The Adjoint Field as an Influence Map: Sensitivity Analysis of FRG Flows

Illustrated on the $O(N)$ model in the fluid-dynamic formulation

Ashutosh Dash

Goethe University, Frankfurt

ERG 2026 · University of Sussex, UK · 2 September

Motivation

A FRG calculation integrates the flow from the UV to the IR to calculate observables $\mathcal{O}$: a mass, f_π, a vertex, a critical exponent, etc.

A recurring question is how strongly $\mathcal{O}$ depends on the choices entering the flow:

- the regulator / shape function (scheme dependence)
- the truncation and retained couplings (systematic error)
- the UV initial data (physical input)
- the field-space resolution (numerical error)

Standard approach

Vary one choice, repeat the full flow, and record the change in $\mathcal{O}$: one run for each check.

The adjoint method provides these sensitivities collectively, at the cost of a single additional (backward) solve.

Cost of a direct parameter scan

For an input with M parameters $p = (p_1, \ldots, p_M)$, a finite-difference gradient perturbs each in turn,

$$\frac{\partial \mathcal{O}}{\partial p_i} \approx \frac{\mathcal{O}(p + \varepsilon e_i) - \mathcal{O}(p)}{\varepsilon}, \qquad i = 1, \ldots, M,$$

requiring one full flow per parameter.

	Direct scan	Adjoint
Gradient with respect to M parameters	M solves	1 solve
Sensitivity to the field at all (σ, t)	not accessible	1 solve
Each additional observable	—	1 backward solve

- The field is a continuum of inputs; a scan cannot resolve it.
- The adjoint cost is independent of the number of parameters M.

Setup: the flow as an initial-value problem

The $0+0$ dimensional $O(N)$ in LPA, written as a conservation law for $u = \partial_\sigma U$ in RG-time $t = \ln(\Lambda/k)$:

$$\partial_t u = L(u), \qquad L(u) \equiv \partial_\sigma\big(Q(\partial_\sigma u) - F(u)\big), \qquad \Gamma^{(2n)} = \partial_\sigma^{2n-1} u\big|_{\sigma=0}.$$

- σ: radial (invariant) mode of the $O(N)$ model; t: RG-time; $\Gamma^{(2n)}$: IR vertices.
- L: the spatial operator: advective flux F and degenerate diffusion Q; the solution develops a steep restoration front.
- A single UV$\rightarrow$IR integration yields the forward trajectory $u(\sigma, t)$.

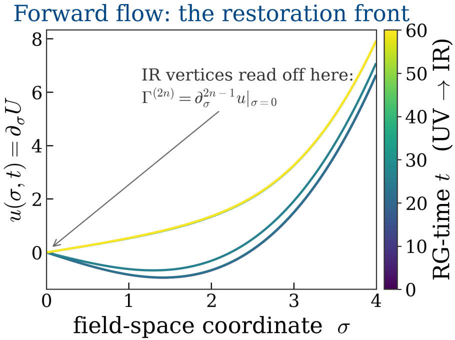

This fluid-dynamic formulation is standard across FRG applications [Grossi and Wink, 2023; Koenigstein et al., 2022]; it serves here as a concrete, fully controlled testbed.

Example $\Gamma^{(4)}$: Finite difference vs. Fit

The usual check is self-convergence under refinement. But, both results are irreconcilable.

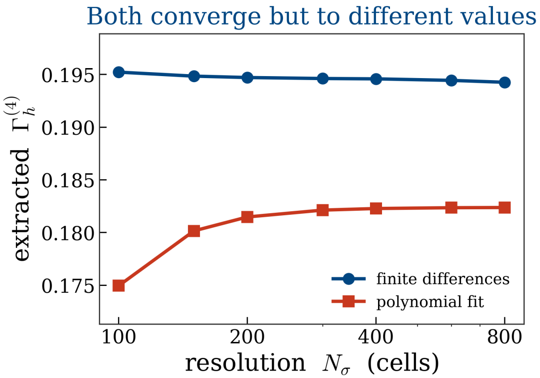

Grid refinement

the value settles $\Rightarrow$ looks converged, tempting to trust.

The adjoint

the error estimate stays **non-zero**.

Grid-convergence is not correctness: the flow converged to a biased value. The adjoint flags it, with no exact answer needed.

