PNNL-38413, Rev 0
DVZ-RPT-121, Rev 0

Web-Based Tools for Data-Informed Remedy Optimization: Software Theory and User Guide

September 2025

Xuehang Song
Frank Lopez Jr.
Mikaela Corney
Paul Tran
Yusuf S. Afzal
Tommy E. Joppich
Sophie N. Baur
Hung Luu
Reem Osman
Christian D. Johnson
Inci Demirkanli

U.S. DEPARTMENT of ENERGY

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <https://www.ntis.gov/about>
Online ordering: http://www.ntis.gov

Web-Based Tools for Data-Informed Remedy Optimization: Software Theory and User Guide

September 2025

Xuehang Song
Frank Lopez Jr.
Mikaela Corney
Paul Tran
Yusuf S. Afzal
Tommy E. Joppich
Sophie N. Baur
Hung Luu
Reem Osman
Christian D. Johnson
Inci Demirkanli

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Summary

This report documents the development and application of two web-based decision-support tools for pump-and-treat (P&T) groundwater remediation systems: PTOLEMY (Pump-and-Treat Optimized Location Evaluation to Maximize Yields) and OPTIMA (Optimization for Pump-and-Treat Implementation, Management, & Assessment). These tools enhance remedy design and management by leveraging advanced computational methods – specifically deep learning and multi-objective optimization – within a user-friendly platform. By integrating data-driven models with established hydrogeological knowledge, PTOLEMY and OPTIMA enable more efficient evaluation of well placement and operational strategies, helping site managers balance multiple remediation objectives under complex conditions. Both tools are implemented as modules within the SOCRATES (Suite Of Comprehensive Rapid Analysis Tools for Environmental Sites) web platform, which provides data access, visualization, and analytics to support remedy optimization across sites in the U.S. Department of Energy Office of Environmental Management complex.

PTOLEMY is a rapid screening module designed to identify promising locations for new extraction wells. It employs a multi-channel three-dimensional convolutional neural network (MC3D-CNN) trained on high-fidelity simulation data to predict the relative performance (in terms of contaminant mass recovery) of potential well sites. Through an interactive web interface, PTOLEMY visualizes the probability of high performance across a site, highlighting areas where an extraction well is likely to yield above-threshold contaminant removal over a multi-year period. PTOLEMY's map-based displays and exportable results support transparent communication of screening analyses. By focusing attention on the most favorable candidate locations, the tool augments traditional engineering judgment and physics-based modeling, providing a data-informed basis for subsequent detailed evaluations.

OPTIMA is a multi-objective optimization module designed to find wellfield layouts and operating schedules that meet various cleanup goals. It quickly evaluates thousands of candidate setups – combinations of well locations, timing, and rates – and returns a small set of best trade-off options for comparison. At its core, OPTIMA uses a U-Net-based surrogate model – a deep-learning emulator of a groundwater flow and transport simulator – to dramatically accelerate scenario evaluations. Coupling this fast surrogate with the NSGA-II (Non-dominated Sorting Genetic Algorithm II) evolutionary algorithm, OPTIMA explores a wide decision space of well locations and schedules to identify Pareto-optimal solutions that trade off key objectives (e.g., minimizing cleanup time, maximizing contaminant mass removal, and minimizing plume extent). The tool outputs a family of optimal configurations and visualizes their trade-offs (Pareto frontiers of cleanup metrics and maps of optimized well placements). Site managers can use these results to understand the range of viable strategies and to select candidate designs for more detailed verification. OPTIMA is currently under active development and not yet fully released; this guide provides early documentation to support planning and gather user feedback.

Both PTOLEMY and OPTIMA are intended as decision support tools to guide and streamline P&T design processes – not as replacements for site-specific modeling or regulatory decision-making requirements. They provide preliminary analyses that must ultimately be confirmed with high-fidelity models and engineering review. The tools are designed for transparency and traceability: They draw from established data sources (e.g., a well-calibrated site model and historical operational records) and allow users to export maps, data tables, and configurations for documentation. By accelerating the exploration of "what-if" scenarios and narrowing down options, PTOLEMY and OPTIMA support more informed, timely, and cost-effective decision-making in P&T remedy management. In summary, this report provides theory, usage guidance, and implementation considerations for these tools, illustrating how they can be applied to improve P&T system outcomes while maintaining quality assurance and alignment with remediation objectives.

Acknowledgments

Funding for this work was provided by the U.S. Department of Energy Hanford Field Office under the Deep Vadose Zone – Applied Field Research Initiative. Pacific Northwest National Laboratory is operated by Battelle Memorial Institute for the U.S. Department of Energy under Contract DE-AC05-76RL01830

Acronyms and Abbreviations

Adam	Adaptive Moment Estimation
CPU	central processing unit
CSV	comma-separated values
DOE	U.S. Department of Energy
eSTOMP	(STOMP family) groundwater simulator
HFO	Hanford Field Office
JPEG	Joint Photographic Experts Group (image format)
MC3D-CNN	multi-channel 3-D convolutional neural network
NSGA-II	Non-dominated Sorting Genetic Algorithm II
OPTIMA	Optimization for Pump-and-Treat Implementation, Management, and Assessment
P2R	Plateau-to-River (groundwater model)
PDF	Portable Document Format
P&T	pump-and-treat
PTOLEMY	Pump-and-Treat Optimized Location Evaluation to Maximize Yields
RMSE	root mean square error
ROD	Record of Decision
SOCRATES	Suite Of Comprehensive Rapid Analysis Tools for Environmental Sites
U-Net	U-Net convolutional neural network architecture
ZIP	ZIP archive (compressed file format)

Contents

Summary ... iii
Acknowledgments ... iv
Acronyms and Abbreviations ... v
1.0 Introduction ... 1
1.1 Purpose and Scope ... 1
1.2 Site Background ... 1
1.3 Overview of PTOLEMY and OPTIMA ... 4
1.4 Intended User and Applications ... 5
1.5 Report Organization ... 5
2.0 PTOLEMY Theory ... 7
2.1 Multi-Channel 3D Convolutional Neural Network Approach ... 7
2.2 Training Data Development and Performance Classification ... 8
2.3 Model Training and Performance Evaluation ... 9
2.4 Decision Criteria and Constraints ... 9
2.5 Current Limitations and Assumptions ... 10
3.0 PTOLEMY User Guide ... 11
3.1 Interface Overview ... 11
3.1.1 Map Panel ... 11
3.1.2 Performance Probability Distribution Graph Panel ... 12
3.2 Interface Overview ... 12
3.2.1 High Performance Thresholds and Contaminants ... 12
3.2.2 Suggested Well Locations ... 13
3.2.3 Adding New Well Locations ... 15
3.3 Analysis Results ... 15
3.3.1 Interactive Chart–Map Linking ... 16
3.3.2 Chart Settings ... 17
3.4 Exporting Results and Reports ... 17
3.5 Summary ... 19
4.0 OPTIMA Theory ... 20
4.1 Framework Overview ... 20
4.2 Surrogate Modeling: U-Net with Physics-Informed Well Drawdown Channel ... 21
4.2.1 U-Net Surrogate for Multi-Variable Prediction ... 21
4.2.2 Input Representation and Steady-State Drawdown Channel ... 21
4.2.3 Training Data and Performance ... 22
4.3 Objective Formulation ... 23
4.3.1 Remediation Objectives ... 23
4.3.2 Decision Variables and Design Representation ... 23

4.3.3 Operational Constraints 23
4.4 NSGA-II Evolutionary Optimization Engine 24
4.4.1 Evolutionary Search Process 24
4.4.2 Surrogate-in-the-Loop Evaluation and Outputs 24
4.5 Current Limitations & Assumptions 24
5.0 OPTIMA Design and User Guide 26
5.1 Interface Overview 26
5.1.1 Simulation Overview Page 26
5.1.2 Workspace Layout 27
5.2 Defining an Optimization Scenario 27
5.2.1 Operational Constraints 27
5.2.2 Objectives Selection 28
5.3 Narrowing Results 29
5.3.1 Full Candidate Set 29
5.3.2 Filtered Subset and Progression 30
5.4 Compare Solutions 31
5.4.1 Initial State 31
5.4.2 Single Solution Selected 32
5.4.3 Current vs. Pinned Time 33
5.4.4 Layout Controls 33
5.5 Exporting Results and Reports 34
5.6 Summary 35
6.0 Conclusions 36
7.0 Continued Development 37
8.0 Quality Assurance 38
9.0 References 39

Figures

Figure 1.1. Extent of Contamination Plumes with Concentrations above Water Quality Standards on the Central Plateau (Source: DOE/RL-2020-60, Rev. 0) ... 2
Figure 1.2. Hydrostratigraphic Units and Generalized Hanford Site Stratigraphy (Source: Hammond and Lupton 2015) ... 3
Figure 2.1. Architecture of the MC3D-CNN Classification Model ... 8
Figure 2.2. Deep Learning Workflow for Well Performance Classification, from Training Data Generation to MC3D-CNN Prediction ... 9
Figure 3.1. PTOLEMY User Interface ... 11
Figure 3.2. Map Options in PTOLEMY ... 12
Figure 3.3. High-Performance Threshold Adjustment Options ... 13
Figure 3.4. Suggested Well Locations on Map ... 13
Figure 3.5. Suggested Well Locations and Distance from Nearest Extraction Well Options ... 14
Figure 3.6. Distance From Nearest Extraction Well Display on Map ... 14
Figure 3.7. Adding a New Well Location to Map ... 15
Figure 3.8. Performance Probability Distribution Graph ... 16
Figure 3.9. Corresponding Chart Lines and Map Well Locations ... 16
Figure 3.10. Changing Graph Y-Axis ... 17
Figure 3.11. Data Export Button ... 18
Figure 3.12. Data Export Options ... 18
Figure 4.1. OPTIMA Computational Framework Illustrating the Integrated Workflow ... 20
Figure 4.2. U-Net Deep-Learning Framework Used by OPTIMA ... 21
Figure 5.1. OPTIMA Simulation Overview (landing page) ... 26
Figure 5.2. Set Parameters – Operational Constraints ... 28
Figure 5.3. Set Parameters – Custom Inputs and Objectives ... 29
Figure 5.4. Refine Results – Full Candidate Set ... 30
Figure 5.5. Refine Results – Filtered Subset Highlighted ... 30
Figure 5.6. Compare Solutions – Initial (no selection) ... 31
Figure 5.7. Compare Solutions – One Solution Selected ... 32
Figure 5.8. Compare Solutions – Current vs. Pinned ... 33
Figure 5.9. Compare Solutions – Collapsed Left Panel ... 34
Figure 5.10. Compare Solutions – Drag to Resize Panels ... 34
Figure 5.11. Download Assets Dialog ... 35

1.0 Introduction

1.1 Purpose and Scope

This report documents the development and application of two complementary decision-support tools – PTOLEMY (Pump-and-Treat Optimized Location Evaluation to Maximize Yields) and OPTIMA (Optimization for Pump-and-Treat Implementation, Management, and Assessment) – designed to enhance groundwater remediation efforts, specifically focusing on pump and treat (P&T) operations, at the Hanford Site's Central Plateau. These web-based tools are implemented as modules within the SOCRATES (Suite Of Comprehensive Rapid Analysis Tools for Environmental Sites) platform, a framework for data access, visualization, and analytics developed by Pacific Northwest National Laboratory to support remedy optimization. These tools were created to improve the effectiveness and efficiency of the 200 West Area pump-and-treat (P&T) system by guiding P&T well placement and operation strategies as pre-screening-level decision support.

PTOLEMY and OPTIMA integrate advanced computational methods (e.g., deep learning surrogates and evolutionary algorithms) with the existing Hanford Site groundwater flow and transport model to support performance-based remedy management. The purpose of this report is to describe both tools, explain their context and rationale, and demonstrate how they work together to inform P&T system optimization. The scope includes an overview of contaminant and aquifer conditions and challenges at the Hanford 200 West Area, the technical architecture and features of each tool, their intended use and limitations, and a roadmap for how these tools fit into the overall P&T optimization workflow. This introduction establishes the decision context and motivations for developing PTOLEMY and OPTIMA and outlines the content and organization of the remainder of the report. It is important to note that OPTIMA is currently under active development and is not yet finalized; this user guide is provided to facilitate planning-level evaluations and gather early user feedback.

1.2 Site Background

The 200 West Area, located on the Hanford Site's Central Plateau, contains several co-mingled groundwater contaminant plumes resulting from historical plutonium separation and waste disposal operations (1940s–1970s). Key contaminants of concern in the Central Plateau's unconfined and semi-confined aquifers include carbon tetrachloride (CCl_4) technetium-99 (Tc-99), nitrate (NO_3), hexavalent chromium, trichloroethene, and others (Figure 1.1). The underlying hydrostratigraphy that controls plume behavior includes unconsolidated Hanford formation sands and gravels, the finer-grained Cold Creek unit, layered Ringold Formation units, and basalt below, forming a complex, heterogeneous aquifer system with discontinuous aquitards and variable hydraulic properties (Figure 1.2). At the 200 West Area, contamination mainly occur across two aquifers – the unconfined Ringold Unit E (Rwie) and the semi-confined Ringold Unit A (Rwia) – which are generally separated by the Ringold Lower Mud Unit (Rlm) aquitard, where the Rlm is locally discontinuous, allowing vertical hydraulic connection. This geologic complexity produces varied groundwater flow paths and contaminant transport rates, complicating remedial design and performance evaluation.

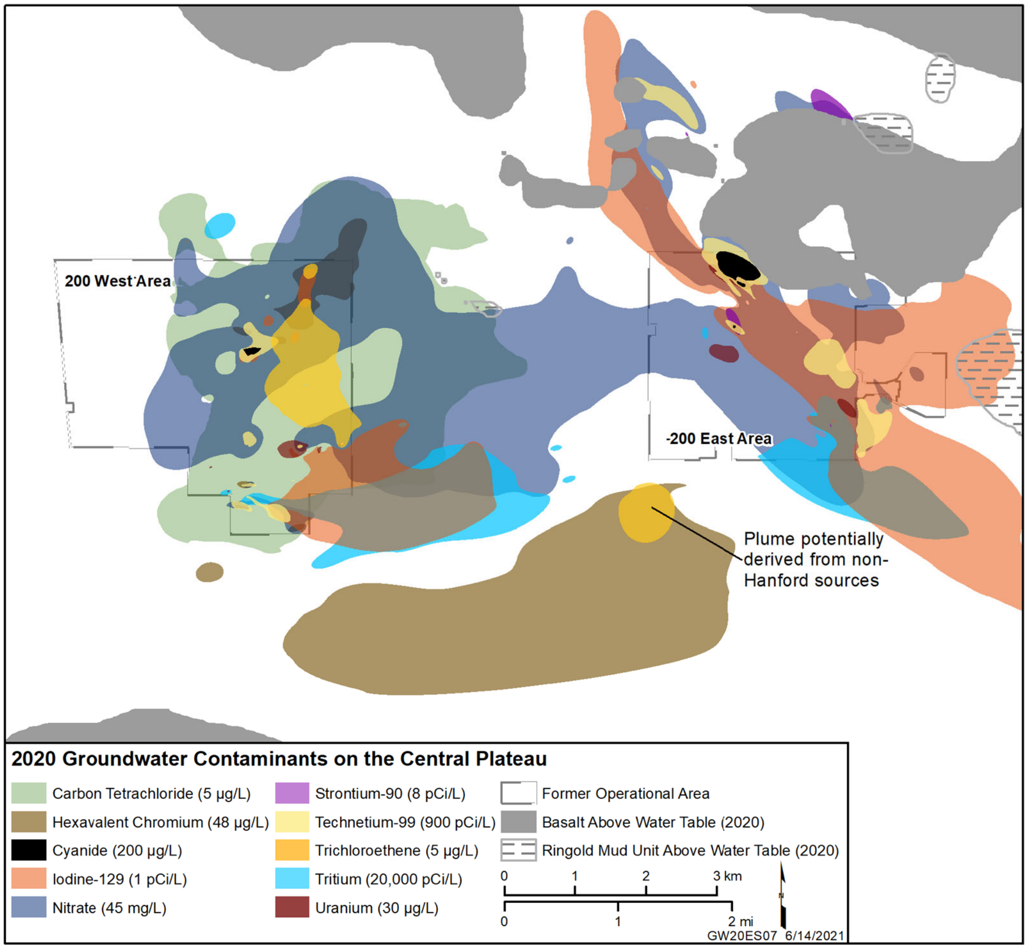

Figure 1.1. Extent of Contamination Plumes with Concentrations above Water Quality Standards on the Central Plateau (Source: DOE/RL-2020-60, Rev. 0)

