Proceedings World Geothermal Congress 2026
Calgary, Canada, June 8 – 11, 2026

Accelerating Geothermal Modeling with Low- and High-Fidelity Fourier Neural Operators

James W. Patterson[1],[2*]
[1] Earth Sciences New Zealand (formerly GNS Science), Wairakei, New Zealand;
[2*] Geothermal Energy and Geofluids Group, ETH Zürich, Switzerland
James.Patterson@eaps.ethz.ch

Keywords: Reservoir Simulation, Surrogate Model, Proxy Model, AI/ML, Fourier Neural Operators, Natural State

ABSTRACT

Geothermal reservoir models are costly to build and calibrate, and generating a single forecast can take hours. Operators tasked with field planning and optimization are constrained by the speed of these forecast simulations, limiting the number of scenarios they can explore. Machine learning can be a powerful tool to speed up computationally expensive tasks, but standard approaches using Neural Networks (NNs) struggle to produce accurate results and require thousands of high-fidelity simulation outputs for training, an impractical expense in geothermal modeling workflows. Fourier Neural Operators (FNOs) have recently emerged as a compelling alternative, outperforming NNs in many physical modeling tasks. FNOs have been shown to produce fast and accurate first-order estimates of oil and gas reservoirs with high spatial fidelity. Furthermore, FNOs can generalize between multiple resolutions, allowing surrogate models trained on low-resolution simulation data to make high-resolution predictions. We present a FNO surrogate model of a geothermal reservoir at multiple resolutions, including natural state production phase simulations. Once trained, these surrogate models generate results in ~0.4 ms, whereas actual simulations require ~5 minutes, enabling rapid scenario testing and accelerated model calibration. These tools are intended to compliment, rather than replace, traditional geothermal simulation workflows and enable faster, more flexible reservoir forecasting.

1 INTRODUCTION

Geothermal energy is a reliable and renewable resource that can provide baseload power generation, yet its commercial development is constrained by geological uncertainty and high upfront costs. Reservoir simulation is a central tool for assessing resource potential and optimizing production, supporting tasks such as natural-state calibration, history matching, and long-term forecasting (O'Sullivan et al., 2001; Grant & Bixley, 2011). These simulations, however, are computationally intensive: building and calibrating a reservoir model can take months, and generating a single forward forecast often requires hours. For operators tasked with field planning, this computational burden limits the number of realizations that can be explored, hindering scenario testing and uncertainty quantification.

Machine learning (ML) methods have been proposed as surrogates for accelerating computationally expensive workflows in subsurface modeling. Standard neural networks (NNs), while effective in some approximation tasks, face two key limitations in physical modeling: (i) they require large volumes of high-fidelity data for training, which is impractical to generate in geothermal workflows, and (ii) they struggle to capture the high-dimensional spatial and temporal fidelity of fluid and heat flow (Rahaman et al., 2019; Karniadakis et al., 2021; Almajid & Abu-Al-Saud, 2022). These limitations reduce their usefulness as reliable forecasting tools in geoscientific applications.

Fourier Neural Operators (FNOs) have recently emerged as a promising alternative. Unlike conventional ML models that learn pointwise mappings, FNOs learn mappings between function spaces, enabling them to capture complex spatiotemporal fields with high accuracy. FNOs are resolution-independent, can generalize across discretizations, and have been shown to provide order-of-magnitude speedups over traditional solvers while retaining high spatial fidelity (Li et al., 2021; Kovachki et al., 2023). Theoretical work has further established FNOs as universal approximators for operators, with favorable error bounds and computational scaling (Kovachki et al., 2021).

These advantages have spurred early applications in the subsurface domain. FNO surrogates have been successfully developed for multiphase flow in petroleum reservoirs (Zhang et al., 2022), inverse modeling in hydrogeology (Guo et al., 2024), and CO_2 storage in realistic geologies (Chandra et al., 2025). More recently, Tang

et al. (2024) demonstrated a multi-fidelity FNO approach for large-scale geological carbon storage, where models trained on abundant low-resolution data and limited high-fidelity simulations achieved similar accuracy with up to 81% lower data-generation cost—a promising paradigm for efficient training regimes in expensive physical simulations. These studies collectively showcase that FNOs can deliver both high accuracy and computational efficiency in modeling complex, high-dimensional flow processes, outperforming traditional NNs and physics-informed methods.

