

Neural Inverse Problems and Degeneracy

SMEFT and Primordial Black Holes

Tae-Geun Kim (金泰瑾)



PARTIQAL 2026

2026.07.14

Inverse Problems


Inverse Problem on SMEFT

Inverse Problem on PBH

Summary

Function

$$f : X \rightarrow Y$$




Well-definedness is all you need. Collisions are allowed:

$$f(x_1) \neq f(x_2) \implies x_1 \neq x_2$$

different parameters can produce the same observable.

Inverse Function

$$f^{-1} : Y \rightarrow X$$



Now y_1 has two preimages. Which one? You need


$$\underbrace{x_1 \neq x_2 \implies f(x_1) \neq f(x_2)}_{\text{injective}} \quad \underbrace{Y = f(X)}_{\text{surjective}}$$

one observable, ambiguous parameters = **degeneracy**.

What is the inverse problem?

A model is a map from parameters C to measurements E . One triplet $(C, \mathcal{M}, E)$ frames every inverse problem.

[Pierre Sabatier, 2000]



C : parameters (causes) · E : measurements (effects)

Direct problem

$$c \in C \xrightarrow{\mathcal{M}} e = \mathcal{M}(c)$$

"Giving $\mathcal{M}$ explicitly."

Inverse problem


$$e \in E \xrightarrow{\mathcal{M}^{-1}} c = ?$$

"Going back from E to C ."

A Century of Inverse Problems

From Abel (1826) to a discipline of its own · Sabatier (2000), "Past and future of inverse problems"

■ Foundations ■ Exact analysis ■ Applied / numerical ■ Inverse methods ■ Field-building



[Concept](#)[Activities](#)[Facilitators](#)[Members](#)[Contact](#)

DEEP-IN SCIENCE

深入 科学

CONCEPT

“DEEP learning for INverse problems (DEEP-IN)” in Sciences Working Group

The essence of discovery in sciences has always been rooted in the reverse engineering of natural phenomena and observational data. This paradigm of deducing the underlying laws of nature from observable outcomes forms the cornerstone of our scientific inquiry. **The DEEP-IN working group is established with the recognition that the elucidation of such complex phenomena demands a fusion of physics insights and advanced deep learning methodologies.**

In response to the evolving landscape of scientific research, our objective is to integrate cutting-edge **deep learning techniques, alongside generative models and other advanced statistical learning methods**, into the toolkit of scientists.

The DEEP-IN working group at **RIKEN-IITEMS** is dedicated to creating an interdisciplinary platform that harnesses the transformative potential of artificial intelligence(AI). This platform is designed to **tackle inverse problems that span a diverse spectrum of sciences, from biology to physics and more in the future.**


