When does array moveout help borehole phase picking? A leave-one-site-out, confound-free benchmark of array versus per-trace deep learning

Isao Kurosawa
e-Technology Innovation, ENEOS Xplora Inc., Tokyo 106-0041, JAPAN

Corresponding author: Isao Kurosawa

E-mail: [kurosawa.isao@eneos.com]

EarthArXiv DOI: *to be assigned*. Software and data archive: Zenodo, doi:10.5281/zenodo.21217615. Code: https://github.com/ISAO9/moirai-l3.

SUMMARY

Deep-learning phase pickers are increasingly deployed on borehole microseismic arrays, yet two questions central to operational monitoring remain poorly quantified: how well such models generalize to a site they were not trained on, and whether the cross-station *moveout* that an array provides actually improves that generalization. We give the first confound-free answer using the AMBER benchmark of eight three-component downhole arrays. A velocity-free, waveform-direct picker is trained under a strict leave-one-site-out (LOSO) protocol — seven sites for training and zero-shot evaluation on the eighth with no fine-tuning — and the contribution of moveout is isolated by an ablation in which the identical two-dimensional U-Net is trained either on the full array (the station axis is convolved, so moveout is available) or per trace (single-station inputs, moveout removed), so that the only difference is the presence of moveout. Across the eight held-out sites the two configurations are statistically indistinguishable in median F1 (0.885 versus 0.878; event-bootstrap 95 per cent confidence intervals overlap at the majority of sites): array moveout confers no *systematic* generalization benefit. The effect is instead strongly site-dependent and, at one site, catastrophic. Where a new array's median P moveout exceeds the largest value seen in training (forge_19, 70 ms versus a 62 ms training maximum) the array P detector collapses (F1 0.18) while the per-trace model recovers it (0.76); probability diagnostics show a detection failure, not a timing error, and a station-shuffle control attributes the collapse to non-transferable cross-station coupling. Because unattended monitoring is governed by worst-case rather than average performance, we

recommend the velocity-free per-trace picker as the safer default and treat array moveout as a conditional, in-distribution enhancement guarded by an explicit out-of-distribution check.

Key words: Neural networks, fuzzy logic; Induced seismicity; Earthquake monitoring; Computational seismology; Time-series analysis.

1 INTRODUCTION

Microseismic monitoring of fluid-injection operations — enhanced geothermal systems, hydraulic fracturing and geological CO_2 storage — is central to managing induced seismicity and imaging reservoir processes (Ellsworth 2013; Grigoli et al. 2017), and is increasingly performed on dense borehole and fibre-optic (DAS) arrays (Karrenbach et al. 2019; Lellouch et al. 2019; Lindsey & Martin 2021). It depends on detecting and timing P and S arrivals in continuous borehole records, often at sampling rates an order of magnitude higher than the regional networks on which most deep-learning (DL) pickers were trained. Established surface-trained models such as GPD (Ross et al. 2018), PhaseNet (Zhu & Beroza 2019) and EQTransformer (Mousavi et al. 2020), distributed through frameworks such as SeisBench (Woollam et al. 2022), have transformed catalogue completeness, but their behaviour on high-frequency downhole arrays is only beginning to be characterized. Systematic cross-domain benchmarking shows that such pickers transfer within a data domain yet degrade on recordings with different characteristics (Münchmeyer et al. 2022), and induced-seismicity applications often require in-domain retraining (Park et al. 2020; Chai et al. 2020). Lim et al. (2025) benchmarked four surface-trained DL pickers on the Preston New Road (PNR-1z) borehole dataset and found that, although the models detect induced seismicity, they do not transfer to the high-frequency borehole regime without fine-tuning. That study frames the problem we address here: rather than asking how well *surface-trained* models transfer, we ask how well a *borehole-trained* model generalizes to an *unseen borehole site*, and whether using the array — specifically the across-station moveout — helps or hurts that generalization.

Two design philosophies are in tension. An array-level picker can, in principle, exploit the coherent sweep of an arrival across depth-ordered stations (moveout) as an additional discriminative cue, much as array methods and analysts do (Rost & Thomas 2002). A per-trace picker ignores this cue and treats each station independently. The array cue is attractive, but it couples the model to the *geometry* of the training arrays; if a new site has a different aperture, station spacing or source–array configuration, a moveout-dependent detector may extrapolate poorly. Whether the net effect across realistic sites is positive, negative or neutral is an empirical question that, to our knowledge, has not been tested in a confound-free way.

We make three contributions. (i) We evaluate a velocity-free, waveform-direct picker under a strict zero-shot LOSO protocol on the eight-site AMBER benchmark, reporting event-bootstrap confidence intervals (Efron & Tibshirani 1994) so that per-site differences are interpretable rather than anecdotal. (ii) We isolate the contribution of moveout with a confound-free ablation — the *same* network, data and training schedule, differing only in whether the station axis is available — avoiding the sampling-rate, architecture and training-domain confounds that make off-the-shelf comparisons inconclusive. (iii) We quantify *when* moveout helps by relating the array advantage to a data-driven measure of each site's across-station moveout, and we diagnose the dominant failure mode as out-of-distribution geometry rather than intrinsic difficulty. The result is a cautionary, operationally actionable finding: across sites there is no systematic benefit to array moveout, the per-trace picker is more robust in the worst case, and we recommend it as the safer default for unattended monitoring.

This work is part of the MOIRAI line of velocity-free monitoring tools; hereafter we refer to the picker simply as the array or per-trace model.

2 DATA: THE AMBER BENCHMARK

We use the AMBER benchmark of three-component downhole microseismic arrays (Verdon et al. 2026), compiled from ten field sub-datasets. After excluding two sub-datasets (Cotton Valley Stage B and FORGE 2022) with incomplete or inconsistent station metadata in our processing, we retain eight clean sites spanning enhanced-geothermal, hydraulic-fracturing, carbon-storage and other induced-seismicity settings: pnr-1, pnr-2 (Preston New Road), mseel_3h, mseel_5h (MSEEL), clearfield_mw4, clearfield_mw6 (Clearfield), aneth (Aneth CCS) and forge_19 (Utah FORGE 2019). All records are processed to a common sampling rate (2000 Hz) and a fixed analysis window of 2048 samples across 12 depth-ordered stations.

Table 1 summarizes the sites. For each we report the number of catalogued events and a data-driven descriptor of array geometry — the *across-station moveout*, defined per event as the difference between the latest and earliest catalogued arrival sample over the stations carrying a pick, converted to milliseconds; we report its per-site median and inter-quartile range (IQR) for P and S. This descriptor is coordinate-free (it is read directly from the labels) and therefore robust to differences in metadata schema between sites. It is also the quantity an array-level moveout cue must learn, which makes it the natural axis along which to ask whether moveout generalizes.

Table 1. Site characteristics. Moveout is the median (IQR) across-station P/S arrival spread over events with at least two picked stations.

Site	Events	P moveout, ms (IQR)	S moveout, ms (IQR)
pnr-1	1258	62.0 (4.0)	105.0

Site	Events	P moveout, ms (IQR)	S moveout, ms (IQR)
mseel_3h	1684	31.5 (12.5)	51.0
mseel_5h	1512	26.8 (17.6)	41.5
clearfield_mw6	1612	25.0 (19.0)	41.5
pnr-2	972	35.0 (3.0)	55.5
clearfield_mw4	1254	25.3 (21.5)	42.3
aneth	298	6.0 (3.0)	18.0
forge_19	1213	70.0 (3.0)	129.0

The eight sites span more than a factor of ten in P moveout (6–70 ms). Two features anticipate the results below: aneth has the smallest moveout (6 ms; stations are nearly redundant), and forge_19 has the largest (70 ms) with an unusually tight IQR (3 ms; an almost constant, steep moveout), exceeding the largest training value available to it under LOSO (pnr-1, 62 ms).

