sides

On the lattice, the Lorentzian and Euclidean formulations are not related by analytic continuation (which requires a complex plane). They are related by something stronger: the **transfer matrix**.

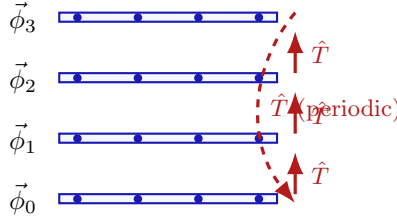
8.1 Construction

Group the 16 field variables by time slice: $\vec{\phi}_{n_t} = (\phi_{(0,n_t)}, \phi_{(1,n_t)}, \phi_{(2,n_t)}, \phi_{(3,n_t)})$. The Euclidean action decomposes as:

$$S_E = \sum_{n_t=0}^3 L_E(\vec{\phi}_{n_t+1}, \vec{\phi}_{n_t}), \quad (8.1)$$

where $L_E(\vec{\phi}', \vec{\phi})$ is the Lagrangian coupling adjacent time slices.

The **transfer matrix** \hat{T} is the N_x -site operator (here $4 \times \infty$ in the field basis) whose matrix element is:



$$Z_E = \text{Tr}[\hat{T}^4] = \int \prod d\vec{\phi}_t \prod \langle \vec{\phi}_{t+1} | \hat{T} | \vec{\phi}_t \rangle$$

$$\langle \vec{\phi}' | \hat{T} | \vec{\phi} \rangle = e^{-L_E(\vec{\phi}', \vec{\phi})}. \quad (8.2)$$

This is manifestly positive and symmetric, so \hat{T} is self-adjoint and positive. Define the Hamiltonian:

$$H = -\frac{1}{a} \ln \hat{T}, \quad E_k \geq 0. \quad (8.3)$$

8.2 Two functions of one operator

The Euclidean and Lorentzian time evolutions are:

$$\text{Euclidean: } \hat{T}^{n_t} = e^{-H n_t a}, \quad \text{Lorentzian: } e^{-i H n_t a}. \quad (8.4)$$

Both involve the *same* Hamiltonian H with the *same* spectrum $\{E_k\}$. The Euclidean correlator uses $e^{-E_k \tau}$ (decays); the Lorentzian correlator uses $e^{-i E_k t}$ (oscillates). Same energies, different encoding.

8.3 The path integral as a derived identity

The partition function is: $Z_E = \text{Tr}[\hat{T}^4] = \int d\vec{\phi}_0 \cdots d\vec{\phi}_3 \prod_{n_t} \langle \vec{\phi}_{n_t+1} | \hat{T} | \vec{\phi}_{n_t} \rangle = \int d^{16} \phi e^{-S_E}$.

The path integral is derived from the transfer matrix by inserting complete sets of states. On the lattice, this is a *theorem*, not a definition.

9 Interactions: The ϕ^4 Theory

9.1 Adding ϕ^4

The free theory ($p(\phi) \propto e^{-\frac{1}{2}\vec{\phi}^T K_E \vec{\phi}}$) is a Gaussian. Add a quartic interaction:

$$\boxed{S_E[\phi] = \frac{1}{2}\vec{\phi}^T K_E \vec{\phi} + \frac{\lambda}{4!} a^2 \sum_{\mathbf{n}} \phi_{\mathbf{n}}^4.} \quad (9.1)$$

The distribution $p(\phi) \propto e^{-S_E}$ is no longer Gaussian: ϕ^4 makes it a **non-Gaussian exponential family**. The Lorentz-invariant version is $\mathcal{L}_E = \frac{1}{2}(\partial_\mu \phi)^2 + \frac{1}{2}m^2 \phi^2 + \frac{\lambda}{4!} \phi^4$.

9.2 What changes

The propagator acquires corrections: $\tilde{G}(k) = 1/(K_E)_{\mathbf{k}} \rightarrow 1/((K_E)_{\mathbf{k}} + \Sigma(k))$, where $\Sigma(k)$ is the *self-energy* from the interaction. The physical mass shifts: $m_{\text{phys}}^2 = m^2 + \Sigma(m_{\text{phys}}^2) \neq m^2$.

Connected correlators (cumulants) of order > 2 are nonzero: the 4th cumulant $\kappa_4 \propto \lambda$, etc.

Sampling requires MCMC (Metropolis, HMC): we can no longer sample by FFT because the distribution is non-Gaussian.

9.3 The source J and generating functionals

Add a source $J_{\mathbf{n}}$ coupled linearly to ϕ :

$$Z_E[J] = \int d^{16}\phi e^{-S_E[\phi] + \vec{J} \cdot \vec{\phi}}. \quad (9.2)$$

Moments from $Z_E[J]$: $\langle \phi_{\mathbf{n}} \rangle_J = \frac{\partial \ln Z_E}{\partial J_{\mathbf{n}}}$.

Cumulants (connected correlators) from $W[J] = \ln Z_E[J]$:

$$\langle \phi_{\mathbf{n}} \phi_{\mathbf{m}} \rangle_c = \left. \frac{\partial^2 W}{\partial J_{\mathbf{n}} \partial J_{\mathbf{m}}} \right|_{J=0} = \text{Cov}(\phi_{\mathbf{n}}, \phi_{\mathbf{m}}). \quad (9.3)$$

The connected two-point correlator is the covariance: $\mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$.

9.4 Feynman diagrams as Taylor expansion

Expand $e^{-\lambda \phi^4 / (4!)}$ in powers of λ :

$$Z_E = Z_0 \left[1 - \frac{\lambda}{4!} \sum_{\mathbf{n}} \langle \phi_{\mathbf{n}}^4 \rangle_0 + O(\lambda^2) \right], \quad (9.4)$$

where $\langle \cdot \rangle_0$ denotes the Gaussian ($\lambda = 0$) expectation. Each term in the expansion is a Gaussian integral, evaluated by the factorisation $\langle \phi^4 \rangle_0 = 3(\langle \phi^2 \rangle_0)^2$ (Isserlis' theorem). The terms are represented as **Feynman diagrams**: lines for propagators G_0 , vertices for interactions λ . On the 4×4 lattice, these are literal sums over 16 sites, not formal expressions.

9.5 Renormalization: the lattice-first view

On the 4×4 lattice, everything is finite: no UV divergences, no infinities. The bare parameters (m^2, λ) are the *natural parameters* $\theta(a)$ of the exponential family. The physical observables (mass m_{phys} , scattering amplitude) are *mean parameters*.

Rendering stability. For any physical observable \mathcal{O} , computed at lattice spacing a and observed at wavelength $\lambda \gg a$:

$$\boxed{\mathcal{O}(\theta(a), a, \lambda) = \mathcal{O}_{\text{phys}} + o(a/\lambda).} \quad (9.5)$$

The physical value $\mathcal{O}_{\text{phys}}$ is independent of a ; the lattice artefacts vanish as $a/\lambda \rightarrow 0$.

Concrete example on the 4×4 lattice. The bare mass is $m^2(a)$ (the coefficient in the action). The physical mass is $m_{\text{phys}}^2 = m^2(a) + \Sigma(a)$, where $\Sigma(a)$ is the self-energy correction from the ϕ^4 interaction. As we refine the lattice ($a \rightarrow a/2$, keeping m_{phys} fixed), we must retune $m^2(a) \rightarrow m^2(a/2) = m_{\text{phys}}^2 - \Sigma(a/2)$. The natural parameter $m^2(a)$ *runs*; the mean parameter m_{phys} does not. This is renormalization — without infinities.

The corrections organise as powers of a/λ :

$$\mathcal{O}(\theta(a), a, \lambda) = \mathcal{O}_{\text{phys}} + c_1 \left(\frac{a}{\lambda}\right)^2 + c_2 \left(\frac{a}{\lambda}\right)^4 + \dots \quad (9.6)$$

(the **Symanzik expansion**). On the 4×4 lattice with $a = 1$: the first correction is c_1/λ^2 , which is small if $\lambda \gg 1$ (i.e., we look at long-wavelength physics).

10 The Hubbard-Stratonovich Transformation (Model 2 \rightarrow Model 3)

10.1 The idea: decouple a quartic into two quadratics

The ϕ^4 interaction makes the exponential family non-Gaussian. The **Hubbard-Stratonovich transformation** (HST) introduces an auxiliary (latent) variable σ that restores Gaussianity — at the cost of adding a new field.

The key identity (for a single site):

$$\boxed{e^{-g\phi^4/2} = \int_{-\infty}^{\infty} \frac{d\sigma}{\sqrt{2\pi g}} e^{-\sigma^2/(2g)} e^{-\sigma\phi^2}.} \quad (10.1)$$

Read this right to left: integrating out σ (a Gaussian integral in σ) reproduces $e^{-g\phi^4/2}$. The quartic ϕ^4 has been “decoupled” into two quadratics: $\sigma^2/(2g)$ and $\sigma\phi^2$.

