Online Resource 1: Supplementary methods and results for TraceRIFT

Shuo Wang[1], Chen Guo[1*], Bingxin Yang[1], Keran Hao[1], Yu Zhang[1], Fuqiang Zhang[1]

[1]School of Information Engineering, Chang'an University, Xi'an, 710064, China.

*Corresponding author(s). E-mail(s): chenguo@chd.edu.cn;

Abstract

Supplementary methods and results for the main article, "TraceRIFT: Reference-preserving ensemble selection for discrete fracture networks from incomplete traces." Prepared for *Earth Science Informatics*.

S1 Optimization objective and guarantee

For a fixed core A and scenarios $y^{(m)}$, let $g_{mi} = \max\{0, \min_{j\in A} d_s(y^{(m)}, y_j) - d_s(y^{(m)}, y_i)\}$. Define $F(S) = M^{-1}\sum_m \max_{i\in S} g_{mi}$ and $F(\varnothing) = 0$. For $U \subseteq V$ and $j \notin V$, the scenario-wise increment is

$$\max\{0, g_{mj} - \max_{i\in U} g_{mi}\} \geq \max\{0, g_{mj} - \max_{i\in V} g_{mi}\}. \tag{S1}$$

Averaging preserves this inequality, proving submodularity. All utilities are nonnegative, so the function is normalized and monotone.

Let the optimal feasible tail be $T^\star = R^\star \cup E^\star$, where $R^\star \subseteq B_{\text{tail}}$ and $E^\star \subseteq \mathcal{C}_{64} \setminus B$, with $|E^\star| = k^\star \leq K$. The enumerated branches include $R^\star$. If $k^\star = 0$, enumeration includes the optimum directly. Otherwise $H(E) = F(R^\star \cup E) - F(R^\star)$ is normalized, monotone and submodular. Cardinality-$k^\star$ greedy gives

$$H(\widehat{E}) \geq \alpha_{k^\star} H(E^\star), \qquad \alpha_k = 1 - (1 - 1/k)^k \geq 1 - 1/e. \tag{S2}$$

Since $F(R^\star) \geq 0$,

$$F(R^\star \cup \widehat{E}) \geq \alpha_{k^\star} F(T^\star) + (1 - \alpha_{k^\star}) F(R^\star) \geq (1 - 1/e) F(T^\star). \tag{S3}$$

Selecting the best branch and applying strictly improving feasible exchanges preserves the bound. For $K \leq 2$, the same argument gives a factor of $3/4$ when the optimum adds two external candidates (Nemhauser et al. 1978).

This result applies to the optimization proposal for the sampled predictive objective. The reference tail is feasible, so the proposal also attains at least the reference objective value. The independent scenario check evaluates the final proposal on separate draws and can restore the reference tail. In that case, the approximation bound relative to the original sampled optimum does not carry over to the returned set.

S2 Implementation settings

Table S1 lists model, sampling and selection settings.

Table S1 Model, sampling and selection settings.

Component	Setting
Observation features	252: geometry 68; axial orientation 45; length/intensity 80; topology 59.
Proposal	Three-seed mixture; five spline transforms, eight bins, two 192-unit layers; batch size 128; learning rate 3×10^{-4}.
Simulator processes	Homogeneous Poisson and toroidal Thomas clustering; one to three axial families; truncated lognormal or gamma radii.
Risk predictor	256 extremely randomized trees; minimum leaf size 2; feature fraction 0.75.
Covariance regularization	Shrinkage 0.10; eigenvalue floor 10^{-4}.
Scenario counts	4,096 for optimization; 4,096 independent draws for the final check.
Candidate budget	80 attempts; exactly 64 retained or explicit construction failure.
Decision	Ordered top ten; core six; tail four; at most two external candidates.
Primary resampling	10,000 paired parent bootstrap replicates; complete parent records travel together.

The solver examines $\sum_{k=0}^{K} \binom{4}{k}$ retained-reference branches. With 54 external candidates, the full feasible-tail count is $\sum_{k=0}^{K} \binom{4}{k}\binom{54}{k}$. Table S2 summarizes the comparison with exhaustive search.

Table S2 Structural search counts and development objective agreement. Objective loss is $100(1 - F_{\text{solver}}/F_{\text{exhaustive}})$, not realized regret.