This work presents FNO surrogate models for modeling geothermal reservoir temperatures and pressures, including natural state and production phases. Two training data sets are made: low-resolution simulation data using coarse gridded simulation results and high-resolution data using finely gridded results. Once trained, the surrogates generate forecasts in fractions of a second rather than minutes, enabling rapid scenario testing and accelerated model calibration. These developments are designed to complement conventional geothermal simulation workflows, enabling faster calibration, more flexible scenario exploration, and ultimately more efficient decision-making in geothermal resource development.

2 METHODOLOGY

2.1 Reservoir Model

The training dataset consists of TOUGH2 reservoir simulation results from a synthetic, 3D box model with varying permeability and upflow source locations/magnitudes. The model has dimensions of 11 km by 11 km and spans elevations -2500 m to 300 m ASL. The sides are no-flow boundaries; the bottom is no flow except for hot upflow sources; the top is open to atmosphere and rain source terms (1 m/year, 10% infiltration rate). Hot upflow at a total rate of 18 kg/s and an enthalpy of 1400 kJ/kg is prescribed at the bottom of the reservoir in a 3 x 3 km square at the center of the model. Low-permeability rock (k = 5E-17 m2) rock forms a 3 km buffer zone from the lateral boundaries, leaving a 5 km by 5 km permeable reservoir in the center of the model. This permeable zone is divided into 1 km x 1 km squares of varying permeability. Vertically, this permeable zone is divided into four regions which are, from bottom to top: 1000 m, 1000 m, 300 m, and 500 m thick. Together, there are 100 variable permeability blocks in the system. Spatially correlated, log-normally distributed permeability fields were generated by sampling multivariate Gaussian random fields with a specified covariance model, using a Cholesky decomposition of the covariance matrix. The log-mean permeability is 1E-15 m2 and standard deviation is 0.61, such that 98% of permeability values fall within 1E-16.5 and 1E-13.5. All permeabilities are isotropic. Correlation lengths of 1000 m in the x-direction and 4000 m in the y- and z- directions were used, resulting in relatively weakly correlated permeability fields in the x-direction.

After the permeability field was created, two reservoir meshes were used: a coarse grid (16 × 16 × 9 cells) with cell dimensions of 500 × 500 m, and a fine grid (32 × 32 × 9 cells) with 250 × 250 m cells; both used the same vertical discretization. Mesh edges align with the variable permeability zones described above, ensuring identical rock permeability fields for both meshes.

Natural state simulations were run for 150,000 years to ensure steady state. The resulting pressures and temperatures at each cell were used as the initial conditions for the production model, in which three production and three injection wells were added to the system. Well locations were drawn from a uniform distribution, with producers constrained to the western half of the permeable reservoir and injectors to the eastern half. A minimum spacing of 1000 m was enforced. Production well mass extraction rates were constant and sampled from a uniform distribution between 30 to 40 kg/s; injection well rates were assumed to be 80% of the production rates, accounting for steam mass loss at the surface. Wellbore models were not included. Production phase results were collected after 10 years of simulations time.

A total of 1000 simulations of natural state and production were run for each mesh. No checks were made to ensure that well placement coincides with sufficient permeability for the prescribed production/injection rate, but a TOUGH2 run failure (due to such incompatibilities) triggered a new permeability/well placement/flow rate realization to replace it. Simulations were run in parallel and results were processed/collected into a single H5 file containing information on reservoir permeability, sources/sinks, and temperatures and permeabilities for natural state and production models.

2.2 FNO Model Training

Natural-state simulations are a key step in reservoir modeling, providing pressure, temperature, and flow fields that are consistent with permeability and boundary conditions. Temperature distributions in natural-state models typically show greater variability than pressure, whereas production simulations exhibit rapid pressure changes in response to injection and production, with temperature evolving more gradually. The results from the

natural state simulations serve as the input data for the production simulations. Therefore, two FNO model types were created: one to predict Natural State temperatures and another to predict Production Phase pressure changes.

The Natural State FNO models were trained using permeability fields as inputs and the end-state temperature distributions as outputs. The permeability field alone is sufficient to determine the equilibrium temperature field, as locations, rates, and enthalpies of hot upflow sources were held constant across model realizations.

