Towards inverse modeling of turbidity currents: The inverse lock-exchange problem

Lutz Lesshafft [a,*], Eckart Meiburg [a], Ben Kneller [b], Alison Marsden [c]

[a] Department of Mechanical Engineering, University of California at Santa Barbara, Santa Barbara, CA 93106, USA
[b] Department of Geology and Petroleum Geology, University of Aberdeen, Aberdeen AB24 3UE, UK
[c] Department of Mechanical and Aerospace Engineering, University of California at San Diego, La Jolla, CA 92093, USA

ARTICLE INFO

ABSTRACT

A new approach is introduced for turbidite modeling, leveraging the potential of computational fluid dynamics methods to simulate the flow processes that led to turbidite formation. The practical use of numerical flow simulation for the purpose of turbidite modeling so far is hindered by the need to specify parameters and initial flow conditions that are *a priori* unknown. The present study proposes a method to determine optimal simulation parameters via an automated optimization process. An iterative procedure matches deposit predictions from successive flow simulations against available localized reference data, as in practice may be obtained from well logs, and aims at convergence towards the best-fit scenario. The final result is a prediction of the entire deposit thickness and local grain size distribution. The optimization strategy is based on a derivative-free, surrogate-based technique. Direct numerical simulations are performed to compute the flow dynamics. A proof of concept is successfully conducted for the simple test case of a two-dimensional lock-exchange turbidity current. The optimization approach is demonstrated to accurately retrieve the initial conditions used in a reference calculation.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

The modern sea floor in deep water is in large part composed of turbidites, the deposits of submarine turbidity currents. Turbidity currents are believed to constitute the principal mechanism of sediment transport from shallow water into the deep sea (see Meiburg and Kneller, 2010 and references therein). Over geological time scales, stacked turbidites may accumulate on abyssal plains and local deep-sea basins, or form submarine fans fed by river deltas. Deeply buried, sandy turbidite sheets represent an important class of hydrocarbon reservoirs, many of which are situated in deep water (Weimer and Slatt, 2007). Many of these turbidites are essentially sheetlike in nature.

There is a large body of experimental work on gravity and turbidity currents that has established the dimensional relations (see the review by Huppert, 2006), and the geometries and properties of their deposits at laboratory scales (e.g. Luthi, 1981; Laval et al., 1988; Middleton and Neal, 1989; Bonnecaze et al., 1993; Garcia, 1994; Alexander and Morris, 1994; Kneller and McCaffrey, 1995; Mulder and Alexander, 2001; Parsons et al., 2002; Al ja'aidi et al., 2004; Baas et al., 2005; Violet et al., 2006). Although the results of many of these experiments have been explicitly applied to submarine deposits at natural scales, such extrapolations are subject to considerable scaling uncertainties, as pointed out by Parsons et al. (2007).

The interaction of turbidity currents with the sea floor, via deposition and erosion of sediment, leads to complex sedimentary features that can only be understood by investigating the flow processes that created them. A recent review of fluid-mechanical research on turbidity currents is given by Meiburg and Kneller (2010).

The application of computational fluid dynamics (CFD) techniques to the simulation of marine turbidity currents is severely restricted due to the length scales involved. High-fidelity simulations that resolve the full Navier–Stokes equations (e.g. Necker et al., 2002, 2005; Cantero et al., 2007; Gonzalez-Juez et al., 2009) are limited to laboratory-scale settings; CFD simulations of field-scale turbidity currents at present are usually based on model equations (see the review by Parsons et al., 2007). Felix (2002) attempted depth-resolved two-dimensional simulations of two historical turbidity currents, based on a Reynolds-averaged Navier–Stokes (RANS) formulation, and further progress in turbulence modeling as well as computational performance may be expected to bring 3D field-scale simulations into the reach of high-fidelity approaches in the future. However, before any CFD method may be used to simulate an actual turbidity current event that took place in the past, and thereby reconstruct its deposit, precise knowledge of the *initial flow conditions* is required. In reality, this need only to be a set of notional conditions at some point along the transport pathway that is sufficient to define uniquely the subsequent behavior of the flow, and its deposition. Such information is generally not available.

* Corresponding author. Current address: Laboratoire d'Hydrodynamique (LadHyX), École Polytechnique, 91128 Palaiseau, France. Fax: +33 1 69 33 52 92.
  E-mail address: lutz.lesshafft@ladhyx.polytechnique.fr (L. Lesshafft).

If information about the thickness and the grain size distribution in a turbidite bed can be inferred from data at a few locations (such as wells), one may attempt to identify initial conditions for a flow simulation that provides a *best fit* in these locations between measured data and numerical predictions of the sediment deposit. The present paper outlines an approach, demonstrated in the context of a simple model problem, to perform such an *inversion* of localized sample data in an automated process. To this end, an optimization problem is formulated, which aims at minimizing the discrepancy between simulation results and reference data. To our knowledge, no published literature on turbidity current inversion exists to date.

Towards the long-term goal of accurate field-scale turbidity current inversion based on well data, the objective of the present study is to introduce a general formulation of the inverse problem, to propose an optimization strategy for its solution, and to present test results for the latter as a first proof of concept. The test presented here is performed under highly simplified, but very controlled conditions. The two-dimensional lock-exchange problem is chosen as a test case. The inversion strategy is used in conjunction with a direct numerical simulation (DNS) approach. The test objective is to reconstruct a sediment deposit that has initially been generated with the same simulation code. The performance of the inversion strategy can thus be assessed independently of any uncertainties associated with the simulation model or with the accuracy of reference data.

The physical setup of the two-dimensional lock-exchange problem and the employed numerical simulation technique are described in Section 2. The inverse problem is introduced in Section 3. This includes the formulation of an optimization objective and the motivation for our choice of the surrogate management framework (SMF) method as a solution strategy. A full description of the SMF implementation is given in Section 4. Inversion results for the present test case are presented and discussed in Section 5, and conclusions are offered in Section 6.

2. Test configuration: two-dimensional lock-exchange problem

2.1. Physical setup

Widely used in fundamental research, the lock-exchange problem represents the elementary archetype of a gravity current. The specific case of lock-exchange turbidity currents, where density differences in the fluid are due to sediment loading, has been investigated experimentally by Bonnecaze et al. (1993), amongst others, and numerically by Necker et al. (2005) and Blanchette et al. (2005). For the purpose of the present study, this generic configuration may serve as a testbed for the proposed inversion approach.

Consider a confined volume of sediment-laden water (the "lock"), submerged in a rectangular tank filled with otherwise clear water. The configuration is sketched in Fig. 1. In an experimental setting, the sediment-laden and the clear water are initially separated by thin plates, and the water inside the lock is stirred, so as to maintain a homogeneous distribution of suspended sediment throughout its volume. At time $t = 0$, the separating plates are suddenly removed, releasing the heavier suspension into the lighter clear water. Under the influence of gravity, the density difference drives a turbidity current along the bottom of the tank. As the current propagates, its sediment is continuously deposited onto the bottom wall, until the current ultimately comes to rest.