K	Branches	Feasible tails	Exact agreement	Mean objective loss (%)
0	1	1	–	–
1	5	217	30/30	0.0000
2	11	8,803	28/30	0.0489
3	15	108,019	27/30	0.0682
4	16	424,270	28/30	0.0161

Mean objective loss over the 120 checked query–budget pairs is 0.0333%, with a maximum of 1.4291%. The 35.2 s development timing covers set optimization for 768 queries across all five budgets.

S3 Evaluation metrics

S3.1 Primary gain and oracle retention

For parent p, let $g_{pF} = R_{pF}^{(R)} - R_{pF}^{(T)}$ and $g_{pM} = R_{pM}^{(R)} - R_{pM}^{(T)}$ denote gains under fragmentation and a missing surface. The primary gain is $G_p = (g_{pF} + g_{pM})/2$. The complete-observation gain is zero by construction, so the all-condition mean is $2G_p/3$. Figure 10d uses this all-condition endpoint, whereas the principal efficacy analysis uses the two-condition mean.

Oracle recall measures whether the selected set retains a minimum-loss candidate from the full pool. Best-member loss is evaluated retrospectively against the recorded parent target. Together with the pool-oracle loss, these quantities distinguish the quality of the available candidates from the loss incurred by retaining a subset.

For selectors using the same pool, the absolute reduction in mean best-member loss equals the reduction in mean retrieval regret. Their relative reductions use different denominators:

$$\delta_{\text{regret}} = \frac{\bar{R}_{\text{Reference}} - \bar{R}_{\text{TraceRIFT}}}{\bar{R}_{\text{Reference}}}, \qquad \delta_{\text{total}} = \frac{\bar{R}_{\text{Reference}} - \bar{R}_{\text{TraceRIFT}}}{\bar{L}_{64}^{\star} + \bar{R}_{\text{Reference}}}.$$

Using the displayed rounded mean support loss $\bar{L}_{64}^{\star} \simeq 0.2721$ and the reported mean regrets 0.090321 and 0.069836 yields total losses of approximately 0.3624 and 0.3419. Their relative reduction is about 5.7%, compared with 22.68% for retrieval regret.

S3.2 Target-wise minima

For target coordinate j, the paired diagnostic is

$$\Delta_{pj} = \frac{1}{2} \sum_{c \in \{F,M\}} \left[\min_{i \in B_{pc}} \frac{|y_{pcij} - y_{pj}^{\star}|}{s_{fj}} - \min_{i \in S_{pc}} \frac{|y_{pcij} - y_{pj}^{\star}|}{s_{fj}} \right]. \tag{S4}$$

The within-set minimizing candidate can differ across coordinates. We average these paired differences over parents, either pooling the two degraded conditions or reporting each separately. The intervals are pointwise estimates for each target.

S3.3 Gain tails and paired condition outcomes

For $g > 0$, the positive and negative exceedance curves are $n^{-1} \sum_p \mathbf{1}\{G_p > g\}$ and $n^{-1} \sum_p \mathbf{1}\{G_p < -g\}$. Both use all n parents as the denominator. The pooled tail summaries classify gains by strict sign. Condition-specific outcomes use the numerical tolerance $\tau = 10^{-15}$: a win has $g_{pc} > \tau$, a loss has $g_{pc} < -\tau$, and a tie lies between these thresholds. These are loss-based classes, so a tie can include different selected lists.

The pooled-parent distribution has approximately 29.3% positive and 18.3% negative gains. The condition-specific win/tie/loss proportions are 18.2/71.7/10.1% under fragmentation and 17.0/71.9/11.2% under a missing surface. Averaging gains before classification differs from averaging the resulting classes. The alluvial diagram cross-tabulates the two condition-specific classifications for the same parents.

Support and Reference-regret summaries divide parents into equal-count bins. They are descriptive analyses of evaluated outcomes. Reference regret is also a component of $G = R^{(R)} - R^{(T)}$, making the association partly dependent on the definition of gain. Pointwise intervals summarize the displayed bins.

S3.4 Nested-prefix support

The prefix oracle is $L^\star_m = \min_{i \le m} d_s(y_i, y^\star)$ for candidates in retained attempt order. Prefixes are nested subsets of the same generated pool, so $L^\star_m$ is nonincreasing for each query. Figure 11 reports the observed mean curve at $m = 8, 16, 24, 32, 48, 64$.

