Data space inversion for efficient uncertainty quantification using an integrated surface and subsurface hydrologic model

Hugo Delottier[1], John Doherty[2] and Philip Brunner[1]

5 [1] Centre for Hydrogeology and Geothermics (CHYN), University of Neuchâtel, Neuchâtel,
Switzerland
[2] Watermark Numerical Computing, Corinda, QLD.

Correspondence to: Hugo Delottier (hugo.delottier@unine.ch)

Abstract. Hydrological numerical modelling is generally designed to provide predictions of uncertain quantities

10 in a decision-support context. In the implementation of decision-support modelling, data assimilation and uncertainty quantification are often the most difficult and time-consuming tasks. This is because the imposition of history-matching constraints on model parameters usually requires a large number of model runs. Data Space Inversion (DSI) provides an alternative (and highly model-run-efficient) method for predictive uncertainty quantification that avoids the need for parameter estimation. It does this by evaluating covariances between

15 model outputs used for history matching (e.g. hydraulic heads) and model predictions based on model runs that sample the prior parameter probability distribution. By focusing on the direct relationship between model outputs under historical conditions and predictions of system behaviour under future conditions, DSI avoids the need to estimate or adjust model parameters. This is advantageous when using such as Integrated Surface and Subsurface Hydrologic Models (ISSHMs). These models are characterised by long run times, a penchant for

20 numerical instability and/or complex parameterisation schemes that are designed to maintain geological realism. This paper demonstrates that DSI provides a robust and efficient means of quantifying the uncertainties of complex model predictions, at the same time as it provides a basis for complementary linear analyses that can explore issues such as data worth. DSI is applied in conjunction with an ISSHM representing a synthetic but realistic stream-aquifer system. Predictions of interest are fast travel times and surface water infiltration. Linear

25 and nonlinear estimates of prediction uncertainty based on DSI are validated against a more traditional approach to prediction uncertainty quantification which requires adjustment of a large number of parameters. A DSI-generated surrogate model is then used to investigate the effectiveness and efficiency of existing and possible future monitoring networks. This demonstrates the benefits of using DSI in conjunction with a complex numerical model to quantify prediction uncertainty and support data worth analysis in complex hydrogeological

30 environments.

1. Introduction

Numerical hydrological models are typically built to make predictions to support decision-making. Note that in this paper we use the terminology "prediction" to refer to any quantity of interest from a numerical model, whether it is estimated in the future or as a by-product of data assimilation. Generally, Bayesian methods are applied to these models so that the uncertainties associated with predictions of management interest can be quantified and reduced. Through this process, the prior uncertainties of these predictions are constrained by the assimilation of local data to provide estimates of posterior prediction uncertainties. A variety of Bayesian methods are available for this purpose. These include linear methods (James et al., 2009; Dausman et al., 2010) as well as linear-assisted methods such as null space Monte Carlo (Tonkin and Doherty 2009; Doherty 2015). More recently, ensemble methods such as those described by Chen and Oliver (2013) and White (2018) have been turned to this task. The attractiveness of ensemble methods lies in their ability to accommodate a large number of parameters with a high degree of model run efficiency.

In many hydrogeological settings, elongate structural or alluvial features that embody interconnected permeability are of high conceptual relevance as they can have a major impact on management-salient predictions. Representation of these features typically requires the use of categorical parameterisation schemes that maintain their inherent connectivity. Despite the efficiency of ensemble methods, it is difficult to adjust and maintain permeability connectedness as parameter fields are adjusted so that model outputs can respect field measurements (Lam et al., 2020; Juda et al., 2022). Theoretical problems arise from the multi-Gaussian assumption on which ensemble methods are based. Practical problems arise from the sometimes problematic behaviour of complex numerical models when endowed with stochastic parameter fields with high contrasts in hydraulic conductivity. These problems are generally exacerbated when solute transport is simulated in addition to flow. It follows that the use of model partner software such as PEST (Doherty 2022) and PESTPP-IES (White 2018) for parameter field conditioning and prediction uncertainty analysis and reduction is not always feasible. This makes it difficult to pursue simulation integrity through the use of physically based numerical models where potentially information-rich site data must be assimilated. Compromises in model structure, parameterisation or process complexity may therefore be required (see e.g. Delottier et al., 2022).

To overcome these problems, methods have been developed to generate posterior distributions of predictions without the need to adjust the complex parameter fields of large, physically based numerical models. See, for example, Satija and Caers 2015; He et al., 2018; Hermans 2017; Sun and Durlofsky 2017. Computational advantages can be gained by establishing a direct link between historical observations and predictions of management interest. To establish this link, a numerical model of arbitrary process complexity is used. This model is equipped with parameter fields that can best express the hydrogeological characteristics of the simulated system. These may include complex structures that are capable of expressing three-dimensional configurations of hydraulic property heterogeneity. Because no adjustment of parameter fields is required, simulation integrity is maintained regardless of the complexity of hydrogeological site conceptualisation, and regardless of the role played by features of structural or alluvial origin that exhibit complex spatial patterns of enhanced and continuous hydraulic conductivity. Instead, the model is used to build a prior probability distribution of predictions of management interest based on samples of realistic hydraulic property fields. The

75   same model runs that are used to explore prior prediction uncertainty are used to construct a joint probability
       distribution that links historical system behaviour with future system behaviour.

       Some of the methods that adopt this approach to posterior predictive uncertainty analysis attempt to develop
       explicit relationships between historical observations of system behaviour on the one hand and predictions of
80   future system behaviour on the other hand (Satija and Caers 2015; Scheidt et al. 2015). In contrast, other
       methods develop a more implicit relationship between these two (Sun and Durlofsky 2017; Lima et al. 2020).
       Once these relationships have been established, predictions of management interest can be directly conditioned
       by real-world measurements of historical system behaviour. The need for manipulation and estimation of
       parameters is thereby obviated. The numerical burden of predictive uncertainty reduction and quantification is
85   thereby reduced to a manageable level, regardless of the complexity of the numerical model and regardless of the
       complexity of the prior parameter probability distribution.

       In this paper, we employ a "data space inversion" (i.e. DSI) methodology that is similar to that described by
       Lima et al. (2020). At the same time, we extend the methodology to consider the value of existing and future
90   data. Assessment of data worth is based on the premise that the value of data increases with its ability to reduce
       the uncertainties of decision-critical predictions. Because it requires uncertainty quantification, assessing the
       value of data using methods that rely on explicit or implicit (through linear analysis) parameter adjustment can
       be computationally expensive, especially when performed in numerical and parameterisation contexts that
       attempt to respect hydrogeological processes and hydraulic property integrity as much as possible. The
95   numerically inexpensive methodology for data worth assessment that is presented herein can support these
       assessments in modelling contexts where it would otherwise be computationally intractable. It is hoped that this
       can support attempts to achieve the goals of decision-support utility and failure avoidance that are set out in
       articles such as Kikuchi 2017 and Doherty and Moore 2020.

100   All of the methods described here can be implemented using PEST (Doherty, 2022) and/or PEST++ (PEST++
       Development Team, 2022). They are therefore readily available to the wider groundwater community.

       The paper is organised as follows: Section 2 describes the theory behind DSI. In the following section, a
       synthetic alluvial stream-aquifer system is introduced. This is then used to (i) validate DSI estimates of
105   prediction uncertainty and (ii) demonstrate the use of DSI in quantifying the value of existing and new data.

2. Theory

2.1 Statistical Linkage between the Past and the Future

110

Let the vector **o** denote outputs of a model that correspond to field measurements of system behaviour (e.g. system states and fluxes). Meanwhile, we denote the actual field measurements of these quantities by the vector **h**. We use the symbol Z to characterise the operator through which model outputs **o** are calculated from model parameters; the latter are denoted by the vector **k**. Therefore:

115

$$\mathbf{h} = \mathbf{o} + \boldsymbol{\varepsilon} = Z(\mathbf{k}) + \boldsymbol{\varepsilon} \tag{1}$$

The vector $\boldsymbol{\varepsilon}$ denotes the noise associated with field measurements **h**. It may also include that part of model-to-measurement misfit that arises from the imperfect nature of numerical simulation. In the present case we

120 combine both of these into a single "noise" term. It is generally assumed that $\boldsymbol{\varepsilon}$ has a mean of **0**; we denote its covariance matrix by C($\boldsymbol{\varepsilon}$).

Ideally, **k** in equation (1) can represent any model inputs that are incompletely known. Conceptually, there is no limit to the number of elements that comprise the vector **k**, and hence to the complexity of model

125 parameterisation. Typically, elements of **k** represent spatially distributed hydraulic and transport properties. They may also include historical system stresses such as pumping and recharge rates that are incompletely known, as well as some incompletely known specifications of other boundary conditions.

Suppose the model is run over a period for which predictions are made. Let the vector **s** contain simulator

130 predictions of management interest. We depict these as being computed from model parameters using the operator Y. Thus:

$$\mathbf{s} = Y(\mathbf{k}) \tag{2}$$

135 Suppose that the model is run N times over its simulation period, and that this simulation period spans both the time over which history-matching is performed (that is, the "calibration period") and the time over which predictions are required (that is, the "prediction period"); depending on the prediction, these two periods may coincide. Suppose further that on each of the N occasions on which the model is run, the vector **k** comprises a different sample of its prior probability distribution of **k**. The set of N realisations of **o** computed by the model

140 over these N model runs can be collected into a matrix **O** by arranging them in columns. Meanwhile, the set of N realisations of **s** can be similarly collected into a matrix **S**. From these realisations, vectors depicting the mean of **O** and the mean of **S** (i.e. **o** and **s**) are calculated as:

$$\underline{\mathbf{o}} = \frac{\sum_{i=1}^{N} \mathbf{o}_i}{N} \tag{3a}$$

145
$$\underline{\mathbf{s}} = \frac{\sum_{i=1}^{N} \mathbf{s}_i}{N} \tag{3b}$$

Let **O** and **S** designate matrices whose columns are comprised of replicates of **o** and **s**. Consider now the covariance matrix:

150
$$C\left(\begin{bmatrix}\mathbf{o}\\ \mathbf{s}\end{bmatrix}\right) = \begin{bmatrix}C_{oo} & C_{os}\\ C_{so} & C_{ss}\end{bmatrix}$$
(4)

Submatrices appearing on the right side of equation (4) can be calculated from realizations of model outputs using the following equations:

155
$$C_{oo} = \frac{(\mathbf{o}-\underline{\mathbf{o}})(\mathbf{o}-\underline{\mathbf{o}})^{\mathrm{T}}}{N-1}$$
(5a)

$$C_{os} = C_{so}^{\mathrm{T}} = \frac{(\mathbf{o}-\underline{\mathbf{o}})(\mathbf{s}-\underline{\mathbf{s}})^{\mathrm{T}}}{N-1}$$
(5b)

$$C_{ss} = \frac{(\mathbf{s}-\underline{\mathbf{s}})(\mathbf{s}-\underline{\mathbf{s}})^{\mathrm{T}}}{N-1}$$
(5c)

Conditioning of predictions by historical observations of system state (see below) often benefits from
160 transforming individual elements o_i and s_i of $\mathbf{o}$ and $\mathbf{s}$ into normally distributed variables o^g_i and s^g_i before calculating the covariance matrices that are depicted in equation (4). Unfortunately, Gaussian transformation of individual elements of $\mathbf{o}$ and $\mathbf{s}$ does not guarantee that the joint probability distribution of $[\mathbf{o}^{\mathrm{gT}}, \mathbf{s}^{\mathrm{gT}}]$ is multi-Gaussian. However, this can be at least partially accommodated by undertaking an ensemble-based conditioning process in which $\mathbf{o}^g$ and $\mathbf{s}^g$ are linked by a surrogate model. This is now discussed. Meanwhile, to make the
165 following equations a little less complex, we drop the "g" superscript.