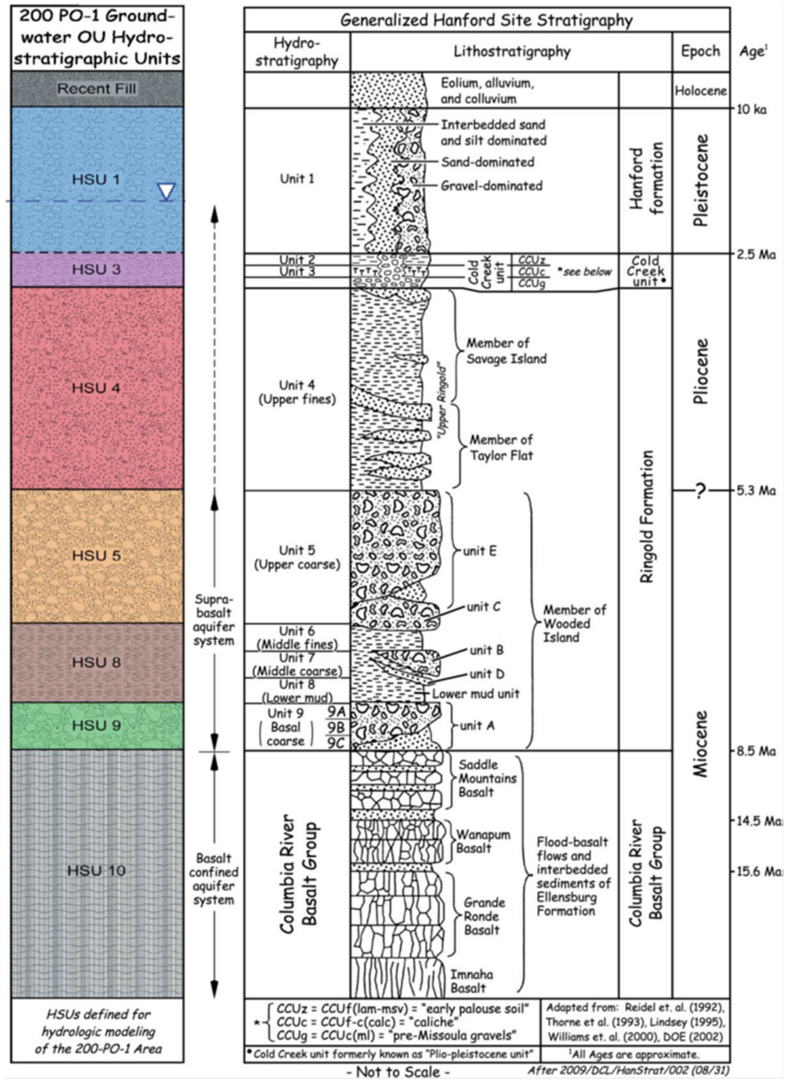

Figure 1.2. Hydrostratigraphic Units and Generalized Hanford Site Stratigraphy (Source: Hammond and Lupton 2015)

To address the spreading plumes, a P&T network was implemented as the primary active remedy for the 200 West Area under a Record of Decision (ROD) finalized in 2008. The 200 West P&T system began operation in 2012 and currently includes approximately 60–70 extraction and injection wells, with an upgraded treatment capacity of ~3,700 gallons per minute (DOE/RL-2008-78). The ROD's remediation objectives target hydraulic containment and mass reduction, aiming for a 95% reduction in total contaminant of concern mass within 25 years (EPA et al. 2008). Early operations met hydraulic containment and mass-removal targets set in the remedial design (e.g., system-wide flow rates and capture zones aligned with the 200-ZP-1 remedial design work plan). However, as the remedy progressed, new data revealed that contaminant and aquifer conditions are less favorable than initially assumed: For example, the observed abiotic degradation half-life of carbon tetrachloride is ~630 years – an order of magnitude longer than the 41-year half-life used in the original design – and significant CCl_4 mass was found in deeper portions of the aquifer previously thought to be less contaminated (PNNL-22062).

These findings indicate that the ROD cleanup levels are unlikely to be met in the planned timeframe under the original well network and operational scheme. In response, regulators and stakeholders have recognized the need for systematic performance optimization of the P&T system to ensure long-term remedial goals are met.

Optimizing the 200 West P&T system is a complex, large-scale problem constrained by both hydrogeologic complexity and computational limitations. A groundwater flow and transport model for the Central Plateau (the Plateau-to-River, or P2R, model) (Budge 2020) is calibrated to site data and used as the basis for evaluating P&T scenarios. While this high-fidelity numerical model can simulate 3D plume evolution under various pumping and injection configurations, it is computationally intensive – a single multi-year transient simulation can take hours to run on one processor.

Formal optimization of well field design would require thousands of such simulations to evaluate candidate solutions, especially under multi-objective or uncertainty-informed analyses. In the case of the 200 West Area, hundreds of decision variables would need to be considered – including where and when to install new extraction wells, screen intervals and depths, pumping schedules and rates, and when to retire or repurpose wells. The decision space is high-dimensional and highly non-linear, as the effectiveness of any configuration depends on complex interactions between pumping-induced flow fields, plume distribution, and aquifer properties. Traditional simulation-optimization approaches (e.g., coupling a numerical flow model with a global search algorithm) face prohibitive computational cost in such settings. Consequently, manual trial-and-error planning remains common – P&T expansion or adjustments are often designed through expert judgment and sequential model runs rather than brute-force mathematical optimization. This manual approach can be time-consuming and may overlook innovative well configurations that yield better performance.

Recognizing these challenges, a formal optimization study for the 200 West P&T system was initiated in fiscal year 2019 (DOE/RL-2019-38). However, applying optimization algorithms to the full P2R model proved very challenging due to high parameter dimensionality, complex site features, and the immense computational expense. There is a critical need for new computational tools that can accelerate the evaluation of well placement scenarios and guide decision-making while maintaining fidelity to the site's known hydrogeology and remedial objectives. To meet these needs, the PTOLEMY and OPTIMA tools were developed.

1.3 Overview of PTOLEMY and OPTIMA

Both PTOLEMY and OPTIMA are delivered through the SOCRATES web application, ensuring consistency with U.S. Department of Energy (DOE) Hanford Field Office (HFO) site data integration and remedy decision support for then Hanford Site. PTOLEMY is a machine-learning-powered well location screening tool that rapidly assesses the relative performance potential of extraction wells across the site. The tool uses deep learning models trained on historical operational data and simulation results from the Hanford Site model to classify which new well locations are likely to be high-performing in terms of contaminant mass recovery.

PTOLEMY provides spatially resolved probability maps showing where new extraction wells have the highest likelihood of achieving significant mass removal over a forecast window. The user interface presents these results as interactive maps and charts, allowing users to adjust performance thresholds, enforce minimum spacing from existing wells, and examine depth-specific performance profiles. By synthesizing complex model outputs and site data into intuitive visual decision aids, PTOLEMY serves as a screening tool to guide engineers and hydrologists on where additional wells might provide the greatest incremental benefit to the P&T system.

OPTIMA is a multi-objective optimization framework that systematically evaluates thousands of well field configurations to identify optimal trade-offs between competing remediation objectives. The tool employs surrogate modeling and evolutionary algorithms to dramatically reduce the computational time required for scenario evaluation – from hours per simulation to seconds. OPTIMA can be configured to optimize across multiple user-defined objectives such as minimizing cleanup time, maximizing contaminant mass removal, minimizing plume extent, or balancing containment goals. The tool enforces practical constraints—provided as inputs— like limits on new well counts, minimum well spacing, and total P&T capacity ceilings to ensure solutions remain realistic. The outcome is a suite of optimized designs representing different strategic balances, presented as Pareto frontiers and decision tables that quantify trade-offs between competing goals.

The two tools work together to accelerate P&T system optimization in a systematic workflow. PTOLEMY identifies promising well locations through rapid screening. OPTIMA evaluates comprehensive well network configurations and identifies optimized designs meeting specified objectives while respecting operational constraints. Finally, top-ranked solutions from OPTIMA are verified through high-fidelity modeling to confirm their performance and compliance with requirements. This integrated approach enables the evaluation of far more alternatives than would be feasible with traditional manual planning or direct application of optimization to full-scale numerical models, supporting more informed and effective decision-making for P&T system modifications.

1.4 Intended User and Applications

The primary users of PTOLEMY and OPTIMA consist of groundwater modelers, remediation engineers, and site decision-makers responsible for designing and managing P&T operations. The tools are intended to inform planning-level evaluations and “what-if” analyses in the context of long-term remedy optimization. PTOLEMY provides a visual, data-driven basis for discussing where new extraction wells might yield the most benefit, while OPTIMA offers a quantitative exploration of different strategic approaches. By making complex modeling results more accessible and comparative, the tools support stakeholder engagement and transparent decision-making.

For example, regulators can review OPTIMA’s Pareto front of solutions to understand the trade-off between cleanup time and mass recovery, ensuring that remedy adjustments align with regulatory objectives and risk considerations. Site operators and engineers can use the tools to prioritize operational changes based on predicted outcomes. It is emphasized that these tools are advisory in nature – they do not supersede the need for detailed design studies, regulatory approval, or field validation. Instead, they enhance the technical basis for decision-making, helping teams converge on smarter solutions faster. The generalizable design of the framework makes it extensible to other large remediation sites facing similar challenges. Ultimately, PTOLEMY and OPTIMA are meant to facilitate a more collaborative and informed planning process among modelers, decision-makers, and regulators, driving more effective and efficient cleanup through the use of advanced analytics.

1.5 Report Organization

The remainder of this report is structured to provide both the theoretical underpinnings and practical guidance for using PTOLEMY and OPTIMA. Section 2.0 presents the technical basis of the PTOLEMY tool, describing the data preparation, model architecture, training process, validation of predictive capability, and current limitations. Section 3.0 serves as a user guide for PTOLEMY, illustrating the user interface, input requirements, and examples of interpreting the well performance probability outputs. Section 4.0 covers the OPTIMA tool theory, including the development of surrogate models, configuration of optimization algorithms, and tool assumptions. Section 5.0 is the OPTIMA user guide, walking through the steps of setting up optimization runs, defining objectives and constraints, and reviewing results. Finally, Section 6.0 summarizes conclusions and Section 7.0 outlines future work, including planned enhancements to further support adaptive groundwater remediation management. Together, these sections provide comprehensive documentation of PTOLEMY and OPTIMA – from scientific foundation through user implementation – in support of more effective P&T optimization at Hanford's Central Plateau.

2.0 PTOLEMY Theory

P&T systems often require targeted placement of new extraction wells to improve contaminant mass removal. Accurately predicting mass-removal performance is challenging because outcomes depend on 3D plume distributions, pumping-induced flow fields, and heterogeneous aquifer properties. Although numerical flow and transport models can represent this complexity, they are time-consuming to calibrate and expensive to run for each potential location. When hundreds of candidates must be evaluated, comprehensive scenario analysis using full-scale models becomes computationally prohibitive.

PTOLEMY addresses this need by leveraging historical and simulated data with deep learning to predict well performance at candidate locations. The tool delivers rapid, screening-level evaluations – minutes rather than hours – while retaining fidelity to complex hydrogeologic conditions. Outputs include spatial probability maps of expected performance, depth-wise profiles, and exportable model output to support design and communication. This section provides a high-level introduction to the deep-learning methods and design principles underlying PTOLEMY; additional methodological details, data preparation steps, and validation protocols are provided elsewhere (Song et al. 2023).

2.1 Multi-Channel 3D Convolutional Neural Network Approach

PTOLEMY formulates well-performance prediction as a spatial pattern-recognition task using a multi-channel 3D convolutional neural network (MC3D-CNN). Rather than explicitly re-solving the governing flow and transport equations for each candidate site, the network uses historical and high-fidelity simulation data to learn which subsurface configurations have historically yielded successful extraction.

The MC3D-CNN treats prediction as a multi-class classification problem, labeling candidate well locations as high, medium, or low performance based on 5-year cumulative contaminant mass recovery. This categorical framing aligns with planning needs, where relative ranking of locations is typically more actionable than precise point predictions. The model processes three complementary 3D inputs ("channels") as co-registered volumetric images:

- Contaminant plume concentration (CCl_4), indicating the spatial distribution of the contaminant plume.
- Hydrostratigraphic units (geologic facies), which influence flow patterns and capture zones.
- Hydraulic conductivity, governing how readily water and contaminants can be extracted.

Each channel enters a dedicated sub-CNN for feature extraction, reflecting that plume structure, stratigraphy, and hydraulic properties carry distinct spatial signatures relevant to well performance. Features from the parallel streams are fused in deeper layers and mapped to a classification output. This integration captures interactions among factors (e.g., high concentrations in low-permeability zones may be less extractable than moderate concentrations in highly permeable zones).

To account for plume dynamics, PTOLEMY can incorporate time-stacked inputs, in which multiple plume snapshots are provided as additional channels. This enables the network to recognize temporal evolution and base predictions on evolving site conditions rather than instantaneous states. Figure 2.1 illustrates the overall MC3D-CNN architecture, with multiple 3D input channels feeding into convolutional layers and culminating in a classification output. During training of the MC3D-CNN, key hyperparameters (e.g., number of layers, filter sizes, dropout rate) were tuned via experimentation to optimize predictive performance.

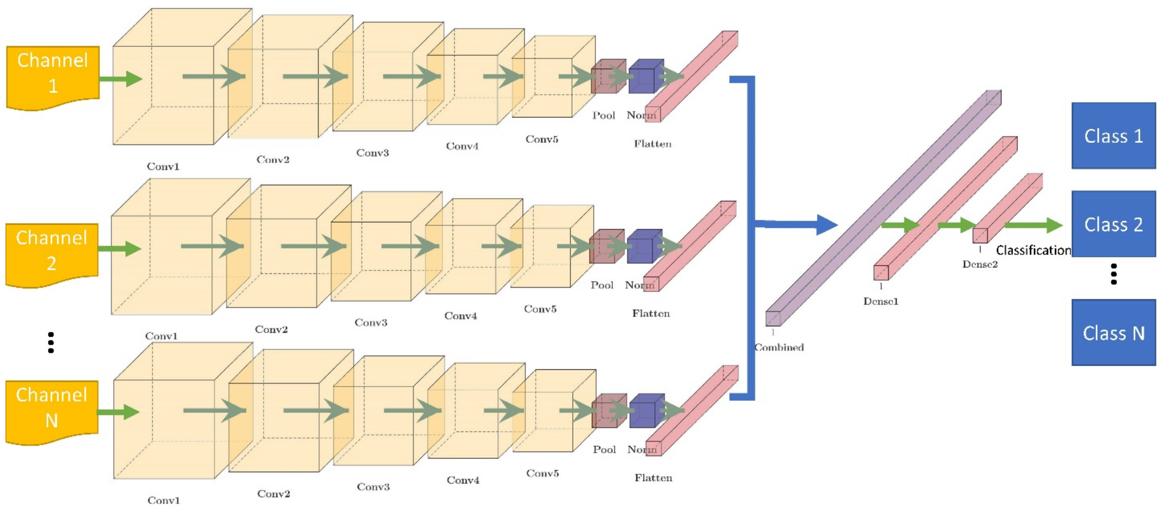

Figure 2.1. Architecture of the MC3D-CNN Classification Model

2.2 Training Data Development and Performance Classification

Training data were derived from the calibrated P2R flow and transport model for the 200 West Area. For each existing well, a ~1000 m × 1000 m × 80 m 3D window centered on the well was extracted, with channels for (a) CCl_4 concentration, (b) hydrostratigraphic unit identifiers, and (c) hydraulic conductivity. Geologic and conductivity fields are static in time, while the plume channel can be sampled at different time steps to represent transient conditions. This pairing allows the model to learn how heterogeneity, aquifer properties, and plume distribution jointly control well performance under realistic settings. Figure 2.2 summarizes the data preparation and training workflow: Multi-channel 3D images are constructed per well and assigned labels based on 5-year mass recovery. This systematic approach ensures that every training example contains the full spatial context around a well location along with its known performance outcome.

Performance labels for each well are defined by 5-year cumulative CCl_4 mass recovery, yielding standardized comparability across locations and times. Wells are categorized into high-, medium-, or low-performance classes using thresholds determined from the empirical distribution of recovery outcomes to produce a balanced classification (approximately one-third per class).

Because a well's performance category can change over time (e.g., early high performance followed by decline as local mass is depleted), the dataset includes multiple time segments per well. Each well contributes several samples, each with a corresponding 5-year future recovery computed at different start times. In total, hundreds of labeled images (e.g., 486 samples for 28 wells over 18 time steps) were generated for training and validation, providing a large dataset with temporal context on performance evolution.

