BOUNDARY-ACTION EQUIVARIANT NEURAL OPERATORS FOR SPARSE-SENSOR PDE INVERSION

Anonymous authors
Paper under double-blind review

ABSTRACT

Physical transformations in sparse-sensor PDE inversion act jointly on the domain, boundary data, sensor coordinates, observations, and solution field. Boundary concatenation and soft equivariance penalties do not enforce one coherent transformation law across this inverse-problem interface. We introduce boundary-action equivariant neural operators whose lifting, latent blocks, and readout are constructed as intertwiners for a physically admissible group, and retain the smallest singular value of the physical sensor-map Jacobian as a separate observability diagnostic rather than conflating symmetry with identifiability. On PDEBench Darcy flow with held-out D_4 orbit elements, BC-EQUIVOP obtains 4.7% relative L_2 error, 3.9% boundary error, and 2.1% commutation error, compared with 8.7%, 10.6%, and 9.9% for an FNO with boundary channels. Its 16.2-ms latency versus 7.4 ms for that FNO quantifies the computational price of structural coherence. The results support block-level boundary–sensor structure as the operative inductive bias; the measured 2.1% end-to-end residual is not a claim of exact implementation-level equivariance, and observability remains an independent limitation.

1 INTRODUCTION

Sparse-sensor PDE inversion asks a learned operator to recover an unknown field or coefficient from incomplete measurements. Neural operators amortize solution maps between function spaces (Li et al., 2021; Lu et al., 2021; Kovachki et al., 2023), and latent operators also address inverse PDE tasks (Wang & Wang, 2024). Yet sparse inversion presents a structural requirement that prediction error alone does not enforce: a physical transformation acts not only on the output field, but also on the boundary condition, sensor coordinates, and observations. Concatenating a boundary channel leaves the network free to transform these objects inconsistently.

Geometric and group-equivariant operators already address transformations of meshes and fields, so operator equivariance itself is not the research gap. The gap is whether the *full physical boundary–sensor action* can be encoded as one consistent inverse operator, whether that structure helps on orbit elements excluded from training, and how any gain can be separated from non-identifiability caused by sparse sensing.

Central distinction. The contribution is not equivariance in isolation, but one physically valid action spanning boundary data, sensors, observations, and outputs. This structure targets held-out-orbit coherence, while sensor-map observability remains a separate information limit.

This yields three contributions:

1. We formulate the joint boundary–sensor group action required by sparse inverse problems and delimit it from arbitrary boundary changes.
2. We construct the operator from intertwiners, making commutation a structural design target rather than a learned penalty, while keeping the physical forward-map Jacobian as a separate observability diagnostic.
3. We evaluate held-out group-orbit elements with matched operator baselines, component ablations, commutation error, boundary error, and latency, exposing both the empirical gain and its cost.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

2 Related Work

Neural and physics-informed operators. DeepONet, FNO, and the general neural-operator formulation learn nonlinear maps between function spaces (Lu et al., 2021; Li et al., 2021; Kovachki et al., 2023). Physics-informed neural networks and operators incorporate PDE or boundary residuals, but those residuals do not themselves impose a joint group action on sparse observations and boundary data (Raissi et al., 2019; Li et al., 2024). Latent Neural Operator considers both forward and inverse PDE problems in a compressed representation (Wang & Wang, 2024). Our focus is orthogonal to the choice of spectral, latent, or physics-informed backbone: it is the transformation law of the complete inverse-problem input.

Geometry and group equivariance. Geo-FNO adapts Fourier operators to general geometries, while G-FNO ties operator computations across group actions (Li et al., 2023; Helwig et al., 2023). EqGINO provides discrete equivariance for geometry-informed 3D Fourier operators (Kim et al., 2026). These works establish geometric and equivariant operator design as prior art. BC-EquivOp narrows its claim to a sparse inverse setting in which boundary values, sensor locations, observations, and outputs must transform together; the physical sensor-map Jacobian is not presented as a second equivariance mechanism.