2.2 The DSI Surrogate Model

Sun and Durlofsky (2017) and Lima et al. (2020) demonstrate use of a covariance-matrix-derived surrogate
170 model that can be used to link $\mathbf{o}$ and $\mathbf{s}$. Conditioning of $\mathbf{s}$ by $\mathbf{h}$ (the measured counterpart of $\mathbf{o}$) is then achieved by adjusting the parameters of this model using standard parameter conditioning methodologies. The former authors use the randomised maximum likelihood method, while the latter authors use an ensemble smoother.

The DSI prediction model is described by the following equation:
175
$$\begin{bmatrix}\mathbf{o}\\ \mathbf{s}\end{bmatrix} = \begin{bmatrix}\mathbf{o}_{\mathrm{p}}\\ \mathbf{s}_{\mathrm{p}}\end{bmatrix} + C^{\frac{1}{2}}\left(\begin{bmatrix}\mathbf{o}\\ \mathbf{s}\end{bmatrix}\right)\mathbf{x}$$
(6)

"Adjustable parameters" of this model comprise the vector $\mathbf{x}$. Meanwhile $\mathbf{o}_{\mathrm{p}}$ and $\mathbf{s}_{\mathrm{p}}$ are prior mean values of $\mathbf{o}$ and $\mathbf{s}$. The prior probability distribution of elements x_i of $\mathbf{x}$ is such that they are independently normal with a standard deviation of 1.0. Thus:

180
$$E(\mathbf{x}) = 0$$
(7a)
and

$$C(\mathbf{x}) = \mathbf{I}$$
(7b)

185 where $E(\mathbf{x})$ denotes expected value and $C(\mathbf{x})$ denotes the prior covariance matrix of $\mathbf{x}$. The square root of the covariance matrix that is depicted in equation (6) is obtained using singular value decomposition. First:

$$C\left(\begin{bmatrix}\mathbf{o}\\ \mathbf{s}\end{bmatrix}\right) = \mathbf{U}\Sigma\mathbf{U}^{\mathrm{T}}$$
(8)

190 where $\mathbf{U}$ is an orthonormal matrix and $\mathbf{\Sigma}$ is a diagonal matrix whose elements are positive singular values. The square root of this matrix is then obtained as:

$$C^{\frac{1}{2}}\left(\begin{bmatrix}\mathbf{o}\\\mathbf{s}\end{bmatrix}\right) = \mathbf{U}\mathbf{\Sigma}^{\frac{1}{2}}\mathbf{U}^{\mathrm{T}} \tag{9}$$

195 Normally $\mathbf{\Sigma}^{\frac{1}{2}}$ is truncated so that unduly low singular values are excluded. The sum of squared singular values (i.e. the diagonal elements of $\mathbf{\Sigma}^{\frac{1}{2}}$) is a measure of the total variability of $\begin{bmatrix}\mathbf{o}\\\mathbf{s}\end{bmatrix}$. Normally truncation of $\mathbf{\Sigma}^{\frac{1}{2}}$ is such that 99% or more of this variability is retained. Once singular value truncation has been accomplished, a linearised inverse problem can be formulated using the following equation. (Without loss of generality, we omit the vector of prior means $\begin{bmatrix}\mathbf{o}_{\mathrm{p}}\\\mathbf{s}_{\mathrm{p}}\end{bmatrix}$ in this and the following equations in order to simplify them.)

200

$$\begin{bmatrix}\mathbf{h}\\\mathbf{s}\end{bmatrix} = C^{\frac{1}{2}}\left(\begin{bmatrix}\mathbf{o}\\\mathbf{s}\end{bmatrix}\right)\mathbf{x} + \begin{bmatrix}\mathbf{\varepsilon}\\\mathbf{0}\end{bmatrix} = \mathbf{M}\mathbf{x} + \begin{bmatrix}\mathbf{\varepsilon}\\\mathbf{0}\end{bmatrix} \tag{10}$$

The matrix operator $\mathbf{M}$ is defined through the above equation. Let the matrix $\mathbf{N}$ be an appropriate "selection matrix" comprised of 1's and 0's such that:

205

$$\mathbf{h} = \mathbf{N}\begin{bmatrix}\mathbf{h}\\\mathbf{s}\end{bmatrix} \tag{11}$$

Then:

$$\mathbf{h} = \mathbf{L}\mathbf{x} + \mathbf{\varepsilon} \tag{12}$$

where

210

$$\mathbf{L} = \mathbf{M}C^{\frac{1}{2}}\left(\begin{bmatrix}\mathbf{o}\\\mathbf{s}\end{bmatrix}\right) \tag{13}$$

In equations (10-12), the vector $\mathbf{h}$ contains measurements of historical system behaviour that correspond to model outputs $\mathbf{o}$ (same as in equation 1). We refer to the model $\mathbf{M}$ herein as the "DSI surrogate model" that links elements of observed system behaviour to elements of its predicted behaviour (normally those that have decision-relevance). As stated above, the former are encapsulated in the vector $\mathbf{o}$ while the latter are

215 encapsulated in the vector $\mathbf{s}$. This surrogate model can be "calibrated" against field measurements of system state encapsulated in $\mathbf{h}$ to obtain a maximum a posteriori (MAP) estimate $\underline{\mathbf{x}}$ of $\mathbf{x}$, from which a MAP estimate $\underline{\mathbf{s}}$ of $\mathbf{s}$ can be readily obtained using equation (6). Alternatively, or in addition, samples of the posterior distribution of $\mathbf{x}$ can be obtained using standard Bayesian methods. In the following analysis we use an ensemble smoother as

220 well as a linearised form of Bayes equation for prediction conditioning. Because the computations embodied in equation (10) are simple, the DSI surrogate model runs extremely fast.

2.3 Imposition of History-Match Constraints

225 In the examples that follow, we apply equations (10) and (12) in a number of different ways. These will now be briefly described.

The first of these applications is similar to that described by Lima et al. (2020). That is, an ensemble smoother is used to directly sample the posterior probability distribution of $\mathbf{x}$. The posterior probability distribution of $\mathbf{s}$ is

230 then sampled by running the DSI surrogate model using these samples of **x**. We use the PESTPP-IES ensemble smoother described by PEST++ development team (2022) for this purpose, invoking its automatic adaptive localisation scheme.

We then use Tikhonov regularised inversion to "calibrate" the surrogate model, thereby obtaining a MAP
235 estimate **x** of **x** from which a MAP estimate **s** of **s** is obtained from a single forward run of the calibrated DSI surrogate model. Calibration is achieved using PEST (Doherty, 2022). Because of equation (6) and the fact that the prior expected value of **x** is **0**, the Tikhonov-regularised solution to the inverse problem described by equation (12) is (Doherty, 2015):

240 $$\underline{\mathbf{x}} = (\mathbf{L^t QL} + \beta^2 \mathbf{I})^{-1} \mathbf{L^t Qh} \qquad (14)$$

In equation (14), **Q** is the observation weight matrix; ideally:

$$\mathbf{Q} = C^{-1}(\boldsymbol{\varepsilon}) \qquad (15)$$
245

In practice, equation (14) is solved iteratively as a value for β is sought that guarantees a model-to-measurement least squares objective function that is equal to, or greater than, a user-specified value. This value is calculated from the statistics of measurement noise (de Groot-Hedlin and Constable, 1991; Doherty, 2003).

250 As explained below, we also use a constrained optimisation process to determine the maximum and minimum values that an individual prediction (i.e. an individual element of **s**) can take subject to the constraint that a least squares objective function does not rise above a user-specified value. This objective function is comprised of appropriately-weighted model-to-field-measurement residuals, as well as departures of elements of **x** from their prior expected values of **0**.
255

Because the prior expected value of **x** is **0**, we can write:

$$\mathbf{0} = \mathbf{x} + \boldsymbol{\eta} \qquad (16)$$

260 where $\boldsymbol{\eta}$ is a random vector with the same statistical properties as **x**. Equation (12) can then be written as:

$$\begin{bmatrix} \mathbf{h} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{L} \\ \mathbf{I} \end{bmatrix} \mathbf{x} + \begin{bmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\eta} \end{bmatrix} \qquad (17a)$$

With obvious definitions for **d**, **J** and $\boldsymbol{\tau}$, equation (17a) can be re-written as:

265 $$\mathbf{d} = \mathbf{Jx} + \boldsymbol{\tau} \qquad (17b)$$

From (7b), the covariance matrix of $\boldsymbol{\tau}$ is:

$$C(\boldsymbol{\tau}) = C\left(\begin{bmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\eta} \end{bmatrix}\right) = \begin{bmatrix} C(\boldsymbol{\varepsilon}) & 0 \\ 0 & \mathbf{I} \end{bmatrix} \qquad (18)$$
270

Let the vector **y** denote the sensitivity of an element s_i of **s** to surrogate model parameters **x**. For the purposes of prediction uncertainty analysis, we define a least squares objective function φ_u as the sum of squared, weighted model-to-measurement misfit residuals and parameter misfit residuals from their preferred values of **0**. Thus:

$$\Phi_u = (\mathbf{h} - \mathbf{o})^\mathrm{T}\mathbf{Q}(\mathbf{h} - \mathbf{o}) + \mathbf{x}^\mathrm{T}\mathbf{x} \tag{19}$$

Suppose that values of the objective function φ_u that exceed a user-specified value of φ_{u0} are deemed to be unlikely (at a certain level of confidence). Vecchia and Cooley (1987) show that solution of the constrained optimisation problem in which s_i is maximized/minimized subject to the constraint that $\varphi_u \leq \varphi_{u0}$ can be obtained through iterative solution of the following equation:

$$\mathbf{x} = \left(\mathbf{J}^\mathrm{T}\mathbf{QJ}\right)^{-1} \left\{ \mathbf{J}^\mathrm{T}\mathbf{Qh} - \frac{y}{2\lambda} \right\} \tag{20a}$$

where:

$$\left(\frac{1}{2\lambda}\right)^2 = \pm \frac{\Phi_{u0} - \mathbf{d}^\mathrm{T}\mathbf{Rd} + \mathbf{d}^\mathrm{T}\mathbf{RJ}\left(\mathbf{J}^\mathrm{T}\mathbf{RJ}\right)^{-1}\mathbf{J}^\mathrm{T}\mathbf{Rd}}{y^t(\mathbf{J}^\mathrm{T}\mathbf{RJ})^{-1}y} \tag{20b}$$

and:

$$\mathbf{R} = \mathrm{C}^{-1}(\boldsymbol{\tau}) \tag{20c}$$

Recall that $\mathbf{d}$ is defined by equation (17b). Maximisation or minimisation of s_i is chosen according to the sign used in equation (20b). In the following example, this process is carried out using the "predictive analysis" functionality of PEST (Doherty, 2022).