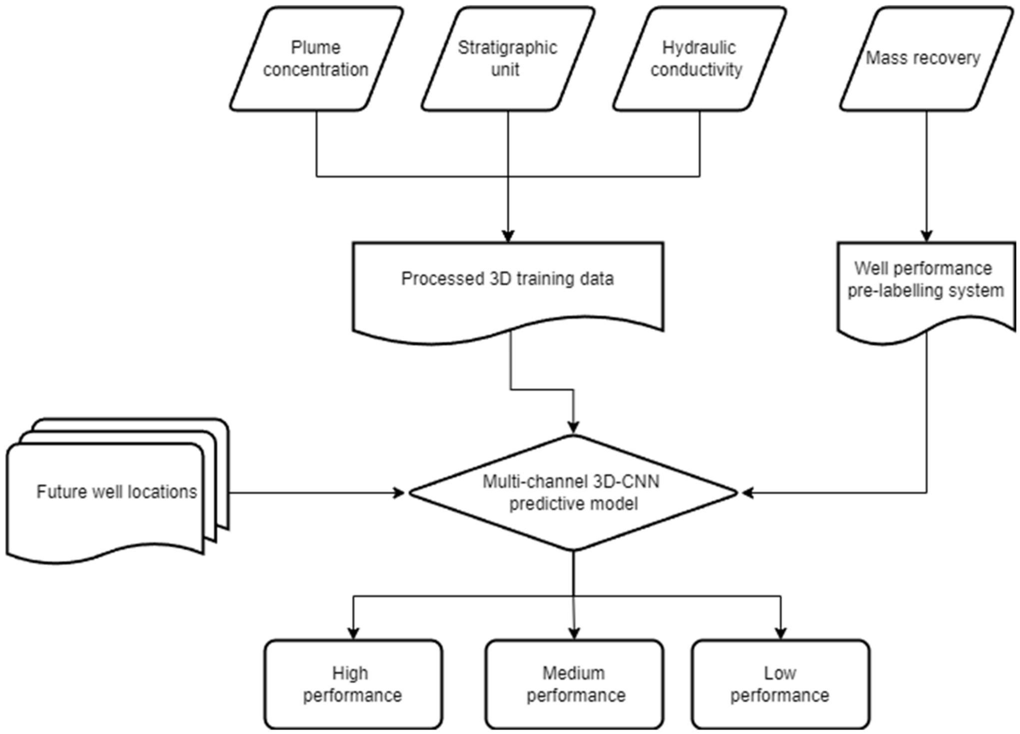

Figure 2.2. Deep Learning Workflow for Well Performance Classification, from Training Data Generation to MC3D-CNN Prediction

2.3 Model Training and Performance Evaluation

The MC3D-CNN was trained using supervised learning with a categorical cross-entropy loss, batch normalization, and dropout to promote generalization and mitigate overfitting. Data were partitioned into training and independent test sets to evaluate performance on unseen wells and future time segments. A grid search over hyperparameters (network depth, convolutional filters, learning rate, etc.) identified the final configuration.

Predictive skill was assessed using standard multi-class metrics derived from the confusion matrix, including overall accuracy, precision for the high-performance class, sensitivity (recall), and specificity. The trained classifier demonstrated >90% accuracy in ordering wells by performance in validation and >80% precision in identifying high-performing locations on held-out tests. These results indicate that the model reliably distinguishes performance categories and, critically for planning, flags high-performing candidates with few false positives – supporting confident use of "high" classifications in screening.

2.4 Decision Criteria and Constraints

While the MC3D-CNN provides a probabilistic ranking of well performance across the domain, PTOLEMY incorporates additional practical criteria to transform these predictions into actionable decision support for remediation planning, including:

1. **Probability threshold**: Users specify the minimum probability for a candidate to be considered "high performing." The default is 50%, highlighting locations with at least a one-in-two likelihood of falling into the high class. Thresholds can be raised (more conservative, fewer candidates) or lowered (more inclusive) to reflect risk tolerance and planning needs.

2. **Minimum spacing from existing wells**: To avoid hydraulic interference and redundant well placement, candidates within a user-defined radius of existing extraction wells are excluded. A default spacing on the order of hundreds of meters is applied and can be adjusted based on site-specific considerations.

Operationally, PTOLEMY identifies all grid locations in the model domain that meet the criteria (exceed the probability threshold and satisfy spacing requirements) and marks these as potential well sites. Clusters of adjacent high-probability cells may be aggregated to suggest an optimal location (representative point) within that high-performing zone. This aggregation process helps avoid suggesting multiple wells in essentially the same high-performing area while preserving information about the spatial extent of favorable conditions.

2.5 Current Limitations and Assumptions

Several limitations and assumptions should guide interpretation and use:

- **Single contaminant and site specificity**: The current model is trained and demonstrated for CCl_4 in the 200 West Area. Extending to other contaminants or sites requires adequate training data (simulations and/or historical records) and model retraining. While the framework is adaptable, each new application entails non-trivial development.
- **Dependence on baseline operations**: Predictions are conditioned on the hydrologic and operational context represented in the training data (e.g., assumed configurations of injection/extraction and pumping rates). The current implementation does not explicitly vary pumping schedules or injection patterns; substantial deviations from the training scenario may alter real-world performance in ways the model does not capture.
- **Uncertainty treatment**: PTOLEMY presently provides deterministic probability estimates without formal propagation of subsurface uncertainty (e.g., plume distribution, hydraulic properties). In practice, the true state is uncertain. A planned enhancement is ensemble prediction using multiple geologic and plume realizations (e.g., geostatistical simulation) to better reflect uncertainty.
- **Training data provenance**: Although extensive, part of the training dataset is derived from a numerical model rather than direct field observations. Performance therefore depends on how well the inputs (plume, geology, hydraulic conductivity) represent actual conditions. The framework does not inherently enforce physical laws beyond what is embedded in the training simulations, so extreme out-of-distribution conditions may reduce reliability – underscoring the importance of representative training ranges.

Despite the limitations described above, the data-driven approach in PTOLEMY offers a powerful and rapid first-pass assessment to narrow down candidate well sites for more detailed evaluation. The tool transforms what would be a computationally prohibitive comprehensive analysis into a manageable screening process, enabling more systematic and informed decision-making about well-placement priorities. Ongoing development is addressing these limitations by incorporating additional data types and adaptive retraining capabilities as more monitoring data become available, ensuring that the tool continues to evolve with improving site understanding.

3.0 PTOLEMY User Guide

3.1 Interface Overview

The PTOLEMY decision-support module is delivered via an interactive web application organized in two panes (Figure 3.1): (1) Performance Probability Distribution Graph Panel (left side), which displays well-performance results; and (2) Map Panel (right side), which displays injection and extraction wells, candidate wells, and contaminant plume contours. Some default analysis options are automatically selected upon opening the module. Users may adjust the High-Performance Threshold and add additional well locations for analysis (described in later sections).

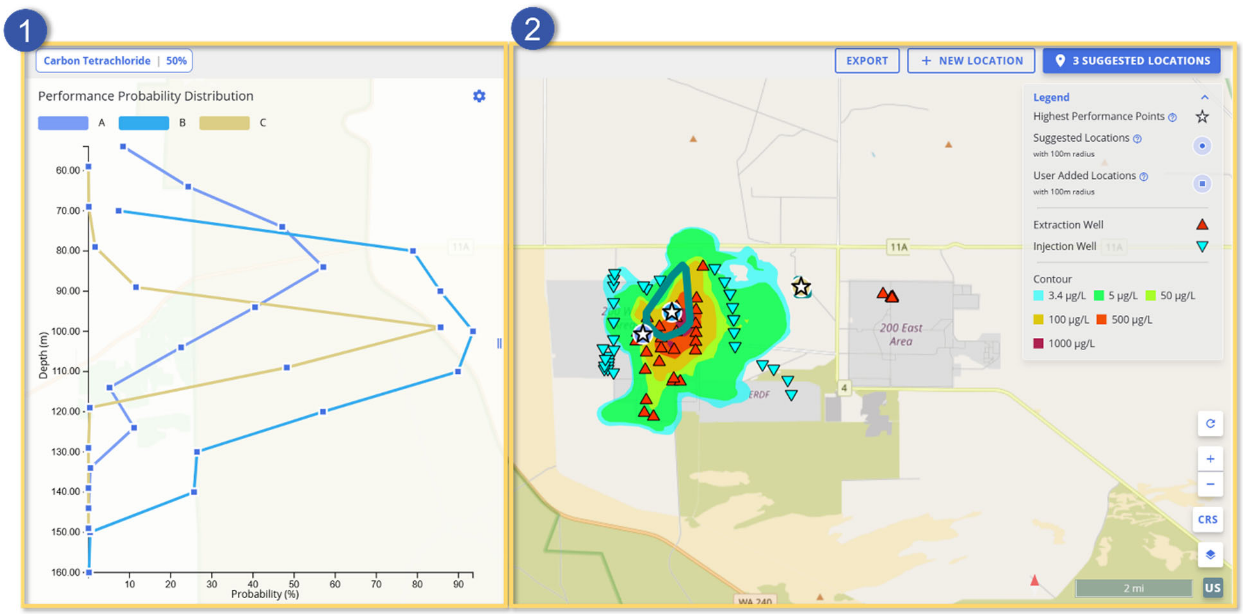

Figure 3.1. PTOLEMY User Interface

3.1.1 Map Panel

The right-hand map panel provides a spatial view of the site and candidate well information. Map options include (Figure 3.2):

1. Reset view – zoom extents to default
2. Zoom in/out navigation
3. Coordinate system – showing cursor location coordinates
4. Layers – including basemaps and module-specific, contaminant plume, and extraction/injection well layers
5. Map units – to switch measurement units

The legend in the upper-right corner of the map panel identifies the icons for candidate well locations, existing extraction/injection wells, and plume contour displays. Contextual layers specific to PTOLEMY can be turned on or off, including current CCl_4 plume concentration contours, locations of existing extraction/injection wells, and any proposed candidate wells. Users can pan and zoom the map to examine different areas of the site in detail.

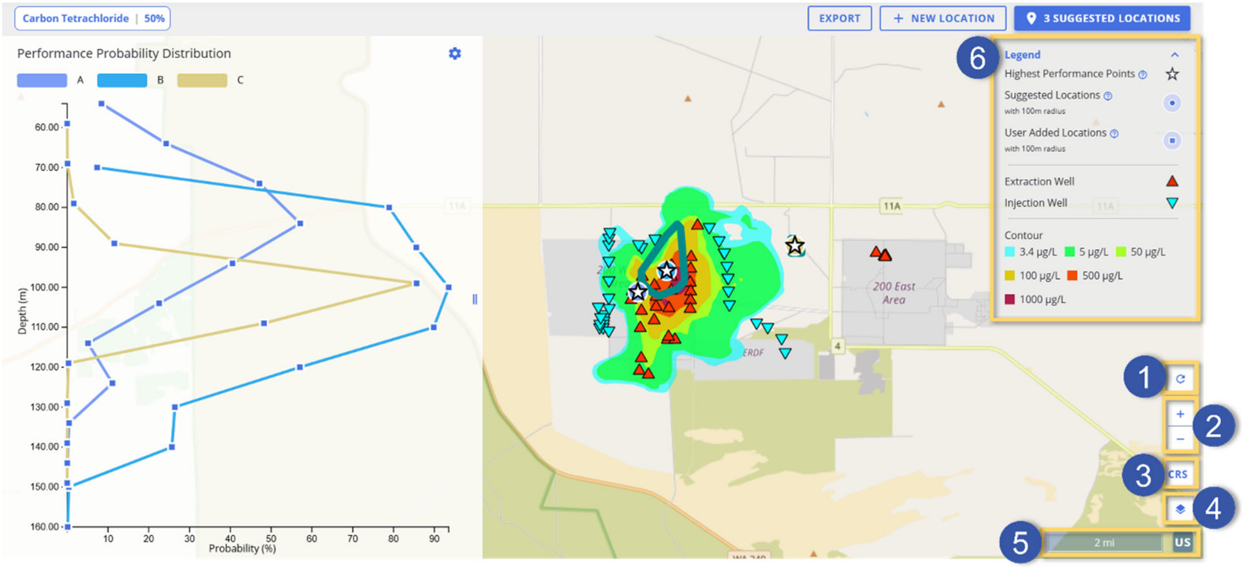

Figure 3.2. Map Options in PTOLEMY

3.1.2 Performance Probability Distribution Graph Panel

The left-hand panel shows the model-predicted performance profiles for any candidate wells under consideration. By default (upon opening the module), this graph is empty until wells are selected or added. Once candidate well locations are chosen (either automatically suggested or user-added), the graph displays the predicted probability of achieving high performance as a function of depth (or elevation) for each well. By default, the chart displays probability on the x-axis and depth (m) on the y-axis, forming a vertical profile of performance likelihood for each well. Each candidate well is represented by a distinct, colored line on the chart, and an identifier (e.g., "Well A," "Well B") appears in the legend to distinguish them. The graph provides insight into how the predicted success of a well may vary with well-screen depth – for instance, a well may have a high (>90%) probability of being high-performing at a depth of 100 m but near 0% at a depth of 160 m, indicating the optimal intake interval for that well.

3.2 Interface Overview

3.2.1 High Performance Thresholds and Contaminants

Upon opening the PTOLEMY module, the default threshold for the likelihood of high performance is set to 50%, meaning that all displayed candidate wells are at least 50% likely to be high performing. The default threshold can be changed by clicking the button in the upper-left corner of the window and adjusting the High-Performance Threshold slider (Figure 3.3). Increasing the threshold (e.g., to 70% or 90%) will apply stricter criteria, resulting in fewer suggested wells (only the most confident high-performing locations remain). Conversely, lowering the threshold will include more candidate wells (casting a wider net). The map and graph update dynamically when the threshold is changed.

Currently, the only contaminant available for analysis is CCl_4.

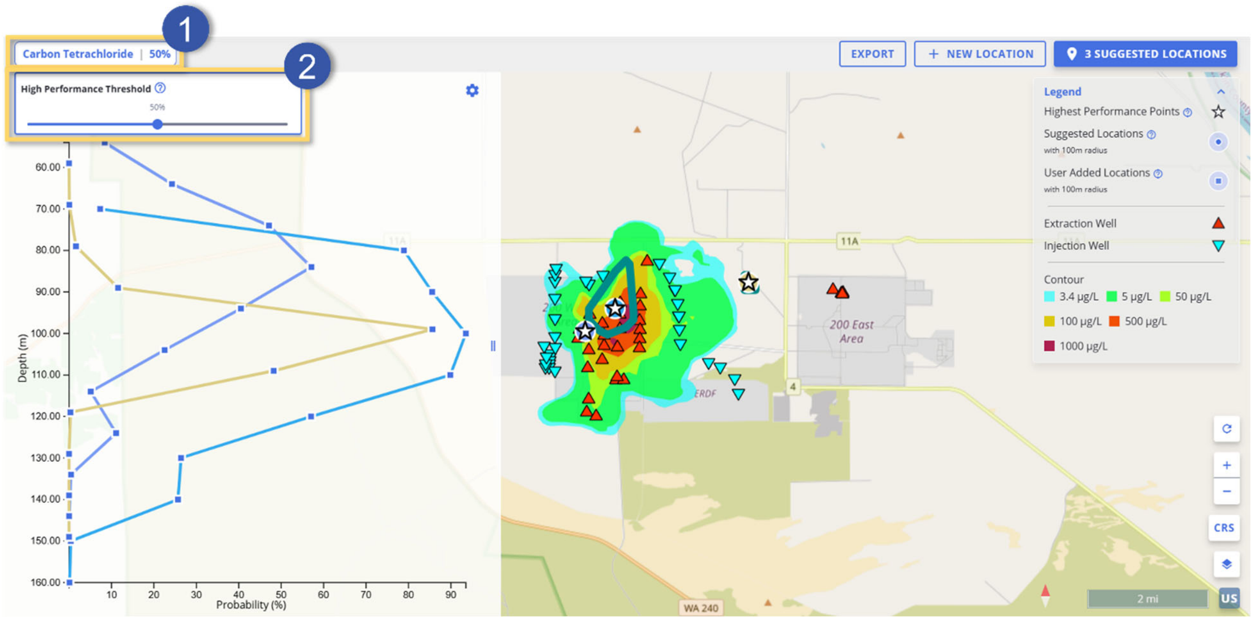

Figure 3.3. High-Performance Threshold Adjustment Options

3.2.2 Suggested Well Locations

Based on the current threshold and spacing settings, PTOLEMY will automatically suggest certain new well locations that meet the criteria. For example, using the default 50% high-performance threshold, three well locations are suggested in Figure 3.4. Stars represent the suggested well locations, and blue outlines represent clusters of high-performing wells. The star is typically placed at a representative central point of each high-performing cluster. This visualization helps the user see both specific suggestions and the broader zones that have high potential. If the user raises the threshold, some stars may disappear (only the highest-probability locations remain); if the threshold is lowered, additional stars may appear as more locations qualify.

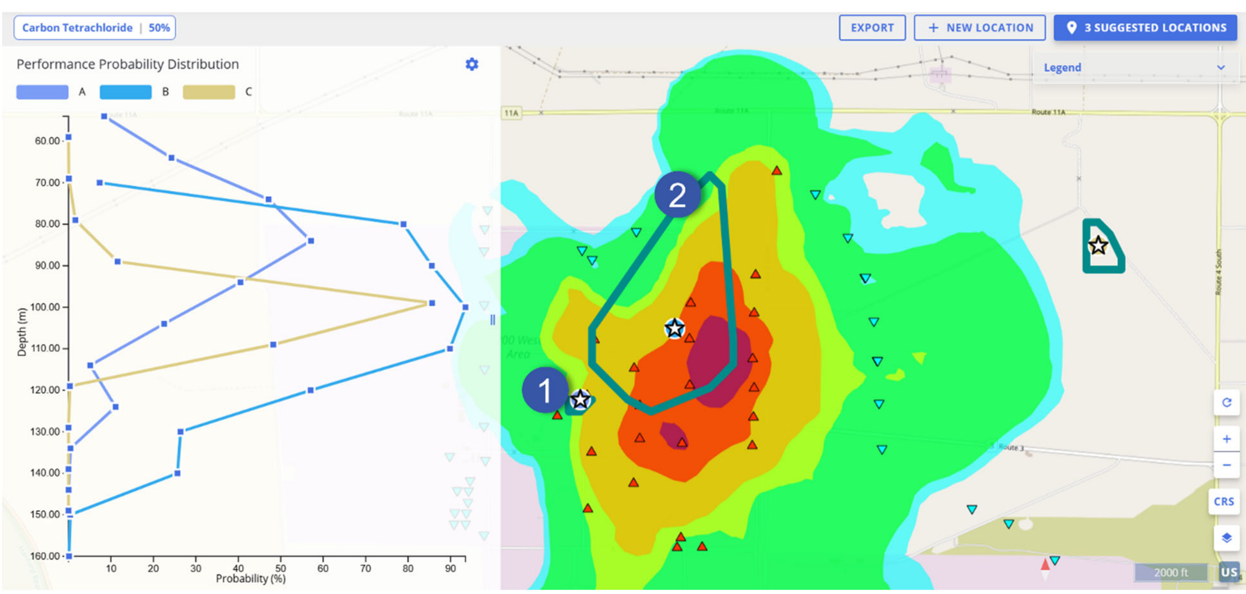

Figure 3.4. Suggested Well Locations on Map

The coordinates of the suggested wells can be viewed by clicking the "[X] Suggested Locations" button in the upper-right (Figure 3.5). This opens a panel that lists each suggested well with its map coordinates. The panel also provides an option to adjust the minimum distance from existing wells used to generate suggestions. The distance from the nearest extraction well can also be adjusted in this window. Hovering over a suggested well location on the map displays the shaded boundary for the specified distance (Figure 3.6). This feature helps users visually ensure that new wells will not overlap with current wells' capture zones.

