COHIBA Petrel

Surface Modeling and Depth Conversion

# Depth surfaces that carry their own uncertainty

The structural modeling engine of the Norwegian Computing Center, running natively inside Petrel. Time to depth conversion now returns a surface *and* the standard deviation map that belongs to it.

| ENGINE | HOST | BUILT WITH |
|---|---|---|
| COHIBA · Norsk Regnesentral | Native Petrel plug-in | Harbour Energy Norge · NR |

# A depth map with no standard deviation map is a decision with no margin

Conventional depth conversion returns one surface. It ties at the control wells and quietly opens where there is no well, and nothing in the deliverable says by how much. Velocity error, travel time interpretation error and well pick error never reach the volumetrics.

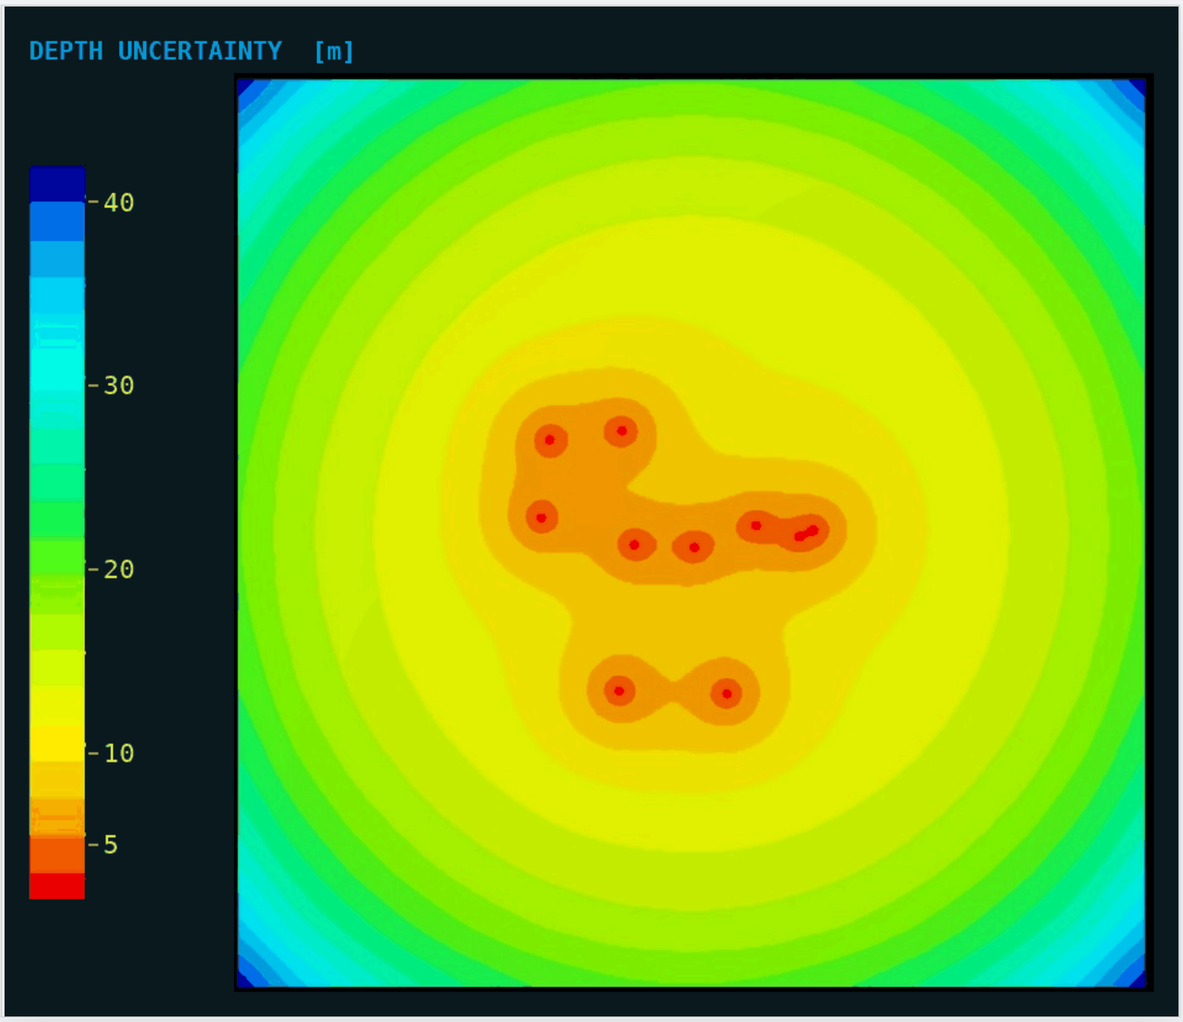


Fig. 01 · Depth uncertainty of a top surface conditioned to ten well points. Around two metres on the wells, tens of metres between and beyond them.

> COHIBA has computed this rigorously for years, but it lived outside Petrel: a command line executable driven by a hand written XML model file. Powerful for specialists, hard to adopt on a working asset team.

# A research grade engine, an industrial interface

COHIBA is developed and maintained by Norsk Regnesentral, the Norwegian Computing Center in Oslo, and is an international reference for probabilistic structural modeling in oil and gas.

DeepSoft built COHIBA Petrel together with Wintershall Dea Norge, now Harbour Energy, and with Norsk Regnesentral. The engine itself is untouched. What changed is everything around it: data mapping, the model file, execution and the round trip of results, all inside Petrel.

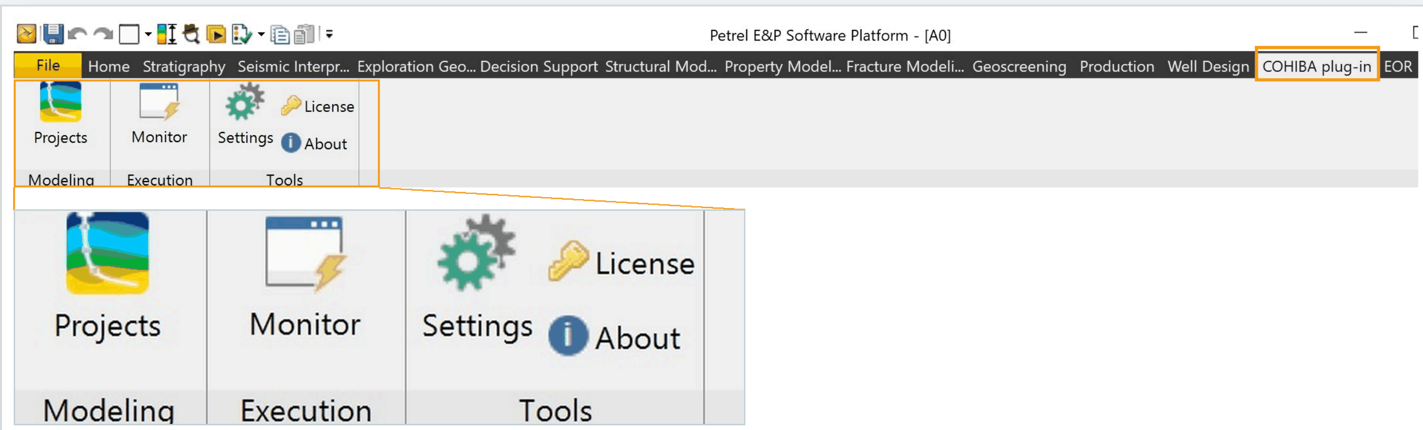


Fig. 02 · A Petrel tab beside Structural Modeling and Property Modeling, not a separate application. Projects, Monitor and Settings sit inside it.

SCIENCE

### Norsk Regnesentral

Engine, algorithms and theory. Oslo, Norway.

FIELD VALIDATION

### Harbour Energy Norge

Requirements and daily use on North Sea assets.