Finally, in the examples that we present below, we apply linear uncertainty analysis to evaluate the worth of various subsets of field data (i.e. elements of $\mathbf{h}$). If we assume a linear relationship between surrogate model outputs (under both history-matching and predictive conditions) and surrogate model parameters, and Gaussian probability distributions for both $\mathbf{x}$ and $\boldsymbol{\varepsilon}$, the posterior uncertainty of a prediction s_i can be calculated from the prior covariances matrix C($\mathbf{x}$) of model parameters on the one hand, and the covariance matrix C($\boldsymbol{\varepsilon}$) of measurement noise on the other hand using the following Bayes-derived equation (Doherty, 2015; Fienen et al., 2010):

$$\sigma^2{}_{si} = \mathbf{y}^\mathrm{T}\mathrm{C}(\mathbf{x})\mathbf{y} - \mathbf{y}^\mathrm{T}\mathrm{C}(\mathbf{x})\mathbf{L}^\mathrm{T}[\mathbf{L}\mathrm{C}(\mathbf{x})\mathbf{L}^\mathrm{T} + \mathrm{C}(\boldsymbol{\varepsilon})]^{-1}\mathbf{L}\mathrm{C}(\mathbf{x})\mathbf{y} \tag{21}$$

The first term on the right side of equation (21) is the prior uncertainty of the prediction s_i. The second term represents predictive uncertainty reduction accrued through history-matching. In implementation of equation (21), the sensitivities embodied in the vector $\mathbf{L}$ can be calculated using finite perturbations of the MAP estimate $\underline{\mathbf{x}}$ of $\mathbf{x}$. Meanwhile C($\mathbf{x}$) is the identity matrix $\mathbf{I}$.

It is important to note that equation (21) includes the values of neither parameters nor observations; it features only sensitivities. Hence, as will be discussed below, it can be easily turned to the task of data worth evaluation. In particular, the ability of a new measurement to reduce the uncertainty of a prediction of interest can be evaluated without actually knowing the value of that measurement.

3. Application

We now discuss application of the DSI methodology in a synthetic alluvial stream-aquifer context. Alluvial river corridors are used worldwide for drinking water supply. Up to 85% of groundwater withdrawals from these

systems come from surface water capture (Scanlon et al., 2023). They provide productive aquifer systems and natural riparian filtration; see, for example, Epting et al. (2022). River-associated aquifers have been formed by millennia of channel meander migration, aggregation and erosional flow processes. These processes produce
320   heterogeneous aquifers in which large-scale features, such as palaeochannels, provide preferential flow paths. These can significantly influence the spatial distribution of exchange fluxes between streams and aquifers. Interactions between surface and subsurface waters are therefore complex. Furthermore, stochastic characterisation of subsurface heterogeneity is strongly non-Gaussian. This makes it difficult to evaluate the statistical properties of management-pertinent model predictions.

325

In our synthetic example, the objective of numerical modelling mimics that of real-world modelling in many alluvial depositional environments with which we are familiar. It is to predict surface water infiltration along a stream, and surface water travel times to water production wells with particular emphasis on the shortest of these times. Both of these variables are salient to management of an alluvial aquifer that is used for drinking water
330   supply. Provision of these predictions by a numerical model can assist water managers in operating a well field in a way that minimizes the potential for bacterial contamination of drinking water (Epting et al., 2018). In this example, we investigate the ability of DSI to associate uncertainties with management-pertinent predictions, and to support data acquisition-strategies which reduce these uncertainties.

335   **3.1 Numerical Model**

Complex interactions and feedback mechanisms between surface and subsurface waters often need to be considered when managing alluvial aquifers, and therefore when simulating them for management purposes. Integrated surface and subsurface hydrologic models (ISSHMs) provide a consistent framework for simulating
340   flow and transport processes that are operative within these aquifers (Brunner et al., 2017). A strength of these models is that they can simulate dynamic feedback between the surface and subsurface water regimes over a wide range of temporal and spatial scales (Paniconi and Putti 2015).

For this reason, and because it is frequently used by the authors in real-world decision-support modelling of
345   alluvial aquifers, the ISSHM modelling platform HydroGeoSphere (HGS) (Aquanty Inc., 2022; Brunner and Simmons 2012; Brunner et al., 2012) is used to perform the numerical simulations documented herein. HGS simulates surface water (SW) and groundwater (GW) flow processes using a globally implicit, finite-difference flow formulation. Simulation of water movement within the 2D surface water domain is based on the diffusion-wave equation, while the three-dimensional, variability-saturated Richards equation is used to simulate
350   subsurface flow processes.

3.2 Model Setup

3.2.1 Geometry

355

The numerical flow model that we deploy to explore and document the capabilities of DSI has a spatial extent of $300 \times 500 \times 30$ m in the x, y and z directions (Figure 1a). An alluvial plain slopes gently towards the stream outlet with a slope of 0.003 m/m in the y-direction. The stream itself flows along the eastern boundary of the model domain in a 3 m deep and 15 m wide channel. The numerical grid is discretised using a 2D unstructured

360 triangular mesh with lateral internodal spacing that ranges between 4 m along the stream and 7 m under the proximal alluvial plain. The 2D mesh was generated using AlgoMesh (HydroAlgorithmics Pty Ltd, 2022). The model has 14 layers. Layer boundaries are at depths of 0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0, 8.0, 11.0, 14.0, 17.0, 20.0, 23.0 and 26.0 metres. The bottom of the model is at a depth of 30 metres. The resulting 3D mesh consists of 112,240 nodes and 204,000 elements.

365

3.2.2 Boundary conditions

The surface water inflow on the upstream side of the stream was conceptualised as a second type (specified flux) boundary corresponding to a constant flow (Q) of 2 $m^3.s^{-1}$. The surface water outflow was implemented as a so-
370 called "critical depth boundary", which allows surface water to leave the overland flow area downstream (Figure 1a). Meanwhile, an areal infiltration rate of 300 $mm.y^{-1}$ is imposed over all nodes of the alluvial plain that borders the river. This top boundary condition represents effective rainfall, as evapotranspiration processes are not explicitly simulated by the model. At the northern and southern ends of the groundwater system, a fluid transfer boundary condition is deployed to ensure a regional groundwater in- and outflow at these upstream and
375 downstream model boundaries respectively. Lateral model boundaries, as well as the basal model boundary, are impermeable.

Two water production wells are represented by flux nodal boundary conditions; each well extracts 400 $m^3.h^{-1}$ from the groundwater system. The well field, therefore, produces a total of 800 $m^3.h^{-1}$ of water. In each well the
380 pump is emplaced at an elevation of 85 m (i.e. at a depth of approximately 14 m). As in all ISSHMs, streamflow simulation is explicit, while the simulation process passes water between the surface and groundwater systems in accordance with heads that are calculated in these two flow domains. This eliminates the need to specify special boundary conditions at the surface water groundwater interface. The surface water domain is directly coupled to the groundwater domain using the dual-node approach with a coupling length of 0.1 m.

385

The transient simulation covers a period of 95 days. Over this period, extraction rates from the two water production wells are varied in order to reproduce a controlled pumping experiment (Figure 1c). During this experiment, well field extraction is shut down for 50 days. The pumping rate is then restored, and simulation continues for another 30 days. This pumping rate reduction experiment introduces a significant transient event to
390 the alluvial system. The response of the system to this event has the potential to provide a significant amount of information on system hydraulic properties. It is worthy of note that a field-scale experiment that is similar to this was conducted in the Aeschau plain, Switzerland, under the supervision of the city of Bern (Schilling et al., 2017).

395 Prior to the 95 day reducing pumping experiment, the model runs for a period of 365 days. This allows attainment of near-steady-state conditions. The average computational time required for the model to simulate the preceding 365 days and the full 95 days of the pumping rate reduction experiment is 10 minutes; simulation is not parallelized. HGS employs adaptive time-stepping; however the maximum time step was limited to a single day in order to ensure computational stability.

400

3.2.3 Parameterisation

The generation of realistic distributions of subsurface hydraulic properties in depositional environments that are characterised by distinct, continuous features of complex geometry such as alluvial channels, is often

405   implemented using rule-based feature-generation codes. Alternatively, packages which implement multipoint geostatistics (Remy et al., 2009; Linde et al., 2015) can be employed. Algorithms that underpin the former codes generate geobodies that can have similar geometric properties to alluvial channels, whereas latter packages employ stochastic image analysis techniques to support reproduction of these complex structures while maintaining their long-distance connectivity.

410

In the present case, we employ the ALLUVSIM alluvial channel simulator (Pyrcz et al., 2009) to generate realisations of the alluvial subsurface. (Note that these do not affect the flow path of the river; conceptually they are remnants of historical river channels). The numerical generation of superimposed and intersecting alluvial channels is controlled by geometric input parameters such as channel depth, width, porosity, starting location

415   and sinuosity (Pyrcs et al., 2009); see Table 1. ALLUVSIM simulations were used to assign sets of alluvial channels to the HGS groundwater flow domain.

