Multi-Agentic Approach for History Matching of Oil Reservoirs

Linar Samigullin[1], **Sergei Shumilin**[1,*], **and Evgeny Burnaev**[1,2]

[1]Skoltech, AI Center, Moscow, Russia
[2]AIRI, Moscow, Russia
[*]s.shumilin@skoltech.ru

ABSTRACT

History matching is a central inverse problem in reservoir engineering, where uncertain reservoir parameters must be calibrated against observations. Although automated history matching can reduce manual effort, practical deployment remains difficult because engineers must still configure heterogeneous workflows involving parameter selection, physically admissible bounds, optimizer choice, hyperparameter tuning, simulator execution, and diagnostic reporting. We propose **PetroGraph**, a multi-agent framework for intelligent reservoir history matching that decomposes this workflow into specialized agents for model review, experimental planning, parameterization, optimization, simulation, and summarization. The system combines large language model agents with domain-specific tools, retrieval-augmented access to simulator documentation, validation of modified ECLIPSE input decks, human-in-the-loop checkpoints, and an OPM Flow-based simulation backend. This design enables users to initiate and steer history matching through natural language while preserving explicit control over selected parameters and optimization settings. We evaluate **PetroGraph** on three reservoir models of increasing complexity: the synthetic SPE1 model, the faulted SPE9 benchmark, and the real-field Norne model. Using weighted normalized root mean square error as the objective, **PetroGraph** reduces the mismatch by 95% on SPE1, 69% on SPE9, and 13% on Norne. These results demonstrate that multi-agent orchestration can automate key decisions in history matching, lower the expertise barrier for operating complex simulation workflows, and provide a flexible foundation for extensible, domain-aware reservoir model adaptation.

Introduction

An important applied domain in which methodological heterogeneity poses a significant challenge is oil reservoir model management. Oil companies typically employ large teams of specialists responsible for updating numerical reservoir models as new field data become available. This process is labor-intensive and strongly dependent on the expertise and experience of reservoir and petroleum engineers.

The task of calibrating a reservoir model against observed field data, such as oil production rate, water production rate, bottom-hole pressure, reservoir pressure, well-log measurements, and other observations, is referred to as reservoir model history matching. History matching consists of iteratively modifying model parameters and evaluating the discrepancy between simulation outputs and observed field data. This process can be automated. Automatic history matching belongs to a class of inverse problems in which the objective is to identify unknown parameters of a mathematical model that allow it to reproduce observed data as accurately as possible. Such problems are typically formulated as optimization problems[1] and involve heterogeneous components, including optimization algorithms, parameterization techniques, and numerical solvers for reservoir flow simulation.

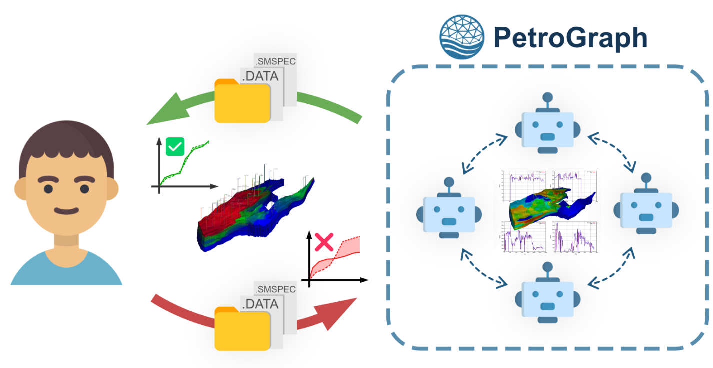

Figure 1. Conceptual overview of **PetroGraph**. A user interacts with a reservoir model through natural language; **PetroGraph** translates the request into a coordinated multi-agent workflow that reviews input data, selects parameters and bounds, runs optimization and simulation, and returns validated model updates and diagnostic feedback.

The central challenge addressed in this study lies in the diversity of available methods and their software implementations.

Multiple optimization techniques are applicable to reservoir inverse problems, and each often requires a number of hyperparameters to be selected or tuned for a specific scenario. This creates a problem of methodological multiplicity, which is the core focus of the present research. Modern artificial intelligence tools, particularly multi-agent systems, provide a way to address this challenge through an LLM-based orchestrator capable of selecting appropriate methods and configuring their parameters in an informed manner.

The application of a multi-agent approach to automatic reservoir history matching is particularly natural because the main components of the history-matching workflow can be decomposed into specialized tasks. These tasks can be represented as independent agents, each with a clearly defined objective. For example, an optimization agent can select a suitable optimization method for a given task and configure its hyperparameters. Moreover, the system can interact with the user through natural language, thereby reducing the expertise required to operate complex history-matching workflows. A functional multi-agent system for oil reservoir history matching can also be extended to related application scenarios, such as gas reservoir history matching or calibration of geological models for mineral deposits, making the approach suitable for multi-domain engineering applications.

In practice, engineers face the difficult problem of selecting an optimal computational pipeline for a given history-matching task. Even experienced engineers may find it challenging to account for the full range of modern methods, software tools, and configuration choices. We aim to address this problem by developing a system that simplifies automatic history matching through a multi-agent methodology. We present **PetroGraph**, a framework for intelligent automatic history matching. **PetroGraph** provides a user-friendly interface that enables engineers without a computer science background to configure, execute, and analyze history-matching workflows more easily.

Related Work

History matching (HM) of oil reservoirs. History matching is an inverse problem that aims to identify a set of reservoir parameters that best fits real-world field data. These parameters are usually divided into two categories: static parameters, such as permeability and porosity, and dynamic or flow-related parameters, such as relative permeability. In the optimization-based formulation of history matching, the reservoir simulator is commonly treated as a black-box function whose outputs are compared with observed data. Usually, an engineer selects the parameters to be optimized based on the analysis of the model. HM has been extensively studied since the 1990s, and today the primary research focus is concentrated on automated history matching[2]. In practice, industrial oil reservoir models may contain tens of millions of cells, which makes hydrodynamic simulation computationally expensive. This motivates reducing the number of simulator runs required during optimization. One of the most advanced approaches to sample-efficient optimization is Bayesian optimization (BO), which builds a surrogate model of the black-box objective function using a Gaussian process model and efficiently queries this model to select the next parameter set for evaluation[1,3].

Multi-agent systems (MASs). MASs are currently demonstrating rapid growth in both the number of publications and the range of practical applications. The combination of a system prompt, an LLM, and a set of tools has become a de facto practical implementation of an agent. The main motivation comes from the idea that a set of diverse cooperating agents can be more efficient than a single actor. The key properties of MASs—the autonomy of agents, their flexibility, and their ability to adapt to environmental changes—allow them to solve complex distributed problems by decomposing them into subtasks executed by individual agents[4].

In this study, the definition of an agent is adopted from the classical work[5]. An LLM-based agent can dynamically form its actions using generalized knowledge acquired through pre-training on large datasets. In 2020–2024, the first works demonstrating a new approach appeared: instead of following a rigidly defined behavioral program, the agent is supplied with a prompt describing its goal and context, while control over its actions is delegated to a language model via a set of tools[6].

MASs in engineering. Current research in this area demonstrates the effectiveness of MASs in solving engineering problems through the cooperation of specialized agents. As the authors of[7] note, this architecture provides a more accurate representation of the real world compared with single-agent systems, since many applied problems naturally involve the interaction of multiple participants making decisions simultaneously. In[8], the authors present a MAS for optimizing automobile design. A set of agents iteratively modifies the shape of the car body and then calculates the aerodynamic characteristics of the resulting shape. This system is a clear example of a hybrid multi-agent system, as it combines agents based on deep learning with agents that perform numerical simulations.

In the context of BIM design automation, it has been shown that LLM agents effectively interpret unstructured design requirements and transform them into API calls for authoring systems, while rule-based validators and geometric modules provide formal verification and correction of models[9]. In automotive engineering, agent-based LLM systems demonstrate the ability to accelerate HARA analysis, requirements generation, and test planning. This is achieved by combining LLM-based semantic inference with rigorous validation of results obtained by traditional methods. This hybrid approach reduces the labor intensity of documentation preparation and improves the reproducibility of safety processes[10].

In manufacturing systems, LLM-enabled MASs are used for process coordination: the LLM agent distributes tasks, generates control instructions, and adjusts the process flow when requirements change. Specialized agents perform calculation procedures and optimization steps, which increase system flexibility and reduce equipment changeover time[11].

In the energy sector, frameworks are being developed in which LLM agents perform semantic analysis of events and generate corrective scenarios, while numerical modules simulate and verify these scenarios before execution. This approach shows potential for faster detection of network disturbances[12].

For human–robot collaboration (HRC) tasks, LLM-oriented approaches enable operators to specify requirements in natural language, and the agent system transforms them into a sequence of controlled steps and robot programs, while digital-twin and verification agents ensure safe execution and compliance with technological constraints[13].

The work[14] addresses the problem of coordinated decision-making at various project stages in the construction industry. The proposed MAS, in which heterogeneous agents representing different disciplines, such as design and finance, engage in a structured dialogue to identify and resolve conflicts, involves exchanging evidence and developing a collective decision. This system demonstrated its effectiveness in four scenarios, significantly improving cost-forecasting accuracy and reducing the number of design flaws. Other applications of MASs for solving design problems in the construction industry are discussed in[15].

The work[16] addresses the problem of integrated support for engineering work, taking into account technical, ethical, and organizational requirements. The authors propose a MAS architecture in which each LLM-based agent plays the role of a subject-matter expert, while a coordinating mechanism ensures collaborative decision-making. The approach has been tested on a number of real-world project proposals, where it has been shown to provide more comprehensive and balanced recommendations than a single agent, improving the quality and completeness of engineering solutions.

The study[17] addresses the low efficiency of conventional LLMs when solving engineering problems that require both domain knowledge and multi-step computations. The authors propose a hybrid MAS architecture in which one agent uses RAG to improve the accuracy of theoretical answers, while another performs mathematical modeling under the control of a central coordinator. This approach provides more robust and interpretable solutions to engineering and biochemical problems, significantly outperforming basic LLMs in terms of stability and completeness of results.

