


Provided for non-commercial research and education use.  
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



ELSEVIER

Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

Environmental Modelling &amp; Software 22 (2007) 1691–1704

Environmental  
Modelling & Software[www.elsevier.com/locate/envsoft](http://www.elsevier.com/locate/envsoft)

# A framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO)

Joshua B. Kollat, Patrick Reed\*

*The Pennsylvania State University, Department of Civil and Environmental Engineering, 212 Sackett Building, University Park, PA 16802-1408, USA*

Received 2 November 2006; received in revised form 30 January 2007; accepted 3 February 2007

Available online 30 March 2007

## Abstract

This study presents a framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO). The VIDEO framework allows users to visually navigate large multi-objective solution sets while aiding decision makers in identifying one or more optimal designs. Specifically, the interactive visualization framework is intended to provide an innovative exploration tool for high-order Pareto-optimal solution sets (i.e., solution sets for three or more objectives). The framework is demonstrated for a long-term ground-water monitoring (LTM) application in which users can explore and visualize tradeoffs for up to four design objectives, simultaneously. Interactive functionality within the framework allows the user to select solutions within the objective space and visualize the corresponding monitoring plan's performance in the design space. This functionality provides the user with a holistic picture of the information provided by a particular solution, ultimately allowing them to make a more informed decision. In addition, the ease with which the framework allows users to navigate and compare solutions as well as design tradeoffs leads to a time efficient analysis, even when there are thousands of potential solutions.

© 2007 Elsevier Ltd. All rights reserved.

*Keywords:* Visualization; Decision support; Monitoring design; Multi-objective optimization; Genetic algorithms; Kriging

## 1. Introduction

The goal of this study is to demonstrate the need for and value of coupling advanced visualization and multi-objective evolutionary algorithms to enhance environmental applications. In support of this goal, the VIDEO (Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization) framework is a new software that enables decision makers to explore large multi-objective solution sets while aiding in the selection of one or more optimal designs. Development of the VIDEO framework was motivated by the increasing emphasis of environmental systems research on multi-objective methods (Chankong and Haimes, 1983; Keeney and Raiffa, 1993; Haimes, 1998; Coello Coello et al., 2002).

Recent innovations in multi-objective evolutionary algorithms (MOEAs) have served to catalyze the development of a broad suite of multi-objective design and decision support methodologies within the environmental and water resources literature (Horn and Nafpliotis, 1993; Ritzel et al., 1994; Cieniawski et al., 1995; Halhal et al., 1997; Loughlin et al., 2000; Reed et al., 2001; Erickson et al., 2002; Reed and Minsker, 2004; Bekele and Nicklow, 2005; Farmani et al., 2005; Tang et al., 2007). MOEA's population-based search allows users to find entire tradeoff (or Pareto) surfaces using a single algorithm run for large, complex problem spaces (Goldberg, 1989; Salomon, 1998; Back et al., 2000).

In general, optimality for multiple objectives is defined by the set of solutions that are globally (locally) better than all other solutions in at least one objective and are termed Pareto-optimal (non-dominated) solutions (Pareto, 1896). The Pareto-optimal front (or non-dominated front) is obtained by plotting these solutions according to their objective values yielding an  $N - 1$  dimensional surface where  $N$  is equal to the total number of

\* Corresponding author. Tel.: +1 814 863 2940; fax: +1 814 863 7304.  
E-mail address: [preed@engr.psu.edu](mailto:preed@engr.psu.edu) (P. Reed).

objectives. High-order Pareto optimization problems as defined by Reed and Minsker (2004) have three or more objectives. They have also been termed “many-objective” problems within the operations research (OR) literature (Fleming et al., 2005). There is a growing body of literature exploring the challenges and benefits associated with solving high-order Pareto optimization problems (Farina and Amato, 2002; Kumar and Ranjithan, 2002; Reed and Minsker, 2004; Bekele and Nicklow, 2005; Fleming et al., 2005; Tang et al., 2006). Please note that in this paper, we use the general term “Pareto-optimal” instead of “non-dominated” when describing approximate multi-objective solution sets because the term is generally appropriate for the best known solution sets for an application.

Although our ability to solve high-order Pareto optimization problems using techniques such as MOEAs reflects more recent computational advances, the systems analysis literature from the past 30 years has recognized that design problems often have several conflicting objectives, especially in the environmental arena (Haimes, 1998). As highlighted by Haimes (1998, pp. 17–18), modern systems research has its origins dating back to Wiener's (1961) text *Cybernetics*, which successfully forecasted a myriad of technological advances that have shaped human–computer interaction as well as design/decision-making processes. Classical multi-objective decision-making methodologies (Chankong and Haimes, 1983; de Neufville, 1990) have generally focused on transformation methods that commensurate multiple objectives into new formulations that are amendable to solution using single objective optimization algorithms. MOEAs represent a break from traditional multi-criteria operations research where it is no longer necessary to transform or commensurate problem structures into single objective formulations and instead new decision support tools can focus on the structure and content of Pareto surfaces.

This paper seeks to: (1) affirm the need and value of combining interactive visualization with high-order Pareto optimization for improved a posteriori decision-making (see Section 2.1) and (2) provide a specific demonstration of the VIDEO framework within an illustrative long-term groundwater monitoring (LTM) design application (see Section 2.2). This study builds on Reed and Minsker (2004) by contributing a visualization framework tailored specifically to many-objective LTM applications. LTM design has long been recognized to have “many objectives” (Moss, 1979) and is an excellent example application for showing that visualization combined with high-order multi-objective solution sets can facilitate discovery and negotiation in the design and decision-making process (Reed and Minsker, 2004). Our use of the terms discovery and negotiation is motivated by the potential of high-order multi-objective solution sets to generate alternatives that capture a broad suite of system behaviors relevant to both modeled and unmodeled objectives (Loughlin et al., 2001), helping decision makers to “discover” system dependencies and/or tradeoffs and exploit this information in the negotiated selection of a solution (Castelletti and Soncini-Sessa, 2006).

In the remainder of the paper, Section 2 provides a more detailed discussion of a posteriori decision-making, the VIDEO framework's components, and the LTM application being used

to demonstrate the framework. Section 3 provides illustrative results for the LTM case study demonstrating how the VIDEO framework facilitates exploration of tradeoffs and negotiated focus on specific solutions. Sections 4 and 5 discuss the implications of using interactive visualization to facilitate a posteriori decision-making.

## 2. Methodology

### 2.1. An overview of a priori and a posteriori multi-objective methods

A priori methods (Coello Coello et al., 2002) seek to model decision maker preference before searching for designs/decisions. A classic example of an a priori method is the “normative” decision-making methodology developed by Keeney and Raiffa (1976, 1993) termed multi-attribute utility analysis (MAUA). MAUA requires extensive surveys of decision maker preferences prior to searching for potential system designs. In MAUA, decision maker surveys must be designed carefully to assure methodological assumptions are satisfied [i.e., preferential and utility independence (for more details see de Neufville, 1990; Keeney and Raiffa, 1993)]. MAUA is an example of a transformative method where the original decision objectives are analytically represented within a utility function that can be optimized using traditional single objective algorithms. In MAUA, all search and decision-making is then in reference to utility. a priori methods such as MAUA have been criticized because they do not condition decision maker preference on potential alternatives, they suffer from decision maker contradictions (or intransitivity), and utility representations of preference are non-unique for groups (i.e., Arrow's Paradox, Arrow, 1963).

In contrast to a priori methods, the VIDEO framework represents an a posteriori decision tool where decision maker preferences for alternatives are expressed after non-dominated or Pareto-optimal alternatives have been identified. In this process, all solutions are initially assumed to have equal preference during the search process. Decision maker preferences are then expressed in the exploration tradeoffs and selection of design solutions. a posteriori decision tools have been criticized in the OR literature due to (1) the mathematical complexity of finding tradeoff solutions and (2) the contention that large solution sets tend to overwhelm/confuse decision makers while providing limited insights into their design preferences (Haimes, 1998; Coello Coello et al., 2002; Zeleny, 2005).

In reference to the criticisms of a posteriori decision-making methods, MOEAs have significantly enhanced our ability to search for and quantify large multi-objective solution sets in the environmental area (Bekele and Nicklow, 2005; Kollat and Reed, 2006; Reed et al., 2007; Tang et al., 2007). Moreover, the second criticism's implication that large solution sets overwhelm decision makers assumes that system expertise and visualizations have limited value. Multi-objective engineering and environmental applications represent complex systems that can be meaningfully visualized in space and/or time. Exploring tradeoffs and visualizing system performance can be useful in two spaces: (1) objective space and (2) design/system space. As an example, a pollution remediation system can have a suite of objectives (minimize cost, maximize reliability, maximize resiliency, minimize redundancy, etc.) each of which shapes the potential real-world system's spatiotemporal performance. There is an increasing number of studies demonstrating that visualization combined with optimization can promote design innovations and provide decision makers with an improved understanding of system behaviors (Mecham, 1997; Balling, 1999; Winer and Bloebaum, 2002; Stump et al., 2003; Reed and Minsker, 2004; Fleming et al., 2005). In the context of this prior work, the VIDEO framework contributes a highly interactive environment for exploring LTM design tradeoffs for up to four objectives and detailed analysis of their consequences in the resultant design space.

### 2.2. VIDEO framework components

#### 2.2.1. Overview of framework

The core of the VIDEO framework has been developed using the Python programming language. Python is a dynamic, object-oriented scripting

language, which integrates easily with other languages, and is virtually platform independent (Chun, 2006). Additionally, Python has a large number of graphical user interface (GUI) frameworks, including Tkinter, PyQt, and wxPython. The VIDEO framework has been developed using Tkinter (Python's standard GUI package), which is built on top of Tcl/Tk and is portable across Windows, Unix, and Mac platforms (Flynt, 2003). Aside from Python's ease-of-use and GUI development capabilities, it allows easy integration with other, lower level (i.e., faster) languages such as C and C++. The VIDEO framework takes advantage of this capability by utilizing the Visualization ToolKit (VTK) (Schroeder, 2001) which is rooted in C++ for the visualization components of the framework, and a spatial interpolator (KT3D) (Deutsch and Journel, 1998) that has been translated from Fortran to C by the authors. The Visualization ToolKit is an open-source software system developed for three-dimensional visualization and image processing developed in C++ and wrapped with various other common programming languages (namely Java, Python, and Tcl/Tk). One of the main strengths of VTK is its fast visualization rendering and excellent interaction capabilities. KT3D has been utilized in the initial version of the VIDEO framework to provide high quality (interpolation) data sets. KT3D has been coded in C and is called by Python to perform the required spatial interpolation on-the-fly during interaction events such as when a user is selecting and comparing solutions. When Pareto-optimal solution sets are large with potentially thousands of solutions (as is the case with the LTM design problem explored in this study), it is more efficient for gridded spatial data sets to be generated on-demand.

