

Interactive Inverse Methodology Applied to Stratigraphic Forward Modelling

C. WIJNS^{1,2}, T. POULET¹, F. BOSCHETTI¹, C. DYT¹ & C.M. GRIFFITHS¹

¹ CSIRO, PO Box 1130, Bentley, WA 6102, Australia

² School of Earth and Geographical Sciences, University of Western Australia, Crawley, WA 6009, Australia

Abstract: An effective inverse scheme that can be applied to complex three-dimensional hydrodynamic forward models has so far proven elusive. In this paper, we investigate an interactive inverse methodology that may offer a way forward. The scheme builds on previous work in linking expert review of alternate output to rapid modification of input variables. This was tested using the SedSim three-dimensional stratigraphic forward modelling program, varying nine input variables in a synthetic example. Ten SedSim simulations were generated, with subtle differences in input, and five dip sections (fences) displayed for each simulation. A geoscientist ranked the lithological distribution in order of similarity to the true sections (the true input values were not disclosed during the experiment). The two or three highest-ranked simulations then acted as seed for the next round of ten simulations, which were compared in turn. After 90 simulations, a satisfactory match between the target and the model was found and the experiment was terminated. Subsequent analysis showed that the estimated input values were ‘close’ to the true values.

Underlying most forward modelling exercises in geosciences is an implicit inverse question. A formal inverse approach to geological modelling has only recently started to develop due mostly to two reasons: first, standard geological modelling is extremely computationally intensive, and second, it is hard to develop numerical cost functions, to drive the inverse search, that can capture enough geological knowledge to make the search meaningful. Creating a cost function is difficult because it involves distilling into a single value an assessment of the geometries, volumes, and positions, as well as determining how geologically reasonable a solution may be.

We employ interactive inversion to tackle the second problem. The concept is of such simplicity as to make it appear as little more than a novelty: replace the numerical evaluation of solution quality with a subjective value provided by an expert user. Underneath the simplicity there is a very powerful concept. We replace what in most cases is an artificial numerical function with the computational power of an expert brain. This provides not only the best ‘geological processing power’ we currently have, but also intuition and expertise, in the form of theoretical knowledge, experience, and *a priori* information. Here, the notion of ‘expert user’ may take different meanings. At the most superficial level, a geologist may have enough expertise to reliably evaluate basic similarities in different geological fields. This is the level we use in this work. At another level, a user who is an expert in some specific forward modelling can employ his or her experience to better interpret the modelling output in terms of the software functionality. As a result, for example, convergence may be im-

proved, and geological understanding may benefit. Yet another level involves a user who is an expert in the optimisation tool used (a genetic algorithm in this work), who may be able to perceive when the inversion itself needs ‘help’, and tune inversion parameters on the run. Boschetti & Moresi (2001) demonstrate applications of these last two cases of expertise. It would be possible to employ all these levels jointly by asking users with different expertise to rank models together.

Attempts to formalize the geological processing power of a human expert have been presented in the literature, but, to our knowledge, are incomplete. Our own experience in attempting to develop a geological image similarity module, accounting for human input, is encouraging (Kaltwasser et al. in press) but still far from being applicable to real world problems. Extensive reference to interactive inversion, standard implementations, and a review of several artistic and engineering applications can be found in Takagi (2001). Boschetti & Moresi (2001) and Wijns et al. (2003) present applications to geoscientific problems.

More recently, we have focussed on visualizing the results of an interactive inversion. The purpose is to enable our inverse strategy to return information to the user in a ‘language’ similar to the one the user employs to feed information to the code. We expect the user to be an expert geoscientist, not an expert in inverse theory. We want the user to understand visually what the inversion has achieved, to obtain further insights into the problem, and, by doing so, to provide more information to the inversion. Such information can be used as further *a priori* input in a subsequent inversion run (e.g., Boschetti et al. 2003).

Variable	True	Step size	Min.	Max.
E-W source location (m)	9500	200	8000	12000
N-S source location (m)	600	90	20	1100
E-W source velocity (m s^{-1})	0	0.25	-1.0	1.0
N-S source velocity (m s^{-1})	0.5	0.2	0.4	2.2
Flow rate (m^3s^{-1})	9	1.5	1	13
Conc. (0.5 mm) (kg m^{-3})	0.025	0.01	0.01	0.1
Conc. (0.2 mm) (kg m^{-3})	0.10	0.03	0.0	0.3
Conc. (0.07 mm) (kg m^{-3})	0.175	0.06	0.0	0.3
Conc. (0.0004 mm) (kg m^{-3})	0.20	0.06	0.0	0.3

Table 1: Parameter values used in the Sedsim simulations.

