DOI: 10.1111/j.1467-8659.2012.03116.x
Eurographics Conference on Visualization (EuroVis) 2012
S. Bruckner, S. Miksch, and H. Pfister
(Guest Editors)

Volume 31 (2012), Number 3

Vismon: Facilitating Analysis of Trade-Offs, Uncertainty, and Sensitivity In Fisheries Management Decision Making

Maryam Booshehrian [1], Torsten Möller [1], Randall M. Peterman [2], Tamara Munzner [3]

[1]School of Computing Science, Simon Fraser University
[2]School of Resource and Environmental Management, Simon Fraser University
[3]Department of Computer Science, University of British Columbia

Abstract

In this design study, we present an analysis and abstraction of the data and task in the domain of fisheries management, and the design and implementation of the Vismon tool to address the identified requirements. Vismon was designed to support sophisticated data analysis of simulation results by managers who are highly knowledgeable about the fisheries domain but not experts in simulation software and statistical data analysis. The previous workflow required the scientists who built the models to spearhead the analysis process. The features of Vismon include sensitivity analysis, comprehensive and global trade-offs analysis, and a staged approach to the visualization of the uncertainty of the underlying simulation model. The tool was iteratively refined through a multi-year engagement with fisheries scientists with a two-phase approach, where an initial diverging experimentation phase to test many alternatives was followed by a converging phase where the set of multiple linked views that proved effective were integrated together in a useable way. Several fisheries scientists have used Vismon to communicate with policy makers, and it is scheduled for deployment to policy makers in Alaska.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics] Applications H.5.2 [Information Interfaces and Presentation] Graphical User Interfaces (GUI), Evaluation/methodology

1 Introduction

Fisheries management involves the regulation of fishing to balance the interests of groups of people who want to catch fish now, and the goal of sustaining viable fish populations for the future. Policy makers make these regulatory decisions in consultation with fisheries scientists, who run extensive computer simulations informed by real-world data. The sophistication of these simulations has steadily increased, and they now generate complex multi-dimensional datasets. The latest simulation models are stochastic and time-dynamic and reflect the variation in environmental influences and uncertainties of many processes [WM04].

However, the complexity of the decision space resulting from the simulations hampers the decision process. Currently, managers need the guidance of the scientists who created the simulations in order to navigate this space and to interpret the simulation model data. The simulation builders would like to allow managers to understand and reason about the decision space on their own. We propose a carefully designed visualization system to enable this transition, created through a three-year collaboration with these scientists.

The contribution of this design study is two-fold. We present an analysis and abstraction of the data, task, and requirements for this fisheries management domain, elucidating the similarities and differences from other domains that also involve multi-dimensional analysis of the relationship between input and output dimensions. We also present the design and implementation of the Vismon interactive visualization tool (see Figure 2). It supports (a) *sensitivity analysis* to check whether small changes in input result in small or large changes in output, and (b) *constraint-based analysis* to rule out parts of the input space based on constraints in the output space. Vismon supports the (c) *analysis of tradeoffs* between indicators *globally distributed* in the output space, whereas many previous systems only support analysis within a local neighborhood. It allows users to (d) *incorporate uncertainty information* into their analysis *at many levels of detail*, with a staged introduction on demand that allows, but does not mandate, drilling down to the full complexity of the underlying probability distributions computed by the stochastic model. Both the abstractions and design were developed and validated iteratively, through a long-term close collaboration [SP06] with fisheries scientists.

2 Fisheries Background

The goal of fish stock assessment is to quantitatively estimate potential outcomes of contemplated management actions. Simulation has been heavily used for stock assessment for decades, with increasing complexity and sophistication [Pet75, Pet04, Pet09, WM04]. Current stochastic simulation models take into account several different sources of uncertainty and risk. Stochastic dynamic models of the natural population implicitly encode a range of alternative hypotheses, taking into account the scientists' limited understanding of the structure and function of aquatic systems and the variability of natural environments. The stochastic models also take into account the challenges of estimating probabilities for uncertain quantities given current monitoring capabilities; real-world data collection (based on harvesting and assessment data) inevitably involves observation and sampling error. The stochastic evaluation of the performance of management options also takes into account both natural variation in catchability and the fact that there may be imperfect compliance with governmental control of human behavior through regulation. A final challenge is communicating complex technical information to decision makers and the public, that is, conveying assumptions, results, and implications to people not actively involved in the analyses.

In this paper, we refer to the managers and stakeholders as *non-expert* users. Although they are experts in the fisheries domain, they are non-expert users of simulation software and statistical data analysis. In contrast, the fisheries scientists are *expert* users, as they are intimately familiar with the data analysis of the simulation systems they have created. The goal of Vismon is to make the data analysis workflow commonly employed by expert users accessible to non-expert users without needing the guidance of the expert user. Vismon is the result of a multi-year two-phase design process in collaboration with various expert users.

3 Data and Task Analysis

We now discuss the data, workflow, and tasks for this design study using a running example, moving from the domain-specific details to abstractions.

3.1 The AYK Simulation Model

The specific simulation model used as a concrete driving example in this paper is for chum salmon populations in the Arctic-Yukon-Kuskokwim (AYK) region of Alaska, U.S.A. [CPZ12]. The three major stakeholder interests at play in this region are sustainability of salmon populations, commercial fishing revenue, and subsistence fishing.

One input option used to help meet managers' objectives is the **escapement target**, which is the desired number of spawning salmon; that is, the number that "escape" being fished. The other is the **harvest rate**, the number of fish that the combination of the commercial and subsistence harvesters should catch *after* the escapement target is met. Each option is set to one of 11 levels, so the simulation covers 121 combinations of these input parameters. Any combination of two options is called a **scenario**. If the user selects a scenario that has not been explicitly simulated, we use bilinear interpolation to predict the outcome of that scenario based on the closest four neighboring simulated scenarios.