10.2 The latent variable model (Model 3)

Apply the HST to the ϕ^4 theory at every lattice site:

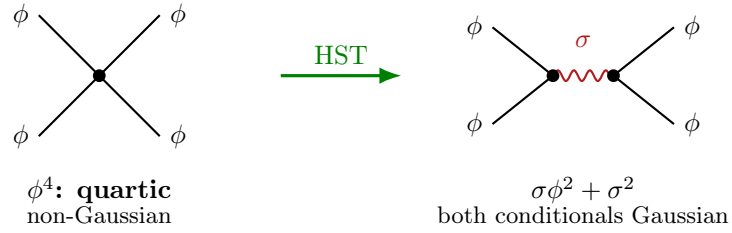
$$\boxed{p(\phi, \sigma) \propto e^{-\frac{1}{2}\vec{\phi}^T K[\sigma] \vec{\phi} - \frac{1}{2g} \sum_{\mathbf{n}} \sigma_{\mathbf{n}}^2}, \quad K[\sigma] = K_E + \text{diag}(\sigma).} \quad (10.2)$$

The joint distribution $p(\phi, \sigma)$ is a **latent variable model**. Both conditionals are Gaussian:

Decoder $p(\phi|\sigma)$: Gaussian with precision $K[\sigma] = K_E + \text{diag}(\sigma)$. The latent field σ modifies the mass at each site: $m^2 \rightarrow m^2 + \sigma_{\mathbf{n}}$.

Encoder $p(\sigma|\phi)$: Gaussian at each site independently, with mean $-g\phi_{\mathbf{n}}^2$ and variance g .

This is like a **restricted Boltzmann machine** (RBM): both conditionals are tractable, enabling Gibbs sampling with no rejection.



	VAE	RBM	HST
Decoder $p(x z)$	Neural net	Bernoulli	$\mathcal{N}(0, K[\sigma]^{-1})$
Encoder $p(z x)$	Intractable	Bernoulli	$\mathcal{N}(-g\phi^2, g)$
Both tractable?	No	Yes	Yes

11 The T^2 Interaction: Emergent Gravity

11.1 From ϕ^4 to $T_{\mu\nu}T^{\mu\nu}$

The ϕ^4 interaction couples a scalar field to itself. For gravity, we need the **stress-energy tensor** $T_{\mu\nu}[\phi]$ — a composite (bilinear) observable built from the field and its derivatives:

$$T_{\mu\nu}[\phi] = \partial_\mu\phi \partial_\nu\phi - \frac{1}{2}\delta_{\mu\nu} [(\partial\phi)^2 + m^2\phi^2]. \quad (11.1)$$

On the 4×4 lattice, $T_{\mu\nu}$ has 3 independent components ($T_{00}, T_{01} = T_{10}, T_{11}$) at each of the 16 sites.

The interaction to decouple is $T_{\mu\nu}T^{\mu\nu}$ — the contraction of the stress tensor with itself. Apply the HST: introduce a latent *symmetric tensor* field $\sigma_{\mu\nu}(\mathbf{n})$ at each site.

11.2 The joint distribution

$$p(\phi, \sigma) \propto \exp \left[-\frac{1}{2}\vec{\phi}^T K[\sigma]\vec{\phi} - \frac{1}{2\kappa^2} \sum_{\mathbf{n}} \sigma_{\mu\nu}(\mathbf{n})\sigma^{\mu\nu}(\mathbf{n}) \right], \quad (11.2)$$

where $K[\sigma] = D_0 + D_1[\sigma]$ is the precision matrix modified by σ , and κ^2 is the gravitational coupling.

Both conditionals are Gaussian:

Decoder $p(\phi|\sigma) = \mathcal{N}(0, K[\sigma]^{-1})$: the scalar field propagates on a background determined by σ . Sampled by FFT.

Encoder $p(\sigma_{\mu\nu}(\mathbf{n})|\phi) = \mathcal{N}(\kappa^2 T_{\mu\nu}[\phi](\mathbf{n}), \kappa^2)$: each component at each site is an independent Gaussian with mean proportional to the local stress tensor.

12 The N -Copy Construction: EM and Gibbs with Explicit N

12.1 N copies of the scalar field

Introduce N independent copies $\phi^{(1)}, \dots, \phi^{(N)}$, each a 16-component vector on the 4×4 lattice, all coupled to one shared latent σ :

$$\mathcal{S}[\{\phi\}, \sigma] = \sum_{a=1}^N \frac{1}{2} \vec{\phi}^{(a)T} K[\sigma] \vec{\phi}^{(a)} + \frac{1}{2\kappa^2} \sum_{\mathbf{n}} \sigma_{\mu\nu} \sigma^{\mu\nu}. \quad (12.1)$$

The first sum is over N copies of ϕ ; the second term is the prior for the **single** σ (no N). The coupling κ is fixed (not scaled with N).

(Physically, $N = N_f \times N_{\text{block}}$ where N_f is the number of field species and N_{block} is the number of fine-grid sites per coarse block on which σ is defined.)

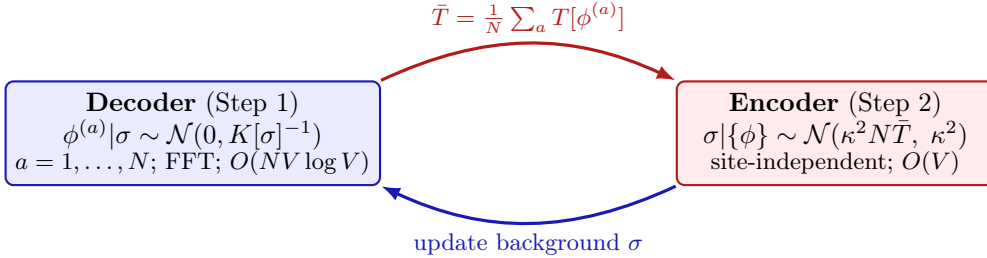
12.2 The Gibbs sampler, the EM algorithm, Einstein's equation, and the hydrodynamic mode

Step 1 (Decoder): sample $\phi^{(a)}|\sigma$ for $a = 1, \dots, N$ by FFT-Gaussian. Each copy is a 16-dimensional Gaussian with precision $K[\sigma]$. Independent, parallel. Cost: $O(NV \log V)$.

Step 2 (Encoder): sample $\sigma|\{\phi\}$ at each site:

$$\sigma_{\mu\nu}(\mathbf{n})|\{\phi\} \sim \mathcal{N}(\kappa^2 N \bar{T}_{\mu\nu}(\mathbf{n}), \kappa^2), \quad \bar{T} = \frac{1}{N} \sum_{a=1}^N T[\phi^{(a)}]. \quad (12.2)$$

Site-independent, trivially parallel. Cost: $O(V)$.



Replace the stochastic encoder with its conditional mean:

E-step: sample N_T copies from the decoder, compute $\langle T_{\mu\nu} \rangle = \frac{1}{N_T} \sum_s T[\phi^{(s)}]$.

M-step: $\sigma_{\mu\nu}^{(t+1)}(\mathbf{n}) = \kappa^2 N \langle T_{\mu\nu} \rangle(\mathbf{n})$ (deterministic, no noise). Iterate with Picard damping.

At convergence: $\bar{\sigma} = \kappa^2 N \langle T \rangle$ — Einstein's equation as the EM fixed point.

A self-consistency equation, not a dynamical law. The M-step $\bar{\sigma}_{\mu\nu} = \kappa^2 N \langle T_{\mu\nu} \rangle_{\bar{\sigma}}$ is a **self-consistency condition**: $\bar{\sigma}$ determines the background on which ϕ propagates (through $K[\bar{\sigma}]$), and $\langle T \rangle_{\bar{\sigma}}$ is the stress tensor computed on that background. Both sides must agree. This is the structure of a **gap equation** (in condensed matter) or an **equation of state** (in thermodynamics): the order parameter must be consistent with the state it induces.

Einstein's equation is *not* the dynamics of a fundamental field. The fundamental fields are the scalars $\phi^{(a)}$, whose dynamics is governed by $K[\sigma]$. Einstein's equation is the **hydrodynamic mode**: the self-consistent macroscopic description that emerges after averaging over the microscopic degrees of freedom.

The analogy is precise: the Navier-Stokes equation is not the dynamics of molecules; it is the self-consistent hydrodynamic description that emerges after averaging over molecular motions. Our scalar fields are the “molecules”; Einstein’s equation is the “Navier-Stokes” of gravity. Both are gap equations (equations of state) for collective variables, not fundamental laws for elementary constituents.