We extend this approach by employing more than one visualization method. Complicated information can be better represented by combining different tools. This method allows the user to provide more information to the inversion than a single measure of quality. What information, i.e., quality measure, the user may decide to provide does not need to be determined before the inversion, but may arise spontaneously during the procedure because of new insight into the problem.

The problem

The inversion has been applied to a simple sedimentation problem where the initial conditions are known exactly. A genetic algorithm (GA) provides an engine for the inversion. Since geological forward modelling is time consuming, the GA population is kept fairly small. Boschetti et al. (2003) discuss convergence issues related to interactive GAs, while a more general discussion of GAs can be found in Wright (1991). Here we use populations of ten individuals. At the end of the exercise, we have collected 90 individuals ranked by a user in terms of similarity to a target image. Intra-generational rankings are simultaneously accumulated in a global ranking order for post-inversion analysis.

The sedimentary model was deliberately chosen to be quite simple, and represents a ramp setting over a 19 km by 19 km area, with a one kilometre grid spacing. The topography slopes down from the south at 0.1 degree. Neither tectonic movement nor sea level change was incorporated. This is a more complex scenario than it may initially seem, as the below sea level accommodation was rapidly filled and a more complex meandering stream and delta formed. Likewise, the low angle of the initial topography leads to a greater variation in possible sediment profiles than would be expected from a steeper gradient. The model incorporates 4 different classes of sediments, including a medium sand (0.5 mm), a fine sand (0.2 mm), a finer sand (0.07 mm) and a clay (0.0004 mm). The only parameters permitted to vary in these simulations are those affecting the input conditions of the fluid laden with sediment. These are the initial east-west and north-south locations

of the fluid source, the initial east-west and north-south velocities of the fluid source, the flow rate of the fluid source, and the initial sediment concentrations of each of the four different grain sizes.

The true values used in the Sedsim simulation are given in Table 1, along with the allowed search ranges. The model was run for 50,000 years. The initial parameters and the step sizes outlined in Table 1 were deliberately chosen so that the target simulation could not be exactly reproduced. The goal of the inverse model is, of course, to retrieve the true (or closest) input parameters from the observations (the output of a stratigraphic simulation). The initial conditions for the trial simulation (Table 1) were chosen by a domain expert as being likely ‘reasonable’ priors. A more formal approach to prior selection may involve analogue studies in a real case, or the complementary interactive methods of Curtis & Wood (2004) (this volume).

Each of the variables was randomly initialized by a genetic algorithm. This is a standard strategy in genetic algorithm inversion, the influence of which, in the final GA result, is discussed in Wright (1991). Ten simulations were run at each generation, and two output sections, displaying lithology and sediment ages, were used to rank the results in order of similarity to the true sections (Figure 1), while the true input values were hidden. Ranking was carried out by evaluating models according to the classes contained in Table 2, which include both specific characteristics and overall resemblance to the target. The two or three highest ranked simulations then acted as seed for the next round of ten simulations, which were compared in turn. Lithology was used exclusively to rank the first six generations, after which it became necessary to analyse the alternating sediment age distribution in order to distinguish between some lithologically similar models. No separate class was adopted for age characteristics (c.f. Table 2): this comparator was used only for a relative rank between two good models that otherwise shared all the same lithological class rankings.

Lithology was chosen as the primary comparator for several reasons. Sedsim’s main purpose is to predict

Class 1	best models
Class 2	good models ("acceptable")
Class 3	acceptable sediment topography
Class 4	acceptable 0.5 to 0.2 mm grain distribution
Class 5	acceptable 0.5 mm grain distribution
Class 6	acceptable 0.07 mm grain distribution
Class 7	acceptable clay distribution

Table 2: Definitions of different classes for ranking.

lithological variations away from well holes, as this is one of the key features that controls the formation and behaviour of a petroleum system, affecting the source rock, migration of hydrocarbons, reservoir formation, and seal quality. It is also one of the most variable characteristics, especially in near-shore and fluvial environments, where local variations in flow or sediment volume can have a marked effect on the lithology. It may have been more representative of the typical Sedsim usage had the lithology been evaluated at several vertical locations representing the lithology identified at a well location, rather than along an entire dip section. The secondary measure of simulation quality was taken to be a display of age layers with successive ages alternately coloured red and blue. These give a good idea of the volume of sediment deposited at each layer as well as the shape of each deposit, roughly akin to what may be seen on a seismic section.