3 METHODS

3.1 Input representation and labels

Each event is a tensor of shape (C, S, T) with $C = 3$ components (N, E, Z), up to $S = 12$ depth-ordered stations and $T = 2048$ samples at $f_s = 2000$ Hz (a 1.024 s window); amplitudes are normalized per station, and station (depth) order is preserved because it carries the moveout. Targets follow a tapered labeller: for a catalogued arrival of phase $\varphi \in \{P, S\}$ at sample t_s^{φ} on station s, the target is a symmetric linear taper of half-width $D = 30$ samples (± 15 ms),

$$y_s^{\varphi}(t) = max\left(0,\ 1 - \frac{|t - t_s^{\varphi}|}{D}\right), \qquad y_s^{N}(t) = clip(1 - y_s^{P} - y_s^{S},\ 0,\ 1), \tag{1}$$

Equation (1) gives the per-station targets, so the network is trained against three soft channels $[P, S, noise]$. Noise-only events (no catalogued picks) carry all-noise targets and provide true-negative supervision.

3.2 Network architecture

The picker (MoiraiPickerL3, ≈ 8 M parameters) is a two-dimensional U-Net (Ronneberger et al. 2015) acting on the (station,time) plane. It is built from GroupNorm + SiLU residual blocks (two 3 × 3 convolutions with a 1 × 1 skip), a stem with base width 32, depth 4 and channel multipliers (1,2,4,8) (widths 32/64/128/256), and a light self-attention layer at the bottleneck. **All down- and up-sampling act on the time axis only** (stem and each encoder stage stride the time axis by 2 via 3 × 3 convolutions; the decoder mirrors this with transposed convolutions and skip connections). The station axis is convolved but **never strided or pooled**, so cross-

station moveout is accessible throughout while the station count is preserved end to end. The input/output channels are 3 (N,E,Z) and 3 (P,S,noise); dropout is 0.1. Because the station axis is never reduced, the same weights process any number of stations $1 \le S \le 12$ unchanged — the property that makes the confound-free ablation (Section 3.5) exact.

3.3 Training

Pixels are dominated by the noise class, so we minimize a class-weighted focal loss (Lin et al. 2017). With logits $z_{c,s,t}$, predicted probabilities $\hat{p}_{c,s,t} = \sigma(z_{c,s,t})$ and binary cross-entropy $BCE(\hat{p}, y) = -[\, y\log\hat{p} + (1-y)\log(1-\hat{p})\,]$,

$$\mathcal{L} = \frac{1}{|\Omega|}\sum_{c,s,t} w_c \left(1-\hat{p}_{c,s,t}\right)^{\gamma} BCE\left(\hat{p}_{c,s,t},\, y_{c,s,t}\right), \qquad \gamma = 2, \quad (w_P, w_S, w_N) = (1,1,0.2). \tag{2}$$

Equation (2) is the per-pixel training objective. (BCE and softmax cross-entropy variants are retained for picker-comparable benchmarking.) Optimization uses AdamW (Adam: Kingma & Ba 2015) with learning rate $\eta = 3 \times 10^{-4}$, decoupled weight decay 10^{-4}, batch size 16 events, and up to 200 epochs with a 5-epoch linear warm-up followed by cosine decay. Training uses mixed precision and gradient clipping, with a guard that skips any non-finite batch so it cannot corrupt the weights. We keep an exponential moving average (EMA) of the parameters (Polyak & Juditsky 1992; Izmailov et al. 2018),

$$\bar{\theta}_t = \alpha\,\bar{\theta}_{t-1} + (1-\alpha)\,\theta_t, \qquad \alpha = 0.999, \tag{3}$$

as in Eq. (3), and report the EMA weights, which were consistently at least as good as the raw weights. Early stopping uses a patience of 25 epochs on the development F1-mean (Section 3.7). The base random seed is fixed (1234); the decisive folds are additionally repeated over three seeds for the error bars in Section 4.2.

3.4 Leave-one-site-out protocol

Generalization is measured zero-shot: for each held-out site the model is trained on the other seven sites and evaluated on the held-out site with no fine-tuning. Model selection (early stopping) uses a development split drawn from the training sites only.

3.5 Confound-free moveout ablation

To isolate the contribution of moveout we train the identical network in two configurations that differ only in whether the station axis carries coherent moveout:

- **array** — full (channel, 12,time) input; the station-spanning convolutions can use moveout.

- **per-trace** — single-station inputs ($S = 1$); moveout is unavailable, while every other factor (architecture, data, optimizer, schedule, augmentation) is identical.

Because the only difference is the presence of moveout, any performance gap is attributable to moveout rather than to sampling rate, architecture or training domain — the confounds that limit off-the-shelf baselines.

3.6 Station-shuffle control

A reviewer may object that collapsing the station axis also changes the receptive field and batch statistics, not only moveout. As a sharper control we additionally train an **array** model on inputs whose station order is independently permuted per sample, applied identically to waveforms and labels: this preserves array size and per-station waveform/label pairing while destroying the coherent across-station moveout. It thus separates the effect of *ordered* moveout from generic *cross-station coupling*. Results on all four decisive sites are in Section 5.3 (Table 3).

3.7 Scoring and uncertainty

A predicted pick for phase φ is the time of the probability peak, $\hat{t}^{\varphi} = arg\max_t \hat{p}^{\varphi}(t)$, declared only where $\max_t \hat{p}^{\varphi}(t) > 0.30$; duplicate peaks within 30 samples (15 ms) are suppressed, and per-trace picks are aggregated to the event. A pick is a true positive (TP) if $|\hat{t}^{\varphi} - t^{\varphi}| \leq \tau$, with primary tolerance $\tau = 20$ ms (40 samples; we also report ± 10 and ± 50 ms). Per phase we report, in Eq. (4),

$$\text{precision} = \frac{TP}{TP+FP}, \quad \text{recall} = \frac{TP}{TP+FN}, \quad F_1 = \frac{2\,\text{precision}\cdot\text{recall}}{\text{precision}+\text{recall}}, \tag{4}$$

and summarize each site by F_1-mean $= \frac{1}{2}(F_1^P + F_1^S)$. On noise-only events any pick is a false alarm, reported as false alarms per noise event. To quantify finite-test uncertainty we compute event-level bootstrap 95 per cent confidence intervals (CIs): events are resampled with replacement ($B = 300$) and the metric recomputed, the CI being the 2.5/97.5 percentiles (Efron & Tibshirani 1994). Training-seed variability is assessed separately by repeating the four decisive folds across three random seeds; we report the mean and standard deviation of the held-out F_1-mean (Fig. 3).

3.8 Moveout descriptor

The geometry descriptor of Section 2 is read directly from the labels: for phase φ and event e,

$$\Delta^{\varphi}_{\text{mo}}(e) = \frac{1}{f_s}\left(\max_{s\in\mathcal{S}(e)} t_s^{\varphi}(e) - \min_{s\in\mathcal{S}(e)} t_s^{\varphi}(e)\right), \tag{5}$$

where $\mathcal{S}(e)$ are the stations carrying a pick; we report its per-site median and IQR. Equation (5) is the quantity an array-level moveout cue must learn, and provides the

axis along which we test whether the array advantage is conditional. We relate the per-site array advantage, $\Delta = F_1$-mean(array) $- F_1$-mean(per-trace), to the site's median P moveout (Fig. 4).

4 RESULTS

4.1 Array-level generalization

Under LOSO the array model generalizes well to seven of eight unseen sites, with held-out F1-mean between 0.85 and 0.99, and S recall is robust everywhere (Fig. 1). The single exception is forge_19, where P collapses (P F1 0.18) while S remains strong (0.86); the collapse is flat across the ±10/20/50 ms tolerances, indicating a *detection* failure rather than a timing error (Section 5.1). One site, pnr-2, shows markedly elevated noise false alarms for the array model (≈ 1.8 picks per noise event, averaged over seeds, versus ≈ 0.15 for per-trace).

4.2 Array versus per-trace

Table 2 reports per-phase F1 for both configurations at the primary tolerance, with the per-site advantage Δ and a verdict from the bootstrap CIs. The result is mixed and site-dependent: the array is significantly better at two sites, the per-trace model is significantly better at two, and the remaining four are statistical ties (Fig. 2).