The VIDEO decision-making software is divided into two primary components, an objective space window and a decision space window (see Fig. 1). The objective space window contains the visualization data associated with the design objectives of the problem, as well as tools that allow the user to

easily manipulate how objective tradeoffs can be visualized. For example, the user can quickly change which of the objectives are displayed, how they are displayed (i.e., plotting axes and color representations) as well as change the display precision of the solutions, and the plotting limits. The decision space component of the software displays the design space associated with the problem and is therefore more problem specific. The decision space component is linked with the objective space component, allowing the user to select specific solutions from within the objective space, and see what these solutions actually represent in terms of a design. Probing tools are provided within the decision space component to allow the user to explore the implications associated with the tradeoffs between their design objectives. These components and tools are discussed in more detail in Section 2.2.3.

### 2.2.2. Generation of multi-objective solution sets

The VIDEO framework is an a posteriori decision support tool that has been designed with the intent of promoting exploration of multi-objective solution sets generated by multi-objective evolutionary algorithms. However, any solution set ranging from lower quality initial approximations generated using Monte Carlo analysis to true Pareto-optimal solutions sets generated through enumeration of every potential solution can be explored within the software. In this study, a 25-well LTM application is used as the test case to demonstrate the VIDEO framework (Kollat and Reed, 2006, 2007). The authors have successfully obtained high quality solution sets for the 25-well LTM application using the Epsilon-Non-dominated Sorted Genetic Algorithm II ( $\epsilon$ -NSGAII) (Kollat and Reed, 2007; Tang et al., 2007), which has been shown to perform as well or better than other top-performing MOEAs (Kollat and Reed, 2006; Tang et al., 2006). These previous studies have motivated our interest in advancing a posteriori decision-making using visualization in combination with large Pareto-optimal sets.

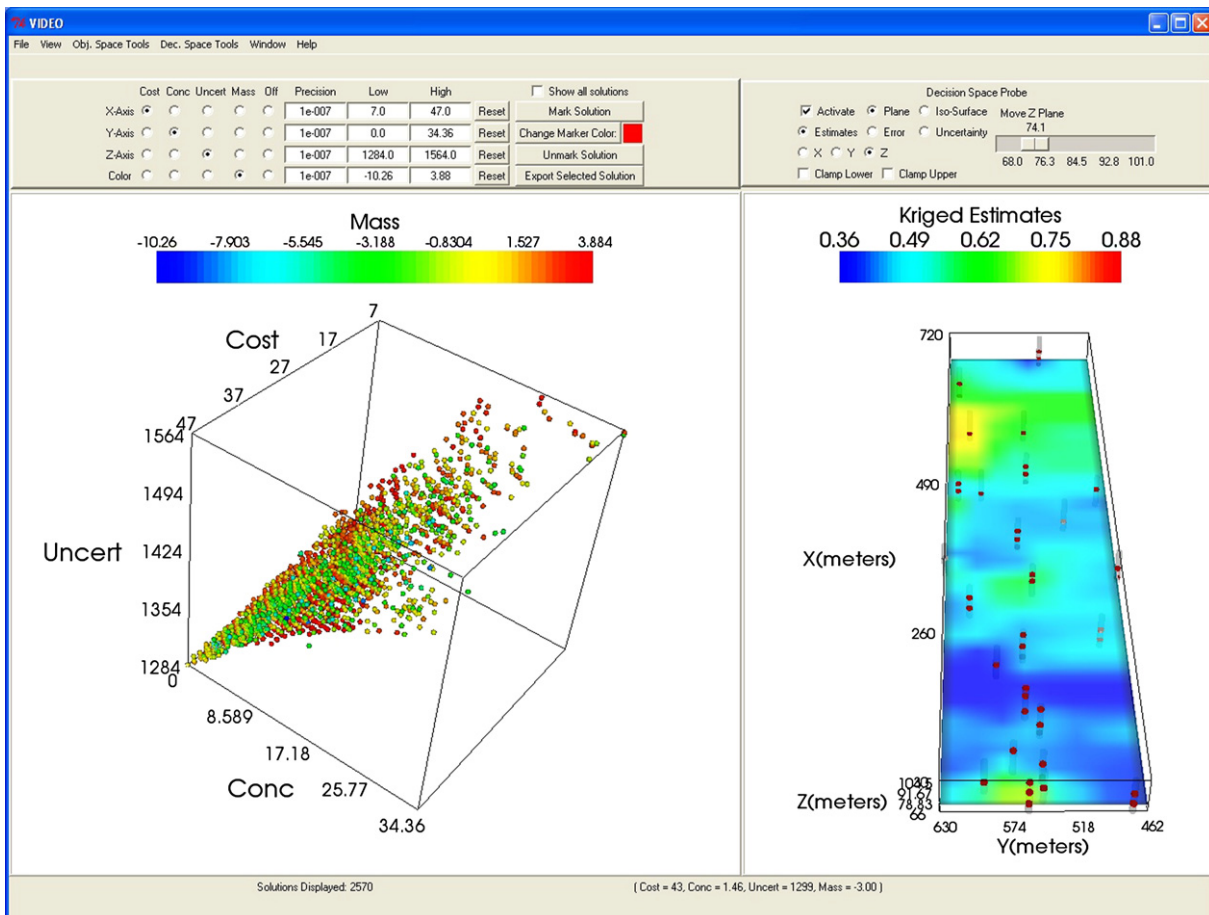


Fig. 1. Overview of VIDEO decision-making software. The interface is divided into an objective space window, a design space window, and a set of tools associated with each of these windows.



### 2.2.3. Interactive visualization

**2.2.3.1. Objective space visualization.** The objective space component of the VIDEO framework currently allows the user to visualize up to four design objectives simultaneously. Up to three design objectives can be plotted on the spatial coordinate axes ( $X$ ,  $Y$ , and  $Z$ ), and the fourth objective can be portrayed using color. Each potential solution is represented as a spherical glyph positioned at the appropriate coordinates and is colored based on a fourth objective. Future versions of the software may allow for visualizing additional objectives through the manipulation of glyph size, shape, or orientation (Balling, 1999). Located immediately above the objective space window are a set of radio buttons that allow the user to change which of the objectives are plotted on each of the coordinate axes and which objective is represented using glyph color. These radio buttons are arranged in a grid fashion with each row of the grid representing the  $X$ ,  $Y$ ,  $Z$ , and color axis, respectively, and each column of the grid representing design objectives 1–4. An additional “off” column is provided to allow the user to turn one or more plotting axes off. This GUI design allows the user to efficiently explore the wide range of plotting possibilities within the objective space, ultimately enhancing their understanding of the problem. For example, by specifying only two plotting axes and switching the others off, two objective subsets of the larger four objectives problem can be plotted to quickly identify tradeoffs between objectives. As another example, glyph color can be utilized as a third objective while two other objectives are plotted on the spatial coordinates, possibly providing enhanced insight into the many relationships that may exist between the design objectives. Alternative plotting techniques are demonstrated in more detail for the four-objective LTM test case in Section 3.1.

**2.2.3.2. Objective space navigation.** The VIDEO framework provides the user with two means of focusing on areas of interest within the objective space. The first is a set of text boxes associated with each of the plotting axes ( $X$ ,  $Y$ ,  $Z$ , and color), which allows the user to numerically specify the plotting limits of each objective. This is useful if the user has problem specific knowledge that allows them to specify thresholds of interest for each of the design objectives. In addition, an interactive tool has been included that allows the user to dynamically resize a plotting box within the actual objective space window in order to focus on sub-regions of interest quickly and efficiently. This interactive tool can be instantiated by pressing “i” while the objective space rendering window is active, or by selecting “Interactive Limits Adjustment” under the “Objective Space Tools” menu.

**2.2.3.3. Objective space thinning.** Previous studies have emphasized the need for understanding objective precision requirements in multi-objective optimization (Laumanns et al., 2002; Kollat and Reed, 2007; Reed et al., 2007). For example, should the concentration estimates of a system be quantified in parts per trillion (ppt), or is parts per million (ppm) a sufficient level of precision? Design objective precision specification has been added into the VIDEO framework so that high resolution Pareto-optimal solution sets can be “thinned” based on the decision maker’s precision requirements. The VIDEO framework uses  $\epsilon$ -dominance (Laumanns et al., 2002) to reduce the precision of the Pareto-optimal solution set. In this approach,  $\epsilon$  values reflecting the required precision for each design objective are specified by the decision maker. Non-dominance sorting is then performed based on the reduced precision rather than the full precision, ultimately resulting in a reduced precision set. For a detailed discussion and analysis of  $\epsilon$ -dominance applied to a LTM problem and its implications for reducing computational demand, please refer to a recent study by Kollat and Reed (2007). Within the VIDEO framework,  $\epsilon$ -dominance values can be specified for each objective, and the solutions displayed in the objective space will be updated to reflect the decision maker’s required precision.

**2.2.3.4. Interactive solution selection.** To facilitate comparison of Pareto-optimal solutions, the user is provided with the ability to manually pick solutions in the objective space window at the click of a mouse button. When the mouse pointer is navigated over the objective space visualization window, a left mouse button click will select the glyph which is underneath the mouse pointer. Once the user has selected a solution, it is highlighted by a bounding box and several events can then occur. First and foremost, the design

associated with the selected solution is displayed in the decision space window. This in and of itself is a very useful way to visually correlate how regions of the objective space map to various types of designs. In addition to displaying the selected solution’s corresponding design, the user also has the ability to extensively probe the data associated with that solution. Users can also “mark” a selected solution, doubling its glyph size and permanently associating a bounding box with the solution. Marked solutions are always displayed in the visualization window until they are unmarked by the user. This allows the user to track various solutions of interest throughout the interactive decision-making process. Solutions which have been marked can subsequently be exported to a file, which can then be read back into the software at a later time to facilitate further comparison.