S4 Wadi Birk transfer analyses

The Wadi Birk catalogue supplies exposure-derived fracture geometry from which the virtual queries are constructed (Panara et al. 2024a,b). The primary analysis includes all 24 selected cells. Post-hoc sensitivity analyses use the 19 nonempty cells or the 17 cells with nonempty withheld surfaces. Table S3 reports these subsets separately.

Table S3 Reference-minus-TraceRIFT effects in primary and nonempty-subset analyses. Positive gains favor TraceRIFT.

Analysis	Mean gain	95% parent interval
Field targets, 24 cells	0.003606	[-0.005004,0.013503]
Held-out traces, 24 cells	0.011310	[-0.013581,0.038994]
Field targets, 19 nonempty cells	0.007995	[-0.002684,0.019706]
Held-out traces, 19 nonempty cells	0.003760	[-0.023113,0.030781]
Nonempty surfaces, 17 cells	0.001162	[-0.002446,0.005132]

Five catalogue cells use the empty-cell target convention. Of the 288 withheld-surface records, 185 contain no catalogue traces and 103 are nonempty. Coverage on empty surfaces follows the boundary-case convention in the evaluation implementation. On the 103 nonempty surfaces, mean trace-length coverage is approximately 3.01%, with zero coverage on 77 surfaces. This one-sided length measure is reported alongside discrepancy and unmatched-length measures to describe both reproduced and missed traces.

Figure 12 uses cell `WB20_I05_J11_K02`, frame 0, and the first missing-surface query in the display order fixed after cell selection. The Reference and TraceRIFT geometries are the best members identified retrospectively within their retained sets for this query. Their reprojections are compared with catalogue traces on the withheld virtual window. The inferential analyses use the full specified cohort of catalogue cells.

S5 Reproducibility materials

The companion release provides two entry points for workflow execution and statistical reproduction. The first runs an observation-to-pool-to-selection example; the second regenerates paper statistics from candidate targets, selection policies, parent identifiers and resampling definitions.

The implementation specifies coordinate conventions, observation masks, target transformations and scales, and candidate deduplication and identity. It also records deterministic tie breaking, the independent proposal check, failure behavior and the software environment. Fold assignments link each query to its training and evaluation roles. Source code is available at https://github.com/SureWangchd/TraceRIFT-CG; version v0.1.1 and its registered model and data assets are available at https://github.com/SureWangchd/TraceRIFT-CG/releases/tag/v0.1.1.

S6 Additional external-baseline comparisons

S6.1 Constrained selection and development lock

The additional comparisons use the archived 64-member pools and the unchanged Reference ranking. All methods return ten members, preserve the ordered six-member core and introduce at most two identities outside the original Reference top ten. Complete observations return that top ten unchanged. Candidate distances use the four-family metric and original training-fold scales from the main study.

The constrained maximal marginal relevance (MMR) baseline follows Carbonell and Goldstein (1998). It starts with the six-member core and fills four tail positions by maximizing

$$\lambda \,\mathrm{Rel}_i - (1-\lambda)\max_{j\in S}\frac{1}{1+d_s(y_i,y_j)}$$

over feasible unselected candidates. MMR-R sets $\mathrm{Rel}_i = 1 - r_i/63$, where r_i is the zero-based Reference rank. MMR-P sets $\mathrm{Rel}_i = 1/(1+\ell_i)$, with ℓ_i the candidate's mean distance to the same 4,096 optimization scenarios used by TraceRIFT. Thus MMR-P receives the same predictive means, covariances and sampled targets, while its selection objective remains relevance minus redundancy.

Constrained farthest-first follows the traversal principle of Gonzalez (1985). Starting from the same core, it selects the feasible candidate maximizing $\min_{j \in S} d_s(y_i, y_j)$ at each step. The core, external quota and final check define adaptations of these methods; the original clustering approximation guarantee is not invoked. Exact score ties use candidate content keys. Retained Reference tail members keep their original order, followed by introduced members in content-key order.

For each MMR variant, λ is selected from $\{0, 0.1, \ldots, 1\}$ on the 256 development parents. The criterion is mean regret after the common check, averaged over the two degraded conditions within each parent. Exact ties favor the larger weight. The selected weights are 0.4 for MMR-R and 0.7 for MMR-P. Both weights and the comparison protocol were frozen before evaluating the new baseline outputs on folds 1–4.