Table 2. Per-phase F1 (EMA, ±20 ms), array versus per-trace, with Δ = F1-mean(array) − F1-mean(per-trace) and the bootstrap-CI verdict. CIs in the text are 95 per cent on F1-mean.

Site	array P	array S	per-trace P	per-trace S	Δ	Verdict
pnr-1	0.989	0.978	0.991	0.985	−0.004	tie
mseel_5h	0.906	0.884	0.889	0.882	+0.010	tie
clearfield_mw6	0.873	0.872	0.871	0.869	+0.003	tie
mseel_3h	0.902	0.884	0.912	0.888	−0.007	tie
pnr-2	0.890	0.871	0.887	0.794	+0.040	**array**
clearfield_mw4	0.860	0.844	0.793	0.756	+0.078	**array**
aneth	0.943	0.835	0.965	0.916	−0.051	**per-trace**
forge_19	0.178	0.857	0.755	0.928	−0.324	**per-trace**

Event-bootstrap 95 per cent confidence intervals on F1-mean (300 resamples) confirm every verdict. Four sites are statistical ties (overlapping intervals): pnr-1 array [0.982, 0.993] versus per-trace [0.982, 0.994]; mseel_3h [0.873, 0.908] versus [0.876, 0.911]; mseel_5h [0.876, 0.915] versus [0.860, 0.898]; clearfield_mw6 [0.853, 0.887] versus [0.854, 0.888]. Two favour the array with disjoint intervals: pnr-2 [0.863, 0.910] versus [0.807, 0.855], driven by S (array S [0.849, 0.902] versus per-

trace [0.762, 0.818]); and clearfield_mw4 [0.832, 0.884] versus [0.740, 0.795], on both phases (P [0.838, 0.890] versus [0.758, 0.811]; S [0.821, 0.884] versus [0.720, 0.785]). Two favour per-trace with disjoint intervals: aneth [0.878, 0.900] versus [0.926, 0.947], driven by S (array S [0.813, 0.850] versus [0.899, 0.928]); and forge_19 [0.516, 0.556] versus [0.815, 0.854], a P-detection collapse (array P [0.168, 0.234] versus per-trace [0.709, 0.777]).

Notably, the two significant differences are both carried by the S phase and point in *opposite* directions — array S is better at pnr-2, per-trace S is better at aneth — which is the signature of a geometry-dependent, not a uniformly beneficial, cue.

Training-seed stability. Repeating the four decisive folds across three seeds confirms every verdict with non-overlapping mean ± standard deviation: forge_19 array 0.458 ± 0.107 versus per-trace 0.808 ± 0.050; aneth 0.885 ± 0.004 versus 0.935 ± 0.012; clearfield_mw4 0.858 ± 0.004 versus 0.752 ± 0.018; pnr-2 0.874 ± 0.001 versus 0.812 ± 0.012 (Fig. 3). The differences therefore exceed both finite-test (bootstrap) and training-seed variability. One feature is itself diagnostic: the array model's seed variance is one to two orders of magnitude larger at forge_19 (± 0.107, runs spanning 0.34–0.54) than at any other site (± 0.001–0.004) — the moveout-based P detector is not merely worse but *unstable* on out-of-distribution geometry, whereas the per-trace model stays high and stable.

4.3 Aggregate behaviour: median tie, divergent worst case

Aggregated over the eight sites the two configurations are indistinguishable in the median (array 0.885 versus per-trace 0.878), so moveout provides no systematic benefit. The difference appears in the *worst case*: the array's worst site is forge_19 at 0.518, whereas the per-trace model's worst site is clearfield_mw4 at 0.774. The site-mean is accordingly lower for the array (0.848 versus 0.880) because forge_19 drags it down. For unattended monitoring, where the worst case governs risk, this asymmetry is the operative result.

5 MECHANISM

5.1 A detection failure, not a timing error

Figure 5 diagnoses the forge_19 P collapse directly. For a representative forge_19 event the predicted P-probability map is weak and follows a shallower apparent moveout than the true, steep arrivals, whereas the S-probability map tracks the catalogued S picks normally. Pooling over all forge_19 events, the maximum P probability at true P arrivals is broadly distributed and largely *below* the 0.30 detection threshold, while for an in-distribution site (mseel_5h) it saturates near 1.0; the S distributions are above threshold for both. The collapse is therefore a loss of P *detection confidence*, consistent with a moveout-dependent P detector that fails to fire on out-of-distribution geometry.

Conversely, when both configurations do detect a pick, their timing is equivalent: matched-pick residuals are tightly centred at zero with near-identical distributions for array and per-trace (P within about ±2 ms, S within about ±3–4 ms; Fig. 6, ~35 000 matched picks per phase). The entire array-versus-per-trace story is thus one of detection robustness, not sub-sample precision.

5.2 A moveout "sweet spot" and out-of-distribution geometry

Relating the array advantage Δ to each site's median P moveout (Fig. 4) reveals a non-monotonic but coherent pattern. The array advantage is positive only at intermediate moveout that is well represented in training (clearfield_mw4 +0.078 and pnr-2 +0.040, at 25–35 ms) and is negative at both extremes:

- **Too little moveout (aneth, 6 ms).** Stations are nearly redundant; the array's cross-station machinery adds capacity without exploitable signal and overfits, hurting generalization (the loss is on S, 0.835 versus 0.916).
- **Edge-of/out-of-distribution moveout (forge_19, 70 ms).** The median P moveout exceeds the largest value the model sees under LOSO (pnr-1, 62 ms) and is unusually consistent (IQR 3 ms); the moveout-dependent P detector extrapolates and collapses (P 0.18), while the per-trace model, which ignores moveout, recovers P (0.76).

We are explicit that moveout magnitude alone is necessary but not sufficient: pnr-1 (62 ms) is a tie rather than a loss, so the failure at forge_19 reflects moveout beyond the training range *together with* its atypical regularity and site-specific waveform character. The actionable point stands: the array advantage does not extrapolate beyond the geometric range it was trained on. Consistent with an extrapolation failure, forge_19 is also where the array P score is least stable across seeds (Section 4.2).

5.3 Station-shuffle control: coupling versus ordered moveout

The shuffle control retains cross-station *coupling/aggregation* (the station-spanning convolutions still operate) but destroys *ordered* moveout (the station axis is permuted per sample). It therefore sits between array (coupling + ordered moveout) and per-trace (neither), and lets us attribute each site's behaviour to one ingredient or the other. Run on all four decisive folds (Table 3), it shows that neither ingredient is uniformly beneficial — each helps at some sites and hurts at others (Fig. 7).

Table 3. Three-way control on the four decisive sites (F1, EMA, ±20 ms).

Site	array P / S / mean	shuffle P / S / mean	per-trace P / S / mean
forge_19	0.178 / 0.857 / 0.518	0.327 / 0.813 / 0.570	0.755 / 0.928 / 0.842

Site	array P / S / mean	shuffle P / S / mean	per-trace P / S / mean
clearfield_mw4	0.860 / 0.844 / 0.852	0.878 / 0.853 / 0.866	0.793 / 0.756 / 0.768
pnr-2	0.890 / 0.871 / 0.881	0.847 / 0.757 / 0.827	0.887 / 0.794 / 0.841
aneth	0.943 / 0.835 / 0.889	0.922 / 0.909 / 0.930	0.965 / 0.916 / 0.941

- **forge_19 (out-of-distribution).** array 0.518 → shuffle 0.570 → per-trace 0.842. Destroying ordered moveout recovers P only slightly (0.18 → 0.33); removing coupling as well recovers most of it (→ 0.76). Cross-station coupling is the dominant cause of the out-of-distribution P-collapse and does not transfer; ordered moveout adds a secondary penalty.
- **clearfield_mw4 (in-distribution, array-favoured).** shuffle 0.866 ≈ array 0.852 ≫ per-trace 0.768. Scrambling moveout does not reduce performance, whereas removing coupling costs ≈ 0.10. Here the advantage is cross-station aggregation, and ordered moveout is irrelevant.
- **pnr-2 (in-distribution, array-favoured via S).** array 0.881 > per-trace 0.841 > shuffle 0.827, with array S 0.871 collapsing to 0.757 under shuffle — below per-trace. Here the advantage *is* coherent moveout: a coupled model fed incoherent moveout is actively misled, worse than independent traces.
- **aneth (6 ms aperture, per-trace-favoured via S).** array 0.889 → shuffle 0.930 → per-trace 0.941, with S 0.835 → 0.909 → 0.916. Scrambling the near-degenerate moveout recovers most of the array's deficit, so that deficit is misuse of an uninformative moveout on a too-small aperture.