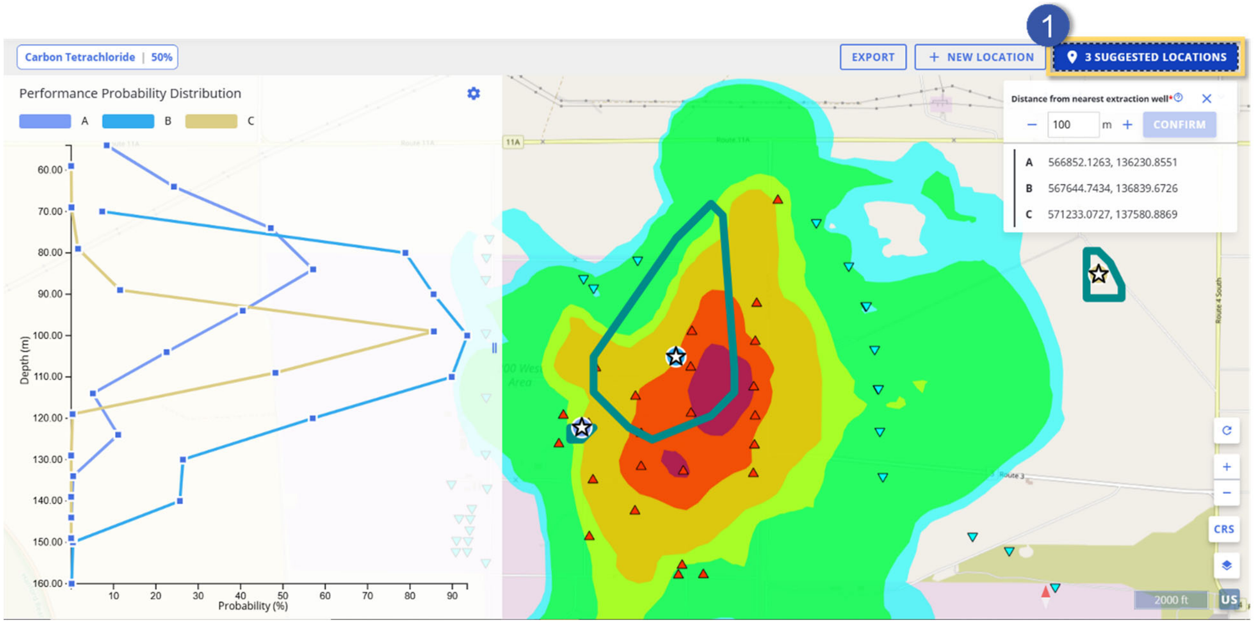

Figure 3.5. Suggested Well Locations and Distance from Nearest Extraction Well Options

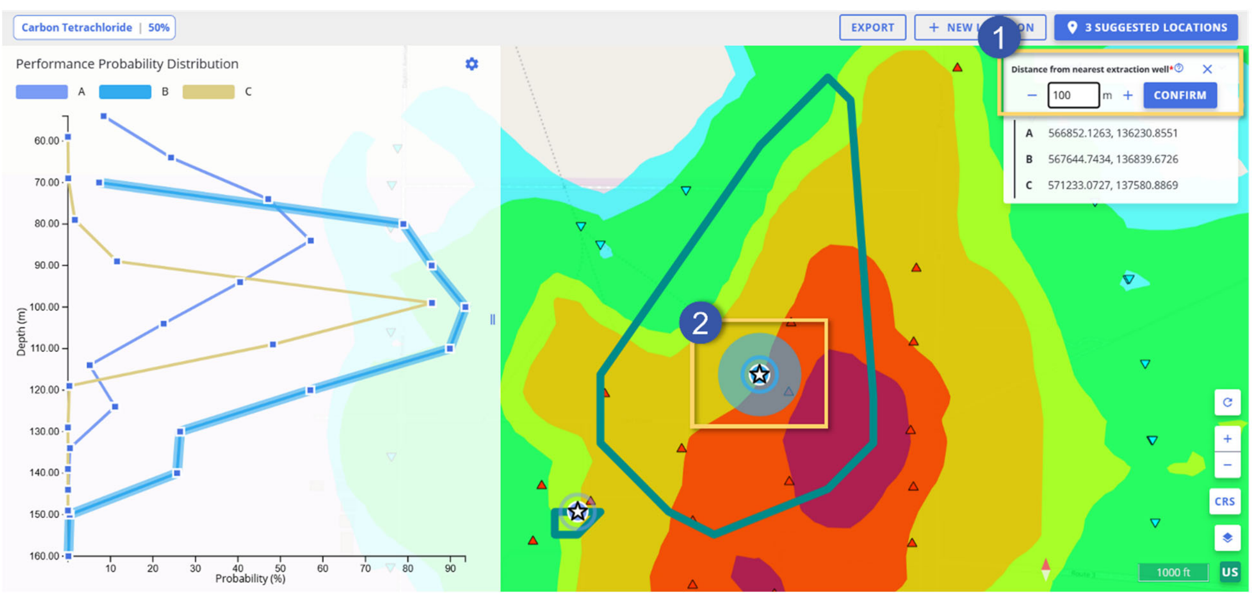

Figure 3.6. Distance From Nearest Extraction Well Display on Map

3.2.3 Adding New Well Locations

PTOLEMY not only suggests wells but also allows users to manually test custom locations. To add a new well location, click the "+ New Location" button located in the upper-right area of the window, then click a location on the map within the blue borders that represent high-performing well areas (Figure 3.7). User-added wells appear on the map as diamond markers, differentiating them from the star icons of auto-suggested wells. The new well is automatically assigned a label (e.g., "Well A" if it is the first, "Well B" second), and it is immediately incorporated into the performance analysis.

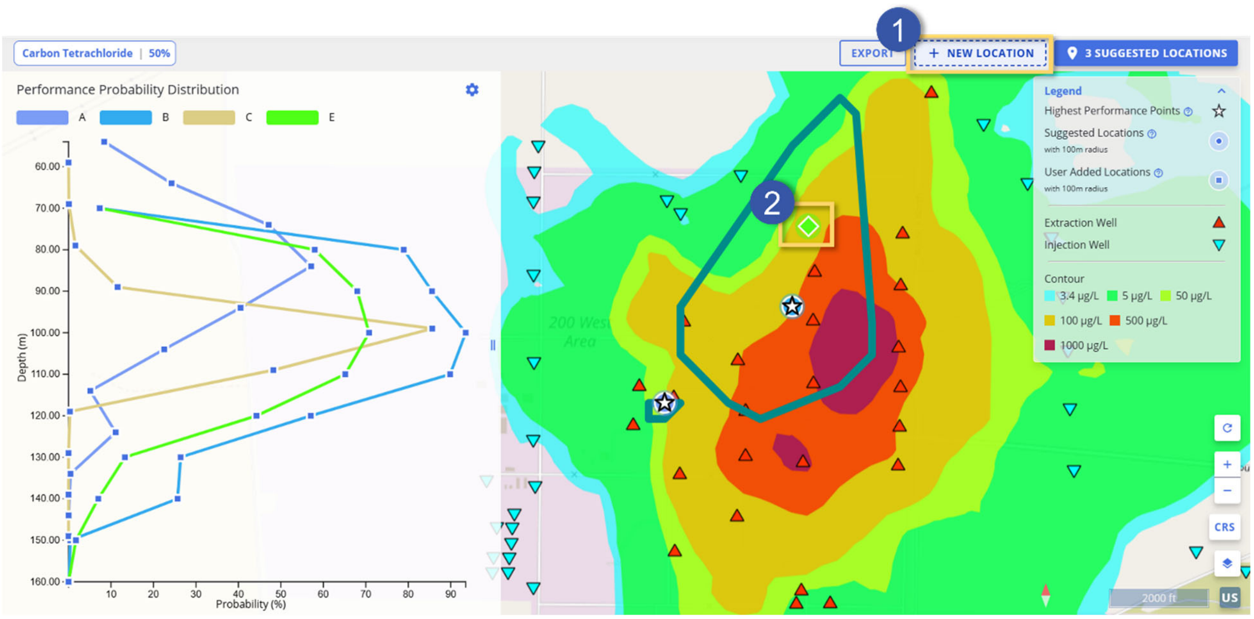

Figure 3.7. Adding a New Well Location to Map

3.3 Analysis Results

The performance results of candidate wells, including both the suggested and user-defined well locations, are displayed in the Performance Probability Distribution panel on the left (Figure 3.8). The well ID is displayed in the legend at the top of the chart. The chart shows that Well A, for example, has a ~90% likelihood of high performance at a depth of 100.0 m. The same well, at a depth of 160.0 m, has a 0% likelihood of high performance. This depth-based performance profile helps in determining the optimal screen interval for each candidate well.

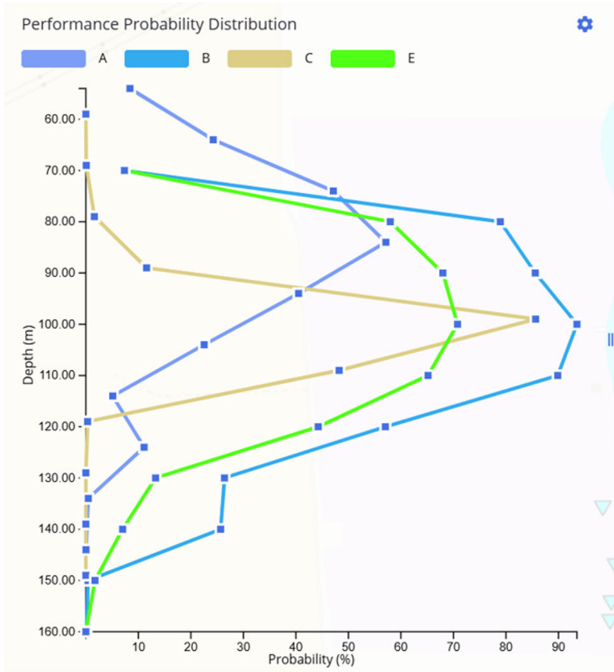

Figure 3.8. Performance Probability Distribution Graph

3.3.1 Interactive Chart–Map Linking

The PTOLEMY user interface is interactive, linking the map and chart for intuitive exploration. Hovering over a line on the chart highlights the corresponding well on the map, and vice versa (Figure 3.9). This linked highlighting makes it easy to connect a graph curve to its spatial location. It is particularly useful when multiple wells are on the graph – the user can identify which curve belongs to which well by simply selecting it in one view or the other. Users can also remove wells from consideration (e.g., using a delete option in the legend or by right-clicking a marker), and both the map and chart will update to reflect the removal.

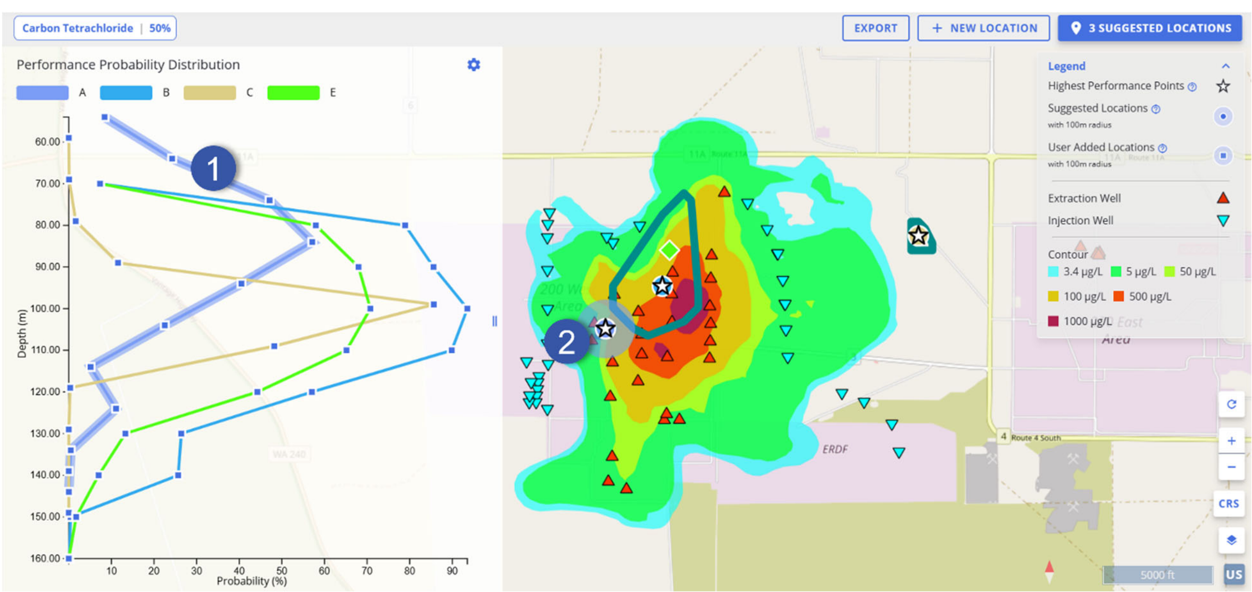

Figure 3.9. Corresponding Chart Lines and Map Well Locations

3.3.2 Chart Settings

The chart's y-axis can be set to "depth" (default setting) or "elevation" by clicking the settings icon the chart panel and selecting the desired y-axis (Figure 3.10). If the user prefers to view elevation above sea level instead of depth below ground, the graph will invert accordingly (since higher elevation corresponds to shallower depth).

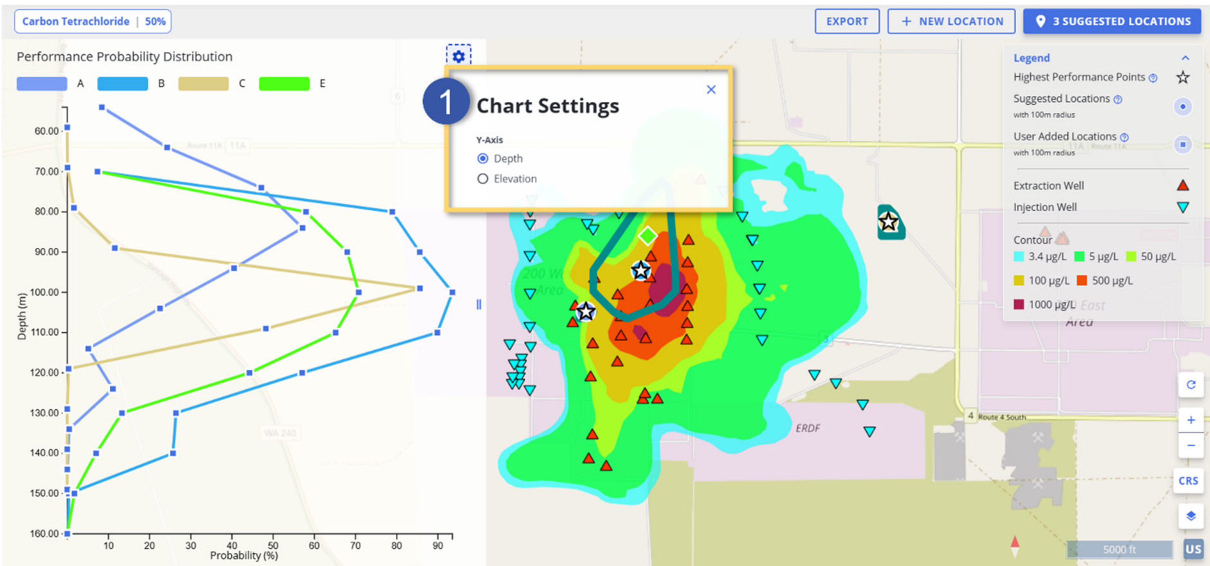

Figure 3.10. Changing Graph Y-Axis

3.4 Exporting Results and Reports

PTOLEMY includes features to export the analysis results for reporting or further use. Data can be exported from PTOLEMY by clicking the "Export" button in the upper-right of the window (Figure 3.11). The export dialog (Figure 3.12) allows users to choose the output format and content to save. Data can be exported in comma-separated value (CSV), Portable Network Graphic (PNG), or Portable Document Format (PDF) formats. Specific locations and types of data can be included or excluded in the export. Supported export formats include:

- **PNG image**: a screenshot of the current map and/or graph view, useful for embedding in reports or presentations.
- **CSV data**: a comma-separated values file containing the underlying data (e.g., the list of suggested well coordinates with their probability values and the depth–probability profile data points for each well).
- **PDF report**: a formatted document capturing the key visuals (map and chart) along with summary information and metadata.