The adjoint field λ: sensitivity of the observable

The adjoint $\lambda(\sigma, t)$ is defined as the sensitivity of $\mathcal{O}$ to the state at every (σ, t):

$$\delta\mathcal{O} = \big\langle \lambda(\cdot, t),\, \delta u(\cdot, t) \big\rangle \qquad \Longleftrightarrow \qquad \lambda(\sigma, t) = \frac{\delta\mathcal{O}}{\delta u(\sigma, t)}.$$

Interpretation. $\lambda(\sigma_0, t)$ is the response of $\mathcal{O}$ to a localized unit perturbation of the state at (σ_0, t), an influence map.

Obtained by a single backward integration of the transposed linearized flow,

$$\frac{d\lambda}{dt} = -A(t)^{\mathsf{T}}\,\lambda, \qquad \lambda(t_{\mathrm{IR}}) = \frac{\partial\mathcal{O}}{\partial u}\bigg|_{t_{\mathrm{IR}}},$$

seeded by the observable, with $A = \partial L/\partial u$.

Reading the map. support concentrates at the physical point $\sigma = 0$ and along the front; negligible at $\sigma_{\max}$.

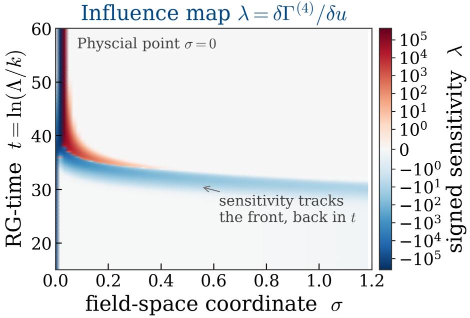

(Adjoint / costate method, Giles and Pierce, 2000; derivation: back-up.)

Application 1: RG freeze-out, per observable

The RG-time weight $S(t) = \|\lambda(\cdot, t)\|_1$ measures how much the observable is still being shaped at scale t.

- flat and low in the UV. Perturbations there wash out, those scales do not affect the observable;
- the rise marks the RG window in which the observable is actually built;
- it saturates near the IR, where it is fixed just before read-out.

The rise sits at a slightly later RG-time for higher vertices: $\Gamma^{(2)} < \Gamma^{(4)} < \Gamma^{(6)}$.

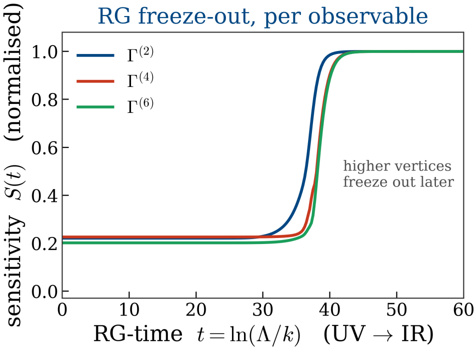

Application 2: a computable error bar on the vertex

Solve the flow on a coarse grid $\Rightarrow$ vertex $\mathcal{O}_h$. Its error is estimated by weighting the flow's residual with the sensitivity λ: [Becker and Rannacher, 2001]

$$\underbrace{\mathcal{O}_{\mathrm{ref}} - \mathcal{O}_h}_{\text{error}} \approx \eta = \int_{\mathrm{UV}}^{\mathrm{IR}} \langle\, \lambda(\sigma, t),\ \underbrace{L_{\mathrm{ref}} - L_h}_{\text{residual } R} \,\rangle\ dt.$$

λ: sensitivity of the vertex to the state (the backward solve).

$R = L_{\mathrm{ref}} - L_h$: residual – how far the coarse flow mis-solves the finer one.

$\langle\cdot,\cdot\rangle$: pairing over field space σ.

Reliable

$$I_{\mathrm{eff}} = \frac{\eta}{\mathcal{O}_{\mathrm{ref}} - \mathcal{O}_h} \approx 1$$

estimate $\approx$ true error.

An error bar on the vertex – computable and localizable without the need of the exact answer.

Where is the error ? Sensitivity $\times$ residual

The error density λR is non-zero only where the observable is sensitive and the scheme is wrong:

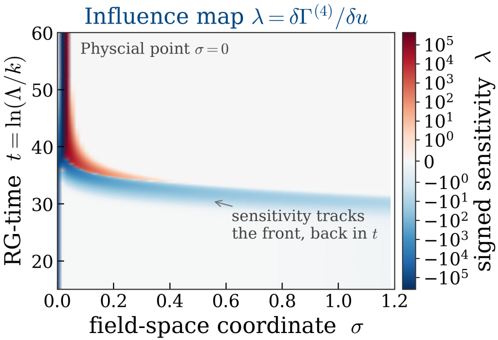

sensitivity λ: broad, along the front

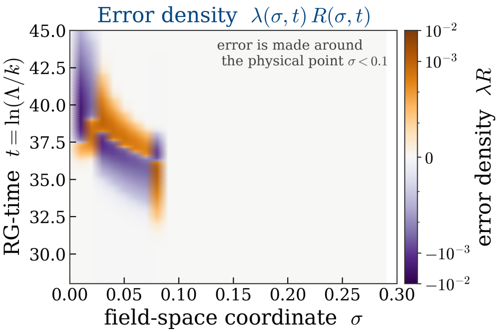

density λR: collapses around the physical point

$\geq 99.9\%$ of the error at $\sigma < 0.1$: the bias made around the physical point.