12.3 Saddle point = LLN; fluctuation = CLT

The encoder output decomposes as:

$$\sigma = \underbrace{\kappa^2 N \langle T \rangle}_{\text{LLN: classical spacetime (EM)}} + \underbrace{O(1/\sqrt{N})}_{\text{CLT: quantum gravity (Gibbs)}} \quad (12.3)$$

Law of large numbers: $\bar{T} = \frac{1}{N} \sum_a T[\phi^{(a)}] \rightarrow \langle T \rangle$ as $N \rightarrow \infty$. The saddle point is the LLN limit. Classical spacetime is the law of large numbers.

Central limit theorem: $\sqrt{N}(\bar{T} - \langle T \rangle) \rightarrow \mathcal{N}(0, \text{Var}(T))$. The $1/\sqrt{N}$ fluctuation is the CLT. Its Gaussianity is a theorem, not an assumption. Quantum gravity is the central limit theorem.

The connected σ - σ correlator (covariance):

$$\langle \sigma(\mathbf{x})\sigma(\mathbf{y}) \rangle_c = \text{Cov}(\sigma(\mathbf{x}), \sigma(\mathbf{y})) \sim \frac{1}{N} \cdot \frac{e^{-m_\sigma |\mathbf{x}-\mathbf{y}|}}{|\mathbf{x}-\mathbf{y}|}. \quad (12.4)$$

This is the **induced propagator** of σ — the graviton correlator. Its decay rate gives the screening mass m_σ ; its amplitude scales as $1/N$. At $N \rightarrow \infty$: the correlator vanishes, only the saddle (classical spacetime) survives.

Algorithm	σ -update	Fluctuation	Physics
EM ($N \rightarrow \infty$)	$\kappa^2 N \bar{T}$ (no noise)	0	Classical spacetime
N -copy Gibbs	$+\mathcal{N}(0, \kappa^2)$	$O(N^{-1/2})$	Quantum gravity

13 Polishing: From Tensor to Metric, and What Clocks Measure

The Gibbs sampler and EM give us the physics: the saddle point $\bar{\sigma} = \kappa^2 N \langle T \rangle$, the $1/\sqrt{N}$ fluctuations, the connected correlator. Now we *interpret* the result.

13.1 What $\sigma_{\mu\nu}$ is and what it is not

In the exact lattice theory, $\sigma_{\mu\nu}(\mathbf{n})$ is a symmetric tensor at each site. It has a Gaussian prior (mass term), no kinetic term, and no gauge invariance. It is *not* a metric. It does not define distances, angles, or curvature. It is a latent variable in a generative model.

13.2 The weak-field polishing: $\sigma \rightarrow g$

We now show, by explicit expansion, that $K[\sigma]$ and $K[g]$ agree to leading order — justifying the reinterpretation of σ as a metric perturbation.

Step 1: expand $K[\sigma]$ (the HST side). The HST coupling $\sigma_{\mu\nu}T^{\mu\nu}[\phi]$ modifies the quadratic action for ϕ . With $T_{\mu\nu} = \partial_\mu\phi\partial_\nu\phi - \frac{1}{2}\delta_{\mu\nu}[(\partial\phi)^2 + m^2\phi^2]$:

$$\begin{aligned}\frac{1}{2}\phi^T K[\sigma]\phi &= \frac{1}{2}(\partial_\mu\phi)^2 + \frac{1}{2}m^2\phi^2 + \sigma^{\mu\nu}T_{\mu\nu} \\ &= \frac{1}{2}(\partial_\mu\phi)^2 + \frac{1}{2}m^2\phi^2 + \sigma^{\mu\nu}\partial_\mu\phi\partial_\nu\phi - \frac{1}{2}\sigma^\alpha{}_\alpha[(\partial\phi)^2 + m^2\phi^2] \\ &= \frac{1}{2}(\delta^{\mu\nu} + 2\sigma^{\mu\nu} - \delta^{\mu\nu}\sigma^\alpha{}_\alpha)\partial_\mu\phi\partial_\nu\phi + \frac{1}{2}m^2(1 - \sigma^\alpha{}_\alpha)\phi^2.\end{aligned}\quad (13.1)$$

The effective “metric” seen by ϕ through the HST coupling is: $g_{\text{eff}}^{\mu\nu} = \delta^{\mu\nu} + 2\sigma^{\mu\nu} - \delta^{\mu\nu}\sigma^\alpha{}_\alpha + O(\sigma^2)$.

Step 2: expand $K[g]$ (the curved-space side). The Euclidean action for a scalar on a curved background $g_{\mu\nu} = \delta_{\mu\nu} + h_{\mu\nu}$ (with $|h| \ll 1$) is:

$$\frac{1}{2} \int \sqrt{g} [g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi + m^2\phi^2]. \quad (13.2)$$

Expand to first order in h . Using $g^{\mu\nu} \approx \delta^{\mu\nu} - h^{\mu\nu} + O(h^2)$ and $\sqrt{g} \approx 1 + \frac{1}{2}h^\alpha{}_\alpha + O(h^2)$:

$$\begin{aligned}\frac{1}{2}\sqrt{g} g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi + \frac{1}{2}\sqrt{g} m^2\phi^2 \\ = \frac{1}{2}(\delta^{\mu\nu} - h^{\mu\nu} + \frac{1}{2}\delta^{\mu\nu}h^\alpha{}_\alpha)\partial_\mu\phi\partial_\nu\phi + \frac{1}{2}m^2(1 + \frac{1}{2}h^\alpha{}_\alpha)\phi^2 + O(h^2).\end{aligned}\quad (13.3)$$

Step 3: match. Comparing (13.1) and (13.3) term by term:

Kinetic term: $2\sigma^{\mu\nu} - \delta^{\mu\nu}\sigma^\alpha{}_\alpha \leftrightarrow -h^{\mu\nu} + \frac{1}{2}\delta^{\mu\nu}h^\alpha{}_\alpha$.

Mass term: $-\sigma^\alpha{}_\alpha \leftrightarrow +\frac{1}{2}h^\alpha{}_\alpha$.

From the mass term: $h^\alpha{}_\alpha = -2\sigma^\alpha{}_\alpha$. Substituting into the kinetic term confirms:

$$\boxed{h_{\mu\nu} = -2\sigma_{\mu\nu} + O(\sigma^2)}. \quad (13.4)$$

The two expansions agree at $O(\sigma) = O(h)$. The HST precision $K[\sigma]$ and the curved Laplacian $K[g] = -\square_g + m^2$ are the same operator at linear order, with $h = -2\sigma$. The “polishing” is not a guess — it is forced by matching the two expansions.

At the saddle point with N copies, $\bar{\sigma}_{\mu\nu} = \kappa^2 N \langle T_{\mu\nu} \rangle$. So the metric perturbation is:

$$\boxed{h_{\mu\nu} = -2\bar{\sigma}_{\mu\nu} = -2\kappa^2 N \langle T_{\mu\nu} \rangle + O(\kappa^4)}. \quad (13.5)$$

The combination $\kappa^2 N$ is the effective gravitational coupling: more matter copies (N larger) means stronger gravity. The $O(\kappa^4)$ corrections are suppressed at weak field.

Step 4: the kinetic term as ϕ propagating on g , with κ explicit. After the identification $h = -2\sigma$, the original HST action

$$\frac{1}{2}(\delta^{\mu\nu} + 2\sigma^{\mu\nu} - \delta^{\mu\nu}\sigma^\alpha{}_\alpha)\partial_\mu\phi\partial_\nu\phi + \frac{1}{2}m^2(1 - \sigma^\alpha{}_\alpha)\phi^2 \quad (13.6)$$

is *rewritten* as:

$$\boxed{\frac{1}{2}\sqrt{g} [g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi + m^2\phi^2] + O(\kappa^4)}. \quad (13.7)$$

The scalar field ϕ “propagates on the metric $g_{\mu\nu} = \delta_{\mu\nu} - 2\kappa^2 \langle T_{\mu\nu} \rangle$ ”: the HST coupling through $T_{\mu\nu}$ is indistinguishable (at order κ^2) from propagation on curved spacetime. The $O(\kappa^4)$ corrections are present in the exact lattice theory but absent in GR — they are part of the irrelevant higher-curvature terms in the heat kernel expansion.

13.3 Clocks, rulers, and the equivalence principle

Why should we interpret σ as a metric? Because *every scalar observable built from ϕ responds to σ in the same way.*

A “clock” is a scalar composite $\mathcal{O}_{\text{clock}}[\phi]$ whose correlator gives a time interval. A “ruler” is a composite $\mathcal{O}_{\text{ruler}}[\phi]$ whose correlator gives a distance. Since every ϕ propagates on the precision matrix $K[\sigma]$, every composite built from ϕ is affected by σ through the same operator.