The topic of context-aware MASs is addressed in the study[18]. This paper presents a detailed overview of context-aware MASs, describes key approaches to their design, and presents a wide range of engineering applications of such systems, including smart cities, supply chains, and other applications where agents must adapt to a dynamic environment. Particular attention is also paid to reliability and security in such heterogeneous environments.

MASs interacting with physical solvers. Several recent studies have explored the development of MASs designed to interface with physical solvers. For instance, the work presented in[19] introduces a multi-agent framework for computational fluid dynamics (CFD) that generates an end-to-end OpenFOAM workflow from a single natural-language prompt. Evaluated on a benchmark dataset covering a wide variety of CFD simulation types, the system achieves a success rate of 88.2%. The authors attribute this significant improvement over existing frameworks—such as MetaOpenFOAM[20], which achieves only a 55.5% success rate on the same benchmark—to the implementation of a Hierarchical Multi-Index RAG system. This system is constructed from OpenFOAM tutorial cases, allowing for more precise and contextually relevant information retrieval.

Frameworks such as AutoFEA[21] and MooseAgent[22] have demonstrated the ability to translate natural-language descriptions of structural engineering problems into executable input files for finite element analysis (FEA) solvers, such as CalculiX and Abaqus. A key innovation in AutoFEA is its use of a GCN–Transformer link-prediction model for retrieval, which accounts for the sequential dependencies between steps in a simulation workflow.

Methodology

Multi Agentic Approach

PetroGraph is a multi-agent system designed to accelerate and lower the entry barrier for the critical industry process of history matching in oil and gas field development.

The process begins with loading the reservoir data in the ECLIPSE 100 format, which consists of a primary .DATA file and its associated files, along with the corresponding simulation results stored in .SMSPEC files. These files contain both the initial simulation outputs and the actual field data against which the history matching will be performed.

The system employs a hybrid interface that combines primary interaction through a chat with manual input fields for parameter editing and selection (see Fig. 2). This approach preserves the familiar chat-based interaction model common to agentic systems while alleviating the need for the user to repeatedly ask the system to change specific parameters. Instead, it enables users to make adjustments directly.

The multi-agent framework consists of several agents linked in an execution chain, which can be represented as a graph (see Fig. 2). These agents are the Reviewer Agent, Planner Agent, Parameterizer Agent, Optimizer Agent, Simulator Agent, and Summarizer Agent. All agents share information through a common state. Detailed descriptions of each agent are provided in the corresponding subsection. The framework is based on the open-source LangGraph framework. A detailed graph architecture is provided in Appendix A.

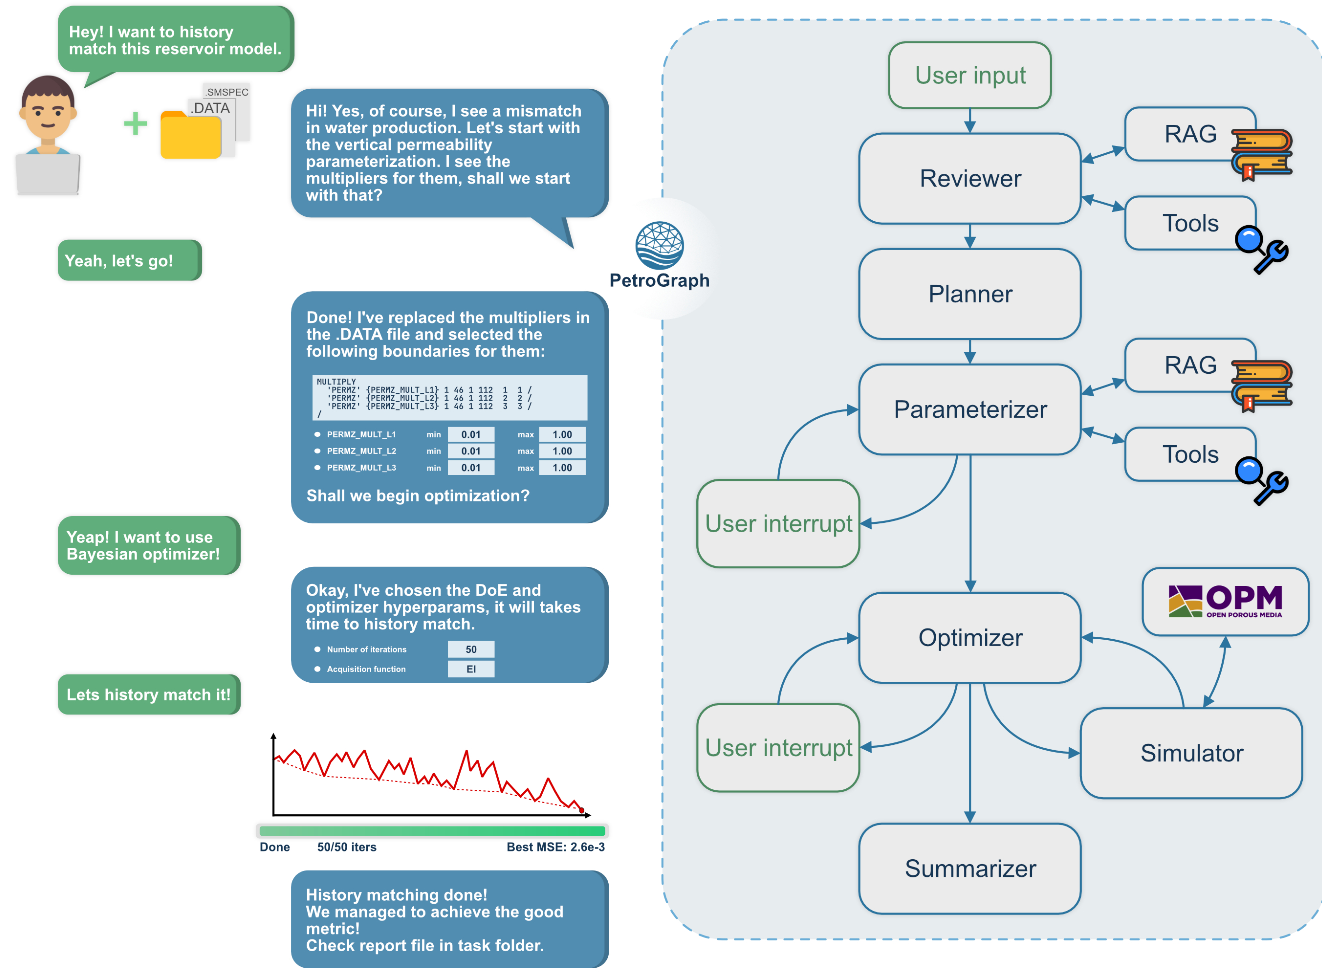

Figure 2. **PetroGraph** system architecture illustrating end-to-end example workflow. The system features six agents with human-in-the-loop: 1. **Reviewer Agent**: examines the reservoir model and creates a textual representation, 2. **Planner Agent**: analyzes data mismatch and proposes a DoE for subsequent agents, 3. **Parameterizer Agent**: finds and sets parameters for history matching and changes keywords content, 2. **Optimizer Agent**: chooses DoE and hyperparameters for optimizer, 3. **Simulator Agent**: evaluates simulations and calculates metric, 4. **Summarizer Agent**: makes summary of reservoir history matching results.

Agents

The **Reviewer Agent** is an LLM-based agent. It is the first to respond: it investigates the provided reservoir model using its tools and then creates a concise textual summary of the data for subsequent agents. This summary includes the model type (blackoil, dissolved gas, etc.), model dimensions, and the presence of multipliers, faults, and other features relevant to history matching.

The **Planner Agent** is an LLM-based agent. It receives global field metrics – provided separately for production, injection, fluid type, and bottom-hole pressure – together with the concise description from the Reviewer Agent. Its task is the Design of Experiment (DoE): it recommends which parameters to select for history matching and defines the optimisation strategy while explicitly accounting for computational cost. For lightweight models the number of iterations and parameters can be generous; for heavier models a conservative approach is warranted in selecting both the number of parameters and iterations.

The **Parameterizer Agent** is an LLM-based agent. It receives the brief reservoir description and the plan from the Planner

Agent. It begins by conducting an in-depth investigation of the input-deck structure, then builds a representation of the model by identifying parameters that can be varied and defining their physically and context based ranges. Template files containing parameter placeholders are created for future value insertion. Throughout this process the agent validates the input deck by performing dry runs with the minimum and maximum substituted values; any error (e.g., a non-monotonic relative permeability curve or an arithmetic expression inside a keyword that does not support it) is relayed back, preventing mistakes that the probabilistic model might otherwise make. The agent also has access to a RAG system based on the description of keywords and sections with examples of the ECLIPSE 100 .DATA file structure from the OPM simulator manual[23]. Once an error-free parameterization is obtained, a Human-in-the-Loop (HITL) checkpoint allows the user to adjust parameters and their bounds on the fly.

The **Optimizer Agent** is also an LLM-based agent. It receives the problem dimensionality and the DoE from the Planner Agent, then selects the DoE parameters – including the number of iterations and the initial sampling strategy – and chooses the optimizer type together with its hyperparameters. After initialisation, another HITL checkpoint occurs, at which the user may modify specific parameter values.

The **Simulator Agent** is the only agent that does not utilise an LLM. During the history-matching loop it operates in an ask-parameters/tell-metric format: it prepares the data for model execution, substitutes the parameter values provided by the optimizer, runs the OPM simulator[24], and calculates the resulting metrics.

