Full Length Article

Bayesian partial pooling to reduce uncertainty in overcoring rock stress estimation

Yu Feng [a,*], Ke Gao [b,**], Suzanne Lacasse [a]

[a] *Norwegian Geotechnical Institute, Oslo, Norway*
[b] *Department of Earth and Space Sciences, Southern University of Science and Technology, Shenzhen, China*

ABSTRACT

The state of in situ stress is a crucial parameter in subsurface engineering, especially for critical projects like nuclear waste repository. As one of the two ISRM suggested methods, the overcoring (OC) method is widely used to estimate the full stress tensors in rocks by independent regression analysis of the data from each OC test. However, such customary independent analysis of individual OC tests, known as no pooling, is liable to yield unreliable test-specific stress estimates due to various uncertainty sources involved in the OC method. To address this problem, a practical and no-cost solution is considered by incorporating into OC data analysis additional information implied within adjacent OC tests, which are usually available in OC measurement campaigns. Hence, this paper presents a Bayesian partial pooling (hierarchical) model for combined analysis of adjacent OC tests. We performed five case studies using OC test data made at a nuclear waste repository research site of Sweden. The results demonstrate that partial pooling of adjacent OC tests indeed allows borrowing of information across adjacent tests, and yields improved stress tensor estimates with reduced uncertainties simultaneously for all individual tests than they are independently analysed as no pooling, particularly for those unreliable no pooling stress estimates. A further model comparison shows that the partial pooling model also gives better predictive performance, and thus confirms that the information borrowed across adjacent OC tests is relevant and effective.

© 2024 Institute of Rock and Soil Mechanics, Chinese Academy of Sciences. Production and hosting by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

1. Introduction

The state of in situ stress is a crucial parameter in many subsurface engineering endeavours. Such examples include rock engineering design (Hudson and Harrison, 1997), geological storage of carbon dioxide and radioactive waste (Zhang et al., 2017; Bao and Burghardt, 2022), hydraulic stimulation design (Liu et al., 2018), borehole stability analysis (Maleki et al., 2014) and seismic hazard assessment (Stein, 1999). The overcoring (OC) method is widely used in stress measurement campaigns in civil and mining engineering (e.g. Fouial et al., 1998; Clément et al., 2009; Li et al., 2019), and is, along with hydraulic fracturing, among the two ISRM suggested methods for rock stress estimation (Haimson and Cornet,

2003; Sjöberg et al., 2003). The main advantage of the OC method is that the full stress tensor comprising six distinct components can be determined from a single OC test with a dedicated measuring device, e.g. the commonly used CSIRO HI and CSIR-type strain cells.

An OC test mainly involves relieving a rock sample from its surrounding rock masses and measuring the response strains along different directions in the sample by the gauges of the strain cell. The measured data from an OC test are usually analysed in a classical linear regression to obtain an estimate of the test-specific stress tensor based on the constitutive assumption that the rock is continuous, homogeneous, isotropic and linearly elastic (CHILE) as

$$\epsilon_i = \boldsymbol{c}_i^\top \boldsymbol{s} + e_i = \mu_{\epsilon_i} + e_i \quad (i = 1, 2, \ldots, n) \tag{1}$$

where ϵ_i is the strain measured by strain gauge i; $\boldsymbol{s} = \begin{bmatrix} \sigma_x & \tau_{xy} & \tau_{xz} & \sigma_y & \tau_{yz} & \sigma_z \end{bmatrix}^\top$ is the vector of the six unknown normal and shear components of the stress tensor referred to an *x-y-z*

* Corresponding author.
** Corresponding author.
E-mail addresses: yu.feng@ngi.no (Y. Feng), gaok@sustech.edu.cn (K. Gao).
Peer review under responsibility of Institute of Rock and Soil Mechanics, Chinese Academy of Sciences.

Cartesian coordinate system; $\boldsymbol{c}_i = \begin{bmatrix} c_{1(i)} & c_{2(i)} & c_{3(i)} & c_{4(i)} & c_{5(i)} & c_{6(i)} \end{bmatrix}^\top$ is the covariate vector for strain gauge i, and it is not directly measurable but transformed from other OC test data (e.g. rock's elastic properties) as given in Appendix A; e_i is the random error between the measured strain ϵ_i and its predicted (i.e. mean) value μ_{ϵ_i} by the theoretical linear stress-strain relationship. The number of strain gauges n is usually 9, 12 or 16 depending on the type of strain cell used. This study follows the convention of rock mechanics to process and present stress tensors under the ENU Cartesian coordinate system (i.e. x east, y north and z upwards).

While OC tests usually go through some quality control procedures (Hakala et al., 2003; Sjöberg et al., 2003), the above OC data analysis for stress estimation inevitably suffers from various sources of uncertainty like any other stress determination method. For instance, considering the delicacy of OC test operation and heterogeneity of natural rocks, measurement errors and model inadequacy (e.g. deviation from the constitutive relation) are deemed non-negligible in practice. Additionally, the limited strain measurements ($n = 9, 12$ or 16) from an OC test may significantly underrepresent the true strain field within the overcored rock sample, thus leading to the statistical sampling error. Although some uncertainty sources have been discussed (Amadei and Stephansson, 1997; Ask, 2003b; Hakala et al., 2003), uncertainty quantification and reduction have largely been ignored in the routine OC data analysis in the sense that test-specific stress estimates obtained from the classical regression model are accepted without questioning their reliability.

Recognising this, recently the first author of the present study extended the classical OC regression model to the Bayesian framework, and demonstrated how uncertainties in the estimated stress tensors and their principal stresses can be probabilistically quantified (Feng et al., 2021b). Following this work, a valuable yet unsolved question is, can the reliability of stress states estimated using the OC method be improved in terms of reduced uncertainty? A natural and virtually no-cost solution to this problem is to incorporate existing stress information into OC data analysis. In the context of OC stress measurement campaign, adjacent OC tests are probably the most practical information source, because they are generally available to us given the common practice of conducting multiple closely spaced (of the order of a metre or less) OC tests along a (sub)horizontal borehole. Fig. 1 exemplifies such OC measurement practice where 18 OC tests were made very closely within 25 m in a horizontal borehole from an underground gallery.
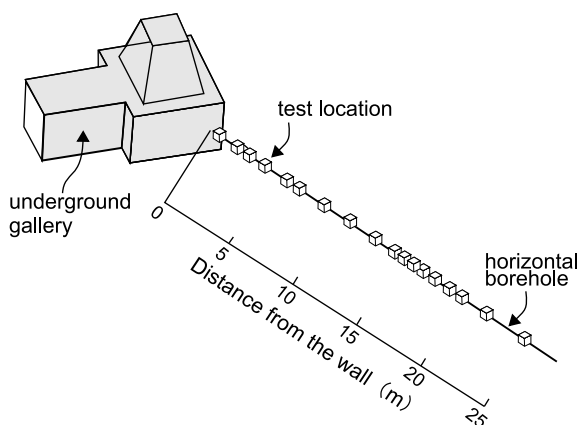
The idea behind this information source is based on the general notion that stress states in close proximity are expected to bear some degree of similarity instead of being completely unrelated, and hence the data of adjacent OC tests may more or less provide certain information to constrain each other's underlying stress states masked by uncertainties. Nevertheless, the customary approach to OC data analysis is to process each OC test independently to generate its respective test-specific stress estimate, thus excluding the possibility of borrowing information across adjacent OC tests. Such independent analysis of each individual data group (e.g. data from each OC test) is statistically known as the no pooling approach. As demonstrated already in Feng et al. (2021b) and later on in this paper, this customary no pooling analysis of OC test data is liable to yield unreliable stress estimates due to various sources of uncertainty. In order for improved stress estimation, a rigorous data pooling approach is needed to allow for borrowing of information across adjacent OC tests.