Each method submits one final proposal to the unchanged independent-check kernel, using the same 4,096 check scenarios and seed as TraceRIFT. A failed check returns the exact Reference list, without trying another proposal or weight. Seeds are $2718281 + 3j + c$ for optimization and $3141593 + 3j + c$ for checking, where j is the global archived parent row and $c \in \{0, 1, 2\}$ is the condition index. Selection runs do not load the query-truth file. Separate evaluation loads truth only after policies have been saved and hashed.

S6.2 Checked and unchecked results

Table S4 reports both final checked outputs and raw proposals. Raw variants retain the same development-selected weights and receive no separate tuning. All six variants satisfy the output, core and external-addition constraints.

Table S4 External-baseline sensitivity to the common predictive check. Mean and P95 are query-level retrieval regret over the two degraded conditions. Total loss is best-member attribute error.

Method	Output	Mean	P95	Total loss
MMR-R	Checked	0.077256	0.289792	0.349355
MMR-R	Raw	0.077710	0.287123	0.349809
MMR-P	Checked	0.078202	0.299304	0.350301
MMR-P	Raw	0.078277	0.300437	0.350376
Farthest-first	Checked	0.088989	0.335183	0.361088
Farthest-first	Raw	0.095151	0.332591	0.367250

The check rejects 831 of 6,428 nonidentity MMR-R proposals, 709 of 6,550 nonidentity MMR-P proposals and 5,120 of 6,552 farthest-first proposals. MMR-R and MMR-P also produce 124 and two identity proposals, respectively. Identity proposals have zero predictive gain and therefore do not pass the strict check, while their output remains unchanged. Without checking, farthest-first mean regret is 0.095151, compared with 0.088989 after checking and 0.090321 for Reference.

The paired intervals in Table S5 use the original 10,000 parent bootstrap draws and seed 1618033. Positive gains favor TraceRIFT. Intervals are pointwise; they are not simultaneous intervals for the full comparison family.

Table S5 Comparator-minus-TraceRIFT degraded mean regret and 95% parent-clustered intervals.

Comparator	Gain	Lower	Upper
Reference	0.020485	0.018320	0.022747
Fixed-donor	0.013603	0.011488	0.015717
Target-guided	0.013748	0.011610	0.015945
MMR-R	0.007420	0.005541	0.009351
MMR-P	0.008366	0.006602	0.010181
Farthest-first	0.019153	0.017066	0.021403
Sequential	-0.000059	-0.000512	0.000402

S6.3 Matched timing and reproducibility

Timing uses the first 20 archived evaluation parents from each held-out fold, with both degraded conditions, giving 160 queries. Each method runs three times per query after warm-up, with execution order rotated. Table S6 reports selection time, including method-specific pairwise distances, relevance scores and final checks. Shared construction of optimization and check scenarios is excluded, as are file access, predictor inference, candidate generation and reprojection. The original Sequential implementation constructs its full exchange path before checking the $K = 2$ endpoint.

Table S6 Matched selection timing on 160 degraded queries, with 480 calls per method. Times are in milliseconds.

Method	Mean	Median	P95
Reference	0.017	0.016	0.032
MMR-R	1.041	0.916	1.612
MMR-P	1.307	1.173	2.024
Farthest-first	1.159	1.072	1.778
Sequential	19.509	18.036	30.393
TraceRIFT	24.712	22.941	36.933

The benchmark ran on Windows with an AMD Ryzen 9 9950X processor, Python 3.11.4 and NumPy 2.4.6. Shared construction of both scenario buffers averaged 47.31 ms per query in this run. The earlier 9.16 ms development-sweep figure in Section S2 averages 768 queries across five budgets and describes a different execution scope.

S6.4 Equivalent implementation and timing

The optional implementation caches candidate columns in column-contiguous buffers and reuses a working maximum buffer. It constructs content tie keys only for candidates that can attain the current best score. The optimization objective, retained-reference branches, greedy choices, strict-improving exchanges and independent check remain unchanged. Arithmetic reductions retain the original column order.

Development verification compared complete nested optimization and check records on all 512 degraded development queries, plus 80 artificial cases spanning $K = 0, \ldots, 4$, exact ties and near-ties. Implementation hashes were frozen before replaying the held-out folds. The truth-free replay reproduced complete records on all 6,552 degraded queries, including branch objectives, exchange decisions, check values and final identities. Outputs and stored objective/check values also matched the original archive exactly. All 3,276 complete-condition outputs remained identical.