For the generation of ALLUVSIM channel sets we adopt a width-to-depth ratio of 1. Furthermore, the depth of channel deposits averages 20m, this being a considerable portion of the modelled subsurface. Recall that the

420   depth of the base of the model is 30m. We do not stack channels at different depths. Therefore each realization of the subsurface contains a small number (up to 10) alluvial channels, each of which is continuous from the north to the south of the model domain while extending from the surface to a depth that approaches the bottom of the alluvium. As is apparent in Table 1, the mean channel x coordinate (about 370m from the western edge of the ALLUVSIM model boundary along its northern boundary) is such that channels tend to occupy parts of the

425   subsurface that are not too far from the course of the present river. Groundwater flow within the near-river part of the model domain is therefore dominated by the presence of these continuous, conductive channels. One of the ALLUVSIM-generated channel distributions (out of a total of 101) was adopted as the "reality" channel distribution. This is displayed in Figure 1b. The other 100 realisations were used to generate realisations of prior hydraulic properties. These are required by the DSI process that is described above.

430

Each realisation of the subsurface contains three facies; each is based on a different realisation of an ALLUVSIM-generated channel set. These facies are channel, non-channel and riverbed deposits. This last facies is mainly present near the banks, while it is absent in the deepest parts in the middle of the stream. In this model we assume that the riverbed deposits have developed in a uniform layer of 50 centimetres, which is therefore

435   assigned to the uppermost model layer. For each HGS model realisation, each of these facies is provided with a spatially uniform set of hydraulic properties (i.e. hydraulic conductivity and porosity). The value of each of these properties is randomly selected from probability distributions that are presented in Table 1. (Intra-facies heterogeneity is purposefully neglected in order to examine DSI's performance in geological settings where hydraulic properties are categorical in nature, with one category exerting a dominant influence on groundwater

440   flow.) Appendix A depicts a number of HGS model parameter realisations, coloured according to horizontal hydraulic conductivity.

A vertical anisotropy of 4 is assigned to all hydraulic conductivities in all realisations of all facies. This is consistent with alluvial depositional systems similar to those that our study attempts to represent. (Gianni et al., 2018; Chen 2000; Ghysels et al., 2018).

Parameters that govern unsaturated flow were also homogeneous, and invariant between realisations. The van Genuchten-Mualem parameters a [m^{-1}] and β [-] are set to 3.48 and 1.75, respectively, these being typical of alluvial gravel (Dann et al., 2009). The residual soil water content is set to 0.095. The Manning's roughness coefficient is set to 1.7×10^{-6} d.m$^{-1/3}$.

3.3 History-matching dataset and model predictions

The dataset used for history matching consists of 95 observations of hydraulic heads in each of eight observation wells, these providing a total of 760 individual observations. These heads are calculated using the HGS "reality parameter field". A random realisation of measurement noise is added to each head measurement. The probability density distribution of synthetic head measurement noise has a mean of 0.0m and a standard deviation of 0.05m. Figure 2 shows "reality heads" together with heads calculated at the same locations using the remaining 100 hydraulic property realisations. There is considerable scatter in these plots. This suggests that the information content of these heads with respect to subsurface hydraulic properties is high. Their information content with respect to predictions of management interest is pursued using the methodologies discussed above.

The predictions of interest are as follows:
- fast travel times (days) at the first well (well #1);
- fast travel times (days) at the second well (well #2);
- rate of surface water infiltration into the groundwater system (l.s^{-1}) along the entire length of the stream.

Travel times are calculated using particles; see Anderson et al. (2015) for details. For each realisation of the alluvial architecture and accompanying hydraulic properties, particles are placed along the riverbed at the top of the subsurface flow domain; see Figure 1b. Most particles leave the model domain through an extraction well. The travel time that corresponds to the 5th percentile of the particle breakthrough curve at a particular well under normal operating conditions is deemed to be the "fast travel time" pertaining to that well. Our focus on fast-moving water acknowledges the likelihood of rapid flow processes being driven by preferential flow paths. These same rapid flow processes can threaten water extraction in the event of surface water contamination.

Fast times and surface water infiltration are calculated at t = 365 days. Thus they pertain to maximum well abstraction rates.

For the synthetic reality, the 5th percentile of travel times is 9.57 and 7.90 days for water production wells #1 and #2, respectively. Meanwhile, surface water infiltration along the entire stream is 130.67 l.s^{-1}.

4. Uncertainty Quantification

485 In this section we document how the posterior (i.e. post-history-matching) uncertainties of the three predictions of interest can be evaluated using four different approaches. One of these requires adjustment of the parameters of the HGS model. The other three approaches require adjustment of parameters of a surrogate model that is built to implement DSI-based predictive uncertainty evaluation in ways that are discussed in Section 2. In all cases parameter adjustment minimizes a least squares objective function that serves as a measure of model-to-

490 measurement misfit. This objective function is calculated as the sum of weighted squared differences between observed and modelled heads in the eight observation wells discussed above. Weights are uniform to reflect the temporal uniformity of measurement noise; each has a value that is equal to the square of the standard deviation of this noise. The expected value of the history-matched objective function is therefore expected to be somewhat greater than 760, this being the number of observations which comprise the calibration dataset. The "somewhat

495 greater" is an outcome of the fact that number of "effective parameters" is unknown prior to solving an ill-posed inverse problem. See Doherty (2015) for details.

4.1 Bayesian history-matching

500 **4.1.1 History matching with the HGS model**

First, we adjust HGS model parameters using the PESTPP-IES ensemble smoother. History-match-constrained parameter fields are then used to make the predictions of interest. The variability of these predictions between these posterior parameter fields is a measure of their posterior uncertainties.

505

The ensemble smoothing process begins with samples of the prior parameter probability distribution. In the present case we use the 100 parameter fields that were obtained in the manner discussed above. These parameter fields are iteratively adjusted until model outputs fit field measurements. The HGS model has 204,000 elements. As is common practice when using parameter ensembles for history-matching, each of these is considered as a

510 separate parameter when undertaking PESTPP-IES-based parameter adjustment.

Piecewise spatial uniformity of the initial parameter fields is lost during the ensemble parameter adjustment process as each parameter is subject to individual adjustment while maintaining a high level of spatial correlation with neighbouring parameters that is inherited from the initial realisations. See Chen and Oliver

515 (2013) for mathematical details of the parameter adjustment process.

The objective function associated with all ensemble realisations was significantly reduced after 3 iterations of the IES parameter adjustment process. Objective function values ranged between 1075 and 16098, except for 30 realisations which suffered excessively slow HGS model solution convergence and where therefore abandoned.

520 Note that each iteration of the IES parameter adjustment process requires as many model runs – in this case HGS model runs - as there are realisations that comprise the ensemble. Figure 2 shows that model-calculated heads at the observation wells are indeed close to "measured" heads. The 70 realisations which remained after 3 IES iterations were used to make the predictions that are described above. The results are plotted in Figure 3a.

525 It can be seen from Figure 3a that the uncertainty of predicted surface water infiltration is significantly reduced by PESTPP-IES based history matching of the HGS model. It can therefore be concluded that the information content of hydraulic heads with respect to this specific prediction is high. In contrast, the uncertainties of first arrival travel time predictions are not significantly reduced; see Figure 3a. The information content of head responses to altered pumping rates with respect to these predictions is therefore lower that it is for surface water

530 infiltration.

An interesting feature of Figure 3a is that some fast travel times calculated for posterior parameter fields exceed those calculated using prior parameter fields. This can be explained by the fact that some posterior parameter fields have lost some of the connectivity exhibited by the prior parameter fields. (See Appendix B). This is an

535 outcome of the PESTPP-IES parameter fitting process. Its parameter adjustment process is only truly Bayesian where prior parameter distributions are Gaussian on a cell-by cell basis (if cell-by-cell parameterisation is employed). Neither the theory on which it is based, nor numerical implementation of that theory in its history-matching algorithm, can guarantee maintenance of long-distance hydraulic property connectedness that cannot be characterised by a multiGaussian distribution. Indeed, history-match-constrained adjustment of connected and

540 categorical parameter fields is still an area of active research (Khambhammettu et al., 2020).

4.1.2 History matching with the DSI surrogate model

The same 100 samples of the prior parameter probability distribution that were used to start the PESTPP-IES

545 data assimilation process were then used to construct a DSI surrogate model using the methodology described in Section 2. Recall that this model is based on an empirical covariance matrix that relates history-matched model outputs to predictive model outputs. A singular value energy level of 0.999 was used in the construction of this surrogate model. This results in the use of 87 singular values (and corresponding eigencomponents) of this matrix, and hence an 87-dimensional **x** vector; see equation (6). Note that the only HGS model runs that were

550 required to construct and history-match the surrogate DSI model were those that were based on the 100 prior hydraulic property field realisations discussed above. No further runs of the HGS model were required to carry out any of the procedures described below.

Sampling of the posterior distribution of surrogate model parameters (i.e. elements of the **x** vector) was

555 implemented using the PESTPP-IES ensemble smoother. However, in contrast to the previous use of PESTPP-IES, its use with the DSI surrogate model incurs a trivial numerical burden because of the extremely fast execution speed of this model. An ensemble comprising 100 realisations of the prior probability distribution of surrogate model parameters was adjusted. After 6 iterations of this process, all realisations yielded good model-to-measurement fits, with objective function values ranging between 814 and 830. Prior and posterior prediction

560 histograms are shown in Figure 3b.

The posterior prediction probability distributions for surface water infiltration calculated by HGS parameter fitting and by DSI surrogate model parameter fitting are very similar. The same cannot be said for fast travel times. The posterior uncertainty of this prediction is lower after DSI surrogate model parameter fitting than after

565 HGS model parameter fitting. This suggests that fitting of DSI surrogate model parameters does not suffer the

same degradation in uncertainty evaluation performance as that which is incurred by PESTPP-IES-based fitting of HGS parameters which tend to loose their continuity as the history-matching process progresses. In contrast, parameters of the surrogate DSI model are not required to maintain any spatial patterns or relationships; they must simply represent observation-to-prediction relationships that are embodied in 100 HGS model outputs that
570 were all calculated using continuous hydraulic property parameter fields.

4.2 Tikhonov-regularised inversion and constrained predictive maximization/minimization

The DSI surrogate model was next calibrated using Tikhonov regularisation; see equation (14). The MAP
575 estimates of **x** that are obtained in this way can then be used to calculate MAP values of predictions. These are listed in Appendix C.