There are two general approaches to pooling data from different sources that may share similar characteristics (e.g. adjacent OC tests herein), namely complete pooling and partial pooling. In complete pooling, data groups from multiple sources are simply combined into one single large dataset for a holistic analysis ignoring variation between groups. In other words, group-specific parameters are assumed identical. Partial pooling is a generalisation of the two extremes of no pooling and complete pooling, and is usually implemented via Bayesian hierarchical modelling (BHM) that allows for variation and hence similarity between groups, thereby enabling borrowing of information across similar groups (Gelman and Hill, 2006; Lunn et al., 2012).

In geotechnical engineering, the complete pooling approach has been widely used to address the issue of limited site-specific data in the sense that a generic database is compiled from geologically similar sites to estimate some common geotechnical parameters. It is increasingly recognised that such complete data pooling ignoring between-site heterogeneity tends to give overconfident estimates of geotechnical parameters that do not well characterise individual sites; in recent years, partial pooling via BHM that explicitly models data heterogeneity/similarity has received a growing interest when combining information from different sources (e.g. Zhang et al., 2016; Lu et al., 2018; Bozorgzadeh et al., 2019; Ching et al., 2021; Xiao et al., 2021). These studies have demonstrated the power of the BHM framework to allow borrowing information across groups (e.g. sites and projects) for improved geotechnical parameter estimates in various geotechnical contexts.

Motivated by these successful applications, this paper presents a novel application of the BHM framework to the important rock engineering problem of in situ stress estimation with the OC method. More specifically, this paper demonstrates how partial pooling analysis of adjacent OC tests via BHM can be performed to borrow information across tests to improve individual test-specific stress estimates compared to the customary no pooling analysis in terms of reduced uncertainty. In the context of rock stress estimation, complete pooling of multiple adjacent OC tests is clearly inappropriate, because it does not allow for variation between test-specific stress states (i.e. assumes identical stress states) whereas stresses are known to show spatial variability even at the small scale of adjacent OC tests. Hence, complete pooling of adjacent OC tests will not be explored in this study. Note that this study is not intended to develop a new stress measurement method compared with existing ones, but rather to enhance the customary interpretation approach (i.e. no pooling regression) to OC tests—a widely used and also ISRM suggested stress measurement method—for improved OC stress estimation.

Fig. 1. Multiple adjacent OC tests conducted along a horizontal borehole (modified after Obara and Sugawara, 2003).

2. Regression models for overcoring stress estimation

2.1. No pooling regression

Below we briefly review the customary no pooling regression model for OC stress estimation formulated in the Bayesian framework (Feng et al., 2021b). This no pooling modelling approach will be benchmarked as the current practice to demonstrate the advantages of the proposed partial pooling model in stress estimation.

Given a set of adjacent OC tests indexed by j ($j = 1, 2, ..., J$), the customary no pooling model is simply a separate regression analysis of each individual OC test j via Eq. (1) to generate its own test-specific stress estimate, and can be written as

$$\left.\begin{aligned} \epsilon_{ij} &= \boldsymbol{c}_{ij}^\top \boldsymbol{s}_j + e_{ij} \\ e_{ij} &\sim \text{Normal}\left(0, \varsigma_j^2\right) \end{aligned}\right\} \tag{2}$$

where ς_j is the standard deviation of the normally distributed errors e_{ij} ($i = 1, 2, ..., n_j$) for each OC test j. Eq. (2) is called the likelihood function, and its unknown parameters, test-specific stress vectors $\boldsymbol{s}_j$ and standard deviations ς_j, require specification of prior distributions in the Bayesian framework. Considering the realistic magnitude range of rock stress tensors and OC strains, the following weakly informative priors are adopted for each test-specific $\boldsymbol{s}_j$ and ς_j ($j = 1, 2, ..., J$) to reflect a lack of prior knowledge about their specific values as

$$\left.\begin{aligned} \boldsymbol{s}_j &\sim \text{MVN}\left([10\ 0\ 0\ 10\ 0\ 10]^\top \text{MPa},\ \text{diag}\left(25^2, 7.5^2, 7.5^2, 25^2, 7.5^2, 25^2\right) \text{MPa}^2\right) \\ \varsigma_j &\sim \text{Half-normal}\left(0, \left(75 \times 10^{-6}\right)^2\right) \end{aligned}\right\} \tag{3}$$

where MVN denotes the multivariate normal distribution and diag($\cdot$) is the diagonal matrix operator. Here, a multivariate prior instead of six separate univariate priors is assigned to $\boldsymbol{s}_j$ in order to allow for the correlation between stress tensor components, and half-normal distribution with positive support is chosen as the prior choice for ς_j. Feng et al. (2021b) have shown that the priors of Eq. (3) are suitable since they do not introduce influential information into OC stress estimation as intended. A detailed interpretation of these priors can be found in Feng et al. (2021b).

2.2. Partial pooling regression

In partial pooling (hierarchical modelling), parameters (e.g. $\boldsymbol{s}_j$) of different data groups are loosely assumed to be similar rather than necessarily independent or identical as implied respectively in the no pooling and complete pooling approaches, and this is why the former approach can be thought of as a generalisation of the latter two. Such similar parameters, formally known as *exchangeable* parameters, can be regarded as being drawn from a higher-level population distribution, and hence in the Bayesian framework can be induced by assigning a common prior distribution with unknown hyperparameters (Lunn et al., 2012; Gelman et al., 2013). As such, each individual group-specific parameter is estimated not only directly from the data of its own group, but also indirectly from the data of all other groups via the parameter population model. It

is this logical borrowing of information across similar data sources that leads to the so-called partial pooling of data.

In partial pooling of adjacent OC tests, exchangeable parameters are their test-specific stress states that are assumed to be similar (neither restrictively independent nor identical). This underlying assumption is reasonable considering the small scale of adjacent OC tests with no observed significant geological discontinuities as defined in this study. Indeed, when performing stress measurement at a much larger scale (especially vertically) and/or encountering significant geological discontinuities, the partial pooling approach should be properly examined taking into account the gradient and discontinuity of the stress field.

While it is possible to have hierarchical (i.e. exchangeable) standard deviations ς_j besides $\boldsymbol{s}_j$, here they are estimated independently as in the no pooling model since they are not the focus of this paper and also we do not have a particular justification for their exchangeability. Therefore, the Bayesian hierarchical regression model for adjacent OC tests $j = 1, 2, ..., J$ is written as

$$\left.\begin{aligned} \epsilon_{ij} &= \boldsymbol{c}_{ij}^\top \boldsymbol{s}_j + e_{ij} \\ e_{ij} &\sim \text{Normal}\left(0, \varsigma_j^2\right) \end{aligned}\right\} \tag{4}$$

where all $\boldsymbol{s}_j$ are assigned a common multivariate normal (MVN) prior with unknown hyperparameters to accommodate their similarity (i.e. exchangeability) and each ς_j is assigned the same prior with no pooling as

$$\left.\begin{aligned} \boldsymbol{s}_j &\sim \text{MVN}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s) \\ \varsigma_j &\sim \text{Half-normal}\left(0, \left(75 \times 10^{-6}\right)^2\right) \end{aligned}\right\} \tag{5}$$

In a fully Bayesian setting, the two unknown hyperparameters of mean vector $\boldsymbol{\mu}_s \in \mathbb{R}^6$ and covariance matrix $\boldsymbol{\Sigma}_s \in \mathbb{R}^{6 \times 6}$ themselves require specification of prior distributions (known as hyperpriors or hierarchical priors), and are respectively given the following weakly informative hyperpriors as

$$\begin{aligned} \boldsymbol{\mu}_s \sim \text{MVN}\big(&[10\ 0\ 0\ 10\ 0\ 10]^\top \text{MPa}, \\ &\text{diag}\left(25^2, 7.5^2, 7.5^2, 25^2, 7.5^2, 25^2\right) \text{MPa}^2\big) \end{aligned} \tag{6}$$

and

$$\left.\begin{aligned} \boldsymbol{\Sigma}_s &= \text{diag}(\varsigma_s)\boldsymbol{\Omega}_s\text{diag}(\varsigma_s) \\ \varsigma_s &\sim \text{Half-normal}\left(0, 5^2\right) \text{MPa} \\ \boldsymbol{\Omega}_s &\sim \text{LKJ}(\eta = 5) \end{aligned}\right\} \tag{7}$$