When exporting, users can select which elements to include. The metadata accompanying an export typically includes the date of analysis, the high-performance threshold used, the minimum spacing used, and descriptions of any user-added well locations. If exporting data, the CSV will list each candidate well with coordinates, depths, and predicted probabilities. If exporting an image, the legend and any annotations visible on-screen will be included.

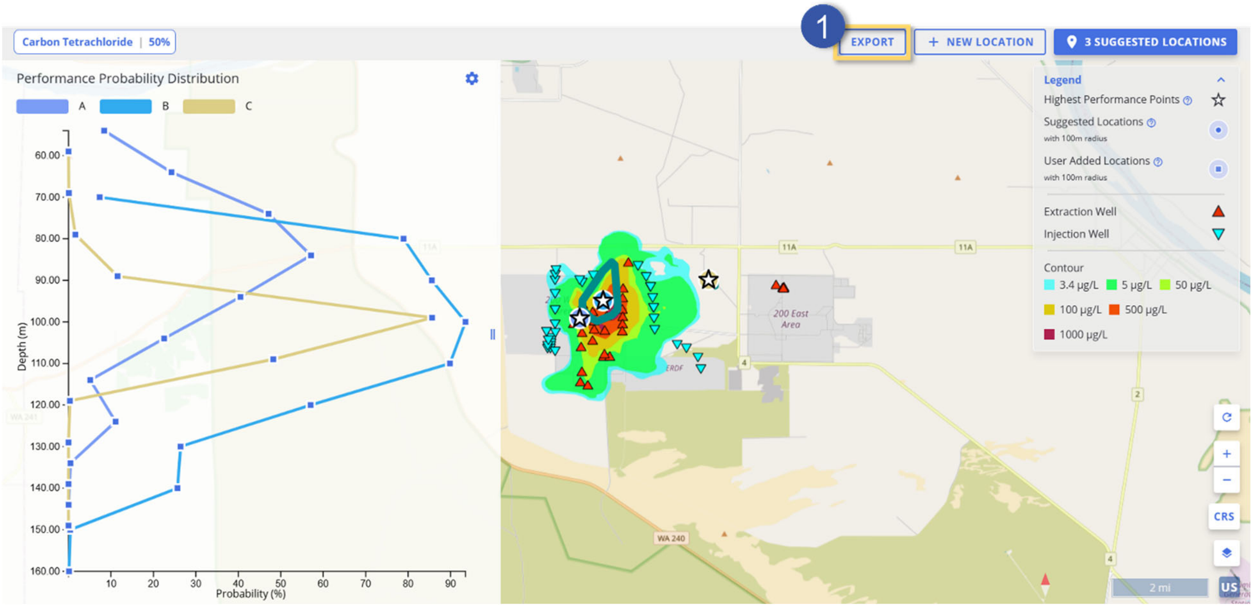

Figure 3.11. Data Export Button

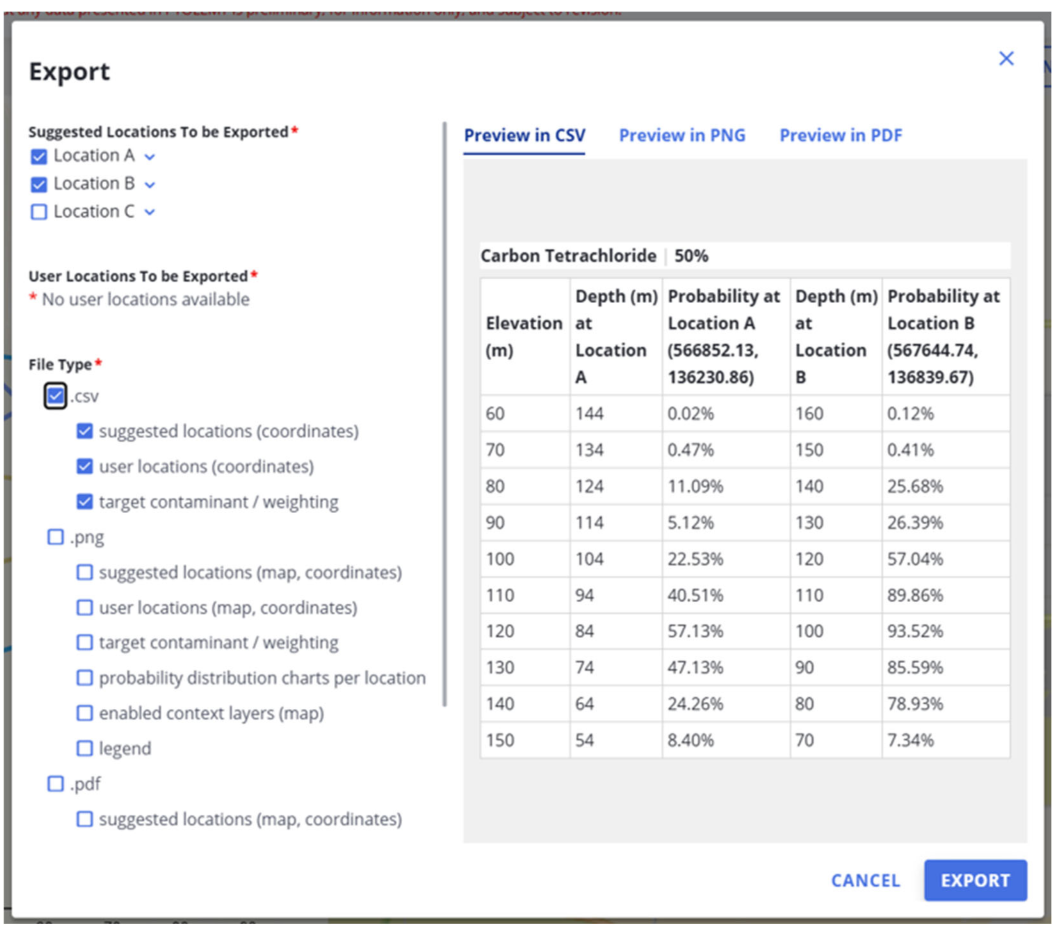

Elevation (m)	Depth (m) at Location A	Probability at Location A (566852.13, 136230.86)	Depth (m) at Location B	Probability at Location B (567644.74, 136839.67)
60	144	0.02%	160	0.12%
70	134	0.47%	150	0.41%
80	124	11.09%	140	25.68%
90	114	5.12%	130	26.39%
100	104	22.53%	120	57.04%
110	94	40.51%	110	89.86%
120	84	57.13%	100	93.52%
130	74	47.13%	90	85.59%
140	64	24.26%	80	78.93%
150	54	8.40%	70	7.34%

Figure 3.12. Data Export Options

3.5 Summary

The PTOLEMY user interface provides an intuitive workflow allowing the user to set the performance threshold, review the suggested well locations on the map, optionally adjust spacing or add new locations, examine the performance probability profiles for the candidate wells, and then export the findings. This interactive tool allows technical staff to quickly evaluate where new extraction wells might yield the highest mass recovery under the complex site conditions – grounded in the validated deep-learning model described in Section 2.0 – yet packaged in an accessible decision-support format. The combination of spatial visualization and depth profiles, along with user-driven scenario adjustment, makes PTOLEMY a useful addition to the groundwater remediation planning toolbox.

4.0 OPTIMA Theory

This section provides a high-level introduction to the deep-learning methods and design principles underlying OPTIMA; additional methodological details, data preparation steps, and validation are provided in Song et al. 2025.

4.1 Framework Overview

OPTIMA formulates P&T design as a multi-objective optimization problem. The framework combines three integrated components: (1) a physics-informed deep-learning surrogate for fast groundwater simulation; (2) an evolutionary optimizer (Non-dominated Sorting Genetic Algorithm II, NSGA-II) to explore design trade-offs; and (3) a setup module for defining scenario inputs, objectives, and constraints, as shown in Figure 4.1. Combining these three components enables users to specify remedial goals and automatically identify wellfield strategies that best balance those objectives through physics-informed modeling, deep-learning surrogates, and evolutionary search.

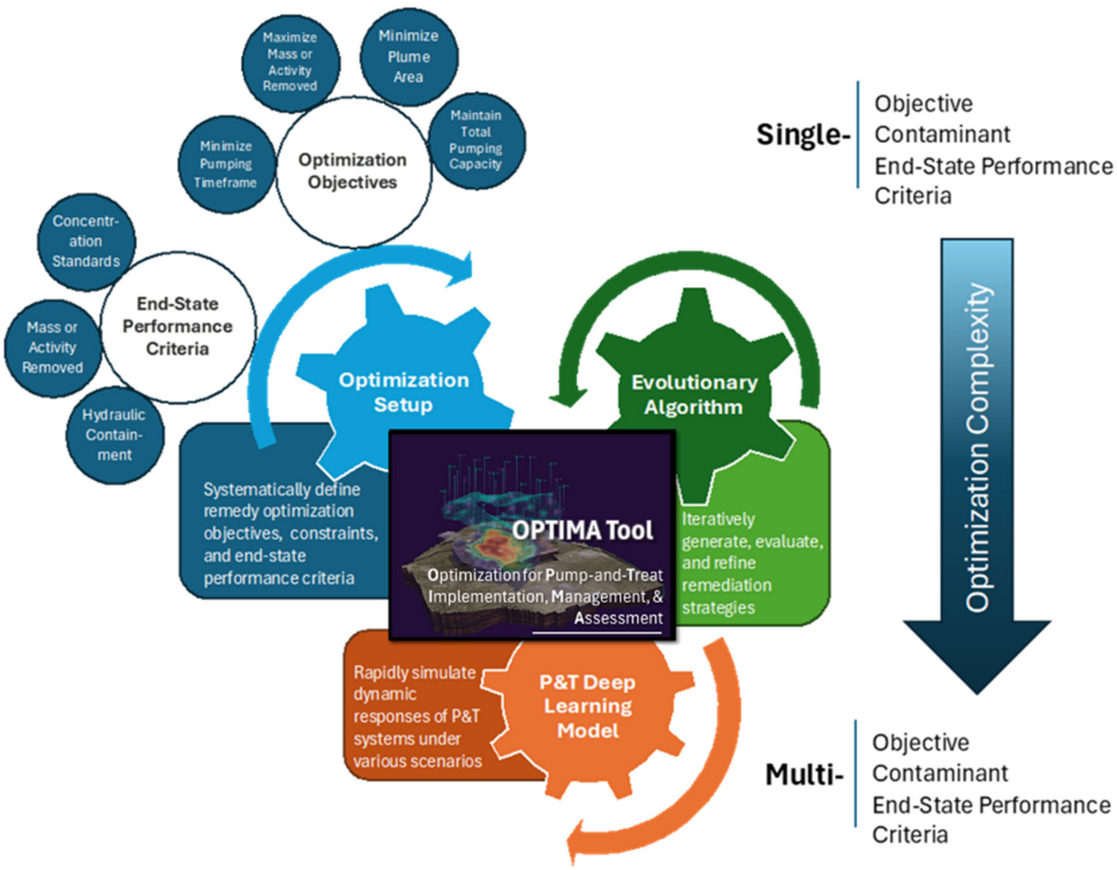

Figure 4.1. OPTIMA Computational Framework Illustrating the Integrated Workflow

At a high level, the user defines decision variables (e.g., which candidate wells to install or retire, extraction/injection mode, and schedule/rate bounds) and selects the remediation objectives of interest (e.g., minimize cleanup time, maximize mass removed, minimize plume extent). The U-Net surrogate, augmented with a steady-state drawdown channel (using the Thiem equation), rapidly predicts transient groundwater response for each candidate scenario. NSGA-II then evolves a population of wellfield configurations toward a Pareto-optimal set that balances the chosen objectives while enforcing site-realistic constraints (e.g., minimum well spacing, drawdown limits, treatment capacity caps). The framework supports multi-contaminant planning by allowing contaminant-specific end-state criteria and objective definitions. The outcome is a family of alternatives (the Pareto front) that makes trade-offs transparent for decision-makers.

4.2 Surrogate Modeling: U-Net with Physics-Informed Well Drawdown Channel

4.2.1 U-Net Surrogate for Multi-Variable Prediction

At OPTIMA's core is a 2D/d3D U-Net convolutional neural network trained to emulate the groundwater flow and transport simulator (Figure 4.2). The U-Net uses a symmetric encoder–decoder with skip connections to capture multi-scale spatial patterns while preserving fine detail. To better retain plume structure, max-pooling layers were replaced with strided convolutions, down-sampling feature maps without discarding spatial information. The output layer employs a custom LeakyClippedTanh activation that constrains predictions to physically realistic ranges: A leaky tanh preserves small gradients for negative values, followed by a clipping step that zeroes any small negative concentrations to ensure non-negative contaminant predictions. The surrogate produces multiple outputs: Separate U-Net models were developed to predict distributions of CCl_4, Tc-99, and hydraulic head, enabling parallel evaluation of co-occurring plumes and groundwater levels under a consistent architecture.

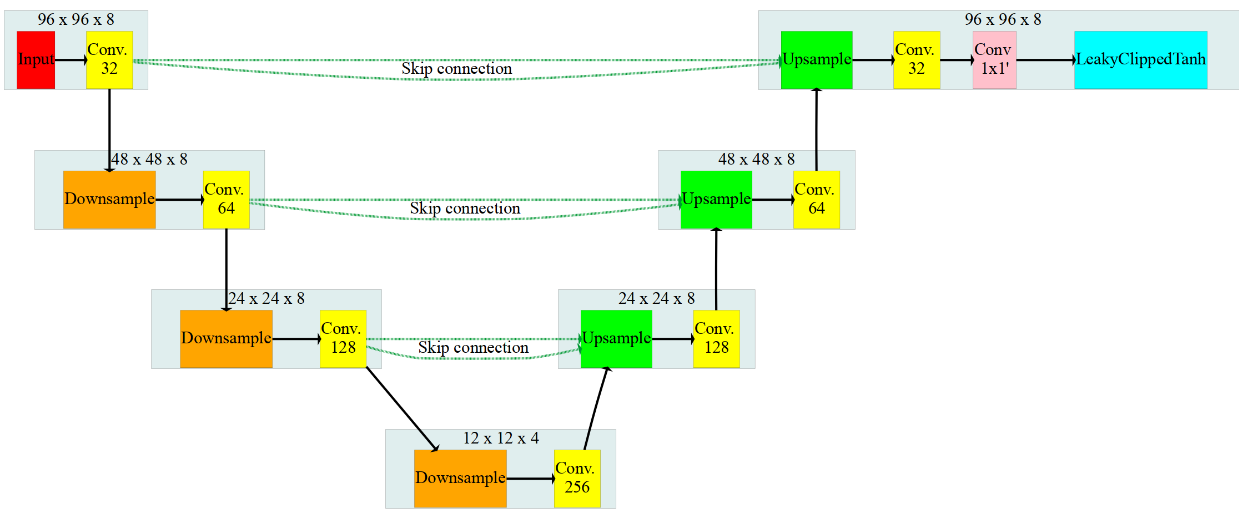

Figure 4.2. U-Net Deep-Learning Framework Used by OPTIMA

4.2.2 Input Representation and Steady-State Drawdown Channel

For the 200 West application, the surrogate operates on a spatial domain gridded at 100 m lateral resolution and multiple vertical layers. Model states are encoded as a 96 × 96 × 8 voxel tensor for plume concentration, plus 2D fields for hydraulic head. At each time step, the network ingests a stacked set of "images": maps of hydraulic head and contaminant concentration at the current time, static hydrogeologic properties, and active well controls.

A key innovation is a physics-informed head-change input that encodes steady-state drawdown (and mounding) induced by continuous extraction and injection. We assume steady-state, homogeneous, confined-aquifer conditions with fully penetrating wells for this computation; although simplified, these assumptions are appropriate for long-term, continuously operated P&T systems where transient buildup/relaxation is minor over planning horizons. Rather than passing sparse well coordinates/rates, OPTIMA pre-computes a steady-state drawdown field using the Thiem equation (Thiem 1906) with superposition. Superposition sums the drawdown (or mounding) effects from multiple wells, reflecting the additive impact of each well's pumping or injection. The resulting field provides an effective first-order representation of the head-change pattern. As a result, the surrogate uses this physics-informed drawdown input as a heuristic feature map to capture first-order head changes without explicitly modeling complex transients. For each model grid cell, the Thiem equation and principle of superposition are applied to compute the total drawdown h at this location:

$$h = h_0 - \sum_{i=1}^{n} \left[\frac{Q_i}{2\pi T} ln \left(\frac{r_i}{r_{0,i}} \right) \right] \tag{4.1}$$

where h_0 is the initial hydraulic head $[\mathrm{L}]$, Q_i is pumping rate $[\mathrm{L}^3\mathrm{T}^{-1}]$, T is transmissivity of the aquifer $[\mathrm{L}^2\mathrm{T}^{-1}]$, r_i is the radial distance from the well to the point of interest $[\mathrm{L}]$, $r_{0,i}$ is the reference/influence radius $[\mathrm{L}]$, and i is the index of well [-]. These parameters are largely based on empirical data. Injecting this physics-based estimate informs the model of first-order hydraulic effects – local head depressions and induced capture zones – without requiring the network to infer them purely from data. This guided input stabilizes training and improves the surrogate's ability to reproduce advection patterns under pumping.