ENGINEERING

### DeepSoft

Ocean plug-in, interface, data mapping and support.

In continuous production use by Harbour Energy, Aker BP and OKEA ASA on time to depth conversion and complex structural modeling.

PROVENANCE

# Every surface is a trend plus a residual

Each horizon is written as a linear combination of known maps, travel time, interval velocity or any other structural map, with coefficients fitted to the well observations, plus a zero mean Gaussian residual whose spatial continuity is set by a variogram.

SURFACE MODEL · TREND PLUS RESIDUAL

$$Z_\ell(\mathbf{x}) = \underbrace{\sum_{p=1}^{P} \beta_p \, f_p(\mathbf{x})}_{\text{trend, large scale}} + \underbrace{\varepsilon_\ell(\mathbf{x})}_{\text{residual, small scale}}$$

$$\varepsilon_\ell \sim \mathcal{N}\left(0,\ \sigma_\ell^2 \, \rho_\ell(\mathbf{h})\right)$$

$Z_\ell$ depth of horizon $\ell$

$f_p$ known maps used as trend terms

$\beta_p$ coefficients fitted from well data

$\sigma_\ell,\ \rho_\ell(\mathbf{h})$ residual SD and variogram


**Fig. 03** · Schematic. Prediction uncertainty collapses at the wells and grows away from them, at a rate set by the variogram.

> Kriging is the best linear unbiased estimator: the prediction honours every conditioning datum while minimising estimation variance.

THE MODEL

# Uncertainty propagates down the whole column

Interval thickness follows the kinematic relation between interval velocity and travel time. Both are uncertain fields with their own trend and residual, typically ten to two hundred metres per second on the velocity and two to twenty milliseconds on the interpreted travel time. Accumulate the intervals and a deep horizon carries the error of everything above it.

TIME TO DEPTH · ACCUMULATED INTERVALS

$$\Delta z_i(\mathbf{x}) = v_i(\mathbf{x})\,\Delta t_i(\mathbf{x})$$

$$Z_L(\mathbf{x}) = Z_0(\mathbf{x}) + \sum_{i=1}^{L} \Delta z_i(\mathbf{x})$$

$v_i$ interval velocity, trend plus residual

$\Delta t_i$ interpreted travel time with its own SD

$Z_0$ reference surface, MSL by default

$Z_L$ surface L, delivered with its SD map

Velocity and travel time uncertainty are propagated jointly through the sum, so each surface gets a realistic SD map. Erosion, onlap and pinch out are resolved afterwards, in a post processing step that can be switched off.

# One surface, several ways to reach it

A surface is rarely defined by a single route through the column. It can be reached by depth converting its own reflector, or by accumulating isochores down from a shallower reflector, or up from a deeper one. COHIBA builds every consistent route and combines them, weighting each by the precision it carries.

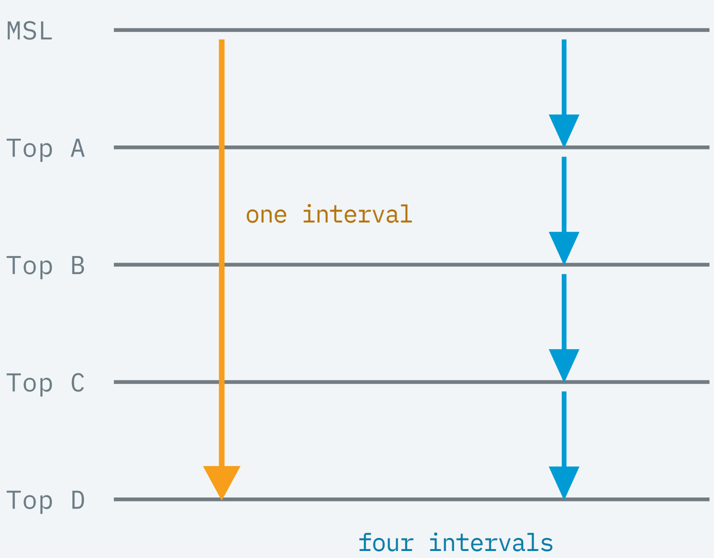


- Depth conversion: travel time times interval velocity
- Isochore stacking: thickness maps from above or below
- Any mix of the two, wherever the data allow it

| ROUTE TO ONE SURFACE | INTERVALS | WEIGHT | DEPTH SD |
|---|---|---|---|
| Direct depth conversion | 1 | 0.62 | 10.0 m |
| Through a second reflector | 10 | 0.15 | 20.1 m |
| Through a third reflector | 19 | 0.09 | 25.3 m |
| Through a fourth reflector | 33 | 0.14 | 20.7 m |

Model weights and depth standard deviation per route, from a mature field model in the Norsk Regnesentral workshop material. The long route carries more weight than the shorter one above it because its residuals are tighter.

> Weights follow the inverse of the residual variance, $\mathbf{W} = \mathbf{C}^{-1}\mathbf{1} \,/\, \mathbf{1}^{\mathsf{T}}\mathbf{C}^{-1}\mathbf{1}$. Two consequences: the combined prediction is tighter than any single route, and every surface becomes dependent on every other, so a well observation on a deep horizon moves the horizons above it. That dependency is computed, not assumed.

# Same data, three estimators, three answers

The residual field is where the estimator shows itself. Simple kriging holds the prior trend fixed, so an error in the trend stays inside the residual. Universal kriging estimates the trend coefficients from the wells alone and removes that bias, at the cost of chasing noise when wells are few. Bayesian kriging carries a prior with an uncertainty attached and lands between the two.