The simulation output is 12 indicators for each scenario, grouped into 3 categories: escapement, subsistence catch, and commercial catch. Each simulation run covers a 100-year time period, and the indicators are statistical measures to characterize the results in each category with four output numbers: the average, median, temporal coefficient of variation (CV, or standard deviation divided by the average), and a risk measure expressed as the percentage of years that something undesirable happened during the 100 simulated years.

The stochastic simulation carries out 500 Monte Carlo trial runs for each scenario. The **high-level dataset** of 1452 elements aggregates the values over the 500 runs into a single number, either the average or the median. We use the term **underlying uncertainty** to mean the information contained in the full dataset rather than the high-level one.

In fisheries science, indicators of outcome need to be either maximized or minimized, but not both. Each of the twelve indicators has a direction of desired change as associated metadata, in addition to its set of quantitative values.

3.2 Data Abstraction

Simulation models are one particular example of multidimensional models where the dimensions are divided into two classes: inputs and outputs. The independent input dimensions to the simulation are known as the **management options**, or **options** for short (two in the AYK model). The dependent output dimensions are known as **indicators** (twelve in the AYK model). The simulation generates the outputs given the inputs by running hundreds of Monte Carlo trials. This full underlying dataset is summarized by a few statistical measures for each indicator (see Figure 1).

The design target for Vismon is two input dimensions and ten to twenty output dimensions. This abstraction covers a significant and interesting part of the possible design space, but not all of it. Our collaborators do not currently run simulations with more than two input dimensions, and do not anticipate needing more than a few dozen indicators in total to reflect the interests of the major stakeholder groups. This choice to constrain the tool to supporting few inputs and a moderate number of outputs has many design implications; for example, it is possible to provide overviews of the entire input space and output space without recourse to complex

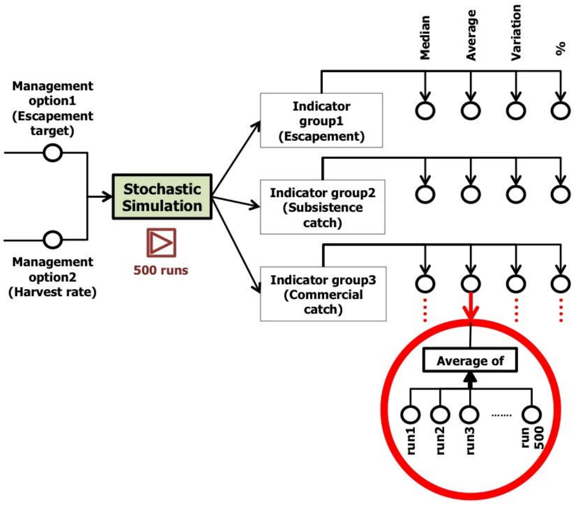
Management option1 (Escapement target)
Management option2 (Harvest rate)
Stochastic Simulation
500 runs
Indicator group1 (Escapement)
Indicator group2 (Subsistence catch)
Indicator group3 (Commercial catch)
Median
Average
Variation
%
Average of
run1
run2
run3
....
run 500

Figure 1: The simulation model used by Vismon takes two inputs, known as management options, and produces groups of output, called indicators. The running example has three groups of indicators, with a median, average, CV value, and a percentage of simulated years with undesireable behaviour. The data underlying each summarized indicator are from 500 Monte Carlo trials.

user interactions for selecting which subsets to view or for exploring only local neighborhoods.

Simulation run times do not affect the interactivity of Vismon, which loads a pre-computed static data file.

3.3 Analysis Task and Previous Workflow

Our collaborators, who are expert users, are working with very mature simulation models that have already been extensively refined and validated [CPZ12], so their interest is in using this model rather than building or verifying it. The analysis workflow established by the scientists contained a detailed list of domain-specific subtasks:

- **W1:** summarize a large number of simulations,
- **W2:** add constraints on value ranges for simulation input (options) and output (indicators) based on stakeholder interests,
- **W3:** select a few candidate combinations of options,
- **W4:** quantify trade-offs in indicators between selected options,

The selection and trade-off analysis should allow the user to observe the following constraints:

- **C1:** avoid sensitive regions of the parameter space where small changes of options yield large changes of indicators,
- **C2:** avoid options with high underlying uncertainty across all Monte Carlo trials.

At the start of our collaboration, Peterman's group was actively engaged in the analysis of simulation results, as they had been for years. Their previous analysis procedure was to use a wide range of individual plots generated with scripts for R and other similar packages. The need to link between these plots has long been recognized; decades ago one of us (Peterman) proposed doing so manually with carefully aligned paper printouts of multiple plots and physical transparency printouts with multiple crosshairs [Pet75]. In addition, although scripts for general-purpose frameworks are a powerful and flexible way to create nearly any individual view showing details at a low level, they require users to know exactly what to specify in advance and was accessible to non-expert users only with close collaboration with scientists. Hence, the overarching goal of our collaborating scientists was:

- **G1:** enable scientists to communicate simulation results to policy makers.
- **G2:** to make this workflow accessible to the non-expert users without the scientists' involvement.

In addition, in the past, only a tiny fraction of the information theoretically available in the dataset was actively considered in the analysis process. The scientists and managers were buried by the quantity of information put out by the models. Essentially, they picked a few points in the parameter space through trial and error and ignored the rest, because they did not have a systematic way to explore the information. Their view of the dataset was narrowly focused; they lacked high-level overviews and other ways to easily synthesize information across a combination of low-level detail views. Exploring the dataset at the level of the aggregate statistical measures (see Figure 1) was very difficult, and understanding the underlying uncertainty expressed in the full details of the Monte Carlo runs was even more so.

Their analysis procedure was most successful in supporting the first two subtasks of summarization (W1) and adding constraints (W2) at a basic level. A very small set of candidate management scenarios was picked (W3) based on past experience, rather than through a data-driven exploration of the simulation output. While the high-level tradeoffs were well known to the scientists and managers, quantifying them for any specific combination of choices was difficult because the relationships are nonlinear (W4). Quantifying trade-offs involved a great deal of cognition and memory, with only minimal help from their perceptual system, in order to synthesize information across multiple individual views. Avoiding sensitive regions (C1) required a great deal of trial and error. Inspecting an individual contour plot showing the values for one indicator could show them regions of rapid change for that indicator where the contour lines were closely spaced, but synthesizing a mental model across all of the indicators was not well supported by the available methods of analysis. Understanding the complexity of the Monte Carlo trials (C2) was also not easily addressed; the scientists typically just worked with the averages because the full dataset was too overwhelming. Communicating results (G1) was only partially addressed. Although their process did support some level of communication between scientists, simulation results were very difficult for policy makers to understand and extremely challenging for stakeholders and the public to grasp (G2).