This process is illustrated in Fig. 2, which displays three snapshots of the normalized total sediment concentration from the numerical simulation of a lock-exchange problem. All flow simulations in this study are restricted to a two-dimensional geometry. The numerical method and specific parameters used in the flow configuration of Fig. 2 are described in the following sections.

2.2. Numerical model

The flow is assumed to be governed by the two-dimensional Navier–Stokes equations. Restricting the analysis to moderate levels of sediment loading, with resulting density variations not larger than 5%, these equations may be written in the Boussinesq
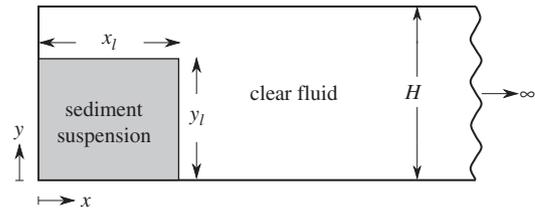

Fig. 1. Lock-exchange configuration, initial condition.
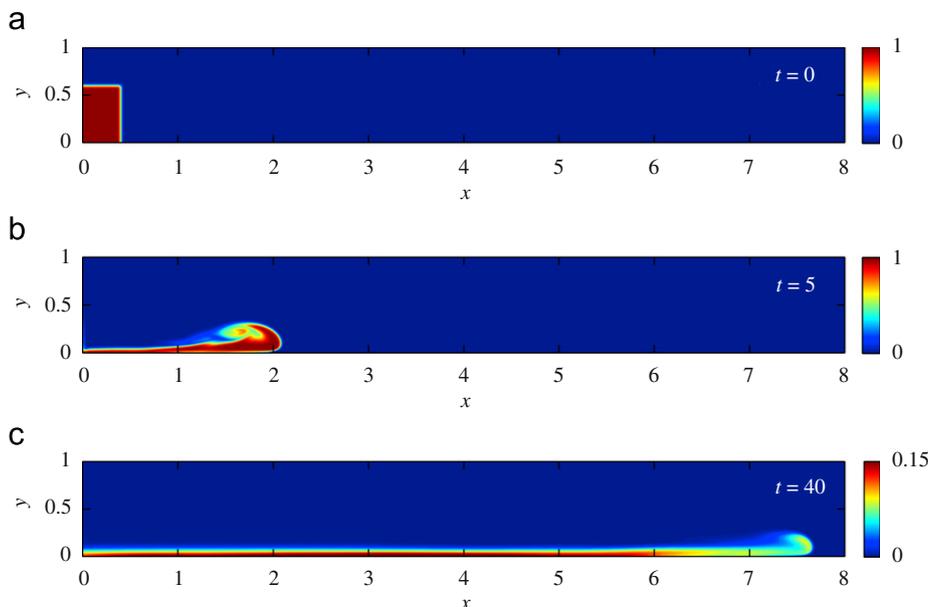

Fig. 2. Snapshots of the total concentration field c_{tot} from the direct numerical simulation of the reference case. The numerical domain extends further downstream to $x = 10$.

approximation (Necker et al., 2005):

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0, \tag{1}$$

$$\frac{\partial u}{\partial t} + u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} = -\frac{\partial p}{\partial x} + \frac{1}{Re}\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right), \tag{2}$$

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} = -\frac{\partial p}{\partial y} - c_{tot} + \frac{1}{Re}\left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2}\right). \tag{3}$$

Eqs. (1)–(3) are given in non-dimensional form. The spatial coordinates x and y are defined in Fig. 1; they are scaled with the tank height H. The horizontal and vertical velocities u and v are normalized with respect to the buoyancy velocity

$$u_b = \sqrt{gH\frac{\rho_p - \rho_f}{\rho_f}c_0}, \tag{4}$$

where g is the gravitational acceleration, ρ_p and ρ_f are the density values of the particle material and of clear water, respectively, and c_0 is a reference value for the volume fraction of particles in suspension (see next paragraph). The Reynolds number is then defined as

$$Re = \frac{u_b H}{\nu} \tag{5}$$

ν being the molecular kinematic viscosity of water. The gravitational term in the vertical momentum equation (3) with these conventions reduces to the total sediment concentration, in the sense of the fraction of volume occupied by sediment particles, $c_{tot}(x,y,t)$, which is normalized with c_0. If N different grain sizes are present, the total sediment concentration is the sum of their individual concentrations c_i:

$$c_{tot}(x,y,t) = \sum_{i=1}^{N} c_i(x,y,t). \tag{6}$$

The reference value c_0 usually is defined as the initial concentration of sediment inside the lock (e.g. Necker et al., 2005). However, for the purpose of the present study, quantitative comparisons will have to be made between simulations with varying initial sediment loading, and it is therefore convenient to use a common reference value c_0 in all simulations. Without loss of generality, c_0 is chosen to be the initial total sediment concentration in the reference case described in the following section.

Sediment concentration in the present context is modeled as a scalar field that evolves in time according to an Eulerian convection–diffusion equation. However, sediment particles do not only follow the fluid motion, but they are constantly subjected to an additional downward gravitational force. The effect of this force is modeled as a constant vertical *settling velocity* v_s, superposed onto the fluid velocity field (Necker et al., 2005). This settling velocity is taken to be identical to the Stokes velocity at which a single particle sinks within a given fluid at rest. If we assume that all particles are spherical and of identical density ρ_p, then the settling velocity is a function of grain size only. In the present framework, all suspended sediment is assumed to be composed of a finite number of grain sizes, each characterized by its settling velocity v_s. One concentration field $c_i(x,y,t)$ is associated with the i-th particle grain size, and for each of the N grain sizes an additional transport equation

$$\frac{\partial c_i}{\partial t} + u\frac{\partial c_i}{\partial x} + (v - v_i^s)\frac{\partial c_i}{\partial y} = \frac{1}{Pe}\left(\frac{\partial^2 c_i}{\partial x^2} + \frac{\partial^2 c_i}{\partial y^2}\right), \quad i = 1,\ldots,N \tag{7}$$

is solved along with the flow equations (1)–(3). The Peclet number Pe is defined with the particle diffusion constant κ as $Pe = u_b H/\kappa$. The concentration fields c_i add up to the total concentration value (Eq. (6)).

Eqs. (1)–(3), (7) are solved numerically on a rectangular domain, discretized by a Cartesian grid with constant cell size. Convective terms are approximated via a third-order ENO ("essentially non-oscillatory") scheme (Osher and Fedkiw, 2003), which is particularly suited for the treatment of propagating sharp fronts. An implicit scheme of second-order centered finite differences is used for all diffusive terms, thereby avoiding time-step restrictions due to low Reynolds or Peclet numbers. Velocities, pressure and concentration values are solved for on a staggered grid. A projection method (e.g. Fletcher, 1991) is used to compute the pressure gradient in (2), (3) in such a way that the continuity Eq. (1) is enforced. This method involves a Poisson problem in p, which is solved numerically with a conjugate gradient algorithm. The equations are advanced in time via a third-order Runge–Kutta scheme, with time steps corresponding to a CFL number of unity.