To evaluate the quality of the history match, we employ the weighted mean normalized root mean square error (wNRMSE). For each well and for production and injection metrics, the weight is defined as the fractional contribution of the corresponding quantity to the total historical production or injection. For the bottom-hole pressure metric, the weight is defined as the reciprocal of the number of wells. The overall objective function is computed as the sum of contributions across all variables with nonzero historical data, including WBHP, WOPR, WWPR, WGPR, WOIR, WWIR and WGIR, with wNRMSE formulated as shown in Eq. 1.

$$\text{wNRMSE} = \text{weight} \cdot \text{NRMSE} = \text{weight} \cdot \frac{\text{RMSE}}{\bar{y}_{\text{hist}}} = \text{weight} \cdot \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(y_{\text{sim}} - y_{\text{hist}}\right)^2}}{\bar{y}_{\text{hist}}} \tag{1}$$

The **Summarizer Agent** is an LLM-based agent. Its primary task is to summarize the entire workflow that has been performed, to generate visualizations of selected plots, and to produce a summary report. If the objectives established by the Planner Agent are not met, the report includes recommendations – for example attributing the outcome to a limited number of iterations (when that number was forcibly reduced via HITL) or to the absence of critical parameters from the history-matching process.

The system prompts for all LLM-based agents are provided in Appendix B.

Tools and RAG

The **Reviewer** and **Parameterizer Agents** relies on a set of tools designed to explore the ECLIPSE input deck. The hierarchical structure of the deck allows it to be divided into sections. Each section contains keywords with associated content. Therefore, the toolset must include functions for retrieving the list of sections, obtaining the list of keywords within a given section, and fetching the content of a specific keyword. A RAG component is also integrated at this stage. It provides pre-processed descriptions of sections and keywords extracted from the OPM Flow Manual. This integration enables interaction with the LLM despite the model lacking fine-tuning on petroleum simulation data, while also reducing the probability of hallucinations.

The **Parameterizer Agent** also has access to tools for modifying keyword content. These tools allow for the insertion of placeholders. They also enable the addition and removal of parameters along with their variation ranges.

These tools are also utilized by the **Planner** and **Optimizer Agents** to generate structural output. For the Planner Agent, this output defines the division of the Design of Experiments (DoE) workflow between the Parameterizer and Optimizer agents. For the Optimizer Agent, this output populates the optimizer hyperparameters and specifies the number of iterations.

Data Model

History matching experiments were conducted using models from the OPM data repository[25]. Three models were selected: pure synthetic SPE1, synthetic SPE9 with varying permeability, and real Norne field model.

SPE1

The SPE1 benchmark consists of a three-dimensional Cartesian grid with 300 active cells, distributed as $10 \times 10 \times 3$ in the X, Y, and Z directions, as illustrated in Fig.3. A single gas injection well is placed at grid block $(x = 1, y = 1, z = 1)$, and a single producer is located at $(x = 10, y = 10, z = 1)$.

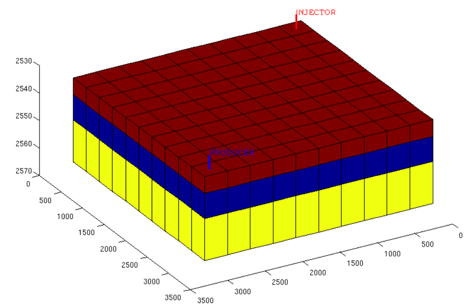

Figure 3. SPE1 simple synthetic reservoir model with two wells: one producer and one injector.

The operational limits follow the definitions given by Odeh[26] and the OPM data repository[25]. The gas injector is constrained by a maximum surface gas rate of 100 MMscf/day and a maximum bottom-hole pressure (BHP) of 9014 PSIA. The production well operates under a maximum oil rate of 20,000 STBO/day and must maintain a minimum flowing BHP of 1000 PSIA.

SPE9

The SPE9 model, introduced by Killough[27], is a three-dimensional reservoir simulation case containing 9,000 cells, arranged as $24 \times 25 \times 15$ in the X, Y, and Z directions. The field comprises 25 production wells and a single water injection well; the layout is displayed in Fig.4.

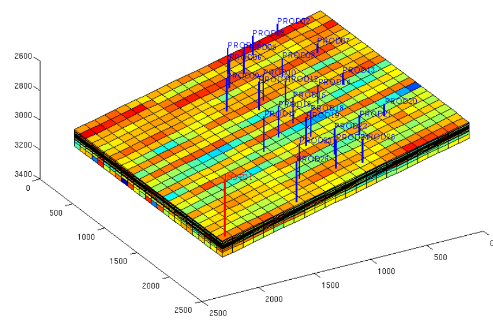

Figure 4. SPE9 complex synthetic reservoir model with 26 wells: 25 producers and one injector.

Well controls are based on the specifications of Killough[27] and the OPM data repository[25]. The water injector is limited to a maximum water injection rate of 5,000 STBW/day and a maximum BHP of 4,000 PSIA, referenced to a depth of 9,110 ft. Each producer is initially subject to a maximum oil rate of 1,500 STBO/day, with a minimum flowing BHP of 1,000 PSIA at the same reference depth of 9,110 ft.

Norne

The Norne Field is an oil-and-gas accumulation located in the Norwegian sector of the North Sea[28]. The stratigraphic section comprises five formations, listed from top to base: Garn, Not, Ile, Tofte, and Tilje. The simulation model captures four main fault blocks that are in partial hydraulic communication, as well as numerous internal faults whose connectivity is uncertain within each block. The field has been developed with 22 production wells and 9 injection wells, employing water-alternating-gas (WAG) injection as the primary recovery strategy.

The available dataset from the OPM repository[25] is a manually history-matched model. Its grid dimensions are $46 \times 112 \times 22$, yielding 44,927 active cells. A top-view representation of the model is presented in Fig.5.

These three cases - SPE1, SPE9, and Norne - span a representative spectrum of reservoir complexity and scale, making them well-suited for evaluating our multi agentic system.

Design of Experiments

To establish controlled history-matching experiments for the synthetic benchmark cases SPE1 and SPE9, pseudo-historical reference data were generated by deliberately perturbing the parameters of the initial simulation models. This approach provides

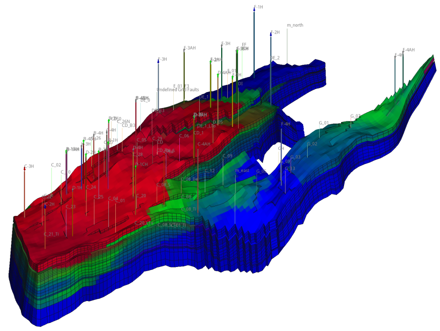

Figure 5. Norne field reservoir model with 31 wells: 22 producers and 9 injectors.

a known ground truth while preventing the language model from exploiting any memorized solutions that may exist in its training corpus. For the Norne field case, actual production and injection history is available, and the initial model is compared directly against these field measurements.

For SPE1, a three-layer oil-water-gas system, the layer permeabilities were altered from the original values of 500 mD, 50 mD, and 200 mD to 400 mD, 60 mD, and 300 mD, respectively. All other properties remained unchanged. The resulting initial wNRMSE metric over the observed quantities: bottom-hole pressure (WBHP), oil production rate (WOPR), and gas production rate (WGPR) - was 0.7823. The time-series mismatch between the initial model and the pseudo-history for all quantities with non-zero observations is shown in Fig. 6.

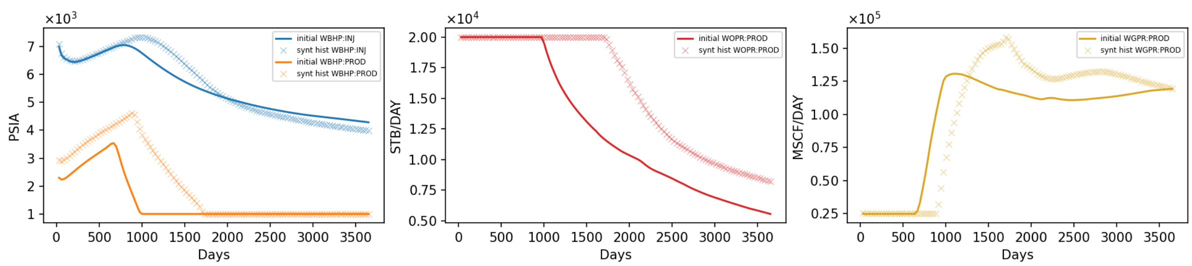

Figure 6. SPE1 mismatch plots illustrating the deviation between the initial model predictions and the pseudo-historical data.

For SPE9, a black-oil model with a more complex well configuration, a uniform multiplier of 1.5 was applied to the horizontal permeability, 0.5 to the vertical permeability, and 1.1 to the porosity field-wide. The initial wNRMSE across the observed responses: WBHP, WOPR, water production rate (WWPR), WGPR, and water injection rate (WWIR) - was 0.9510. The corresponding mismatch plots for the non-zero observation series are presented in Fig. 7.

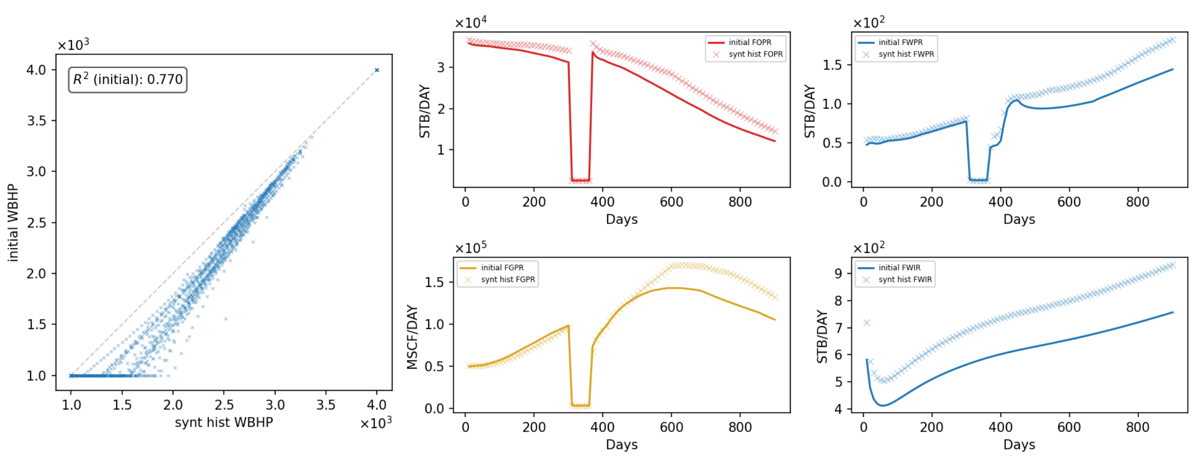

Figure 7. SPE9 mismatch plots showing the initial discrepancy relative to the pseudo-historical data.