The mechanism is therefore heterogeneous rather than single-cause: cross-station coupling drives both the only catastrophic failure (forge_19, where it fails to transfer) and the largest in-distribution gain (clearfield_mw4); coherent moveout is decisive at exactly one site (pnr-2) and is mildly harmful where the aperture is too small to supply it (aneth). No single array ingredient is safe across geometries. The per-trace model, which uses neither, avoids the worst case and ties the array in the median — which is precisely why we recommend it as the default. This also refines Fig. 4: the moveout axis is a proxy for how far the test geometry sits from training, not the direct causal variable.

6 DISCUSSION

Our findings refine, rather than contradict, the intuition that an array is informative. Cross-station moveout is genuinely useful where it is both present and in-distribution, but it is a double-edged cue: it provides no systematic benefit in the median, it is

uninformative when the array is too tight, and it is actively harmful when a new site's geometry lies beyond the training range.

These results connect microseismic phase picking to the broader problem of distribution shift in machine learning (Quiñonero-Candela et al. 2009). Cross-domain benchmarks have shown that surface-trained pickers transfer within a data domain but degrade on recordings with different characteristics (Münchmeyer et al. 2022), and Lim et al. (2025) showed specifically that surface-trained models do not transfer to the high-frequency borehole regime without fine-tuning. Our experiment isolates a complementary and more subtle boundary: even for a borehole-trained model evaluated on borehole data, the array's moveout cue is itself a training-domain feature that fails to transfer once the test geometry leaves the range spanned by the training sites. The failure is therefore not about waveform bandwidth or sampling rate — the usual transfer confounds — but about *geometry*, and it is invisible to any evaluation that does not hold out a site with out-of-range moveout. This is why the confound-free, leave-one-site-out design matters: an in-domain train/test split, or a comparison confounded by architecture and sampling rate, would have reported the array as uniformly competitive and hidden the one regime in which it is dangerous.

For operational monitoring the relevant criterion is worst-case reliability under unattended, zero-shot deployment to new sites. On that criterion the velocity-free per-trace picker is the safer default: it matches the array in the median, avoids the catastrophic out-of-distribution P collapse (0.52 → 0.77 worst case), and produces far fewer noise false alarms where the array over-triggers (pnr-2: ≈ 0.15 versus ≈ 1.8 picks per noise event). Where a site's geometry is known to resemble the training distribution, an array model can be used to recover the modest in-distribution gain, ideally gated by a moveout check against the training range.

In practice this suggests a simple deployment rule. Because the moveout descriptor of Eq. (5) is read directly from a preliminary catalogue and requires no velocity model, an operator can compare a new site's median P moveout with the training range before choosing a configuration: inside the range, the array model may be used for its in-distribution gain; at or beyond the range, the per-trace model is the safer default and the array output should be distrusted. The same descriptor provides an interpretable, physically meaningful trigger for retraining or fine-tuning, in contrast to opaque confidence scores that remain high precisely where the array detector fails (Section 5.2). More broadly, the velocity-free, per-trace design aligns with the goal of unattended, site-agnostic monitoring for CO2 storage and geothermal operations, where new arrays are commissioned faster than site-specific training data can be accumulated.

Limitations. The benchmark comprises eight sites; the non-monotonic Δ-versus-moveout relationship is suggestive rather than definitive, and a larger site population would sharpen it. The moveout descriptor is a one-dimensional proxy for geometry;

coordinate-based aperture and source–array distance, where available, would complement it. A fair comparison against a *fine-tuned* external picker (rather than the off-the-shelf baseline, which is confounded by sampling rate and training domain) is left to the fine-tuned numbers of Lim et al. (2025) or future work. Training-seed error bars for the decisive folds (Fig. 3) confirm every verdict, and the station-shuffle control (Section 5.3, all four decisive folds) shows the mechanism is heterogeneous — coupling drives the out-of-distribution collapse, coherent moveout helps at only one site.

7 CONCLUSIONS

Under a strict, confound-free, leave-one-site-out test on eight borehole arrays, cross-station moveout provides no systematic generalization benefit for velocity-free P/S phase picking: array and per-trace configurations tie in the median, and the array's only large effect is a catastrophic, geometry-driven P-detection collapse at the one site whose moveout exceeds the training range. The array advantage, where it exists, is confined to intermediate, in-distribution moveout. Because unattended monitoring is governed by the worst case, we recommend the velocity-free per-trace picker as the safer default, and we frame array moveout as a conditional, in-distribution enhancement to be used with an explicit out-of-distribution guard.

ACKNOWLEDGMENTS

The author thanks the AMBER team (Verdon, Lim, Leung, Lapins, Rodríguez-Pradilla, Read and Werner) for assembling and openly releasing the downhole benchmark used here, and the maintainers of SeisBench for the open tooling on which this work builds. This research was conducted independently at IVXA, Japan. The author declares no competing interests.

DATA AVAILABILITY

The AMBER benchmark data are publicly available under CC-BY-4.0 (Verdon et al. 2026; doi:10.5281/zenodo.18944111; code at https://github.com/kelleuseis/AMBER_Public); surface-trained baselines are distributed through SeisBench (Woollam et al. 2022). The training, evaluation and analysis code, the EMA-best model weights, and the scripts that regenerate every figure and table are released in a public archive with a permanent Zenodo DOI (doi:10.5281/zenodo.21217615) and mirrored on GitHub (https://github.com/ISAO9/moirai-l3). The figures in this paper are regenerable from the released result logs without retraining.

REFERENCES

- Allen, R.V., 1978. Automatic earthquake recognition and timing from single traces, *Bull. seism. Soc. Am.*, **68**(5), 1521–1532.

- Beroza, G.C., Segou, M. & Mousavi, S.M., 2021. Machine learning and earthquake forecasting—next steps, *Nat. Commun.*, **12**, 4761.
- Chai, C., Maceira, M., Santos-Villalobos, H.J., Venkatakrishnan, S.V., Schoenball, M., Zhu, W., Beroza, G.C. & Thurber, C., 2020. Using a deep neural network and transfer learning to bridge scales for seismic phase picking, *Geophys. Res. Lett.*, **47**(16), e2020GL088651.
- Clarke, H., Verdon, J.P., Kettlety, T., Baird, A.F. & Kendall, J.-M., 2019. Real-time imaging, forecasting, and management of human-induced seismicity at Preston New Road, Lancashire, England, *Seismol. Res. Lett.*, **90**(5), 1902–1915.
- Efron, B. & Tibshirani, R.J., 1994. *An Introduction to the Bootstrap*, Chapman & Hall/CRC, New York.
- Ellsworth, W.L., 2013. Injection-induced earthquakes, *Science*, **341**(6142), 1225942, doi:10.1126/science.1225942.
- Grigoli, F., Cesca, S., Priolo, E., Rinaldi, A.P., Clinton, J.F., Stabile, T.A., Dost, B., Garcia Fernandez, M., Wiemer, S. & Dahm, T., 2017. Current challenges in monitoring, discrimination, and management of induced seismicity related to underground industrial activities: a European perspective, *Rev. Geophys.*, **55**(2), 310–340, doi:10.1002/2016RG000542.
- Ioffe, S. & Szegedy, C., 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift, in *Proceedings of the 32nd International Conference on Machine Learning*, PMLR **37**, 448–456.
- Izmailov, P., Podoprikhin, D., Garipov, T., Vetrov, D. & Wilson, A.G., 2018. Averaging weights leads to wider optima and better generalization, in *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Karrenbach, M., Cole, S., Ridge, A., Boone, K., Kahn, D., Rich, J., Silver, K. & Langton, D., 2019. Fiber-optic distributed acoustic sensing of microseismicity, strain and temperature during hydraulic fracturing, *Geophysics*, **84**(1), D11–D23.
- Kettlety, T., Verdon, J.P., Werner, M.J. & Kendall, J.-M., 2020. Stress transfer from opening hydraulic fractures controls the distribution of induced seismicity, *J. geophys. Res.*, **125**(1), e2019JB018794.
- Kingma, D.P. & Ba, J., 2015. Adam: a method for stochastic optimization, in *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- LeCun, Y., Bengio, Y. & Hinton, G., 2015. Deep learning, *Nature*, **521**(7553), 436–444, doi:10.1038/nature14539.
- Lellouch, A., Yuan, S., Spica, Z., Biondi, B. & Ellsworth, W.L., 2019. Seismic velocity estimation using passive downhole distributed acoustic sensing records: examples from the San Andreas Fault Observatory at Depth, *J. geophys. Res.*, **124**(7), 6931–6948.