**2.2.3.5. Decision space probing.** As mentioned previously, the decision space component of the VIDEO framework is highly problem specific. However, the demonstration version of the framework utilizes an LTM design problem where the objectives are to effectively monitor a contamination plume using a pre-determined set of well locations while minimizing system cost, error, and uncertainty. Specific details of the LTM test case used in this study are described in Section 2.3. The VIDEO framework represents the LTM decision space by plotting the available well locations throughout the sampling domain. Each well location is represented by a transparent cylinder, and within each cylinder, the actual sampling locations available along the well’s vertical axis are marked. When a user selects a solution in the objective space, the wells which are associated with the selected solution are highlighted in red. If the decision space probing functionality is activated, the selected solution will be Kriged on-the-fly to produce maps of concentration estimates, estimation error, and estimation uncertainty throughout the sampling domain.

For the LTM test case examined in this study, Quantile Kriging was used to obtain contaminant estimates throughout the sampling domain by using the data at the sampled locations. Estimation error was calculated by comparing the estimates obtained when all sampling locations are utilized, to the data obtained for a particular Pareto-optimal solution in which only a subset of the data points are utilized. In addition, since Kriging is a minimum error variance estimator, the local uncertainty associated with each estimate throughout the sampling domain is also available.

Following Kriging, the user is provided with the ability to interactively move a probing plane throughout the sampling domain in order to explore the implications of the design. When using a plane probe, the user can choose an  $X$ – $Y$ ,  $X$ – $Z$ , or  $Y$ – $Z$  plane depending on the geometry of the decision space domain. This plane can then be moved throughout the domain either through the use of a slider bar (available in the decision space tool frame) or by manually moving the plane throughout the domain using the mouse. The plane will display the underlying data by using red to represent areas of highest estimates, error, or uncertainty (depending on which is selected) and blue to represent areas of lowest estimates, error, or uncertainty. Data associated with different objective space solutions can be directly compared using the plane probe by setting the plane position to some location in the design domain, and subsequently selecting and comparing various solutions in the objective space.

An additional probing option allows the user to plot and manipulate a colored iso-surface representing a constant data value throughout the domain. The data value, which the iso-surface represents, can be controlled through the use of a slider bar, and the color of the iso-surface changes with changing data values based on the data range (i.e., red represents high values and blue represents low values). While a probing plane provides a 2D slice of the decision space for all data values, an iso-surface provides a volumetric view of the decision space at a constant data value. One useful way in which to utilize the iso-surface functionality is to explore how the iso-surfaces change at a given data value for various objective space solutions.

The color mapping of both the plane probe and the iso-surface can be normalized to a minimum and maximum data value through the use of the “Clamp Lower” and “Clamp Upper” color limit check boxes. This ensures that the color scale used to represent the underlying data is normalized, allowing for efficient and accurate comparison between the maps generated by various solutions. However, the solutions which should be used to clamp the color limits are highly problem specific, and often difficult to identify. Taking this into consideration, the ability to manually clamp the color limits has been

included within the VIDEO framework. Clamping the color limits ensures that when navigating the objective space, probed data in the decision space is always represented on the same color scale, making direct comparison between various solutions fast and accurate.

### 2.3. Monitoring test case description

The LTM test case used to demonstrate the VIDEO framework represents a simulated perchloroethylene (PCE) contamination plume originating from an underground storage tank. The contaminant plume is based on a hypothetical 50 million node flow and transport simulation (Maxwell et al., 2000) through actual hydrogeology located at the Lawrence Livermore National Lab (LLNL) in Livermore, California. The site hydrogeology has been comprehensively characterized so that the 1-m resolution simulation reflects a highly heterogeneous alluvial aquifer. PCE concentration data are available at 25 pre-determined well sampling locations and there are 1–3 sampling points available along each of the well's vertical axes. However, in this study, it is assumed that if a well is sampled, then all available locations along its vertical axis are sampled, yielding a total of 25 decision variables (or  $2^{25}$  possible well sampling schemes). The sampling domain extends 690 m in the  $X$  direction, 168 m in the  $Y$  direction, and 38.5 m in the  $Z$  direction. Additional details of the LTM test case can be found in Reed et al. (2004).

#### 2.3.1. Design objectives

Four design objectives were chosen for the LTM test case, each of which was minimized. The design objectives included Cost, Concentration Error (Conc), Uncertainty (Uncert), and Mass Error (Mass). The “Cost” objective reflects the normalized cost of sampling the system. There are 25-well locations with 1–3 sampling locations at each well for a total of 47 available sampling locations. This implies that the “Cost” objective can range from zero (no locations sampled) to 47 (all locations sampled). The “Conc” objective reflects the concentration estimation error between a Kriged map of the plume utilizing all available sampling locations, and a Kriged map of the plume which utilizes a subset of the sampling locations. The “Uncert” objective reflects the uncertainty associated with the Kriged map of the contaminant plume by summing the estimation variances attained for each estimation location. The “Mass” objective reflects the error between the total mass of PCE estimated by Kriging the domain based on all available well sampling locations, and the estimated mass of PCE obtained by Kriging the domain based on a subset of well sampling locations. Readers interested in the actual equations used to quantify each of these objectives can refer the publications Kollat and Reed (2006, 2007).

#### 2.3.2. Spatial interpolation

PCE concentration estimates were obtained at unsampled locations throughout the sampling domain of the LTM test case using Quantile Kriging (QK). Kriging provides a minimum error variance estimate of contaminant concentration at an unsampled location provided the data at the sampled locations. Quantile Kriging has been chosen in this study based on its effectiveness in providing high quality plume interpolations despite highly variable PCE concentrations and preferential sampling (Reed et al., 2004). QK extends Ordinary Kriging by transforming contaminant concentrations to quantile space according to their rank ordering to perform the spatial interpolation (Journel and Deutsch, 1997). This is done to reduce the influence of extreme concentration values on the mean and variance of the data. KT3D, a three-dimensional Kriging library included as part of the GSLIB software package available in Fortran (Deutsch and Journel, 1998), was used to perform the Kriging for this study. A C translation of KT3D developed by the authors has been integrated into the VIDEO framework to produce on-the-fly spatial interpolations of the LTM decision space.

Before the Kriging is performed, a variogram analysis is conducted to reveal the spatial correlation structure of the data, and a model is chosen which will best represent this structure when computing estimates (Goovaerts, 1997). The variogram model chosen for this study was a spherical model with nugget = 0.005 and range = 100 m. In this study, estimates were obtained across a grid defined by 34 blocks in the  $X$  domain, 7 blocks in the  $Y$  domain, and 7 blocks in the  $Z$  domain, for a total of 1666 estimate locations. QK

assumes a locally stationary concentration mean within local estimation neighborhoods at each grid location. In this study, the shape of the search neighborhood was an ellipsoid where the closest 24 data points within the search neighborhood were used in estimation calculations. An octant search was used to help ensure that the closest data points were well distributed about the estimation point (for more details see Deutsch and Journel, 1998). This is especially important since the wells for this test case contained multiple sampling locations along their axes.

#### 2.3.3. Generating the Pareto-set

A very close approximation to the true Pareto-optimal solution set was generated for the LTM application presented in this paper using a Master-Slave parallelization of the  $\epsilon$ -NSGAI (Tang et al., 2007). A constraint was placed on the LTM test case such that if a well sampling scheme contained too few sampling points to fully Krige the entire sampling domain (as defined by the Kriging parameters described in Section 2.3.2), the scheme was considered infeasible and was penalized by the algorithm similarly to Kollat and Reed (2006). The  $\epsilon$ -NSGAI was permitted to run for 3.2 million function evaluations and a Pareto-optimal approximation set containing 2570 solutions was generated. Previous studies and experience indicate that this approximation set very likely represents at least 90% of the true Pareto-optimal solution set (Tang et al., 2007). The Pareto-set generated by the  $\epsilon$ -NSGAI contained solutions with “Cost” ranging from 7 to 47, “Conc” ranging from 0 to 34.4, “Uncert” ranging from 1284 to 1564, and “Mass” ranging from  $-10.3$  to 3.9 (negative values result because the mass error objective is scaled using a log transform). This Pareto-optimal solution set is provided for demonstration purposes within the initial version of the VIDEO framework and is subsequently used to demonstrate the framework throughout the remainder of this paper. The VIDEO framework has been designed to complement MOEAs or any other Pareto optimization algorithms.

## 3. VIDEO case study

### 3.1. Exploring and understanding objective tradeoffs

Visualizing “many-objective” Pareto-optimal surfaces can be facilitated by representing three objectives using the traditional spatial coordinates, and additional objectives by changing representation characteristics. Representation of additional objectives can be done in a number of ways including the use of color, shape, size, and orientation (Balling, 1999; Stump et al., 2003). For example, in the VIDEO framework, a fourth design objective is portrayed through the use of color, where blue indicates low objective values and red indicates high objective values. The VIDEO framework's visualization of the four-objective Pareto-optimal set is shown in Fig. 2A, with the “Cost”, “Conc”, and “Uncert” objectives plotted on the  $X$ ,  $Y$ , and  $Z$  axes, respectively, and the “Mass” objective plotted using color. However, as noted in Section 2.2.3, this can easily be changed through the manipulation of the plotting axes radio buttons. For example, starting with the configuration shown in Fig. 2A, simply switching the “Mass” objective to the  $Z$ -axis, and the “Uncert” objective to the color axis (see Fig. 2B) results in a configuration which provides new insight into the tradeoffs between the objectives because trends in color emphasize the “Uncert” objective.