Considering their tasks at a generic level, we conjectured that many useful scenarios might not even be considered in the candidate set, and conversely that too much analysis time was being spent exploring candidates later found to be unsuitable. Our goal was to allow managers to explore the data on their own and together with different stakeholders. Hence, we needed to create an interface that would allow anyone to quickly find relevant scenarios and to compare several of them.

3.4 Design Requirements

Two key design requirements were to allow all workflow tasks (W1-W4) and constraints (C1, C2) to be carried out within a single tool, and to support interleaving these tasks in any order. For example, the tool should allow uncertainty information to be incorporated into the process of selecting of the candidate actions.

The requirement of encouraging but not forcing uncertainty analysis led us to the design goal of creating views that could all be used in a straightforward way with only the information from the high-level simplified dataset. The combination of these two goals led us to a strategy where views could be augmented with information from the full underlying uncertainty dataset at different levels of complexity. For example, a user should be able to start exploring scenarios using only the high-level average values, and then later incorporate uncertainty information to see which of them are uncertain or risky. The user should also be able to include that information in the initial exploration, so that when faced with scenarios that have the same average values, they can prefer the more certain ones.

4 Related Work

The previous work most relevant to Vismon pertains to systems for multi-dimensional analysis. One of the earliest systems for exploring multi-dimensional parameter spaces was HyperSlice [vWvL93], later extended as the Prosection Matrix [TS98]. These systems required complex navigation of high-dimensional input spaces that is not necessary given our target of only two input dimensions, and they did not support uncertainty or trade-offs analysis.

The inspirational works of Piringer et al. [PBK10] and Berger et al. [BPFG11] address some of the same problems as Vismon. We share some aspects of their solutions in HyperMoVal and its follow-on system, including our choice of linked views in general, and sensitivity analysis via contour plot matrices in particular. Although trade-off analysis was not supported by HyperMoVal [PBK10], it was added in their follow-on work [BPFG11]. They used a Pareto front [EMKH10] technique that allowed the user to change the position of a single scenario in order to understand the trade-off of two indicators, supporting local trade-off analysis. Similarly, HyperMoVal does not support uncertainty analysis at all, whereas the follow-on system does provide simple uncertainty visualization such as box plots. A major difference between these two systems and Vismon is their support for validating simulations by comparing measured data to simulated data, a task that is not required in our domain. Another is their emphasis on navigating the space of many input dimensions, adding complexity that is not necessary for our target of only two inputs.

Two previous systems have focused on parameter selection for image segmentation algorithms [PBCR11, TWSM*11]. Both of these systems handle the case of many input dimensions rather than many output dimensions, for a key difference from Vismon. Pretorius et al. [PBCR11] support sensitivity analysis for this complex case by presenting segmented images at the leaves of a tree encapsulating input parameter settings; users can make visual judgements of the difference in quality between nearby segmentations in local neighborhoods. Their system does not support uncertainty analysis. On the other hand, the Tuner system [TWSM*11] does handle sensitivity analysis in a similar fashion to Vismon, but only with two outputs rather than many. Tuner's trade-off analysis is constrained to a local analysis, and uncertainty analysis is covered only to a limited extent. A strength of both systems is the incorporation of sampling into the analysis pipeline; however, in Vismon the simulation of the data and its analysis are decoupled, so handling sampling issues is not a requirement for our problem.

A number of other systems, such as Design Galleries [MAB*97] and the more recent FluidExplorer [BM10] use clustering to present a comprehensive overview of the variations in the high-dimensional output space. However, neither sensitivity analysis, trade-offs analysis, nor uncertainty is covered in their systems. Classifying the parameter space into classes of different outcomes would be interesting future work to augment Vismon's current capabilities.

Previous work on specific visual encoding and interaction techniques used in Vismon is discussed in the context of the design decisions covered in Sections 5 and 6.

5 Vismon Interface

Vismon is built using multiple linked views, as shown in Figure 2. The three main data abstractions used in Vismon are options (the two input dimensions), indicators (the many output dimensions), and scenarios (a specific combination of the two input options, each of which has associated with it a value for each output dimension). Each of the three main views has a different visual encoding to emphasize different aspects of these elements and the relationship between them. The color coding for scenarios is the same across all of them. Each main view is itself composed of small-multiple charts, with linked highlighting between analogous items on mouseover. All views provide ways to show uncertainty at multiple levels of complexity on demand, but do not force uncertainty analysis on users who want to start simply.

The Constraint pane on the left has sliders that show

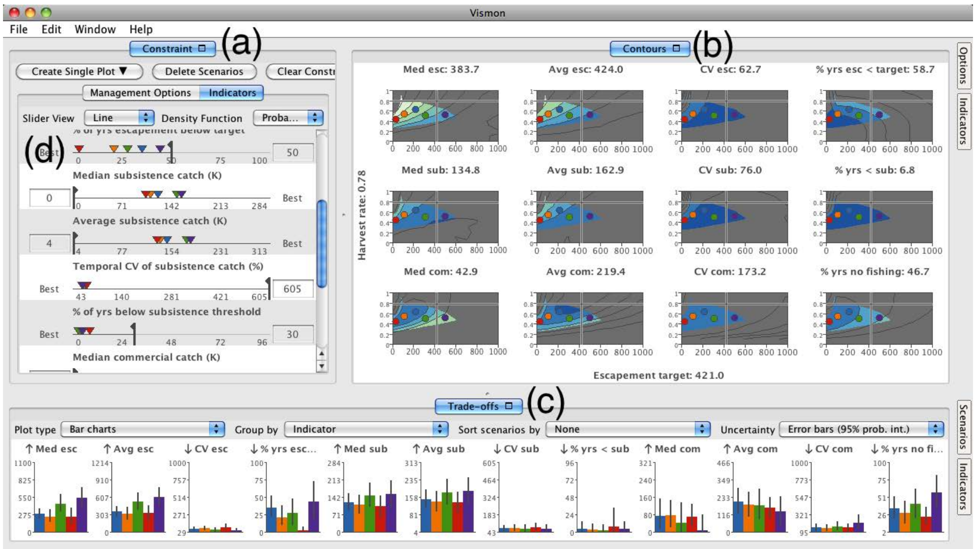

Slider View: Line
Density Function: Proba...
Plot type: Bar charts
Group by: Indicator
Sort scenarios by: None
Uncertainty: Error bars (95% prob. int.)