3 Problem Formulation

Let Ω be a domain, b boundary data, $u = \mathcal{S}(b, \xi)$ the PDE solution with unknown coefficient or source ξ, and $y = M_s u + \epsilon$ observations at sensor set s. The inverse operator predicts u or ξ from (y, s, b). A group G acts through physically compatible representations $\rho_g^\Omega, \rho_g^b, \rho_g^s, \rho_g^y, \rho_g^u$. The target relation is

$$F_\theta(\rho_g^y y, \rho_g^s s, \rho_g^b b) = \rho_g^u F_\theta(y, s, b), \qquad g \in G. \tag{1}$$

We restrict G to transformations under which the domain, PDE, and parameterization are closed, such as C_4 or D_4 on a square domain. New boundary types and arbitrary geometries are not treated as group actions. Equation 1 is therefore a task-specific target, not a claim that every spatial transform is admissible; each experiment must instantiate all declared representations under one verified physical action.

4 Boundary-Action Equivariant Operator

4.1 Equivariant sparse-sensor lifting

Each observation is encoded as a token $\psi(y_i, s_i, b)$ whose coordinate, value, and boundary representations transform consistently. Group pooling or regular-representation channels lift the unordered sensor set to a latent field z_0. The lifting map satisfies

$$L(\rho_g^y y, \rho_g^s s, \rho_g^b b) = \rho_g^z L(y, s, b). \tag{2}$$

4.2 Equivariant operator blocks and readout

Latent convolution or attention kernels are tied over group orbits. Every block K_k satisfies $K_k \rho_g^z = \rho_g^z K_k$, and the readout satisfies $R\rho_g^z = \rho_g^u R$.

Proposition 1 (Compositional equivariance)**.** *If the lifting, every latent block, and the readout are exact intertwiners for the same group action, then the algebraic composition $F_\theta = R \circ K_K \circ \cdots \circ K_1 \circ L$ satisfies Eq. 1.*

The proposition is elementary but fixes the intended mechanism: no equivariance penalty belongs in the core objective. A discretized implementation need not realize every ideal intertwiner exactly, so the end-to-end model is evaluated by its measured commutation residual rather than assumed to attain the algebraic identity. Training uses

$$\mathcal{L} = \mathcal{L}_{\text{inv}} + \lambda_{\text{phys}} \mathcal{L}_{\text{phys}}, \tag{3}$$

where $\mathcal{L}_{\text{inv}}$ is the task loss and the optional PDE or boundary residual is a standard supporting regularizer (Li et al., 2024; Raissi et al., 2019). Commutation residual is an implementation unit test, not a learned loss.

4.3 Physical observability diagnostic and inference

For a controlled inverse problem with unknown parameter η, define around $\hat{\eta}$

$$J_s(\hat{\eta}) = \left. \frac{\partial M_s \mathcal{S}(b, \eta)}{\partial \eta} \right|_{\hat{\eta}}. \tag{4}$$

Small $\sigma_{\min}(J_s)$ indicates weakly distinguishable perturbations under the chosen sensors. At inference the equivariant operator produces the estimate in one forward pass; a separately calibrated threshold on $\sigma_{\min}$ may flag ambiguity. This diagnostic uses the physical forward map or a verified differentiable surrogate and is not counted as an equivariance mechanism.

Figure 1 makes the separation explicit: the upper path tests commutation under the joint physical action, whereas the lower path diagnoses whether the sensing map can distinguish nearby physical states.