No-slip conditions are imposed for the velocities at the left and bottom boundaries of the domain, and free-slip conditions at the top surface. An infinitely long tank is modeled by prescribing convective outflow conditions (Ruith et al., 2004) at the right domain boundary $x = x_{max}$. These boundary conditions let flow structures and suspended sediment exit the numerical domain with negligible spurious reflections. Sediment is also allowed to leave the flow domain through the bottom boundary, where it is assumed to settle out and form the bottom deposit. In the current model, topographical changes of the bottom boundary due to deposited sediment are neglected. Neumann conditions are imposed for all concentration fields at the left boundary. For a detailed description of the flow solver, the interested reader is referred to the dissertation of Hall (2009).

2.3. Reference configuration

The lock-exchange configuration that will be used as a test case for the inverse problem is defined by the following parameters:

- lock dimensions $x_l = 0.4$, $y_l = 0.6$, normalized with the tank height,
- three sediment grain sizes with settling velocities $v_1^s = 0.001$, $v_2^s = 0.005$, $v_3^s = 0.01$, normalized with u_b, and initial concentrations $c_1 = 0.5$, $c_2 = 0.25$, $c_3 = 0.25$, normalized such that their sum is 1,
- Reynolds number Re = 2000 and Peclet number Pe = 20 000.

The sum of the non-dimensional initial concentration values c_i is equal to 1, because all concentration values are scaled with respect to the initial total concentration in the present reference configuration.

The flow simulation is performed on a numerical domain of dimensions $0 \leq x \leq 10$ and $0 \leq y \leq 1$, with constant grid spacing $\Delta x = \Delta y = 0.01$. Grid convergence has been confirmed in test calculations with $\Delta x = \Delta y = 0.005$. The value of the Reynolds number has been chosen to be on the lower end of typical values from laboratory experiments. It is large enough to allow the formation of complex vortical flow structures, but sufficiently small to permit direct numerical simulations at a reasonable computational cost. The total number of required grid points roughly increases linearly with Re. The Peclet number, which governs the diffusion of particles, should in theory be infinite, unless indeed the particles were small enough to be subjected to Brownian motion. This parameter is kept at a finite value only to ensure numerical stability. One order of magnitude larger than the Reynolds number, its influence on the simulation results may be expected to be negligible.

The initial total concentration field c_{tot} of the reference configuration is shown in Fig. 2a. As the lock is released, potential energy

of the suspension phase is converted to kinetic energy during the early stage of the flow (Fig. 2b), forming a *current head* that propagates along the bottom (Fig. 2c). In the simulation, sediment contained in the computational cells adjacent to the bottom boundary is allowed to leave the domain at its settling velocity v_s. The non-dimensional rate of change of the bottom deposit height h_i associated with the i-th grain size is given by

$$\frac{\mathrm{d}h_i}{\mathrm{d}t} = \frac{c_0 c_i v_i^s}{\sigma_i}, \tag{8}$$

where σ_i is introduced as the packing density of settled sediment. The effects of σ_i, i.e. the trapping of fluid within a porous sediment layer, are not taken into account in the present investigation, and c_0 remains an arbitrary scaling parameter within the limits of the Boussinesq approximation. In the following, the scaled deposit heights

$$\tilde{h}_i(x,t) = \frac{\sigma_i}{c_0} h_i(x,t) = \int_0^t c_i(x,0,\tau) v_i^s \, \mathrm{d}\tau \tag{9}$$

will be reported. The time integration in Eq. (9) is performed during runtime, and the final deposit heights are thus obtained at the end of the simulation.

As the finest grain size ($v_1^s = 0.001$) takes a very long time to completely sink to the ground even after the current has virtually ceased, a special numerical procedure is introduced at the very end of the simulation: The computation is stopped after 99.9% of the medium grain size sediment has been deposited, at time $t = 83$, and all remaining suspended sediment is added to the bottom deposit at its current x-location. The convective horizontal sediment flux, defined as

$$F_c(x,t) = \int_0^1 c_{tot}(x,y,t) u(x,y,t) \, \mathrm{d}y \tag{10}$$

at $t = 83$ has a maximum value along x of 2.9×10^{-5}, only 0.024% of its overall maximum of 0.122 at $t = 1.46$. Further horizontal sediment transport after $t = 83$ may therefore be neglected.

The resulting final deposit profiles are displayed in Fig. 3. As may have been expected, the largest grains settle out fairly close to the flow origin (triangles), medium-size grains are transported somewhat further downstream (squares), whereas some of the finest grains (circles) remain in suspension all the way down to the final run-out length of the current. Close inspection of the flow evolution shows that local irregular features of the deposit profiles (i.e. deviations from monotonic decay) are associated with large-scale vortices in the flow field. As a side comment, we note that such features could not be captured by a depth-averaged simulation approach.
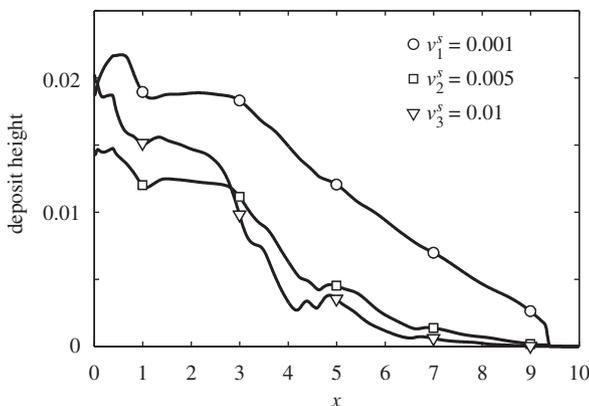

Fig. 3. Final deposit profiles in the reference case for each of the three sediment grain sizes, characterized by their settling velocities v_i^s. The symbols mark deposit heights in the five control points that are used for the inversion.

3. Inverse problem formulation

3.1. Objective

The long-term objective of turbidity current inversion is to reconstruct the global distribution of turbidite bed properties, based on available well-log data in only a few locations. The present paper aims at an optimal reconstruction of the deposit profiles displayed in Fig. 3, based on their values in *five control points*. The true deposit heights for each grain size are assumed to be known at streamwise locations $x = 1, 3, 5, 7$ and 9. Reconstruction of the full profiles is then attempted by performing a number of simulations of the kind described in Section 2.3 with a variety of parameter settings, and by comparing the final deposit in the control points to the objective values. Identifying suitable physical parameters for an optimal match constitutes an *optimization problem*.