It is interesting to note at this point that by solving for the high-order Pareto-optimal surface for four objectives, all of the sub-problems based on three objectives, two objectives, or even a single objective are implicitly solved at the same time. Non-domination sorting can be performed on the

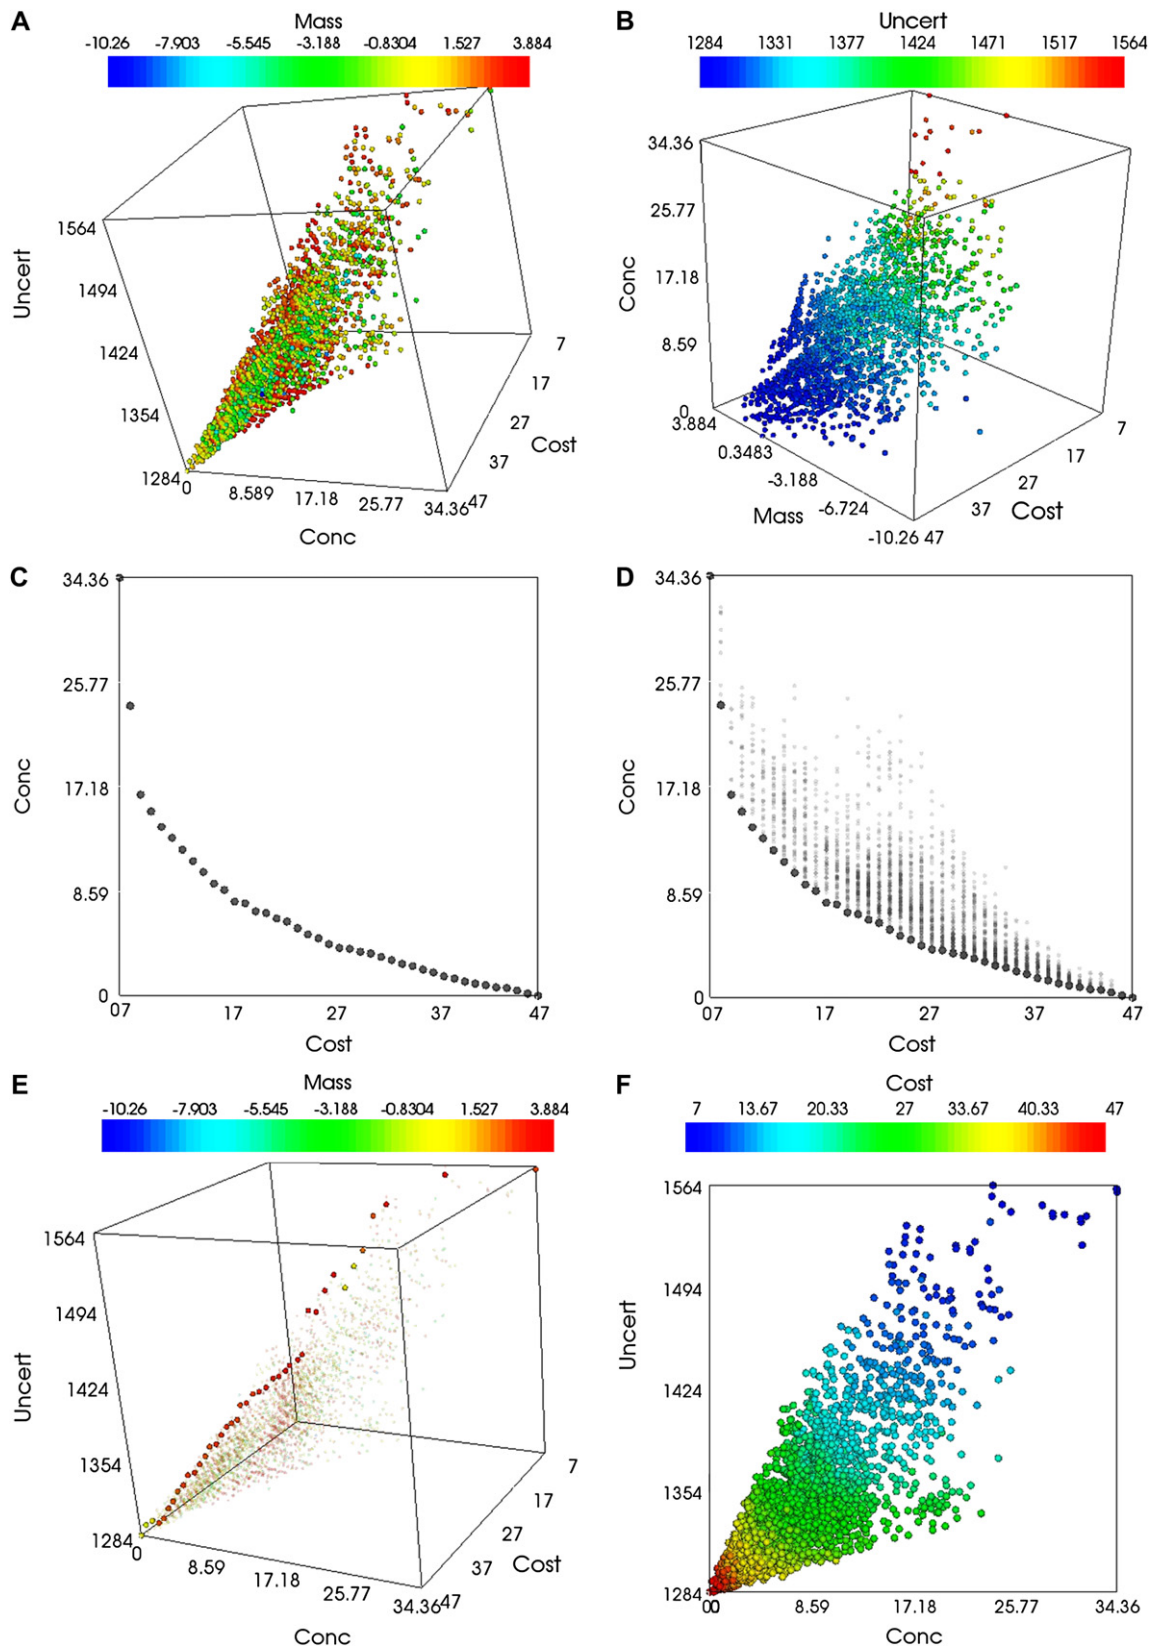


Fig. 2. Various examples of how a four-objective Pareto-optimal solution set can be viewed within the VIDEO framework. A shows the full four-objective Pareto set with “Cost”, “Conc”, and “Uncert”, plotted on X, Y, and Z, and “Mass” plotted as color. B reverses the “Uncert” and “Mass” objectives such that “Mass” is plotted on Z and “Uncert” is plotted using color. C shows the two-objective subset Cost–Conc and D and E show this subset in the context of the full four-objective set. F shows how a third objective (Cost) can be plotted using color to reveal additional information and insight.



high-order Pareto-surface in terms of fewer objective subsets. For example, the two-objective tradeoff between “Cost” and “Conc” is shown in Fig. 2C and is obtained by simply isolating solutions which dominate the full, high-order Pareto-set with respect to these two objectives, but are non-dominated with respect to one another. Fig. 2D shows the “Cost” versus “Conc” non-dominated subset highlighted with the remaining solutions composing the full four-objective Pareto-set displayed using reduced marker sizes with transparency added (both features which are available within the VIDEO framework). When manipulating the plotting axes within the VIDEO framework, the user can choose whether the whole high-order Pareto-set should be displayed (Fig. 2D), or the non-dominated subset corresponding with the objectives selected for plotting (Fig. 2C) by simply selecting a check box which tells the software to display the full set, with the current non-dominated set highlighted. It is also interesting to note that non-dominated subsets can be viewed in the context of the full four-objective space. For example, the “Cost” versus “Conc” Pareto subset shown in Fig. 2C and D can be viewed in the “Cost” versus “Conc” versus “Uncert” versus “Mass” objective space (see Fig. 2E). This can aid in revealing trade-off relationships between objective subsets, and the high-order Pareto-optimal space.

Other interesting ways of viewing the Pareto-optimal surface are illustrated in the following examples. The three-objective Pareto-front associated with “Cost”, “Conc”, and “Uncert” can be viewed by simply turning the “Mass” objective off. If viewing the three-dimensional surface associated with this three-objective Pareto-front is difficult, two of the objective can be plotted on the  $X$  and  $Y$  axes, and the third objective portrayed using color. Any combination of objectives can be explored in this manner, and color can always be used to represent any objective. As another example, if “Conc” and “Uncert” are plotted on the  $X$  and  $Y$  axes, respectively, and “Cost” is plotted using color (see Fig. 2F), the relative distributions of solutions at each cost level can be viewed very easily.

### 3.2. Exploring and understanding the design space

Before one begins exploring the LTM decision space, it is prudent to identify likely solutions with which the color limits should be clamped so that the color scale used to probe the decision space is normalized (see “Decision Space Probing” in Section 2.2.3). For the LTM test case explored in this study, it is anticipated that the highest “Cost” solution (i.e., the solution which samples from all available locations) will generally result in the lowest “Conc”, “Uncert”, and “Mass” objective values. Thus, the lower color limit was clamped to the “Cost” = 47 solution. Likewise, the upper color limit was clamped to a “Cost” = 7 solution, which is anticipated to represent the highest “Conc”, “Uncert”, and “Mass” objective values.

After normalizing the color limits, the best place to begin in understanding the LTM test case is to examine the highest cost solution because this solution will theoretically provide

the most accurate picture of the contamination plume. The objective values associated with this solution are “Cost” = 47, “Conc” = 0.0, “Uncert” = 1284, and “Mass” = 0.0. The highest cost solution is marked in Fig. 3A (remember that the user can simply click on the solution within the objective space window to make a selection). The Kriged maps associated with the high cost solution are shown in Fig. 3B through E. Fig. 3B shows the quantile ranked estimates of the contamination plume using the plane probe at  $z = 85$  m. In this figure, we can clearly see the contaminant source at the southern boundary of the domain. At this elevation, there is also a region of high concentration toward the center of the sampling domain and a region of high concentration at the northwest boundary of the sampling domain. It is useful to note at this point that the plane probe can be moved throughout the sampling domain. Its current location ( $z = 85$  m) was chosen for illustrative purposes because it provides a good picture of the contaminant source, the general shape of the plume, and several other areas of high concentration that may be of interest. Fig. 3C displays an iso-surface plotted at quantile concentration estimates of 0.70. Again, this data value was chosen somewhat arbitrarily, but to a large extent because it represents a relatively high quantile estimate (or spatial volumes containing high PCE concentrations). The estimation error is shown in Fig. 3D at the same elevation used in B. The estimation error is quantified as the difference between the estimates obtained by sampling from all locations, and the estimate obtained using a subset of sampling locations. Since the highest cost solution (i.e., all locations are sampled) has been selected in this case, the error is zero throughout the domain. The estimation uncertainty is shown in Fig. 3E at  $z = 85$  m. In this figure, the uncertainty is very low near the contamination source at the southern boundary of the domain (because of the preferential sampling which occurs in this region). In addition, the uncertainty is low around each sampling location because Kriging is an unbiased estimator (meaning that the estimates are true to the data values at the sampled locations). The area of highest uncertainty occurs at the north-east boundary of the domain where there are no well sampling locations available.