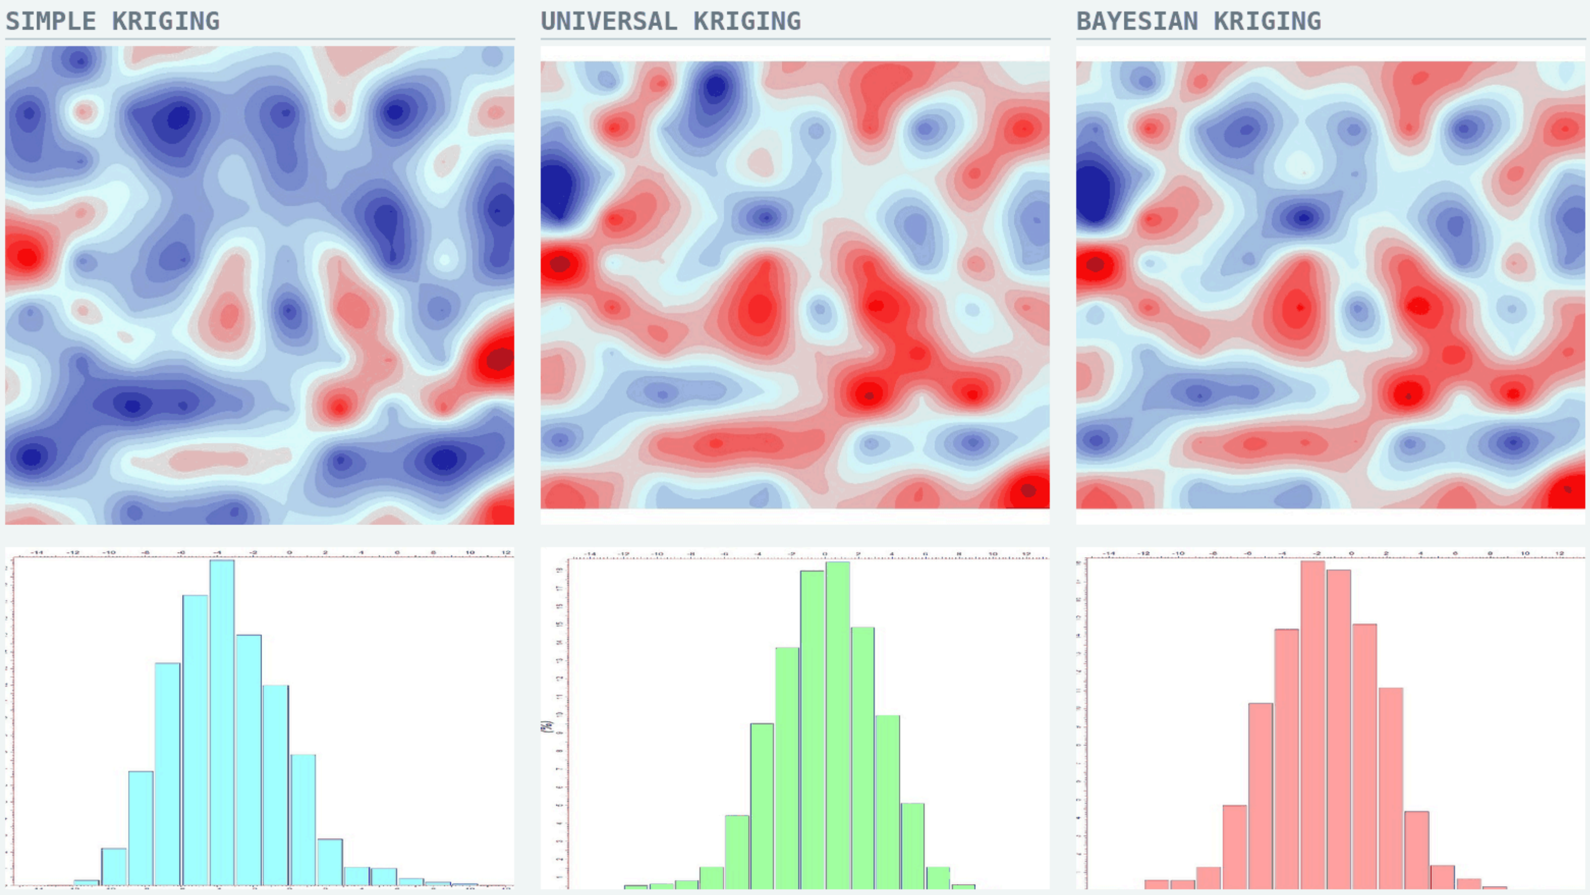


**Fig. 04** · Residual field and residual histogram for one top surface. Same wells, same variogram, three estimators. Norsk Regnesentral COHIBA course dataset.

Histogram off centre: the trend error stayed in the residual

Histogram centred: trend coefficients taken from the wells

Nearly centred, and the prior is still in the model

> Switching between the three is one dropdown in Modeling settings, so a comparison like this is an afternoon of work, not a project.