Figure 2: The Vismon interface: (a) Constraint pane, top-left, shows the list of management options and indicators in separate tabs; (b) Contour Plot Matrix pane, top-right, shows the contour plots of indicators as functions of the two management options and supports scenario selection; (c) Trade-offs pane, bottom, shows detail with the indicators for the selected scenarios; (d) Separate sliders are assigned to management options and indicators in Constraint pane.

ranges for, and allow constraints to be imposed on, individual options and indicators, with optional histograms showing the underlying data distributions for a richer view. The location of the average or median of each scenario with respect to the value ranges is shown with a colored triangle. The Contours pane on the right has a contour plot matrix with one plot for each active indicator showing values with respect to options axes, with scenarios shown as colored dots. The Trade-offs pane on the bottom shows details about indicator values for the active scenarios through rectangular or radial charts. The view can either show a chart for each scenario with marks for indicators, or vice versa.

5.1 Constraint Pane

The Constraint pane shows a tab for the management options and a tab for indicators, both as one-dimensional ranges. In both cases, the base small-multiple view shows a range slider [AWS92, TSDS96], with both a moveable handle for quick interactive positioning and a text box for precise numerical entry when the user knows a value of interest in advance. The sliders allow the user to restrict the active range of any input option or output indicator to avoid unacceptable regions, which changes the shape of the shaded permissible scenario region in the Contours pane plots.

The results of moving the input option sliders are not surprising; a straight line sweeps out horizontally or vertically to change the rectangular size of the active region, because these values correspond with the underlying grid used for both simulation computation and the contour plot display axes. However, changing the undesirable range of the output indicators leads to complex and non-obvious shapes for the active region, as shown in Figure 2. With just a few minutes of exploring with these sliders, the analyst can get the gist of how constraining the different input and output dimensions affects the set of possible scenarios (helping to achieve W2).

Options have bidirectional sliders with both a minimum and maximum handle, and two text boxes. Indicator sliders have only a single handle, since their directionality is known from the metadata. The label `Best` appears instead of a text box on the side that is the most desirable direction, and the handle also has a small flag pointing in that direction as a subtle visual cue.

The plain sliders for the indicators convert to scented widgets [WHA07] on demand from the user, showing histograms of distributions in the underlying dataset in order to provide more guidance on what choices to make when setting the ranges (Figure 3).

There are two choices, either or both of which can be shown. The simpler choice, `MC Trials`, shows a histogram with the distribution of all values for this indicator across all the Monte Carlo trials. Figure 3a (as well as the video) shows an example for the `Average commercial catch` indicator, where the slider bar has been moved from the default position of 0 to the value of 100K. We can see this is an indicator where the maximum value is the most desirable because the `Best` label is on the right side of the slider. The geometric intuition is straightforward: the user can see in advance whether a small or a large part of the distribution will be filtered out when the slider bar is moved to a particular position, rather than using a trial and error process where the slider is moved and then the results are scrutinized. Figure 3b shows the result of drilling down even further by clicking on the red triangle representing the thus-

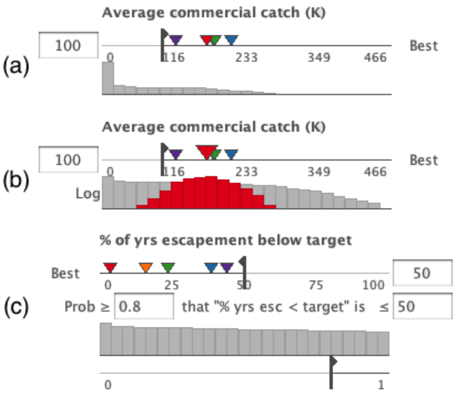
(a)
Average commercial catch (K)
100
0
116
233
349
466
Best
(b)
Average commercial catch (K)
100
0
116
233
349
466
Best
Log
(c)
% of yrs escapement below target
Best
0
25
50
75
100
50
Prob ≥ 0.8 that "% yrs esc < target" is ≤ 50
0
1

Figure 3: The Constraint pane sliders become scented widgets showing histograms on demand. (a) MC Trials shows the distribution of all indicator values across all Monte Carlo trials (500 in our example data). (b) Selecting a scenario shows its distribution vs. the full distribution, log-scaled vertically. (c) Probabilistic Objectives allows the user to set a second probabilistic constraint.

far best scenario as chosen by the manager in the video. The histogram has a colored overlay allowing the user to compare the distribution of the trials just for the chosen scenario with that of the full dataset of all simulated scenarios. The Log label appears on the left to show that in this mode the vertical axis is now log-scale rather than linear, to ensure that the overlay details are fully visible.

The more complex choice, Probabilistic Objectives, allows a sophisticated user to reason about all the Monte Carlo trials, not just their average (following constraint C2). It uses a two-part filter with a second slider and histogram. The base slider still sets a limit on the value of the target indicator. The second slider allows the user to set a probabilistic limit corresponding to the percentage of Monte Carlo simulation trials that are above that indicator limit for indicators that need to be maximized, or below for those that need to be minimized. The second slider allows the user to change this probability value interactively from the default of 0%, meaning that no possibilities have been ruled out, up to a higher number. In Figure 3c, the user has set the probability to 80% that the % of years escapement below target indicator is smaller than 50; that is, a high probability of avoiding low escapement. The plots in the Contours pane will show which scenarios have been ruled out by this limit by crossing them out with X's, as illustrated in Figure 4d.