Diagnosis, and a cure

Near $\sigma = 0$ the advective flux carries a spurious $1/\sigma$ term, biasing $\Gamma^{(4)}$, independent of the mesh. A fit rather than finite-difference vertex-extraction method cures it.

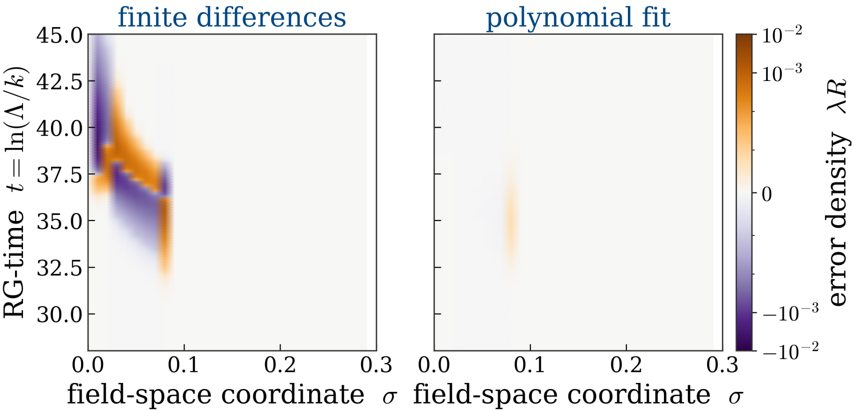

Changing the extraction method collapses the error density by orders of magnitude: the adjoint quantifies it. (Details: backup.)

The true error: convergence to the reference

Relative error of $\Gamma^{(2,4,6)}$ vs the exact[1] ϕ^4 $O(4)$ values: FD vs polynomial fit:

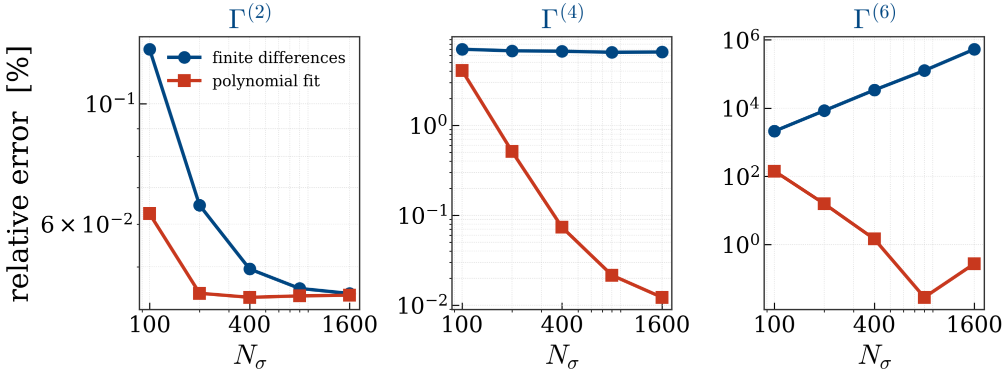

- $\Gamma^{(2)}$: accurate both ways ($< 0.1\%$). $\Gamma^{(4)}$: FD plateaus at $\sim 7\%$, fit converges to $\sim 10^{-2}\%$.
- $\Gamma^{(6)}$: FD **diverges** ($\propto \Delta\sigma^{-5}$ noise); the fit converges to $\sim 0.1\%$.

[1]See Koenigstein et al., 2022 for the exact analytical values.

Outlook: one backward solve, many directions

The $O(N)$ study is a *validated template*. The same adjoint λ — one backward solve — opens directions we could not cover here:

Beyond error estimation

- regulator-shape sensitivity $\rightarrow$ gradient-based *regulator design*
- goal-oriented *adaptive meshing* — cells where the observable cares
- second-order adjoint (Hessian) $\rightarrow$ Newton optimisation & UQ

Validation & scale-up

- benchmark-free checks by *method agreement* (apparent convergence)
- per-observable, per-scale, per-ingredient *error attribution* — a quantitative complement to the modular LEGO® systematic-error programme (Ihssen et al., 2024)
- from the $O(N)$ testbed toward *functional QCD*

The adjoint turns "how sensitive / how accurate / what to change" into *one computation* — ready to plug into production FRG.