3.2. Cost function

The optimization objective is formally expressed as finding the *global minimum* of a cost function $J(\boldsymbol{\alpha})$, which depends on a set of physical simulation parameters $\boldsymbol{\alpha}$, and which measures the discrepancy between the final deposit obtained with these parameters and the reference solution in the five control points. The cost function in the present study is defined as a sum of L^2-norms:

$$J(\boldsymbol{\alpha}) = \sum_{i=1}^{3} \left(\sum_{j=1}^{5} [\tilde{h}_i(\boldsymbol{\alpha}, x_j) - \tilde{h}_i^{ref}(x_j)]^2 \right)^{1/2}. \tag{11}$$

In this notation, $\tilde{h}_i(\boldsymbol{\alpha}, x_j)$ is the deposit height (Eq. (9)) of the i-th sediment grain size computed with the parameters $\boldsymbol{\alpha}$ in the five control points x_j. The reference solution (Fig. 3) is denoted as $\tilde{h}_i^{ref}$. The cost function is non-negative and has a global minimum $J(\boldsymbol{\alpha}) = 0$ where $\boldsymbol{\alpha}$ precisely corresponds to the parameters of the reference case. For all practical purposes, it seems safe to assume that $J(\boldsymbol{\alpha}) > 0$ for all other parameter combinations.

3.3. Free parameters

The choice of free parameters $\boldsymbol{\alpha}$ to be varied in the inverse problem is guided by what can reasonably be assumed to be known from an ill-documented laboratory experiment: known are the height of the tank (reference length H) and the fluid used (viscosity v, hence Re is defined). The particle grain sizes, and therefore the dimensional settling velocities, are known from the documented deposit measurements in the control points. Unknown are:

- dimensions x_l, y_l of the lock,
- initial concentrations c_1, c_2, c_3 of each grain size.

These are the five free parameters that are retained for the inverse problem.

Note that the sum of c_1, c_2 and c_3 in general will not be equal to 1 in the following. The reference values c_0 and u_b, used to make the governing equations (1)–(3) non-dimensional, are *not* required to denote the total concentration and the buoyancy velocity for each initial condition. Instead, they are assumed to represent the same dimensional value in all cases, and it is therefore consistent to use identical values of the non-dimensional parameters v_i^s, Re and Pe in all simulations.

3.4. Choice of method

A well-suited method for the problem at hand needs to be chosen from the abundant numerical optimization techniques

found in the literature. Comprehensive reviews are given by Bewley (in press) and Kolda et al. (2003). A useful first categorization is provided by the distinction between gradient-based and gradient-free techniques.

Gradient-based optimization methods require both the cost function value $J(\boldsymbol{\alpha})$ and its gradient $J'(\boldsymbol{\alpha})$ to be computed at each evaluation. Starting from a given point in parameter space, steepest descent or conjugate gradient methods are then commonly applied in order to arrive at a local minimum. In many applications, the gradient can be efficiently computed by solving the adjoint flow equations (Gunzburger, 1995), provided that the flow dynamics may be reasonably represented in terms of linear fluctuations around a steady mean flow. If no such steady state can be defined, as in the case of the present flow problem, an adjoint-based gradient calculation becomes exceedingly expensive with regard to computational time and storage requirements. Alternatively, the cost function gradient in a low-dimensional parameter space may be approximated via a finite-difference method. At any given point in parameter space, one additional full simulation per free parameter is then required, where each parameter in turn is varied by a small increment.

A particular *gradient-free* optimization strategy, proposed by Booker et al. (1999), has become known as the *surrogate management framework* (SMF) method. Rather than searching for a local minimum along a descent path, the SMF approach first aims at providing a *global* approximation of a largely unknown cost function, by interpolating between a number of points in parameter space, where the cost function has been evaluated. Optimization is then easily and inexpensively performed on the interpolation function (the cost function "surrogate"), and additional evaluations are carried out in new points of interest. The SMF method relies on this surrogate-based search step for efficient exploration of the parameter space that is global in nature and can be customized to suit particular applications.

The efficiency of the SMF and closely related Kriging-based techniques has been demonstrated in many fluid mechanical engineering contexts, such as helicopter rotor optimization (Serafini, 1998; Booker et al., 1999), aeronautical shape design (Chung and Alonso, 2002), airfoil noise reduction (Marsden et al., 2004a, 2004b, 2007) and blood flow in complex vessel geometries (Marsden et al., 2008).

Unlike in many engineering optimization tasks, e.g. aerodynamic design or noise reduction, convergence to a *local* minimum far away from the *global* minimum of the cost function represents no success in the present inverse problem. With regard to the long-term objective of turbidite modeling, only the best possible match with well-log data can provide valid predictions with high probability. Its global approach and its demonstrated efficiency for numerically expensive flow applications make the SMF method particularly appealing for turbidity current inversion.

4. Surrogate management framework

4.1. General concepts

Before any optimization can be performed, the SMF method requires the initial construction of a cost function *surrogate*. The first step therefore consists of evaluating the true cost function for a set of points in parameter space. A surrogate is then found as an interpolating function that exposes the trends between all currently known points. *Kriging* interpolation is the common method of choice for the construction of a surrogate (Jones, 2001), and is used in the present study. A pseudo-random, evenly spaced distribution of initial points across a restricted region of the parameter space is generated with a *latin hypercube sampling*

algorithm (McKay et al., 1979). Following the initial surrogate construction, all further function evaluations are restricted to points that lie on a discrete mesh in parameter space.

SMF optimization relies on alternating application of two complementary procedures, named the *search* and *poll* steps. The exploratory *search* step performs a minimum search on the surrogate, and it returns points of interest for subsequent cost function evaluations. In its simplest form, the search step returns a single optimal value of the current surrogate function. If multiple points may be evaluated in parallel, the search step can be configured to return several local minima, as well as exploratory points in poorly sampled regions of the parameter space. After each new cost function evaluation, the Kriging surrogate is updated and, if the previous search was successful, a new search step is performed.

If the search procedure fails to identify a new optimal point, a *poll* step is initiated. By evaluating the cost function on neighboring mesh points around the current optimal point, the poll step either confirms its local optimality, or returns an improved point. If an improved point is found, a new search step is attempted. Otherwise, the algorithm has successfully converged to a local optimum on the current mesh. If higher precision is desired, the optimization may be continued on a refined parameter mesh. Convergence of the SMF method to a local minimum is formally guaranteed by the poll step in combination with mesh refinement, as discussed by Audet and Dennis (2003), Audet (2004) and Lewis and Torczon (1999, 2000).

4.2. Implementation

The parameter space for the present problem is spanned by the parameters $\boldsymbol{\alpha} = (x_l, y_l, c_1, c_2, c_3)$. Each dimension is discretized over the range $0.1 \leq \alpha_i \leq 1$ with uniform step size $\Delta\alpha_i = 0.05$. Zero values are not included, because the lock volume must be non-zero, and all three grain sizes must be present in the initial configuration. A number of 50 initial points is generated via latin hypercube sampling. These are not required to lie on the mesh. Fifty flow simulations are then performed prior to all optimization. While this fairly large number of initial points represents a significant computational effort (~ 500 CPU hours on AMD Opteron 265), it ensures that all regions of the parameter space are amply sampled, and that the vicinity of the true global minimum may be retrieved by the initial surrogate with some confidence. The Matlab toolbox DACE (Lophaven et al., 2002) is used to perform the Kriging interpolation.