In order to demonstrate the VIDEO framework’s ability to effectively combine visualization and interaction, a compromise solution is selected from the remaining 2569 Pareto-optimal designs. As shown in Fig. 2B, plotting “Uncert” as color and “Mass” on the  $Z$ -axis emphasizes “Uncert” trends which may not otherwise be apparent. This plotting configuration was utilized in Fig. 3E to select a compromise solution for comparison to the highest cost solution selected previously. The interactive plotting limits window was adjusted to display solutions between “Cost” = 16 to 32, “Conc” = 3.3 to 17.7, and “Mass” = -3.5 to 3.9. In addition, the precision adjustment feature was used to “thin” the objective space, making it easier to locate a compromise solution. The precision settings for each objective were “Cost” = 2.0, “Conc” = 2.0, “Uncert” = 5.0, and “Mass” = 2.0. These settings were selected to thin the space based on the nominal ranges of each design objective. After selecting and comparing several



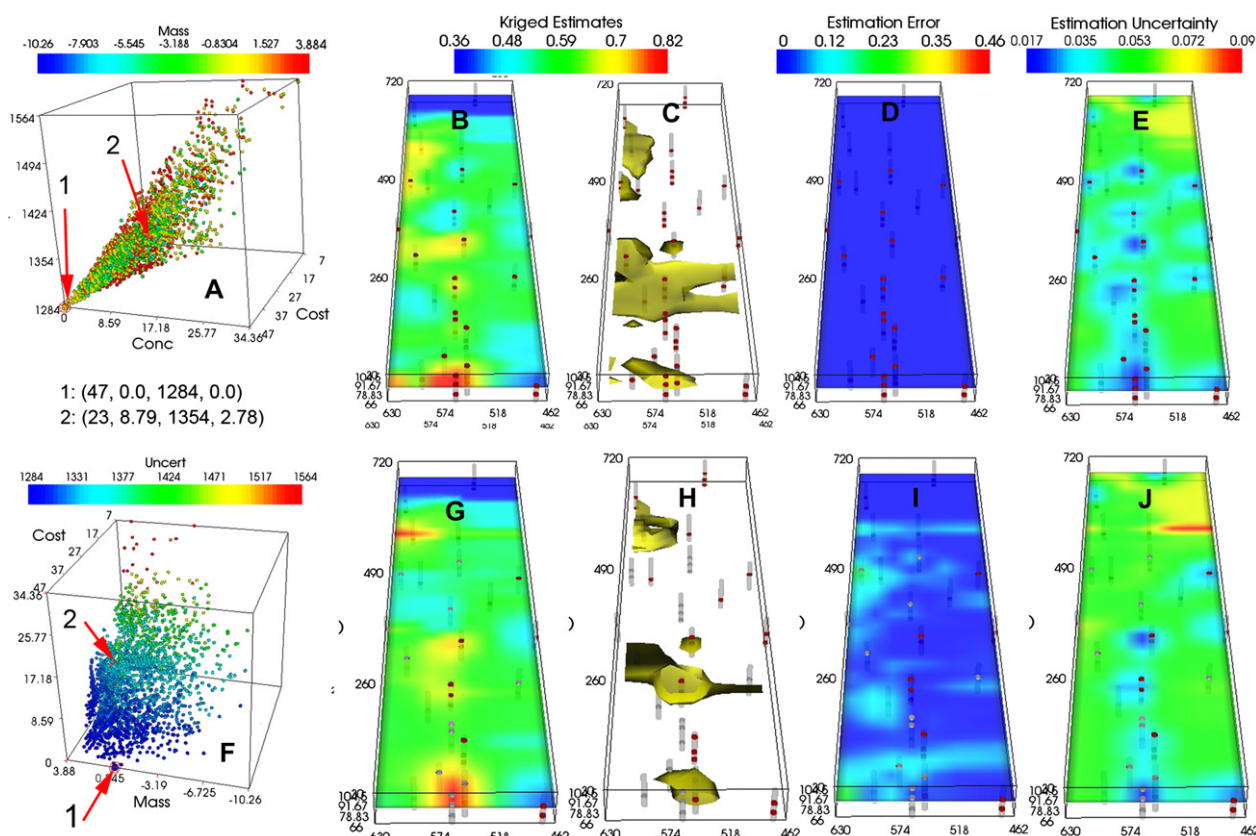


Fig. 3. Example of solution comparison and decision space probing functionality of the VIDEO framework. A and F show alternative views of the locations of the two selected solutions in the objective space. B and G show the associated plume estimates at  $z = 85$  m. C and H show an iso-surface plotted at concentration = 0.70. D and I display the estimation error at  $z = 85$  m. E and J show the spatial extent of estimation uncertainty at  $z = 85$  m.

solutions from the Pareto-set subject to the above constraints, the solution shown in Fig. 3F was selected for comparison because of its desirable characteristic of providing a fairly accurate picture of the plume but at a lower cost. The selected solution represents the objective values: “Cost” = 23, “Conc” = 8.79, “Uncert” = 1354, and “Mass” = 2.78. Compared to the high cost solution examined previously, this represents a 60% decrease in “Cost”, but only a 25% increase in “Conc” and “Uncert” and a 20% increase in “Mass”. The concentration estimates, estimation error, and associated uncertainty are shown in Fig. 3G through J (similarly to Fig. 3B through D). Fig. 3G shows that this particular solution seems to capture the source of the contamination relatively well. However, the solution tends to overestimate the contaminant concentration at the northwest boundary of the plume and underestimate the concentration at the center-west boundary. Fig. 3H provides insight into the volumetric data associated with this solution. In comparing Fig. 3C with H, the iso-surfaces plotted at quantile concentration estimates of 0.7 differ widely between the two solutions. However, both solutions do reveal the high contaminant concentration, which occurs near the center of the domain, mostly because the well at this location is sampled in the second solution and the adjacent wells are not. The estimation error presented in Fig. 3I further reveals this overestimation at the northwest boundary. In addition, high error in the southwest corner can be observed in Fig. 3I. Moving back to Fig. 3G reveals that

this solution does not capture what appears to be an area of high concentration just west of the source. The uncertainty map shown in Fig. 3J shows band-like region of very high uncertainty occurring at the north-east corner of the domain.

The reader should again note that the VIDEO framework allows for interactive probing of the entire volume of the decision space. In the above example, relocation of the probing planes may reveal characteristics of the comparison solution (shown in Fig. 3F), which make it more (or less) desirable as a possible compromise solution. Generally, the objective of selecting and comparing solutions would be to find a lower cost solution which accurately represents important aspects of the contamination plume while remaining within acceptable error and uncertainty bounds throughout the sampling domain. The VIDEO framework provides a solution marking feature which allows the user to mark (and unmark) solutions as they are navigating the objective space. Solutions which have been marked are doubled in size and are bounded with a user specified color. In addition, solutions which have been marked are locked for visualization, meaning that they will never be turned off (unless the user “unmarks” them) regardless of whether or not they are within the plotting limits, or within the precision specification, or dominated by a subsequent non-domination subset sort. This allows the user to efficiently search the space in a variety of ways, all the while tracking solutions of interest. A solution export feature then allows the user to export the marked solutions to a file which can

subsequently be read back into the software to further refine the selections. This process of selection and refinement is referred to as negotiated design selection and is illustrated in Section 3.3.

### 3.3. Negotiated design selection

The process of selecting various solutions of interest and subsequent refinement based on decision maker preferences is referred to as negotiated design selection. In this section, this process will be illustrated using one approach, although many others are possible because the VIDEO framework provides a multitude of visualization tools which can provide many means of selecting interesting solutions. We begin by identifying two-objective subsets of the larger four-objective LTM test case. The two-objective tradeoff representing “Cost” versus “Conc” (Cost–Conc) is shown in Fig. 4A. The solutions in this figure have been marked in red using the “Mark Non-Dominated Subset” option available under the “Objective Space Tools” menu. The color used to mark the solutions can be specified under the “Mark Solution” button available in the objective tools. Keeping the Cost–Conc tradeoff solutions marked from the previous set, the two-objective Cost–Uncert tradeoff is shown in Fig. 4B as solutions marked in green. This process can be repeated for all possible two-objective tradeoff subsets: Cost–Mass, Conc–Mass, and Uncert–Mass. Fig. 4C shows all of the two-objective tradeoffs marked within the Cost–Conc space as shown in Fig. 4A where the Cost–Conc tradeoff is marked in red, the Cost–Uncert tradeoff in green, the Cost–Mass tradeoff in blue, the Conc–Mass in purple, and the Uncert–Mass in orange. If you look closely at Fig. 4A through C, you will see that some solutions which are highlighted occur on more than one two-objective tradeoff. The VIDEO framework handles this by placing multiple marking boxes around solutions which occur on multiple non-dominated subsets. This allows the user to quickly identify intersections which exist between multiple tradeoffs.

At this point, all two-objective LTM tradeoffs have been identified and marked within the full four-objective test case. Manipulating the plotting axes within the framework allows us to see what these tradeoffs look like in the context of the full four-objective space. Fig. 4D shows the full four-objective Pareto-set plotted similarly to Fig. 2A with “Cost”, “Conc”, and “Uncert” plotted on the spatial axes and “Mass” represented by the color of the solutions (where blue represents low “Mass” and red is high “Mass”). In Fig. 4D only the marked solutions are highlighted and the remaining solutions are shown in the background for locational perspective. One of the most interesting features identified from this figure is that the solutions within the Cost–Conc tradeoff (red) and the Cost–Uncert tradeoff (green) generally have very high “Mass” error (red) solutions. Another very interesting feature is that the solutions between these tradeoffs appear as a geometric compromise region. In addition, based on the color of the actual solutions in the compromise region, it appears that they generally exhibit a lower “Mass” error (more blue and

green solutions) than the Cost–Conc and Cost–Uncert tradeoffs. These compromise solutions which are shown in closer detail and marked in Fig. 4E are now the focus of further investigation.