Following calibration of the DSI surrogate model, equation (21) was used to calculate the prior and posterior standard deviations of the uncertainties of the three predictions of interest based on an assumption of surrogate
580 model linearity. Sensitivities of model outputs that correspond to members of the calibration dataset, and that correspond to model predictions, were calculated using finite difference perturbations from calibrated surrogate model parameter values. Linear-calculated standard deviations are then used to plot the probability density distributions for the three predictions to be compared with ensemble-calculated standard deviations (Figure 3c).

585 Next, posterior uncertainties of the three predictions of interest were calculated using the constrained maximisation/minimisation procedure that is described by equation (20). The objective function constraint provided to this equation was that which enables calculation of Scheffe 95% confidence intervals (see Vecchia and Cooley, 1987). The values obtained through this process agree reasonably well with maximum and minimum predictions obtained through Bayesian history matching of the DSI surrogate model (Figure 3c).
590

Differences between evaluated predictive uncertainties that appear in Figure 3 (statistics are available in Appendix C) are generally small; however some are large enough to warrant discussion. Differences between linear and nonlinear uncertainty estimates can be at least partly attributed to approximations that calculation of these estimates require. Linear estimation of posterior predictive uncertainty requires not only an assumption of
595 linearity of the DSI model; it also assumes Gaussianality of the prior probability distributions of DSI model parameters (i.e. elements of the **x** vector of equation 6). Meanwhile, nonlinear analysis using ensembles is prone to sampling errors incurred through the use of only 100 realisations; it is also affected by problems that PESTPP-IES has in maintaining continuity of hydraulic properties (see above). Another factor that may degrade the accuracy of ensemble-calculated posterior uncertainties is failure of model outputs based on some parameter
600 realisations to fit measured heads to within limits that are set by the statistical properties of measurement noise. Nevertheless, in spite of these approximations, it is pleasing to note that the ratio of prior to posterior uncertainty is reasonably consistent between linear and nonlinear estimates of posterior predictive uncertainty. This increases confidence in the data worth analysis which is the subject of the next section.

605 **5. Data worth**

Quantifying the effectiveness and efficiency of existing and planned data acquisition and monitoring strategies can (and often should) be an important outcome of decision-support groundwater modelling. This is because quantification of uncertainty brings with it an ancillary benefit, this being quantification of the extent to which existing or contemplated data can reduce the uncertainties of one or more decision-critical model predictions. The worth of data increase in proportion to their ability to achieve this outcome.

5.1 Value of existing data

The worth of data that already comprise the calibration dataset can be evaluated according to two different metrics. The value of subsets of existing data can be assessed by evaluating the extent to which the posterior standard deviations of decision-critical model predictions are increased by their omission from this dataset. The worth of these subsets can also be assessed by evaluating the extent to which the prior standard deviation of these predictions is reduced by including a particular subset as the only members of the history-matching dataset. In the present section, we evaluate the worth of existing data according to the first of these metrics only. This is implemented using linear analysis based on equation (21). It is worth noting that, because the surrogate model runs so fast, the ensemble smoother could also have been used to sample the posterior distribution of our predictions with the omission of certain subsets of the history-matching dataset.

Using equation (21), data is notionally omitted in a calibration dataset simply by setting their associated weights to values of zero. This was done for each of the observation wells that are featured in Figure 1a. The outcomes of this analysis are presented in Figure 4.

Inspection of Figure 4 reveals that members of the existing observation network are more informative of predictive flux than they are of fast travel times. Furthermore, the cost of omitting OBS110 from the existing observation network appears to be greater than that of omitting any other observation well from the network. This implies a certain degree of uniqueness of information that is forthcoming from this well. This is not surprising as OBS110 is located upstream and closer to sources of surface water infiltration than most of the other observation wells. It is therefore close to the path of much of the water that flows from the river to production wells.

Of particular interest is the high information content of OBS17 data with respect to the prediction of induced infiltration. This is attributable to the fact that prior realisations of hydraulic properties display spatially uniformity within each stratigraphic unit (that is channels, riverbed and other alluvium). Therefore, while pumping-induced surface water infiltration is sensitive to the structure and disposition of paleochannels which convey water from the river to production wells, it is also sensitive to aquifer properties that can control inflow/outflow of water to/from the northern and southern boundaries of the model domain. The difference in head between OBS17 and the southern fixed-head boundary is informative of these properties. The importance of these head measurements would be reduced if stratigraphic unit properties were internally heterogeneous.

5.2 Value of a new observation point

Equation (21) can be readily turned to evaluation of the issue of whether it is worth supplementing the existing observation network of eight observation wells with a new observation well. In our study, contenders for the new
650 observation point are each of the 146 sites that are represented by black crosses in Figure 1a. We establish the worth of including a new point in the monitoring network by including heads at that point in an expanded history-matching dataset, thereby assuming that observations from this well were available during the calibration period that has formed the basis for all investigations that have been discussed so far. The reduction in the uncertainty standard deviation of management-pertinent predictions that is accrued through expansion of the
655 history-matching dataset in this way is a measure of the worth of the hypothesized new data. Recall from the discussion in Section 2 of this paper, that equation (21) does not require the values of observation to assess their worth. It requires only the sensitivities of corresponding model outputs to model parameters.

We base our exploration of data worth on a DSI surrogate model. However, HGS model outputs on which this
660 model is based must be extended to include those at the candidate observation wells. This enables the construction of an expanded covariance matrix that links model outputs at these sites to predictions of interest. In our case, these were available from archived output files of the 100 HGS model runs which form the basis for results that are reported in previous sections of this paper. In other cases another suite of large model runs may be required for the construction of a new surrogate model. Note, however, that once these model runs have been
665 completed, all further analyses are conducted using the DSI surrogate model.

To verify this linear approach to data worth assessment, reductions in prediction uncertainty that would result from updating the calibration dataset to include heads from each of the 146 potential observation wells that are shown in Figure 1a were calculated using an alternative (and much more laborious) nonlinear approach.
670

For each of the 100 prior realisations of HGS parameter fields, heads were calculated at each of the 146 trial observation points. For each of these observation points, for each of these parameter realizations, a new DSI model was constructed using the methodology discussed above. The history-matching dataset for each of these models included the existing dataset (comprised of eight observation wells), as well as the expanded dataset that
675 is pertinent to that well and that particular realisation of the prior HGS parameter field. Predictions were the same for each DSI surrogate model (i.e. fast travel times to the production wells and surface water infiltration). Each of these 14,600 DSI surrogate models was then history-matched against its respective history-matching dataset to yield the posterior uncertainties of our predictions of interest. This was done using the PESTPP-IES ensemble smoother. For each trial observation well, the "total" posterior uncertainty of a prediction of interest
680 was calculated by accumulating prediction realisations over all 100 DSI surrogate models that pertained to each expanded observation well. These were compared with the posterior uncertainties calculated using the original DSI surrogate model that was based on a history-matching dataset that did not include the new well.

For both of these approaches, the percent reduction in the standard deviations of the uncertainties of our
685 predictions was calculated for each of the 154 observation wells. These were then interpolated over the model grid for presentation purposes; see Figure 5.

Areas of high and low worth of installation of a new observation well are broadly similar between the upper and lower sets of maps that comprise Figure 5, especially for the surface water infiltration prediction.

690

It is clear from Figure 5 that the collection of new hydraulic head data in the northern part of the model has the potential to reduce the uncertainty associated with predictions of surface water infiltration. This also applies to fast travel times to the northern extraction well, and to a lesser extent for fast travel times to southern extraction well. The nonlinear analysis also suggests that acquisition of head data in the river corridor between the existing

695   OBS110 and OBS83 observation wells may yield reductions in uncertainties of fast travel time predictions for both water production wells. To a lesser extent (especially for the southern production well), this is also found in the linear analysis. This makes sense, as this zone is crossed by most of the particles that arrive at the production wells.

700   Both of the methods for data worth analysis that are described above are based on approximations. Nonlinear analysis suffers from the limited number of realisations on which it is based. Obviously, linear analysis suffers from an assumption of DSI model linearity. Nevertheless, despite these approximations, the two approaches yield results that are in broad agreement with each other. Furthermore, they are both readily deployable in real-world contexts where hydraulic property distributions and processes are complex; this applies especially to DSI-

705   based linear analysis.

6. Discussion and Conclusions

The purpose of this paper is to document and demonstrate the use of data space inversion (DSI) as a means of

710   quantifying and reducing the posterior uncertainties of predictions made by complex models with complex parameter fields. Both of these model attributes make application of traditional uncertainty analysis difficult. Model complexity increases model run time; it can also increase the propensity of a model to exhibit unstable numerical behaviour when endowed with a stochastic parameter field. The complexity of parameterisation, especially when it involves the use of continuous, connected hydraulic properties, can violate the assumptions

715   upon which these properties are adjusted in order for model outputs to respect measurements of system behaviour.

Data space inversion addresses the first of these challenges by replacing a numerical model with a fast-running surrogate model. This model is tuned to the decision-support role for which the original, complex model was

720   built. This is because the DSI surrogate model is designed to replicate the ability of a numerical model to simulate past measurements of system behaviour and to make predictions of future system behaviour that are of interest to management. Both of these complex model outputs can be calculated using parameter fields of arbitrary complexity that represent aspects of the subsurface that are critical to past and future groundwater behaviour. In many cases, these will include structural or alluvial features that can rapidly transport water and

725   dissolved contaminants to points of environmental impact.

A strength of the DSI methodology is that it does not require adjustment of hydraulic property fields in order for model outputs to replicate the past so that they can be used as a basis for posterior sampling of future

groundwater behaviour. Instead, the parameters and predictions of the surrogate model (rather than the original
730 model) are subjected to adjustment and posterior sampling. The nature of these surrogate model parameters is
such that they embody the parameterisation complexity of the original model without actually replicating it.
Instead, they are formulated to replicate the effects of that complexity on the model outputs used for both history
matching and system management. Thus, the complex nuances of system hydraulic properties are implicitly
taken into account as predictions of future system behaviour are conditioned by measurements of past system
735 behaviour.

This paper demonstrates that the use of a surrogate DSI-based model can extend beyond that of sampling the
posterior probability distribution of one or a number of predictions of management interest. Like conventional
numerical models, it can be used for rapid assessment of the worth of existing or new data. The simplest (and
740 most rapid) form of data worth analysis relies on an assumed linear relationship between surrogate model
parameters and surrogate model outputs. There may be many circumstances where this linearity assumption is
more applicable than that of a linear relationship between conventional model parameters and conventional
model outputs, as the latter relationships are bypassed in definition and construction of the DSI surrogate model
and its parameters. DSI-based evaluation of data worth may therefore be more reliable than evaluation of data
745 worth using a complex model. In fact, if a complex model has many parameters (as it should), and its run times
are high (as they often are), calculation of sensitivities that are required for linear analysis may not be possible.
Meanwhile, by construction, the parameters of a DSI surrogate model can take the complex dispositions and
connectivity relationships of real-world hydraulic properties into account.