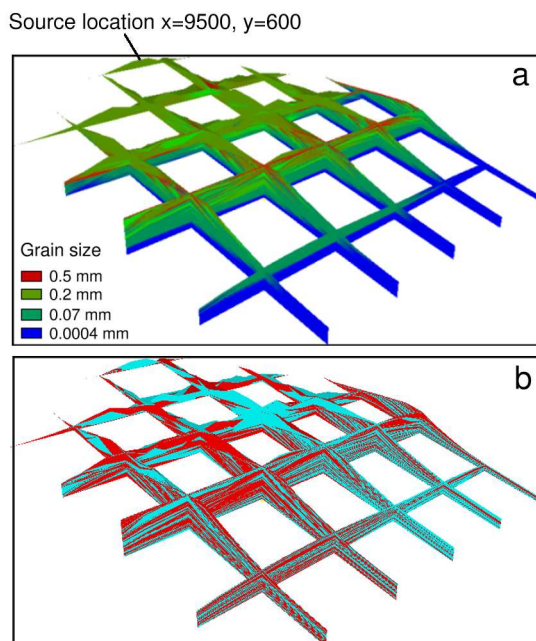


Figure 1: Target simulation with every fourth fence (section) shown for (a) lithology (grain size) and (b) 'alternating age'.

Sedsim stratigraphic forward modelling

Three-dimensional stratigraphic forward modelling enables the combined influence of a variety of inter-dependent basin processes to be studied at geological time scales. The results reflect possible changes in sediment distribution over time as a function of the changing depositional environment. Past studies have demonstrated the value of this approach (Griffiths et al. 2001; Griffiths & Dyt 2001; Griffiths & Paraschivoiu 1998; Koltermann & Gorelick 1992; Martinez & Harbaugh 1993). Any computer modelling is only as good as the validity of the input data and the algorithms used in the program. The selection of prior information for computer models should be no more onerous than that required for a conceptual geological model of an area. However, the need for quantitative data forces the geologist to a greater degree of commitment than may otherwise be the case.

Sedsim is a three-dimensional stratigraphic forward modelling program developed originally at Stanford University in the 1980s and extensively modified and extended in Australia since 1994. Sedsim flow and sedimentation programs are linked to modules including subsidence, sea level change, wave, storm, and geostrophic current transport, compaction, slope failure, and carbonates. The program models sediment erosion, transport, and deposition, and predicts clastic and carbonate sediment distribution on a given bathymetric surface. The conceptual background to Sedsim is described by Tetzlaff & Harbaugh (1989). Sedsim is controlled by a parameter input file, and files describing relative sea-level change, initial topography/bathymetry, and tectonic movement for each grid cell over time.

Basic principles of Sedsim operation

Sedsim uses an approximation to the Navier-Stokes equations in three dimensions. The full Navier-Stokes equations describing fluid flow in three dimensions are currently impossible to solve due to limitations in computer speed (it would take longer to simulate a flow than the real event). Sedsim instead simplifies the flow by using isolated fluid elements to represent continuous flow (Tetzlaff & Harbaugh 1989, Chapter 2). This Lagrangian approach to the hydrodynamics allows for a massive increase in speed of computation and simplification of the fluid flow equations. Simulations over geological periods can at best hope to capture the mean conditions and create a general pattern of sediment distribution, rather than capture the exact timing of each individual pulse of material.

Fluid elements travel over an orthogonal grid describing the topographical surface, reacting to the local topography and conditions such as the flow density and the density of the medium through which the element is passing (e.g., air, sea water, or fresh water). Fluid elements are treated as discrete points with a fixed volume,