Fig. 4E focuses the objective space window on three solutions of particular interest which have been identified as compromise solutions between the Cost–Conc and Cost–Uncert tradeoffs. These solutions are also marked in Fig. 4D to provide a reference of where they are located in the full four-objective space. First we will focus on the dark blue solution (labeled 1) which represents the lowest “Mass” objective value. In Fig. 4E, solution 1 intersects the Cost–Mass, the Conc–Mass, and the Uncert–Mass tradeoffs. Remember that the VIDEO framework provides a means of identifying this by retaining multiple bounding boxes of various colors indicating that the solution is intersected by multiple tradeoffs. The objective values expressed as a percentage of the maximum objective value of this solution are shown in column one of Table 1. Solution 1 shown in Fig. 5A reduces sampling costs by 30% while only increasing concentration estimation error by 20% and uncertainty by 14%, while the increase in mass estimation error is actually very close to 0%. The relative reduction of each of the objectives’ values for the second and third solutions highlighted in Fig. 4E is also shown in Table 1. In the table, each of the solutions 1–3 is associated with a correspondingly lower “Cost”. However, the solution which samples from 33 points in the domain actually has a lower “Conc” error value than the solution which samples from 35 points. In addition, the “Cost” = 33 solution is also better than the “Cost” = 35 solution in terms of the “Uncert” objective. Even more interesting is the fact the third solution which samples from only 32 points yields even further improvements in the “Uncert” objective. In terms of the “Mass” objective, the third solution presents a compromise between the “Mass” error of the first and second solutions.

Now that the relative magnitudes of the objective values for the three solutions have been compared, their corresponding maps of concentration estimates, estimation error, and estimation uncertainty can be probed and compared using VIDEO’s design space features. Fig. 5 presents these maps for each of the three solutions compared (rows 1–3 of Fig. 5) as well as the maps produced when all available well locations are sampled (designated by the row labeled “Full Cost”). Similarly to Fig. 3, each plane probe is positioned at elevation  $z = 85$  m, and the iso-surface is plotted at a quantile concentration value of 0.7. In comparing the contaminant estimate planes of the three solutions, solution 2 appears to produce a map which most closely represents the information provided by the full test case. However, shifting to an iso-surface view (column 2 of Fig. 5) reveals that the volumetric extent of high concentrations as predicted by solution 2 is underestimated. Comparing the estimation error maps in column 3 of Fig. 5 reveals that solution 1 exhibits the lowest error at  $z = 85$  m. However, moving to the uncertainty map associated with solution 1 reveals a region of very high uncertainty present at the north end of the domain because the well at this boundary is not sampled in this solution. This likely indicates a preference

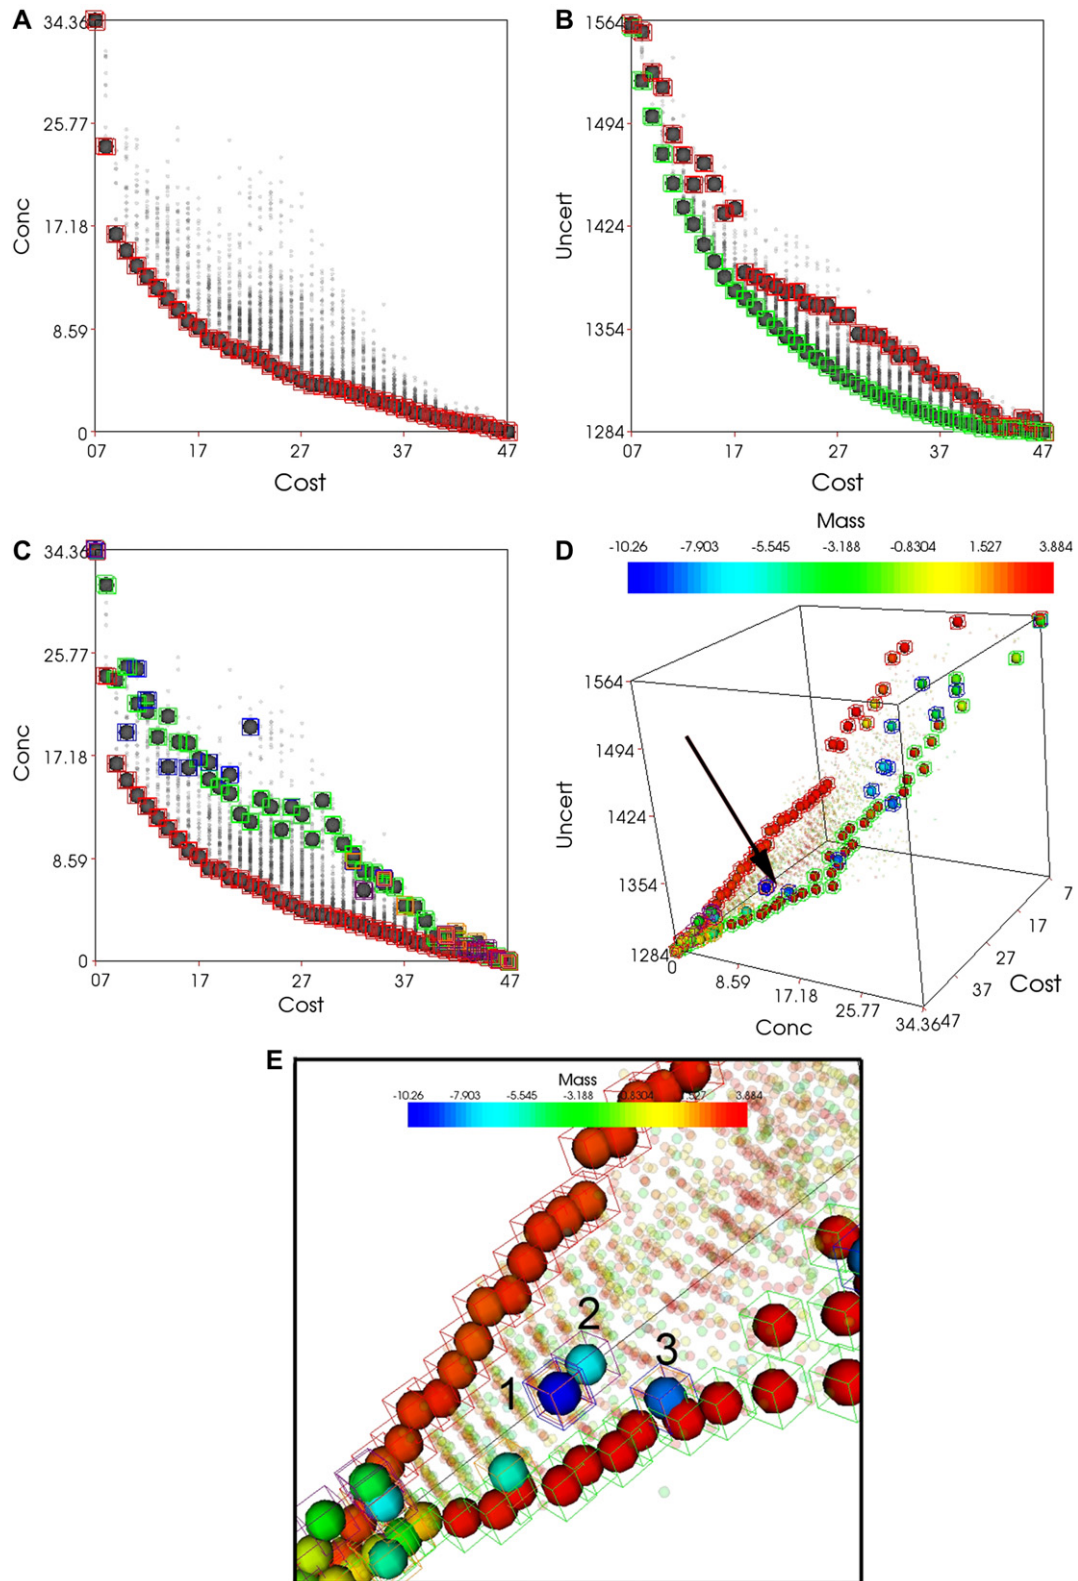


Fig. 4. Negotiated design selection within the VIDEO framework using two objective subsets of the larger four-objective problem. A shows the tradeoff Cost–Conc marked in red and B, the tradeoff Cost–Uncert marked in green. C shows the other two-objective subsets Cost–Mass, Conc–Mass, and Uncert–Mass marked in blue, purple, and orange, respectively. D shows the marked two-objective subsets in the context of the full four-objective space. E shows the selection of three solutions of interest, which are explored more fully in the negotiated design selection section of the study.



Table 1  
Summary of objective values expressed as a percentage of the maximum for each of the three solutions chosen as negotiated designs

Three solutions of interest	1st (%)	2nd (%)	3rd (%)
“Cost” as percentage of maximum	70	65	63
“Conc” as percentage of maximum	20	17	24
“Uncert” as percentage of maximum	14	13	10
“Mass” as percentage of maximum	0	24	13

for solution 2 since quantifying the extents of the contamination is important. Comparing the uncertainty maps of solutions 2 and 3 reveals a preference toward solution 2 since the uncertainty generally decreases throughout many portions of the domain. Based on the above analysis, either solutions 1 or 2 could be suitable as a negotiated design. In addition, it is interesting to note that solution 2 may have been difficult to locate using traditional tools because it is not a solution which is intersected by any of the two-objective tradeoff subsets. However, the VIDEO framework provided an efficient means of locating this solution which would certainly be of interest to a decision maker.

#### 4. Discussion

Tools such as MOEAs are allowing engineers and scientists to solve “many-objective” problems (Fleming et al., 2005) through the generation of high dimensional Pareto-optimal solution sets. Relative to traditional multicriterion decision methods, MOEAs inherently shift methodological focus toward the challenges posed by providing decision makers with the ability to explore and understand large tradeoff solution sets. Our development of the VIDEO framework is based on the hypothesis that design expertise and interactive visualization of both objective tradeoffs and their design space consequences can maximize the value and validity of using MOEAs as a posteriori decision tools. Although exact quantitative metrics can be used to assess the quality of MOEA search (Zitzler et al., 2003; Kollat and Reed, 2006; Tang et al., 2007), it is much more subjective to judge the value of high-order Pareto-optimal sets in the design and decision-making process.