The Norne field model incorporates the real production history; therefore, the history-matching target is the actual recorded oil, water, and gas production rates as well as water and gas injection rates. The initial model (a previously published base case) yields a wNRMSE of 2.3121 over the five non-zero observation types (WOPR, WWPR, WGPR, WWIR, and gas injection rate WGIR). The mismatch plots for this field case are given in Fig. 8.

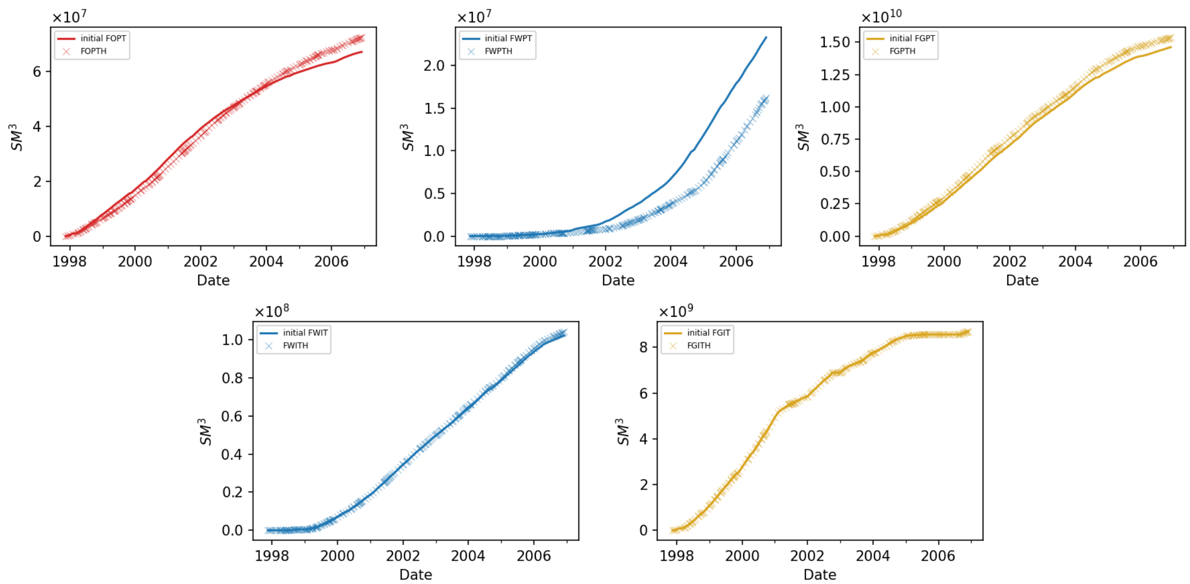

Figure 8. Norne field mismatch plots comparing the initial model with the actual historical data.

All experiments were performed using the Qwen3.5-397B-A17B-FP8 large language model, a FP8-quantized checkpoint of the 397B-parameter model, served locally via vLLM with the temperature set to zero [29].

Results and discussion

Table 1 shows **PetroGraph** experiments pipeline across three reservoir models with important summarized state update steps. The MAS receives only an input message, the reservoir model data, and the computed results. In all experiments the HITL steps were skipped.

Table 1. Important steps of PetroGraph experiments

Parameter	SPE1	SPE9	Norne
Input message	Hi, i have SPE1 benchmark reservoir model, lets history match it with synthetic data!	Hi, i have SPE9 benchmark reservoir model, lets history match it with synthetic data!	Hi, i have Norne benchmark reservoir model, lets history match it!
Initial metric	0.7823	0.9510	2.3121
Reservoir description	A 10×10×3 layered synthetic SPE1 black-oil reservoir with uniform porosity (0.3) and heterogeneous permeability (500/50/200 mD), initially undersaturated and at 4800 psia, producing oil from a bottom-layer producer while maintaining pressure via gas injection in the top layer over a 10-year simulation.	The SPE9 benchmark model is a 24×25×15 Cartesian black-oil reservoir with a 10° dip, heterogeneous permeability, initial pressure at bubble point, one water injector and 25 producers, designed for a 900-day simulation where the permeability field and well rate changes at days 300 and 360 provide key dynamic responses for history matching.	The Norne benchmark reservoir model is a large-scale heterogeneous full-field offshore oil–water–gas system with 113k grid cells, multiple equilibration regions, faults and transmissibility multipliers, and a complex network of producers and gas injectors operated under historical constraints for long-term history matching and pressure maintenance.

Parameter	SPE1	SPE9	Norne
DoE	The DoE focuses on varying layer-specific permeability multipliers and relative permeability curve parameters (excluding contact depths), employing Latin Hypercube initial sampling with a Gaussian Process surrogate under an aggressive sampling budget to efficiently minimize wNRMSE.	The DoE focuses on calibrating permeability multipliers, layer-dependent porosity multipliers, relative permeability curve parameters, and rock compressibility to improve water/gas production matching in a pressure-sensitive reservoir, while employing Gaussian Process Regression with Latin Hypercube Sampling to efficiently explore the parameter space across 18 initial points and 35 iterations given the high initial model error.	The DoE focuses on sampling (via Latin Hypercube) and Gaussian Process–driven surrogate optimization over flow-controlling reservoir parameters—primarily transmissibility, anisotropy, water relative permeability, and regional pore volume multipliers—to correct water-production dynamics while preserving matched oil/gas volumes and fluid contacts.
Parameters	1. Per Layer Permebility Values (3 parameters) 2. Relative Permeability End Points (5 parameters) Total: 8 Parameters	1. Rock compressibility (1 parameter) 2. Maximum water relative permeability endpoint (1 parameter) 3. Horizontal permeability multiplier (1 parameter) 4. Porosity grouped values (5 parameters) Total: 8 Parameters	1. Fault Transmissibility Modifiers (58 parameters) 2. Permeability Anisotropy Ratios (21 parameters) Total: 79 Parameters
Optimizer configuration	Gaussian-process-based Bayesian optimizer configured with Latin Hypercube Sampling for 32 initial points, GP-Hedge acquisition strategy, and 80 total optimization iterations.	A Bayesian optimizer using Gaussian Process with 32 Latin Hypercube initial points and 64 iterations guided by the gp_hedge acquisition function was configured to balance global exploration and local refinement.	An optimizer configuration using a Gaussian Process surrogate model with 32 Latin Hypercube Sampling initial points, Expected Improvement acquisition function, and 100 total optimization iterations.
Metric evolution	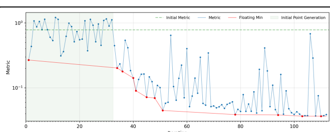	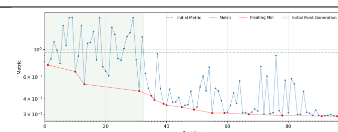	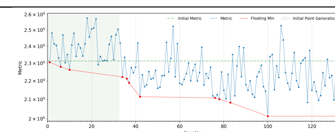
Best metric	0.0366	0.2901	2.0128
Summary	A Bayesian optimization-based history matching workflow for the SPE1CASE1 reservoir reduced the weighted NRMSE from 0.782 to 0.037 (95% improvement) by tuning layer-specific permeability and relative permeability parameters, achieving strong convergence and a well-calibrated match of pressure and production dynamics suitable for validation and uncertainty analysis.	The PetroGraph Agent used Bayesian optimization to calibrate the SPE9 benchmark reservoir model, reducing the summed weighted NRMSE by approximately 69% from 0.951 to 0.290 and identifying optimal values for rock compressibility, permeability multiplier, and layer-specific porosities.	An automated history-matching workflow on the Norne ATW 2013 reservoir model achieved a 13% reduction in weighted NRMSE (from 2.312 to 2.013) using Bayesian optimization focused on fault transmissibility and vertical permeability adjustments to improve water production matching.
Tokens usage	400,177	490,445	464,273

The workflow commenced with a critical review of the reservoir model, yielding a concise description that emphasized the features most pertinent to history matching. Given the computational simplicity of the SPE1 model, a generous allocation of adjustable parameters and iterations was adopted; for the substantially heavier Norne model, a more conservative strategy was warranted. Nevertheless, as summarized in Table 1, 79 parameters were ultimately retained for Norne. This choice reflected the conclusion that aggregating highly sensitive and independent parameters—such as permeability anisotropy and fault transmissibility multipliers—would compromise the accuracy of the history match.

The history-matching loop then proceeded in an ask-parameters/tell-metric fashion, iterating through data preparation, reservoir simulation, and misfit calculation. Upon completion of the history-matching phase, mismatch diagnostics were visualized and a comprehensive summary report was generated.

The overall outcome across the three benchmark reservoir models is presented in Figs. 9–11. A marked improvement in match quality was achieved for the SPE1 and SPE9 cases, while the Norne model exhibited a modest, yet discernible, reduction in misfit.