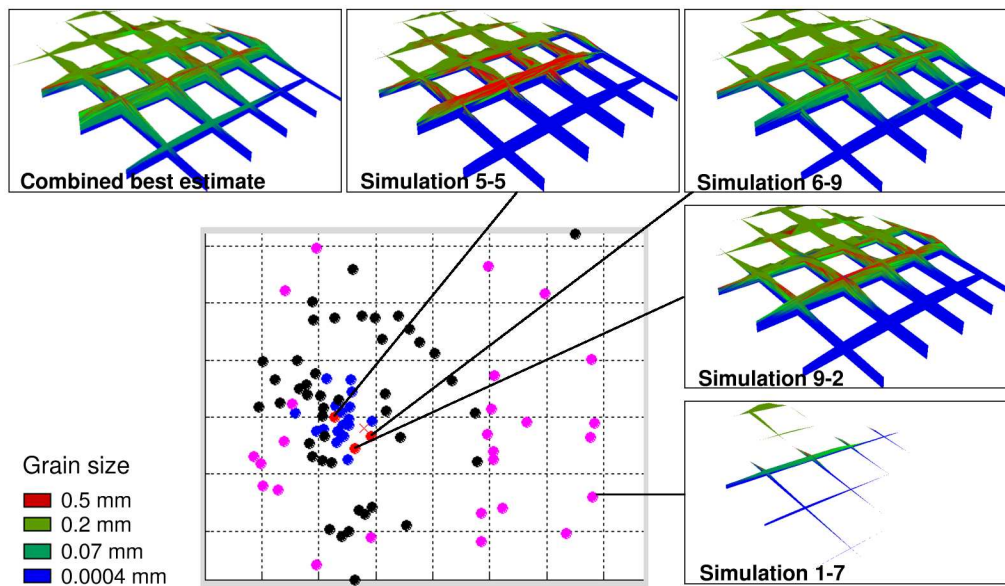


Figure 2: Sammon's map, with each point assigned a colour according to its quality ranking. Axes represent distance, but there is no physical scale. The three 'best' points (simulations very similar to the target) are shown in red, 'good' points (somewhat similar) are shown in blue, 'mediocre' points are grey, and 'very bad' points are coloured magenta. Simulation results for the three best models are mapped onto their corresponding points, and the true (target) solution is plotted on the map by an \times . The 'combined best estimate' is obtained by averaging the three best (red) points.

an approach known as 'marker-in-cell'. Several simplifications are made to the Navier-Stokes equations (Tetzlaff & Harbaugh 1989). The most important of these is that the flow is expected to be uniform in the vertical direction (i.e., the whole of the fluid element has the same velocity), and that the friction experienced by the fluid element is controlled by Manning coefficients. The net result of these simplifications is that the Navier-Stokes equations are modified into non-linear ordinary differential equations. These equations are now solved using a modified Cash-Karp Runge Kutta scheme (Press et al. 1986) that ensures stable and accurate fourth-order-in-time solutions.

Although Sedsim has the capability to model many aspects of the depositional process, the current experiments used only a small subset of the modules available. Modules switched off were: internal nested grids, slope failure, wave, storm, and geostrophic current, sea level change, tectonic subsidence or uplift, syn-depositional compaction, isostasy, and carbonates. Even though the model used here was an extremely simplified version of a typical Sedsim simulation, the depositional environment modelled is typical of a river-dominated delta in a cool, temperate, shallow lacustrine setting with no significant subsidence over a 50 ka period.

The simulations

Apart from the variables discussed above and listed in Table 1, all other input values remained the same for

all 90 simulations. Each simulation took four hours under MSDOS on an Intel computer (Fujitsu LifeBook[®]) running at 850 MHz with 256Mb RAM.

Visualization and analysis of the solution space

The goal of the inversion is to optimize nine parameters controlling a sedimentation process. This results in a nine-dimensional search space. A number of tools have been proposed in the literature to visualize high-dimensional (nD) spaces (Buja et al. 1996). In this work we employ multi-dimensional scaling (MDS) and a self-organizing map (SOM). Both methods attempt to plot nD vectors on a 2D surface in such a way that topology is best respected, i.e., in such a way that points close to one another in the original nD space are plotted in close proximity in the 2D map. Obviously, with the exception of trivial cases, respecting the topology perfectly is impossible. MDS and SOM differ in the way they approximate such an impossible mapping. In a nutshell, MDS yields a better overall picture of the distance relationships between the nD points, while SOM provides a better picture of local relationships, acting more like a clustering algorithm. Accordingly, different kinds of information can be extracted by these two tools. Using the accumulated global rankings of simulation outputs, we visualize and analyse, *a posteriori*, the effectiveness of the inversion.




Figure 3: Sammon's map. Points are labelled with a quality ranking according to similarity to the target, and 1 to 20 refer to good points. Classes 1 and 2 contain the best simulation results, but class 4 points are close to the good points.