Outlook: one backward solve, many directions

The $O(N)$ study is a validated template. The same adjoint λ, opens directions we could not cover here:

Beyond error estimation

- regulator-shape sensitivity $\rightarrow$ gradient-based **regulator design**.
- goal-oriented **adaptive meshing**, cells where the observable cares.
- adjoint of an adjoint (Hessian) $\rightarrow$ Newton optimisation & uncertainty quantification.

Validation & scale-up

- per-observable, per-scale, per-ingredient **error attribution**.
- from the $O(N)$ testbed towards **functional QCD**.

Back-up

Derivation of the adjoint

Back-up: setup

Forward problem: an evolution with a spatial operator L:

$$\partial_t u = L(u), \qquad u(\cdot, 0) = u_0, \qquad t \in [0, T].$$

(FRG: $\partial_t u + \partial_\sigma F[u] = \partial_\sigma Q[\partial_\sigma u]$, $u = \partial_\sigma U$, $t =$ RG-time, $T = t_{\mathrm{IR}}$.)

Observable: a functional of the final state:

$$\mathcal{O} \;=\; g\big(u(\cdot, T)\big).$$

(For $\Gamma^{(4)}$, g is the finite stencil giving $\partial_\sigma^3 u|_0$: a linear functional.)

Goal: the full sensitivity field $\dfrac{\delta \mathcal{O}}{\delta u(\sigma, t)}$: how much the field at every (σ, t) affects the observable.

Back-up: dynamics as a constraint

Adjoin the equation of motion with a multiplier field $\lambda(\sigma, t)$ — the adjoint (or costate):

$$\mathcal{L} \;=\; g\big(u(T)\big) \;-\; \int_0^T \big\langle \lambda,\ \partial_t u - L(u) \,\big\rangle \, dt.$$

- $\big\langle a,\, b\big\rangle = \int a\, b\, d\sigma$ is the spatial inner product.
- On any true solution the bracket vanishes, so $\mathcal{L} = \mathcal{O}$ for every λ.
- λ is a Lagrange multiplier: a conjugate field enforcing the EOM.

λ is free; we choose it later to eliminate the unknown terms.

Back-up: vary the input

Perturb the initial data $u_0 \to u_0 + \delta u$. To first order the perturbation $\delta u(\sigma, t)$ obeys the linearized (tangent) equation

$$\partial_t \, \delta u = A(t) \, \delta u, \qquad A(t) \equiv \frac{\partial L}{\partial u}\bigg|_{u(t)}.$$

Since $\mathcal{L} = \mathcal{O}$ for any λ, vary $\mathcal{L}$. With $\delta g = \big\langle \partial_u g, \, \delta u(T) \big\rangle$ and $\delta L = A \, \delta u$:

$$\delta \mathcal{O} \; = \; \big\langle \partial_u g, \, \delta u(T) \big\rangle \; - \; \int_0^T \big\langle \lambda, \, \partial_t \, \delta u \big\rangle \, dt \; + \; \int_0^T \big\langle \lambda, \, A \, \delta u \big\rangle \, dt.$$

The problem

$\delta u(\sigma, t)$ is unknown: one tangent solve per perturbation. We want δu as a common factor. Two steps achieve this.

Back-up: transpose the operators onto λ

Step 1: integrate by parts in time (shift ∂_t from δu to λ):

$$\int_0^T \langle \lambda,\, \partial_t\, \delta u \rangle\, dt = \langle \lambda(T),\, \delta u(T) \rangle - \langle \lambda(0),\, \delta u_0 \rangle \; - \; \int_0^T \langle \partial_t \lambda,\, \delta u \rangle\, dt.$$

Step 2: transpose in space (A^T defined by $\langle Aa,\, b \rangle = \langle a,\, A^\mathsf{T} b \rangle$):

$$\int_0^T \langle \lambda,\, A\, \delta u \rangle\, dt = \int_0^T \langle A^\mathsf{T} \lambda,\, \delta u \rangle\, dt.$$

Grouping by where δu appears:

$$\delta \mathcal{O} = \langle \partial_u g - \lambda(T),\, \delta u(T) \rangle + \langle \lambda(0),\, \delta u_0 \rangle + \int_0^T \langle \partial_t \lambda + A^\mathsf{T} \lambda,\, \delta u \rangle\, dt.$$

δu is now factored out everywhere and λ is still free.