The sliders are not only controls but also displays, even when not augmented by the histograms. They act as legends that document the range of each option or indicator; the sliders have the same visual range on the screen but cover very different regions of data space. They also show the full name for options and indicators, rather than the short names used in the other panes to save space. Most importantly, the scenario triangles show the distribution of the scenarios with respect to these ranges in a high-precision way using spatial position. That distribution would require more mental effort to glean from the Contour plots, where it is encoded more indirectly and with lower precision as the color of the contour band in which the scenario dot is embedded.

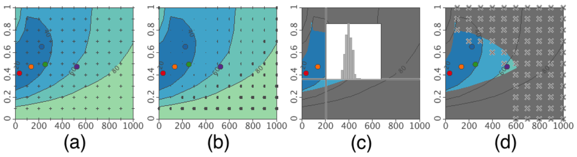

Figure 4: Contour plots can show (a) small points showing the underlying 11 × 11 grid of our example data that is the basis for isocontour interpolation, (b) the points size coded in two directions to summarize the underlying uncertainty as two numbers, (c) a histogram showing the full probability distribution for the point under the cursor, (d) the X'd out region from probabilistic constraints, for comparison to the greyed-out region from the deterministic constraints.

5.2 Contours Pane

The Contours pane contains a contour plot matrix that has one two-dimensional plot for each active output indicator. Each contour plot is drawn using Marching Squares based on simulation values given at an 11 × 11 grid. Each (x,y) location in the contour plot represents a scenario. Again, a scenario is characterized by two independent variables, the parameter settings used for the input management option choices, and has many dependent variables, the output indicators. The small-multiple views are linked with a crosshair that appears at the same (x,y) location in each of them when the cursor moves across any of them, and the exact numeric value for the indicator at that point is shown in each title bar. Numeric legends on contour lines are automatically shown when the plot size is sufficiently large, as shown in Figure 4; these are different in each plot, since each indicator is separately normalized to the color map.

The plots are all linked to the Constraint pane sliders that provide data-driven constraints on the active region within each of them. All plots have the same two axes of the options, and the demarcation between the colored active region and the greyed-out restricted region is the same in all. The plots show the high-level dataset: either the average or the median of the underlying simulation runs (500 in our example data) for their indicator. The plots resize dynamically to fit within the pane as it resizes or the number of plots to show changes as indicators are de- or re-activated, so that they are always visible side by side without the need to scroll. By default, all indicator plots are shown; Figure 2 shows the full set of 12 in the example dataset. They can be turned on and off with a right-mouse popup menu when the cursor is over a plot, or through the control pane which is accessible through the Options tab on the top right of the pane.

The static array of contour patterns provides an overview of the high-level dataset that is focused on the individual indicators (helping to achieve W1). Moving the cursor across the plot allows fast comparison between the indicator values for a single scenario because of the dynamically linked crosshairs. We chose to keep a contour plot matrix at the heart of the system because they were both familiar and effective, and facilitate sensitivity analysis by allowing users to note regions where isocontours fall close together as places to avoid (facilitating constraint C1).

One main use of this view is to guide the user in selecting a small set of candidate scenarios (part of W3), which can be compared in detail in the Trade-offs pane. Clicking within a contour plot selects the scenario at that point. Its location is marked with a colored dot in all plots in this pane and a colored triangle along each indicator range on the Constraint pane sliders. Marks representing selected scenarios are small and show an identifier that is unique for each scenario, so they are coded with high-saturation colors in different hues. This shared color coding acts as a link across different views.

The user can also explore some of the underlying uncertainty data in the Contour pane (facilitating constraint C2), as shown in Figure 4a-c. The 11×11 grid of our example data through which the contours are interpolated is the set of 121 pre-computed scenarios, which can be shown on demand as small points (Figure 4a). These points can be size coded with two additional numbers that summarize the underlying Monte Carlo trials in terms of the same 95% confidence interval information that is used for the error bars described in the next section (Figure 4b). Uncertain regions are clearly indicated by large dots, and strongly asymmetric intervals can be seen where the dots have visibly different aspect ratios. The user can also turn on a histogram showing the full distribution over all Monte Carlo trials at the point under the crosshair (Figure 4c). The histogram updates as the cursor moves, and can be displayed either in the plot containing the cursor or in all of the linked plots.

5.3 Trade-offs Pane

The Trade-offs pane allows a detailed assessment of the trade-offs between a small set of scenarios with a set of small-multiple bar chart plots (supporting W4).

These plots support two kinds of analysis. The default mode is to group outputs by indicator, showing one plot for each indicator with a different colored bar for each scenario, allowing easy comparison of how indicators change across scenarios. The opposite mode is to group by scenario, where each plot shows a single scenario with the bar heights showing all of its indicators. Conversely, this mode allows easy comparison of indicator values within a particular scenario, and the profiles of entire scenarios with each other. When the plots are grouped by indicator, the bars in all of them can be sorted by the value of any indicator, rather than the default based on the order in which the scenarios were created.

The plots support four different levels of showing the underlying uncertainty information, as shown in Figure 5a-d.

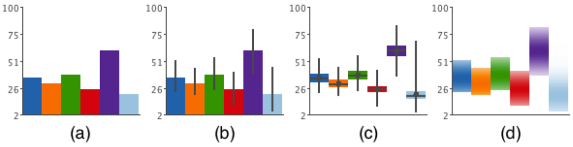

Figure 5: The Trade-offs pane can show uncertainty information on bars in four ways. (a) None. (b) Error bars. (c) Box plots. (d) Shaded distributions.