750 This paper also demonstrates more complex uses to which a DSI surrogate model can be put. Tikhonov-
regularised inversion and constrained, nonlinear prediction optimisation are demonstrated in Section 4. Section 5
demonstrates how more complex assessments of data worth can be made than that which relies on an assumption
of surrogate model linearity. While the numerical costs associated with these assessments are high, they are far
from prohibitive.
755

However, together with strength comes weakness. It is acknowledged that while the DSI methodology enables
rapid and effective posterior uncertainty analysis in contexts that may otherwise render such analyses
approximate at best and impossible at worst, a modeller is entitled to feel a sense of frustration at not being able
to "see for him/herself" the parameter fields that give rise to predictive extremes. Not only may an understanding
760 of these fields add to a modeller's understanding of a system; they may accomplish the same thing for decision-
makers and stakeholders as well.

Code availability

765   All software and the source code that was used to implement analyses that are described herein can be downloaded from the PEST web site at https://www.pesthomepage.org/programs and the PEST++ web site at https://github.com/usgs/pestpp. Data used for this study are available in a public data repository under the doi: 10.5281/zenodo.7707975.

Author contribution

770

HD and PB designed the synthetic numerical experiment. JD implemented the DSI method in PEST. JD performed the linear analysis using the DSI-based model. HD performed the non-linear analysis using the DSI-based model. HD writes the first draft. JD and PB collaborate with HD on this first draft. PB secured the funding for HD.

775   **Competing interests**

The authors declare that they have no conflict of interest

Acknowledgements

780   The authors acknowledge the funding provided by the Swiss National Science Foundation (SNF, grant number 200021_179017) and the GMDSI project (see https://www.gmdsi.org). The authors would like to thank Jeremy T. White for his helpful feedback on our paper.

References

Anderson MP, Woessner WW and Hunt RJ.: Applied groundwater modeling: simulation of flow and advective transport. Academic, Cambridge, MA, 2015.

Aquanty Inc. HydroGeoSphere Theory Manual.Waterloo, ON. p. 101, 2022.

Brunner, P., Doherty, J., and Simmons, C. T.: Uncertainty assessment and implications for data acquisition in support of integrated hydrologic models. Water Resources Research, 48, W07513, doi:10.1029/2011WR011342, 2012.

Brunner, P., and Simmons, C. T.: HydroGeoSphere: a fully integrated, physically based hydrological model. Groundwater 50, 170–176, doi:10.1111/j.1745-6584.2011.00882.x, 2012.

Brunner, P., Therrien, R., Renard, P., Simmons, C. T., and Hendricks Franssen, H. J. : Advances in understanding river-groundwater interactions. Reviews of Geophysics, 55, 818–854, doi:10.1002/2017RG000556, 2017.

Chen, X. : Measurement of streambed hydraulic conductivity and its anisotropy. Environmental Geology, 39,12, doi:10.1007/s002540000172, 2000.

Chen, Y. and Oliver, D. S.: Levenberg–marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. Computational Geosciences, 17(4):689–703, doi : 10.1007/s10596-013-9351-5, 2013.

Dann, R., Close, M., Flintoft, M., Hector, R., Barlow, H., Thomas, S and Francis, G.: Characterization and estimation of hydraulic properties in an alluvial gravel vadose zone. Vadoze Zone Journal, 8(3):651—663, 2009.

Dausman, A.M., J. Doherty, C.D. Langevin, and M.C. Sukop.: Quantifying data worth toward reducing predictive uncertainty. Groundwater 48, no. 5: 729–740, doi: 10.1111/j.1745-6584.2010.00679.x, 2010.

de Groot-Hedlin, C and Constable, S.: Occam's inversion to generate smooth, two-dimensional models from magnetotelluric data. Geophysics (12):1613-1624, doi : 10.1190/1.1442813, 1991.

Delottier, H., Therrien, R., Young, N. L., and Paradis, D.: A hybrid approach for integrated surface and subsurface hydrologic simulation of baseflow with Iterative Ensemble Smoother. Journal of Hydrology, 606. doi:10.1016/j.jhydrol.2021.12740, 2022.

Doherty, J.: Ground water model calibration using pilot points and regularisation. Groundwater, doi : 10.1111/j.1745-6584.2003.tb02580.x, 2003.

Doherty, J.: Calibration and uncertainty analysis for complex environmental models. Published by Watermark Numerical Computing, Brisbane, Australia. 227pp. ISBN: 978-0-9943786-0-6. Downloadable from www.pesthomepage.org, 2015.

Doherty, J. and Moore, C.: Decision support modeling: data assimilation, uncertainty quantification, and strategic abstraction. Groundwater 58, 327–337. doi. org/10.1111/gwat.12969, 2020.

Doherty, J.: Manual for PEST. Watermark Numerical Computing. Downloadable from http:www.pesthomepage.org, 2022.

Epting, J., Huggenberger, P., Radny, D., Hammes, F., Hollender, J., Page, R.M., Weber, S., Bänninger, D., and Auckenthaler, A.: Spatiotemporal scales of river-groundwater interaction – The role of local interaction

processes and regional groundwater regimes. Sci. Total Environ. 618, 1224–1243.
825   https://doi.org/10.1016/j.scitotenv.2017.09.219, 2018.

Epting, J., Vinnaa, L.R., Piccolroaz, S., Affolter, A. and Scheidler, S.: Impacts of climate change on Swiss alluvial
       aquifers--A quantitative forecast focused on natural and artificial groundwater recharge by surface
       water infiltration. Journal of Hydrology X. 17(100140), doi:10.1016/j.hydroa.2022.100140, 2022.

Fienen, M. N., Doherty, J., Hunt, R. J., and Reeves, J.: Using prediction uncertainty analysis to design hydrologic
830   monitoring networks:Example applications from the Great Lakes water availability pilot project(Vol.
       2010–5159. Reston, Virgina, USA: USGS, 2010.

Ghysels, G., Benoit, S., Awol, H., Jensen, E. P., Debele Tolche, A., Anibas, C., and Huysmans, M.:
       Characterization of meter-scale spatial variability of riverbed hydraulic conductivity in a lowland river
       (Aa River, Belgium). Journal of Hydrology, 559, 1013–1027, doi:10.1016/j.jhydrol.2018.03.002, 2018.

835 Gianni, G., Doherty, J., and Brunner, P : Conceptualization and calibration of anisotropic alluvial systems: Ptifalls
       and biases. Groundwater, doi:10.1111/gwat.12802, 2018.

He, J., Sarma, P., Bhark, E., Tanaka, S., Chen, B., Wen, X-H., Kamath, J.: Quantifying expected uncertainty
       reduction and value of information using ensemble-variance analysis. SPE J. 23, 428–448,
       doi:10.2118/182609-PA, 2018.

840 Hermans T.: Prediction-focused approaches: an opportunity for hydrology. Groundwater 55:683–687, doi:
       10.1111/gwat.12548, 2017.

HydroAlgorithmics Pty Ltd. AlgoMesh User Guide. Melbourne, VIC. 2022.

James, S.C., J. Doherty, and A.-A. Eddebarh.: Practical Postcalibration uncertainty analysis: Yucca Mountain,
       Nevada, USA. Groundwater 47, no. 6: 851–869., doi:10.1111/j.1745-6584.2009.00626.x, 2009.

845 Juda, P., Straubhaar, J., and Renard, P .: Comparison of three recent discrete stochastic inversion methods and
       influence of the prior choice. Compt. Rendus. Géosci. 355, doi: 10.5802/crgeos.160, 2022.

Khambhammettu, P., Renard, P., and Doherty, J.: The Traveling Pilot Point method. A novel approach to
       parameterize the inverse problem for categorical fields. Advances in Water Resources, 138, 103556.
       doi:10.1016/j.advwatres.2020.103556, 2020.

850 Kikuchi, C.: Toward increased use of data worth analyses in groundwater studies. Groundwater 55 (5), 670–673.,
       doi: 10.1111/gwat.12562, 2017.

Lam, D.-T., Kerrou, J., Renard, P., Benabderrahmane, H., and Perrochet, P.: Conditioning multi-gaussian
       groundwater flow parameters to transient hydraulic head and flowrate data with iterative ensemble
       smoothers: a synthetic case study. Front. Earth Sci. 8:202, doi: 10.3389/feart.2020.00202, 2020.

855 Lima, M.M., Emerick, A.A. and Ortiz, C.E.P.: Data-space inversion with ensemble smoother. Comput. Geosci. 24,
       1179-1200, doi:10.1007/s10596-020-09933-w, 2020.

Linde, N., P. Renard, T. Mukerji, and J. Caers, Geological realism in hydrogeological and geophysical inverse
       modeling: A review, Adv. Water Resour., 86, 86–101, doi:10.1016/j.advwatres.2015.09.019, 2015.

Paniconi, C., and Putti, M.: Physically based modeling in catchment hydrology at 50: survey and outlook. Water
860   Resour. Res. 51, 7090–7129, doi: 10.1002/2015WR017780, 2015.

PEST++ Development Team. Manual for PEST++. Downloadable from https://github.com/usgs/pestpp, 2022.

Pyrcz, M.J., Boisvert, J.B. and Deutsch C V. ALLUVSIM: A program for event-based stochastic modeling of
       fluvial depositional systems. Comput Geosci ;35:1671–85, doi :10.1016/j.cageo.2008.09.012, 2009.

Remy, N., Boucher, A. and Wu, J.: Applied Geostatistics with SGeMS: A User's Guide. Cambridge University
865   Press, 2009.

Satija, A. and Caers, J.: Direct forecasting of subsurface flow response from non-linear dynamic data by linear
least-squares in canonical functional principal component space. Adv. Water Resour. 77, 69–81,
doi:10.1016/j.advwatres.2015.01.002, 2015.

Scanlon, B.R, Fakhreddine, S., Rateb, A., de Graaf, I., Famiglietti, J., Gleeson, T., Grafton, R.Q.; Jobbagy, E.,
870   Kebede, S., Kolusu, S.R. and others: Global water resources and the role of groundwater in a resilient
water future. Nature Reviews Earth & Environment. 1—15, doi :10.1038/s43017-022-00378-6, 2023.