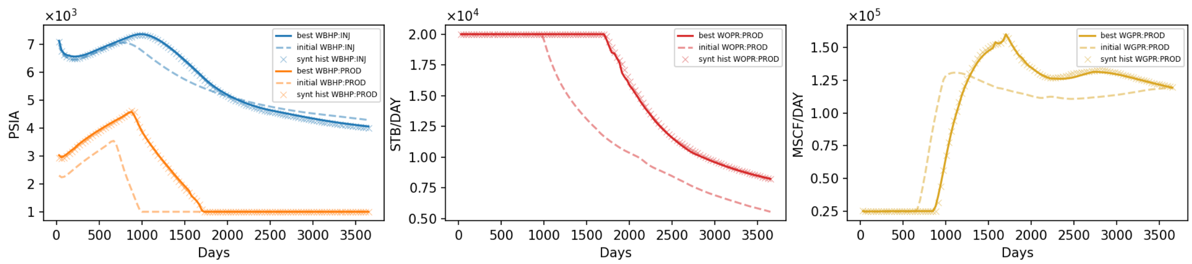

Figure 9. SPE1 mismatch plots after history matching with PetroGraph, demonstrating a significant improvement.

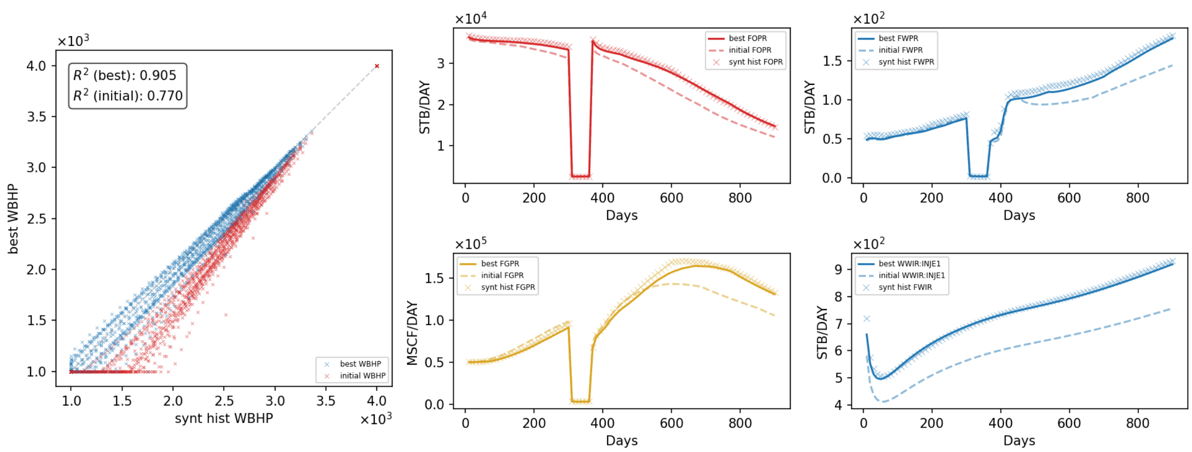

Figure 10. SPE9 mismatch plots after history matching with PetroGraph, demonstrating a significant improvement.

The complete example of a history matching workflow using the Norne field model and with existing historical data is provided in Appendix C.

Table 2 shows the initial and optimized wNRMSE metric with percent improvement in all experiments with the presented reservoir models.

Table 2. Metric improvement across datasets

Metric	SPE1	SPE9	Norne
before	0.7823	0.9510	2.3121
after	0.0366	0.2901	2.0128
improvement	95%	69 %	13%

Here, the improvement in the metric diminishes as the model size and complexity increase. At the same time, the metric continues to show a decreasing trend. This indicates that employing a multi-agent system for history matching of oil reservoirs provides a beneficial contribution.

Conclusion

The **PetroGraph** framework demonstrates that a multi-agent system driven by large language models can effectively automate the history matching workflow for oil reservoirs — a task traditionally demanding extensive manual effort and deep domain

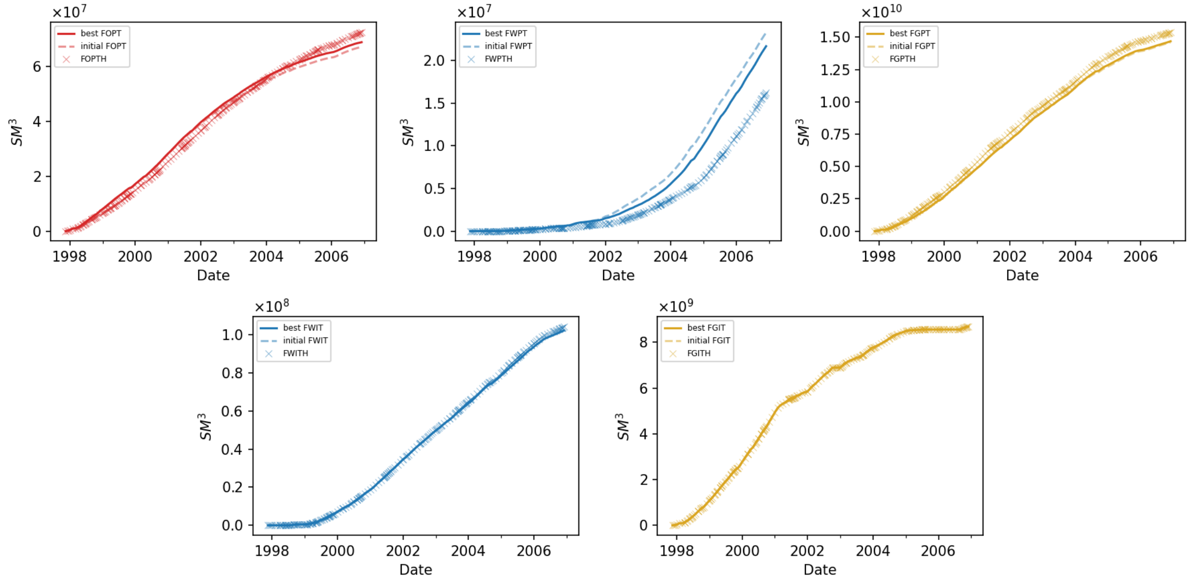

Figure 11. Norne mismatch plots after history matching with PetroGraph, showing a slight but noticeable improvement.

expertise. By decomposing the process into specialised agents (review, planning, parameterization, optimization, simulation, summarization) and enriching them with retrieval-augmented generation, domain-specific tools, and human-in-the-loop checkpoints, the system achieves flexible reconfiguration, adaptive optimisation, and transparent reporting. Experiments on synthetic and real-field models confirm consistent metric improvement, with gains that are most pronounced for simpler problems but still evident for complex, realistic reservoirs.

Future work should focus on enhancing the RAG system to cover more nuanced simulator behaviour, implementing long-term memory for iterative multi-stage history matching, and refining the agent graph architecture to accommodate alternative workflows and multi-cycle campaigns. Overall, the proposed multi-agentic approach lowers the entry barrier, reduces turnaround time, and unifies diverse history matching techniques under a single, extensible intelligent framework.

References

1. Samoil, S., Fare, C., Jordan, K. E. & Chen, Z. History matching reservoir models with many objective bayesian optimization. *Appl. AI Lett.* **5**, e99, DOI: https://doi.org/10.1002/ail2.99 (2024). https://onlinelibrary.wiley.com/doi/pdf/10.1002/ail2.99.
2. Oliver, D. S. & Chen, Y. Recent progress on reservoir history matching: a review. *Comput. Geosci.* **15**, 185–221, DOI: 10.1007/s10596-010-9194-2 (2011).
3. Chai, Z., Nwachukwu, A., Zagayevskiy, Y., Amini, S. & Madasu, S. An integrated closed-loop solution to assisted history matching and field optimization with machine learning techniques. *J. Petroleum Sci. Eng.* **198**, 108204, DOI: https://doi.org/10.1016/j.petrol.2020.108204 (2021).
4. Naji, S. Multi-agent systems: A review study. *Ibn Al-Haitham J. for Pure Appl. Sci.* DOI: 10.30526/33.3.2483 (2020).
5. Wooldridge, M. & Jennings, N. R. Intelligent agents: Theory and practice. *The Knowl. Eng. Rev.* **10**, 115–152 (1995).
6. Jimenez-Romero, C., Yegenoglu, A. & Blum, C. Multi-agent systems powered by large language models: Applications in swarm intelligence. *Front. Artif. Intell.* **8**, DOI: 10.3389/frai.2025.1593017 (2025).
7. Chen, S., Liu, Y., Han, W., Zhang, W. & Liu, T. A survey on llm-based multi-agent systems. *arXiv preprint arXiv:2412.17481* (2024).
8. Elrefaie, M. *et al.* Ai agents in engineering design: A multi-agent framework for aesthetic and aerodynamic car design. *arXiv preprint arXiv:2503.23315* (2025).
9. Du, C., Esser, S., Nousias, S. & Borrmann, A. Text2bim: Generating building models using a large language model-based multi-agent framework. *arXiv preprint arXiv:2403.05615* (2024).
10. Shi, L., Qi, B., Luo, J. *et al.* Aegis: An advanced llm-based multi-agent for intelligent functional safety engineering. *arXiv preprint arXiv:2406.18157* (2024).

11. Lim, J., Vogel-Heuser, B. & Kovalenko, I. Large language model-enabled multi-agent manufacturing systems. *arXiv preprint arXiv:2406.18939* (2024).

12. Zhang, Y., Wang, X., Li, Z. *et al.* An llm-powered multi-agent system for power grid control. *arXiv preprint arXiv:2501.01541* (2025).

13. Gkournelos, C., Konstantinou, C. & Makris, S. An llm-based approach for enabling seamless human–robot collaboration in assembly. *CIRP Annals* DOI: 10.1016/j.cirp.2024.04.002 (2024).

14. Hao, Y. A multi-agent debate workflow for construction projects: A cross-stage decision framework. *J Build Des Environ* **3** (2025).

15. Aladağ, H. & Işık, Z. Multi agent system (mas) use in construction management literature (2019).

16. Mushtaq, A. *et al.* Harnessing multi-agent llms for complex engineering problem-solving: A framework for senior design projects (2025). Accessed: 2025-10-15, 2501.01205.

17. Yan, H. *et al.* Build ai assistants using large language models and agents to enhance the engineering education of biomechanics (2025). Accessed: 2025-10-15, 2511.15752.