Eq. (6) is mathematically equivalent to six separate univariate normal hyperpriors, and such form allows for a single-line representation of six probability density functions and also a ready extension to incorporate possible correlation between the mean stress components when such prior information is available. In order to express prior information more flexibly and robustly for
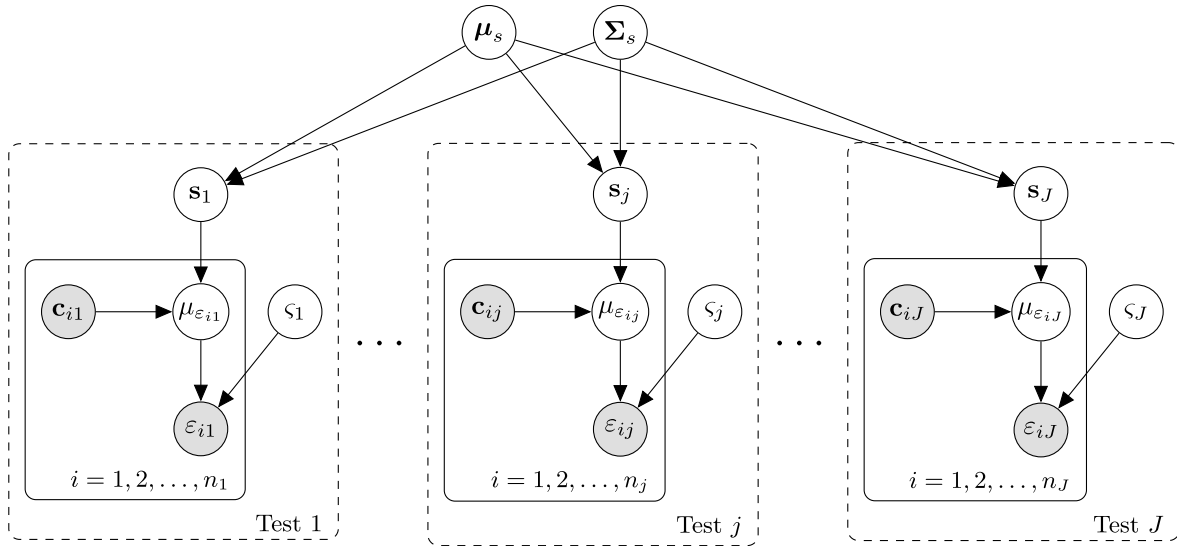
Fig. 2. Bayesian hierarchical model for partial pooling of adjacent OC tests. Nodes of open circles and filled circles represent unobservable variables and observable variables, respectively, directed edges indicate conditional dependence, and solid plates indicate repetition for each i.

covariance matrix, it is recommended to avoid directly specifying a single prior on it, but rather to specify separate priors on its decomposed standard deviations and correlation matrix (Barnard et al., 2000; Gelman et al., 2013). Eq. (7) states that the covariance matrix Σ_s is decomposed into a standard deviation vector ς_s and a correlation matrix Ω_s, and then a weakly informative prior distribution is separately assigned to Ω_s and each component of ς_s. In Eq. (7), LKJ denotes the LKJ distribution named after its authors of Lewandowski et al. (2009), and it is a common prior choice for correlation matrix and is essentially a multivariate generalisation of the beta distribution controlled by a single shape parameter η. A detailed discussion on the prior specification for covariance matrix can be found in statistics texts (Gelman et al., 2013; Kruschke, 2014; McElreath, 2019) as well as in stress estimation works (Feng et al., 2020, 2021a).

Fig. 2 illustrates the proposed Bayesian hierarchical model for partial pooling of adjacent OC tests using the directed acyclic graph. It is seen that for each OC test j, its test-specific stress parameter s_j is learned from its own data and is also informed by the data from all the other OC tests via the higher-level MVN population distribution parameterised by μ_s and Σ_s, thus enabling borrowing of information across adjacent OC tests. Such partial pooling of adjacent OC tests allows information to be borrowed by all individual test-specific stress parameters simultaneously rather than by only one certain stress parameter.

Note that in Bayesian inference, prior choices for model parameters are never unique in terms of distributional type and hyperparameter values. When prior information is not available, weakly informative priors that are based on some contextual knowledge are now strongly advocated as the default prior choice, because the traditional flat/vague priors placing roughly equal probabilities over an unrealistically wide parameter range are known to be not robust and may lead to biased posterior results in case of limited data and/or complex models (Gelman et al., 2017; Lemoine, 2019; Gelman and Yao, 2021). Throughout this study, the adopted weakly informative priors provide only some realistic soft ranges for unknown parameters, and are generally non-informative to the posterior results. Such choices of weakly informative priors in the context of stress analysis have been demonstrated to be suitable (i.e. not significantly influence posteriors) in the authors' previous works (Feng et al., 2021a, b).

2.3. Bayesian posterior computation

Given prior distributions for unknown parameters $p(\theta)$ and the likelihood function for observed data $p(y\,|\theta)$, the joint posterior distribution of parameters can be expressed, according to Bayes' rule, $p(\theta\,|y) \propto p(\theta)p(y\,|\theta)$. In BHM, the model parameters θ themselves may have hierarchical prior distributions $p(\theta\,|\varphi)$ conditioned on unknown hyperparameters φ; in this case, the joint posterior distribution of all unknown parameters is expressed as $p(\theta, \varphi\,|y) \propto p(\varphi)p(\theta\,|\varphi)p(y\,|\theta)$.

However, posterior distributions usually do not have a closed analytical form, particularly for complex models like hierarchical models. Hence, Bayesian computation often entails numerical approximation methods such the commonly used Markov chain Monte Carlo (MCMC) simulation method (Gelman et al., 2013). In this study, all Bayesian models were fitted using the MCMC method (more specifically, the Hamiltonian Monte Carlo algorithm) to obtain parameters' posterior distributions. Three Markov chains with different initial values were simulated in parallel to have confirmed their convergence. Each chain consists of 3000 draws after the first 5000 iterations were discarded in the warm-up period, and hence a total of 9000 posterior draws were used to approximate the posterior distribution of each parameter.

3. Data description

To demonstrate the performance of the partial pooling model for in situ stress estimation, this paper employs five case studies with sets of adjacent OC tests made at the Äspö Hard Rock Laboratory (HRL) of the Swedish Nuclear Fuel and Waste Management Company. The Äspö HRL, operated since 1986, is a renowned geoscientific research facility for testing and developing safe technological solutions for a final geological repository of spent nuclear fuel (see its overview in Fig. 3a), and it has conducted extensive high-quality stress measurements that have been well documented publicly.