4.2.3 Training Data and Performance

The surrogate was trained on an extensive dataset generated with eSTOMP. Approximately 10,000 numerical scenarios were simulated, each spanning 12 years (2025–2037) with varied well placements, pumping rates, and schedules. From these runs, >120,000 snapshot pairs of system state were collected (annual head and concentration maps) as training examples. Each sample consisted of input features (heads, plumes, and computed drawdown at time t) and outputs at t + 1. The model was trained on 70% of the samples (15% validation, 15% testing) using Adam (Adaptive Moment Estimation) optimizer (learning rate tuned in the 10^{-4}–10^{-2} range) with mini-batches of 32–256. Early stopping prevented overfitting, and performance was tracked via root mean square error (RMSE) on the validation set.

The final surrogate demonstrated high predictive accuracy and temporal stability over the 12-year horizon. Deployed autoregressively (feeding each prediction back as the next input), error growth remained minimal, indicating reliable tracking of plume evolution. For example, predicted heads had RMSE on the order of 0.02–0.07 m relative to the physics-based model; CCl_4 predictions had RMSE increasing from ~0.3 to 1.6 µg/L over 12 years – below the 3.4 µg/L cleanup criterion; Tc-99 predictions had RMSE ~0.17–0.76 pCi/L, well under the 900 pCi/L drinking-water standard. Speed is equally important: A single 12-year plume simulation via U-Net inference takes ~0.5–0.6 s on CPU, versus minutes for the full 3D model on high-performance computing – roughly a 10^{3x} speed-up. This enables the optimizer to evaluate thousands of scenarios per run. In short, the surrogate preserves the essential physics and accuracy needed for decision-making while enabling near-real-time scenario evaluation – a prerequisite for efficient, formal optimization in a large decision space.

The surrogate model will continue to be updated to remain consistent with the site contractor's latest model (e.g., P2R).

4.3 Objective Formulation

4.3.1 Remediation Objectives

OPTIMA supports flexible, multi-objective setups. Common objectives include minimizing time to cleanup, maximizing total mass removed, and minimizing plume extent (area above a threshold), applied to one or multiple contaminants. For example, three parallel goals can be considered together: (1) minimize time to reach the target concentration for CCl_4; (2) maximize cumulative CCl_4 mass extracted over the planning horizon; and (3) minimize the Tc-99 area above its regulatory limit. Pursuing these simultaneously reveals trade-offs – e.g., aggressive pumping may speed CCl_4 cleanup at the expense of Tc-99 behavior, or vice versa. Users can mix objectives across contaminants or metrics (e.g., operating cost), and OPTIMA yields a Pareto set rather than a single "best" answer.

4.3.2 Decision Variables and Design Representation

Decision variables encode where, when, and how wells are deployed/operated. Each candidate solution represents a full multi-year configuration and schedule. Key elements include (a) well location selection – which new extraction wells to install (from candidates) and which existing wells to decommission/replace; (b) installation timing – when wells are drilled/retired; and (c) pumping assignment – rates or rate categories per well, subject to plant capacity. Time is parameterized in discrete epochs (e.g., allow new wells at most once every 2 years over a 12-year horizon, yielding 6 windows). In each window, only a limited number of wells can be added, and each new well must operate for a minimum period (e.g., 5 years) before it is eligible for shutdown. This structure reflects practical planning intervals and reduces search dimensionality while aligning with real decision points. Each configuration is evaluated by the surrogate to predict outcomes (cleanup times, mass removed, plume size), which the optimizer uses to rank solutions.

4.3.3 Operational Constraints

All candidates must satisfy engineering and regulatory limits. Key constraints (user-configurable) include:

- **Treatment capacity limit**: Total extraction is capped by plant throughput; injection capacity limits may also apply.
- **Minimum well spacing**: New wells must be separated from each other and existing wells; exclusion zones can be enforced.
- **Phased installation caps**: A maximum number of new wells per time window/year and a maximum total over the horizon.
- **Per-well flow bounds**: Each well's extraction rate is bounded by design and discharge limits; total rates must respect plant capacity.
- **Well retirement time**: Minimum on-time before retirement.

Together, the objectives, decision variables, and constraints define a high-dimensional decision space that OPTIMA searches for improved P&T strategies.

4.4 NSGA-II Evolutionary Optimization Engine

4.4.1 Evolutionary Search Process

OPTIMA uses NSGA-II to navigate the complex, non-linear, and partially discrete decision space. NSGA-II (Deb et al. 2002) maintains an evolving population of candidate solutions and steers it toward the Pareto frontier. Each solution receives a non-domination rank (how many others dominate it across all objectives) and a crowding distance to preserve diversity along the front – ensuring a broad spread of high-performing solutions instead of convergence to a single point.

The algorithm begins with an initial population (random or heuristic). Each individual encodes a multi-year P&T configuration. In each generation, crossover recombines parts of two parents (e.g., portions of installation schedules or subsets of locations) to produce a child design; mutation perturbs aspects of a design (e.g., installation time, toggling a candidate location) to introduce variability. All solutions are evaluated by the surrogate, ranked via non-dominated sorting, and filtered by crowding distance within tiers. Top solutions are retained as parents for the next generation. Over dozens to hundreds of generations, the population evolves toward an approximate Pareto front, improving objectives and covering the trade-off extremes.

4.4.2 Surrogate-in-the-Loop Evaluation and Outputs

For each candidate, the U-Net surrogate rapidly simulates the multi-year outcome (e.g., time to cleanup, mass recovered, plume area). Because a 12-year forecast runs in <1 s, NSGA-II can evaluate thousands of scenarios per run. OPTIMA makes these trade-offs explicit: The Pareto front can be visualized as parallel-coordinate plots, where improving one objective worsens another. Decision-makers can interrogate the set via interactive visualizations to select a strategy aligned with priorities. Then, top designs can be re-simulated with the full physics-based model for verification.

By combining advanced machine learning with evolutionary algorithms, OPTIMA yields a family of optimal strategies and a quantitative view of the performance envelope (e.g., cleanup time vs. mass removed) at a fraction of the computational cost of physical-based numerical models – substantially expanding the scope and efficiency of P&T optimization.

4.5 Current Limitations & Assumptions

Several limitations and assumptions should guide interpretation and use, including:

- **Site-specific training and validity domain**: The surrogate is trained for the Hanford 200 West Area; accuracy is bounded by the training domain. Applying it to very different hydrogeology or far-out-of-distribution well configurations can degrade fidelity. Retraining/recalibration is required for new sites or substantially changed conditions. Optimization is deterministic; it does not explicitly propagate parametric or future uncertainty. Results should be treated as conditional on the calibrated model and used as a pre-screening guide. Recommended solutions should be validated with high-fidelity simulations and field data before implementation.
- **Operational simplifications**: The current implementation keeps the injection network static, treats pumping rates as pre-set rather than fully continuous control. Dynamic pumping strategies are not modeled; once turned on, a well typically operates continuously until retired, and wells are not cycled on/off. These choices reduce dimensionality but may exclude some beneficial strategies. As such, OPTIMA's suggestions should be interpreted in the context of these constraints – highlighting where

to add/remove wells and when, under relatively stable operation. Future enhancements can relax these simplifications to explore a broader operational space.

Despite these limits, OPTIMA provides a fast-turnaround, data-driven tool for pre-screening P&T strategies. Used appropriately – within its validated domain and with subsequent verification – it can rapidly narrow options and illuminate trade-offs in complex remediation decisions.

5.0 OPTIMA Design and User Guide

OPTIMA is under active development (alpha) and not yet released for operational use. The current build is in internal testing and debugging, and UX mockups are available for HFO feature discussions. The descriptions and screenshots below reflect the alpha interface and may evolve before production. This user guide is included now to (1) document the core workflow, objectives, constraints, and data provenance; and (2) support planning-level engagement with HFO to identify priority features and usability needs. The descriptions and screenshots below reflect the current alpha interface, which may evolve in future versions.

5.1 Interface Overview

5.1.1 Simulation Overview Page

OPTIMA is delivered as an interactive web application. Upon navigating to the OPTIMA web interface, users are presented with a Simulation Overview landing page listing all saved simulation runs. This page serves as a run registry, allowing users to manage and initiate optimization scenarios (Figure 5.1).

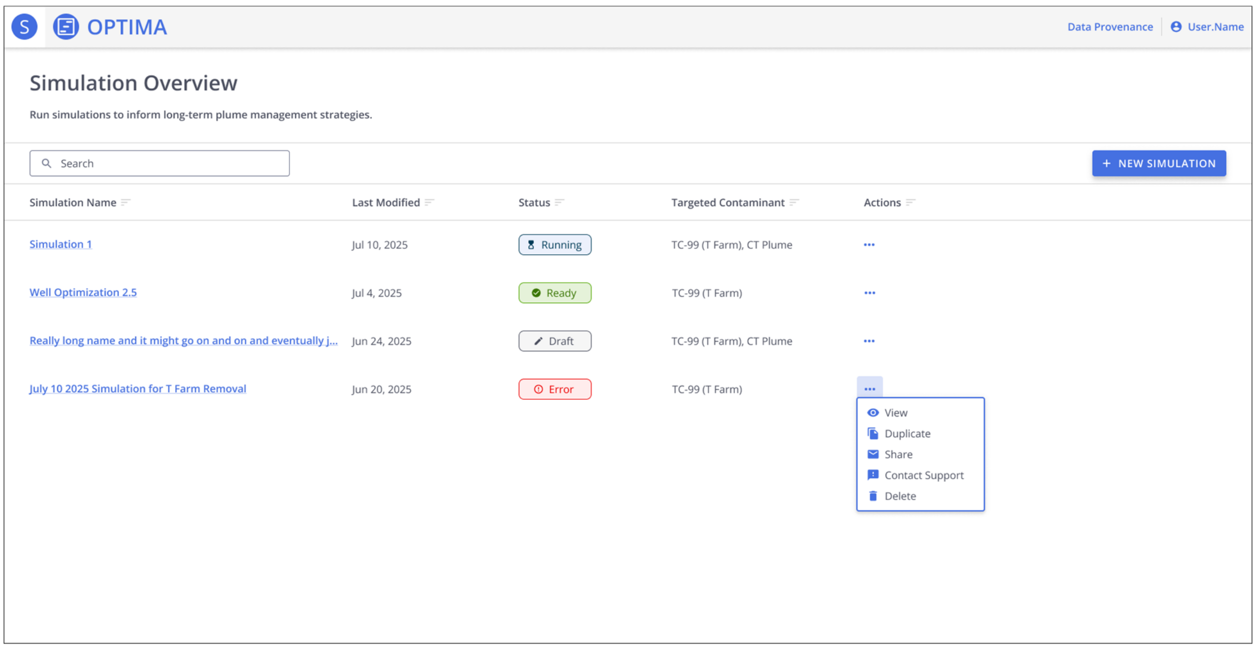

Figure 5.1. OPTIMA Simulation Overview (landing page)

The dashboard lists existing simulations with key details: simulation name, last modified date, status (e.g., Running, Ready, Draft, Error), targeted contaminants, and an Actions menu for each entry. The Actions menu provides options to View results, Duplicate the run, Share it, Contact Support, or Delete the simulation. A search bar filters simulations by name or metadata. To begin a new optimization, click New Simulation (top-right). Each simulation entry remains accessible for reproducibility and traceability.

5.1.2 Workspace Layout

After the user clicks New Simulation (or viewing an existing simulation), OPTIMA launches an interactive workspace composed of three coordinated panels, plus supporting navigation controls:

- **Solution & Objective Summary** (left): A vertical sidebar listing candidate solutions (numbered) alongside a miniature parallel-coordinates chart; selecting a line/solution drives the other panels.
- **Pumping Schedule** (center): A Gantt-style timeline showing when each well is installed, active, or retired under the selected constraints.
- **Plume Visualization** (right): Spatial plume snapshots by depth layer at the current timeline position, with an optional Pinned column for side-by-side time comparison.
- **Timeline Controls** (bottom): Scrubber for play/pause and jumps to specific years; the marker controls which plume snapshots are displayed.
- **Header** (top): Data provenance and account settings.

These components are illustrated in Figure 5.6 through Figure 5.11, beginning with the Compare Solutions workspace, where the Solution Summary, Pumping Schedule, and Plume Visualization panels are shown in operation. Sections 5.2 through 5.6) step through the typical workflow: defining an optimization scenario, narrowing results, comparing solutions in detail, and exporting the findings.

5.2 Defining an Optimization Scenario

To configure a new optimization run, proceed through the Set Parameters wizard after clicking New Simulation. The wizard collects the scenario definition – operational constraints and optimization objectives – before running the solver. The form is divided into two main parts: Well Network Operational Constraints (left) and Objectives (right, including contaminant selection). By default, the wizard loads with recommended (default) values that can use pre-computed runs for instant results. Users may modify inputs to customize the scenario. (Custom runs may require additional processing time.)

5.2.1 Operational Constraints

In this first step, define global operational constraints for P&T optimization (Figure 5.2):

- **Total treatment capacity**: Overall pumping capacity the system can handle.
- **Frequency of installation years**: Interval for adding wells (e.g., every year). A related field sets the maximum number of wells per installation year (e.g., up to 5).
- **P&T system retirement year**: Planned year to cease active pumping (end of active remediation). Optionally, specify post-retirement simulation years to track plume evolution after shut-down.
- **Individual well max flow rate**: Upper bound per well (ensures no single well exceeds a practical operational rate).
- **Minimum distance between wells**: Minimum spacing in meters (e.g., 50, 100, 150 m) to avoid clustering and respect siting limits.

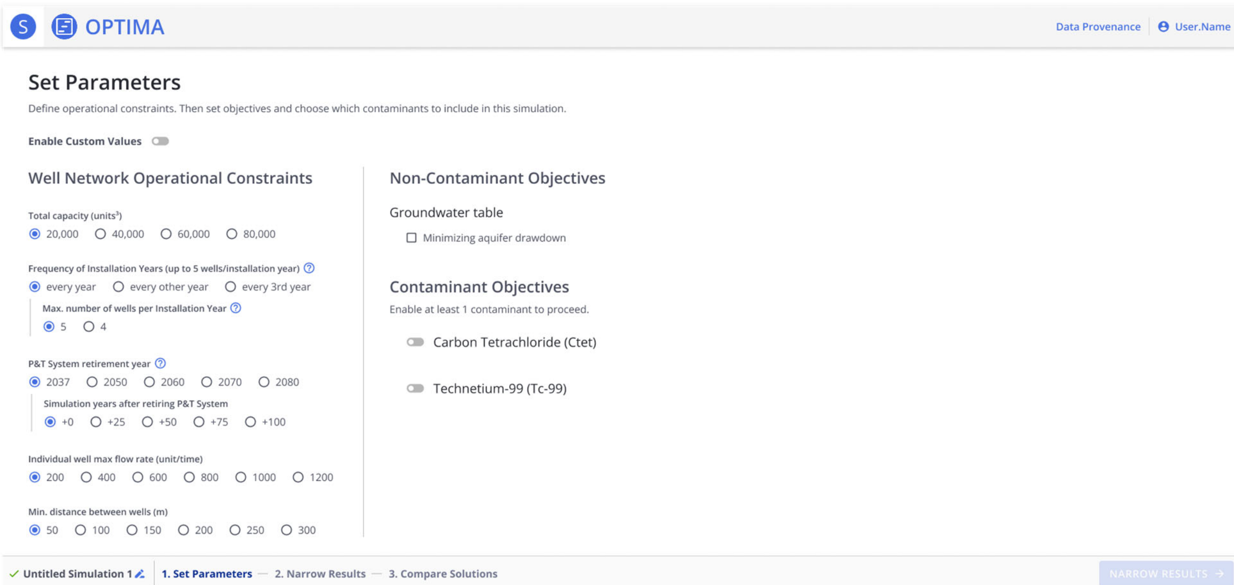

Figure 5.2. Set Parameters – Operational Constraints

By default, Enable Custom Values is off (standard defaults). In default mode, OPTIMA uses pre-computed optimization results for immediate exploration. If toggled on, all fields accept arbitrary custom inputs – offering flexibility but triggering a new optimization run (banner warns runtime may increase by up to ~1 hour). Users can adjust any parameter before proceeding.

5.2.2 Objectives Selection

When custom values are enabled (warning banner shown), all parameters can be modified freely. The Objectives panel becomes active once at least one contaminant is enabled. The user specifies which contaminants to include and what objectives to pursue for each. The current implementation supports two contaminants: CCl_4 and Tc-99. For each enabled contaminant, toggle one or more objective functions and set acceptable concentration thresholds that define plume area (extent above the threshold). Available objectives include:

- Minimize plume area (above threshold)
- Maximize mass removed (over the simulation period)
- Minimize remedy time (time to meet remediation concentration targets)
- Minimize aquifer drawdown (non-contaminant objective)

At least one contaminant must be enabled to proceed. Thresholds are user-set (with suggested normal ranges shown for guidance). After defining constraints and objectives, click Narrow Results →. With defaults, OPTIMA retrieves pre-computed optimal solutions immediately. With custom inputs, a new run is launched using NSGA-II with the surrogate; the Simulation Overview shows status Running, changing to Ready upon completion. In either case, the next step is to explore and narrow the candidate results.