Table S7 gives a separate paired implementation benchmark on the same 160-query sample. Each path was warmed up per query; original and optimized calls then alternated across three repeats. The benchmark ran serially after replay workers exited, with numerical-library threads set to one. Times include public-interface validation, complete result records and the unchanged independent check, excluding scenario construction, equivalence assertions and file access.

Table S7 Paired implementation timing, with 480 calls per implementation. Times are in milliseconds.

Implementation	Mean	Median	P95
Original	19.063	17.816	27.643
Buffer reuse	5.534	5.258	7.572

The contemporaneous means give a 3.44-fold selection-stage speedup. This implementation benchmark uses the same hardware and Python/NumPy versions as the multi-method run, with its own warm-up and execution order. Exact equivalence preserves the reported scientific results and finite-scenario proposal guarantee.

Online Resource 2 supplies the frozen protocols, development grid, checked and raw policies, parent-level metrics, both timing runs, equivalent-selector code, replay records and independent audits. The source arrays are the four registered statistical assets from release v0.1.1. Their recorded SHA-256 values are checked before use. Verification covers content-identity preservation under storage permutation, the MMR-R $\lambda = 1$ Reference limit, the $\lambda = 0$ farthest-first limit, output feasibility and independent recomputation of the distances and paired intervals. The supplied scripts and audit records document the executed analyses.

S7 Sensitivity to the Reference ranking

S7.1 Alternative rankings and fixed selection protocol

This retrospective analysis changes only the Reference ranking on the existing cached pools and evaluation cohort. It compares three rankings, each with Reference, TraceRIFT and Sequential outputs. The original Borda ranking is a reproduction control, with equal fusion of raw TARGET and learned PyCoRE rank positions. The PyCoRE score is the negative root-mean-square of four learned nonnegative concept energies. The target-consistency-only ranking directly reuses the archived raw TARGET ranking by negative regularized Mahalanobis distance to its fold-honest predicted target. That predictor is distinct from the risk predictor used for scenario coverage (implementation identifier A13); the latter's original means and covariances remain unchanged in all comparisons.

The observable-discrepancy-only ranking reproduces the existing public-evidence score

$$s_i^{(O)} = -\left(\frac{1}{4}\sum_{k=1}^{4} u_{ik}^2\right)^{1/2},$$

where the archived u_{ik} are the geometry, orientation, topology and quality/missingness discrepancy components. The cached single-precision values are promoted to double precision before the calculation. The four components retain their existing definitions and enter equally, without additional fitting or normalization. This public score differs from learned PyCoRE concept energy, the implementation's arithmetic-mean total discrepancy, and the four-family target-attribute distance in main-text Eq. (2). Exact score ties use ascending candidate content hashes.

Each ranking supplies its own ordered top ten and first-six core. All methods return ten distinct candidates from the same 64-member pool. TraceRIFT and Sequential preserve that core and allow at most two identities outside that ranking's top ten; complete observations return its original ordered top ten exactly. The original training-fold scales, 4,096 optimization scenarios, 4,096 independent check scenarios and query-specific seeds from Section S6 are reused. No model is retrained, and no budget or selection parameter is retuned.

The protocol and implementation hashes were frozen before truth-free selection. On all 256 development parents, the equivalent TraceRIFT implementation reproduced the original implementation's complete optimization and check records under all three rankings. On the 3,276 held-out parents, the Borda control reproduced all three archived selectors' ordered outputs exactly, including TraceRIFT objective values, check statistics, rollback flags, replacement counts and seeds. Both development and held-out policy files were saved and hashed before separate evaluation loaded query truth.

S7.2 Absolute results and paired effects

Tables S8 and S9 report all nine combinations and all prespecified paired contrasts. Every combination shares the same pool-oracle error floor, whose degraded mean is 0.272099. Primary effects average the two degraded observations within each parent.

Intervals reuse the original 10,000 parent bootstrap draws and seed 1618033; they are pointwise two-sided 95% percentile intervals. Gains are calculated within each Reference ranking.