The Äspö HRL consists of four main rock types: Småland granite, diorite, greenstone and apite; it also involves several narrow, steeply dipping transmissive major fractures trending WNW-NNW (Ask, 2004). The five OC test sets were all conducted in relatively continuous rocks using the CSIRO HI strain cell with the location of
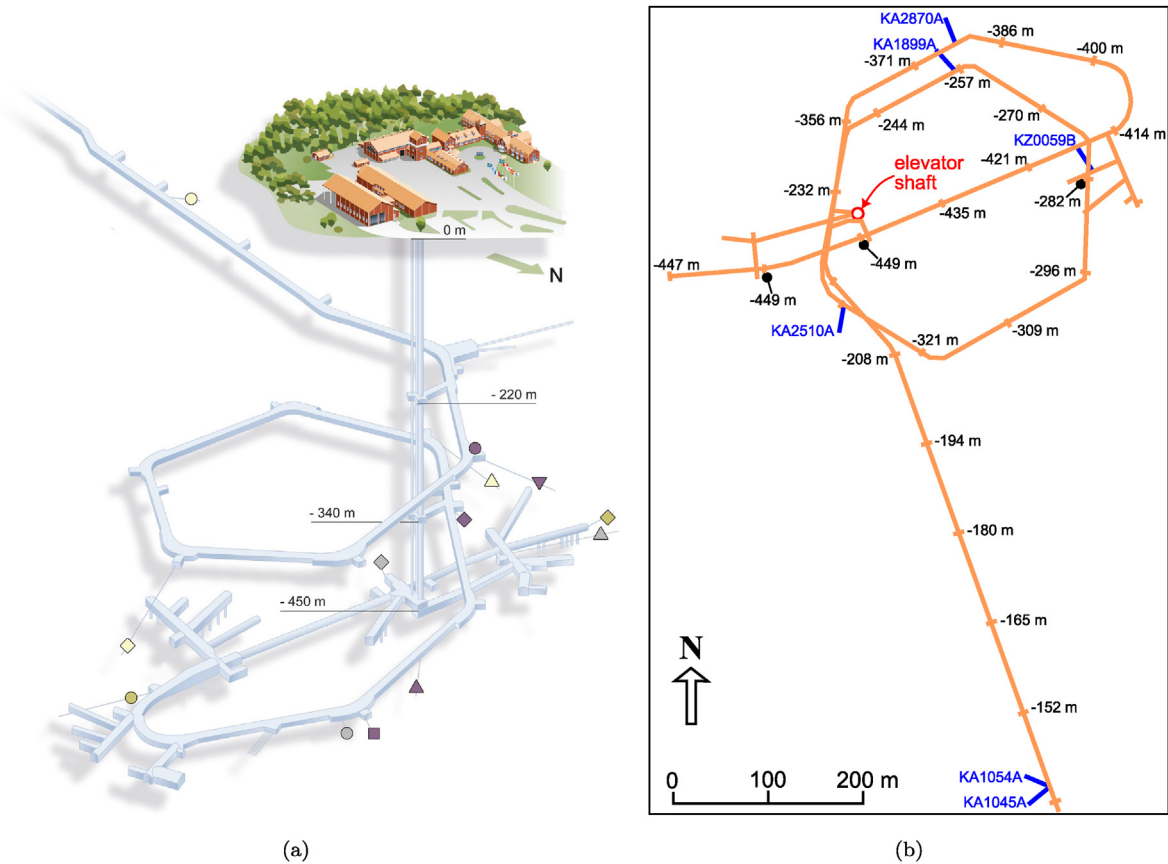
Fig. 3. (a) Overview of the Äspö HRL (from Osterholz et al., 2022), and (b) Map of the Äspö HRL showing the locations (blue lines) of the (sub)horizontal OC stress measurement boreholes used in this study.

their (sub)horizontal measurement boreholes that are not crossed by major fractures, as shown in Fig. 3b. We emphasize that in this study, adjacent OC tests are explicitly defined on the individual borehole level which is not encountered by considerable discontinuities, such that their underlying test-specific stress states can be reasonably assumed to be statistically exchangeable.

An overview of the employed OC tests is provided in Table 1, and all raw OC test data are available in Section 7 with the KA2510A OC test data presented as an example in Table B1. Note that boreholes KA1045A and KA1054A are made next to each other and hence their seven OC tests are analysed as a single set denoted as KA10X, and the average between-test spacing in each measurement borehole is about only 0.4 m. A detailed description of all OC stress measurements using the CSIRO HI at the Äspö HRL is provided by Ask (2003a).

4. Results

4.1. Posterior estimation

Once the posterior distribution of a parameter is obtained, its mean and spread (e.g. 95% credible interval (CI) and standard deviation) can be respectively taken as a point estimate and the associated uncertainty of that parameter. Taking KA10X OC test 1 as an example, Fig. 4a and b displays the trace plots of the three simulated Markov chains of its test-specific stress component σ_x from the no and partial pooling models, respectively, and the three well-mixed and stable chains visually indicate their convergence to the target posterior distribution; Fig. 4c shows the histograms of

Table 1
Overview of the adjacent OC test sets.

Borehole	Depth (m)	Trend (°)	Plunge (°)	Number of tests	Strain cell
KA1045A[a]	143	206.4	−5.3	4	9-gauge CSIRO HI
KA1054A[a]	143.7	293.8	−5.1	3	9-gauge CSIRO HI
KA2510A	334.7	190.8	2.5	6	12-gauge CSIRO HI
KA1899A	256.6	317	1.5	5	12-gauge CSIRO HI
KA2870A	379.3	338.7	1.2	5	12-gauge CSIRO HI
KZ0059B	416.5	340.2	2.0	5	12-gauge CSIRO HI

[a] OC tests from the KA1045A and KA1054A boreholes are treated as a single adjacent test set denoted as KA10X.

the two simulated posterior distributions, from which their posterior means and 95% CIs are calculated. Compared to the no pooling model, the much narrower 95% CI from the partial pooling model suggests less uncertainty in its σ_x estimate.

For the five sets of adjacent OC tests analysed, the posterior stress estimates for the KA10X and KA2510A test sets are discussed in detail here, and the results for the other three test sets are given in Appendix C. Figs. 5 and 6 show the posterior test-specific stress tensor estimates for KA10X and KA2510A OC tests from the no and partial pooling models, and the results are arranged by stress tensor component.