Five new test points are selected in each *search* step. Up to three new points are local minima of the surrogate function. These are identified by conducting line searches (`fmincon` in Matlab) from 100 randomly distributed seeding points. The remaining two new test points, or more if less than three minima are found, are placed in undersampled parameter regions. A formal selection criterion for these points proposed by Cox and John (1997) (see also Jones, 2001) makes use of the fact that the Kriging algorithm not only returns interpolation values, but also provides a local measure of uncertainty ("mean square error", e.g. Jones, 2001). The mean square error σ is large in regions far away from sample points. In order to place new points in regions where the uncertainty is high, and where at the same time the predicted cost function value is low, the criterion of Cox and John (1997) seeks minima of the weighted combined objective

$$\tilde{f}(\boldsymbol{\alpha}, \lambda) = (1-\lambda)f(\boldsymbol{\alpha}) - \lambda\sigma(\boldsymbol{\alpha}). \tag{12}$$

Both the surrogate function $f(\boldsymbol{\alpha})$ and the mean square error $\sigma(\boldsymbol{\alpha})$ are renormalized to vary between 0 and 1 before the weighted function $\tilde{f}(\boldsymbol{\alpha})$ is formed. The weight factor λ is successively stepped up from 0

(pure cost function minimization) in increments of 0.1, until a total of five new points has been found. Direct neighbors are not allowed among new test points in a *search* step.

If, after evaluation of the cost function in all new test points from the *search* procedure, no improved point is found, a *poll* step is performed, i.e. the cost function is evaluated in several points around the current optimum. In this work, the polling set is chosen using the generalized pattern search (GPS) method. Recent variants of the SMF method have also been proposed with alternate polling strategies, including mesh adaptive direct search (MADS, see Audet and Dennis, 2006). A formal requirement for convergence is that the polling directions form a *positive spanning basis* of the parameter space (Davis, 1954; Booker et al., 1999). This condition implies that $N+1$ polling points are required in an N-dimensional parameter space. Among the many possible choices, a polling stencil could be selected randomly. However, in the present implementation, the choice of polling directions is guided by physical considerations.

Given the full deposition profiles from a simulation with the current optimal parameter setting α_0, a first-order estimation of the corresponding profiles for a slightly varied parameter setting $\alpha_0 + \varepsilon$ may be obtained by rescaling the deposit *height* according to ε, while the deposit *shape* is assumed to remain unchanged. Variations in x_l and y_l affect the height of all three grain size deposits simultaneously, variations in c_1, c_2 and c_3 affect the height of the respective deposits individually. For instance, consider a simulation result in which all three deposit curves lie entirely below the objective points shown in Fig. 3. Clearly, improvement in such a case must be sought by increasing the total volume of the lock, the sediment concentration values, or all of the above. In order to select the most promising candidate points for a *poll* step, the deposit curves are estimated on all $3^5 - 1 = 242$ mesh points neighboring the current optimal parameter setting. Estimations are obtained by linearly rescaling the deposit height, according to the relative change in each of the five parameters. Predictions of the cost function values are then obtained by evaluating Eq. (11) for the deposit estimates, and all neighbor points are ranked accordingly.

Following Booker et al. (1999), the first five polling directions must be linearly independent, and the sixth direction is given by their negative sum. The best five linearly independent neighbor points, as ranked according to their estimated cost function value, are therefore selected as polling points, while the sixth is then generated automatically. Note that this sixth point generally is not a direct neighbor of the current optimum.

5. Results

A successful inversion technique returns the correct initial flow conditions as a final result with acceptable accuracy. An *efficient* and *robust* technique must be successful quickly and reliably. Success and efficiency of the SMF strategy applied to our test problem are visualized in Fig. 4. The symbols connected by a solid line represent the cost function value at the current best evaluated point as a function of optimization steps. The numerical current optimal values, together with the corresponding parameter settings, are reported in Table 1. Open symbols in Fig. 4 denote results from a *search* step, solid symbols mark *poll* results. Each *search* step involves five, each *poll* step six cost function evaluations, as explained in Section 4.2. The value at the "zeroth" step is the best value obtained during the 50 initial simulations. The first three optimization steps invoke the *search* algorithm, which yields improved points during the first two steps, but is unsuccessful in step 3. During the fourth step, as traced by the solid line, the *poll* algorithm is called as a consequence. It can be seen from Table 1 that the optimal setting obtained in step 2 is already a direct mesh
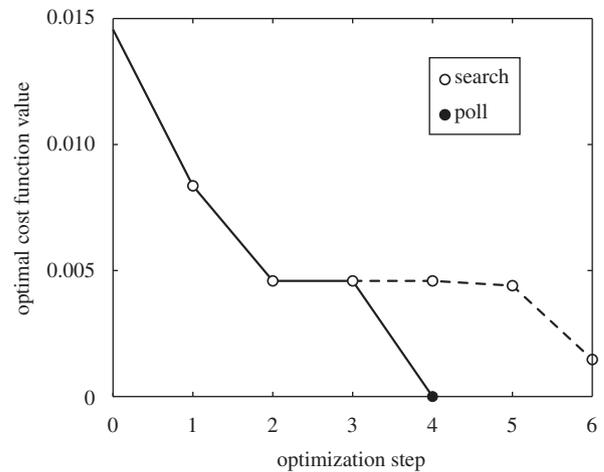

Fig. 4. Cost function values of the current optimal parameter setting as a function of optimization steps. The optimization starts from the full set of 50 initial simulations. Solid line: combination of *search* and *poll* steps; dashed line: *search* steps only.

Table 1
Numerical values of current optimal parameter settings, optimization based on 50 initial simulations. The associated cost function values $J(\alpha)$ correspond to Fig. 4. Upper part: combination of *search* and *poll* steps (solid line in Fig. 4); lower part: *search* steps only (dashed line in Fig. 4).

Step	x_l	y_l	c_1	c_2	c_3	$J(\alpha)$
0	0.23	0.64	0.78	0.73	0.15	0.0146
1	0.4	0.45	0.65	0.3	0.25	0.0084
2	0.35	0.55	0.55	0.3	0.25	0.0046
3	0.35	0.55	0.55	0.3	0.25	0.0046
4	0.4	0.6	0.5	0.25	0.25	0.0000
4	0.35	0.55	0.55	0.3	0.25	0.0046
5	0.35	0.6	0.7	0.25	0.3	0.0044
6	0.35	0.6	0.6	0.3	0.3	0.0015

neighbor of the exact reference configuration $(x_l, y_l, c_1, c_2, c_3) = (0.4, 0.6, 0.5, 0.25, 0.25)$. The *poll* algorithm (step 4) returns this point directly among the five best estimates, and the inversion is completed.