After polishing ($K[\sigma] \rightarrow -\square_g + m^2$), every scalar composite propagates on the *same* metric $g_{\mu\nu}$. Clocks tick at rates determined by g_{00} ; rulers measure lengths determined by g_{ij} . The geometry is not imposed from outside: it is the **emergent interpretation of what scalar-field instruments measure**.

This is the **equivalence principle**: all matter couples to the same metric. In the HST framework, this is a consequence of the T^2 coupling, which ensures that every ϕ species sees σ through the same stress-tensor interaction. The T^2 interaction “recalibrates” all scalar composites to read the same geometry.

13.4 The saddle-point vs. fluctuation distinction

	Saddle point (EM)	Fluctuation (Gibbs)
Object	$\bar{\sigma} = \kappa^2 N \langle T \rangle$	$\delta\sigma = \sigma - \bar{\sigma}$
Physics	Classical metric (GR)	Quantum gravity corrections
Polishing	Yes: $\bar{\sigma} \rightarrow g_{\mu\nu}$	No: $\delta\sigma$ stays a tensor
Geometry	Riemannian, with curvature	No geometry; random variable
Mass	None (saddle is a profile)	$m_\sigma > 0$ (from prior)
Statistics	LLN	CLT ($O(1/\sqrt{N})$)

The beauty is in the rendering. The health is in the engine.

13.5 Diffeomorphism invariance: a consequence of the polishing

After polishing, the effective action $S_{\text{eff}}[g]$ depends on σ only through the metric $g_{\mu\nu} = \delta_{\mu\nu} + h_{\mu\nu}$. The action $\int \sqrt{g} R$ is invariant under **diffeomorphisms**: smooth coordinate transformations $x^\mu \rightarrow x'^\mu(x)$. This is the gauge symmetry of general relativity — the freedom to choose coordinates without changing the physics.

Diffeomorphism invariance is *not* a property of the exact lattice theory (which has a fixed lattice and no coordinate freedom). It is an **emergent symmetry** of the polished, long-wavelength theory, valid to $O(a^2/\lambda^2)$. The lattice breaks $\text{Diff}(M)$ just as it breaks Lorentz invariance — but at long wavelengths, both are restored.

13.6 At the quantum level, geometry does not exist

The $1/\sqrt{N}$ fluctuation $\delta\sigma = \sigma - \bar{\sigma}$ is *not* a random metric. It is not that “the metric fluctuates” or “spacetime is fuzzy.” It is that **the concept of a metric does not apply** at the level of quantum fluctuations.

The analogy is thermodynamic. Temperature is a well-defined property of a macroscopic system: it is a saddle-point quantity, the Lagrange multiplier of the energy constraint, meaningful only in the thermodynamic limit. At the level of individual molecules, temperature does not exist — not because it “fluctuates,” but because the concept is undefined. A single molecule does not have a temperature; it has a kinetic energy. Temperature emerges only after averaging over $\sim 10^{23}$ molecules (the law of large numbers).

Similarly: pressure, density, viscosity — all macroscopic (hydrodynamic) concepts that are meaningless at the molecular level. They are properties of the saddle point, not of the fluctuations.

The metric $g_{\mu\nu}$ is exactly like temperature. It is the saddle-point value of the latent tensor $\sigma_{\mu\nu}$, meaningful only after averaging over N matter copies and N_{block} spatial sites. At the level of individual $\sigma_{\mu\nu}(\mathbf{n})$ samples, there is no metric, no geometry, no curvature, no diffeomorphism invariance. There is just a random tensor at each lattice site, drawn from a Gaussian conditional.

Concept	Exists at saddle (macro)	Does not exist at fluctuation (micro)
Gravity	Metric $g_{\mu\nu}$	Just a tensor $\sigma_{\mu\nu}$
Thermodynamics	Temperature T	Just kinetic energies
Fluid mechanics	Pressure p , velocity \mathbf{v}	Just molecular positions

13.7 Strong-field departure from Einstein’s equation

The polishing $h_{\mu\nu} = -2\sigma_{\mu\nu}$ is valid at $O(\kappa^2)$ (weak field). At strong field ($\kappa^2 N \langle T \rangle \sim O(1)$), the $O(\kappa^4)$ terms in the expansion $K[\sigma] \approx K[g] + O(\sigma^2)$ are no longer negligible. The exact lattice saddle-point equation $\bar{\sigma} = \kappa^2 N \langle T \rangle_{K[\bar{\sigma}]}$ differs from Einstein’s equation $G_{\mu\nu} \propto \langle T_{\mu\nu} \rangle_{K[g]}$ by corrections that grow with the field strength.

Near a would-be singularity (where GR predicts $R \rightarrow \infty$), the corrections are $O(1)$ and the polishing breaks down entirely. The exact lattice theory has no singularity (every sum is finite, every eigenvalue of $K[\sigma]$ is positive), but the polished theory (GR) does. The singularity is an artefact of the polishing, not a property of the exact theory.

This is the deepest sense in which the lattice theory is healthier than GR: not just technically (finite sums vs. infinite integrals), but physically (no singularities, no information paradoxes, no breakdown of predictivity). The pathologies of quantum gravity are pathologies of the polishing, absent in the engine.

13.8 The effective action after polishing

After polishing, we integrate out the N copies of ϕ . Each is a Gaussian integral over 16 variables on the curved background:

$$\int d^{16} \phi^{(a)} e^{-\frac{1}{2} \vec{\phi}^{(a)T} K[g] \vec{\phi}^{(a)}} = (2\pi)^8 e^{-\frac{1}{2} \text{Tr} \ln K[g]}. \quad (13.8)$$

The N copies (independent given g) multiply:

$$S_{\text{eff}}[g] = \underbrace{\frac{1}{2\kappa^2} \sigma^2[g]}_{\substack{\text{prior mass} \\ O(1): \text{suppressed}}} + \underbrace{\frac{N}{2} \text{Tr} \ln(-\square_g + m^2)}_{\substack{\text{induced action} \\ O(N): \text{dominant}}}. \quad (13.9)$$

The prior is $O(1)$ (one σ); the Tr log is $O(N)$ (N copies). At large N , we only need the Tr log: the physics is entirely in the log-determinant of the curved Laplacian.

14 The Heat Kernel Expansion: Why GR Is the Leading Term

Now expand the Tr log of the curved Laplacian.

14.1 The Seeley-DeWitt expansion

For slowly varying $g_{\mu\nu}$ (curvature radius $\gg a$):

$$\begin{aligned} \frac{1}{2} \text{Tr} \ln(-\square_g + m^2) &= \underbrace{\frac{a_0}{a^D} \int \sqrt{g}}_{\text{cosmological}} + \underbrace{a_1 \int \sqrt{g} R}_{\text{Einstein-Hilbert}} \\ &+ \underbrace{a_2 a^2 \int \sqrt{g} (c_1 R^2 + c_2 R_{\mu\nu} R^{\mu\nu})}_{\text{4-derivative}} \\ &+ \underbrace{a_3 a^4 \int \sqrt{g} R^3}_{\text{6-derivative}} + \underbrace{a_4 a^6 \int \sqrt{g} R^4}_{O(a^6)} + \dots \end{aligned} \quad (14.1)$$

On the 4×4 lattice: this is $\frac{1}{2} \sum_{i=1}^{16} \ln \lambda_i(-\square_g + m^2)$, expanded for slowly varying g .

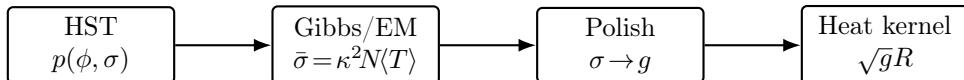
The a_1 term gives Newton's constant: $\frac{1}{16\pi G_N} = \frac{N}{2} a_1$. More matter \rightarrow stronger gravity.

14.2 The tower: each additional curvature power is suppressed

Term	Deriv.	Scaling	$\lambda=10a$	Status
Λ (cosmological)	0	$1/a^D$	$\sim 10^D$	Relevant
$\sqrt{g} R$ (E-H)	2	a^0	1	Marginal: this is GR
$\sqrt{g} R^2$	4	a^2	10^{-2}	Irrelevant: 1%
$\sqrt{g} R^3$	6	a^4	10^{-4}	0.01%
$\sqrt{g} R^4$	8	a^6	10^{-6}	Negligible
$\sqrt{g} R^5$	10	a^8	10^{-8}	Negligible

At solar-system scales: R^2 correction $\sim 10^{-70}$. GR is the *leading non-trivial term* in a systematically improvable expansion. Non-renormalizability is not a problem at fixed $a > 0$.