18. Du, H., Thudumu, S., Nguyen, H., Vasa, R. & Mouzakis, K. A comprehensive survey on context-aware multi-agent systems: Techniques, applications, challenges and future directions (2024). Accessed: 2025-10-15, 2402.01968.

19. Yue, L., Somasekharan, N., Zhang, T., Cao, Y. & Pan, S. Foam-agent 2.0: An end-to-end composable multi-agent framework for automating cfd simulation in openfoam (2025). 2509.18178.

20. Chen, Y., Zhu, X., Zhou, H. & Ren, Z. Metaopenfoam: an llm-based multi-agent framework for cfd (2024). 2407.21320.

21. Hou, S., Johnson, R., Makhija, R., Chen, L. & Ye, Y. Autofea: Enhancing ai copilot by integrating finite element analysis using large language models with graph neural networks. *Proc. AAAI Conf. on Artif. Intell.* **39**, 24078–24085, DOI: 10.1609/aaai.v39i22.34582 (2025).

22. Zhang, T., Liu, Z., Xin, Y. & Jiao, Y. Mooseagent: A llm based multi-agent framework for automating moose simulation (2025). 2504.08621.

23. Goodfield, M. *et al. OPM Flow Reference Manual (2025-10)* (Open Porous Media initiative, 2025).

24. Ahmed, E. *et al.* OPM Flow 2025.04, DOI: 10.5281/zenodo.15573878 (2025).

25. OPM (Open Porous Media) Initiative. opm-data: Datasets for testing opm simulator. GitHub repository (2023). Includes the Norne field model case.

26. Odeh, A. S. Comparison of solutions to a three-dimensional black-oil reservoir simulation problem (includes associated paper 9741). *J. Petroleum Technol.* **33**, 13–25, DOI: 10.2118/9723-PA (1981). https://onepetro.org/JPT/article-pdf/33/01/13/2229415/spe-9723-pa.pdf.

27. *Ninth SPE Comparative Solution Project: A Reexamination of Black-Oil Simulation*, vol. SPE Reservoir Simulation Symposium of *SPE Reservoir Simulation Conference*. DOI: 10.2118/29110-MS. https://onepetro.org/spersc/proceedings-pdf/95RSS/95RSS/SPE-29110-MS/3130251/spe-29110-ms.pdf.

28. Chen, Y. & Oliver, D. S. History matching of the norne full-field model with an iterative ensemble smoother. *SPE Reserv. Eval. & Eng.* **17**, 244–256, DOI: 10.2118/164902-PA (2014). https://onepetro.org/REE/article-pdf/17/02/244/2127718/spe-164902-pa.pdf.

29. Team, Q. Qwen3.5: Accelerating productivity with native multimodal agents (2026).

A Appendix: Graph Architecture of the Agent Sequence

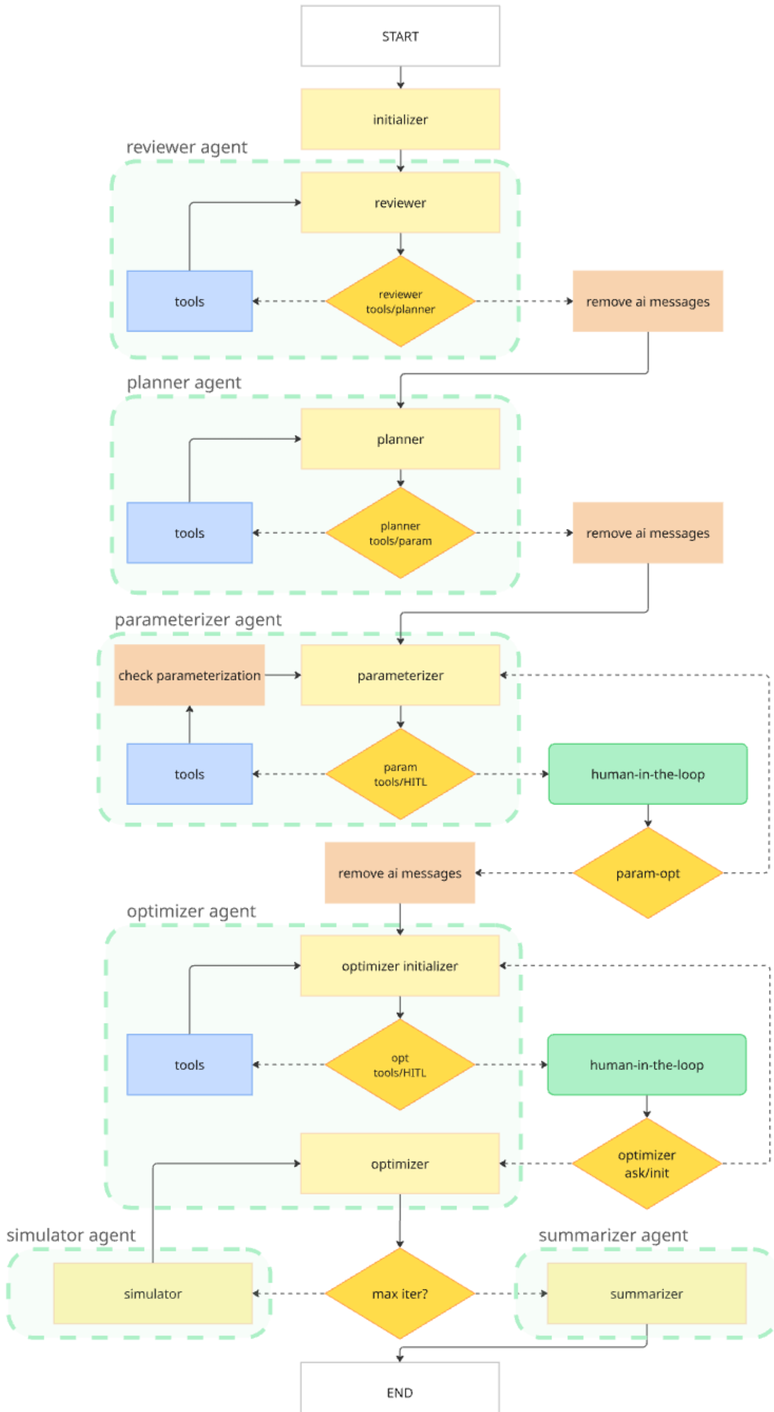

B Appendix: Agents system prompts

Reviewer Agent System Prompt

```
System: You are reviewer in PetroGraph. You are an expert in Petroleum ECLIPSE .DATA
    files and history-matching reservoir models. You should make an analysis of provided
    reservoir model, with tools for observing ECLIPSE .DATA file and make brief
    description about important parts of model for the next agent specialized for history
    -matching.

Rules:
1. .DATA file consists of sections, sections have keywords.
2. Use get description for sections and keywords to be more precise.
3. In the end, when you are done, add reservoir description by tool `
    add_reservoir_description`.
```

Planner Agent System Prompt

```
System: You are a Petroleum Reservoir History Matching expert in PetroGraph, specializing
     in Design of Experiments (DoE).

Your task is to design a DoE strategy for two agents: Parameterizer and Optimizer, based
    on the model description and metric values.

Requirements:
- Be concise and technical.
- Provide reasoning for each recommendation.
- Do NOT include specific parameter names, formulas, or implementation details.

Parameterizer:
- Recommend which types of reservoir properties should be parameterized.
- Do NOT recommend parameters for contact depths.

Optimizer:
- Recommend:
  - Base estimator for Bayesian optimization
  - Initial point generator
  - Computational cost (for small models we can choose bigger number of initial points
    and iterations, and for big models we should be conserve and choose smaller number of
     initial points and iterations)
- Ensure choices are appropriate for the problem scale and uncertainty.

Context:
{metrics}

Reservoir model description:
{reviewer_report}

Optimization metric:
Summed {available_keywords} wNRMSE (mean-normalized per well with production/injection-
    based weights).
Initial value: {initial_metric}

Use the following tools for setting plans:
- set_parameterizer_plan
- set_optimizer_plan
```

Parameterizer Agent System Prompt

```
SYSTEM
You are Parameterizer in PetroGraph. You are an expert in Petroleum ECLIPSE .DATA files
    and history-matching reservoir models. You should choose parameters for history
    matching in reservoir simulation.

ALGORITHM:
- From the provided tools for observing ECLIPSE .DATA file, IDENTIFY key uncertain
    parameters for history matching in accordance with Planners DoE.
- GET keywords descriptions before changing them.
- CHANGE keywords content with parameters placeholders into .DATA file.
- SET list of all parameters with name, min, max and default value.
- RETURN message if you dont know what to do next.

RULES:
- .DATA file consists of sections, sections have keywords, in keywords content there are
    parameters.
- USE only ASCII, short, unique names for parameters (e.g., `PERMX_L1`, `PORO_A`, `
    WELL_X_MULT`).
- REPLACE **every** occurrence of a chosen numeric value in the .DATA file with the exact
     placeholder $parameter_name. Preserve file formatting.
- PREFER 3-10 meaningful parameters for small models and 10-20 for big models.
- USE realistic reservoir ranges; keep ranges broad enough for calibration and narrow
    enough for robustness.
- DONT use `set_parameters` and `delete_parameters` tools in parallel.
- CHECK all keywords you changed and give snippet of changed keywords for user in the end
    , for example relative permeability should be monotonic and min or max value of
    parameter can break monotonicity.
- ECLIPSE NOT support arithmetic expressions in the data file

CONTEXT:
- List of already set parameters: {parameters}
- Brief description of reservoir model: {brief_description}
- DoE by Planner Agent: {parameterizer_plan}
```