Switching to the *poll* procedure as soon as one *search* step fails to provide an improved point is consistent with the original SMF algorithm described by Booker et al. (1999). However, one might choose to iterate the *search* further, in order to explore more thoroughly different regions of the parameter space. For comparison, results from a continued sequence of *search* steps are displayed in Fig. 4 as open symbols connected by a dashed line. Since each step provides five new known points for the construction of a Kriging-based cost function surrogate, the *search* routine alone may always be expected to yield further convergence. In the present case, improved points are indeed found during steps 5 and 6 (see Fig. 4 and Table 1). However, after step 6 the surrogate predictions appear to become erratic, with all local minima far removed from the reference configuration. Calculations were stopped after step 6. A known problem of the Kriging algorithm is that it requires a set of fairly evenly spaced sample points. If significant "clustering" of nearby points occurs, the interpolation tends to become ill-conditioned, resulting in spurious amplitude variations in less densely sampled regions (Booker, 2000). It appears that, after the sixth *search* step, too many sample points have accumulated in the vicinity of the true optimum, which, as a paradoxical consequence, is not retrieved anymore by the Kriging surrogate. As another surprising observation, we note that all improved points obtained from *search* steps, including
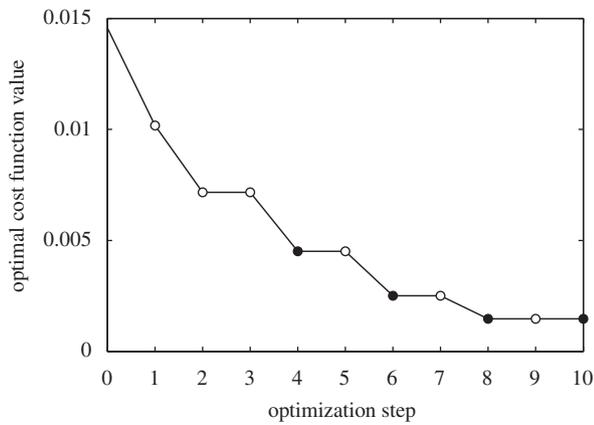
Fig. 5. Cost function values of the current optimal parameter setting as a function of optimization steps. The optimization starts from a *reduced set* of 25 initial simulations. Open symbols: search steps; solid symbols: poll steps.

Table 2
Numerical values of current optimal parameter settings, optimization based on 25 initial simulations. The associated cost function values $J(\boldsymbol{\alpha})$ correspond to Fig. 5.

Step	x_l	y_l	c_1	c_2	c_3	$J(\boldsymbol{\alpha})$
0	0.23	0.64	0.78	0.73	0.15	0.0146
1	0.3	0.55	0.45	0.25	0.3	0.0102
2	0.25	0.6	0.65	0.3	0.35	0.0072
3	0.25	0.6	0.65	0.3	0.35	0.0072
4	0.3	0.6	0.65	0.3	0.3	0.0045
5	0.3	0.6	0.65	0.3	0.3	0.0045
6	0.3	0.65	0.6	0.3	0.3	0.0025
7	0.3	0.65	0.6	0.3	0.3	0.0025
8	0.35	0.6	0.6	0.3	0.3	0.0015
9	0.35	0.6	0.6	0.3	0.3	0.0015
10	0.35	0.6	0.6	0.3	0.3	0.0015

Fig. 6. Final deposit profiles returned by the inversion with 25 initial simulations (thick lines), compared to the reference solution (thin lines and symbols). (a) Smallest grain size with $v_1^s = 0.001$; (b) medium grain size with $v_2^s = 0.005$; (c) large grain size with $v_3^s = 0.01$.

steps 1 and 2, have been found as local minima of the *mixed* objective function (12) with $\lambda > 0$, as opposed to predicted minima of the pure cost function ($\lambda = 0$).

The fast convergence of the SMF method within only four optimization steps, and especially the Kriging-based detection of a direct neighbor of the true optimum in step 2, may be suspected to be coincidental. Furthermore, the initial sampling with 50 points may be considered extensive by some readers, although this number corresponds to an average of only $50^{1/5} = 2.19$ points per parameter dimension. In order to further test the efficiency of the present inversion approach, a second series of calculations is performed, starting from a surrogate function that is based on only 25 initial points. The first half of the same set of initial points as before is used; the best out of all 50 points is among these. The optimization results are presented in Fig. 5, numerical values are given in Table 2.

Just as in the previous series (compare Figs. 4 and 5), the first two search steps yield significant improvement, but the third search is unsuccessful. Subsequent convergence is only achieved through poll steps, each one followed by an unsuccessful search step. Poll step number 10 finally fails to return an improved point, and, according to the rule specifications in Section 4.2, the optimization loop is stopped. The failure of poll step number 10 is an interesting result, because this step includes a simulation with initial conditions $(x_l, y_l, c_1, c_2, c_3) = (0.35, 0.6, 0.55, 0.3, 0.3)$, which by all measures is closer to the reference configuration than the "optimal" point reported in Table 2. However, the cost function in this case is found not to decrease monotonically in the direction towards the global minimum, and can therefore be said to be non-smooth on the scale of the current grid spacing.
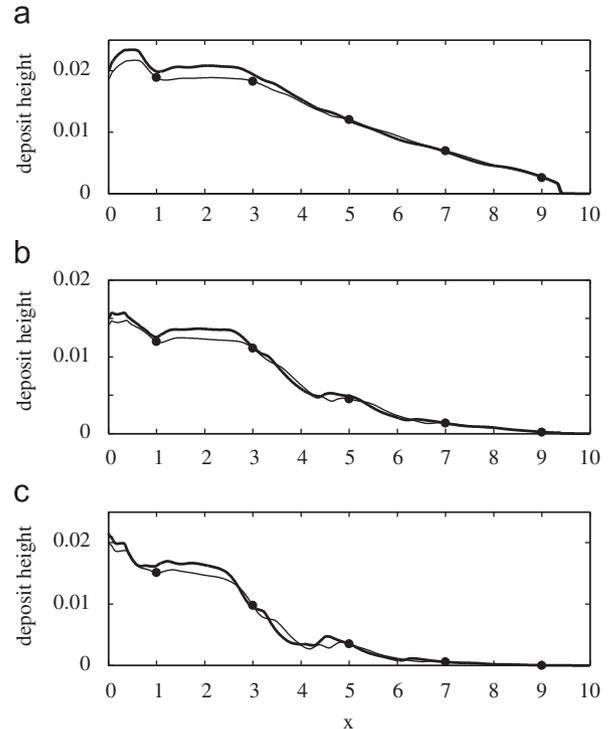
With only 25 initial cost function evaluations, the optimization strategy does not converge to the exact reference configuration. It must be pointed out that the original SMF method by Booker et al. (1999) prescribes refining of the parameter grid upon convergence of the search and poll routines on a given grid. This element is important in the formal convergence analysis by Audet and Dennis (2003). Grid refinement is omitted in our study, due to computational cost considerations.

The final deposit profiles returned by the optimization strategy based on 25 initial evaluations are compared to the reference configuration in Fig. 6. Although no perfect matching is achieved, the results are certainly very close, and for practical purposes the inversion may be considered successful.