The Production Phase FNO models were trained on permeability fields and production/injection well locations and flow rates as inputs. Transient pressure changes after 10 years were output. The production/injection flow rate field contained only six non-zero values (3 producers, 3 injectors), making it the kind of high-frequency feature that is poorly captured by the truncated Fourier basis and leading to model instability and poor generalization. Therefore, source/sink input fields were smoothed with a Gaussian kernel using SciPy (Virtanen et al. 2020) to distribute these point sources over a small region near the wells, making the inputs more compatible with the FNO's spectral representation. Note that production and injection wells were located away from the model boundaries and therefore the source/sink pad consists only of zeroes. Input data were also padded using the Torch padding function (torch.nn.functional.pad) in replicate mode, extending the domain by replicating boundary element values to padding cells.

The FNO model was constructed using the Python NeuralOperator package. The FNO model was configured to retain up to (4, 16, 16) Fourier modes in the vertical and horizontal directions, 32 hidden channels, and a projection channel ratio of 2. Data were split into training (80%), validation (10%), and test data (10%). Training was performed using the Adam optimizer with an initial learning rate of 1E-3, decayed via a cosine annealing schedule down to 1E-4. For Natural State temperatures, a mean squared error was used as the loss function; for Production Phase pressures, a weighted mean squared error was used instead, where cells with a non-zero source/sink term was given double weight. Tests showed this weighting scheme to significantly improve production pressure predictions near wells without significantly sacrificing overall field prediction performance.

The two FNO surrogate model types (natural state temperature and production phase pressure) were trained on one of two datasets: 800 realizations of the coarse model (LF: low-fidelity) and 800 realizations of the fine model (HF: high-fidelity). All FNO models were trained over 100 epochs without early stopping, which was observed to be sufficient for training and validation losses to converge. Both FNO models were tested using the 100 fine model test realizations withheld from training and validation to ensure consistent comparison. That is, the LF model was used to predict HF data, despite never being trained on any HF data. Note that the same permeability field realizations are used for both coarse and fine mesh simulations, resulting in paired output. It is therefore important to use the same train/validation/test splits for both LF and HF FNO models, otherwise the LF model will be tested on permeability fields it has already "seen" during training.

3 RESULTS

3.1 Natural State Temperature

TOUGH2 Natural State simulations results show a temperature distribution between 10 and 310 °C across the 100 test data realizations. Hot inflow from the bottom of the reservoir flows upwards, following high permeability channels created from the spatially correlated permeability distributions and heating up those elements. Low permeability regions may remain cooler, as flow is diverted elsewhere. High-permeability elements may also facilitate downward convective flow, becoming cooler than the surrounding elements. Thus, Natural State temperature is a non-linear and complex problem, where spatial connectivity of elements to the surface or to the hot upflow region is critical. Each high-resolution Natural State simulation took ~2 minutes on average.

The mismatch between TOUGH2 results and the FNO model predictions is calculated for every element within the permeable reservoir (Figure 1). Although the Low-Fidelity FNO model was never trained on results from the finely meshed reservoir model, it performs relatively well, with 90% of cells within -24.3 and +38.8 °C of the test data's simulations. The Low-Fidelity predictions are biased by +4.7 °C. The High-Fidelity FNO performs slightly better, with 90% of predictions within -29.4 and +25.3 °C, and with a bias of less than +0.1 °C. The mismatches between simulations and surrogate predictions generally fall around the 1:1 line for both FNO models, with temperature mismatch highest at the middle temperatures/depths of the model and lower at the top and bottom boundaries. Each FNO surrogate model prediction takes ~0.2 ms, 600,000 times faster than the simulations.

Visual inspection of the results shows that both FNO models capture the general shape of the hot upflow zone, but significant areas of temperature mismatch exist. TOUGH2 simulation permeability and temperatures are shown together with FNO model predictions and differences between the two for a given slice of the reservoir

(Figure 2). These regions are relatively contiguous in space (often >3000 m wide) and are somewhat consistent across the Low- and High-Fidelity FNOs.