However, the results in Figs. 2–5 do demonstrate how the VIDEO framework can allow environmental engineers and scientists to better understand LTM tradeoffs while seeking a negotiated compromise solution. Perhaps the most interesting result in the prior section is the demonstration that the two-objective Cost–Uncert and Cost–Conc tradeoffs bound a compromise region in the four-objective space (see Fig. 4D). The value of the compromise region within the full four-objective space is particularly interesting given that to date, a vast majority of environmental applications using MOEAs have focused on two-objective formulations due to their simple interpretations as well as their reduced computational constraints. As demonstrated in Fig. 4D, the two-objective non-dominated sets by themselves do not capture the full range of design alternatives available. This study demonstrates that as we overcome the computational constraints posed by

representing our problems with “many-objective” formulations, we should also advance our ability to move beyond traditional two-objective cost–benefit analyses.

High-order Pareto optimization supports emerging decision and design philosophies that seek to allow experts and decision makers to “shop” (Balling, 1999) through sets of “alternative” solutions that will promote design innovations and provide decision makers with an improved understanding of system behaviors. Moreover, our ability to mathematically abstract design into a suite of functions or norms that represent our design or decision support objectives is limited. It is inevitable that some aspects of performance or design will be unmodeled but remain important in the decision process (Loughlin et al., 2001). The VIDEO framework synergistically supports decision makers in better understanding modeled and unmodeled objectives by allowing decision makers to interactively explore high dimensional objective spaces with large solution sets to understand their design tradeoffs while also providing spatial analysis of the design consequences of the LTM tradeoffs.

As noted in Section 2.2, the VIDEO tool has been developed within a general object-oriented Python programming framework. Although this study has demonstrated the VIDEO framework on a LTM application, the tool can be readily adapted to other applications. The objective space probing tools are general and can be very easily adapted to any multi-objective application. The largest challenge and potential limit in applying the VIDEO framework is to develop appropriate probing tools and visualizations for designs or decisions (maps, movies, three-dimensional mesh representations, etc). For each new application area, the combination of powerful, full featured scripting languages like Python with advanced visualization tools such as the VTK provides a set of general tools for developing very complex visualizations and animations. For applications using Geographic Information Systems and computationally expensive design simulators (e.g., Bekele and Nicklow, 2005), on-the-fly probing of designs would require that the Pareto-optimal solutions be simulated offline with their relevant output saved in appropriate file formats. The output files could then be quickly interrogated using decision space probing tools.

The VIDEO framework presented in this study demonstrates the need and value for environmental and water resources professionals to consider high-order Pareto optimization as a new problem class. MOEAs when combined with advanced visualization tools can serve to elucidate complex and potentially unknown dependencies between our objectives (both modeled and unmodeled) for environmental systems. Moreover, the environmental area is clearly being shaped by the emergence of spatiotemporal simulation and statistical tools that are closely coupled with geospatial information systems and geodatabases to support integrated assessment and management of impacted systems (deVoil et al., 2006; Dorner et al., 2007; Schlüter and Rüger, 2007). The emergence of these tools further supports the value and need for a posteriori decision tools that couple advanced visualizations and MOEAs (for an excellent example see Bekele and Nicklow, 2005).



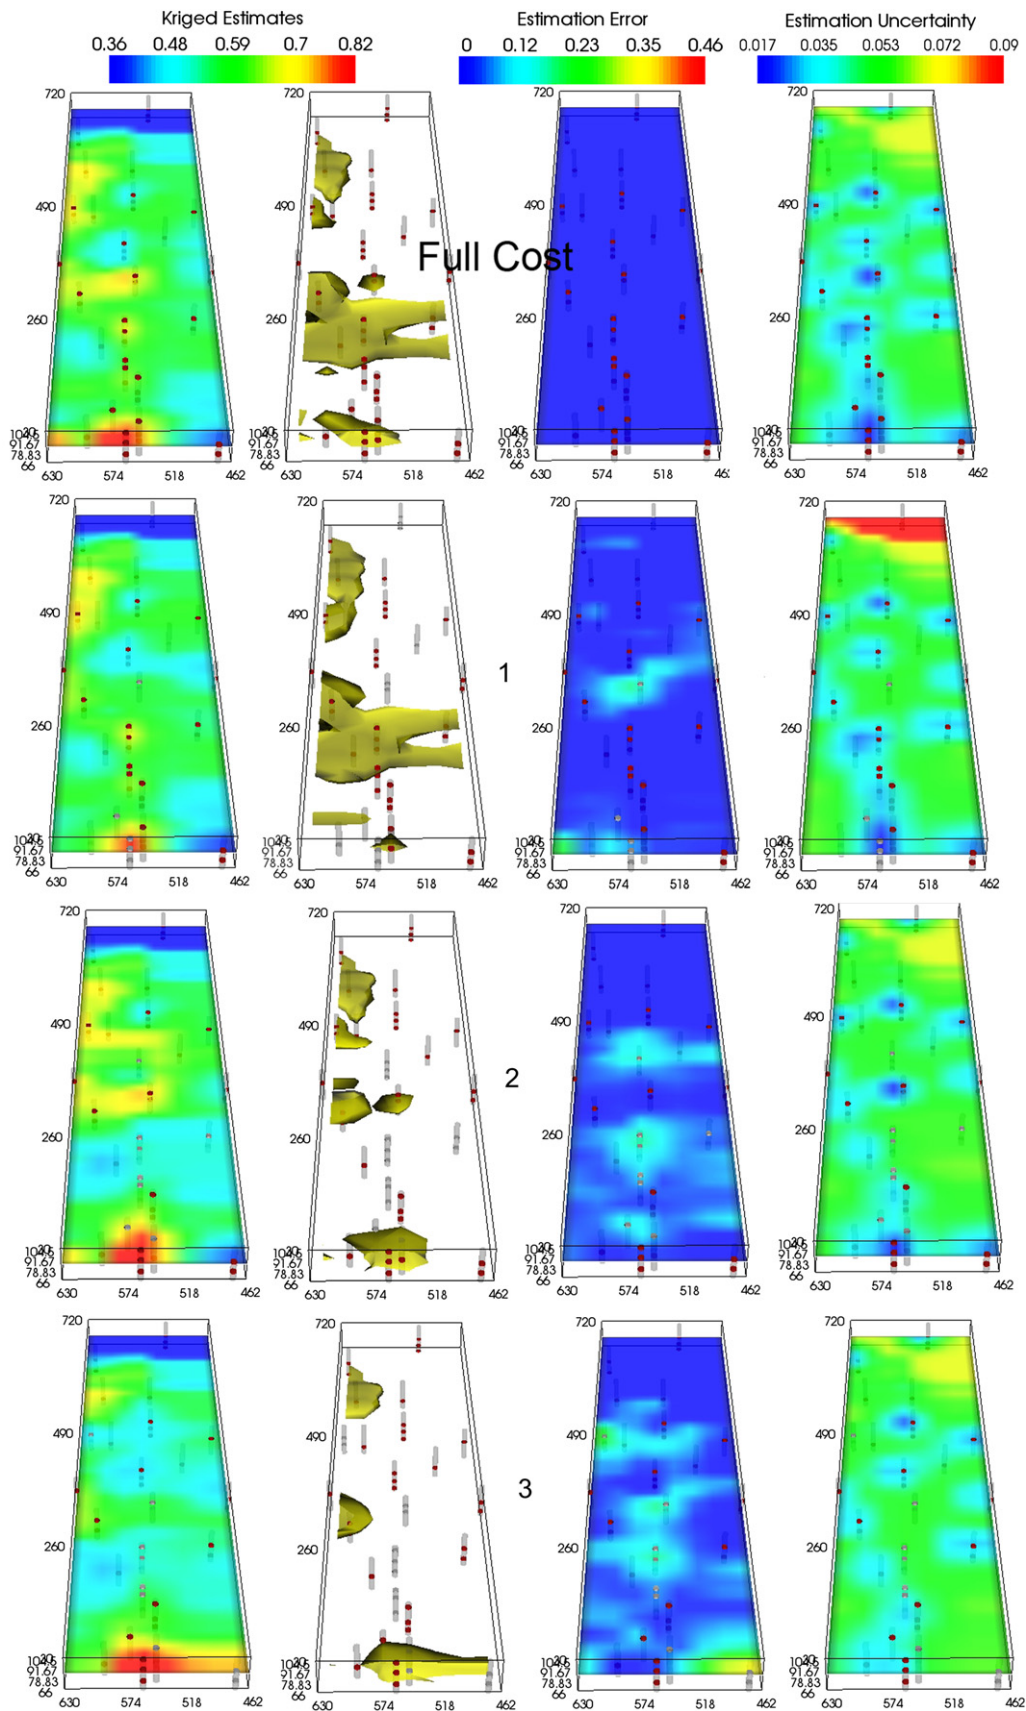


Fig. 5. VIDEO comparison of solutions 1-3, which were selected for further consideration. Columns 1-4 of the figure represent the concentration estimates at  $z = 85$  m, a concentration iso-surface at 0.70, and Kriging error and uncertainty at  $z = 85$  m, respectively. The first row of the figure represents the maps associated with the highest cost solution and the remaining rows, the maps associated with solutions 1-3, respectively.

## 5. Conclusions

This study presents a framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO). The VIDEO framework has been developed to affirm the need and value of combining interactive visualization with high-order Pareto optimization for improved a posteriori decision-making as demonstrated in this study using a long-term groundwater monitoring design application. The core of the VIDEO framework has been developed using the Python programming language. The VIDEO framework has been developed using Tkinter (Python's standard GUI package), which is built on top of Tcl/Tk and is portable across Windows, Unix, and Mac platforms. Aside from Python's ease-of-use and GUI development capabilities, it also allows for easy integration with other, lower level (i.e., faster) languages such as C and C++. The VIDEO framework takes advantage of this capability by utilizing the Visualization ToolKit (Schroeder, 2001), which provides the VIDEO framework with fast visualization rendering and excellent interaction capabilities. The VIDEO framework allows users to visually navigate large multi-objective solution sets while aiding decision makers in identifying one or more optimal designs. Specifically, the interactive visualization framework is intended to provide an innovative exploration tool for high-order Pareto-optimal solution sets (i.e., solution sets for three or more objectives).

The framework is demonstrated for a long-term groundwater monitoring application in which users can explore and visualize tradeoffs for up to four design objectives, simultaneously. Interactive functionality within the framework allows the user to select solutions within the objective space and visualize the corresponding monitoring plan's performance in the design space. If a spatial estimation algorithm is used in the design objective formulation (e.g., in this study, Quantile Kriging was used to provide interpolated maps of contaminant concentration estimates, and their associated error and uncertainty), the estimation can be performed on-the-fly when a solution is selected and the corresponding maps can then be displayed in the decision space. Although Quantile Kriging is used to demonstrate the framework, it should be noted that it is adaptable to any spatiotemporal evaluation of LTM designs (e.g., interpolators, smoothers, or filters). This functionality provides the user with a holistic picture of the information provided by a particular solution, ultimately allowing them to make a more informed decision. In addition, the ease with which the framework allows users to navigate and compare solutions as well as design tradeoffs leads to a time efficient analysis, even when there are thousands of potential solutions.