Multi-dimensional scaling

MDS works by calculating all the pair-wise distances between points in the nD space. Points are then positioned on a 2D surface in such a way that pair-wise distances in the 2D map match the original distances as much as possible. Among possible MDS implementations, we have adopted Sammon's mapping (Sammon 1969), which emphasizes the preservation of small distances.

As explained above, we cannot expect this mapping to be optimal. Consequently, we cannot expect the map to be always easily interpretable. The highly non-linear mapping inherent in the MDS approach may result in a series of projections, rotations, and stretches that may hinder the interpretation of data relations. Fortunately, this is not the case in our test. Figure 2 shows an MDS map containing all 90 points collected during the inversion. Different symbols indicate the simulation quality. The 'best' points (ranks 1-3) are coloured red, while other 'good' points (ranks 4-20) are blue, 'mediocre' points are grey, and 'very bad' points are magenta. Points belonging to these first two sets plot close to one another, which may suggest the existence of an approximately convex area in the original nD search space containing good models. It is tempting to infer something

about the uniqueness of the solution. However, the limited number of points, as well as the approximate mapping we are using, compels us to use caution in drawing such conclusions. This map should be used as a tool to suggest patterns that need to be properly verified by further analysis.

An obvious (though not sufficient) test to verify convexity is to generate an input by the weighted average of all good models, and judge whether the resulting geological model is also good. The simulation output from the weighted average model, shown in Figure 2 as the 'combined best estimate', weakly confirms the convexity hypothesis. Since an MDS map is not reversible, it is not possible to take a generic point on the 2D map and transform it into the original nD space.

The above analysis refers only to the subjective ranking of image similarity. In producing this ranking, the user found it useful to assign the solutions to one or more of the classes in Table 2. The first two classes represent the good models, as mentioned above. The other classes represent models that are not good overall, but that possess a good feature. Because of the non-linearity inherent in the physics of the model, relationships between input parameters and classes are not obvious beforehand. This is information we hope to obtain from the inversion process.




Figure 4: Sammon's map. Points are labelled with a quality ranking according to similarity to the target, and 1 to 20 refer to good points. Classes 1 and 2 contain the best simulation results, and class 6 points are relatively far from the good points.

The MDS map can be used to visualize the comparative locations of the different classes within the parameter space. Figures 3 and 4 contain the points belonging to, respectively, classes 4 and 6, plotted with grey circles. Figure 3 suggests that '0.5 to 0.2 mm distribution' is an important factor in judging the quality of a model, as it is rare to have a bad simulation within this class. On the other hand, Figure 4 suggests that models are not very sensitive to '0.07 mm grain distribution' since both good and bad models fit this criterion. Simulations such as these also indicate the sensitivity of simulation results to the various input parameters. The Sammon's map in Figure 3 shows that the quality of the simulation is strongly dependent on the accuracy of predicting the correct concentration of coarser grain material, with class 4 points strongly matching with class 1 and 2 (best and good models). This makes sense in several ways. The coarse material drops out of suspension from a fluid source first, and hence is typically deposited in either a fluvial or near shore environment. Its presence has a strong effect on subsequent physical processes in the area. It is also capable of forming steep sediment profiles, again altering the hydraulic surface significantly. Fine grained materials tend to form lower slope angles, and drop out of suspension typically in deeper, calmer, water. This has less effect on subsequent flows, which

is why we see in Figure 4 that the quality of a solution is very poorly correlated with class 6, the concentration of fine grained material.

The MDS visualisation provides some other information. First, models closely resembling the target seem to be located within a relatively small sub-domain of the initial search space. Second, this domain seems to be approximately convex. Third, the visual analysis of the models, underlying the interactive inversion approach, apparently captures important aspects of the geological interpretation, since we can not expect the previous two features to arise completely by chance. Although we cannot exclude the existence of other areas in the parameter space containing good models, it is natural to focus the search on the smaller sub-domain identified above, should we wish to further improve our analysis.