- Li, J., Zhu, W., Biondi, E. & Zhan, Z., 2023. Earthquake focal mechanisms with distributed acoustic sensing, *Nat. Commun.*, **14**, 4181, doi:10.1038/s41467-023-39639-3.
- Lim, C.S.Y., Lapins, S., Segou, M. & Werner, M.J., 2025. Deep learning phase pickers: how well can existing models detect hydraulic-fracturing induced microseismicity from a borehole array?, *Geophys. J. Int.*, **240**(1), 535–549, doi:10.1093/gji/ggae386.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P., 2017. Focal loss for dense object detection, in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 2980–2988, doi:10.1109/ICCV.2017.324.
- Lindsey, N.J., Martin, E.R., Dreger, D.S., Freifeld, B., Cole, S., James, S.R., Biondi, B.L. & Ajo-Franklin, J.B., 2017. Fiber-optic network observations of earthquake wavefields, *Geophys. Res. Lett.*, **44**(23), 11792–11799, doi:10.1002/2017GL075722.
- Lindsey, N.J. & Martin, E.R., 2021. Fiber-optic seismology, *Annu. Rev. Earth Planet. Sci.*, **49**, 309–336, doi:10.1146/annurev-earth-072420-065213.
- Mousavi, S.M., Ellsworth, W.L., Zhu, W., Chuang, L.Y. & Beroza, G.C., 2020. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking, *Nat. Commun.*, **11**, 3952, doi:10.1038/s41467-020-17591-w.
- Mousavi, S.M. & Beroza, G.C., 2022. Deep-learning seismology, *Science*, **377**(6607), eabm4470, doi:10.1126/science.abm4470.
- Münchmeyer, J., Woollam, J., Rietbrock, A., Tilmann, F., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J. & Soto, H., 2022. Which picker fits my data? A quantitative evaluation of deep learning based seismic pickers, *J. geophys. Res.*, **127**(1), e2021JB023499, doi:10.1029/2021JB023499.
- Park, Y., Mousavi, S.M., Zhu, W., Ellsworth, W.L. & Beroza, G.C., 2020. Machine-learning-based analysis of the Guy-Greenbrier, Arkansas, earthquakes: a tale of two sequences, *Geophys. Res. Lett.*, **47**(6), e2020GL087032.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J. *et al.*, 2019. PyTorch: an imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems* **32**, pp. 8024–8035.
- Perol, T., Gharbi, M. & Denolle, M., 2018. Convolutional neural network for earthquake detection and location, *Sci. Adv.*, **4**(2), e1700578, doi:10.1126/sciadv.1700578.
- Polyak, B.T. & Juditsky, A.B., 1992. Acceleration of stochastic approximation by averaging, *SIAM J. Control Optim.*, **30**(4), 838–855.
- Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A. & Lawrence, N.D. (eds), 2009. *Dataset Shift in Machine Learning*, MIT Press, Cambridge, MA.

- Ronneberger, O., Fischer, P. & Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation, in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, LNCS **9351**, pp. 234–241, doi:10.1007/978-3-319-24574-4_28.
- Ross, Z.E., Meier, M.-A., Hauksson, E. & Heaton, T.H., 2018. Generalized seismic phase detection with deep learning, *Bull. seism. Soc. Am.*, **108**(5A), 2894–2901, doi:10.1785/0120180080.
- Ross, Z.E., Yue, Y., Meier, M.-A., Hauksson, E. & Heaton, T.H., 2019. PhaseLink: a deep learning approach to seismic phase association, *J. geophys. Res.*, **124**(1), 856–869, doi:10.1029/2018JB016674.
- Rost, S. & Thomas, C., 2002. Array seismology: methods and applications, *Rev. Geophys.*, **40**(3), 1008, doi:10.1029/2000RG000100.
- Smith, J.D., Ross, Z.E., Azizzadenesheli, K. & Muir, J.B., 2022. HypoSVI: hypocentre inversion with Stein variational inference and physics-informed neural networks, *Geophys. J. Int.*, **228**(1), 698–710.
- Tarvainen, A. & Valpola, H., 2017. Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results, in *Advances in Neural Information Processing Systems* **30**, pp. 1195–1204.
- Verdon, J., Lim, C.S.Y., Leung, K., Lapins, S., Rodríguez-Pradilla, G., Read, E. & Werner, M.J., 2026. The AI-Ready Downhole Microseismic Benchmark Database (AMBER) (Version v1) [Data set], Zenodo, doi:10.5281/zenodo.18944111.
- Wamriew, D., Pevzner, R., Maltsev, E. & Pissarenko, D., 2021. Deep neural networks for detection and location of microseismic events and velocity model inversion from microseismic data acquired by distributed acoustic sensing array, *Sensors*, **21**(19), 6627, doi:10.3390/s21196627.
- Woollam, J., Münchmeyer, J., Tilmann, F., Rietbrock, A., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J. & Soto, H., 2022. SeisBench—a toolbox for machine learning in seismology, *Seismol. Res. Lett.*, **93**(3), 1695–1709, doi:10.1785/0220210324.
- Zhan, Z., 2020. Distributed acoustic sensing turns fiber-optic cables into sensitive seismic antennas, *Seismol. Res. Lett.*, **91**(1), 1–15, doi:10.1785/0220190112.
- Zhu, W. & Beroza, G.C., 2019. PhaseNet: a deep-neural-network-based seismic arrival-time picking method, *Geophys. J. Int.*, **216**(1), 261–273, doi:10.1093/gji/ggy423.
- Zhu, W., McBrearty, I.W., Mousavi, S.M., Ellsworth, W.L. & Beroza, G.C., 2022. Earthquake phase association using a Bayesian Gaussian mixture model, *J. geophys. Res.*, **127**(6), e2021JB023249, doi:10.1029/2021JB023249.