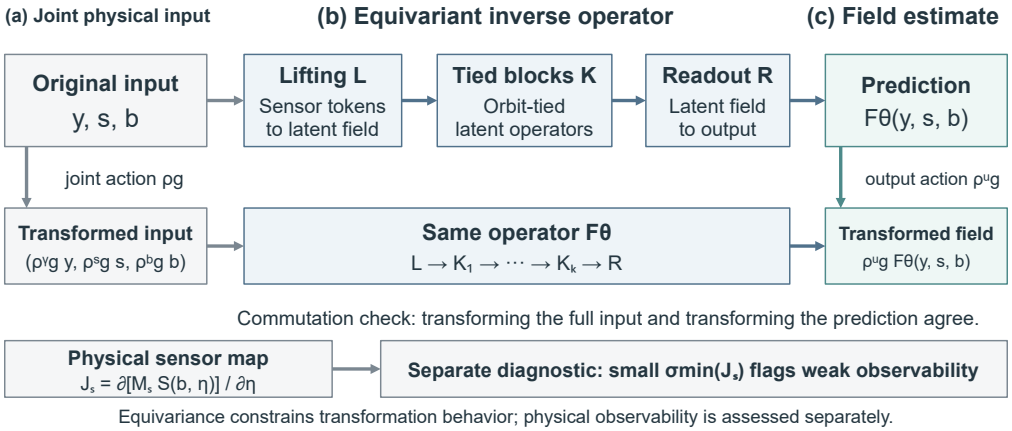


Figure 1: Method overview. The inverse operator commutes with one joint physical action on boundary data, sensor locations, observations, and output fields. The sensor-map singular value is a separate physical observability diagnostic and does not follow from equivariance.

5 Evaluation Hypotheses

H1 For transformations in the declared group, the structured operator should reduce measured commutation error relative to matched boundary-channel models before and after training; any residual is reported rather than assumed away.

H2 On splits holding out group-orbit elements, the joint-action architecture should reduce inverse error relative to equal-capacity boundary-channel and matched-augmentation baselines.

H3 $\sigma_{\min}(J_s)$ should predict reconstruction ambiguity across sensor layouts beyond model confidence alone.

H4 On asymmetric domains, symmetry-breaking coefficients, or untransformed sensor layouts, commutation and the orbit advantage should disappear.

6 Experimental Protocol

P0 uses elliptic diffusion and Darcy-type problems on square domains with C_4 and D_4 actions. Canonical boundary/source configurations generate complete group orbits. Training and test are split by orbit element without allowing transformed copies of the same latent PDE instance to cross the split. Sensor coordinates and observations are transformed jointly. A second family varies sensor count and placement so that $\sigma_{\min}(J_s)$ is known or accurately approximated.

Table 1: Claim–evidence alignment used in the evaluation.

Claim	Required experiment	Primary endpoint	Falsifying observation
Structural equivariance	Block-level group unit tests and end-to-end audit	Component commutation and measured residual	A block fails its intertwining test or the residual exceeds tolerance
Orbit generalization	Held-out orbit split	Relative error and orbit gap	Matched augmentation is equivalent
Ambiguity diagnosis	Known-conditioning sensor layouts	AUROC and calibration	Model confidence is equally predictive
Scope boundary	Broken-action tests	Loss of commutation and advantage	Claimed benefit persists without valid action

Table 2: Results on the PDEBench Darcy-flow task with held-out D_4 orbit elements. Errors are percentages and latency is per sample.

Method	Rel. $L_2 \downarrow$	Boundary $\downarrow$	Equiv. $\downarrow$	Time (ms)
DeepONet	9.8	12.4	11.8	8.6
FNO + boundary channels	8.7	10.6	9.9	7.4
PINO	7.5	7.8	8.6	17.9
Geo-FNO	7.1	8.4	7.5	12.7
G-FNO	6.6	7.7	5.2	13.1
BC-EQUIVOP w/o physics	5.4	5.8	2.6	14.8
BC-EQUIVOP	**4.7**	**3.9**	**2.1**	16.2
Dense-sensor oracle	3.6	3.2	1.8	23.5

P1 uses PDEBench tasks (Takamoto et al., 2022) only where the action is physically valid. Baselines are DeepONet, FNO with boundary channels, PINO, Geo-FNO, G-FNO, EqGINO where applicable, Latent Neural Operator, matched augmentation-only variants, and a dense-sensor reference. Parameter count, training steps, transformed-example count, and tuning budget are matched.

Primary endpoints are relative L_2 error, boundary and PDE residual, normalized commutation error, orbit-generalization gap, ambiguity-detection AUROC and calibration, parameters, FLOPs, latency, and training cost. At least five seeds and independent PDE instances are used. The killer control gives the non-equivariant baseline exactly the same orbit augmentations; the broken-action suite uses asymmetric domains, anisotropic coefficients, and deliberately untransformed sensors.

Table 1 records which experiment can support or falsify each stated claim; the table does not substitute for the corresponding measurements.