Name	Affiliation	Position	Fields	Email	Last Name
☀️Akinori Tanaka	RIKEN-AIP/IITEMS	Senior Research Scientist	Machine Learning, Mathematical Physics	akinori.tanaka@riken.jp	Tanaka
☀️Catherine Beauchemin	RIKEN-IITEMS	Deputy Director	Virophysics	cbeau@riken.jp	Beauchemin
☀️Koji Hashimoto	Kyoto University	Professor	Machine Learning Physics	koji@scphys.kyoto-u.ac.jp	Hashimoto
☀️Lingxiao Wang	RIKEN-IITEMS	Research Scientist	QCD, Deep Learning	lingxiao.wang@riken.jp	Wang
☀️Masato Taki	RIKEN-IITEMS/Rikkyo Uni.	Team Director	Mathematical Physics, Machine Learning	5066866@rikkyo.ac.jp	Taki
Akira Harada	NITIC	Assistant Professor	Space Science	harada@icrr.u-tokyo.ac.jp	Harada
Álvaro Pastor Gutiérrez	RIKEN-IITEMS	SPDR	High-Energy Physics	alvaro.pastor@riken.jp	Gutiérrez
Andreas Ipp	TU Wien	Senior Researcher	Lattice Gauge Theory	andreas.ipp@tuwien.ac.at	Ipp
Enrico Rinaldi	Quantinuum	Lead R&D Scientist	Quantum Simulation	erinaldi.work@gmail.com	Rinaldi
Fu-Peng Li	Fudan Univerisity	Postdoctoral Fellow	Nuclear Physics, AI for Science	fpli@fudan.edu.cn	Li
Gert Aarts	Swansea University	Professor	Lattice QCD, Machine Learning	g.aarts@swansea.ac.uk	Aarts
Haozhao Liang	UTokyo	Associate Professor	Nuclear Physics	hzliang@g.ecc.u-tokyo.ac.jp	Liang
Jinyang Li	KEK/RIKEN-IITEMS	JRA	Quantum Computing	jinyang.li@riken.jp	Li
Kenji Fukushima	UTokyo	Professor	Hadron Physics, QCD	uku@nt.phys.s.u-tokyo.ac.jp	Fukushima
Misaki Ozawa	Univ. Grenoble Alpes	Junior Chair	Statistical Physics	si.Ozawa@univ-grenoble-alp.fr	Ozawa
Qianteng Zhu	RIKEN-IITEMS/SJTU	IPA	Lattice Gauge Theory, Machine Learning	zhuqianteng@sjtu.edu.cn	Zhu
Sho Ozaki	Hirosaki University	Assistant Professor	Medical Physics	sho.ozaki@hirosaki-u.ac.jp	Ozaki
Shuzhe Shi	Tsinghua University	Assistant Professor	High Energy Nuclear Physics	huzhe-shi@tsinghua.edu.cn	Shi
Steffen Backes	RIKEN-TRIP	Expert Technician	Condensed Matter Physics	steffen.backes@riken.jp	Backes
Tae-Geun Kim	Fudan Univerisity	Postdoctoral Fellow	AI for Science	tgkim@fudan.edu.cn	Kim
Takumi Doi	RIKEN-IITEMS	Senior Research Scientist	Lattice QCD, Hadron Physics	doi@ribf.riken.jp	Doi
Tetsuo Hatsuda	RIKEN-IITEMS	Division Director	Hadron Physics, QCD	tetsuo.hatsuda@gmail.com	Hatsuda
Xiao-Yang Wang	RIKEN-IITEMS	Postdoctoral Researcher	Quantum Computing	xiaoyang.wang@riken.jp	Wang
Xingyu Guo	SCNU	Lecturer	Quantum Simulation	guoxy@m.scnu.edu.cn	Guo
Yan Lyu	RIKEN-IITEMS/Berkeley	Postdoctoral Fellow	Lattice QCD, Hadron Physics	yan.lyu@riken.jp	Lyu
Yang-Yang Tan	UTokyo	Postdoctoral Researcher	QCD, Deep Learning	yangyang-tan@foxmail.com	Tan
Yuki Fujimoto	Niigata University	Assistant Professor	Neutron Star, Machine Learning	fujimotophy@gmail.com	Fujimoto
Zhenyu Zhu	Rochester Institute of Technology(RIT)	Postdoctoral Fellow	Neutron Star, Nuclear Astrophysics	zhenyu.zhu@rit.edu	Zhu

Can one hear the shape of a drum?

A membrane Ω clamped along its rim Γ vibrates; its pure tones $F = Ue^{i\omega t}$ are the eigenvalues of the fixed-boundary Helmholtz problem: $\nabla^2 U + \lambda U = 0$ in Ω , $U = 0$ on Γ . [Mark Kac, 1966]

CAN ONE HEAR THE SHAPE OF A DRUM?

Forward is a solvable Laplace problem, but can the inverse be solved?




Can one hear the shape of a drum?

Not exactly, but you still hear its area, perimeter, and holes. [Gordon-Webb-Wolpert, 1992]

SO, CAN YOU HEAR THE SHAPE?

No, not the exact shape — but you do hear a lot

1 YOU CAN HEAR




$$\rightarrow N(\lambda) \sim \left(\frac{\text{Area}}{4\pi} \right) \lambda$$

- ✓ Area
- ✓ Perimeter
- ✓ number of holes

the spectrum fixes these

2 YOU CANNOT HEAR

same sound, different shape



isospectral pair
(Gordon-Webb-Wolpert, 1992)

3 THE VERDICT

NO — shape is not unique

but you DO hear:

- ✓ area
- ✓ perimeter
- ✓ holes



area
perimeter
holes

uniqueness fails, invariants survive

The same triplet, three times. Only $(C, \mathcal{M}, E)$ change; the inverse and its **degeneracy** do not. [Sabatier, 2000]