DEEPSOFT



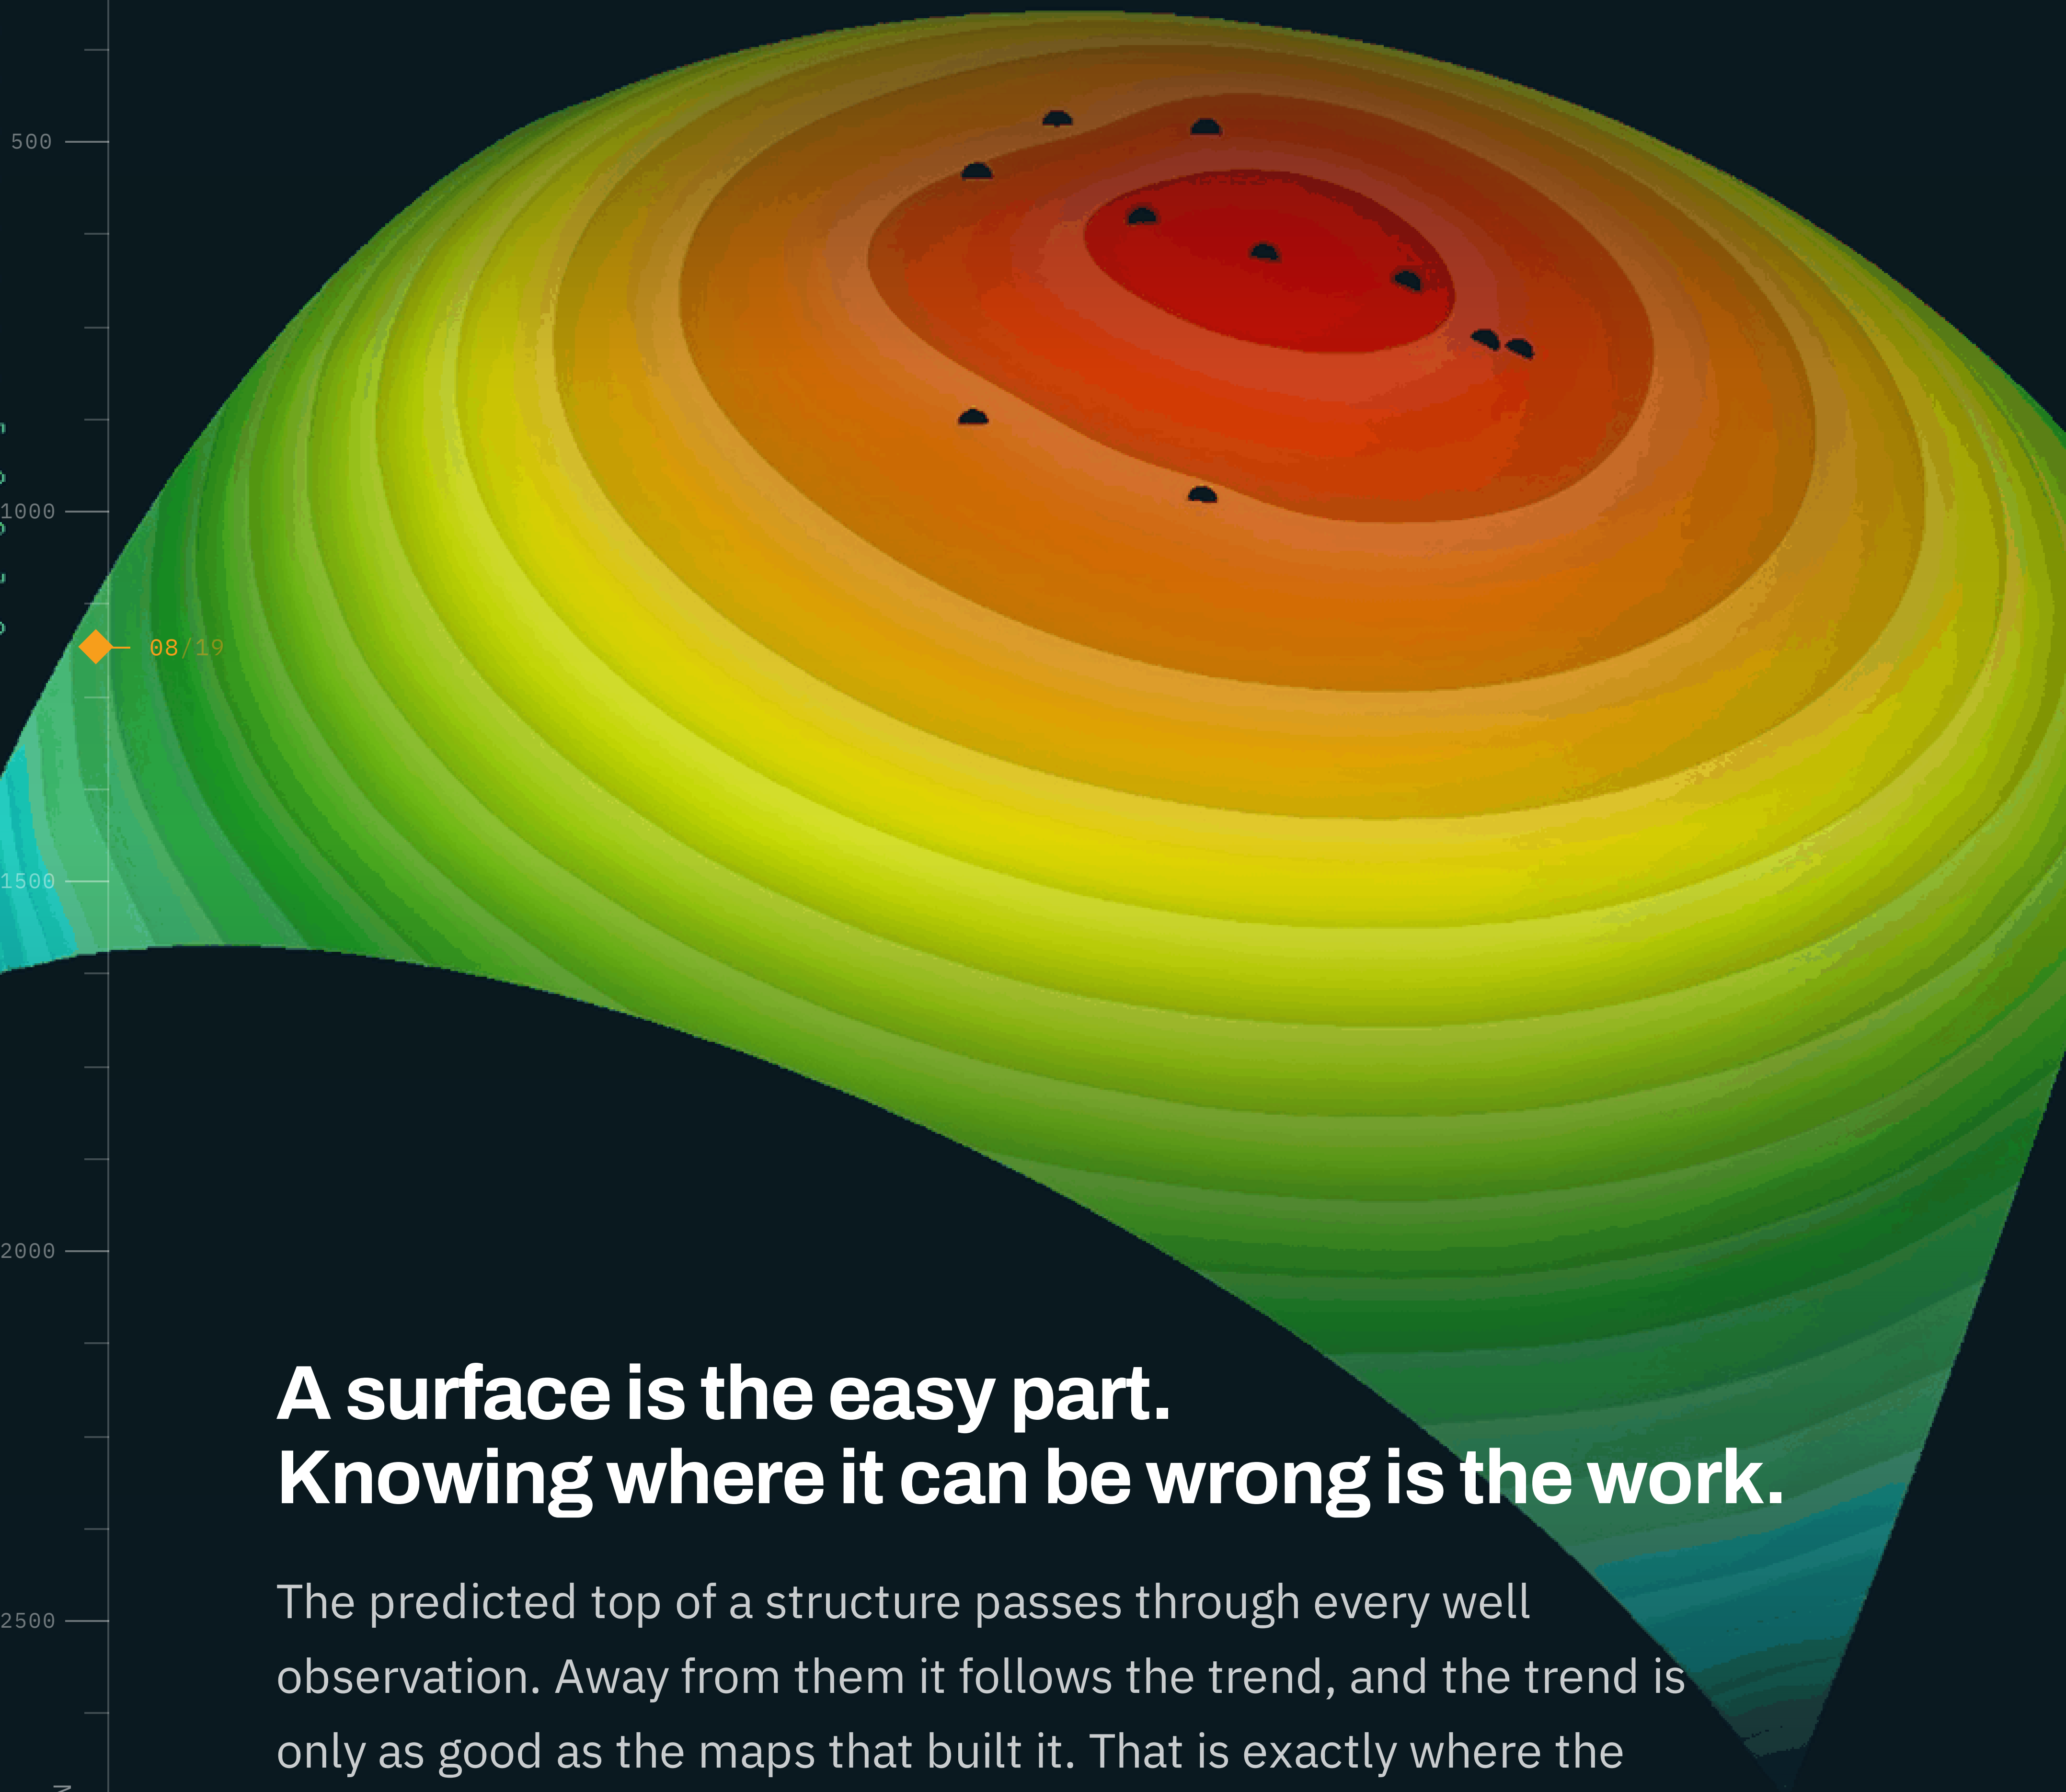

## A surface is the easy part.
## Knowing where it can be wrong is the work.

The predicted top of a structure passes through every well observation. Away from them it follows the trend, and the trend is only as good as the maps that built it. That is exactly where the standard deviation map earns its place.

Fig. 05 · Predicted depth surface with well observations in black. Norsk Regnesentral COHIBA course dataset.

DEEPSOFT

# Four stages, and not one command line

**01 · PETREL PROJECT**

Travel time surfaces, velocity and isochore maps, well tops, zone logs, uncertainty logs, area maps

**02 · COHIBA PETREL**

Project settings, Modeling settings, Wells, Surfaces and Intervals, Volumes, Advanced settings. The plug-in writes the COHIBA model file and exports the data the engine needs

**03 · COHIBA ENGINE**

Trend fitting, kriging prediction, FFT stochastic simulation, zonation control, volumetrics. Progress and warnings streamed to the Monitor

**04 · BACK IN PETREL**

Depth surfaces and realizations, trend and residual SD maps, QC point sets, diagnostic and volume tables, imported automatically into the project tree

> No XML editing, no working directory to manage, no manual reimport. The plug-in owns the whole round trip.