FIGURES

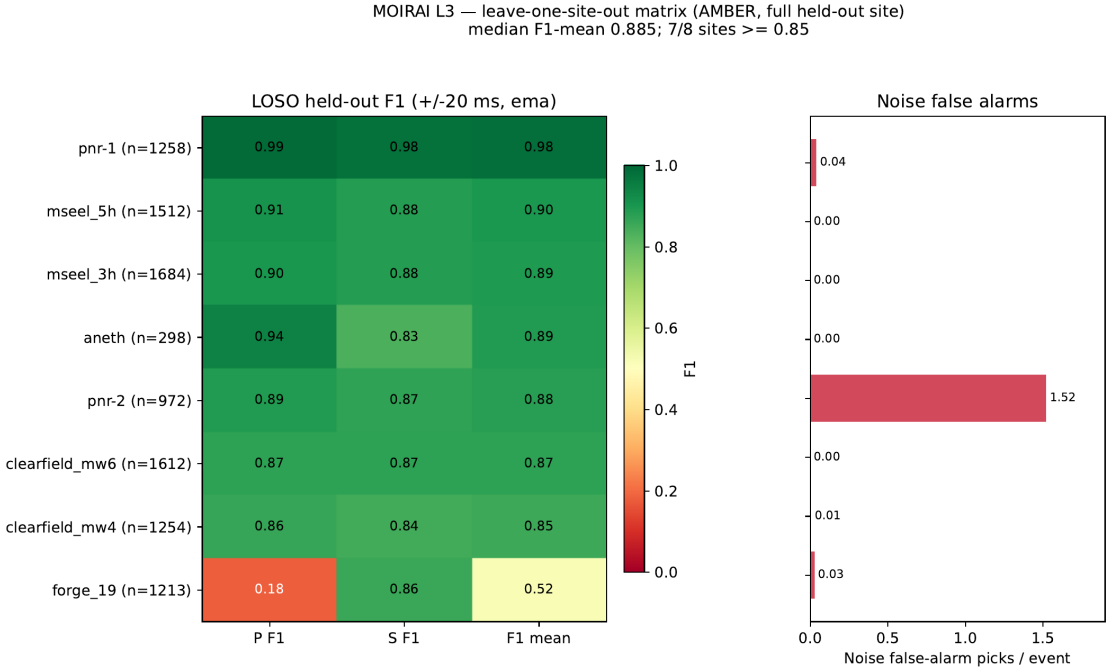

Figure 1. Leave-one-site-out held-out F1 (EMA weights, ±20 ms). Each row is a held-out site; columns are P, S and F1-mean. Seven of eight sites are strong; forge_19 shows a P-detection collapse. The right panel shows noise false alarms per noise event.

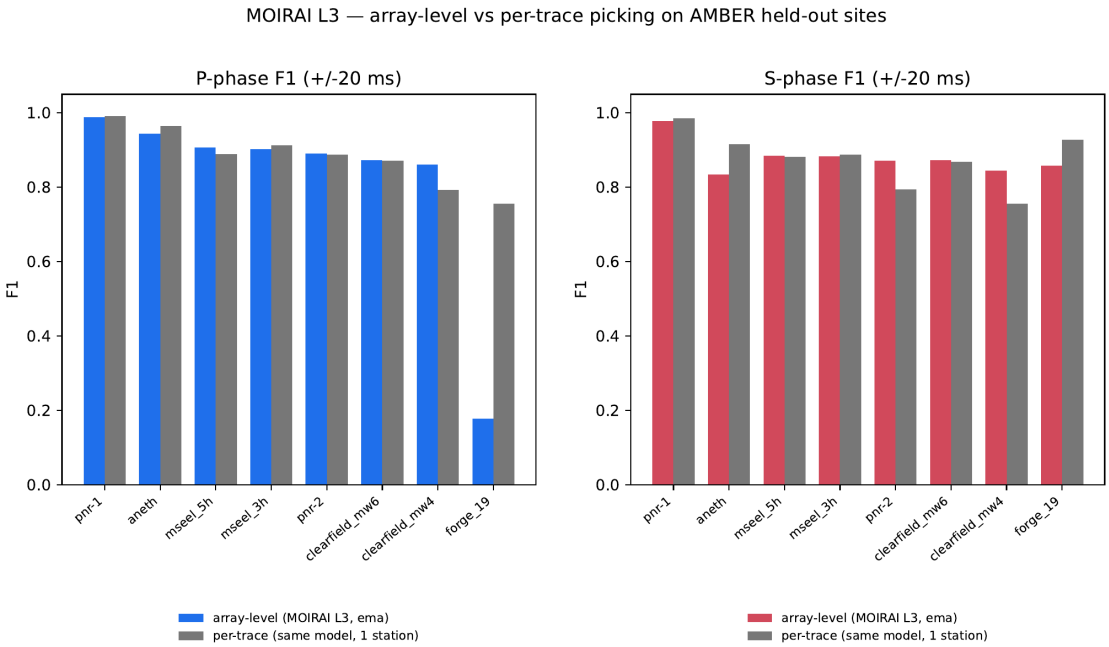

Figure 2. Array versus per-trace per-phase F1 across the eight held-out sites (EMA, ±20 ms).

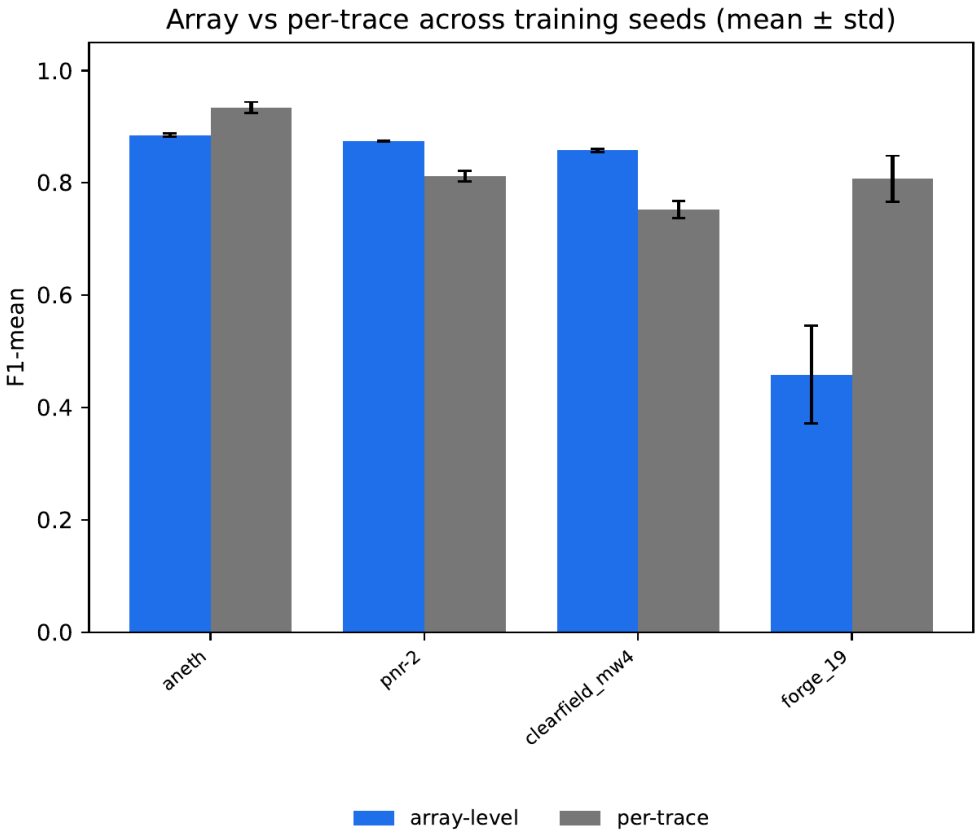

Figure 3. Array versus per-trace F1-mean on the four decisive held-out sites: mean ± standard deviation over three training seeds. All verdicts are non-overlapping; the array variance is largest at forge_19.