The simplest possibility is *none* (Figure 5a), so that analysts who want to do only high-level analysis are not forced to deal with uncertainty. The default *error bar* mode superimposes a simple error bar showing the 95% confidence interval on top of the mark (Figure 5b). In *box plot* mode (Figure 5c) the high-level aggregate number is still shown explicitly, but with less salience. Both error bars and box plots are part of the standard arsenal of visualization techniques for uncertain data [OM02, HKDP06, PKRJ10]. The *shaded distribution* mode shows the uncertainty information in full detail (supporting C2) by using a saturation map that encodes the full distribution as normalized density (Figure 5d); this visual encoding was first introduced by Jackson with the name of density strip [Jac08]. The high-level aggregate number is not the most visually salient aspect of the display, but this visual encoding conveys the most information.

5.4 General Functionality

Vismon is implemented in Java 1.6, with diagrams drawn using custom Java2D graphics code. The tick marks on plot axes dynamically adapt to use the available space [TLH10]. Our implementation achieves interactive response on current hardware with the datasets in use by our collaborators, after start-up preprocessing of several seconds.

6 Process and Validation

The requirements for and design of Vismon were created through an iterative two-phase process, through engagement with target users in fisheries science. It has already been successfully deployed for communication between scientists and policymakers, and will soon be deployed for use by policymakers in Alaska.

6.1 Two-Phase Process

We embarked on a two-stage design process for Vismon: a diverging phase to test alternatives and refine our understanding of the requirements, and a converging phase to recombine the successful elements into a focused and useable system based on the fine-tuned requirements.

We chose a framework featuring multiple linked views as the obvious starting point given known visualization design principles [BWK00, LM10]. We reasoned that this baseline capability would speed up the previous analysis process immensely, since our target users were essentially doing linking and brushing [BC87] by hand in their previous workflow.

In the diverging design phase, we assembled many views, both existing and novel, within a testbed linked-view framework. These views included contour plots and scatterplots as single views or a full matrix, a correlation matrix, parallel coordinates, histograms, box plots, star glyphs, and an experimental pixel-based view to show probability distributions. In the terminology of Lloyd and Dykes [LD11], this prototype/testbed functioned as a *data sketch*, allowing the users to explore real data with many different visual encoding and interaction techniques.

In formative testing over the course of 18 months, we gathered feedback from the target users about many versions of the interactive prototype. We were able to distinguish between the more and less effective views for their purposes, iteratively refined the more promising views, and iteratively refined our understanding of the tasks, requirements, and use case. Section 3 lays out this final understanding; it was not nearly so clearly understood at the start of this design study.

In the converging design phase, we redesigned the entire interface for usability based on the combination of previously gathered user feedback and a cognitive walkthrough that incorporated our validated requirements. We carefully considered the trade-offs between the power of multiple views and the cognitive load of excessive window management, and considered screen real estate as a scarce resource. The final design had a core set of three main views with critical functionality that were always visible on screen, with a very small set of optional pop-up views for uncommon operations. This design proved to meet the needs of the scientists (as expressed in goal G1) much better than the first version, and was finalized after only a short round of further refinement. Section 5 describes the capabilities of this final version.

6.2 Lessons Learned

We were able to simplify some aspects of the interface as we moved to a more precise data abstraction: for instance, the constraint sliders started as two-way sliders, but after we realized that indicators did not need to be both maximized and minimized, we substituted one-way sliders.

Other changes were based on a better understanding of the task abstraction. For example, noting that users had created individual scatterplots as part of their previous analysis process, we conjectured that a scatterplot matrix might be even more useful than a collection of separate plots located at haphazard locations; we also provided a correlation matrix with small boxes color-coded by correlation as a compact overview. However, these views were not used; we eventually understood that finding correlations between the indicators was not a central task for sensitivity analysis.

Several more exotic visual encodings were not effective in this domain. The familiar bar charts were strongly preferred by many of the users. Some were willing to use star glyphs, which were left active in the final tool. Parallel coordinates were firmly rejected by these users.

The histogram sliders underneath the constraint sliders evolved in response to the failure of an experimental visual encoding of uncertainty that proved to be incomprehensibly complex. The original *probability statement plot* encoding was a two-dimensional plot with the horizontal axis encoding indicator values from minimum to maximum, and the vertical axis showing probabilities from 0 to 1. Within the box, a pixel had a greyscale value indicating the percentage of management option combinations that satisfied the x axis as the boundary value and the y axis value as the probability value. The final solution does not attempt to show this entire two-dimensional space simultaneously: the Probabilistic Objectives histogram shows only a single one of those curves, and changing the top slider changes which curve is shown. We realized that the users did not need to internalize the full details of the global uncertainty distribution; they only needed to inspect local regions.

6.3 Validation and Deployment

A fisheries scientist with extensive experience in building and analyzing simulation models was our primary target user and source of domain information (Peterman, a coauthor of this paper). Over the course of 30 months, 15 interviews and feedback sessions took place. We also solicited feedback from four more fisheries scientists who work directly with fisheries policymakers for the AYK region, as arms-length target users who were not directly involved in the design process. There were a total of five individual sessions of a few hours each where they used versions of the Vismon first-phase prototype with Peterman's data in individual sessions, and one additional session where one used a near-final Vismon second-phase prototype with his own dataset.

Vismon is being deployed through a staged process, in a similar spirit to the staged development approach of LiveRAC [MMKN08]. The primary target user is the gatekeeper to other fisheries scientists, and they in turn are the gatekeepers to fisheries policymakers. The two interesting deployment milestones for Vismon match our goals G1 and G2.

Several versions of the Vismon prototype have been deployed at the G1 (scientists' communication with policymakers) level by Peterman, including a demonstration to 40 research biologists and high-level fisheries managers from the Alaska Department of Fish and Game. The response of these policymakers was highly positive, verifying that the goal of facilitating communication between technical and nontechnical users was achieved. More specifically, the visual framework allowed managers to ask new questions, promoted discussion and debate, and built trust between managers and scientists for the data analysis process.