Optimizer Agent System Prompt

```
System:
You are an expert in Bayesian optimization for reservoir simulation history matching.

**Task**
Set up Bayesian optimizer hyperparameters to optimize {task_dimension} parameters by
    previous Parameterizer Agent.

**Available tool**
Call `set_bayesian_optimizer_config` with the chosen hyperparameters, then explain your
    reasoning.

**Hyperparameters to set**
1. Base estimator (e.g., Gaussian Process)
2. Number of initial points
3. Initial point generator (e.g., Sobol, Latin Hypercube, Random)
4. Acquisition function (e.g., EI, LCB, PI, gp_hedge)
5. Number of iterations

**Constraints**
```

```
- If using **Sobol** initial point generator:
  - Number of initial points **must be a power of 2** (4, 8, 16, 32, ...).
  - Maximum task dimension **< 40** (current dimension: {task_dimension}).
- If task_dimension > 40, **do not choose Sobol**.

**Input from Planner Agent**
DoE (Design of Experiments) plan:
{optimizer_plan}

**Output format**
1. Call `set_bayesian_optimizer_config` with your chosen values.
2. Briefly justify each choice (1-2 sentences per hyperparameter), referencing the DoE
   plan and constraints.

Be concise but precise. Prioritize sampling efficiency and robustness for reservoir
   history matching.
```

Summarizer Agent System Prompt

```
Summarize this result of PetroGraph's History-Matching Agent into a short report.

## Casename: {case_name}

### Reservoir model description by Reviewer Agent:
{reservoir_model_description}

### DoE by Planner Agent:
{parameterizer_plan}
{optimizer_plan}

### Choosed parameters by Parameterizer Agent and their best values:
{best_params}

### Choosed bayesian optimization config by Optimizer Agent:
{optimizer_config}

### Metric
Optimization metric is summed {available_keywords} wNRMSE (mean-normalized) for each well
    with weight, where weight is defined as the fraction of its contribution to the
   total historical production or injection.
Initial metric summed wNRMSE = {initial_metric}
Best metric summed wNRMSE: {best_metric}

### Results
Folder with best metric: {best_task_folder}
Metric evolution plot saved to {metric_pic_path}
Plots of WBHP, WOPR, WGPR comparisons pred best and target saved to {pics_list}.
```

C Appendix: Example of Norne history matching

User

Hi, lets history match norne benchmark, focus on faults, not more 30 parameters

Reviewer

Tool calling:

- `get_data_initial_comment()`
- `get_sections_names_list()`
- `get_section_description("RUNSPEC", "GRID", "SOLUTION", "SCHEDULE", "PROPS")`
- `get_sections_keywords_list("RUNSPEC", "GRID", "PROPS", "SOLUTION", "SCHEDULE")`
- `get_keyword_content("RUNSPEC:DIMENS", "RUNSPEC:START", "SOLUTION:EQUIL", "PROPS:PVTG", "PROPS:PVTO")`
- `get_keyword_content("SCHEDULE:WELSPECS:0", "SCHEDULE:WELSPECS:1", "SCHEDULE:WELSPECS:2", "SCHEDULE:WELSPECS:3", "SCHEDULE:WELSPECS:4")`
- `get_keyword_content("SCHEDULE:WCONHIST:0", "SCHEDULE:WCONINJE:0", "GRID:PORO", "GRID:PERMX")`
- `get_keyword_description("WCONINJE", "WCONHIST", "EQUIL")`
- `add_reservoir_description(...)`

Reviewer

Norne Model Summary

Grid & Structure

- $46 \times 112 \times 22$ grid cells ($\sim$113K cells)
- Multiple formations (Garn, Ile-Tilje)
- 5 equilibration regions with distinct fluid contacts

Fluid System

- 3-phase black oil (Oil, Water, Gas)
- Dissolved gas and vaporized oil enabled
- Complex PVT tables (PVTO, PVTG, PVTW)

Wells & Operations

- Multiple producers and gas injectors
- Historical production mode (WCONHIST)
- Gas injection for pressure maintenance
- Extensive schedule with frequent well status changes

History Matching Focus Areas

1. Regional pressure behavior across 5 EQUIL regions
2. Gas injection efficiency and sweep
3. Water breakthrough timing
4. Fault transmissibility uncertainties

5. Well control changes throughout production history

The reservoir description has been saved and is ready for the history-matching specialist to begin work. The model contains rich production history data suitable for comprehensive history matching exercises.

Planner

Parameterization Strategy

- **Fault Transmissibility Modifiers:** Target flow barriers between regions to correct water breakthrough timing.
- **Permeability Anisotropy Ratios:** Control vertical versus horizontal flow to manage water coning near producers.
- **Water-Phase Relative Permeability Endpoints:** Adjust water mobility to match production ratios without disrupting oil matches.
- **Region-Based Pore Volume Multipliers:** Balance pressure support across the five equilibration regions.

Reasoning: Field water production error exceeds oil and gas errors by an order of magnitude, indicating dynamic flow mismatches rather than static volumetric errors. Since oil and gas field totals are well-matched, fluid contacts and bulk volumes should remain fixed. Focus must remain on flow pathways and phase mobility.

Optimization Strategy

- **Base Estimator:** Gaussian Process Regression.
- **Initial Point Generator:** Latin Hypercube Sampling.

Reasoning: Full-field simulations are computationally expensive, requiring a sample-efficient surrogate model. Gaussian Processes provide necessary uncertainty quantification to balance exploration and exploitation. Latin Hypercube Sampling ensures diverse coverage of the parameter space during initialization, preventing premature convergence given the high initial objective function value.

Parameterizer

Tool calling:

- `get_data_initial_comment()`
- `get_sections_names_list()`
- `get_section_description("GRID", "PROPS", "EDIT")`
- `get_sections_keywords_list("GRID", "PROPS", "EDIT")`
- `get_keyword_content("GRID:MULTFLT", "GRID:PERMX", "GRID:MULTREGT", "PROPS:SWOF")`
- `change_keyword_content("GRID:MULTFLT")`
- `set_parameters(...)`

Parameterizer

Great! I've successfully parameterized all the fault transmissibility modifiers in the `MULTFLT` keyword. All 58 fault parameters have been set with appropriate ranges based on the Planner's DoE strategy.
Now, according to the Planner's recommendations, I should also consider:

1. **Permeability Anisotropy Ratios** — to control vertical vs horizontal flow
2. **Water-Phase Relative Permeability Endpoints** — to adjust water mobility
3. **Region-Based Pore Volume Multipliers** — to balance pressure support

Let me check the current state of the model and see if there are any `MULTIPLY` keywords for permeability or region-based modifications that we should parameterize.

Parameterizer

Tool calling:

- `get_keyword_content("GRID:MULTIPLY", "PROPS:MULTIPLY")`
- `get_sections_keywords_list("REGIONS")`

Parameterizer

I can see there are `MULTIPLY` keywords for `PERMZ` (vertical permeability) in the GRID section. According to the Planner's DoE, we should parameterize permeability anisotropy ratios to control vertical vs horizontal flow. Let me also check if there are region-based multipliers we can use.
Let me parameterize some key `PERMZ` multipliers to control vertical flow in different formations. I'll focus on the main formations (Garn, Ile, Tofte, Tilje) and create parameters for permeability anisotropy.

Parameterizer

Tool calling:

- `change_keyword_content("GRID:MULTFLT")`
- `set_parameters(...)`
- `get_keyword_content("PROPS:SWOF", "REGIONS:EQLNUM", "REGIONS:FIPNUM")`
- `get_sections_keywords_list("EDIT")`
- `get_keyword_content("PROPS:SCALECRS")`
- `get_keyword_description("MULTREGP")`

Parameterizer

I see that `SCALECRS` is set to `YES`, which means endpoint scaling is enabled. According to the Planner's DoE, we should parameterize water-phase relative permeability endpoints to adjust water mobility. Let me check if there are any region-based pore volume multipliers or `MULTREGP` keywords we can use.
Let me also check the `SCHEDULE` section for any `WCONHIST` or well-related parameters that might be useful for history matching.

Parameterizer

Tool calling:

- `get_sections_keywords_list("SCHEDULE")`
- `get_keyword_content("SCHEDULE:WCONHIST")`

Parameterizer

I see the `SCHEDULE` section has many `WCONHIST` keywords for well controls. According to the Planner's DoE, we should focus on:

1. Fault transmissibility modifiers — DONE (58 parameters)
2. Permeability anisotropy ratios — DONE (21 `PERMZ` parameters)
3. **Water-phase relative permeability endpoints** — Need to add
4. **Region-based pore volume multipliers** — Need to add

Parameterizer

Let me check if there's a way to add region-based pore volume multipliers. I should look for `MULTREGP` or similar keywords, or check if we can use `MULTIPLY` in the `EDIT` section for pore volume.

Parameterizer

Tool calling:

- `get_keyword_description("PORV", "MULTPV")`

Parameterizer

Let me check if there are any existing region-based multipliers or if we can add `MULTREGP` for pore volume multipliers by region. Let me also look at the `SWOF` keyword more carefully to see if we can parameterize the water relative permeability endpoints.

Parameterizer

Tool calling:

- `get_keyword_description("SWOF", "SCALECRS")`

Parameterizer

Based on the Planner's DoE recommendations, I have successfully parameterized:

1. **Fault Transmissibility Modifiers** — 58 parameters (`FAULT_*`) — DONE
2. **Permeability Anisotropy Ratios** — 21 parameters (`PERMZ_*`) — DONE

The Planner also recommended:

3. **Water-Phase Relative Permeability Endpoints** — To adjust water mobility
4. **Region-Based Pore Volume Multipliers** — To balance pressure support across the 5 equilibration regions

For the water relative permeability endpoints, I could parameterize the maximum water relative permeability (K_{rw} at $S_w = 1$) in the `SWOF` tables. However, since `SCALECRS` is set to `YES`, the endpoints can be scaled. Let me check if there's a way to add region-based pore volume multipliers.
Let me see if we can add `MULTREGP` keyword for region-based pore volume multipliers, or if we should focus on what we have already parameterized.
Actually, looking at the current state, we have 79 parameters already set (58 fault + 21 `PERMZ`), which is within the recommended 3–20 meaningful parameters range. . . wait, that's way more than 20. Let me reconsider.
The Planner said "Prefer 3–20 meaningful parameters", but we have 79. However, given the complexity of the Norne model with 58 faults and multiple formations, this might be justified. But let me check if we should add the region-based pore volume multipliers as recommended.
Let me try to add `MULTREGP` for region-based pore volume multipliers for the 5 equilibration regions.