	Drum spectral geometry	SMEFT UV $\rightarrow$ Wilson coeff.	PBH Hawking evaporation
C · model (unknown)	shape Ω	UV couplings g	mass function $\psi(M)$
$\mathcal{M}$ · forward map	Helmholtz solve	EFT matching	Hawking operator
E · data	tones $\{\lambda_n\}$	Wilson coeffs c	photon spectrum $S(E)$
Inverse · $E \rightarrow C$	tones $\rightarrow$ shape	WC $\rightarrow$ UV	spectrum $\rightarrow \psi$
Degeneracy	isospectral drums	degenerate UV completions	FSR floor

Inverse Problems

Inverse Problem on SMEFT


Inverse Problem on PBH

Summary




Álvaro Pastor Gutiérrez

RIKEN iTHEMS



Lingxiao Wang

RIKEN iTHEMS & UTokyo



Victor Maura

King's College London


Direct searches are limited

ATLAS SUSY Searches* - 95% CL Lower Limits

July 2019

ATLAS Preliminary

$\sqrt{s} = 13$ TeV



*Only a selection of the available mass limits on new states or phenomena is shown. Many of the limits are based on simplified models, c.f. refs. for the assumptions made.

ATLAS Preliminary, $\sqrt{s} = 13$ TeV, Run 2; representative 95% CL limits in simplified models.

What does each row mean?

Model

What is produced and how it decays.

Signature

What the detector sees: jets, leptons, photons, or missing momentum.

Mass limit

The mass range excluded at 95% CL for that model.


These are conditional limits, not universal mass bounds.

And the reach ends near the collision energy: heavier states are simply out of range.


≡ WHY LOOK INDIRECTLY? ≡

No new particles yet, so read their low-energy imprint


① Direct: make it



② Indirect: feel it ✓




③ One common language




WHAT IS AN EFFECTIVE FIELD THEORY?

Low-energy physics without the full high-energy theory

1 Two separated scales



2 Integrate out heavy modes



3 A tower of operators

$$\mathcal{L}_{\text{eff}} = \mathcal{L}_{\text{light}} + \sum \frac{c_i O_i}{\Lambda^n}$$

ordered by dimension = relevance

O_1



O_2

O_3

$(E/\Lambda)^n$
: smaller

4 Model-independent

many UV theories, one EFT



same operator form

unknown physics lives in the c_i

At low energy, a heavy mediator looks like a contact interaction.

$$g^2(p) \frac{1}{p^2 + m_\varphi^2} \xrightarrow{p^2 \ll m_\varphi^2} \frac{C_{g^2}(p)}{\Lambda_{\text{UV}}^2}$$

Λ_{UV}

Effective UV scale. Sets the interaction strength with C .

m_φ

Mass of the heavy mediator that cannot be produced on shell.

$p^2 \ll m_\varphi^2$

The propagator becomes local and loses its momentum dependence.

$$\frac{g^2}{m_\varphi^2} \approx \frac{C}{\Lambda_{\text{UV}}^2}$$

The UV mediator is replaced by a Wilson coefficient in the low-energy theory.

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \sum_i \frac{C_i}{\Lambda^2} \mathcal{O}_i^{(6)}$$

dimension-6 Warsaw basis · $U(2)^5$ flavour · tree-level Granada matching at $m_\varphi = 1$ TeV · scalar & fermion extensions only (no heavy vectors)




which g ? many share one c

Forward $g \rightarrow c$: fixed by the matching, trivial | **Inverse** $c \rightarrow g$: degenerate

Warsaw basis: Grzadkowski et al. 2010 · Granada: de Blas et al. 2017 · tree-level matching via $\mathcal{L}_{\text{sme}}$: Gargalionis et al. 2024 · $U(2)^5$: Greljo et al. 2022

$$c = M\mu(g)$$


$M \in \mathbb{R}^{39 \times 60}$ constant $\cdot$ $\mu(g)$: coupling monomials, degrees 2 to 4



Forward one matrix product | **Inverse** $\text{rank } M \leq 39 < 60$, so $\ker M \neq 0$

Granada dictionary: de Blas et al. 2017 $\cdot$ tree-level matching via 1sme: Gargalionis et al. 2024





EFFECTIVE RANK

[Thibeault et al., Nat. Phys. (2024)]

$$p_i = \sigma_i / \sum_j \sigma_j$$

$$r_{\text{eff}} = \exp\left(-\sum_i p_i \ln p_i\right)$$

$$r_{\text{eff}} \approx 17.6$$

- rank $M = 39$, but $\dim \ker M = 21$ leave no trace.
- σ spans $16.3 \rightarrow 0.18$: only ~ 18 directions are effectively constrained.
- Inverse is doubly hard**: hard null space + soft tail.