The self-organizing map

The last, and probably most important aspect of our analysis, is to understand which input parameters most control the variability in the geological process. We explore this important element via the use of Kohonen's self-organizing map (SOM) (Kohonen 2001). This algorithm spreads nD points over a 2D plane in such




Figure 5: SOM representing the clustering of model outputs (total distance) and the distribution of the magnitudes of the nine variables. Variables correspond to those in Table 1. Only the good points (quality ranking 1 to 20) are labelled on the map.

a way that neighbourhood topology is respected, i.e., two points lying close to one another in the higher-dimensional space should lie close in the 2D plot. In doing so, it acts as a classification algorithm that separates all the input data into clusters according to similarity. Depending on the topology of the nD points, more than one data point may be assigned to a 2D cell. We prevent this from happening by choosing a map containing many more cells than data points.

The SOM is usually displayed by assigning a colour to each cell, which represents the distance between neighbouring cells in the original nD space. These distances show how the map needs to stretch in order to accommodate the complex nD topology. By doing so, it implicitly displays basic data clusters, i.e., sets of data points that are all in close proximity. An example is given in the top left image of Figure 5. Unfortunately, a SOM is more reliable in its representation of local cluster relations. For global relations, the MDS described above is better suited.

An attractive feature of a SOM is that it allows a display of the magnitude of each original dimension (i.e., each model parameter) at each cell location. This is shown in the remaining images in Figure 5. The labels indicate the location on the SOM of the 20 best models

generated during the inversion. Here also, they cluster close to one another. By analysing the variability (and values) of each dimension at the locations corresponding to the best models, we can estimate which input parameters most control the geological process. When good models show a wide range of values for a specific variable, such as for the north-south source velocity, we conclude that the solution is not very sensitive to this variable. On the other hand, if all good models share a fairly constant value of a specific parameter, such as for the flow rate or 0.2 mm grain concentration, we cannot ascertain whether the solution is extremely sensitive to this parameter, or whether the genetic algorithm has fixed upon one value. Our further work will focus on the choice of sample points to avoid the latter scenario.

Discussion

The results show that, at least in this case, convergence to visually similar sections occurred within very few generations. This offers hope that such a technique may be a practical approach to the inversion of complex 3D stratigraphic forward models. The danger of getting trapped in local minima is always present, in GA approaches as much as in other inverse strategies. One




Figure 6: Convergence graphs of the nine variables, for selected generations, with the underlying true values shown as vertical lines. (a) Source location in X and Y coordinates, (b) initial velocities in the X and Y directions, (c) flow rate, and (d) grain concentrations.

approach to testing for the existence of these in a GA environment is to introduce larger value ‘jumps’ every n^{th} generation and test for survivability.

There are obviously many different ways one could extract final value estimates from the resulting predictions. One could take modal values from the ten most successful of all generations. Figure 2 shows that, even without using a weighted average of predicted values, the best results from simulations in the sixth and ninth generations (6-9 and 9-2) are visually similar to the target simulation. In the case of the combined best estimate model shown in Figure 2, a study carried out (before the model was run) using between 3 and 20 best simulations, with weights corresponding to their quality, suggested that the optimal choice was three. Thus the combined best estimate was a weighted average of the input values from the three models most visually similar to the target, all shown in Figure 2.

Some interesting features emerge from the predicted input values associated with these visually similar sections. First, the simulations do not improve markedly after the sixth generation. Second, there is a convergence to source (X) values, source velocity (V_x) val-

ues, and flow rate (Q) values that are different from the target values. The stable estimated values are not randomly distributed, but are systematically offset by up to 20% (Figure 6). The most obvious explanation for this is that the GA has settled on a local numerical minimum that gives a result that is visually similar to the target, i.e., this minimum is beyond the resolution of a human interpreter. This may indicate that the suggested convexity of the solution space of good models is valid only at the visual level, which is a limitation of the visual ranking method.

It is clear that the comparison component could be automated given a suitable section comparator as described by Griffiths et al. (1996) and Griffiths & Duan (2000), and the process sped up considerably. This may be desirable considering the observation that the convergence slows after the sixth generation. Obviously, visual comparison is very effective for quick convergence in the initial generations. In fact, it is probably more effective than numerical comparisons at this first stage, where a human can quickly evaluate multiple criteria with ease when outputs are very different. However, as outputs become more similar, it is harder

for a human to rank them. For this reason, age characteristics were used to distinguish between lithologically similar results. More types of representation could have been used, e.g., water depth facies and porosity, but the challenge for a human evaluator increases accordingly. Thus at this stage, when model results become similar, it may be desirable to substitute, or otherwise integrate, a numerical evaluation. Another approach could involve capturing the decision making process of the modeller and using that to guide the inversion (e.g., Kaltwasser et al. in press). Nevertheless, this has been a valuable demonstration of the possibility of combining a relatively simple, subjective comparison scheme with a GA approach to taking geological priors and testing their validity.