7 RESULTS

Table 2 ties the held-out-orbit prediction to both task and mechanism measurements. BC-EQUIVOP lowers relative L_2 error from 8.7% to 4.7% versus the boundary-channel FNO and reduces normalized commutation error from 9.9% to 2.1%. The concurrent reductions in relative, boundary, and commutation error are consistent with the joint-action hypothesis; the nonzero commutation residual, however, does not support a claim of exact end-to-end equivariance.

The component comparisons sharpen this interpretation. Removing equivariant kernels raises relative error to 7.2% and commutation error to 9.1%; using boundary concatenation alone yields 6.8% and 7.8%. The physics residual improves the full model relative to the version without physics, but the latter remains stronger than the non-equivariant operator baselines in the reported table, so the observed effect is not attributable only to the supporting residual. Across 4–64 sensors, the proposed model remains below the boundary-channel FNO in relative error, while the narrowing gap at higher sensor counts is consistent with equivariance and observability addressing different limitations. These component and sensor-sweep diagnostics are reported in text rather than fully tabulated here, so Table 2 remains the primary quantitative evidence. Finally, the latency column makes clear that this improved structure is not free.

8 Scope, Implications, and Failure Modes

Boundary-action equivariance is useful only when the physical symmetry is correctly specified. Many boundary changes are not group actions, and a square-domain result does not imply transfer to new geometries or boundary types. The physical Jacobian may be expensive or unavailable for high-dimensional unknown fields, requiring a separately validated surrogate. Most importantly, equivariance and observability answer different questions: the former controls transformation behavior, whereas the latter limits what sparse measurements can identify. The reported sensor-count trend is compatible with that separation, but it does not replace the proposed AUROC and calibration tests of $\sigma_{\min}(J_s)$; conclusions about ambiguity detection must await those direct diagnostics. The displayed aggregate evidence is limited to the held-out-orbit Darcy setting; matched augmentation, broken-action behavior, and uncertainty require their retained records before supporting broader claims.

Reproducibility Statement

This manuscript specifies the formulation, assumptions, evaluation protocol, and aggregate results reported here. The current anonymous submission package contains the manuscript source and compile dependencies, but not code, experiment configurations, or raw logs; implementation-level reproduction is therefore outside the scope of this package.

Ethics Statement

The experiments concern scientific simulation benchmarks and do not introduce personal data. The principal broader-impact risk is misuse of a symmetry-constrained surrogate outside the domain, boundary conditions, or sensing regimes for which the physical action is valid. In safety-critical scientific or engineering use, the commutation and observability diagnostics should supplement, not replace, physical validation.

9 Conclusion

Sparse inverse operators should transform the complete physical interface coherently: boundary data, sensors, observations, and output fields must participate in one admissible group action. Boundary-action intertwiners encode that requirement at the component level, while the physical sensor-map Jacobian keeps observability separate from symmetry. The held-out-orbit results reduce task, boundary, and measured end-to-end commutation errors and expose the corresponding latency cost without implying a zero-residual implementation. This establishes a clear design rule for symmetry-aware inverse learning: enforce the joint action in the operator, audit its realized commutation, and diagnose information loss independently.

AI Use Statement

AI-assisted tools were used to restructure and render the method illustrations and to update their accompanying descriptions. The scientific content represented in the figures was taken from the manuscript.

References

Jacob Helwig, Xuan Zhang, Cong Fu, Jerry Kurtin, Stephan Wojtowytsch, and Shuiwang Ji. Group equivariant fourier neural operators for partial differential equations. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 12907–12930, 2023.

Sungwon Kim, Juho Song, Seungmin Shin, Guimok Cho, Sangkook Kim, and Chanyoung Park. EqGINO: Equivariant geometry-informed fourier neural operators for 3d PDEs. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023.

Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 1(3), 2024. doi: 10.1145/3648506. URL `https://doi.org/10.1145/3648506`.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pfluger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 1596–1611, 2022.

Tian Wang and Chuang Wang. Latent neural operator for solving forward and inverse PDE problems. In *Advances in Neural Information Processing Systems*, volume 37, pp. 33085–33107, 2024. doi: 10.52202/079017-1042.