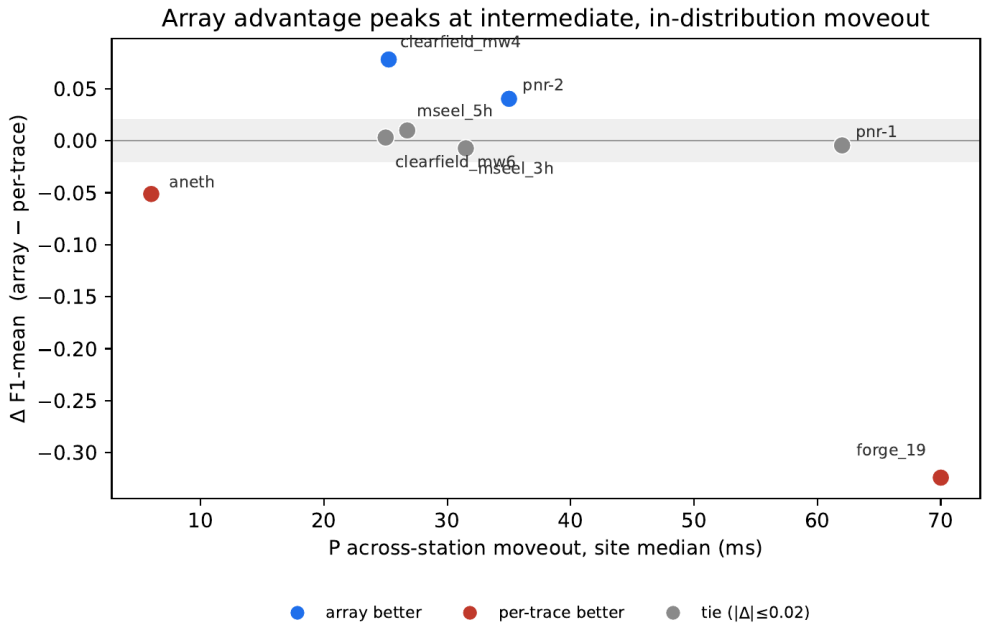

Figure 4. Per-site array advantage Δ = F1-mean(array) − F1-mean(per-trace) versus the site median P moveout. The advantage peaks at intermediate, in-distribution moveout and turns negative at both extremes.

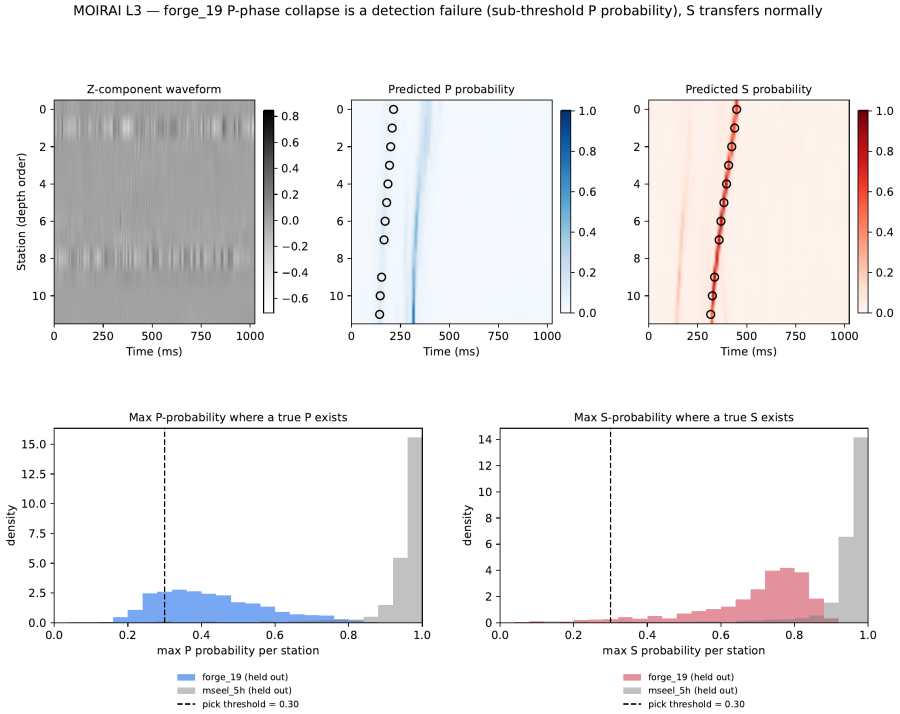

Figure 5. Diagnosis of the forge_19 P collapse. The predicted P-probability is weak and follows a shallower apparent moveout than the true arrivals, and the maximum P probability at true arrivals lies largely below the 0.30 threshold (a detection failure), whereas S transfers normally.

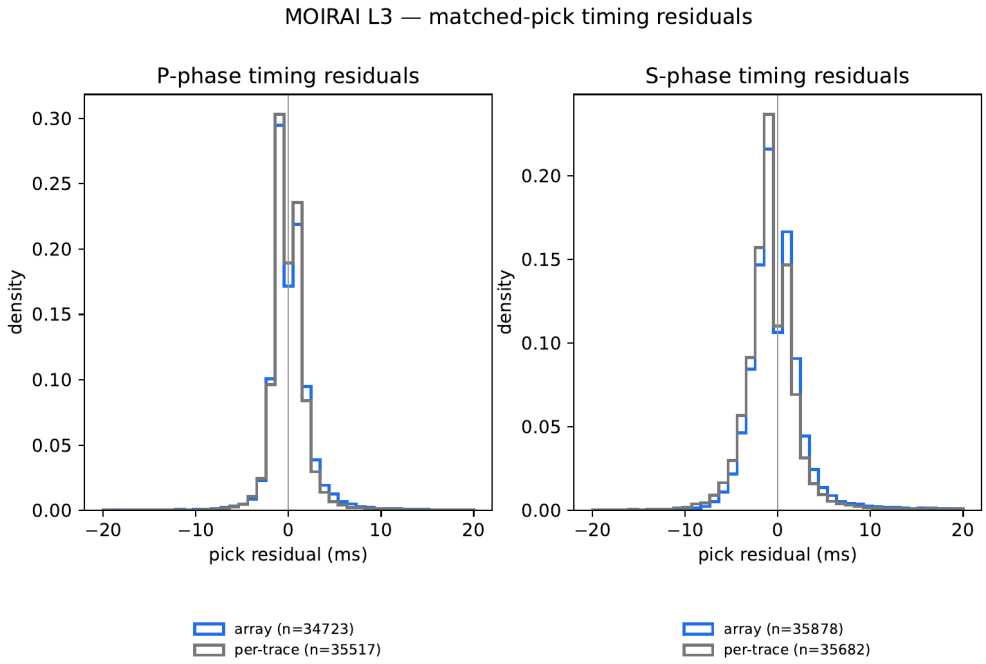

Figure 6. Matched-pick timing residuals for array and per-trace; the distributions are near-identical, so the array-versus-per-trace difference is one of detection, not timing.

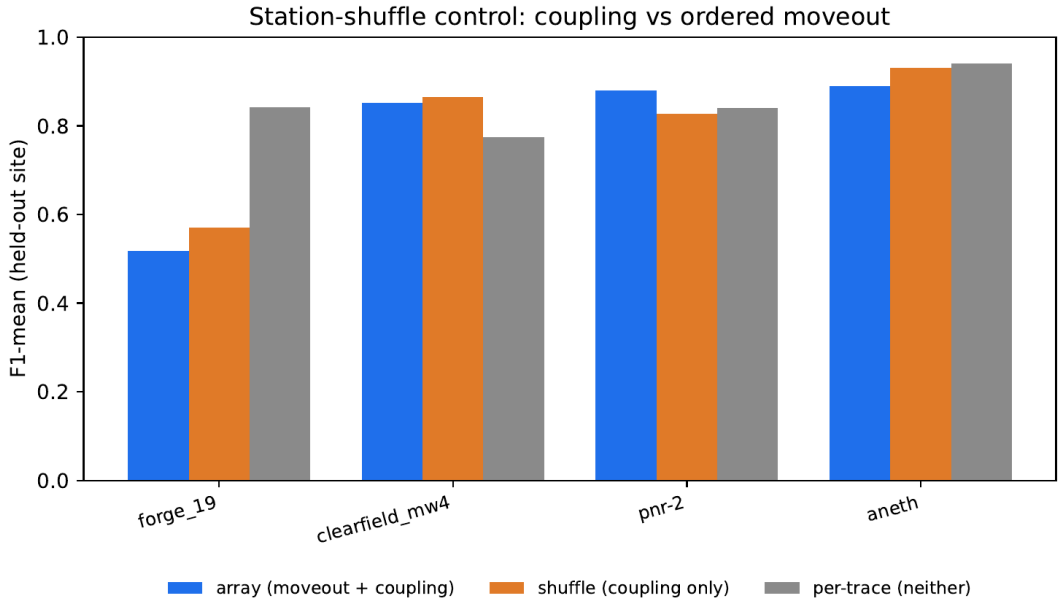

Figure 7. Station-shuffle control: array (moveout + coupling), shuffle (coupling only) and per-trace (neither) F1-mean on the four decisive sites; neither ordered moveout nor cross-station coupling helps uniformly.