Scheidt C, Jeong C, Mukerji T and Caers J.: Probabilistic falsification of prior geologic uncertainty with seismic
amplitude data: application to a turbidite reservoir case. Geophysics; 80M89–M12,
doi:10.1190/geo2015-0084.1., 2015.

875 Schilling, O. S., Gerber, C., Purtschert, R., Kipfer, R., Hunkeler, D., and Brunner, P.: Advancing physically-based
flow simulations of alluvial systems through atmospheric noble gases and the novel ^{37}Ar tracer
method.Water Resour. Res. 53, 10465–10490. doi: 10.1002/2017WR020754, 2017.

Sun, W. and Durlofsky, L. J.: A new data-space inversion procedure for efficient uncertainty quantification in
subsurface flow problems. Math. Geosci. 49(6), 679–715, doi:10.1007/s11004-016-9672-8, 2017.

880 Tonkin, M.J., and J. Doherty.: Calibration-constrained Monte Carlo analysis of highly parameterized models using
subspace techniques. Water Resources Research 45:W00B10, doi :10.1029/2007WR006678, 2009.

Vecchia, A.V. and Cooley, R.L.: Simultaneous confidence and prediction intervals for nonlinear regression
models with application to a groundwater flow model. Water Resour Res, 23 (7), 1237-1250,
doi:10.1029/WR023i007p01237, 1987.

885 White, J.T.: A model-independent iterative ensemble smoother for efficient history-matching and uncertainty
quantification in very high dimensions. Environmental Modelling and Software 109, 191-201,
doi:10.1016/j.envsoft.2018.06.009, 2018.
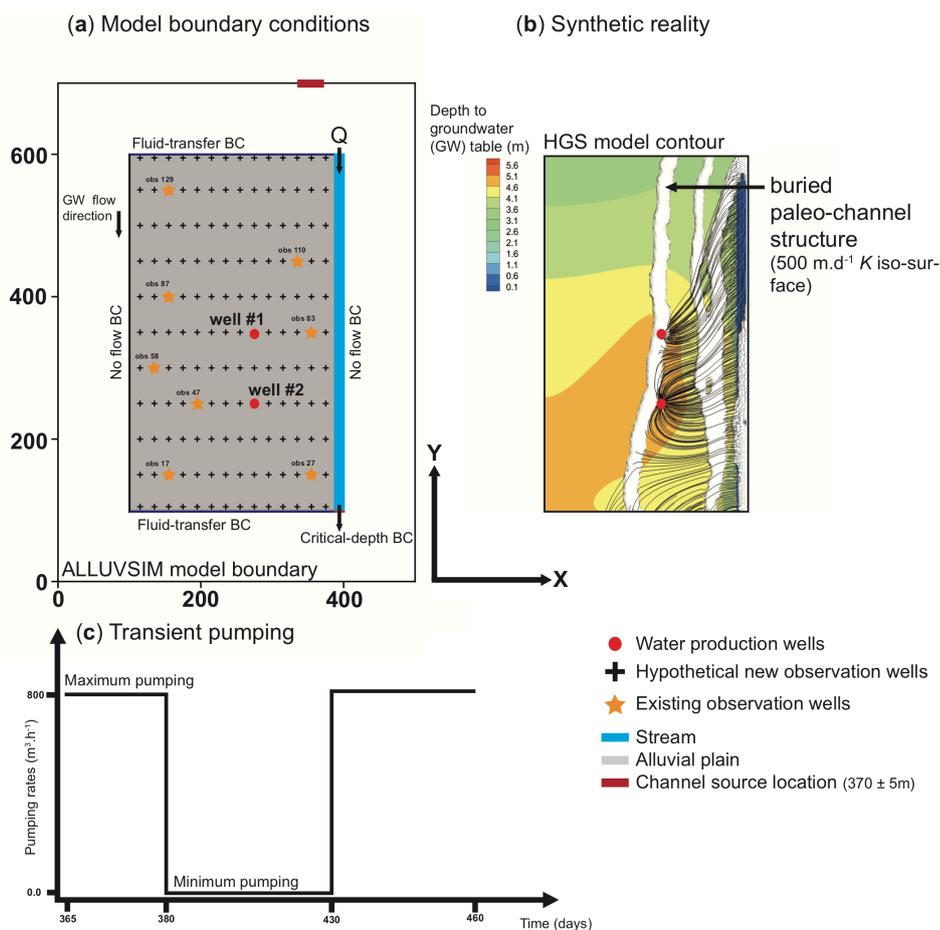
Figure 1: (a) Model boundary conditions and observation points that are discussed herein. (b) "Synthetic reality" showing "true" alluvial channels that govern hydraulic properties such as hydraulic conductivity and porosity. Depth to water table under the maximum pumping regime is also shown. So too are the paths of particles (black lines converging towards water production wells) that are used to calculate travel times to water production wells. (c) Transient pumping rates for the two production wells over a period of 95 days; note that the pumping rate is reduced to zero for a period of 50 days.

Table 1: Model parameters used for the synthetic reality, and prior parameter means and uncertainties used for prior ensemble generation.

Parameter type	Unit	Synthetic reality	Priors mean	Prior standard deviation
HydroGeoSphere				
Riverbed Porosity	(-)	0.4	0.20	0.05 (log)
Riverbed Hydraulic conductivity	(m.d^{-1})	2.5	5.0	0.20 (log)
Aquifer Porosity	(-)	0.09	0.10	0.05 (log)
Aquifer Hydraulic conductivity	(m.d^{-1})	176	150.0	0.10 (log)
Paleo-Channel Porosity	(-)	0.3	0.3	0.05 (log)
Paleo-Channel Hydraulic conductivity	(m.d^{-1})	866	1000.0	0.10 (log)
ALLUVSIM				
Channel level elevation	(m)	90	-	
Channel depth	(m)	20	20	3
Channel source location	(m)	370	370	5
Width-to-depth ratio	(-)	1	-	
Channel sinuosity	(-)	0.1	0.1	0.015
Net-to-gross ratio	(-)	0.05	-	
Maximum number of streamlines	(-)	10	-	

900
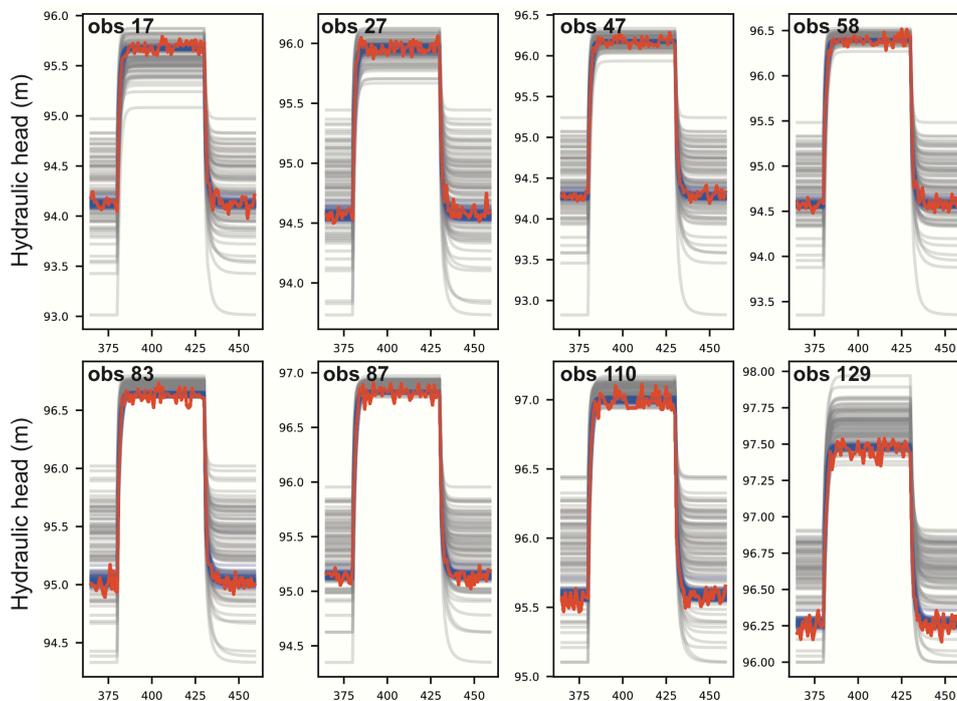
Figure 2: Observations of hydraulic heads (red); these are calculated by the "reality realisation" of the HGS parameter field. Heads calculated by the remaining simulations are shown in grey. Heads calculated using posterior hydraulic property realisations (see below) are shown in blue; there are 70 of these.