Table S8 Absolute degraded results for each Reference ranking and selector. Mean and P95 are retrieval regret; total is best-member attribute loss. Exact and near recall are fractions of the 6,552 degraded queries; near recall uses the absolute tolerance 0.05.

Ranking	Selector	Mean	P95	Total	Exact	Near
Borda	Reference	0.090321	0.347001	0.362420	0.378816	0.538462
Borda	TraceRIFT	0.069836	0.265895	0.341935	0.398199	0.576007
Borda	Sequential	0.069777	0.266553	0.341876	0.397741	0.576007
Target only	Reference	0.106272	0.395621	0.378371	0.384005	0.509005
Target only	TraceRIFT	0.075007	0.276816	0.347106	0.399878	0.559371
Target only	Sequential	0.075254	0.279912	0.347353	0.402930	0.560134
Observable RMS	Reference	0.104000	0.325971	0.376099	0.249847	0.423382
Observable RMS	TraceRIFT	0.079746	0.255355	0.351845	0.298535	0.488400
Observable RMS	Sequential	0.079967	0.255436	0.352066	0.297619	0.487332

Table S9 Within-ranking paired degraded regret contrasts. Positive values favor the second selector named in the contrast. Bounds are pointwise 95% parent-bootstrap intervals.

Ranking	Contrast	Gain	Lower	Upper
Borda	Reference − TraceRIFT	0.020485	0.018320	0.022747
Borda	Reference − Sequential	0.020544	0.018373	0.022844
Borda	Sequential − TraceRIFT	-0.000059	-0.000512	0.000402
Target only	Reference − TraceRIFT	0.031265	0.028450	0.034140
Target only	Reference − Sequential	0.031018	0.028233	0.033841
Target only	Sequential − TraceRIFT	0.000246	-0.000340	0.000825
Observable RMS	Reference − TraceRIFT	0.024254	0.022108	0.026427
Observable RMS	Reference − Sequential	0.024032	0.021914	0.026179
Observable RMS	Sequential − TraceRIFT	0.000221	-0.000329	0.000769

Both alternative rankings change the ordered top ten on every held-out query. Over the degraded conditions, their mean top-ten identity overlaps with Borda are 5.93/10 for target consistency and 2.98/10 for observable discrepancy. Thus the analysis changes the protected hypotheses as well as the reference tail. Both coverage selectors reduce mean regret relative to each ranking's own Reference, while all three Sequential-minus-TraceRIFT intervals include zero. Absolute TraceRIFT regret is lowest with Borda among these three tested rankings.

The independent check rejects four of 6,551 nonidentity proposals for each coverage selector under Borda, two of 6,552 under target-only ranking, and none of 6,552

under observable-discrepancy ranking. Each Borda coverage selector also has one identity proposal, which fails the strict positive-gain check without changing its output. Reference itself and complete observations do not undergo this check. All outputs satisfy their ranking-specific complete-list, ordered-core, output-size and two-addition constraints.

Online Resource 3 supplies the frozen protocol, code, development and held-out policies, all absolute and paired results, fold- and condition-specific summaries, ranking overlaps, check records and independent audits. It also supplies the source data and scripts for the updated main-text Figures 6 and 7, whose eight-method comparisons concern the original Borda Reference.

References

Carbonell J, Goldstein J (1998) The use of MMR, diversity-based reranking for reordering documents and producing summaries. In: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 335–336, https://doi.org/10.1145/290941.291025

Gonzalez TF (1985) Clustering to minimize the maximum intercluster distance. Theoretical Computer Science 38:293–306. https://doi.org/10.1016/0304-3975(85)90224-5

Nemhauser GL, Wolsey LA, Fisher ML (1978) An analysis of approximations for maximizing submodular set functions—i. Mathematical Programming 14:265–294. https://doi.org/10.1007/BF01588971

Panara Y, Menegoni N, Finkbeiner T, et al (2024a) High-resolution analysis of 3D fracture networks from digital outcrop models, correlation to plate-tectonic events and calibration of subsurface models (Jurassic, Arabian Plate). Marine and Petroleum Geology 167:106998. https://doi.org/10.1016/j.marpetgeo.2024.106998

Panara Y, Menegoni N, Finkbeiner T, et al (2024b) High-resolution digital outcrop models of late Jurassic outcrop, Wadi Birk, Saudi Arabia. Mendeley Data, Version 1, https://doi.org/10.17632/9bpcypmh2j.1