14.3 The complete pipeline



Each arrow discards information: lattice structure, prior mass, higher curvature terms. The exact theory is rougher but healthier; the polished theory is more beautiful but more pathological.

15 The Hamiltonian: Model 1 (The RNN)

Everything so far has been the path integral (Lagrangian) side — Models 2 and 3. Now we tell the *real* physics story: **Model 1**, the Hamiltonian, which is a recurrent neural network (RNN) with unitary weight matrix.

15.1 The classical Hamiltonian

For the free scalar field, the classical Hamiltonian density is (in continuum notation):

$$\mathcal{H} = \pi \dot{\phi} - \mathcal{L} = \frac{1}{2}\pi^2 + \frac{1}{2}(\partial_x \phi)^2 + \frac{1}{2}m^2 \phi^2 = T + V, \quad (15.1)$$

where $\pi = \partial \mathcal{L} / \partial \dot{\phi} = \dot{\phi}$ is the conjugate momentum. The Lagrangian is $T - V$; the Hamiltonian is $T + V$. Note: the Hamiltonian has the same form as the Euclidean Lagrangian $\mathcal{L}_E = \frac{1}{2}(\partial_\tau \phi)^2 + \frac{1}{2}(\partial_x \phi)^2 + \frac{1}{2}m^2 \phi^2$. This is not a coincidence.

On the 4×4 lattice (4 spatial sites), the classical Hamiltonian is:

$$H = \sum_{n_x=0}^3 \left[\frac{1}{2}\pi_{n_x}^2 + \frac{1}{2} \left(\frac{\phi_{n_x+1} - \phi_{n_x}}{a} \right)^2 + \frac{1}{2}m^2 \phi_{n_x}^2 \right]. \quad (15.2)$$

This is a function of 4 coordinates ϕ_{n_x} and 4 momenta π_{n_x} : an 8-dimensional phase space.

15.2 Quantization: operators and commutators

Promote ϕ_{n_x} and π_{n_x} to operators satisfying:

$$[\hat{\phi}_{n_x}, \hat{\pi}_{n_y}] = i \delta_{n_x n_y}, \quad [\hat{\phi}_{n_x}, \hat{\phi}_{n_y}] = [\hat{\pi}_{n_x}, \hat{\pi}_{n_y}] = 0. \quad (15.3)$$

This is the same as $[\hat{x}, \hat{p}] = i$ in quantum mechanics, but with one pair per spatial site.

15.3 Diagonalization in momentum space

Why momentum? The Hamiltonian has translation symmetry: $[\hat{H}, \hat{T}_a] = 0$ where \hat{T}_a shifts all fields by one site. The eigenstates of \hat{T}_a are momentum states e^{ikn_x} . By the spectral theorem, \hat{H} and \hat{T}_a share eigenstates. Momentum is the natural label because the Hamiltonian is translation-invariant.

Define Fourier modes: $\hat{\phi}_k = \frac{1}{\sqrt{4}} \sum_{n_x} \hat{\phi}_{n_x} e^{-ikn_x a}$, $\hat{\pi}_k = \frac{1}{\sqrt{4}} \sum_{n_x} \hat{\pi}_{n_x} e^{-ikn_x a}$, with $k \in \{0, \frac{2\pi}{4a}, \frac{4\pi}{4a}, \frac{6\pi}{4a}\}$ (4 allowed momenta on a 4-site lattice). The Hamiltonian becomes:

$$\hat{H} = \sum_k \left[\frac{1}{2} |\hat{\pi}_k|^2 + \frac{1}{2} \omega_k^2 |\hat{\phi}_k|^2 \right], \quad \omega_k^2 = \hat{k}^2 + m^2. \quad (15.4)$$

This is **4 independent harmonic oscillators**, one per momentum.

15.4 Creation and annihilation operators

For each mode k , define:

$$\hat{a}_k = \sqrt{\frac{\omega_k}{2}} \hat{\phi}_k + \frac{i}{\sqrt{2\omega_k}} \hat{\pi}_k, \quad \hat{a}_k^\dagger = \sqrt{\frac{\omega_k}{2}} \hat{\phi}_{-k} - \frac{i}{\sqrt{2\omega_k}} \hat{\pi}_{-k}. \quad (15.5)$$

These satisfy $[\hat{a}_k, \hat{a}_{k'}^\dagger] = \delta_{kk'}$. The Hamiltonian becomes:

$$\boxed{\hat{H} = \sum_k \omega_k \left(\hat{a}_k^\dagger \hat{a}_k + \frac{1}{2} \right) = \sum_k \omega_k \left(\hat{n}_k + \frac{1}{2} \right)}, \quad (15.6)$$

where $\hat{n}_k = \hat{a}_k^\dagger \hat{a}_k$ is the **number operator**: it counts particles with momentum k .

15.5 Fock space

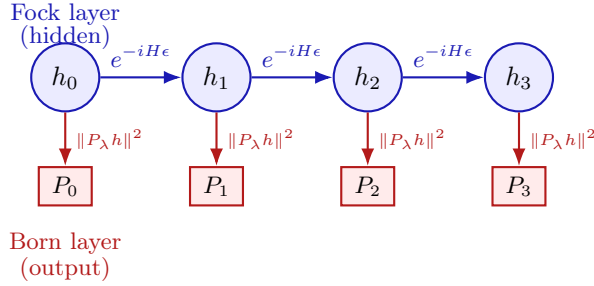
The vacuum $|0\rangle$ is the state with no particles: $\hat{a}_k|0\rangle = 0$ for all k . The one-particle states are $|k\rangle = \hat{a}_k^\dagger|0\rangle$, with energy ω_k . Multi-particle states: $|k_1, k_2\rangle = \hat{a}_{k_1}^\dagger \hat{a}_{k_2}^\dagger|0\rangle$, etc.

The Hilbert space $\mathcal{F} = \bigoplus_{n=0}^{\infty} \mathcal{H}_n$ (the **Fock space**) is the direct sum over all particle numbers. On the 4-site lattice, k takes 4 values, and each mode can hold $0, 1, 2, \dots$ particles (bosonic).

15.6 The Schrödinger equation: the RNN

The state evolves as $|h_t\rangle = e^{-iHt}|h_0\rangle$ (the **Schrödinger equation**). This is a *unitary* evolution: $\|h_t\| = \|h_0\|$. The Hamiltonian H generates time translations.

This is **Model 1: a recurrent neural network**. The hidden state h_t is a vector in Fock space. The weight matrix is $W = e^{-iH\epsilon}$ (unitary). At each time step, $h_{t+\epsilon} = W h_t$. The Born rule $P(\lambda) = \|P_\lambda h_t\|^2$ is the output layer, producing the probabilities of observed outcomes. The RNN runs deterministically in the hidden (Fock) layer; randomness appears only at the output (Born) layer.



15.7 From Hamiltonian back to path integral

The path integral is recovered by time slicing: $\langle \phi_f | e^{-iHT} | \phi_i \rangle = \int \prod_{n_t} \mathcal{D}\phi_{n_t} e^{iS_M[\phi]}$. Insert $\mathbf{1} = \int d\vec{\phi}_{n_t} |\vec{\phi}_{n_t}\rangle \langle \vec{\phi}_{n_t}|$ between each pair of $e^{-iH\epsilon}$. Use the Trotter formula. The result is the Lorentzian path integral (4.1) that we started with.

On the Euclidean side: $\langle \phi_f | e^{-H\tau} | \phi_i \rangle = \int \prod d\phi e^{-S_E}$. Same procedure, same transfer matrix, but with $e^{-H\tau}$ instead of e^{-iHt} .

The circle is closed: Hamiltonian \rightarrow time slicing \rightarrow path integral \rightarrow Wick rotation \rightarrow Euclidean path integral \rightarrow transfer matrix \rightarrow Hamiltonian.

15.8 Why the path integral gives physical answers

The path integral is not fundamental — the Hamiltonian is. But the path integral encodes the same spectrum: the poles of the Lorentzian propagator are the energy eigenvalues ω_k ; the decay rates of the Euclidean correlator are the same ω_k . The equivalence is guaranteed by the transfer matrix construction.