First, Fig. 5 interestingly shows that KA10X OC tests 2−7 yield similar no pooling estimates of τ_{xy} of around −1.5 MPa with their 95% CIs substantially overlapping between −3 and 0 MPa, while OC test 1 gives a significantly different τ_{xy} estimate of −3.5 MPa accompanied by a much wider 95% CI of (−8.3, 1.6) MPa. Such obviously extreme no pooling estimates are also observed for the
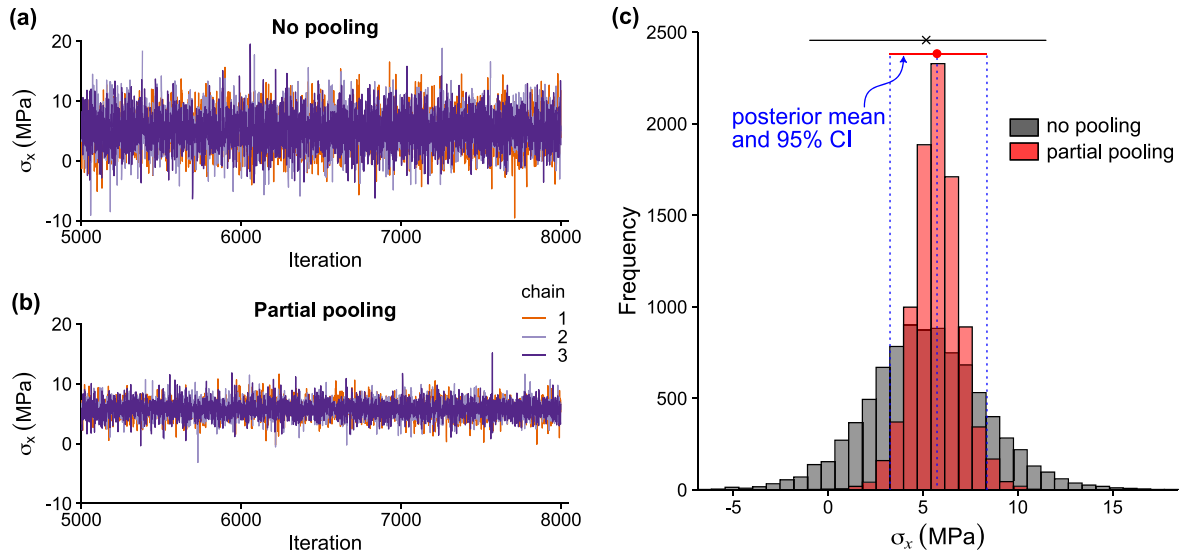
Fig. 4. Example trace plots and histograms of the simulated posterior draws of σ_x for KA10X OC test 1.
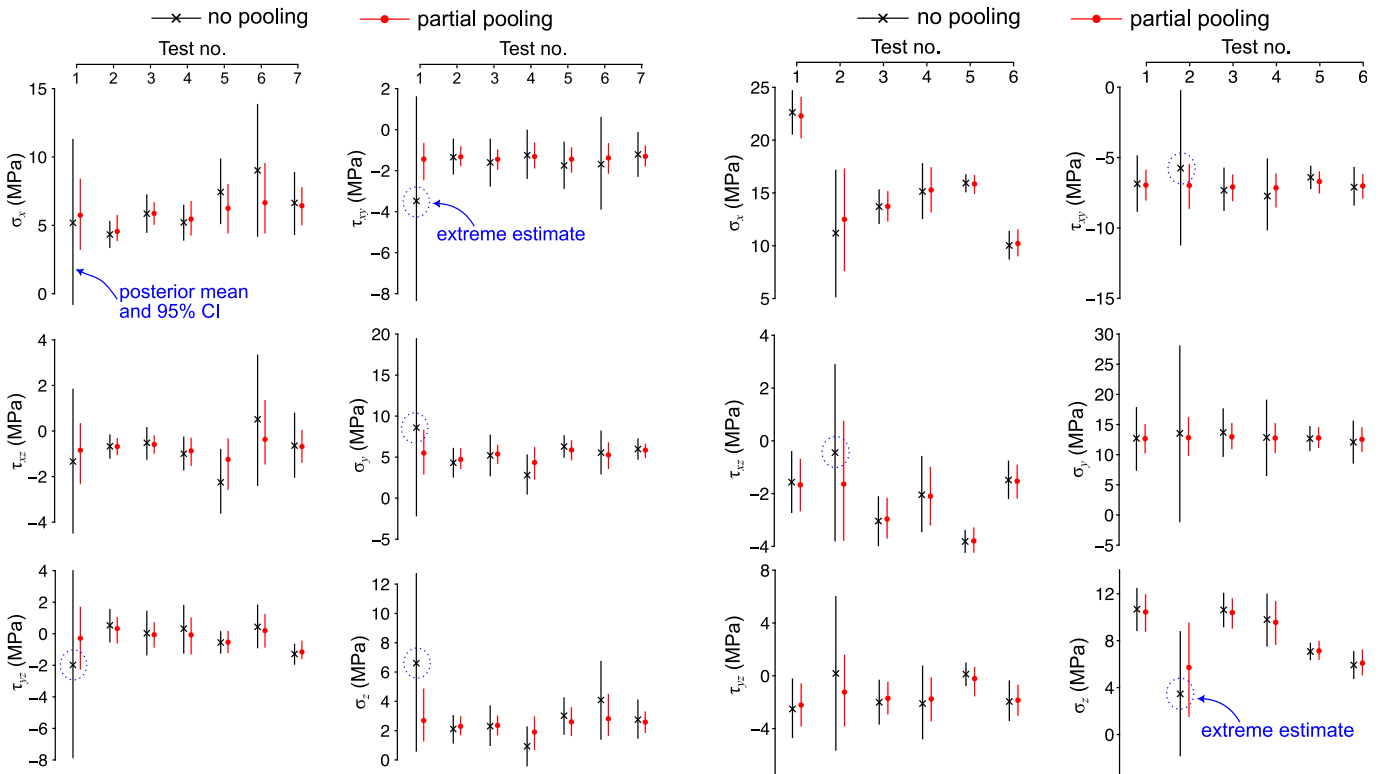

Fig. 5. Posterior estimates of test-specific stress tensors for KA10X OC tests from the no and partial pooling models.

Fig. 6. Posterior estimates of test-specific stress tensors for KA2510A OC tests from the no and partial pooling models.

stress components σ_y, σ_z and τ_{yz} for OC test 1. Despite the natural stress variability in space, this single extreme stress tensor estimate for OC test 1 appears fairly suspect given the overall comparability between the other six test-specific stress tensor estimates at adjacent locations. Indeed, for OC test 1, the wide 95% CIs indicate that the sources of uncertainty discussed in Section 1 have brought in a large combined error which leads to its extreme stress tensor estimate. For the KA2510A set shown in Fig. 6, the no pooling model also generates a similarly extreme stress tensor estimate for OC test 2, whose τ_{xy}, τ_{xz} and σ_z components exhibit substantial differences

and considerably larger uncertainties than the other five adjacent OC tests.

The extreme stress estimates discussed above statistically highlight the insufficiency of the customary no pooling analysis of OC test data for in situ stress estimation, thus calling for an approach that can incorporate additional information to reduce uncertainty in test-specific stress estimates. The insufficiency of the no pooling model arises from the fact that it estimates each individual test-specific stress parameter based solely on the data from

the associated single OC test that is subject to various uncertainty sources. It is worth noting that identifying the unreliable test-specific stress estimates would have been not possible without uncertainty quantification, and this is exactly the main contribution of Feng et al. (2021b) by extending the classical OC regression model to the Bayesian framework.

For each KA10X and KA2510A OC test in Figs. 5 and 6, the partial pooling model generally gives a different yet less uncertain estimate than that of the no pooling model, which is manifested in the narrower hierarchical 95% CIs for all the six stress tensor components. In particular, the two extreme no pooling stress estimates for KA10X OC test 1 and KA2510A OC test 2 become greatly more reliable when partial pooling is applied. A similar but less notable overall phenomenon is observed for the other three adjacent OC test sets (i.e. KA1899A, KA2870A and KZ0059B) as presented in Fig. C.1.

These results demonstrate that partial pooling of adjacent OC tests via BHM indeed allows information to be borrowed across these tests, and hence yields improved test-specific stress estimates with reduced uncertainties simultaneously for all individual tests than they are customarily analysed independently (i.e. no pooling). The uncertainty reduction brought by partial pooling is particularly notable for those highly unreliable stress estimates from no pooling. Hence, when adjacent OC tests are available at a site, they should be analysed preferably using partial pooling instead of no pooling for improved stress estimation. Fortunately, as noted earlier, adjacent OC tests as defined in this study are usually available in OC stress measurement campaigns.

4.2. Overall uncertainty reduction

Since a full stress tensor is composed of six distinct components, it is more intuitive to use a scalar value to compare the overall uncertainty in the estimated six-dimensional stress vector $\boldsymbol{s}$ from the two models considered. Here, we adopt the extensively used scalar metric of multivariate dispersion called *effective standard deviation* for this purpose, which is defined as the $2d$th root of the determinant of the covariance matrix of a d-dimensional random variable. For the six-dimensional stress vector $\boldsymbol{s}$, the effective standard deviation of its posterior distribution is computed as (Feng et al., 2021a):

$$SD_{\text{eff}} = \sqrt{|\text{cov}(\boldsymbol{s})|^{1/6}} = \sqrt{\left|\text{cov}\left(\left\{\boldsymbol{s}^{(k)}\right\}_{k=1}^{K}\right)\right|^{1/6}} \tag{8}$$

where $\text{cov}(\cdot)$ and $|\cdot|$ denote the operator for covariance matrix and matrix determinant, respectively; and $\boldsymbol{s}^{(k)}$ is the kth MCMC draw ($k = 1, 2, …, K$) from the posterior distribution of $\boldsymbol{s}$.