## Acknowledgements

The authors of this work were partially supported by the United States National Science Foundation under the CAREER grant CBET-0640443. Any opinions, findings and conclusions or recommendations expressed in this paper are those

of the writers and do not necessarily reflect the views of the United States National Science Foundation. The authors would also like to acknowledge Dr. Raymon Masters of Penn State's Graduate Education and Research Services (GEARS). Ray's visualization and VTK expertise was instrumental in providing us the necessary tools to begin developing the VIDEO framework.

## References

- Arrow, K., 1963. *Social Choice and Individual Values*. Yale University Press, New Haven, CT.
- Back, T., Fogel, D., Michalewicz, Z., 2000. *Handbook of Evolutionary Computation*. IOP Publishing Ltd and Oxford University Press, Bristol, UK.
- Balling, R., 1999. Design by shopping: a new paradigm. In: *Proceedings of the Third World Congress of Structural and Multidisciplinary Optimization*, Buffalo, NY, pp. 295–297.
- Bekele, E.G., Nicklow, J.W., 2005. Multiobjective management of ecosystem services by integrative watershed modeling and evolutionary algorithms. *Water Resources Research* 41, W10406, doi:10.1029/2005WR004090.
- Castelletti, A., Soncini-Sessa, R., 2006. A procedural approach to strengthening integration and participation in water resource planning. *Environmental Modeling & Software* 21 (10), 1455–1470.
- Chankong, V., Haimes, Y., 1983. *Multiobjective Decision Making: Theory and Methodology*. North-Holland, New York, NY.
- Chun, W.J., 2006. *Core Python Programming*, second ed. Prentice Hall PTR, New York, NY.
- Cieniawski, S.E., Eheart, J.W., Ranjithan, S.R., 1995. Using genetic algorithms to solve a multiobjective groundwater monitoring problem. *Water Resources Research* 31 (2), 399–409.
- Coello Coello, C., Van Veldhuizen, D.A., Lamont, G.B., 2002. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, NY.
- Deutsch, C.V., Journel, A.G., 1998. *GSLIB: Geostatistical Software Library and User's Guide*. Oxford University Press, New York, NY.
- Dorner, S., Shi, J., Swayne, D., 2007. Multi-objective modelling and decision support using a Bayesian network approximation to a non-point source pollution model. *Environmental Modeling & Software* 22 (2), 211–222.
- Erickson, M.A., Mayer, A., Horn, J., 2002. Multi-objective optimal design of groundwater remediation systems: application of the niched Pareto genetic algorithm (NPGA). *Advances in Water Resources* 25 (1), 51–56.
- Farina, M., Amato, P., 2002. On the optimal solution definition for many-criteria optimization problems. In: Keller, J., Nasraoui, O. (Eds.), *Proceedings of the 2002 NAFIPS-FLINT International Conference*. IEEE Service Center, Piscataway, New Jersey, pp. 233–238.
- Farmani, R., Savic, D.A., Walters, G.A., 2005. Evolutionary multi-objective optimization in water distribution network design. *Engineering Optimization* 37 (2), 167–183.
- Fleming, P.J., Purshouse, R.C., Lygoe, R.J., 2005. Many-objective optimization: an engineering design perspective. In: Coello Coello, C., Hernandez, A., Zitzler, E. (Eds.), *Evolutionary Multi-Criterion Optimization*. Springer Lecture Notes in Computer Science, Guanajuato, Mexico, pp. 14–32.
- Flynt, C., 2003. *Tcl/Tk: A Developer's Guide*, second ed. Morgan Kaufmann, New York, NY.
- Goldberg, D.E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Company, Reading, MA.
- Goovaerts, P., 1997. *Geostatistics for Natural Resources Evaluation*. Oxford University Press, New York, NY.
- Haimes, Y., 1998. *Risk Modeling, Assessment, and Management*. John Wiley & Sons, Inc., New York, NY.
- Halhal, D., Walters, G.A., Ouazar, D., Savic, D.A., 1997. Water network rehabilitation with structured messy genetic algorithm. *Journal of Water Resources Planning and Management* 123 (3), 137–146.
- Horn, J., Nafpliotis, F., 1993. *Multiobjective Optimization Using the Niched Pareto Genetic Algorithm*. IlliGAL Report No. 93005. University of Illinois, Urbana, IL.

- Journel, A.G., Deutsch, C.V., 1997. Rank order geostatistics: a proposal for a unique coding and common processing of diverse data. In: Baafi, E.Y., Schofield, N.A. (Eds.), *Proceedings of the Fifth International Geostatistics Congress*. Kluwer Academic Publishers, Wollongton, Australia.
- Keeney, R., Raiffa, H., 1976. *Decisions with Multiple Objectives*. Wiley, New York, NY.
- Keeney, R.L., Raiffa, H., 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Cambridge University Press, Cambridge, UK.
- Kollat, J.B., Reed, P., 2006. Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. *Advances in Water Resources* 29 (6), 792–807.
- Kollat, J.B., Reed, P.M., 2007. A computational scaling analysis of multiobjective evolutionary algorithms in long-term groundwater monitoring applications. *Advances in Water Resources* 30 (3), 408–419.
- Kumar, S.V., Ranjithan, S.R., 2002. Evaluation of the constraint method-based multiobjective evolutionary algorithm (CMEA) for a three-objective optimization problem. In: Langdon, W.B., et al. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002)*. Morgan Kaufmann, New York, NY, pp. 431–438.
- Laumanns, M., Thiele, L., Deb, K., Zitzler, E., 2002. Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary Computation* 10 (3), 263–282.
- Loughlin, D.H., Ranjithan, S.R., Baugh Jr., J.W., Brill Jr., E.D., 2000. Application of genetic algorithms for the design of ozone control strategies. *Journal of the Air and Waste Management Association* 50, 1050–1063.
- Loughlin, D.H., Ranjithan, S.R., Brill Jr., E.D., Baugh Jr., J.W., 2001. Genetic algorithm approaches for addressing unmodeled objectives in optimization problems. *Engineering Optimization* 33, 549–569.
- Maxwell, R., Carle, F.S., Tompson, F.B., 2000. *Contamination, Risk, and Heterogeneity: On the Effectiveness of Aquifer Remediation*. Livermore, CA.
- Mecham, M., October 1997. Raytheon Integrates Product Development. *Aviation Week & Space Technology* 6, 50.
- Moss, M.E., 1979. Some basic considerations in the design of hydrologic data networks. *Water Resources Research* 15 (6), 1673–1676.
- de Neufville, R., 1990. *Applied Systems Analysis: Engineering Planning and Technology Management*. McGraw-Hill Publishing Company, New York, NY.
- Pareto, V., 1896. *Cours D'Economie Politique*. Rouge, Lausanne.
- Reed, P., Ellsworth, T., Minsker, B.S., 2004. Spatial interpolation methods for nonstationary plume data. *Ground Water* 42 (2), 190–202.
- Reed, P., Kollat, J.B., Deviredy, V., 2007. Using interactive archives in evolutionary multiobjective optimization: a case study for long-term groundwater monitoring design. *Environmental Modelling & Software* 22 (5), 683–692.
- Reed, P., Minsker, B.S., 2004. Striking the balance: long-term groundwater monitoring design for conflicting objectives. *Journal of Water Resources Planning and Management* 130 (2), 140–149.
- Reed, P., Minsker, B.S., Goldberg, D.E., 2001. A multiobjective approach to cost effective long-term groundwater monitoring using an Elitist Nondominated Sorted Genetic Algorithm with historical data. *Journal of Hydroinformatics* 3 (2), 71–90.
- Ritzel, B.J., Eheart, J.W., Ranjithan, S.R., 1994. Using genetic algorithms to solve a multiple objective groundwater pollution containment problem. *Water Resources Research* 30 (5), 1589–1603.
- Salomon, R., 1998. Evolutionary algorithms and gradient search: similarities and differences. *IEEE Transactions on Evolutionary Computation* 2 (2), 45–55.
- Schlüter, M., Rüger, N., 2007. Application of a GIS-based simulation tool to illustrate implications of uncertainties for water management in the Amudarya river delta. *Environmental Modelling & Software* 22 (2), 158–166.
- Schroeder, W.J. (Ed.), 2001. *The Visualization Toolkit User's Guide*. Kitware, Inc..
- Stump, G., Simpson, T.W., Yukish, M., Harris, E.N., 2003. Design space visualization and its application to a design by shopping paradigm. In: *Proceedings of ASME 2003 Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. ASME, Chicago, Illinois.
- Tang, Y., Reed, P., Kollat, J.B., 2007. Parallelization strategies for rapid and robust evolutionary multiobjective optimization in water resources applications. *Advances in Water Resources* 30 (3), 335–353.
- Tang, Y., Reed, P., Wagener, T., 2006. How efficient and effective are evolutionary multiobjective algorithms at hydrologic model calibration? *Hydrology and Earth System Sciences* 10, 289–307.
- deVoil, P., Rossing, W.A.H., Hammer, G.L., 2006. Exploring profit-sustainability trade-offs in cropping systems using evolutionary algorithms. *Environmental Modelling & Software* 21 (9), 1368–1374.
- Wiener, N., 1961. *Cybernetics: or Control and Communication in the Animal and the Machine*. M.I.T. Press, New York, NY.
- Winer, E.H., Bloebaum, C.L., 2002. Development of visual design steering as an aid in large-scale multidisciplinary design optimization. Part 1: method development. *Structural and Multidisciplinary Optimization* 23, 412–424.
- Zeleny, M., 2005. The evolution of optimality: de novo programming. In: Coello Coello, C., Hernandez, A., Zitzler, E. (Eds.), *Evolutionary Multi Criterion Optimization: Third International Conference (EMO 2005)*. Lecture Notes in Computer Science. Springer-Verlag, Guanajuato, Mexico, pp. 1–13.
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., Grunert da Fonseca, V., 2003. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation* 7 (3), 117–132.