15.9 The interacting Hamiltonian: a polynomial of creation operators

Add the ϕ^4 interaction. Express $\hat{\phi}_{n_x}$ in terms of creation/annihilation operators:

$$\hat{\phi}_{n_x} = \sum_k \frac{1}{\sqrt{2\omega_k}} \left(\hat{a}_k e^{ikn_x a} + \hat{a}_k^\dagger e^{-ikn_x a} \right). \quad (15.7)$$

On the 4-site lattice, k takes 4 values, so this is a sum of 4 terms. The interaction Hamiltonian is:

$$\hat{H}_I = \frac{\lambda}{4!} \sum_{n_x=0}^3 \hat{\phi}_{n_x}^4 = \frac{\lambda}{4!} \sum_{n_x} \left(\sum_k \frac{\hat{a}_k e^{ikn_x a} + \hat{a}_k^\dagger e^{-ikn_x a}}{\sqrt{2\omega_k}} \right)^4. \quad (15.8)$$

Expanding the fourth power produces terms with all combinations of \hat{a} and \hat{a}^\dagger :

$\hat{a}^\dagger \hat{a}^\dagger \hat{a}^\dagger \hat{a}^\dagger$: creates 4 particles (from vacuum).

$\hat{a}^\dagger \hat{a}^\dagger \hat{a}^\dagger \hat{a}$: creates 3, destroys 1 (net creation of 2).

$\hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a}$: creates 2, destroys 2 (**scattering**: $2 \rightarrow 2$).

And so on, down to $\hat{a} \hat{a} \hat{a} \hat{a}$ (destroys 4).

The full Hamiltonian $\hat{H} = \hat{H}_0 + \hat{H}_I$ is no longer diagonal in the Fock basis: particles scatter, are created, and annihilated. The matrix $\langle n'_0, n'_1, n'_2, n'_3 | \hat{H} | n_0, n_1, n_2, n_3 \rangle$ has off-diagonal elements connecting different occupation-number states.

15.10 The Dyson series: another route to Feynman diagrams

The time-evolution operator in the interaction picture:

$$\hat{U}(t, 0) = T \exp \left(-i \int_0^t \hat{H}_I(t') dt' \right) = \mathbf{1} - i \int_0^t \hat{H}_I(t') dt' + (-i)^2 \int_0^t \int_0^{t'} \hat{H}_I(t') \hat{H}_I(t'') dt'' dt' + \dots \quad (15.9)$$

Each term is a product of creation/annihilation operators at different times, contracted by the free propagator $\langle 0 | T \{ \hat{\phi}(x) \hat{\phi}(y) \} | 0 \rangle = G_0(x, y)$.

This is the **Dyson series**: the same Feynman diagrams obtained from the Euclidean Taylor expansion of $e^{-\lambda \phi^4 / (4!)}$, but derived from the Hamiltonian side. Lines are propagators G_0 ; vertices are interaction points λ ; the sum over intermediate states is the sum over Fock-space occupation numbers. On the 4×4 lattice, each diagram is a finite sum over 4 momenta and 4 time slices.

15.11 Rendering stability from the Hamiltonian side

The Hamiltonian spectrum $\{E_n\}$ defines the physical observables: masses ($E_1 - E_0$), scattering amplitudes (matrix elements of \hat{U}), decay rates. These are *mean parameters*: properties of the eigenvalues and eigenvectors of \hat{H} , not of the bare couplings $m^2(a), \lambda(a)$ in the Lagrangian.

As the lattice spacing a changes, the bare couplings $m^2(a), \lambda(a)$ must be adjusted to keep the spectrum fixed:

$$E_1(m^2(a), \lambda(a), a) - E_0(m^2(a), \lambda(a), a) = m_{\text{phys}} = \text{const}. \quad (15.10)$$

This is **rendering stability**: the natural parameters run with a ; the mean parameters do not.

15.12 Observables and degenerate observables

An **observable** \hat{O} is a Hermitian operator on Fock space. Its expectation in state $|h_t\rangle$ is $\langle \hat{O} \rangle_t = \langle h_t | \hat{O} | h_t \rangle$.

A **degenerate observable** has repeated eigenvalues: multiple eigenstates give the same measurement outcome. The projection onto the eigenspace of outcome λ is P_λ , and the Born probability is $P(\lambda) = \|P_\lambda h_t\|^2$.

Degeneracy arises from **partial access**: the observer measures only a subset of the degrees of freedom (e.g., the field at one site, not all 4 sites; or the total energy, not the mode-by-mode occupation numbers).

15.13 Partial access and thermalization

When the observer has access only to a subsystem A (e.g., one spatial site), the full state $|h_t\rangle$ is projected to a **reduced density matrix**:

$$\rho_A = \text{Tr}_B(|h_t\rangle\langle h_t|), \quad (15.11)$$

where B is the complement of A (the other 3 sites). Even if $|h_t\rangle$ is a pure state, ρ_A is generically mixed.

For a generic energy eigenstate (or a state evolved for long time), the reduced density matrix approaches a **thermal** form: $\rho_A \approx e^{-\beta_{\text{eff}} H_A} / Z_A$. This is the **eigenstate thermalization hypothesis** (ETH): partial access to a quantum system produces effective thermality.

This connects to our three models: Model 1 (the RNN) runs unitarily on the full Fock space. But the observer’s partial access (degenerate observables) produces an effective thermal state — which is Model 2 (the EBM $e^{-\beta H} / Z$). The Euclidean path integral is not just a computational trick; it is the natural description of what a partial-access observer sees.

15.14 Composite operators and macroscopic averages

A **composite operator** is built from products of the fundamental field: $\hat{\phi}_{n_x}^2, (\partial_\mu \hat{\phi})^2, \hat{T}_{\mu\nu}[\phi]$. These are the observables that “instruments” (clocks, rulers, detectors) actually measure.

A **macroscopic average** is a composite operator averaged over a large region or many copies:

$$\bar{O} = \frac{1}{N_{\text{avg}}} \sum_{i=1}^{N_{\text{avg}}} \hat{O}_i. \quad (15.12)$$

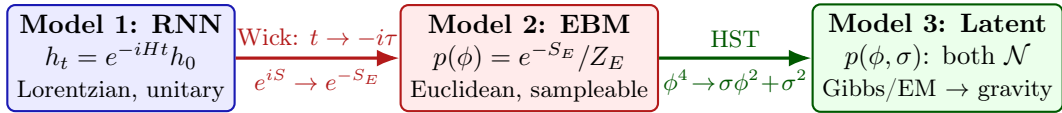
By the **law of large numbers**, $\bar{O} \rightarrow \langle \hat{O} \rangle$ as $N_{\text{avg}} \rightarrow \infty$: the macroscopic average converges to the quantum expectation. By the **central limit theorem**, the fluctuation is $O(1/\sqrt{N_{\text{avg}}})$ and Gaussian.

This is exactly the saddle-point / fluctuation decomposition of σ in the N -copy construction: the EM saddle $\bar{\sigma} = \kappa^2 N \langle T \rangle$ is the macroscopic average (LLN); the $O(1/\sqrt{N})$ Gibbs fluctuation is the quantum correction (CLT). The classical metric is the macroscopic average of a composite operator. Quantum gravity is the fluctuation around that average.

16 Summary: Three Generative Models on 16 Numbers

Concept	Continuum	4×4 lattice
Model 1: RNN (Lorentzian / Hamiltonian)		
Hamiltonian	$H = \frac{1}{2}\pi^2 + \frac{1}{2}(\partial_x\phi)^2 + \frac{1}{2}m^2\phi^2$	4 harmonic oscillators
Particles	$a_k^\dagger 0\rangle$	4 creation operators
Time evolution	$ h_t\rangle = e^{-iHt} h_0\rangle$	Unitary 16×16 matrix
Output	Born rule $P = \ P_\lambda h_t\ ^2$	Projection + squared norm
Model 2: EBM / Exponential Family (Euclidean)		
Distribution	$p(\phi) \propto e^{-S_E}$	$\mathcal{N}(0, K_E^{-1})$: 16-dim Gaussian
Propagator	$G(k) = 1/(k^2 + m^2)$	$(K_E^{-1})_{\mathbf{k}}$: covariance
Log-normaliser	$\ln Z = -\frac{1}{2}\text{Tr} \ln K_E$	Sum of 16 logarithms
Mass	Decay rate of $G_E(\tau)$	Fit exponential on 4-site lattice
Model 3: Latent Variable Model (HST / Gravity)		
Latent field	$\sigma_{\mu\nu}$	$3 \times 16 = 48$ extra variables
Both conditionals	Gaussian (like RBM)	FFT-Gaussian + site-independent
Saddle (EM, LLN)	Einstein: $\bar{\sigma} = \kappa^2 N \langle T \rangle$	Deterministic update at each site
Fluctuation (CLT)	Graviton: $\langle \sigma\sigma \rangle_c \sim 1/N$	Encoder noise $\sim \kappa^2$
Polishing	$\sigma \rightarrow g_{\mu\nu} = \delta + h$	$h = -2\kappa^2 \langle T \rangle + O(\kappa^4)$
Heat kernel	$\sqrt{g} R$ from $\text{Tr} \log$	Sum of 16 log-eigenvalues

The passage from Model 1 to Model 2 is the **Wick rotation** ($t \rightarrow -i\tau$): it converts the unitary RNN into a sampleable energy-based model. The passage from Model 2 to Model 3 is the **Hubbard-Stratonovich transformation**: it introduces a latent variable that makes the model conditionally Gaussian.