Fig. 7 shows the effective standard deviation of each posterior test-specific $\boldsymbol{s}$ for each of the five adjacent OC test sets from the two models considered. Compared with no pooling, partial pooling of adjacent OC tests leads to varying degrees of reduction in the overall uncertainty in all individual estimated test-specific stress tensors. For the five adjacent OC test sets (28 tests in total) being analysed, most test-specific stress tensor estimates have seen an overall uncertainty reduction by at least 15% from no pooling to partial pooling. This is a clear indication that the degree of improvement in OC stress estimation by partial pooling is
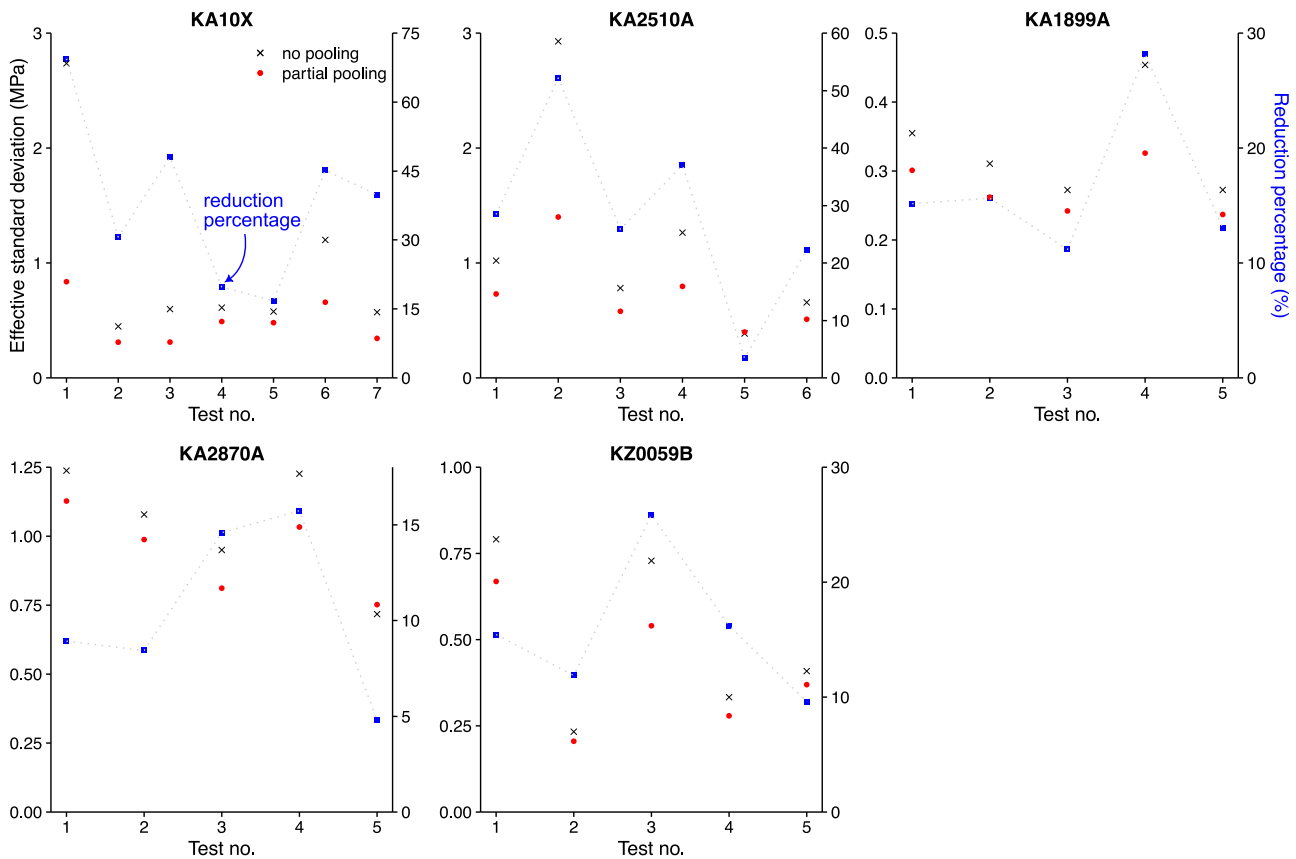

Fig. 7. Overall uncertainty reduction for test-specific stress tensor estimates between no pooling and partial pooling modelling.
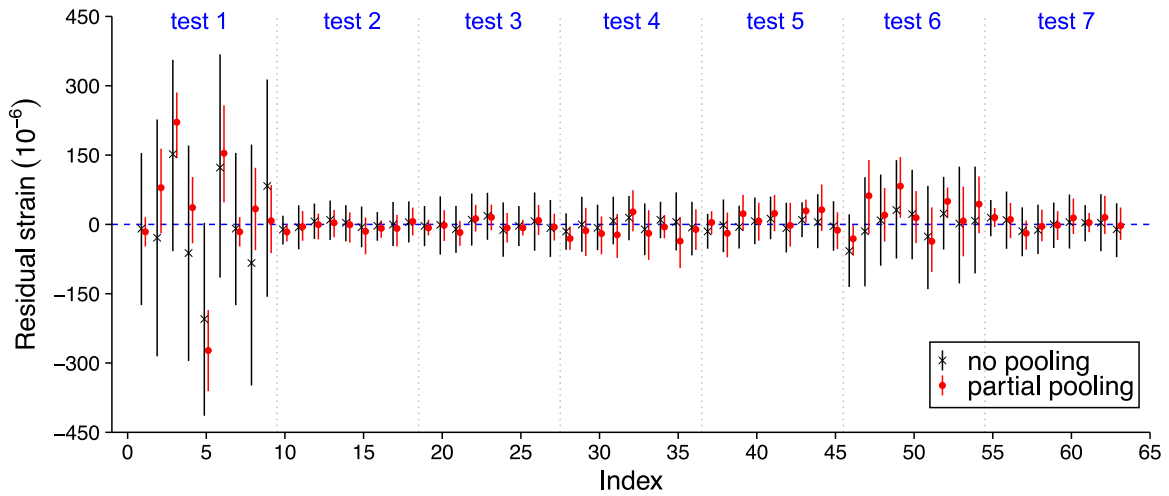
Fig. 8. Residuals from the no and partial pooling models for the KA10X OC test set.

meaningful. For the KA10X and KA2510 OC test sets under focus, their relative overall uncertainty reduction is particularly notable. For example, the overall uncertainty in the estimated stress tensor for KA10X test 1 and KA2510A test 2 are reduced remarkably by 70% and 52%, respectively, which is expected from our earlier discussion on Figs. 5 and 6.

The observed overall uncertainty reduction for all five case studies further confirms the assertion made in Section 4.1 that partial pooling of adjacent OC tests can reduce uncertainty in all individual test-specific stress tensors estimated from the customary no pooling analysis, particularly for those unreliable ones.

5. Model comparison

Although not necessarily of interest in OC data analysis, a predictive model comparison is presented in this section to give some additional insights into the two models and their benefits from a different perspective. Fig. 8 illustrates the posterior means and 95% CIs of the residual strains (i.e. difference between the measured and fitted strains) from the no pooling and partial pooling models for the KA10X OC test set (7 tests with each having 9 measured strains). It should be noted that the residual plot is equivalent to the common plot of fitted versus measured values, and here residuals themselves have posterior distributions as the fitted values are calculated based on uncertain stress parameters.

In Fig. 8, residual estimates from no pooling are generally closer to zero but accompanied by larger uncertainty. At first sight, these smaller absolute residuals suggest that no pooling gives a better fit to the OC test data than partial pooling, yet unfortunately this is a sign of overfitting to the limited data within each OC test. To demonstrate this, these two models are compared based on their out-of-sample predictive accuracy (generalisation error equivalently). A common measure of predictive accuracy for probabilistic models is the expected log predictive density (elpd), which can be estimated by two general types of approaches: cross-validation, and information criteria (e.g. AIC, DIC and more recently WAIC) that can be viewed as approximations to different versions of the former (Gelman et al., 2014; Vehtari et al., 2017). Here, the widely advocated method, leave-one-out cross-validation (LOO-CV), is adopted to evaluate the elpd of the two Bayesian models for the five adjacent OC test sets used in this study.

Given some data $y_1, y_2, \ldots, y_n$, LOO-CV evaluates the elpd by repeatedly holding out one data point y_i ($i = 1, 2, \ldots, n$) and refitting the model with parameters θ to the remaining data y_{-i}. The Bayesian LOO-CV estimate of the elpd is defined as

$$\text{elpd}_{\text{loo}} = \sum_{i=1}^{n} \ln p(y_i | y_{-i}) = \sum_{i=1}^{n} \ln \int p(y_i | \theta) p(\theta | y_{-i}) d\theta \quad (9)$$

where $\ln p(y_i | y_{-i})$ and $p(y_i | \theta)$ denote the log predictive density and likelihood for the new (left-out) data point y_i, respectively; and $p(\theta | y_{-i})$ is the posterior distribution of θ based on data y_{-i}. In the context of OC data analysis, it is the measured strains ϵ_i that are left out for comparison to their predicted values. A detailed discussion on Bayesian model comparison can be found in standard statistics texts (e.g. Gelman et al., 2013, 2014; Vehtari et al., 2017) and in the relevant geoscience works (Bozorgzadeh and Bathurst, 2019; Feng et al., 2022).