DEEPSOFT

# Wells enter as constraints, not as an afterthought

Well tops, measured depth logs, zone logs, depth uncertainty logs and fault logs are mapped per well or inherited from a global default. Zone logs tell the engine where a surface has to move to restore correct zonation.

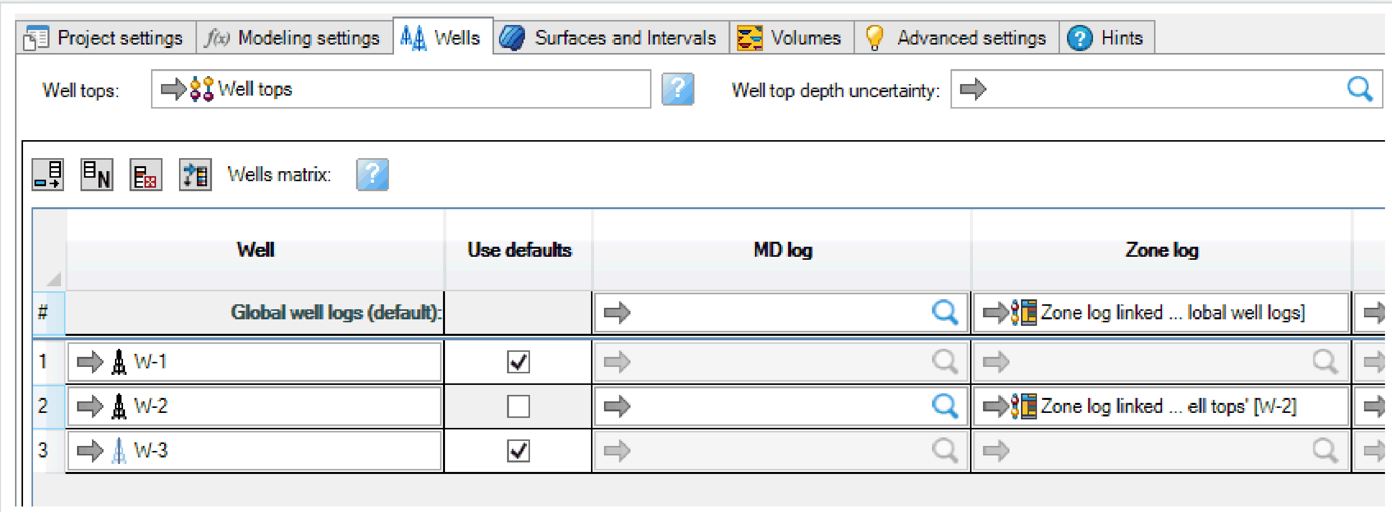


Fig. 06 · Wells tab. Global defaults on top, per well overrides below.

- Well tops, with a per well depth uncertainty log when the pick is soft
- Zone logs along the whole trajectory, not only the entry pick
- Deep directional resistivity, as a constraint away from the borehole
- Fault logs, to keep picks that fall on a fault out of the estimation

> A horizontal well carries its own TVD uncertainty, growing by roughly five metres per thousand metres drilled laterally. The engine treats the path as uncertain rather than exact, and can be allowed to move it within that uncertainty.

# The zone log does not check the model. It moves it.

In the prior model below, the well runs above the zones it is supposed to be inside, and the log agrees with the modelled surfaces along 13.9 per cent of its length. Conditioning on the log itself, and not only on the entry pick, pulls every surface into a geometry that reproduces the whole trajectory.

PRIOR MODEL

CONDITIONED TO THE ZONE LOG

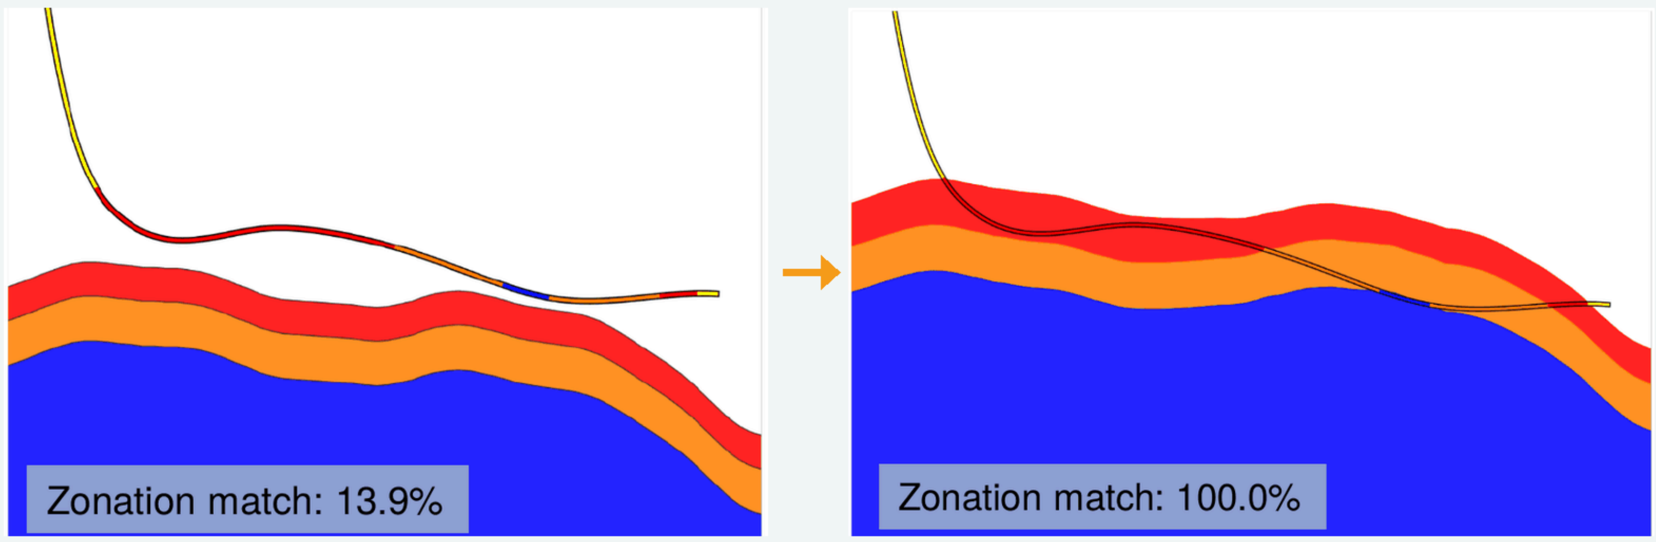


**Fig. 07** · Norsk Regnesentral synthetic real time case study. The coloured trajectory is the zone log, the bands are the modelled zones. Method: Abrahamsen and Benth, Mathematical Geology 33(6), 2001.

- Where a surface has to pass to keep the zonation, the engine places a target point and then treats it as an ordinary well point
- How many target points appear is a setting: on one well, raising the threshold from 0.80 to 0.98 took it from 11 points and 80.6 per cent match to 32 points and 98.3 per cent

> This is what makes a horizontal well useful for structural modelling and not only for production. Every metre drilled inside a zone becomes an observation.