Gravity emerges as the effective dynamics of the latent variable: the saddle point (EM, law of large numbers) gives Einstein's equation; the $1/\sqrt{N}$ fluctuation (Gibbs, central limit theorem) gives quantum gravity. The polishing ($\sigma \rightarrow g_{\mu\nu}$) is forced by matching the HST precision to the curved Laplacian at $O(\kappa^2)$: clocks and rulers (scalar composites) all see the same geometry.

Every concept has an explicit, computable, finite-dimensional realisation on the 4×4 lattice. The continuum theory is the limit $a \rightarrow 0$, $V \rightarrow \infty$; but the physics is already present at $V = 16$.

A The Stress-Energy Tensor: Three Perspectives

The stress-energy tensor $T_{\mu\nu}[\phi]$ plays a central role in this tutorial: the T^2 interaction is decoupled by the HST, the saddle point $\bar{\sigma} = \kappa^2 N \langle T \rangle$ is Einstein’s equation, and the polishing $K[\sigma] \rightarrow K[g]$ works because $T_{\mu\nu}$ is the correct coupling. But why $T_{\mu\nu}$? Three perspectives.

A.1 Perspective 1: Noether current from spacetime symmetry

The free scalar Lagrangian $\mathcal{L} = \frac{1}{2}(\partial_\mu\phi)^2 - \frac{1}{2}m^2\phi^2$ is invariant under spacetime translations $x^\mu \rightarrow x^\mu + a^\mu$. By **Noether’s theorem**, every continuous symmetry yields a conserved current. The conserved current associated with translation invariance is the **canonical stress-energy tensor**:

$$T_{\mu\nu}^{\text{can}} = \partial_\mu\phi \partial_\nu\phi - \delta_{\mu\nu} \mathcal{L} = \partial_\mu\phi \partial_\nu\phi - \frac{1}{2}\delta_{\mu\nu}[(\partial\phi)^2 - m^2\phi^2]. \quad (\text{A.1})$$

Conservation: $\partial^\mu T_{\mu\nu}^{\text{can}} = 0$ (on shell, i.e., when ϕ satisfies the Klein-Gordon equation). This gives 4 conserved charges: $P_\nu = \int dx T_{0\nu}^{\text{can}}$ is the total energy-momentum.

Kinematic vs. dynamic. The conservation law $\partial^\mu T_{\mu\nu} = 0$ is **kinematic**: it follows from the symmetry alone, regardless of the specific dynamics (free or interacting, scalar or spinor, with or without gravity). It is a consequence of translation invariance, period.

The T^2 interaction introduces **dynamics**: it determines *how* the stress tensor sources the latent field σ , producing the self-consistent saddle (Einstein’s equation) and the quantum fluctuations around it. The dynamics distorts the spacetime symmetry: $\sigma_{,\mu\nu} \neq 0$ breaks translation invariance (different points see different σ), generating curvature.

So the kinematic property ($\partial^\mu T_{\mu\nu} = 0$) tells us *what* is conserved; the dynamic interaction (T^2) tells us *how* the conserved quantity sources geometry. The hydrodynamic mode (the EM saddle) is “protected” by the conservation law: the saddle must respect $\partial^\mu \langle T_{\mu\nu} \rangle = 0$, which constrains the metric perturbation $h_{\mu\nu}$ to satisfy the linearised Einstein equation (which has the Bianchi identity $\nabla^\mu G_{\mu\nu} = 0$ built in).

A.2 Perspective 2: Hilbert definition (variation of the matter action)

Given a matter action $S_{\text{matter}}[\phi, g]$ on a curved background $g_{\mu\nu}$, the **Hilbert stress-energy tensor** is defined as:

$$T_{\mu\nu} = -\frac{2}{\sqrt{g}} \frac{\delta S_{\text{matter}}}{\delta g^{\mu\nu}}. \quad (\text{A.2})$$

For the free scalar on a curved background, $S_{\text{matter}} = \frac{1}{2} \int \sqrt{g} (g^{\mu\nu} \partial_\mu\phi \partial_\nu\phi + m^2\phi^2)$:

$$T_{\mu\nu} = \partial_\mu\phi \partial_\nu\phi - \frac{1}{2}g_{\mu\nu} [g^{\alpha\beta} \partial_\alpha\phi \partial_\beta\phi + m^2\phi^2]. \quad (\text{A.3})$$

Evaluated at $g_{\mu\nu} = \delta_{\mu\nu}$ (flat space), this coincides with the canonical tensor.

Why this is relevant for the HST. The polishing section showed that $K[\sigma] \approx K[g]$ at $O(\kappa^2)$, with $h_{\mu\nu} = -2\sigma_{,\mu\nu}$. The coupling $\sigma^{\mu\nu} T_{\mu\nu}[\phi]$ in the HST action is precisely the first-order variation of $S_{\text{matter}}[\phi, g]$ with respect to g :

$$S_{\text{matter}}[\phi, \delta + h] = S_{\text{matter}}[\phi, \delta] - \frac{1}{2} \int h^{\mu\nu} T_{\mu\nu}[\phi] + O(h^2). \quad (\text{A.4})$$

The HST coupling $\sigma \cdot T$ is the linearisation of “ ϕ propagates on g ” — the same calculation we did in the polishing section, now understood as the definition of $T_{\mu\nu}$ via the Hilbert variational principle. The two perspectives (Noether and Hilbert) give the same tensor because the action is diffeomorphism-invariant.

A.3 Perspective 3: T^2 from the Wilsonian RG flow

The T^2 interaction can also be understood as generated by the **Wilsonian renormalization group** (RG). Starting from a free scalar field and integrating out high-momentum shells $\Lambda' < |k| < \Lambda$, the effective action at scale Λ' acquires all operators consistent with the symmetries — including $T_{\mu\nu}T^{\mu\nu}$.

Concretely: the one-loop effective action from integrating out fast modes generates a tower of operators organised by dimension. The leading operators are the kinetic term $(\partial\phi)^2$ and the mass $m^2\phi^2$ (renormalization of existing couplings). The subleading operators include $(\partial\phi)^4$, $\phi^2(\partial\phi)^2$, and $T_{\mu\nu}T^{\mu\nu}$ — all of which are quartic in ϕ and its derivatives.

The T^2 interaction is thus not introduced “by hand”: it is *generated* by the RG flow whenever the theory has spacetime (translation/rotation) symmetry. The HST then decouples this RG-generated interaction into the latent field $\sigma_{\mu\nu}$, and gravity emerges as the induced dynamics of σ .

In the lattice-first view: the exact lattice theory already contains T^2 (and all higher interactions consistent with the lattice symmetries) at any nonzero coupling. The question is not “why T^2 ?” but “why is T^2 the *leading* interaction that couples to a rank-2 tensor?” The answer: because $T_{\mu\nu}$ is the unique rank-2 conserved current of the free theory (Noether), and the RG flow generates the lowest-dimension interaction built from it.

B The Heat Kernel Expansion: A Closer Look

B.1 How the Tr log arises

The $\text{Tr} \ln K$ appears in two ways:

Exact (in the HST). Integrating out ϕ from the Gaussian conditional $p(\phi|\sigma) = \mathcal{N}(0, K[\sigma]^{-1})$ gives $\int d^V\phi e^{-\frac{1}{2}\phi^T K \phi} = (2\pi)^{V/2}/\sqrt{\det K} = \text{const} \times e^{-\frac{1}{2}\text{Tr} \ln K}$. On the 4×4 lattice: $\text{Tr} \ln K = \sum_{i=1}^{16} \ln \lambda_i$ (sum of 16 logarithms of the eigenvalues of the 16×16 matrix K). This is exact, not an approximation.

Approximate (in general QFT). For a non-Gaussian theory $p(\phi) \propto e^{-S[\phi]}$, expand around the saddle $\bar{\phi}$: $S[\phi] = S[\bar{\phi}] + \frac{1}{2}(\phi - \bar{\phi})^T S''[\bar{\phi}](\phi - \bar{\phi}) + \dots$. The one-loop correction is the Gaussian integral over the fluctuation, giving $\frac{1}{2}\text{Tr} \ln S''[\bar{\phi}]$. This is exact at one loop; higher loops add corrections.

In both cases, the object to expand is $\text{Tr} \ln K$ where K is a $V \times V$ positive-definite matrix.