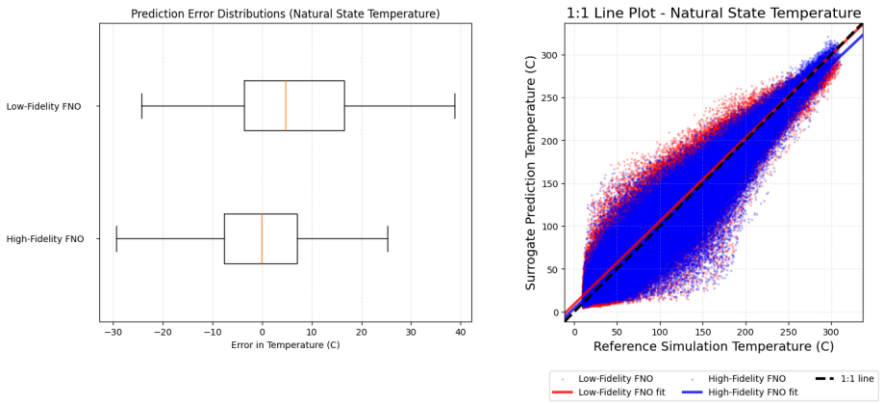


Figure 1. Boxplots and 1:1 line plot of Natural State temperature test data mismatch for each FNO surrogate model. Boxes span the 25th to 75th percentiles; whiskers span the 5th to 95th percentiles.

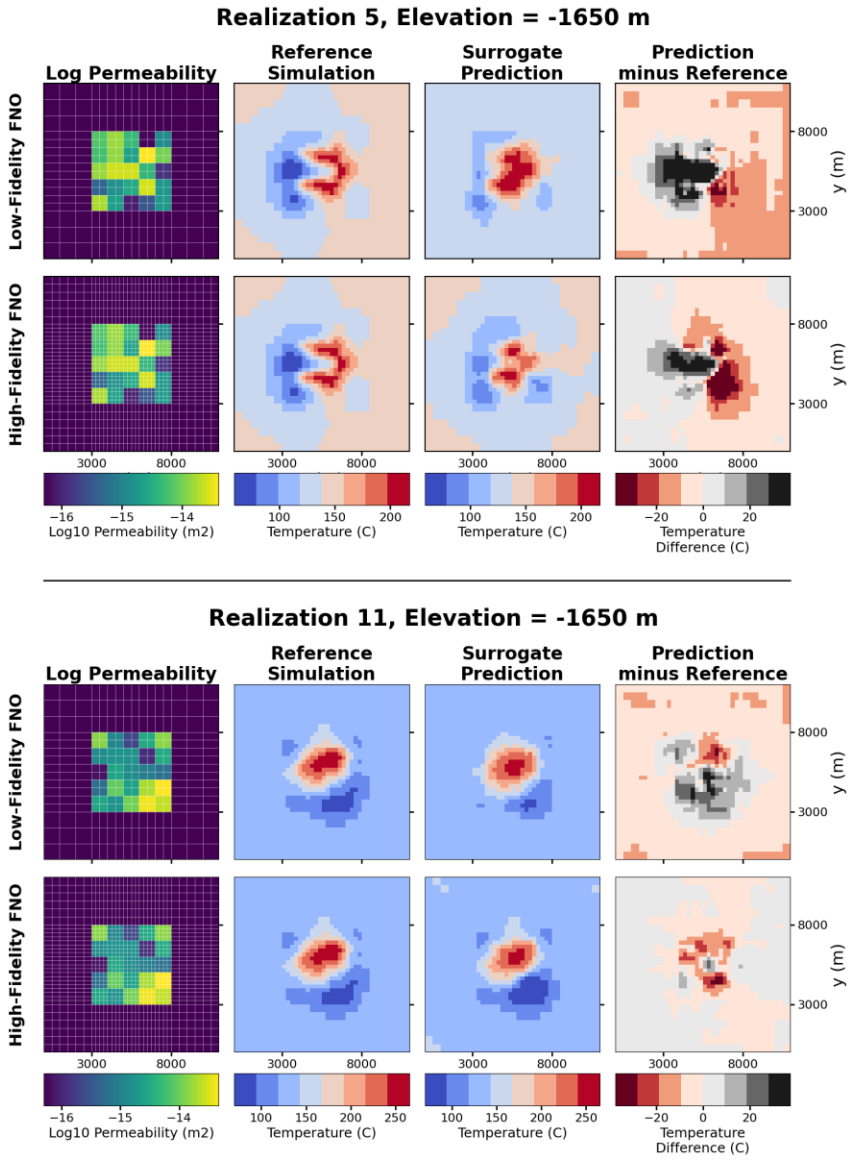


Figure 2. Horizontal slices through the reservoir (map-view) for two test data model realizations, showing permeability, TOUGH2 simulated temperatures, FNO model predicted temperatures, and the difference between them. Rows correspond to Low- and High-Fidelity FNOs. Note that gridlines on permeability graph show the mesh of the FNO training data, while predictions are all made on the high-fidelity mesh.