Back-up: choose λ to eliminate the unknowns

The interior δu and the final $\delta u(T)$ are arbitrary; make their coefficients vanish:

Adjoint problem (backward in time)

$$\boxed{\partial_t \lambda = -\,A(t)^{\mathsf{T}}\,\lambda}, \qquad \boxed{\lambda(\cdot, T) = \frac{\partial g}{\partial u}\Big|_T}$$

Everything then collapses to the boundary term at $t = 0$:

$$\boxed{\delta\mathcal{O} \;=\; \big\langle \lambda(\cdot, 0),\; \delta u_0 \big\rangle} \qquad \Longrightarrow \qquad \frac{\delta\mathcal{O}}{\delta u(\sigma, 0)} = \lambda(\sigma, 0).$$

- Solve the forward problem once (store $u(t)$ to build $A(t)$).
- Solve the adjoint backward once, seeded by the observable.
- Read off the gradient with respect to the whole initial field.

Back-up

Polynomial fit method

Back-up: robust vertex extraction (Savitzky–Golay)

The problem with finite difference extraction: The finite-difference stencils $\Gamma^{(4)} = (u_2 - 2u_1)/\Delta\sigma^3$, $\Gamma^{(6)} = (u_3 - 4u_2 + 5u_1)/\Delta\sigma^5$ difference a few cell values and divide by $\Delta\sigma^5$: condition number $\sim \Delta\sigma^{-5}$, so any small scheme error is amplified and $\Gamma^{(6)}$ grows as $\Delta\sigma \to 0$.

Fix (local odd least-squares fit, Savitzky and Golay, 1964)**.** Near the physical point u is odd and smooth, $u(\sigma) = \sum_{k=0}^{K} c_k\, \sigma^{2k+1}$, so

$$\Gamma^{(2)} = c_0, \qquad \Gamma^{(4)} = 6\, c_1, \qquad \Gamma^{(6)} = 120\, c_2.$$

Fit $\{c_k\}$ by least squares to the cell averages over a **fixed physical window** $\sigma \le \sigma_{\text{fit}}$ (refinement adds points $\Rightarrow$ better conditioned), with the cell-averaged odd-monomial basis.

Result. $\Gamma^{(6)}$ turns from diverging (5000% at $N_\sigma{=}1600$) to convergent ($\sim 0.3\%$); and $\Gamma^{(4)}$ drops to $\sim 10^{-4}$.

Back-up: the fit-consistent adjoint seed

The fit is **linear** in the cell values: collect the fit basis in a matrix A (column $k =$ the monomial σ^{2k+1} averaged over each cell), with $\mathbf{u} = (u_1, \ldots, u_M)$ the near physical point, cell values. Then

$$\mathcal{O} = \text{scale} \cdot c_k, \qquad \mathbf{c} = (A^\mathsf{T} A)^{-1} A^\mathsf{T}\, \mathbf{u} \quad \text{(least-squares fit).}$$

Hence the adjoint input is one row of the (pseudo-)inverse,

$$\boxed{\frac{\partial \mathcal{O}}{\partial u_i} = \text{scale} \cdot \left[A\,(A^\mathsf{T} A)^{-1} e_k\right]_i}$$

- Bounded and spread over the fit window, unlike the FD seed's $\Delta\sigma^{-(2n-1)}$ spike at $\sigma = 0$.
- Validated: $I_{\text{eff}} = 1.02$ (robust across resolution).
- Consequence: η_{fit} is $\sim 20\times$ smaller than η_{FD}, the adjoint explains why the fit reduces the error.

References I

Roland Becker and Rolf Rannacher. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica*, 10:1–102, 2001.

Michael B. Giles and Niles A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65:393–415, 2000.

Eduardo Grossi and Nicolas Wink. Resolving phase transitions with discontinuous Galerkin methods. *SciPost Phys. Core*, 6:071, 2023. doi: 10.21468/SciPostPhysCore.6.4.071.

Friederike Ihssen, Jan M. Pawlowski, Franz R. Sattler, and Nicolas Wink. Towards quantitative precision in functional QCD I. *arXiv preprint*, 2024.

Adrian Koenigstein, Martin J. Steil, Nicolas Wink, Eduardo Grossi, Jens Braun, Michael Buballa, and Dirk H. Rischke. Numerical fluid dynamics for FRG flow equations: Zero-dimensional QFTs as numerical test cases. *Phys. Rev. D*, 106:065012, 2022.

Abraham Savitzky and Marcel J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964.