4 $\mathcal{B}_0 = \{\lambda_S, \kappa_{S_3}, \lambda_\Xi, \lambda_{\Xi_1}\}$: each g_a^2 leaves no WC trace (zero diagonals).

7 C_φ narrow row: 8 monomials (incl. $\lambda_{\Theta_1}^2, \lambda_\varphi^2$) feed one WC $\rightarrow$ 7 indifferent directions.

10 dispersed redundancies: over-fed WCs,
no single-monomial reading.


$$4 + 7 + 10 = 21$$

No measurement of c resolves these directions. **This is the degeneracy.**


Best Linear Unbiased Estimator

$$\hat{\mu} = M^+ c$$

Moore-Penrose pseudo-inverse: no linear unbiased estimator beats its per-monomial R^2 . [Penrose 1955; Gauss-Markov: Christensen 2011]



26/60 reach $R^2 \geq 0.95$ · 24/60 have $R^2 < 0$ · **escape needs a nonlinear or biased estimator**




loud WCs dominate, quiet informative ones ignored

$$\tilde{\omega}_w = (\hat{\sigma}_w^2 + \tau)^{-p} / \sum_{w'} (\hat{\sigma}_{w'}^2 + \tau)^{-p}$$

normalized inverse-variance weights ($p = 0.5$): every WC weighs equally

Variance weighting rescues quiet WCs



LOUD WC $C_{quqd}^{(1)}, \sigma_w \approx 1.4$

MSE $R^2 = 1.00 \approx$ VW $R^2 = 1.00$

loss choice: no effect


QUIET WC $C_{qq}^{(1)}, \sigma_w \approx 0.01$

MSE $R^2 = 0.85 \rightarrow$ VW $R^2 = 1.00$

loss choice: decisive

Variance weighting bites **only where the scale gap does**: the quiet WCs.

But the inverse?



Well-recovered loud, $\sigma_w \sim 0.8$

High-variance couplings recover cleanly, $R^2 \approx 0.99$.

Degenerate quiet, $\sigma_w \sim 0.06$


The MLP **collapses to** ≈ 0 , worse than the mean:


$$\kappa_{\Xi} \quad R^2 = -1.5$$

$$y_{qq\omega_1} \quad R^2 = -1.4$$

VW loss fixes the **scale**, but a plain MLP cannot resolve the **degenerate directions**. $\rightarrow$ **a new architecture is needed.**


Piecewise-Linear Embedding $\times$ Cross-Attention (X)






Piecewise-Linear Embedding [Gorishniy et al. 2022]

- A raw scalar gets only **one linear projection**: the tabular bottleneck of Transformers.
- PLE lifts each coupling into a per-feature $\text{ReLU}(g_i - b_k)$ basis, a rich token, reaching **GBDT-level** accuracy.




Learnable WC queries [Carion et al. 2020, DETR]

- A shared MLP head makes all WCs **compete** in one pooled representation.
- Instead, **one learnable query per WC**: each attends only to its sparse sector of M (≤ 10 monomials/WC).



Both near-exact (**MLP** 0.999, **PLEX** 0.9999), yet PLEX cuts per-WC error $\sim 10\text{--}50\times$.





01 Train and freeze the forward surrogate

$$f_\theta : g \rightarrow c, \quad f_\theta \approx \mathcal{M}$$

02 Initialize with the inverse model

$$g_0 = h_\varphi(c) \in \mathcal{M}^{-1}(c)$$

03 Null-aware spectral update



$$\delta_t = V \operatorname{diag}\left(\frac{\sigma_k}{\sigma_k^2 + \lambda}\right) U^T \tilde{r}_t, \quad \delta_t \leftarrow \delta_t \cdot \min\left(1, \frac{\tau}{\|\delta_t\|_\infty}\right)$$

$$g_{t+1} = \max(g_t - \eta \delta_t, 0)$$


where

$$r_t = f_\theta(g_t) - c, \quad J_t = (\partial f_\theta / \partial g) |_{g_t}, \quad \tilde{r}_t = W^{1/2} r_t, \quad \tilde{J}_t = W^{1/2} J_t = U \Sigma V^T$$


W : diagonal inverse-variance weighting in WC space. Small σ_k directions receive vanishing spectral weight.