3.2 Production Phase Pressure Change

After the Natural State simulation, the three production and three injection wells are added to the TOUGH2 model and a 10 year Production Phase is simulated. The high-resolution simulations take 3 minutes on average. Pressure changes within the permeable reservoir are calculated and used to quantify the mismatch between TOUGH2 simulation outputs and FNO predictions (Figure 3). FNO surrogate predictions for the Production Phase each take ~0.2 ms, similar to those of the Natural State, again showing a speedup of nearly 900,000 times. When all model elements are compared, the FNO model predictions line up very well, but included in these measurements are many elements far away from the wells, where pressure differences are lower by default, artificially lowering the average mismatch. For production simulations, predicting pressure drop immediately near feed points is critical. Therefore, a mask is applied to calculate the prediction error in a small region near feed points (within 3 elements of the feed point element) and again for only the feed point elements themselves. Together, these show that both FNO surrogate models make poorer predictions nearer to feed points.

In contrast to the Natural State temperature performance, Production Phase pressure change performance is highly dependent on the training data type, with the Low-Fidelity FNO surrogate performing relatively poorly and High-Fidelity FNO surrogate performing relatively well. The Low-Fidelity FNO predicts 50% of feed point pressures within -40.9 and +38.9 bar, and the 90% whiskers extend to -73.6 and +84.3 bar. The High-Fidelity trained FNO has 50% of feed points between -3.4 and +4.1 bar and 90% of mismatches between -20.3 and +27.0 bar. The Low-Fidelity FNO production phase predictions tend to underpredict pressure changes, both positive and negative (Figure 3, right). This follows from the fact that injection/production into larger elements will result in smaller pressure changes, a difference commonly observed in models with coarser meshes.

Visual inspection of TOUGH2 model outputs and FNO model predictions shows the improved performance of the High-Fidelity FNO model (Figure 4). The Low-Fidelity FNO model has broader regions of elevated mismatch and the High-Fidelity FNO model shows spatially limited, minor mismatches in pressure change.

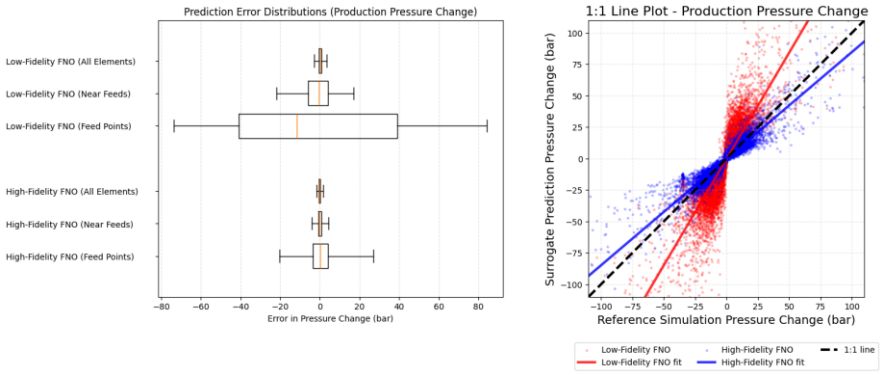


Figure 3. Boxplots of Production Phase pressure difference test data mismatch for each FNO surrogate model. Boxes span the 25th to 75th percentiles; whiskers span the 5th to 95th percentiles. Includes boxplots for all elements in the permeable reservoir, elements near the feed points (within 2 elements), and only the feed point elements themselves. 1:1 Line Plot shows elements at and near the feed points.