6. Conclusion

A numerical inversion strategy has been proposed for the reconstruction of complete deposits from turbidity currents, based on reference data in a small number of control points. The method consists of an iterative optimization procedure that invokes successive numerical simulations of the fluid mechanical and depositional processes in a turbidity current. This automated procedure identifies suitable initial conditions for the flow solver in each iteration step, and thus drives the simulation results towards improved agreement with the reference data.

The strategy has been implemented and tested. The implementation involves a direct numerical simulation (DNS) method that includes a model to account for sediment deposition; the inversion is performed by use of a *surrogate management framework* optimization strategy (Booker et al., 1999).

A two-dimensional turbidity current at low Reynolds number Re = 2000 and with three different grain sizes in suspension has been chosen as a test case to demonstrate the viability of the approach. Reference data for the inversion are generated from a

numerical simulation with a given set of initial conditions, characterized by five parameters. The inversion objective is to reconstruct the deposit thickness and grain size composition across the entire deposit, based on the reference information in five control points as only input. In a first scenario, the inversion is started with 50 initial "blind shot" simulations, and the complete deposit profiles are precisely retrieved within only four subsequent optimization steps. This extraordinary convergence rate commends our strategy, but it must be suspected to be due in part to lucky circumstances. Therefore a second inversion attempt has been presented, based on a reduced number of only 25 initial "blind shots". The final inversion result in this case is computed in the eighth optimization step, the automated procedure exits after two subsequent unsuccessful steps. Agreement with the reference data in this second example is not perfect, but sufficiently close for the test to be deemed successful.

The tests presented in this study demonstrate that the underlying strategy is viable for the inversion of numerically computed turbidity current deposits, and that the SMF method represents a suitable optimization procedure for this purpose. However, a few shortcomings of the present SMF implementation have been exposed: first, the *search* routine in both test examples leads to improvement during the first few optimization steps, but is of little benefit later on. This behavior seems to be due to a degeneration of the Kriging surrogate function, which is known to occur as sample points become increasingly clustered (Booker, 2000). Clustering of sample points is an unavoidable result of the *poll* routine as well as of optimization convergence. Some approaches to remedy this problem have been proposed (Audet et al., 2000; Booker, 2000), and should be further explored in order to improve convergence of the present inversion method. Secondly, in the case with 25 initial simulations, the non-smooth behavior of the cost function in the vicinity of the global minimum seems to prevent full convergence of the present optimization procedure. In principle, the SMF method should be particularly well-suited for handling non-smooth cost functions (e.g. Audet et al., 2008). Ongoing development of surrogate functions (Wang et al., 2010) and polling strategies, like MADS (Audet and Dennis, 2006) or non-cubic lattices (Belitz and Bewley, 2008), may be expected to improve the robustness of the present method, and will be tested in future work.

The principal roadblock for an application of the present inversion strategy to field-scale problems is the necessity for a forward simulation method that is reliable and efficient at very high Reynolds numbers. The direct numerical simulation approach employed in the present study can be trusted to yield reliable results, since it involves only a minimum of modeling assumptions, but it is clearly not appropriate for field-scale applications. Simulations based on depth-averaged equations are computationally inexpensive, but their reliability and robustness in the context of turbidity currents remains debatable to date. From a computational fluid dynamics perspective, large-eddy simulation (LES) and Reynolds-averaged Navier–Stokes (RANS) models represent established techniques for high Reynolds number calculations at acceptable computational cost. We plan to explore these approaches for further progress on turbidity current inversion. While large-scale 3D simulations involving turbulence closure and complex bathymetry will continue to present a difficult challenge, both in terms of computational cost and modeling accuracy, the performance of the inversion strategy is not related in any obvious way to the internal complexity of the forward model. However, more realistic forward simulations will probably involve a larger number of free parameters, which clearly will add to the expense and difficulty of the inversion.

One particular difficulty ensues from the coupling between deposits from successive flow events via erosion. Conceptually, the effect of erosion can be taken into account by inverting multiple events simultaneously, although such an approach may prove prohibitive due to the increased number of free parameters, and iterative strategies may be preferable. When actual well data is to be used as input to the inversion, significant difficulties arise from data uncertainty. Since it is rarely possible to make bed-for-bed correlations in the subsurface, a geostatistical approach may be required, rather than a deterministic one, to characterize *distributions* of turbidite properties at the well locations. Inversion could then be based on a finite number of flow archetypes, with the objective to match the probability density functions of the relevant properties at the well locations.

References

Al ja'aidi, O., McCaffrey, W., Kneller, B., 2004. Factors influencing the deposit geometry of experimental turbidity currents; implications for sand-body architecture in confined basins. In: Lomas, S., Joseph, P. (Eds.), Confined Turbidite Systems. Geological Society of London, pp. 45–58.

Alexander, J., Morris, S., 1994. Observations on experimental, nonchannelized, high-concentration turbidity currents and variations in deposits around obstacles. J. Sediment. Res. A Sediment. Petr. Proc. 64, pp. 899–909.

Audet, C., 2004. Convergence results for pattern search algorithms are tight. Optim. Eng. 5 (2), 101–122.

Audet, C., Béchard, V., Le Digabel, S., 2008. Nonsmooth optimization through mesh adaptive direct search and variable neighborhood search. J. Global Optim. 41 (2), 299–318.

Audet, C., Booker, A.J., Dennis, J.E., Frank, P.D., Moore, D.W., 2000. A surrogate-model-based method for constrained optimization. AIAA Paper 2000–4891.

Audet, C., Dennis, J.E., 2003. Analysis of generalized pattern searches. SIAM J. Optim. 13, 889–903.

Audet, C., Dennis Jr., J.E., 2006. Mesh adaptive direct search algorithms for constrained optimization. SIAM J. Optim. 17 (1), 2–11.

Baas, J., McCaffrey, W., Haughton, P., Choux, C., 2005. Coupling between suspended sediment distribution and turbulence structure in a laboratory turbidity current. J. Geophys. Res. 110, C11015.

Belitz, P., Bewley, T., 2008. A lattice-based approach to derivative-free optimization. APS DFD Meeting Abstracts, November.

Bewley, T.R., in press. Numerical Renaissance: Simulation, Optimization, and Control. Renaissance Press, San Diego, CA, USA. Preprint available at: 〈http://numerical-renaissance.com/〉.

Blanchette, F., Strauss, M., Meiburg, E., Kneller, B., Glinksy, M.E., 2005. High-resolution numerical simulations of resuspending gravity currents: conditions for self-sustainment. J. Geophys. Res. 110, C12022.

Bonnecaze, R.T., Huppert, H.E., Lister, J.R., 1993. Particle-driven gravity currents. J. Fluid Mech. 250, 339–369.

Booker, A.J., 2000. Well-conditioned kriging models for optimization of computer models. Technical Report, Mathematics and Computing 002, Boeing Phantom Works, Seattle, WA, USA.