Vismon is freely available, with written and video tutorials, at http://www.vismon.org. Scientists have begun to use it at the G1 level using only these training materials, for example a recent presentation to the Fraser River panel of the Pacific Salmon Commission. Two scientists have now gained sufficient confidence in Vismon to plan for its de-

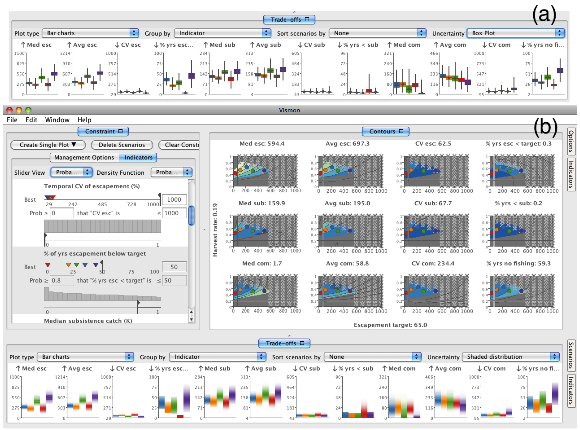

Plot type: Bar charts
Group by: Indicator
Sort scenarios by: None
Uncertainty: Box Plot

Slider View: Proba...
Density Function: Proba...

Temporal CV of escapement (%)
Best 1000
Prob ≥ 0 that "CV esc" is ≤ 1000

% of yrs escapement below target
Best 50
Prob ≥ 0.8 that "% yrs esc < target" is ≤ 50

Med esc: 594.4
Avg esc: 697.3
CV esc: 62.5
% yrs esc < target: 0.3
Med sub: 159.9
Avg sub: 195.0
CV sub: 67.7
% yrs < sub: 0.2
Med com: 1.7
Avg com: 58.8
CV com: 234.4
% yrs no fishing: 59.3
Harvest rate: 0.19
Escapement target: 65.0

Plot type: Bar charts
Group by: Indicator
Sort scenarios by: None
Uncertainty: Shaded distribution

Figure 6: Case study.

ployment at the G2 (analysis by policymakers) stage, as part of the AYK Sustainable Salmon Initiative. It will be a central tool for scientists and managers to carry out model-based analysis to consider alternative harvesting policies. The commitment has been made that "the software will be used by the Panel at Phase 2 workshops and stakeholder meetings"; the first such workshop is scheduled for April 2012 in Alaska. Workshop participants will be trained to build and run a management strategy evaluation model, and will use Vismon to visualize its output.

7 Case Study

This case study summarizes the interactive Vismon session shown in more detail in the accompanying video, using the driving example dataset described in Section 3. It incorporates the kinds of information shown in demonstrations by the fisheries scientists, as discussed above.

The target user is a fisheries manager in Alaska who hopes to make an informed policy decision by finding a scenario with low uncertainty that best suits her objectives. On startup, no constraints have yet been set, so the entire rectangular region of each plot in the Contours pane is fully colored to show that it is active. The manager decides to add constraints using the sliders in the Constraint pane on the left, to reduce the size of the active region by eliminating scenarios that produce unacceptable values of particular indicators (Figure 2). The resulting active region in the contour plots showing the feasible scenarios is much smaller than the full original set, thereby simplifying the complexity faced by the manager. Some indicators within the active region are dark blue, showing that they are in the highly preferred range, while others are the lighter green color indicating unfavorable values.

The manager then explores a few management options by clicking the mouse in a few locations within the active region, and the Vismon window updates to include the Trade-offs pane showing detailed information about those scenarios. Figure 6a shows the Trade-offs pane after she switches to showing more detailed uncertainty information as box plots rather than error bars. She then digs even deeper by looking at the shaded distributions (bottom of Figure 6b). By assessing the trade-offs for the chosen scenario, and considering her objectives of having high commercial catch while having low percentage of years with low escapement target and no commercial fisheries, she decides that the red scenario looks promising.

She scrolls down in the Constraint pane to look at the commercial catch indicators, and then sorts the charts in the Trade-offs pane by the Average commercial catch indicator. She also changes the settings to show the distributions for each indicator (Figure 3b).

Figure 3c shows the display after she turns on the second set of histograms underneath the sliders in the Constraint, and sets some probabilistic limit. The Contours plots now have many crossed out locations, indicating the option combinations that have been ruled out by this setting for the probabilistic acceptable values. She notes that none of the scenarios are in the crossed out region that is unacceptable according to this probabilistic constraint (Figure 6b). She concludes that the red scenario is indeed the best alternative, given her analysis of the underlying Monte Carlo information in addition to the summarized version of the dataset.

8 Conclusion and Future Work

Over the past three years we worked closely with fisheries scientists to develop a tool that allows the analysis of fisheries models to move from their builders into the hands of managers and policymakers. We engaged in a requirements analysis that allowed us to articulate and abstract data and tasks for this domain. The resulting Vismon system facilitates the analysis of multidimensional data with two independent inputs, 10 to 20 outputs, and includes a multilevel view of the models underlying uncertainty, expressed by multiple Monte Carlo runs. Sensitivity of a policy decision is encoded by close contour lines in the contour plots. A comprehensive trade-off analysis allows users to compare several alternative policy options.

An iterative design and validation process for Vismon has led fisheries scientists to deem it ready for deployment with policymakers. Future work will focus on facilitating this deployment, adapting it to more closely fit the needs of policymakers, and observing the learnability and utility of methods for uncertainty visualization.

9 Acknowledgements

This project was funded by National Science and Engineering Research Council of Canada and Simon Fraser University. We thank Sean Cox, Andy Cooper, Mike Jones, Matt Catalano, and Catherine Michielsons for feedback on Vismon. We thank Matt Brehmer, Jessica Dawson, Stephen Ingram, and Michael Sedlmair for helpful comments.