DEEPSOFT

# The column, its trends and its variograms, in one matrix

Surfaces carry reflector status, well tops, travel time map and time uncertainty. Intervals carry their type, velocity or thickness, the trend maps with coefficient mean and SD, and the residual variogram.

| | Interval | Type | + | Trend value(s) / Coeff. mean / Coeff. SD | | | | | Residual variogram function | Residual variogram range |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | MSL_TopA | Velocity | + | Surface | vel_MSL_TopA | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 2 | TopA_TopB | Velocity | + | Surface | vel_TopA_TopB | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 3 | TopB_TopC | Velocity | + | Surface | vel_TopB_TopC | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 4 | TopC_TopD | Thickness | + | Surface | C | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 5 | TopD_TopE | Thickness | + | Surface | D | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 6 | TopE_TopF | Thickness | + | Surface | E | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 7 | TopF_TopG | Thickness | + | Surface | F | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |
| 8 | TopG_TopH | Thickness | + | Surface | G | 1.000 | 0.050 | - | Generalized exponential | 2000.000 |

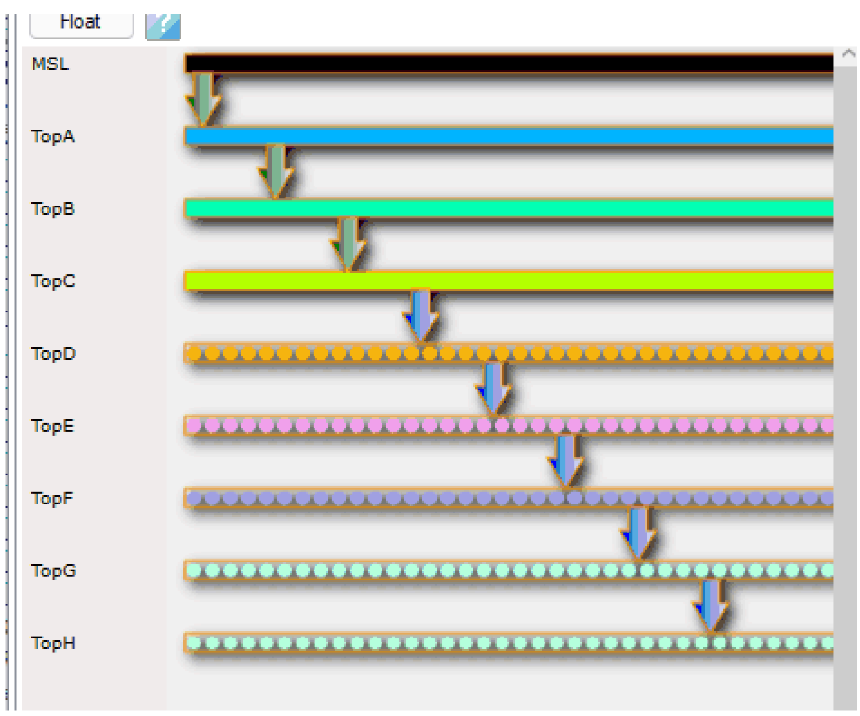


Fig. 08 · Surfaces and Intervals. Velocity driven intervals on top, thickness driven below, and the column panel the plug-in draws from the same definition.

> Solid lines are reflectors, dotted lines are non reflecting surfaces, arrows show which interval drives which surface.

# Prediction or simulation, decided in one field

Set mode to **prediction** for a deterministic surface with its SD map, or to **simulation** with a seed and a number of realizations. Each realization draws the trend coefficients from their posterior distribution, draws the residual fields with a fast Fourier transform algorithm, then conditions the result to the well data by simple kriging.

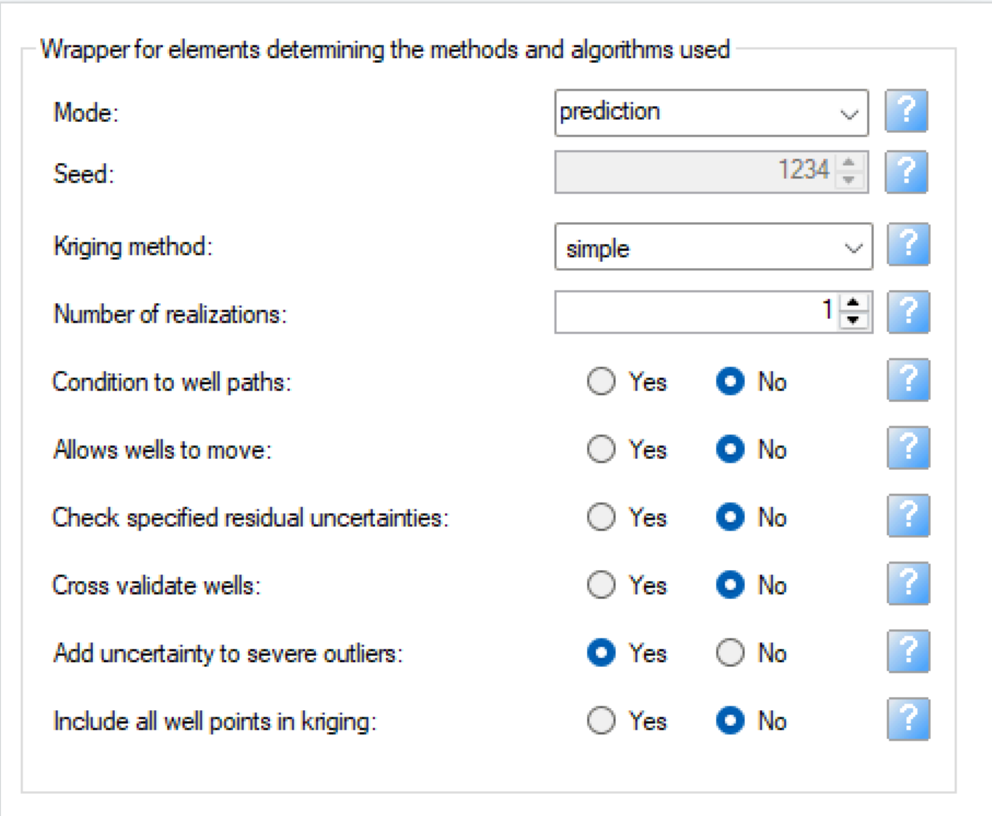


Fig. 09 · Modeling settings. Mode, seed, kriging method, realizations.

Erode and onlap, average of crossing surfaces

Cross validation of wells, severe outlier handling

Small adjustment at zonation points

> The FFT step reproduces the specified variogram almost exactly, which sequential Gaussian simulation and turning bands do not (Abrahamsen et al., ECMOR XVI, 2018). Estimators are one dropdown apart, so changing one is an experiment, not a project.

DEEPSOFT

# The engine log, in front of the geoscientist

Start, stop and status sit next to the raw COHIBA output. Section timings, warnings and the suggested tasks section are read directly in the Monitor, so a run that ended with a zonation conflict or an extreme trend outlier says so before anyone maps the result.

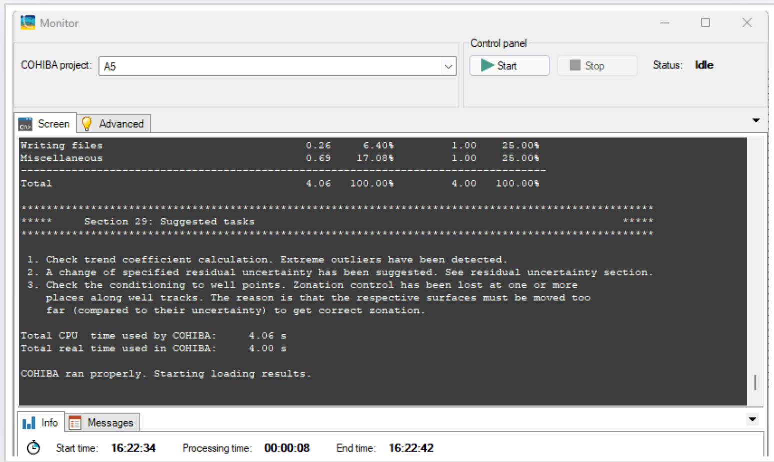


Fig. 10 · Monitor. Engine output, timings, warnings and suggested tasks.