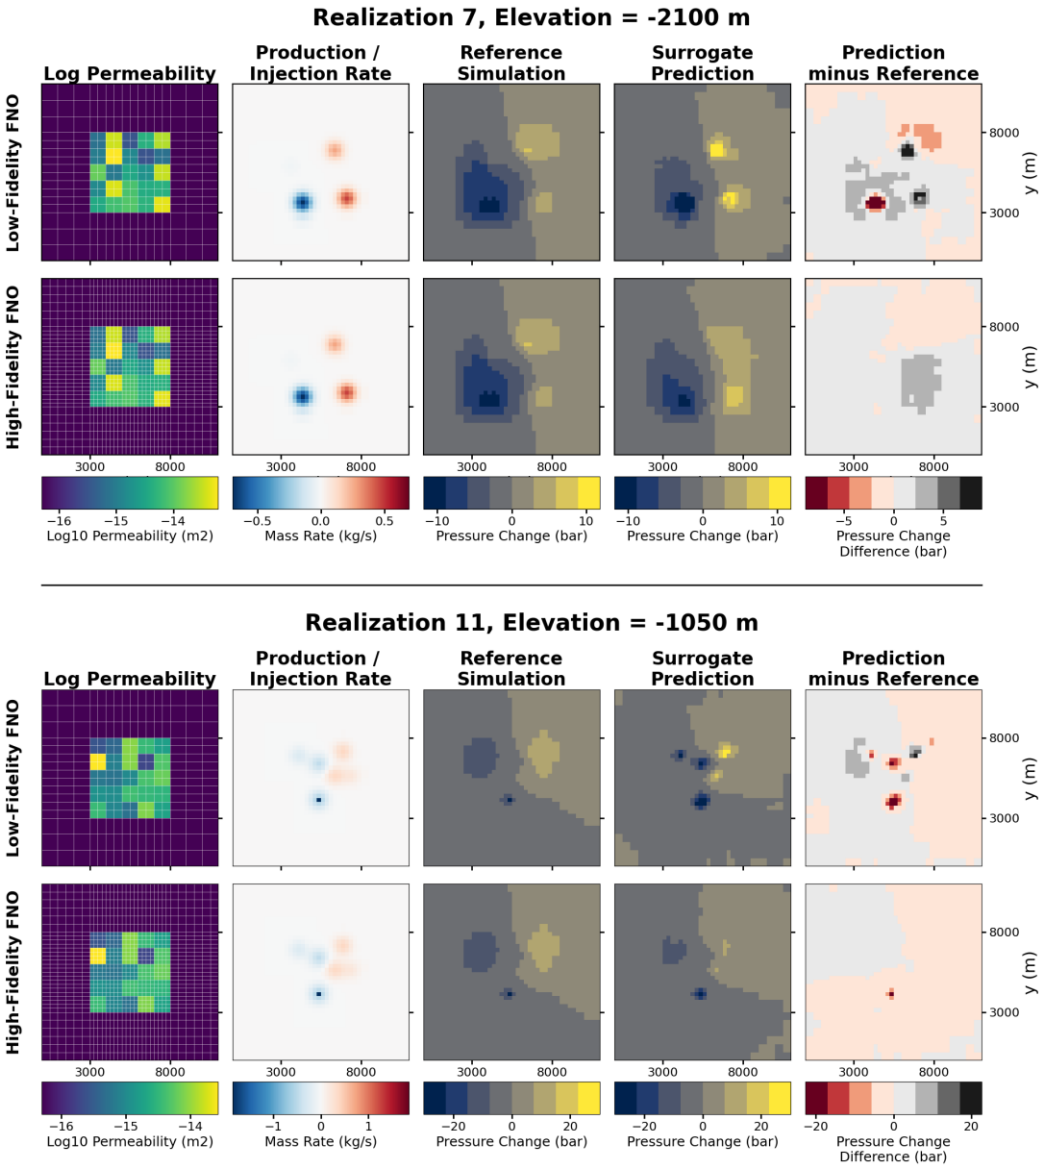


Figure 4. Horizontal slices through the reservoir (map-view) for two test data model realizations, showing permeability, Gaussian smoothed well production/injection rates, TOUGH2 simulated pressure changes, FNO model predicted pressure changes, and the difference between them. Rows correspond to Low- and High-Fidelity FNOs. Note that gridlines on permeability graph show the mesh of the FNO training data, while predictions are all made on the high-fidelity mesh.