References

[AWS92] AHLBERG C., WILLIAMSON C., SHNEIDERMAN B.: Dynamic queries for information exploration: An implementation and evaluation. In *Proc. ACM Conf. Human Factors in Computer Systems (CHI)* (1992), pp. 619–626. 5

[BC87] BECKER R. A., CLEVELAND W. S.: Brushing scatterplots. *Technometrics 29* (1987), 127–142. 7

[BM10] BRUCKNER S., MÖLLER T.: Result-driven exploration of simulation parameter spaces for visual effects design. *IEEE Trans. Visualization and Computer Graphics (Proc. Vis) 16*, 6 (2010), 1468–1476. 4

[BPFG11] BERGER W., PIRINGER H., FILZMOSER P., GRÖLLER E.: Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. *IEEE Computer Graphics Forum (Proc. EuroVis) 30*, 3 (2011), 911–920. 4

[BWK00] BALDONADO M. Q. W., WOODRUFF A., KUCHINSKY A.: Guidelines for using multiple views in information visualizations. In *Proc. ACM Advanced Visual Interfaces (AVI)* (2000), pp. 110–119. 7

[CPZ12] COLLIE J. S., PETERMAN R. M., ZUEHLKE B.: A fisheries risk-assessment framework to evaluate trade-offs among management options in the presence of time-varying productivity. *Canadian Journal of Fisheries and Aquatic Sciences 69*, 2 (February 2012), 209–233. 2, 3

[EMKH10] ESKELINEN P., MIETTINEN K., KLAMROTH K., HAKANEN J.: Pareto navigator for interactive nonlinear multiobjective optimization. *Operations Research Spectrum 32*, 1 (2010), 211–227. 4

[HKDP06] HAGEN H., KERREN A., DANNENMANN P., POTTER K.: Methods for presenting statistical information: The box plot. *Visualization of Large and Unstructured Data Sets 4* (2006), 97–106. 7

[Jac08] JACKSON C. H.: Displaying uncertainty with shading. *The American Statistician 62*, 1 (2008), 340–347. 7

[LD11] LLOYD D., DYKES J.: Human-centered approaches in geovisualization design: Investigating multiple methods through a long-term case study. *IEEE Trans. Visualization and Computer Graphics (Proc. InfoVis) 17*, 12 (2011), 2498–2507. 8

[LM10] LAM H., MUNZNER T.: *A Guide to Visual Multi-Level Interface Design From Synthesis of Empirical Study Evidence*, vol. 1 of *Synthesis Lectures on Visualization Series*. Morgan Claypool, 2010. 7

[MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUML W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proc. ACM Conf. Computer Graphics and Interactive Techniques (SIGGRAPH)* (1997), pp. 389–400. 4

[MMKN08] MCLACHLAN P., MUNZNER T., KOUTSOFIOS E., NORTH S.: LiveRAC - interactive visual exploration of system management time-series data. In *Proc. ACM Conf. Human Factors in Computing Systems (CHI)* (2008), pp. 1483–1492. 8

[OM02] OLSTON C., MACKINLAY J. D.: Visualizing data with bounded uncertainty. In *Proc. IEEE Symp. Information Visualization (InfoVis)* (2002), pp. 37–40. 7

[PBCR11] PRETORIUS A. J., BRAY M.-A. P., CARPENTER A. E., RUDDLE R. A.: Visualization of parameter space for image analysis. *IEEE Trans. Visualization and Computer Graphics (Proc. InfoVis) 17*, 12 (2011), 2402–2411. 4

[PBK10] PIRINGER H., BERGER W., KRASSER J.: HyperMoVal: interactive visual validation of regression models for real-time simulation. *IEEE Computer Graphics Forum (Proc. EuroVis) 29*, 3 (2010), 983–992. 4

[Pet75] PETERMAN R. M.: New techniques for policy evaluation in ecological systems: Methodology for a case study of pacific salmon fisheries. *Journal of the Fisheries Research Board of Canada 32*, 11 (1975), 2179–2188. 2, 3

[Pet04] PETERMAN R. M.: Possible solutions to some challenges facing fisheries scientists and managers. *ICES Journal of Marine Science: Journal du Conseil 61*, 8 (2004), 1331–1343. 2

[Pet09] PETERMAN R. M.: Fisheries science in the future. In *The Future of Fisheries Science in North America*, Beamish R. J., Rothschild B. J., (Eds.), vol. 31. Springer Netherlands, 2009, pp. 167–184. 2

[PKRJ10] POTTER K., KNISS J., RIESENFELD R., JOHNSON C. R.: Visualizing summary statistics and uncertainty. *IEEE Computer Graphics Forum (Proc. EuroVis) 29*, 3 (2010), 823–832. 7

[SP06] SHNEIDERMAN B., PLAISANT C.: Strategies for evaluating information visualization tools: Multi-dimensional in-depth long-term case studies. In *Proc. AVI Workshop on BEyond time and errors: novel evaLuation methods for Information Visualization (BELIV)*. ACM, 2006. Article 6. 2

[TLH10] TALBOT J., LIN S., HANRAHAN P.: An extension of Wilkinson's algorithm for positioning tick labels on axes. *IEEE Trans. Visualization and Computer Graphics (Proc. InfoVis) 16*, 6 (2010), 1036–1043. 7

[TS98] TWEEDIE L., SPENCE R.: The prosection matrix: A tool to support the interactive exploration of statistical models and data. *Computational Statistics 13*, 1 (1998), 65–76. 4

[TSDS96] TWEEDIE L., SPENCE R., DAWKES H., SU H.: Externalising abstract mathematical models. In *Proc. ACM Conf. Human Factors in Computer Systems (CHI)* (1996), pp. 406–412. 5

[TWSM*11] TORSNEY-WEIR T., SAAD A., MÖLLER T., WEBER B., HEGE H.-C., VERBAVATZ J.-M., BERGNER S.: Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration. *IEEE Trans. Visualization and Computer Graphics (Proc. Vis) 17*, 12 (2011), 1892–1901. 4

[vWvL93] VAN WIJK J. J., VAN LIERE R.: HyperSlice visualization of scalar functions of many variables. In *Proc. IEEE Conf. Visualization (Vis)* (1993), pp. 119–125. 4

[WHA07] WILLETT W., HEER J., AGRAWALA M.: Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics (Proc. InfoVis) 13* (2007), 1129–1136. 5

[WM04] WALTERS C. J., MARTELL S. J. D.: *Fisheries ecology and management*. Princeton University Press, 2004. 1, 2