Conclusion

An interactive inversion scheme, which combines human evaluation of model outputs with a genetic algorithm for the exploration of parameter space, is effective in converging towards a target section when using 3D hydrodynamic modelling. The use of human interaction may have a particular advantage in ranking model outputs near the start of the process, where multiple criteria can be quickly evaluated. At the end of the inversion, we visualize the parameter space for all results, using, in this case, multi-dimensional scaling and self-organizing maps, to draw conclusions about the relative importance of different physical parameters that control the sedimentation process. The visualization may also be used to delimit areas in parameter space where further investigations should be concentrated.

References

- Boschetti, F. & Moresi, L. 2001. Interactive inversion in geosciences. *Geophysics*, 66 (4), 1226–1234.
- Boschetti, F., Wijns, C. & Moresi, L. 2003. Effective exploration and visualisation of geological parameter space. *Geochemistry, Geophysics, Geosystems*, 4 (10), 1086, doi:10.1029/2002GC000503.
- Buja, A., Cook, D. & Swayne, D. 1996. Interactive high-dimensional data visualization. *Journal of Computational and Graphical Statistics*, 5, 78–99.
- Curtis, A. & Wood, R. 2004. Input uncertainty and conditioning in siliciclastic process modelling. In: Curtis, A. & Wood, R. (Eds.) *Geological Prior Information*. Geological Society, London, Special Publications, 239, xx–xx.
- Griffiths, C., Duan, T. & Mitchell, A. 1996. How to know when you get it right: A solution to the section comparison problem in forward modelling. In: *Proceedings of Numerical Experiments in Sedimentology*. May 1996, University of Kansas.
- Griffiths, C. & Paraschivoiu, E. 1998. Three-dimensional forward stratigraphic modelling of early Cretaceous sedimentation on the Leveque and Yampi Shelves, Browse Basin. *APPEA Journal*, 38 (1), 147–158.
- Griffiths, C. & Duan, T. 2000. Quantitative comparison of observed stratigraphy and that predicted from forward modelling. In: *Proceedings of the 31st International Geological Congress*, 6–17 August 2000, Rio de Janeiro, CDROM.
- Griffiths, C. & Dyt, C. 2001. Six years of SedSim exploration applications (abstract). *AAPG Bulletin*, 85, 13.
- Griffiths, C., Dyt, C., Paraschivoiu, E. & Liu, K. 2001. SedSim in hydrocarbon exploration. In: Merriam, D. & Davis, J. (Eds.), *Geologic Modeling and Simulation*. Kluwer Academic, New York, 71–97.
- Kaltwasser, P., Boschetti, F. & Hornby, P. (in press). Measure of similarity between geological sections accounting for subjective criteria. *Geophysical Research Letters*.
- Kohonen, T. 2001. *Self-organizing maps*. 3rd edition. Series in Information Sciences, vol. 30. Springer, New York, 501 pp.
- Koltermann, C. & Gorelick, S. 1992. Palaeoclimatic signature in terrestrial flood deposits. *Science*, 256, 1775–1782.
- Martinez, P. & Harbaugh, J. 1993. Simulating nearshore environments. *Computer Methods in the Geosciences*, vol. 12. Pergamon Press, New York, 265 pp.
- Press, W., Flannery, B., Teukolsky, S. & Vetterling, W. 1986. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, New York, 818 pp.
- Sammon, J. 1969. A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, C-18, 401–409.
- Takagi, H. 2001. Interactive evolutionary computation: Fusion of the capacities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89 (9), 1275–1296.
- Tetzlaff, D. & Harbaugh, J. 1989. *Simulating Clastic Sedimentation*. Computer Methods in the Geosciences. Van Nostrand Reinhold, New York, 202 pp.
- Wijns, C., Boschetti, F. & Moresi, L. 2003. Inverse modelling in geology by interactive evolutionary computation. *Journal of Structural Geology*, 25(10), 1615–1621, doi:10.1016/S0191-8141(03)00010-5.
- Wright, A.H. 1991. Genetic algorithms for real parameter optimization. In: Rawlins, G. (Ed.), *Foundations of Genetic Algorithms*. Morgan Kaufmann, San Mateo, 205–218.