Table 2 summarises the elpd estimates and their difference for the two Bayesian models evaluated on each of the five adjacent OC test sets employed. Larger negative elpd indicates higher expected predictive accuracy. For the KA10X test set, the LOO-CV elpd estimates for the no pooling and partial pooling models are −341.6 and −328.3 with their standard errors being 7.5 and 7.2, respectively, and the difference between the two elpd values (no pooling − partial pooling) is −13.3 with a standard error of 3.8. Since this elpd difference is large considering the log scale and is 3.5 standard errors away from zero, it can be asserted that partial pooling has a remarkably better predictive performance than no pooling for the KA10X test set. The same assertion can be made for the KA2510 test set. For the remaining three test sets (KA1899A, KA2870A and KZ0059B), the elpd differences are relatively small and roughly 1 to 1.6 standard errors away from zero, indicating that

Table 2
Model comparison using leave-one-out cross-validation.

Model	$elpd_{(SE^*)}$				
	KA10X	KA2510A	KA1899A	KA2870A	KZ0059B
No pooling	$-341.6_{(7.5)}$	$-423.4_{(10.1)}$	$-282.5_{(4.0)}$	$-355.3_{(5.9)}$	$-301.3_{(5.9)}$
Partial pooling	$-328.3_{(7.2)}$	$-410.1_{(9.9)}$	$-278.7_{(4.8)}$	$-353.0_{(5.9)}$	$-298.7_{(6.6)}$
Difference**	$-13.3_{(3.8)}$	$-13.3_{(3.7)}$	$-3.8_{(2.4)}$	$-2.3_{(2.2)}$	$-2.6_{(2.2)}$

* Standard error of estimated elpd.
** No pooling - Partial pooling.

partial pooling outperforms no pooling only to a moderate extent. Yet overall, all five OC test sets indicate that the partial pooling model has a superior predictive performance than the customary no pooling model.

The model comparison shows that although giving a better fit to the OC test data, the no pooling model fails to predict unseen new data (i.e. strains) as accurately as the partial pooling model, thus signalling a tendency of overfitting. Essentially in partial pooling, additional information borrowed across adjacent OC tests serves as regularisation to sacrifice some goodness of fit for reduced uncertainty in the estimated stress parameters (bias-variance trade-off in machine learning parlance), whose combined effect leads to an overall reduction in the prediction error and hence prevention of overfitting to the noisy data. This suggests that the information contained between adjacent OC tests is indeed relevant and effective in the five case studies, otherwise its introduced bias will overtake the effect of reduced stress uncertainty to thus yield an increased prediction error. The predictive comparison results further confirm the superiority of partial pooling over no pooling when analysing multiple adjacent OC tests.

It is worth noting that the goal of OC data analysis is the estimation of test-specific stress state parameters; hence, the no pooling and partial pooling models were actually evaluated based on the reliability of estimated stress parameters (see Section 4) rather than the predictive accuracy of OC strains. In predictive modelling cases, it is recommended that the models be compared and selected based on their predictive performance on the response variables of interest.

6. Discussion

Uncertainty in the estimated stresses by the OC method can be reduced by several aspects, such as improving the measuring device, following careful test procedures and performing quality control procedures. In fact, these time-consuming and/or costly measures have already been integrated into the development and routine application of the OC method, and they mainly target to minimise the impact of measurement errors on the estimated test-specific stresses at the pre-analysis stage. Despite being effective, these measures cannot eliminate measurement errors considering the delicate OC test operation in complex underground conditions. Moreover, other sources of uncertainty discussed in Section 1 still persist. Consequently, the no pooling analysis of individual OC tests without additional information incorporated is liable to yield unreliable test-specific stress estimates as confirmed in the five case studies.

Given the information source of adjacent OC tests that is usually available, the proposed Bayesian partial pooling approach offers a practical solution at virtually no additional cost for improved OC stress estimation at the analysis stage. In this regard, this approach constitutes a highly valuable complement to the current practice of the OC method for rock stress estimation. Nevertheless, partial pooling for information borrowing via BHM is not without restrictions when applied to interpretation of OC tests to obtain test-specific stress estimates.

We emphasize again the proposed approach should be applied to adjacent OC tests as defined in this study, that is, those close to each other at the individual borehole scale which is not crossed by considerable discontinuities (also see Section 3). This application restriction is to ensure that the adjacent OC tests analysed can be reasonably assumed to sample from the same stress field continuum within the small local rock volume involved, and thus the exchangeability/similarity assumption underlying BHM can be honoured to a great extent. In other words, the proposed partial pooling approach is not intended for interpreting multiple OC tests at the larger cross-borehole or even site scale, where discontinuities and/or stress gradients may play a significant role and must be taken into account properly.

Various probabilistic assumptions made about the proposed partial pooling model like the normal distribution for regression errors may be formally checked and compared with other alternatives for potential further model improvement. Nonetheless, this is a different and boarder topic called (statistical) model checking and selection, and is beyond the scope of this paper focusing on comparing the proposed and customary approaches for interpretation of OC tests under some same probabilistic settings.

In this study, we used weak prior distributions for unknown parameters to demonstrate how the Bayesian partial pooling model allows borrowing information across adjacent OC tests in stress estimation. It is worth noting that other sources of stress information may be available in practice (e.g. faulting regime, vertical stress-depth relation and borehole breakouts), and then the proposed model may be extended, with modification, to integrate such additional information in forms of informative priors to further reduce OC stress estimation uncertainty. For example, if there exists a major fault that controls the general stress regime in a large rock volume containing the set of adjacent OC tests being analysed, the faulting stress regime in terms of stress orientations and relative magnitudes may be encoded into informative priors for the parameters of the so-called stress population (i.e. μ_s and Σ_s). Integrating different sources of stress information to further enhance OC stress estimation under the BHM framework warrants a future investigation.

7. Conclusions

This paper discussed various sources of uncertainty associated with OC stress estimation in rocks, and highlighted that the customary no pooling approach that analyses each individual OC test independently is liable to yield unreliable test-specific stress tensor estimates under the impact of these uncertainty sources. To address this important problem, we proposed that a practical and no-cost solution may be to incorporate into OC data analysis additional stress information implied in adjacent OC tests that are usually available in OC measurement campaigns. Therefore, a Bayesian partial pooling (hierarchical) model was presented for combined analysis of adjacent OC tests.

Five case studies of adjacent OC tests from the Äspö HRL demonstrated that partial pooling of adjacent OC tests indeed allows borrowing of information across adjacent tests, and hence yields improved test-specific stress estimates with reduced uncertainties simultaneously for all individual involved tests than they are customarily analysed independently with no pooling. The case studies also revealed that the overall uncertainty reduction by partial pooling is particularly notable for those unreliable no pooling stress tensor estimates.

This paper also performed a formal predictive model comparison for the five adjacent OC test sets, and revealed that partial pooling gives not only improved stress parameter estimates but also better strain prediction. The model comparison further confirmed that the information borrowed across adjacent OC tests is relevant and effective, thereby preventing the tendency of overfitting in no pooling modelling.

Data and code availability

The overcoring test data and the codes of the Bayesian no pooling and partial pooling models used in this study are available at https://doi.org/10.5281/zenodo.7772400.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jrmge.2023.05.003.

References

Amadei, B., Stephansson, O., 1997. Rock Stress and its Measurement. Springer, Dordrecht, Netherlands.

Ask, D., 2003a. Analysis of Overcoring Stress Data at the Äspö HRL, Sweden - Analysis of Overcoring Rock Stress Measurements Preformed Using the CSIRO HI. Technical Report Swedish Nuclear Fuel and Waste Management Co Stockholm, Sweden.