Refinement amplifies the surrogate



Refining against its own biased forward, **MLP+TTR drops below its init** (0.959 $\rightarrow$ 0.935); **PLEX+TTR, against a faithful surrogate, rises** (0.988 $\rightarrow$ 0.999). The basin selector matters as much as the refinement.

Reconstruction solved, UV capped



NNLS: non-negative least squares, a sign-constrained linear baseline. [Lawson & Hanson 1974]

What solved the SMEFT inverse?

A good basin selector + a good refinement

01

Basin selector

PLEX inverse + VW loss

Selects a strong initial UV basin. It does not pretend to identify the null directions.



02

Null-aware refinement

NAGN: Jacobian filter + trust region

Refines only identifiable directions and keeps the surrogate bounded.

RAW UV R^2

0.856 $\rightarrow$ 0.865

+

ROUND-TRIP R^2

0.989 $\rightarrow$ 0.999

The null space is not removed. The method stops forcing an answer where the data are silent.

Inverse Problems

Inverse Problem on SMEFT

Inverse Problem on PBH

Summary

Jeonghwan Park

Yonsei University

Min Gi Park

Yonsei University

Yeji Park


CTPU, IBS

Hyunjoo Jung

Yonsei University

Seong Chan Park

Yonsei University & KIAS






NOT COLLAPSED STARS

- **Origin:** gravitational collapse of large density fluctuations in the early universe, not stellar death.
- **Mass function:** a full distribution $\psi(M)$ set by the formation model, not a single stellar mass scale.
- **Dark matter:** a leading candidate across a wide mass window.


[image: Science Tokyo]

- **Hawking radiation** is thermal emission at $T_H \propto \frac{1}{M}$, filtered by greybody factors.
- Particles lighter than T_H are emitted, then **decay and hadronize** into **secondary photons**.




Log Normal

$$\psi(M|\sigma, M_c) = \frac{1}{\sqrt{2\pi}\sigma M} \exp\left(-\frac{\ln^2(M/M_c)}{2\sigma^2}\right)$$




Critical Collapse

$$\psi(M|\alpha, M_f) = \frac{\alpha^2 M^\alpha}{M_f^{1+\alpha} \Gamma(1/\alpha)} \exp\left(-\left(\frac{M}{M_f}\right)^\alpha\right)$$




(Smooth) Power Law

$$\begin{aligned} \psi(M|M_{\min}, \alpha, \beta) \\ = \frac{\beta M_{\min}^{\alpha-1}}{\Gamma((\alpha-1)/\beta)} M^{-\alpha} \exp\left(-\left(\frac{M_{\min}}{M}\right)^\beta\right) \end{aligned}$$




$$S(E) = \mathfrak{H}[\psi](E) = \int_{M_{\min}}^{M_{\max}} \frac{d^2 N_{\gamma}^{\text{sec}}}{dE dt} \psi(M) dM$$

$\psi(M)$




$S(E)$



- Consider an operator $G : \mathcal{F} \rightarrow \mathcal{G}$, where $f(x) \in \mathcal{F}$ and $g(y) \in \mathcal{G}$ are functions.


Discretized Input Function



GLAON (GLA + Operator Network)


: Gated Linear Attention encoder
+ Transformer decoder

[Gated Linear Attention: Yang et al., ICML (2024)]



Stacy: combine known mass functions


$$\psi_{\text{Stacy}} = \frac{|\beta|}{M_s \Gamma(\alpha)} (M/M_s)^{\alpha\beta-1} \exp(-(M/M_s)^\beta)$$





GBP: Stacy + evaporation tails

$$\psi_{\text{GBP}} = \frac{\beta}{M_s B(\alpha, \gamma)} (M/M_s)^{\alpha\beta-1} [1 + (M/M_s)^\beta]^{-(\alpha+\gamma)}$$



both tails power-law, $\psi \propto M^{\alpha\beta-1}$ & $M^{-1-\beta\gamma}$; Stacy = $\gamma \rightarrow \infty$ limit