- Trend coefficient diagnostics and outlier detection
- Residual uncertainty suggestions when the specified SD does not match the data
- Zonation control reports, per well and per surface
- Outliers graded by probability: flagged below five per cent, given an added measurement error below one per cent, removed below two in a thousand

> In the reference eight surface model shipped with the manual, a full run takes about four seconds of CPU. Results load automatically when the engine exits cleanly.

# Everything returns as native Petrel objects

- Depth surfaces, one predicted surface or one per realization
- Trend and residual standard deviation maps attached to each surface
- Fitted trend surfaces for every horizon
- QC point sets: well points, zonation points and surface target points
- Diagnostic tables: trend coefficients, well mismatches, zonation status and volumes

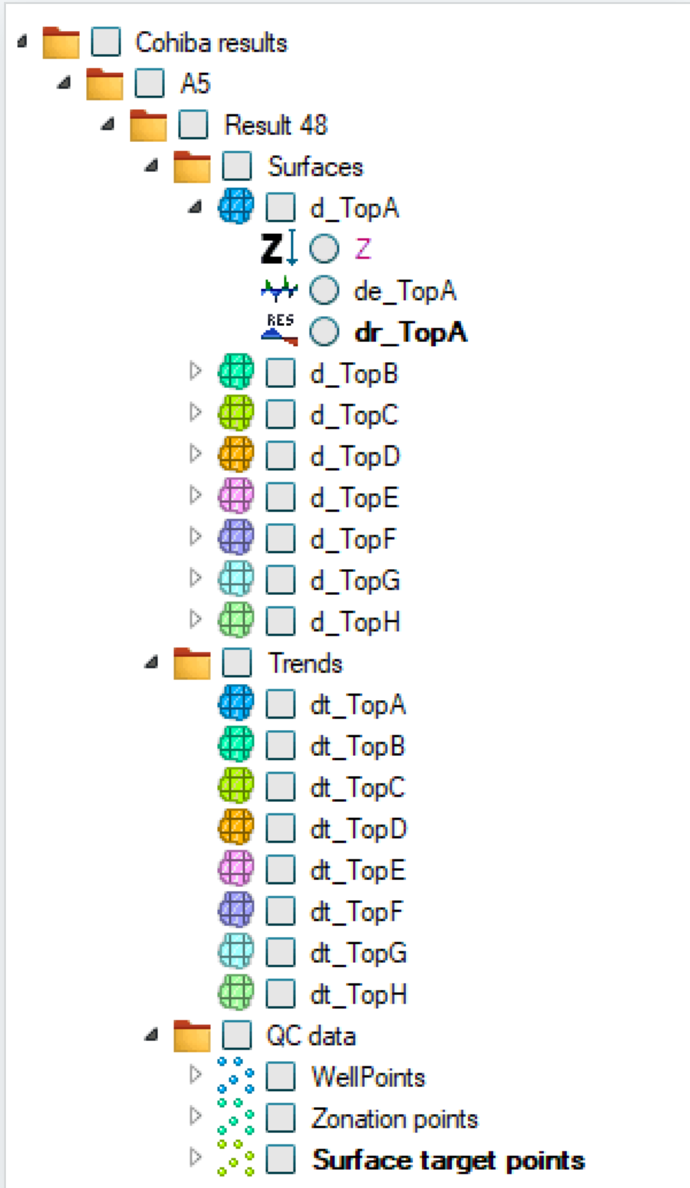


**Fig. 11** · Results tree after a run. Surfaces, trends and QC data.



> Surface target points are where the constrained algorithm decided a surface has to pass to keep the zone log correct. Nothing is dumped to disk for the user to reimport: the tree is built inside the project, grouped by project and by run, so two runs can be compared side by side.

# Gross rock volume, same uncertainty model

Define a reservoir between a top and a base surface, add contacts and an optional area map, and the engine reports the volume for every realization. Trapped volume only, isolated volume filtering and connectivity to a point are switches in the tab.

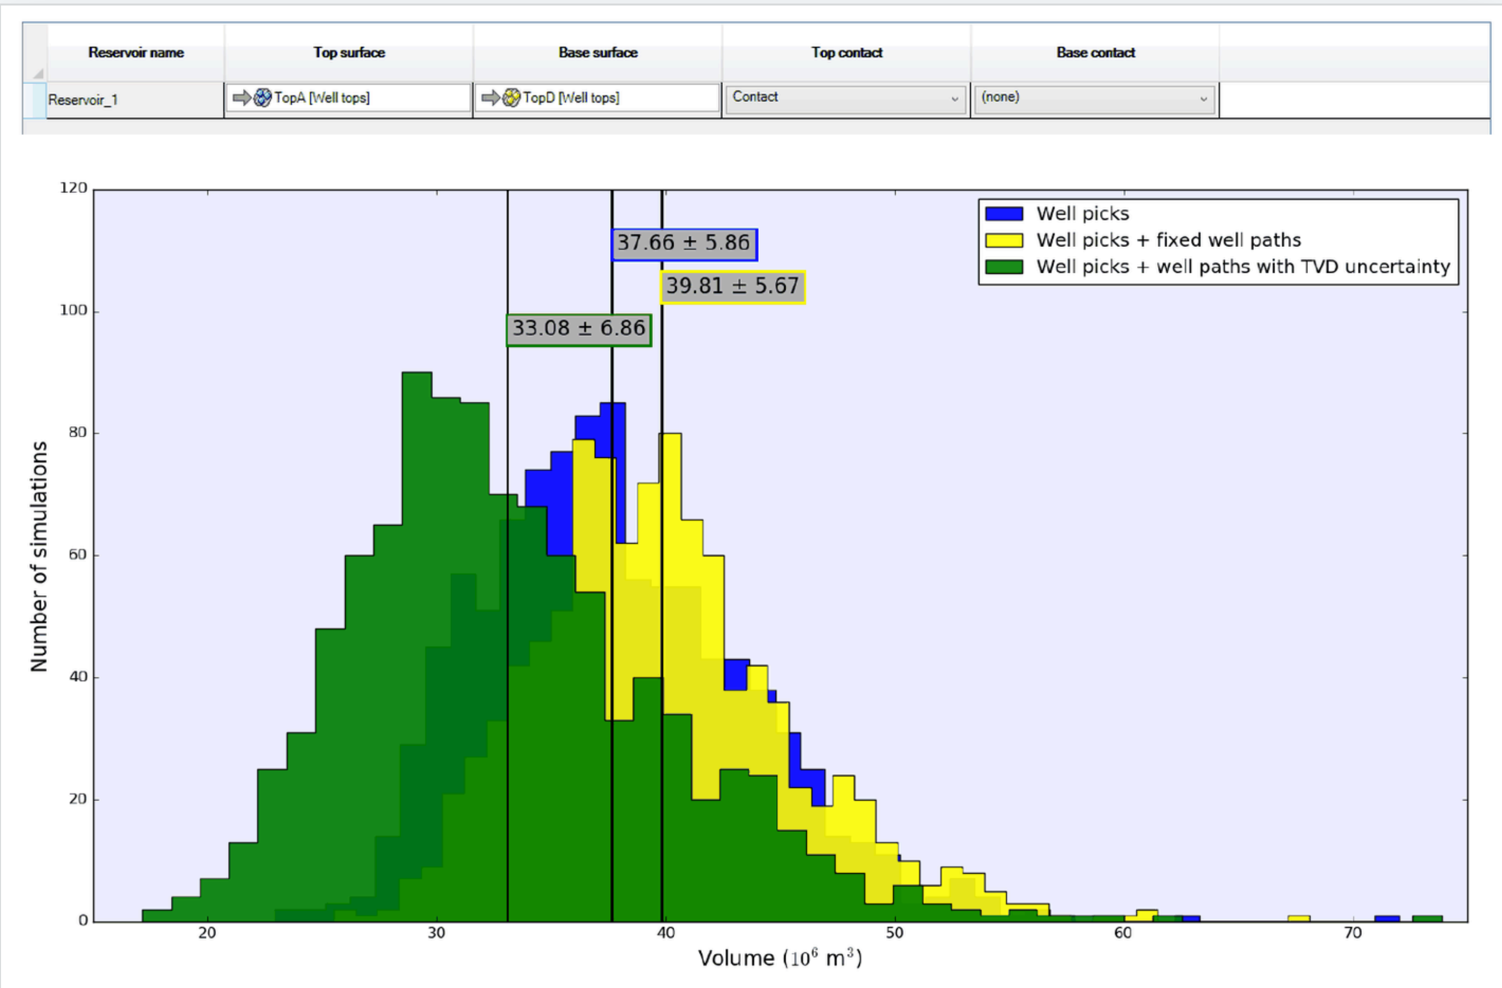


**Fig. 12** · Reservoir definition in the Volumes tab, and gross rock volume over a thousand realizations under three conditioning choices. Norsk Regnesentral synthetic real time case study.