B.2 The Schwinger proper-time representation

The key identity: for any positive operator K ,

$$\boxed{\ln K = - \int_0^\infty \frac{ds}{s} (e^{-sK} - e^{-s})}, \quad (\text{B.1})$$

where s is the **proper time** (Schwinger parameter). Taking the trace:

$$\text{Tr} \ln K = - \int_0^\infty \frac{ds}{s} (\text{Tr} e^{-sK} - V). \quad (\text{B.2})$$

On the 4×4 lattice: $\text{Tr} e^{-sK} = \sum_{i=1}^{16} e^{-s\lambda_i}$ is a sum of 16 decaying exponentials. The integral over s gives $\sum_i \ln \lambda_i = \text{Tr} \ln K$.

B.3 Why it is called the “heat kernel”

The operator e^{-sK} is the **heat kernel**: the solution to the heat equation

$$\frac{\partial}{\partial s} u(s, \mathbf{x}) = -K u(s, \mathbf{x}), \quad u(0, \mathbf{x}) = \delta(\mathbf{x} - \mathbf{y}). \quad (\text{B.3})$$

If $K = -\Delta + m^2$ (the Laplacian plus mass), this is literally the heat equation: $u(s, \mathbf{x})$ describes the diffusion of heat from a point source at \mathbf{y} , with s playing the role of time. The “proper time” s is not physical time; it is the fictitious time of a diffusion process on the lattice (or on a curved manifold, after polishing).

On the 4×4 lattice: e^{-sK} is a 16×16 matrix whose (i, j) entry $[e^{-sK}]_{ij}$ gives the “heat” at site i at time s , starting from a unit source at site j .

B.4 The asymptotic expansion

For small s (short proper time), $\text{Tr} e^{-sK}$ has an asymptotic expansion in powers of s :

$$\text{Tr} e^{-sK} \sim \frac{1}{(4\pi s)^{D/2}} \sum_{n=0}^{\infty} a_n s^n \quad (s \rightarrow 0^+), \quad (\text{B.4})$$

where D is the dimension and a_n are the **Seeley-DeWitt coefficients**. On a finite lattice, the expansion is in powers of s/a^2 (the ratio of proper time to lattice scale squared).

Substituting into the proper-time integral:

$$\text{Tr} \ln K \sim -\frac{a_0}{(4\pi)^{D/2}} \cdot \frac{1}{a^D} - \frac{a_1}{(4\pi)^{D/2}} \cdot \ln \frac{1}{m^2 a^2} - \frac{a_2}{(4\pi)^{D/2}} \cdot a^0 - \dots \quad (\text{B.5})$$

Each coefficient a_n is a local functional of the background (the metric $g_{\mu\nu}$ after polishing, or the tensor $\sigma_{\mu\nu}$ before polishing).

B.5 The unpolished expansion ($K[\sigma]$) vs. the polished expansion ($K[g]$)

The heat kernel expansion works for *any* positive operator K , polished or not. The polishing map $h_{\mu\nu} = -2\sigma_{\mu\nu} + O(\sigma^2)$ induces a **term-by-term correspondence** between the unpolished coefficients $a_n[\sigma]$ (polynomials in σ and its lattice derivatives) and the polished coefficients $a_n[g]$ (geometric invariants built from the curvature):

n	Unpolished $a_n[\sigma]$	Polished $a_n[g]$	Connection
0	$\sum_{\mathbf{n}} f_0(\sigma_{\mathbf{n}})$	$\int \sqrt{g}$	$\sqrt{g} \approx 1 + \frac{1}{2} h^\alpha{}_\alpha = 1 - \sigma^\alpha{}_\alpha$
1	$\sum_{\mathbf{n}} [c \Delta \sigma + c' \sigma \Delta \sigma + \dots]$	$\frac{1}{6} \int \sqrt{g} R$	$R \sim \partial^2 h \sim \partial^2 \sigma$
2	$\sum_{\mathbf{n}} [(\Delta \sigma)^2 + (\partial \sigma \partial \sigma)^2 + \dots]$	$\frac{1}{180} \int \sqrt{g} (R_{\mu\nu\rho\sigma}^2 - R_{\mu\nu}^2 + \dots)$	$R^2 \sim (\partial^2 \sigma)^2$

The correspondence is systematic: each $\partial^2 \sigma$ maps to a curvature R (since $R \sim \partial^2 g \sim \partial^2 \sigma$), each $(\partial \sigma)^2$ maps to $(\partial g)^2 \sim \Gamma^2$ (Christoffel symbols squared), and each power of σ without derivatives maps to a factor of \sqrt{g} . The polishing is an invertible change of variables (at weak field) that repackages the lattice polynomials into geometric invariants.

The key point: the unpolished $a_n[\sigma]$ are computable directly on the lattice (they are sums over 16 sites involving σ and its finite differences). The polished $a_n[g]$ are the same numbers,

rewritten in geometric language. The geometry is an *interpretation* of the lattice computation, not an independent input.

At strong field ($\sigma \sim O(1)$), the $O(\sigma^2)$ correction in $h = -2\sigma + O(\sigma^2)$ matters: the unpolished and polished coefficients differ, and the geometric interpretation breaks down. The lattice computation (using $a_n[\sigma]$) remains valid; the geometric interpretation (using $a_n[g]$) does not. This is another sense in which the lattice theory is more fundamental than GR.

$$a_0[g] = \int \sqrt{g} d^D x \quad (\text{volume}) \quad (\text{B.6})$$

$$a_1[g] = \frac{1}{6} \int \sqrt{g} R d^D x \quad (\text{Einstein-Hilbert}) \quad (\text{B.7})$$

$$a_2[g] = \frac{1}{180} \int \sqrt{g} (R_{\mu\nu\rho\sigma} R^{\mu\nu\rho\sigma} - R_{\mu\nu} R^{\mu\nu} + \frac{5}{2} R^2 + 6\Box R) d^D x \quad (4\text{-derivative}) \quad (\text{B.8})$$

The polishing maps $a_n[\sigma] \rightarrow a_n[g]$: the unpolished coefficients become geometric invariants. We care about the polished expansion because it gives us general relativity ($a_1 \propto \int \sqrt{g} R$), but the unpolished expansion is equally well-defined and more fundamental (no approximation needed).

B.6 Example 1: gravity ($K = -\Box_g + m^2$, spin 0)

For N minimally coupled scalars on a curved background:

$$\frac{N}{2} \text{Tr} \ln(-\Box_g + m^2) \rightarrow \frac{N}{12} \cdot \frac{1}{(4\pi)^{D/2}} \int \sqrt{g} R + \dots \quad (\text{B.9})$$

Comparing with the Einstein-Hilbert action $\frac{1}{16\pi G_N} \int \sqrt{g} R$:

$$\frac{1}{16\pi G_N} = \frac{N}{12(4\pi)^{D/2}} \cdot I_1, \quad (\text{B.10})$$

where I_1 depends on the UV structure (on the lattice: $I_1 \sim 1/a^{D-2}$). Newton's constant is inversely proportional to N : more matter \rightarrow stronger induced gravity.

B.7 Example 2: gauge theory ($K = -D_\mu^2 + m^2$, spin 1)

For a gauge field A_μ with covariant derivative $D_\mu = \partial_\mu - igA_\mu$, the operator is $K = -D_\mu D^\mu + m^2$. The heat kernel expansion gives:

$$\frac{1}{2} \text{Tr} \ln(-D^2 + m^2) \rightarrow \frac{1}{12(4\pi)^{D/2}} \int \text{tr}(F_{\mu\nu} F^{\mu\nu}) + \dots \quad (\text{B.11})$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu - ig[A_\mu, A_\nu]$ is the field strength. The a_1 coefficient generates the **Yang-Mills action** $\int \text{tr}(F^2)$ — Maxwell's equations (for $U(1)$) or the gluon dynamics (for $SU(3)$) are induced by the matter loops, just as Einstein's equation is.

B.8 The universal pattern

Auxiliary field	Rank	Operator K	a_1 generates
σ (scalar)	0	$-\Delta + m^2 + \sigma$	$(\partial\sigma)^2$: scalar kinetics
A_μ (photon)	1	$-D_\mu^2 + m^2$	$F_{\mu\nu}F^{\mu\nu}$: Maxwell/YM
$\sigma_{\mu\nu}$ (graviton)	2	$-\square_g + m^2$	$\sqrt{g}R$: Einstein- Hilbert

At every rank, the same mechanism: an auxiliary field with a Gaussian prior, no bare kinetic term, coupled to matter through its conserved current. The matter loops (Tr log) generate the kinetic term via the a_1 Seeley-DeWitt coefficient. The induced action is Maxwell (F^2), Einstein ($\sqrt{g}R$), or a scalar kinetic term ($(\partial\sigma)^2$) — depending on the rank of the auxiliary field.