Log Normal




Critical Collapse





Smooth Power-law









Log Normal




Critical Collapse




Smooth Power-law






Degeneracy: accurate reconstruction




WHY TEST-TIME REFINEMENT BREAKS HERE

Refining psi node-by-node forgets it is a function, so you get a wiggly implausible mass function


① Recall the idea




② But psi is discrete



③ Pointwise update



④ Result: wiggly psi



Decoder-manifold Constrained Laplace Approximation

01 Decoder-manifold MAP

$$\begin{aligned}\psi(m) &= D_\varphi(m), \quad m_0 = E_\varphi(S_{\text{obs}}) \\ \hat{\delta} &= \underset{\delta}{\operatorname{argmin}} \left\| \mathfrak{H}_\theta(D_\varphi(m_0 + \delta)) - S_{\text{obs}} \right\|_2^2 + \lambda \|\delta\|_2^2 \\ \hat{m} &= m_0 + \hat{\delta}\end{aligned}$$

- ψ = decoder output
- Encoder sets the initial m_0
- MAP fits S_{obs}

02 Rank-truncated Laplace in memory space

$$\begin{aligned}J_m &= \frac{\partial(\mathfrak{H}_\theta \circ D_\varphi)}{\partial m} \Big|_{\hat{m}} = U\Sigma V^T \\ C_m &= (\lambda I + J_m^T J_m / \sigma_n^2)^{-1} \\ \tilde{m}^{(k)} &= \hat{m} + V_r(\lambda I + \Sigma_r^2 / \sigma_n^2)^{-1/2} \varepsilon^{(k)}, \quad \varepsilon^{(k)} \sim \mathcal{N}(0, I_r)\end{aligned}$$

- Curvature $J_m \rightarrow$ posterior width
- Keep top- r directions
- Sample latent m , not ψ


03 Decode and refine each draw


$$\begin{aligned}m_*^{(k)} &= \underset{m}{\operatorname{argmin}} \left\| \mathfrak{H}_\theta(D_\varphi(m)) - S_{\text{obs}} \right\|_2^2 + \lambda \|m - \tilde{m}^{(k)}\|_2^2 \\ \psi^{(k)} &= D_\varphi(m_*^{(k)}), \quad k = 1, \dots, K\end{aligned}$$

- Re-fit each draw near $\tilde{m}^{(k)}$
- Every $\psi^{(k)}$ stays smooth
- Ensemble $\rightarrow$ uncertainty band


DeCLA: UNCERTAINTY ON THE DECODER MANIFOLD

Sample in a latent memory space, decode each draw into a smooth mass function





LogOscillatory



Plain TTR wiggles → **stay on the manifold** + **quantify the rest**

01

Decoder manifold

$$\psi = D_{\varphi}(m)$$

Every draw is a decoder output, so ψ stays a smooth, plausible function, no pointwise wiggles.



02

Constrained Laplace UQ

DeCLA: latent Laplace + refine

Samples and refines in latent memory, then decodes an ensemble, a band where the FSR floor leaves ψ unresolved.

NAIVE TTR

wiggly $\psi(m_i)$, fits data not physics



DeCLA

smooth ψ ensemble with honest band


The FSR floor is not removed. DeCLA reports a band instead of forcing one ψ where the data are silent.

Inverse Problems

Inverse Problem on SMEFT

Inverse Problem on PBH

Summary



[My Webpage](#)

- **Degeneracy forces a bias.** Every inverse problem is a triplet $(C, \mathcal{M}, E)$ with a built-in **degeneracy** (isospectral drums, the SMEFT null space, the PBH FSR floor). There the data are silent, so a **bias (prior)** is required to answer at all.
- **SMEFT** ($c \rightarrow g$), a **data-driven bias**: NAGN refines only the identifiable directions and holds the null space at the learned inverse's prior, a biased **point estimate**.
- **PBH** ($S(E) \rightarrow \psi$), a **manifold bias**: DeCLA confines ψ to the learned decoder manifold (plausible mass functions) and reports the residual uncertainty as a **band** on it.
- **Takeaway.** A bias is unavoidable, so make it physically meaningful, and stay honest about what the data cannot decide.