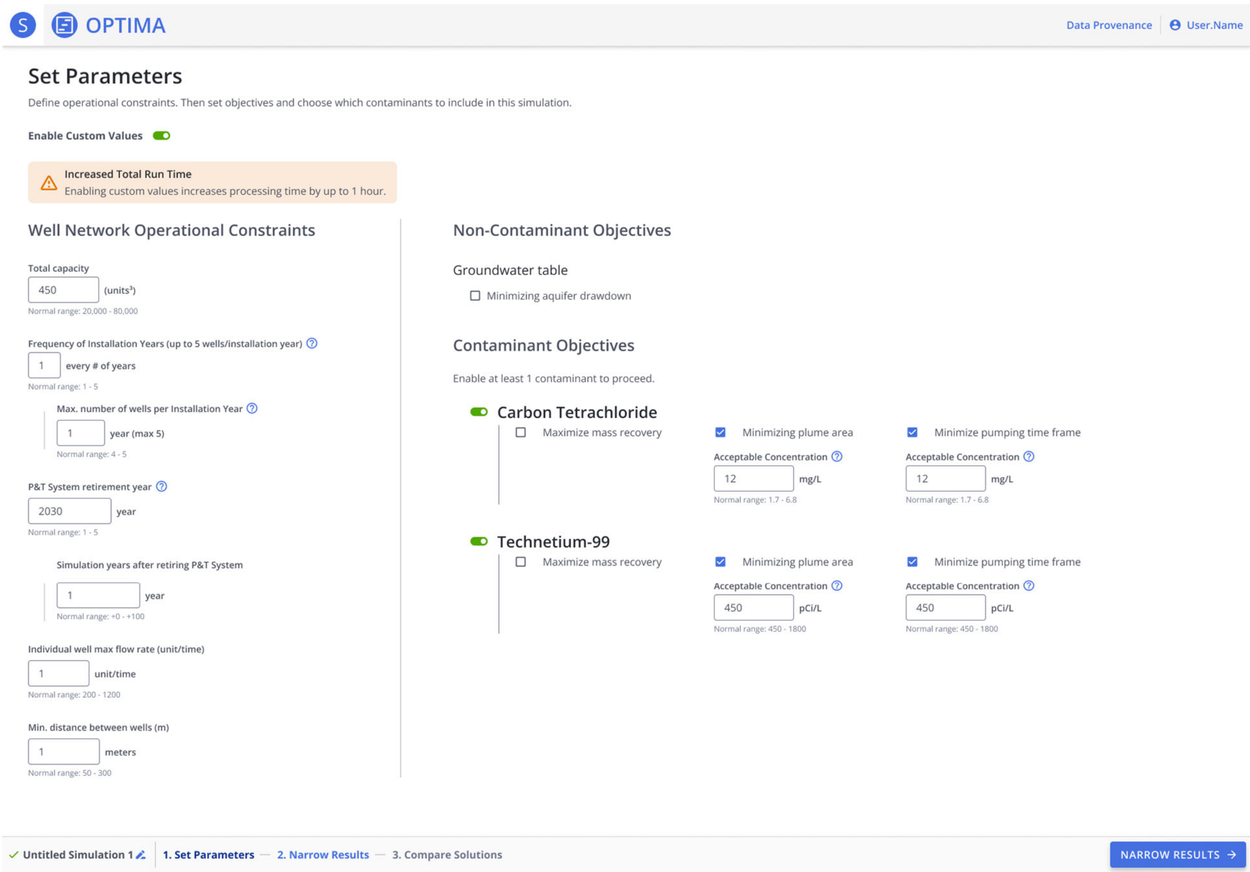

Figure 5.3. Set Parameters – Custom Inputs and Objectives

5.3 Narrowing Results

After a run is ready, the Refine Results view presents all feasible final candidate solutions from the evolutionary search. Users interactively filter to a manageable subset before detailed comparison. The default interface is a parallel-coordinates plot of the entire candidate set.

5.3.1 Full Candidate Set

Each polyline represents one wellfield design across four objective axes: Mass Removed (kg), Tc-99 Plume Area (km^2), and Remedy Time (years). (Example: "Showing 117/117 possible solutions.") Use the Min/Max boxes beneath each axis to filter ranges (Figure 5.4). Lines outside any active range fade; those meeting all filters remain bright. This reveals trade-offs (e.g., lower cost vs. larger plume area; minimal plume vs. longer time/higher cost).

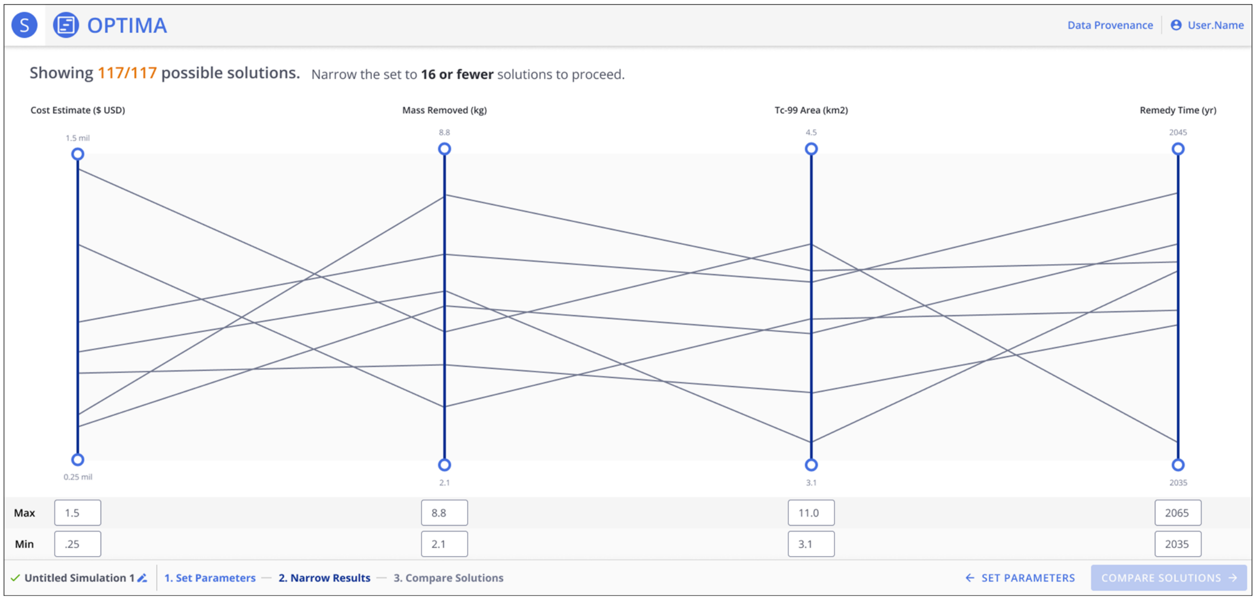

Figure 5.4. Refine Results – Full Candidate Set

5.3.2 Filtered Subset and Progression

Filtering narrows candidates (e.g., “Showing 3/7 possible solutions” where 3 meet the combined criteria). Remaining solutions are highlighted (others fade for context) (Figure 5.5). Reduce to ≤16 solutions before proceeding – then click Compare Solutions →. If more than 16 remain, the interface prompts for further narrowing.

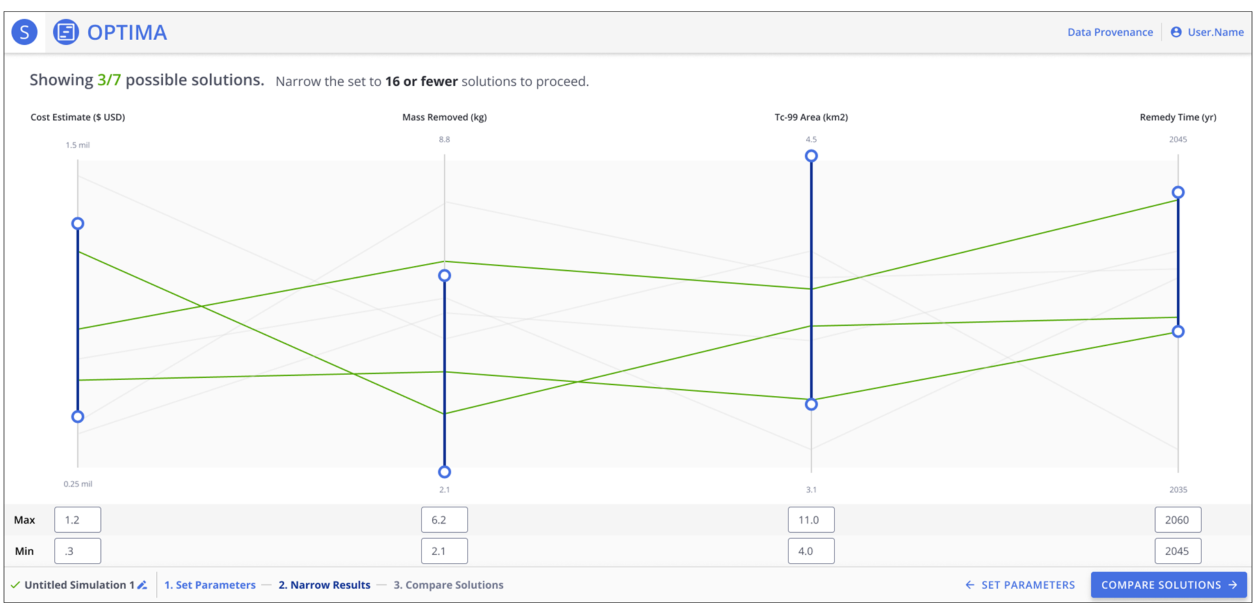

Figure 5.5. Refine Results – Filtered Subset Highlighted

5.4 Compare Solutions

The Compare Solutions workspace supports in-depth operational and spatial review for up to 16 shortlisted designs. The layout uses the three coordinated panels (Solution list, Pumping schedule, Plume visualization) with tools to examine timeline feasibility and plume evolution.

5.4.1 Initial State

The left panel lists candidate solutions with a mini parallel-coordinates plot. The center and right panels display a prompt – "Select a solution from the left to begin exploring results." Click any solution to populate the pumping schedule and plume visualizations.

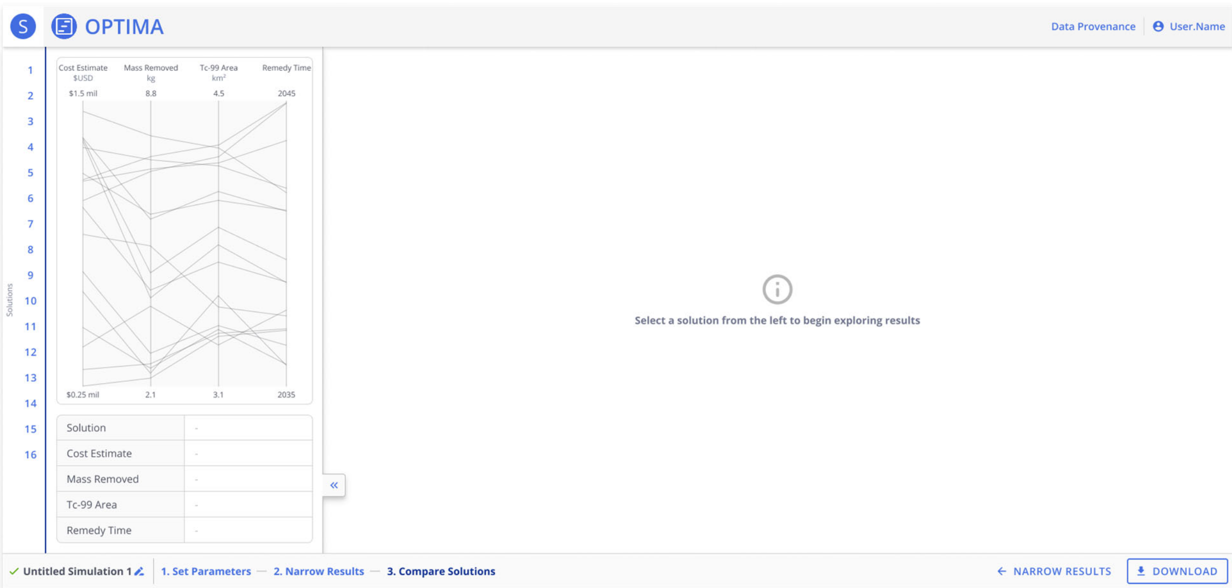

Figure 5.6. Compare Solutions – Initial (no selection)

5.4.2 Single Solution Selected

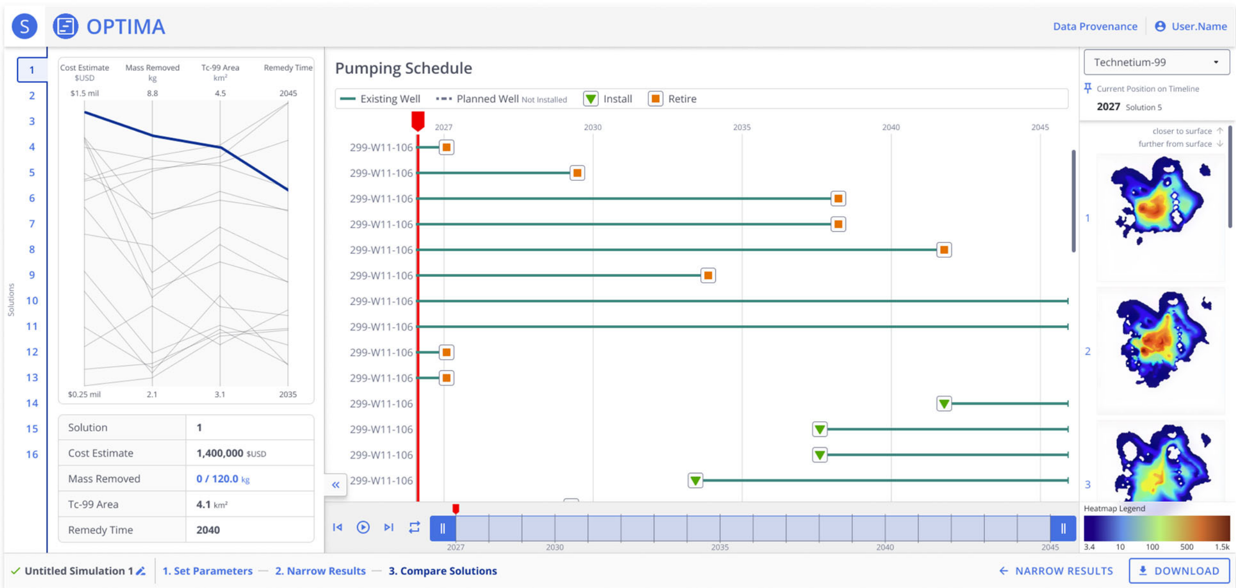

Figure 5.7. Compare Solutions – One Solution Selected

The left panel highlights the selected solution's polyline and displays a summary card with objective values (e.g., Mass Removed, Tc-99 Area, Remedy Time). The center panel presents the pumping schedule as per-well Gantt bars over time (e.g., 2025–2045), using consistent visual conventions:

- Solid green = active period of an existing well.
- Dashed green = scheduled period for a planned well (future).
- Green triangle = install event (year a new well comes online).
- Orange square = retire event (year a well is shut off).

The bottom scrubber controls the current year, and a red vertical marker synchronizes the schedule and plume views (e.g., showing an installation in 2027).

The right panel (Plume Visualization) shows small-multiple maps by depth for the current time and the selected contaminant (chosen via dropdown). Warm colors indicate higher concentrations, and the legend is aligned to reporting thresholds. These maps reveal where concentrations exceed the threshold and how the selected design influences plume containment and reduction.

5.4.3 Current vs. Pinned Time

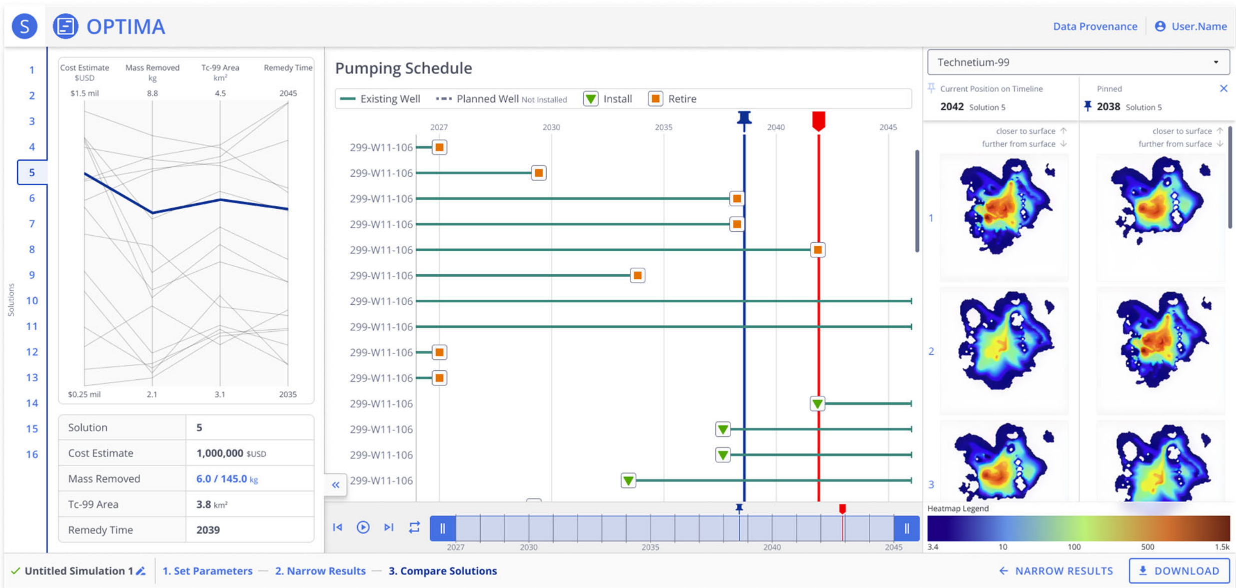

Figure 5.8. Compare Solutions – Current vs. Pinned

Pin a second time snapshot to compare two moments for the same solution (e.g., 2038 pinned vs. 2042 current in Solution 5). The schedule shows two vertical markers, and the plume panel renders two columns – Pinned and Current – with a consistent legend, enabling clear before/after assessment (e.g., Tc-99 area shrinkage, persistence of hot spots). The user can clear or adjust the pinned time at any point.