SUPPORTING INFORMATION

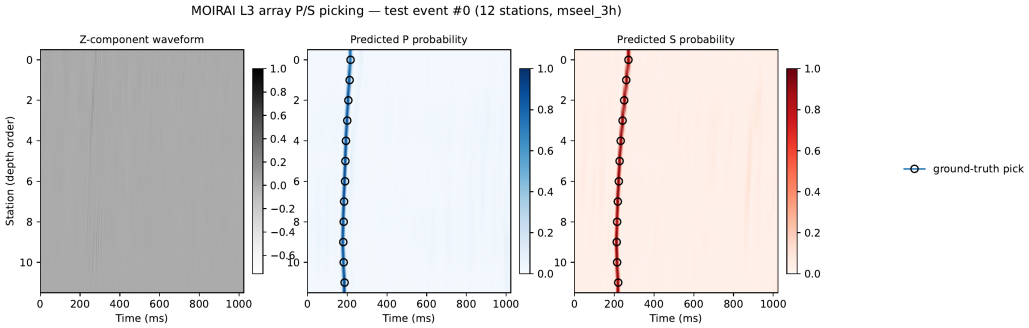

Figure S1. Example per-station P/S/noise probability panel for a representative event.

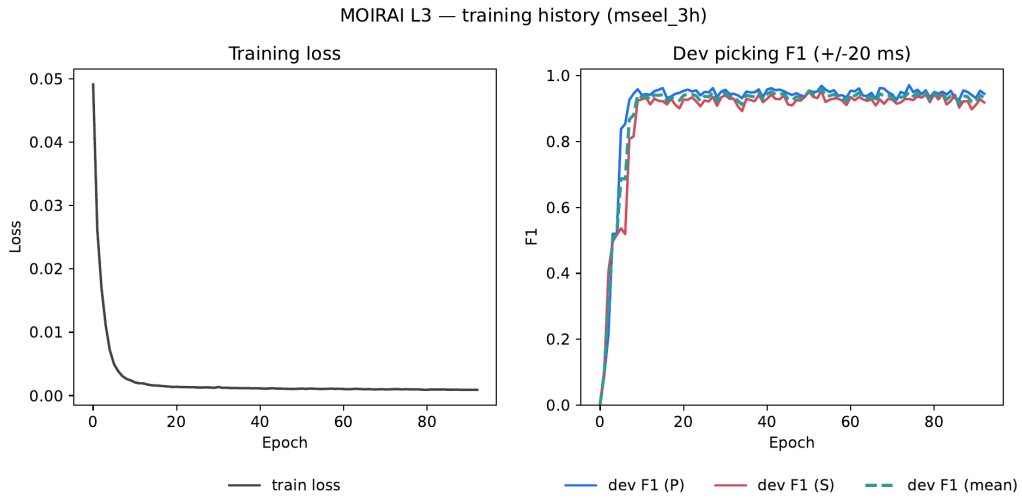

Figure S2. Training and development curves.

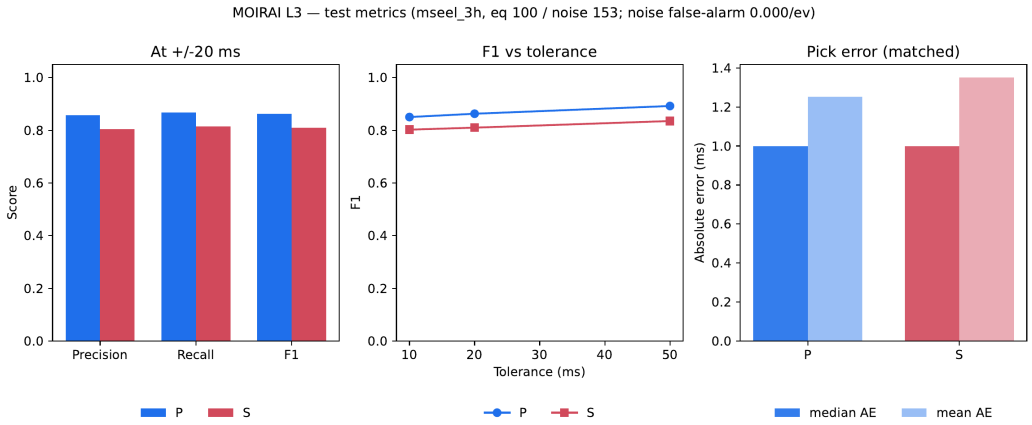

Figure S3. Summary of held-out metrics.

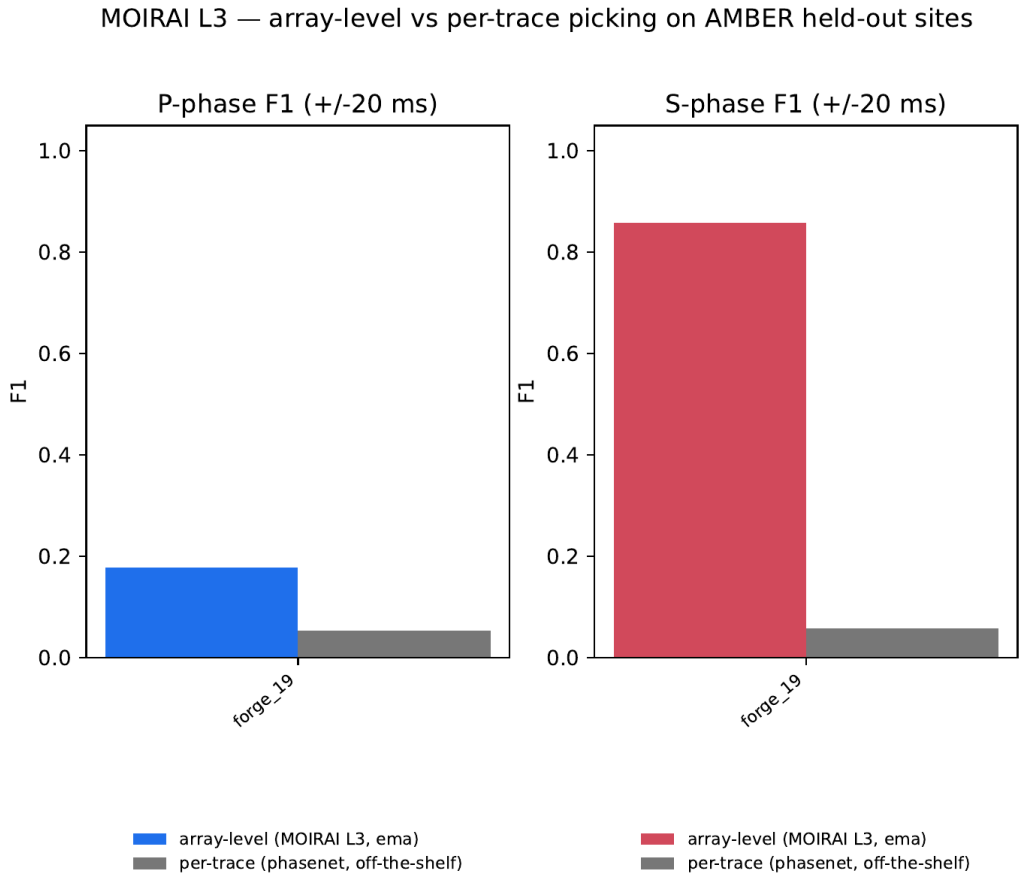

Figure S4. Off-the-shelf surface-trained baseline applied to a borehole site (motivation only; confounded by sampling rate, architecture and training domain).