905
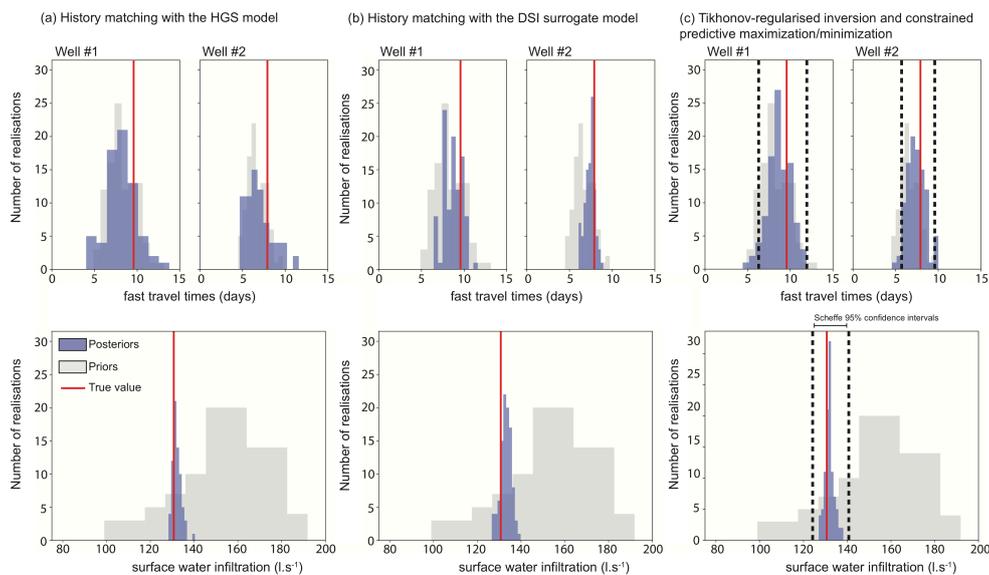
Figure 3: Posterior distributions of model predictions of fast travel times (days) and surface water infiltration (l.s⁻¹) with (a) application of IES (after 3 iterations) to the HydroGeoSphere (HGS) model ensemble (b) application of IES (after 6 iterations) to the surrogate model ensemble and (c) application of equation (21) (linear analysis) and equation (20) (nonlinear analysis) with the MAP estimate of the DSI surrogate model. The dashed lines in (c) represent values for the Scheffe 95% confidence intervals.
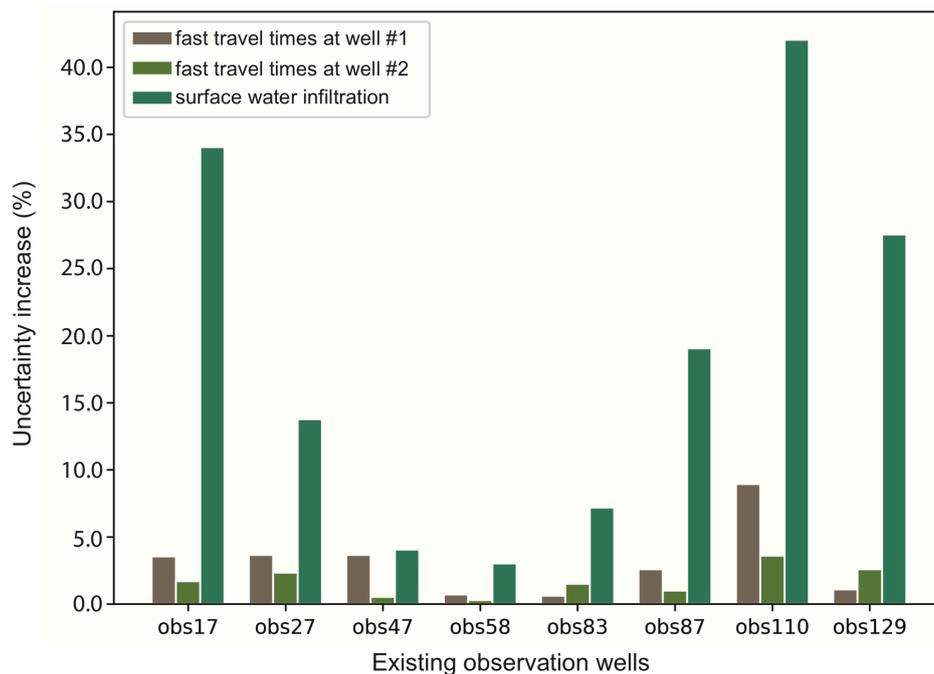
915

Figure 4: Percent increase in the standard deviation of posterior uncertainty of each of the three predictions discussed herein when one existing observation well is removed from the current observation dataset.
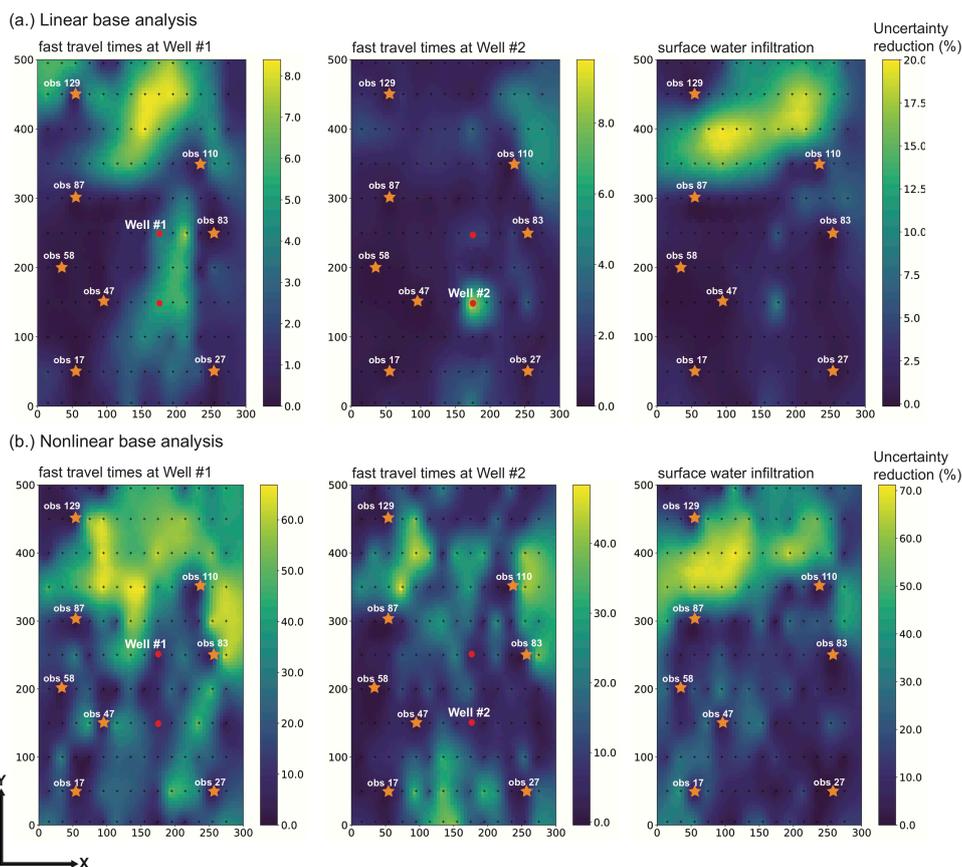
920

Figure 5: Percent decrease in posterior uncertainty of each of the three predictions accrued through supplementing the existing calibration dataset with head measurements gathered in an extra well following (a) linear analysis and (b) nonlinear analysis.

925

Appendices
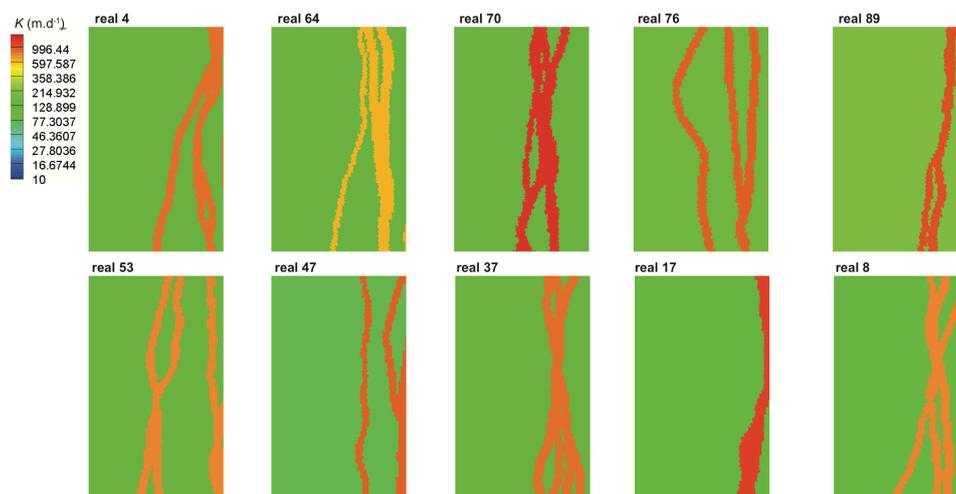
930   **Appendix A. Prior realizations of alluvial structural features generated from ALLUVSIM and mapped into the HGS model for hydraulic conductivity (*K*) parameterisation.**
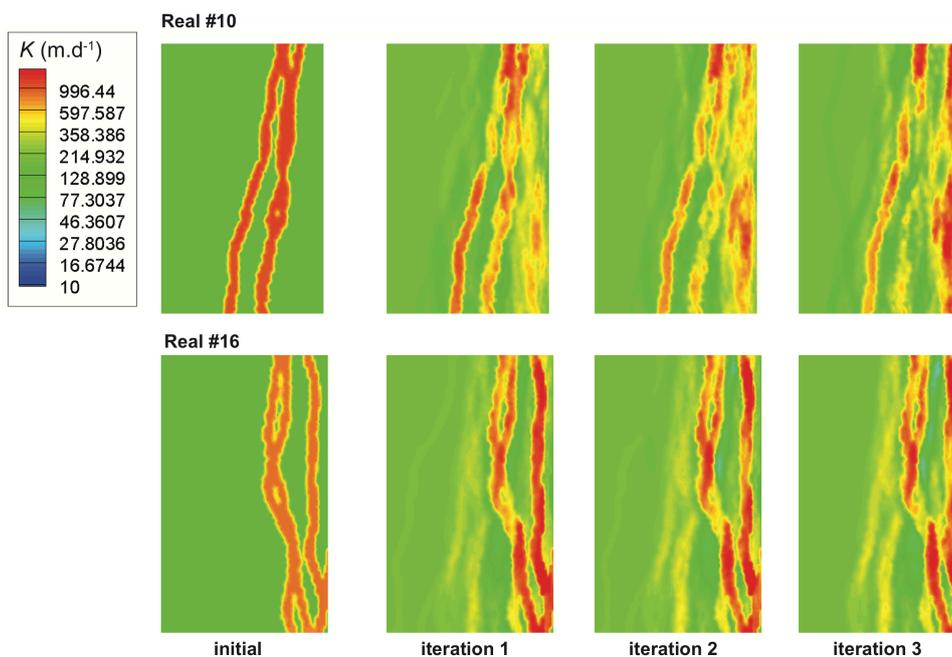
Appendix B. Posterior estimates (for realisations number 10 and 16) of hydraulic conductivity (*K*) after applying the Iterative Ensemble Smoother over three iterations by adjusting the HGS model parameters

935   **with the existing dataset.**

Appendix C: Posterior uncertainty statistics of distributions shown in Figure 3.

940

		Bayesian history-matching on HGS model	Bayesian history-matching on DSI model	MAP estimates on DSI model	Nonlinear analysis on DSI model	Linear analysis on DSI model
HGS model calls		1000	100	100	100	100
fast travel times for well #1 ("reality" = **9.57** days)	Min	4.82	6.59		6.24	
	Max	11.86	10.50		11.89	
	Mean	8.13	8.65	8.63		
	Std	1.63	1.00			1.44
fast travel times for well #2 ("reality" = **7.90** days)	Min	4.67	6.23		5.65	
	Max	9.95	8.52		9.54	
	Mean	6.82	7.36	7.42		
	Std	1.37	0.52			0.94
surface water infiltration ("reality" = **130.67** l.s^{-1})	Min	128.30	128.33		124.06	
	Max	135.89	137.98		140.58	
	Mean	131.94	133.27	131.90		
	Std	1.67	2.21			2.15