> Volume above a contact is a non linear function of the surface, so a predicted surface cannot give you its distribution. Only an ensemble can. And what you condition on moves that distribution: treat the well paths as exact and the mean gross rock volume is 39.8 million cubic metres, carry their TVD uncertainty and it is 33.1, about seventeen per cent lower and with a wider spread.

# One workstep, and COHIBA is part of the pipeline

The COHIBA project executor appears under Processes, Ocean plug-ins, in the Petrel Workflow Editor. Drop it into a workflow, point it at a project, and it runs headless, logging to the workflow console.

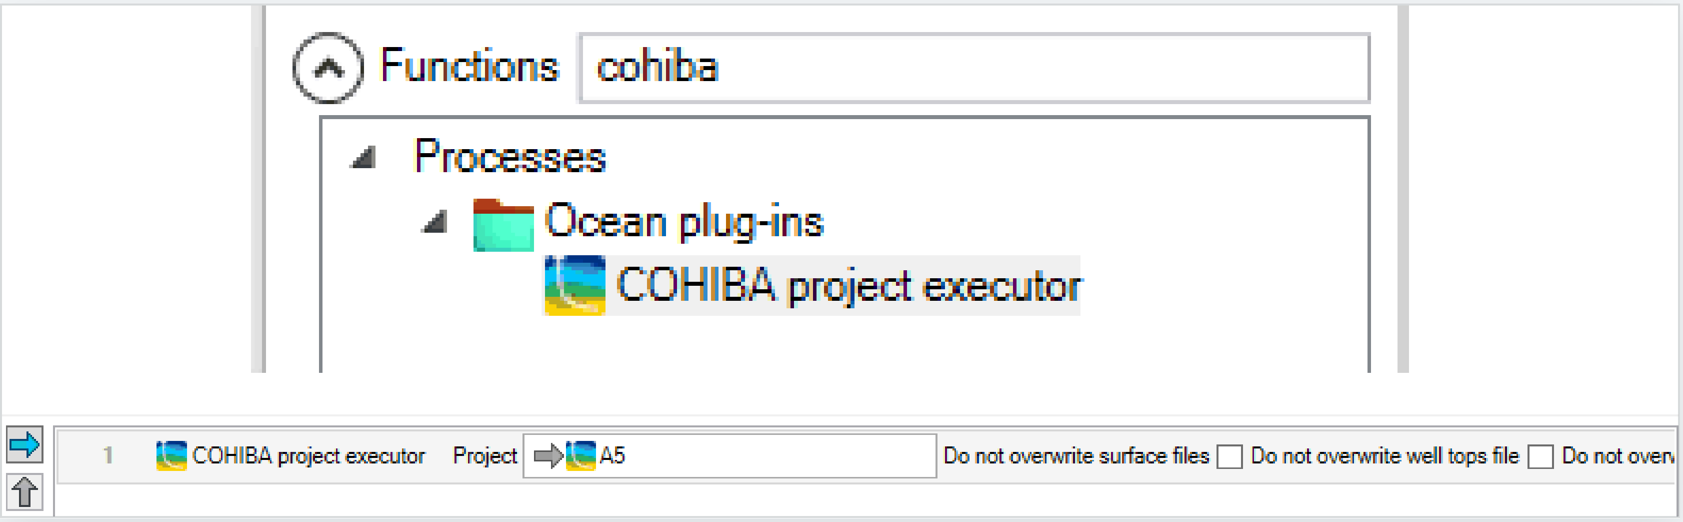


Fig. 13 · COHIBA project executor inside the Petrel Workflow Editor.

- Batch execution across several COHIBA projects, one after another, with no dialogs in the way
- Chaining with other Petrel processes: seismic QC upstream, depth conversion post processing downstream
- Parameterized uncertainty studies driven by Petrel workflow variables, combined with Run in background



> This is the path to large ensembles: hundreds of parameterized runs without a single dialog opening.

# What the asset team actually gets

**0** XML

MODEL FILES WRITTEN OR EDITED BY HAND

**100**%

ZONATION MATCH AFTER CONDITIONING TO ZONE LOGS, AGAINST 13.9 % ON THE UNCONDITIONED MODEL

**3** operators

NORTH SEA OPERATORS RUNNING IT IN PRODUCTION TODAY

Every predicted surface arrives with its standard deviation map. Every volume arrives with the ensemble behind it. Well planning stops arguing about one surface and starts arguing about a range, which is the argument worth having.

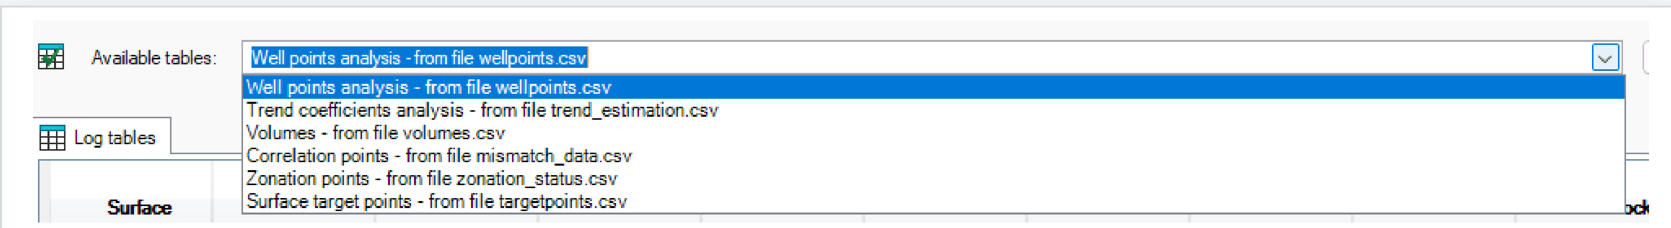


Fig. 14 · Diagnostic tables loaded back into Petrel: trend coefficients, well mismatches and volumes.

> The zonation figure is from the Norsk Regnesentral synthetic real time case study. Numbers for a specific asset come out of running it on your own data, which is what an evaluation licence is for.

DEEPSOFT

# It runs in your Petrel, on your data, in your hands

COHIBA Petrel is licensed under a commercial lease, with technical support direct from the DeepSoft engineering team. Ask for an evaluation and build the first project yourself: your column, your wells, your time maps.

SITE

deepsoft.com.br

CONTACT

contact@deepsoft.com.br

LICENSING

Plug-in lease with engineering support

COHIBA Petrel

Surface Modeling and Depth Conversion

DEEPSOFT

COHIBA is a product of Norsk Regnesentral and is licensed separately by NR. COHIBA Petrel is the Petrel integration and needs a COHIBA licence to run, so the two are not sold as one bundle. Most teams that come to us already hold one. If yours does not, we introduce you to NR and set the plug-in up around it.