Booker, A.J., Dennis, J.E., Frank, P.D., Serafini, D.B., Torczon, V., Trosset, M.W., 1999. A rigorous framework for optimization of expensive functions by surrogates. Struct. Optim. 17, 1–13.

Cantero, M.I., Balachandar, S., Garcia, M.H., 2007. High resolution simulations of cylindrical gravity currents. J. Fluid Mech. 590, 437–469.

Chung, H.-S., Alonso, J.J., 2002. Design of a low-boom supersonic business jet using cokriging approximation models. AIAA Paper 2002–5598.

Cox, D.D., John, S., 1997. SDO: a statistical method for global optimization. In: Alexandrov, N., Hussaini, M. (Eds.), Multidisciplinary Design Optimization: State of the Art. SIAM, Philadelphia, PA, pp. 315–329.

Davis, C., 1954. Theory of positive linear dependence. Am. J. Math. 76, 733–746.

Felix, M., 2002. Flow structure of turbidity currents. Sedimentology 49 (3), 397–419.

Fletcher, C.A., 1991. Computational Techniques For Fluid Dynamics, vol. 2. Springer, 484 pp.

Garcia, M., 1994. Depositional turbidity currents laden with poorly sorted sediment. J. Hydraul. Eng. 120, 1240–1263.

Gonzalez-Juez, E., Meiburg, E., Constantinescu, S.G., 2009. Gravity currents imping-ing on bottom-mounted square cylinders: flow fields and associated forces. J. Fluid Mech. 631, 65–102.

Gunzburger, M., 1995. Flow Control. Springer, Berlin, 381 pp.

Hall, B., 2009. Submarine channel and sediment wave formation by turbidity currents: Navier–Stokes based linear stability analysis. Ph.D. Thesis, University of California at Santa Barbara, USA, 146 pp.

Huppert, H., 2006. Gravity currents: a personal perspective. J. Fluid Mech. 554, 299–322.

Jones, D.L., 2001. A taxonomy of global optimization methods based on response surfaces. J. Global Optim. 21, 345–383.

Kneller, B., McCaffrey, W., 1995. Modelling the effects of salt-induced topography on deposition from turbidity currents. In: Travis, C., Harrison, H., Hudeac, M., Vendeville, B., Peel, F., Perkins, R. (Eds.), Salt, Sediment and Hydrocarbons. SEPM Gulf Coast Section, Houston, pp. 137–145.

Kolda, T.G., Lewis, R.M., Torczon, V., 2003. Optimization by direct search: new perspectives on some classical and modern methods. SIAM Rev. 45 (3), 385–482.

Laval, A., Cremer, M., Beghin, P., Ravenne, C., 1988. Density surges: two-dimensional experiments. Sedimentology 35, 73–84.

Lewis, R.M., Torczon, V., 1999. Pattern search algorithms for bound constrained minimization. SIAM J. Optim. 9, 1082–1099.

Lewis, R.M., Torczon, V., 2000. Pattern search methods for linearly constrained minimization. SIAM J. Optim. 10, 917–941.

Lophaven, S.N., Nielsen, H.B., Søndergaard, J., 2002. DACE: a MATLAB Kriging toolbox. Technical Report, IMM-TR-2002-12, Danmarks Tekniske Universitet, Copenhagen.

Luthi, S., 1981. Some new aspects of two-dimensional turbidity currents. Sedimen-tology 28, 97–105.

Marsden, A.L., Feinstein, J.A., Taylor, C.A., 2008. A computational framework for derivative-free optimization of cardiovascular geometries. Comput. Meth. Appl. Mech. Eng. 197 (21-24), 1890–1905.

Marsden, A.L., Wang, M., Dennis Jr., J.E., Moin, P., 2004a. Optimal aeroacoustic shape design using the surrogate management framework. Optim. Eng. 5 (2), 235–262.

Marsden, A.L., Wang, M., Dennis Jr., J.E., Moin, P., 2004b. Suppression of airfoil vortex-shedding noise via derivative-free optimization. Phys. Fluids 16 (10), L83–L86.

Marsden, A.L., Wang, M., Dennis Jr., J.E., Moin, P., 2007. Trailing-edge noise reduction using derivative-free optimization and large-eddy simulation. J. Fluid Mech. 572, 13–36.

McKay, M.D., Conover, W.J., Beckman, R.J., 1979. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 21, 239–245.

Meiburg, E., Kneller, B., 2010. Turbidity currents and their deposits. Annu. Rev. Fluid Mech. 42, 135–156.

Middleton, G., Neal, W., 1989. Experiments on the thickness of beds deposited by turbidity currents. J. Sediment. Petrol. 59, 297–307.

Mulder, T., Alexander, J., 2001. Abrupt change in slope causes variation in the deposit thickness of concentrated particle-driven density currents. Mar. Geol. 175, 221–235.

Necker, F., Härtel, C., Kleiser, L., Meiburg, E., 2002. High-resolution simulations of particle-driven gravity currents. Int. J. Multiphase Flow 28, 279–300.

Necker, F., Härtel, C., Kleiser, L., Meiburg, E., 2005. Mixing and dissipation in particle-driven gravity currents. J. Fluid Mech. 545, 339–372.

Osher, S., Fedkiw, R., 2003. Level Set Methods and Dynamic Implicit Surfaces. Springer, New York, 296 pp.

Parsons, J., Schweller, W., Stelting, C., Southard, J., Lyons, W., Grotzinger, J., 2002. A preliminary experimental study of turbidite fan deposits. J. Sediment. Res. 72, 619–628.

Parsons, J.D., Friedrichs, C.T., Traykovski, P.A., Mohrig, D., et al., 2007. The mechanics of marine sediment gravity flows. In: Nittrouer, C., Austin, J., Field, M., Kravitz, J., Syvitski, J., Wiberg, P. (Eds.), Continental Margin Sedimentation: From Sediment Transport to Sequence Stratigraphy. Wiley & Sons, pp. 275–337.

Ruith, M.R., Chen, P., Meiburg, E., 2004. Development of boundary conditions for direct numerical simulations of three-dimensional vortex breakdown phenom-ena in semi-infinite domains. Comput. Fluids 33, 1225–1250.

Serafini, D.B., 1998. A framework for managing models in optimization for computationally expensive functions. Ph.D. Thesis, Rice University, Houston, TX, USA, 188 pp.

Violet, J., Sheets, B., Pratson, L., Paola, C., Beauboeuf, R., Parker, G., 2006. Experiment on turbidity currents and their deposits in a model 3d subsiding minibasin. J. Sediment. Res. 75, 820–843.

Wang, Q., Moin, P., Iaccarino, G., 2010. A high order multivariate approximation scheme for scattered data sets. J. Comput. Phys. 229 (18), 6343–6361.

Weimer, P., Slatt, R.M., 2007. Introduction to the petroleum geology of deepwater setting. In: AAPG Studies in Geology. American Association of Petroleum Geologists.