Ask, D., 2003b. Evaluation of measurement-related uncertainties in the analysis of overcoring rock stress data from Äspö HRL, Sweden: a case study. Int. J. Rock Mech. Min. Sci. 40, 1173—1187.

Ask, D., 2004. New Developments of the Integrated Stress Determination Method and Application to the Äspö Hard Rock Laboratory. PhD thesis. Royal Institute of Technology, Sweden.

Bao, T., Burghardt, J., 2022. A Bayesian approach for in-situ stress prediction and uncertainty quantification for subsurface engineering. Rock Mech. Rock Eng. 55, 4531—4548.

Barnard, J., McCulloch, R., Meng, X.-L., 2000. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. Stat. Sin. 10, 1281—1311.

Bozorgzadeh, N., Bathurst, R.J., 2019. Bayesian model checking, comparison and selection with emphasis on outlier detection for geotechnical reliability-based design. Comput. Geotech. 116, 103181.

Bozorgzadeh, N., Harrison, J.P., Escobar, M.D., 2019. Hierarchical Bayesian modelling of geotechnical data: application to rock strength. Geotechnique 69, 1056—1070.

Ching, J., Wu, S., Phoon, K.-K., 2021. Constructing quasi-site-specific multivariate probability distribution using hierarchical Bayesian model. J. Eng. Mech. 147, 04021069.

Clément, C., Merrien-Soukatchoff, V., Dünner, C., Gunzburger, Y., 2009. Stress measurement by overcoring at shallow depths in a rock slope: the scattering of input data and results. Rock Mech. Rock Eng. 42, 585—609.

Feng, Y., Bozorgzadeh, N., Harrison, J.P., 2020. Bayesian analysis for uncertainty quantification of in situ stress data. Int. J. Rock Mech. Min. Sci. 134, 104381.

Feng, Y., Gao, K., Mignan, A., Li, J., 2021a. Improving local mean stress estimation using Bayesian hierarchical modelling. Int. J. Rock Mech. Min. Sci. 148, 104924.

Feng, Y., Harrison, J.P., Bozorgzadeh, N., 2021b. A Bayesian approach for uncertainty quantification in overcoring stress estimation. Rock Mech. Rock Eng. 54, 627—645.

Feng, Y., Mignan, A., Sornette, D., Gao, K., 2022. Investigating injection pressure as a predictor to enhance real-time forecasting of fluid-induced seismicity: a Bayesian model comparison. Seismol. Res. Lett. https://doi.org/10.1785/0220220309.

Fouial, K., Alheib, M., Baroudi, H., Trentsaux, C., 1998. Improvement in the interpretation of stress measurements by use of the overcoring method: development of a new approach. Eng. Geol. 49, 239—252.

Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis, Third edition. Chapman and Hall/CRC, Boca Raton, FL.

Gelman, A., Hill, J., 2006. Data Analysis Using Regression and Multilevel/hierarchical Models. Cambridge University Press, Cambridge, UK.

Gelman, A., Hwang, J., Vehtari, A., 2014. Understanding predictive information criteria for Bayesian models. Stat. Comput. 24, 997—1016.

Gelman, A., Simpson, D., Betancourt, M., 2017. The prior can often only be understood in the context of the likelihood. Entropy 19, 555.

Gelman, A., Yao, Y., 2021. Holes in Bayesian statistics. J. Phys. G Nucl. Part. Phys. 48, 014002.

Haimson, B.C., Cornet, F.H., 2003. ISRM suggested methods for rock stress estimation - Part 3: hydraulic fracturing (HF) and/or hydraulic testing of pre-existing fractures (HTPF). Int. J. Rock Mech. Min. Sci. 40, 1011—1020.

Hakala, M., Hudson, J.A., Christiansson, R., 2003. Quality control of overcoring stress measurement data. Int. J. Rock Mech. Min. Sci. 40, 1141—1159.

Hudson, J.A., Harrison, J.P., 1997. Engineering Rock Mechanics: An Introduction to the Principles. Pergamon, Oxford, UK.

Kruschke, J.K., 2014. Doing Bayesian Data Analysis: a Tutorial with R, JAGS, and Stan, Second edition. Academic Press, London, UK.

Lemoine, N.P., 2019. Moving beyond noninformative priors: why and how to choose weakly informative priors in Bayesian analyses. Oikos 128, 912—928.

Lewandowski, D., Kurowicka, D., Joe, H., 2009. Generating random correlation matrices based on vines and extended onion method. J. Multivariate Anal. 100, 1989—2001.

Li, P., Cai, M., Guo, Q., Miao, S., 2019. In situ stress state of the northwest region of the Jiaodong Peninsula, China from overcoring stress measurements in three gold mines. Rock Mech. Rock Eng. 52, 4497—4507.

Liu, J., Ding, W., Gu, Y., Xiao, Z., Dai, J., Dai, P., Chen, X., Zhao, G., 2018. Methodology for predicting reservoir breakdown pressure and fracture opening pressure in low-permeability reservoirs based on an in situ stress simulation. Eng. Geol. 246, 222—232.

Lu, S., Zhang, J., Zhou, S., Xu, A., 2018. Reliability prediction of the axial ultimate bearing capacity of piles: a hierarchical Bayesian method. Adv. Mech. Eng. 10. https://doi.org/10.1177/1687814018811054.

Lunn, D., Jackson, C., Best, N., Thomas, A., Spiegelhalter, D., 2012. The BUGS Book: a Practical Introduction to Bayesian Analysis. CRC Press, Boca Raton, FL.

Maleki, S., Moradzadeh, A., Ghavami Riabi, R., Sadaghzadeh, F., 2014. Comparison of several different methods of in situ stress determination. Int. J. Rock Mech. Min. Sci. 71, 395—404.

McElreath, R., 2019. Statistical Rethinking: a Bayesian Course with Example in R and Stan, Second edition. Chapman and Hall/CRC, New York, NY. https://doi.org/10.1201/9781315372495.

Obara, Y., Sugawara, K., 2003. Updating the use of the CCBO cell in Japan: overcoring case studies. Int. J. Rock Mech. Min. Sci. 40, 1189—1203.

Osterholz, H., Turner, S., Alakangas, L.J., Tullborg, E.-L., Dittmar, T., Kalinowski, B.E., Dopson, M., 2022. Terrigenous dissolved organic matter persists in the energy-limited deep groundwaters of the Fennoscandian Shield. Nat. Commun. 13, 4837.

Sjöberg, J., Christiansson, R., Hudson, J.A., 2003. ISRM Suggested Methods for rock stress estimation - Part 2: overcoring methods. Int. J. Rock Mech. Min. Sci. 40, 999—1010.

Stein, R.S., 1999. The role of stress transfer in earthquake occurrence. Nature 402, 605—609.

Vehtari, A., Gelman, A., Gabry, J., 2017. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Stat. Comput. 27, 1413—1432.

Xiao, S., Zhang, J., Ye, J., Zheng, J., 2021. Establishing region-specific N—Vs relationships through hierarchical Bayesian modeling. Eng. Geol. 287, 106105.

Zhang, C., Chen, Q., Qin, X., Hong, B., Meng, W., Zhang, Q., 2017. In-situ stress and fracture characterization of a candidate repository for spent nuclear fuel in Gansu, northwestern China. Eng. Geol. 231, 218—229.

Zhang, J., Juang, C., Martin, J., Huang, H., 2016. Inter-region variability of Robertson and Wride method for liquefaction hazard analysis. Eng. Geol. 203, 191—203.

Yu Feng obtained his BSc and MSc degrees in Mining Engineering from University of Science and Technology Beijing, in 2013 and 2016, respectively, and his PhD in Geotechnical Engineering from University of Toronto, Canada, in 2021. He is now a postdoctoral researcher at Norwegian Geotechnical Institute (NGI), a world-renowned research institute in geotechnical engineering, especially offshore geotechnics. His research interests include (1) uncertainty and risk assessment in systems of geotechnics and geoscience, (2) crustal stress, (3) offshore wind site characterization, (4) induced seismicity in geo-energy and (5) geologic carbon storage.