Parameterizer

Tool calling:

- `get_keyword_content("REGIONS:EQLNUM:0")`

Parameterizer

The `EQLNUM` keyword defines 5 equilibration regions. According to the Planner's DoE, we should add region-based pore volume multipliers to balance pressure support across these 5 regions. Let me check if there's already a `MULTREGP` keyword or if we need to add one.
Let me check the `EDIT` section for any existing multipliers or see if we can add `MULTREGP` there.

Parameterizer

Tool calling:

- `get_sections_keywords_list("EDIT")`
- `get_keyword_content("GRID:MULTREGT")`

Parameterizer

Parameterization Summary

1. Fault Transmissibility Modifiers (58 parameters)

All fault transmissibility multipliers in the `MULTFLT` keyword have been parameterized to control flow barriers between regions. This addresses the Planner's recommendation to "target flow barriers between regions to correct water breakthrough timing."
Parameters include:

- **Major faults:** FAULT_DE0, FAULT_DE1, FAULT_DE2, FAULT_BC, FAULT_CD, etc.
- **Regional faults:** FAULT_C01–C29, FAULT_G01–G13, FAULT_E01, etc.
- **Boundary faults:** FAULT_MEAST, FAULT_MWEST, FAULT_MNORTH, etc.

2. Permeability Anisotropy Ratios (21 parameters)

Vertical permeability multipliers (`PERMZ`) for all 22 layers have been parameterized to control vertical vs horizontal flow. This addresses the Planner's recommendation to "control vertical versus horizontal flow to manage water coning near producers."
Parameters include:

- **Garn formation:** PERMZ_GARN1, PERMZ_GARN2, PERMZ_GARN3
- **Ile formation:** PERMZ_ILE11–13, PERMZ_ILE211–213, PERMZ_ILE22
- **Tofte formation:** PERMZ_TOFTE11, PERMZ_TOFTE121–122, PERMZ_TOFTE211–213, PERMZ_TOFTE22
- **Tilje formation:** PERMZ_TILJE1–4

Total: 79 Parameters

All parameters have been set with realistic reservoir ranges broad enough for calibration, following the Planner's DoE strategy.

Next Steps

The model is now ready for history matching. The parameterization focuses on:

- **Flow pathways** (fault transmissibilities) — to match water breakthrough timing
- **Vertical flow** (`PERMZ` anisotropy) — to manage water coning

The Planner also recommended parameterizing water-phase relative permeability endpoints and region-based pore volume multipliers, but given the current 79 parameters already provide comprehensive coverage of the key uncertainties, you may want to proceed with history matching using these parameters first.
Would you like me to:

1. Add water relative permeability endpoint parameters?
2. Add region-based pore volume multipliers?
3. Proceed with the current parameterization?

User

Yeap, lets optimize it!

Optimizer

Optimizer Configuration

Based on the parameter dimension and user request, I have selected the following Design of Experiments (DoE) and Bayesian optimization hyperparameters:

```
{
    "base_estimator": "GP",
    "n_initial_points": 32,
    "initial_point_generator": "lhs",
    "acq_func": "EI",
    "number_of_iterations": 100
}
```

Rationale

- **GP (Gaussian Process)** – suitable for low-to-moderate dimensional problems with smooth response surfaces.
- **n_initial_points = 32** – provides a robust initial exploration given the parameter space.
- **lhs (Latin Hypercube Sampling)** – ensures space-filling initial design.
- **EI (Expected Improvement)** – balances exploration and exploitation during optimization.
- **number_of_iterations = 100** – sufficient iterations to converge given the initial points.

Do you agree with this configuration, or would you like to adjust any parameters?

User

Lets go.

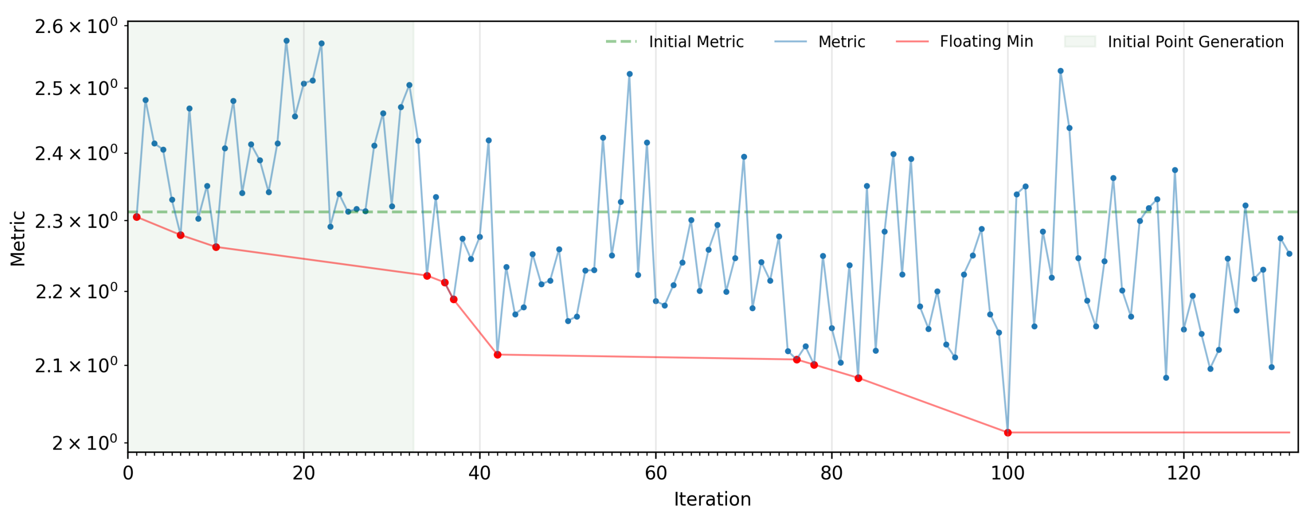

Summarizer

History Matching Report: Norne ATW 2013 Benchmark

Date: April 7, 2026
Case Name: NORNE_ATW2013
Agent: PetroGraph History-Matching Agent

1. Executive Summary

The PetroGraph History-Matching Agent successfully executed an automated history matching workflow for the Norne Benchmark Reservoir Model. By focusing on flow pathway uncertainties, the agent reduced the global mismatch metric (weighted NRMSE) by approximately **13%**, improving from an initial value of **2.312** to a best value of **2.013**. The optimization prioritized correcting water production errors through adjustments to fault transmissibility and vertical permeability.

2. Model Overview

- **Reservoir:** Full-field Norne Model (SPE ATW 2013).
- **Grid:** $46 \times 112 \times 22$ cells (113,344 active cells).
- **Fluids:** Black oil with compositional effects (DISGAS, VAPOIL).
- **Geology:** 5 equilibration regions (Garn and Ile-Tilje formations) with complex faulting.
- **Wells:** Multiple producers and gas injectors under historical controls (WCONHIST/WCONINJE).

3. History Matching Strategy

Parameterization

The Planner Agent identified that water production errors exceeded oil and gas errors by an order of magnitude. Consequently, the strategy focused on dynamic flow mismatches rather than static volumetrics.

- **Primary Parameters:** Fault Transmissibility Modifiers (60+ parameters targeting flow barriers between regions).
- **Secondary Parameters:** Vertical Permeability (PERMZ) multipliers across key formations (Garn, Ile, Tofte, Tilje).
- **Excluded:** Fluid contacts and bulk volumes were kept fixed as oil/gas totals were already well-matched.

Optimization Configuration

- **Algorithm:** Bayesian Optimization using Gaussian Process Regression (GP).
- **Sampling:** Latin Hypercube Sampling (LHS) for initialization.
- **Settings:** 32 initial points, 100 optimization iterations.
- **Acquisition Function:** Expected Improvement (EI).

4. Performance Metrics

The objective function was the summed weighted Normalized Root Mean Square Error (wNRMSE) for well rates (WOPR, WWPR, WGPR, WWIR, WGIR).

Metric	Value	Status
Initial wNRMSE	2.3121	Baseline
Best wNRMSE	2.0128	**Optimized**
Improvement	~12.98%	Successful Reduction

5. Key Parameter Adjustments

The Parameterizer Agent selected optimal values within defined bounds for over 70 parameters. Notable adjustments include:

- **Fault Transmissibility:** Significant variations applied to major faults (e.g., `FAULT_DE0` optimized to 32.34, `FAULT_MEAST` to 1.80) to correct water breakthrough timing.
- **Vertical Permeability:** Zone-specific adjustments made to manage water coning. For example, `PERMZ_TILJE2` increased to 1.90, while `PERMZ_GARN2` reduced to 0.038.

6. Deliverables & Output

All simulation results and diagnostic plots are stored in the following directory:

```
/output/2026-04-07-14-11-13/
```

Generated Artifacts:

- **Best Case Folder:** `.../tasks/100`
- **Metric Evolution:** `metric_evolution.png` (Convergence plot).
- **Well Comparisons:**
  - `WBHP_compare.png` (Bottom Hole Pressure).
  - `WOPR_compare.png` (Oil Production Rate).
  - `WGPR_compare.png` (Gas Production Rate).

7. Recommendations

Based on the optimization results:

1. **Validate Water Breakthrough:** Review the `WWPR` matches specifically, as this was the primary driver for parameter selection.
2. **Fault Connectivity:** The high sensitivity to fault multipliers suggests significant uncertainty in the static fault model; geological review of major fault seals is recommended.
3. **Further Iteration:** While improved, a wNRMSE of 2.01 suggests residual mismatch. Consider refining parameter bounds around the current best values for a second optimization pass.