5.4.4 Layout Controls

The left rail can be collapsed (chevron at panel edge) to expand the schedule and plume views within a full-width window (Figure 5.8). The panels can also be resized via the splitter between schedule and plume views. Enlarging the plume maps supports spatial diagnostics (edge containment, hot spots, morphology) while the schedule remains visible for temporal reference (Figure 5.9).

When the browser window is narrowed (e.g., laptop, tiled windows), the interface reflows automatically: The summary and mini-chart stack, the schedule remains scrollable, plume maps stay legible, and the timeline scrubber continues to function (Figure 5.10). This provides additional working space via responsive re-layout, complementing the manual collapse/resize controls.

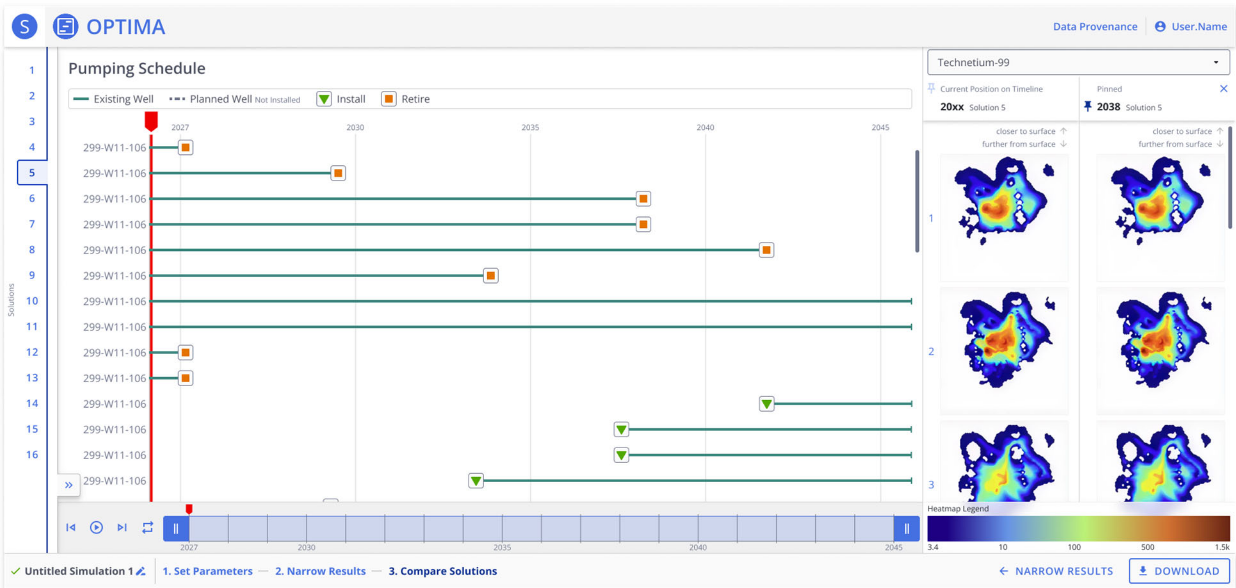

Figure 5.9. Compare Solutions – Collapsed Left Panel

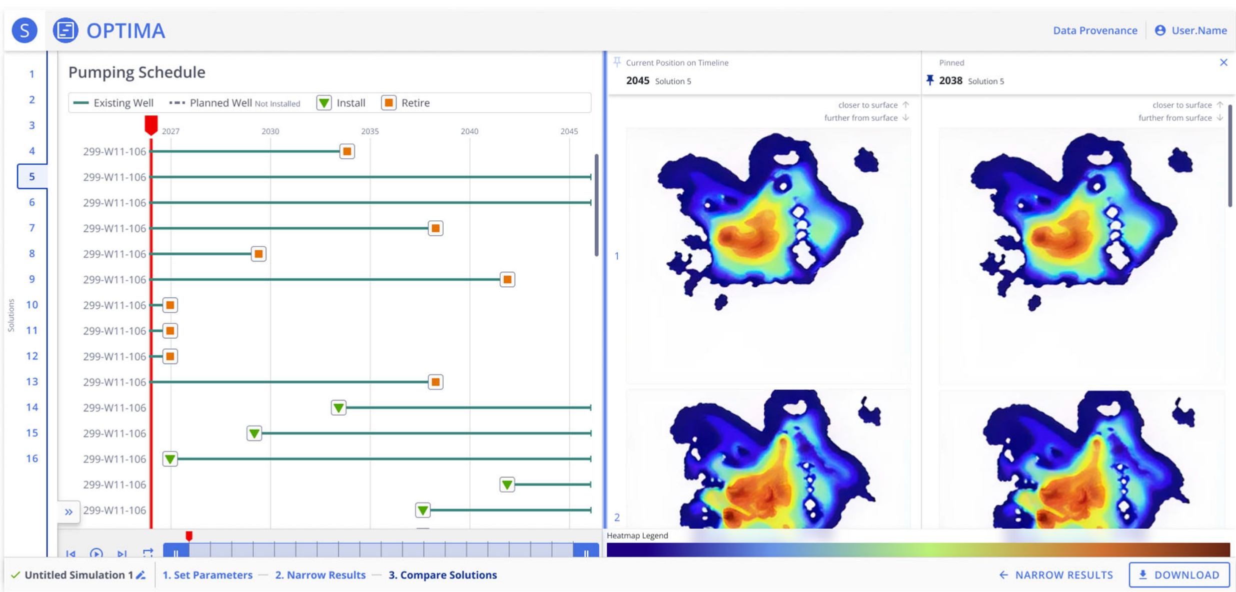

Figure 5.10. Compare Solutions – Drag to Resize Panels

5.5 Exporting Results and Reports

OPTIMA includes tools to export analysis outputs for reporting, quality assurance, and downstream use. Click the Download button in the upper-right/header to open the Download Assets dialog (Figure 5.11). The dialog provides a search field and a Select All option. Each shortlisted solution appears as an expandable section listing available assets:

- **Parallel Coordinates Chart (.png)**: Image of the subset/solution highlight.
- **Well Schedule (PDF)**: Gantt-style pumping schedule report for the solution.

- **Plume Map (JPEG)**: Plume snapshot(s) for a selected time or a composite.
- **Data Table (.xlsx/.csv)**: Objective values, installation/retirement schedule, volumes pumped, and related metrics.

Each item displays an estimated file size. Users may export one solution or multiple solutions simultaneously; OPTIMA packages the selections as a compressed ZIP for download. When exporting images, legends and on-screen annotations are included; when exporting data tables, fields capture the values displayed in the interface (objectives, schedules, and flow summaries) for reproducibility.

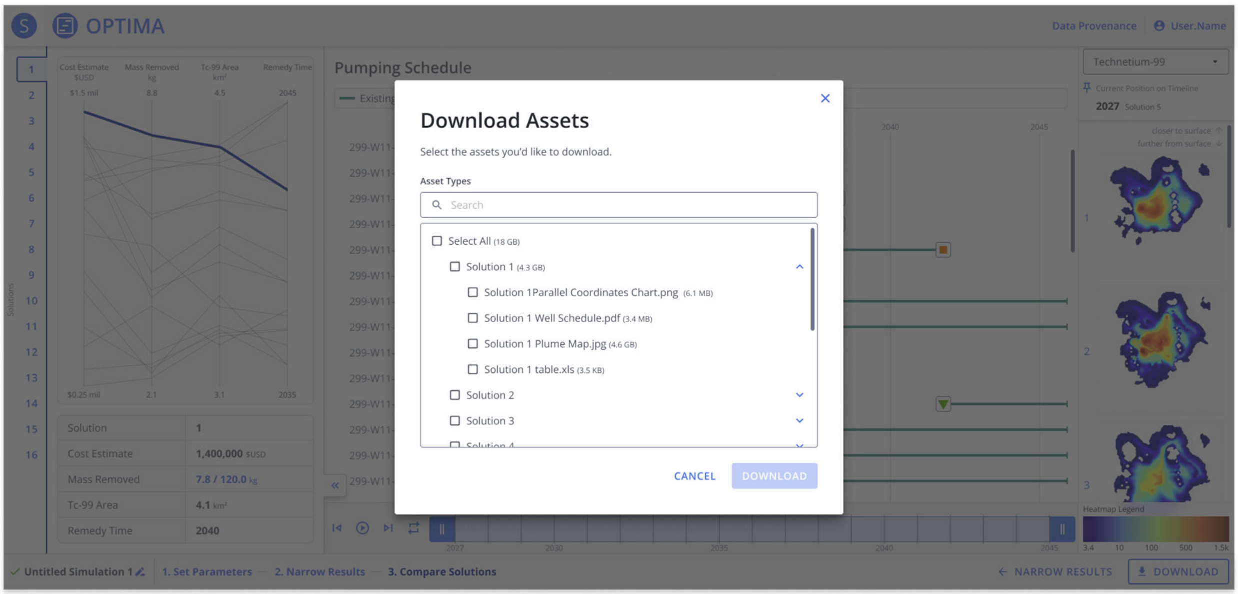

Figure 5.11. Download Assets Dialog

5.6 Summary

OPTIMA provides a reproducible, interactive workflow that moves from defining constraints and objectives → generating a broad set of candidate solutions → narrowing results → in-depth operational and spatial comparison → exporting findings. By coupling a physics-informed surrogate model with a multi-objective optimization algorithm, OPTIMA makes large, complex design spaces navigable. It directly links high-level trade-offs – cleanup time, contaminant mass removed, plume extent – to implementable wellfield schedules and plume outcomes.

Although still in an alpha stage, OPTIMA demonstrates how advanced computation can be delivered in an accessible decision-support format. It enables rapid "what-if" exploration (e.g., adding wells, adjusting pumping regimes) and immediate visualization of cleanup impacts. Early use indicates that OPTIMA can significantly reduce the time and effort required to evaluate P&T design alternatives. As development continues, the tool is expected to complement traditional modeling and engineering analysis by accelerating scenario screening and helping practitioners select remedies that balance cost, time, and performance.

Note: OPTIMA is intended to guide scenario exploration, not to serve as the final design tool. All results should be confirmed with detailed numerical modeling and professional engineering judgment.

6.0 Conclusions

This report has documented the development and application of two complementary decision-support tools – PTOLEMY and OPTIMA – designed to improve the effectiveness of P&T remediation at the Hanford 200 West Area.

- PTOLEMY provides a probabilistic framework for evaluating candidate extraction wells. It combines outputs from the site's flow-and-transport model with deep learning to estimate the likelihood that candidate wells will achieve higher mass recovery. The tool supports efficient prescreening of well placement options and is designed for reproducibility and transparency.
- OPTIMA extends this capability through multi-objective optimization of full wellfield design and operation. By embedding a physics-informed surrogate model within an NSGA-II evolutionary algorithm, OPTIMA identifies candidate strategies that balance competing objectives such as plume containment, mass removal, cleanup time, and well drawdown.
- Together, the tools enable a linked workflow: PTOLEMY identifies where wells are likely to be effective, and OPTIMA evaluates how entire wellfields perform under operational constraints. Both tools substantially reduce computational effort compared with direct simulation while maintaining fidelity to site conditions.

Overall, PTOLEMY and OPTIMA provide a science-based, user-oriented approach to guiding P&T optimization. They enhance transparency in decision-making by linking quantitative performance metrics directly to implementable well configurations. Note that OPTIMA remains under active internal development and has not yet been formally released; this guide was prepared to support stakeholder engagement and provide technical transparency ahead of the tool's final release.

7.0 Continued Development

PTOLEMY is operational now, and OPTIMA remains under active internal development and validation prior to formal release. Continued development is planned to expand capability, increase robustness, and improve user experience:

- **Current status and validation**: PTOLEMY will receive periodic model updates/retraining as new monitoring data become available. OPTIMA is under active internal development and validation (pre-release), with progressive testing against the site decision model.
- **Broader contaminants and site coverage**: Current surrogate models focus on CCl_4 and Tc-99 in the 200 West Area. Future training will extend OPTIMA to additional contaminants, and hydrogeologic settings, and expand PTOLEMY to new plume conditions.
- **Enhanced optimization features**: Planned additions include support for expanded operational constraints (e.g., variable pump rates, screen lengths, injection scenarios), including uncertainty quantification, and improved visualization of solution trade-offs.
- **User interface refinements**: Ongoing work will make the OPTIMA interface more flexible, including more side-by-side solution comparisons, advanced filtering, and enhanced reporting.
- **Integration with adaptive remedy management**: A longer-term goal is to embed both tools within an iterative management framework, where new data, updated forecasts, and changing objectives can trigger periodic re-evaluation of remedy strategies.
- **Performance improvements**: Continued refinement of surrogate models and optimization algorithms will further reduce run times, enabling exploration of larger and more complex scenario spaces.
- **Integration within SOCRATES**: Future enhancements will focus not only on the individual tools, but also on expanding their integration within SOCRATES to leverage additional modules (e.g., plume analytics, monitoring data visualization) and to support broader applications for the DOE Office of Environmental Management.

8.0 Quality Assurance

This work was performed in accordance with the Pacific Northwest National Laboratory (PNNL) Nuclear Quality Assurance Program (NQAP). The NQAP complies with the United States Department of Energy Order 414.1D, Quality Assurance. The NQAP uses NQA-1-2012, Quality Assurance Requirements for Nuclear Facility Application as its consensus standard and NQA-1-2012 Subpart 4.2.1 as the basis for its graded approach to quality.

9.0 References

Budge, T. 2020. *Model Package Report: Plateau to River Groundwater Model*, Version 8.3. CP-57037-Rev. 2. CH2M Hill Plateau Remediation Company, Richland, WA. https://doi.org/10.2172/1601635

Deb, K., A. Pratap, S. Agarwal, and T. Meyarivan. 2002. "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II." *IEEE Transactions on Evolutionary Computation* 6(2): 182–197. https://doi.org/10.1109/4235.996017

DOE/RL-2008-78. 2009. *200 West Area 200-ZP-1 Pump-and-Treat Remedial Design/Remedial Action Work Plan*. Rev. 0, REISSUE. U.S. Department of Energy, Richland Operations Office, Richland, WA. https://pdw.hanford.gov/document/0096137

DOE/RL-2019-38. 2019. *200-ZP-1 Operable Unit Optimization Study Plan*, Rev. 0. U.S. Department of Energy, Richland Operations Office, Richland, WA. https://pdw.hanford.gov/document/AR-03236

DOE/RL-2020-60. 2019. *Hanford Site Groundwater Monitoring Report for 2020*. Rev. 0. U.S. Department of Energy, Richland Operations Office, Richland, WA. https://www.hanford.gov/files.cfm/DOE_RL-2020-60_R0_Clean.pdf

EPA, DOE, and Ecology. 2008. *Record of Decision Hanford 200 Area 200-ZP-1 Operable Unit Superfund Site, WA*. EPA/541/R-08/115, U.S. Environmental Protection Agency, U.S. Department of Energy, and Washington State Department of Ecology, Seattle, WA. http://pdw.hanford.gov/arpir/index.cfm/viewDoc?accession=00098825

Hammond, T. B. and D. Lupton. 2015. *Development of the Hanford South Geologic Framework Model, Hanford Site, Washington*. ECF-HANFORD-13-0029, Rev. 2. CH2M Hill Plateau Remediation Company, Richland, WA.

PNNL-22062. 2012. *Abiotic Degradation Rates for Carbon Tetrachloride and Chloroform: Final Report*. RPT-DVZ-AFRI-012, Pacific Northwest National Laboratory, Richland, WA. https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-22062.pdf

Song, X., H. Ren, Z. Hou, et al. 2023. "Predicting Future Well Performance for Environmental Remediation Design Using Deep Learning." *Journal of Hydrology* 617(February): 129110. https://doi.org/10.1016/j.jhydrol.2023.129110

Song, X., I. Demirkanli, Z. Hou, et al. 2025. "Integrating Analytical Solutions and U-Net Model for Predicting Groundwater Contaminant Plumes in Pump-and-Treat Systems." *Advances in Water Resources* 202(August): 105002. https://doi.org/10.1016/j.advwatres.2025.105002

Thiem, G. 1906. *Hydrologische Methoden*. Leipzig: J. M. Gebhardt's Verlag. https://cir.nii.ac.jp/crid/1130000797191951616

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354
1-888-375-PNNL (7665)

www.pnnl.gov