4 DISCUSSION

4.1 Natural State Temperature

The Low- and High-Fidelity trained FNO surrogates performed similarly in predicting natural state temperatures, with only slight improvements in the High-Fidelity models. The overprediction bias in the LF FNO may be explained by differences in the simulation data: the average temperature of the coarse simulations is 79.6 °C, while the average temperature of the fine simulations is 71.7 °C. This is an common feature of geothermal simulations with different meshes. The marginal improvement in the range of the 90% error bars from 63.1 °C in the LF FNO to 54.7 °C in the HF FNO may imply a limitation of the FNO representation presented here rather than a dependence on model fidelity. FNOs excel in diffusive problems, where properties are global and smooth, but hot fluid upflow is a largely advective process which is local and directional. This is a fundamental limitation in the method, but there are other potential adjustments that could be made to improve these results:

- Objective function & weighting: augmenting the mean squared error to give extra weighting to the interior model where upflow occurs (rather than the permeable reservoir as a whole) could improve fitting in those areas where predictions are most important.

- Non-uniform grids: the TOUGH2 mesh has irregular vertical and semi-irregular horizontal spacing commonly used in geothermal modeling, but ill-suited for FNO models. Using regular gridding or accounting for the irregular grids, as in Lingsch et al. (2023), could improve results.

- FNO architecture: according to initial testing, increasing the Fourier modes does not appear to improve prediction performance, but explicitly including the hot upflow regions in a separate channel may improve predictions in those areas.

Despite its shortcomings, the FNO models presented capture the hot upflow's broad strokes of location (following vertical channels of high permeability) and magnitude, even when these advective forces were relatively complex. For example, model realizations with particularly high and well-connected permeabilities resulted in convective cells of downward flowing, cooler water - and these downflows were accurately predicted by both FNO models. In a reservoir modeler's workflow, this is a useful first-order approximation of a reservoir's behavior that can save significant computation and human time.

Additionally, though the FNO method's performance may be capped for these problems, the relatively accurate performance of the Low-Fidelity FNO temperature predictions may suggest that even coarser "Very-Low-Fidelity" models could be used to quickly generate initial guesses of High-Fidelity models, significantly reducing simulation time for a computationally expensive part of the reservoir simulation process.

4.2 Production Phase Pressure Change

The diffusive nature of pressure changes is fundamentally better suited to the FNO method, but sharp point-sources arising from production/injection wells are not. This is evident by the biased results from the Low-Fidelity FNO, which predicted consistently lower pressure changes. Because the Low-Fidelity FNO model was trained on coarsely meshed data with larger elements, it therefore predicts lower pressure changes. This is a fundamental difficulty in mixing resolutions. The Gaussian smoothing of the source/sink input channel significantly improved prediction quality near the wells for both models, but performance at the element containing the source/sink itself remains difficult to predict. Those feed point elements' pressures are particularly important for reservoir modelers as they become the reservoir pressure that wellbore models ultimately use to calculate wellhead fluid properties.

There are various techniques that could be applied to improve predictions here:

- Hybrid architectures: implement U-Nets to capture the near-well, high-frequency structure that Fourier models miss.

- Physics-informed loss: add diffusion PDE residual terms that penalize mismatch in flux balance around wells.

- Mixed-Fidelity training: Early results, not shown here, indicate that fine-tuning the Low-Fidelity FNO on High-Fidelity simulation data, similar to Tang et al (2023), can significantly improve the Production Phase pressure prediction with a small number of high-resolution simulations.

Although the Low-Fidelity FNO underpredicted pressure changes in response to injection or production, it did so relatively consistently. This suggests it may still be used for rapid prototyping for comparative well location optimization: the magnitude of the pressure change may be underestimated, but the relative change compared to other realizations may still be useful.

5 CONCLUSION

Fourier Neural Operators (FNOs) have been applied to create surrogate models for geothermal reservoirs. This study demonstrates the capability of the FNO method to estimate Natural State temperature fields and Production Phase pressure changes in a simple 3D reservoir model. The surrogate models effectively capture key features of the reservoir, including the location and magnitude of hot upflows and the pressure drawdown from multiple production and injection wells, but also contain significant mismatches to the testing data.

FNO surrogate models trained on Low-Fidelity and High-Fidelity data were compared. Results show that FNO model performance in predicting Natural State temperatures appears to be largely fidelity-independent, with low-fidelity trained FNO models performing slightly worse to those trained on high-fidelity data. This may suggest that even coarser training data may be used to make first-order predictions of finely gridded reservoir models. Production Phase pressure change predictions, on the other hand, were shown to be highly dependent on

training data fidelity, particularly at well feed points. Further investigation is needed to reduce the relatively large mismatches between FNO prediction and simulation data that arise at well feed points.

Future work will be focused on two key areas that will form the basis of a full peer-reviewed journal publication: first, a comparative benchmark against other surrogate models (e.g. UNet and CNNs); and second, the use of Mixed-Fidelity training strategy, where the Low-Fidelity FNO surrogate is fine-tuned on a portion of the High-Fidelity data, improving results without requiring large amounts of computationally heavy simulations. Further improvements, such as physics informed losses and others suggested above, may also improve surrogate models performance with even fewer simulations needed, offering a promising new direction for efficient and rapid geothermal reservoir evaluation.

Fast and accurate surrogate models unlock the possibility to conduct rapid uncertainty quantification and optimization, such as determining optimal well placement or injection strategies. Such work would be prohibitively slow with traditional simulators and would accelerate decision-making and lead to more efficient and reliable resource management.

REFERENCES

Almajid, M. M., & Abu-Al-Saud, M. O. (2022). Prediction of porous media fluid flow using physics-informed neural networks. Journal of Petroleum Science and Engineering, 208, 109205. https://doi.org/10.1016/j.petrol.2021.109205

Chandra, A., Koch, M., Pawar, S., Panda, A., Azizzadenesheli, K., Snippe, J., … & Hohl, D. (2025). Fourier neural operator–based surrogates for CO_2 storage in realistic geologies. arXiv preprint arXiv:2503.11031. https://doi.org/10.48550/arXiv.2503.11031

Grant, M. A., & Bixley, P. F. (2011). Geothermal Reservoir Engineering (2nd ed.). Academic Press.

Guo, Q., Koch, J., Yin, M., Fu, X., & Cardiff, M. (2024). Reduced geostatistical approach with a Fourier neural operator surrogate for inverse modeling of hydraulic tomography. Water Resources Research, 60, e2023WR034939. https://doi.org/10.1029/2023WR034939

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. Nature Reviews Physics, 3(6), 422–440. https://doi.org/10.1038/s42254-021-00314-5

Kovachki, N., Lanthaler, S., & Mishra, S. (2021). On universal approximation and error bounds for Fourier neural operators. Journal of Machine Learning Research, 22(290), 1–76. http://jmlr.org/papers/v22/20-1036.html

Kovachki, N. B., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2023). Neural operator: Learning maps between function spaces with applications to PDEs. Journal of Machine Learning Research, 24(171), 1–97. http://jmlr.org/papers/v24/21-1524.html

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Stuart, A., Bhattacharya, K., & Anandkumar, A. (2021). Fourier neural operator for parametric partial differential equations. In International Conference on Learning Representations (ICLR). https://openreview.net/forum?id=c8P9NQVtmnO

Lingsch, L., Michelis, M. Y., de Bézenac, E., Perera, S. M., Katzschmann, R. K., & Mishra, S. (2023). Beyond Regular Grids: Fourier-Based Neural Operators on Arbitrary Domains. arXiv. https://arxiv.org/abs/2305.19663v4 arxiv.org

O'Sullivan, M. J., Pruess, K., & Lippmann, M. J. (2001). State of the art of geothermal reservoir simulation. Geothermics, 30(4), 395–429. https://doi.org/10.1016/S0375-6505(01)00005-0

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., … & Courville, A. (2019). On the spectral bias of neural networks. In Proceedings of the 36th International Conference on Machine Learning (pp. 5301–5310). PMLR. https://proceedings.mlr.press/v97/rahaman19a.html

Tang, H., Kong, Q., & Morris, J. P. (2024). Multi-fidelity Fourier Neural Operator for fast modeling of large-scale geological carbon storage. Journal of Computational Physics. https://doi.org/10.1016/j.jhydrol.2024.130641

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., … & van Mulbregt, P. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Zhang, K., Zuo, Y., Zhao, H., Ma, X., Wang, Y., & Li, G. (2022). Fourier neural operator for solving subsurface oil/water two-phase flow PDEs. SPE Journal, 27(3), 1815–1830. https://doi.org/10.2118/209139-PA