WHITE PAPER

AI@Scale in Upstream Oil & Gas Production Optimization

Exploring the use of AI and data-driven optimization to unlock value, with help from BCG and Solution Seeker's Production Compass offering

February 2025

By Håvard Holmås, Vidar Gunnerud, Michael Buffet, Odd Arne Sjåtil, Félix Campos & Meghan Janzen

Introduction

To enable reliable supply while minimizing costly new investments, it's becoming increasingly important for upstream oil and gas companies to ensure that they maximize value from existing assets.

Although these companies have been at the forefront of technology development and implementation for decades, recent experience strongly suggests that harvesting value from AI at scale requires a shift in mindset that we are yet to see in the industry.

Furthermore, AI and GenAI in upstream operations have so far mostly been applied to improve process flow or to pockets of asset integrity and maintenance. Legacy technology still dominates when it comes to modeling of the physical flow of hydrocarbons through the production system. **However, the value unlocked from improved insights in the latter domain of production optimization is potentially orders of magnitude higher than any other application.**

This article explores the value potential of a typical 5-10% production uplift, and looks closely at the success criteria for AI and data-driven production optimization. We also introduce the unique offering from BCG and our technology partner Solution Seeker in this space, Production Compass, which has now been road-tested in live projects across the sector and is based on more than 15 years of research and co-development with leading oil and gas companies globally.

Adopting a data-driven approach to O&G production optimization — why and how?

The ongoing energy transition calls for disruptive action within the oil and gas sector. Long-term decline in demand suggests that only the most cost-efficient barrels of oil will remain profitable, emphasizing the need for producers to stay low on the demand-supply curve. The International Energy Agency projects a global consumption figure of approximately 20 million barrels of oil per day by 2050[1] — a fraction of current demand. In addition to economic pressures, reducing emissions and minimizing CO2 footprint is now critical to maintaining the industry's license to operate. As environmental expectations grow, companies know they'll need to prioritize sustainability to align with regulatory and societal demands. Compounding these challenges, the scarcity of skilled talent requires innovative approaches to working practices, and renewed investment in efforts to cultivate an attractive brand that can draw and retain the workforce of the future.

Seen in this context, maximizing value from existing assets is critical as it displays several attractive characteristics when compared with investing in new developments. First, time to value is typically a period of weeks or months as opposed to years (or even decades). Second, OpEx and CapEx requirements are negligible relative to typically multi-billion-dollar investments in green- or brownfield developments. Third, emissions and nature footprint of production optimization activities are typically significantly lower than the lifecycle footprint of a new development – they may even be net positive.

At the same time, production optimization attracts several significant challenges — in fact, they're issues that the industry has been grappling with for decades. These can be grouped into two main categories:

Decision inputs: Optimizing drainage strategy and the flow of hydrocarbons from the reservoir, through wells, pipelines and processing equipment to receiving facilities is inherently a highly complex problem. It requires real time handling of hundreds or thousands of operational data points, optimization parameters, and boundary conditions, often hampered by poor data quality and accuracy. Additionally, the modeling of multiphase flow such as oil, gas, water and particles in pipelines is recognized as one of the most demanding mathematical challenges in the world. Consequently, the derived input parameters for decision making often leave huge uncertainties. As an example, the error in estimates of a fundamental property such as how much oil/gas/water is produced from each well in a field typically is in the order of 10-100%!

Decision making: Our experience working with asset and centralized production organizations around the world points to multiple challenges in the end-to-end decision-making process, with various root causes. Examples include the fact that production optimization has traditionally gained less management attention than CapEx-intensive activities such as new developments or drilling. Second, a day-to-day focus on firefighting and delivering on planned production leaves little room in the organization to address optimization opportunities. Third, silos between disciplines such as production, reservoir and process engineering can hinder holistic, value-based optimization. Finally, we may find a tendency to rely on experience rather than data and structured analysis.

The above factors often lead to overly conservative decision-making — which means oil and gas companies are leaving a significant number of barrels on the table.

The good news is that **these types of challenges are a sweet spot for AI and a data-driven approach**: AI is at its best when dealing with large collections of data, and behaviors and trends that are difficult to predict. It also holds huge potential when it comes to focusing attention on the highest value opportunities, through the use of smart alarms and proactive and automated advice that improves the quality and speed of decision making.

As an example, we can imagine the typical situation of a production engineer who monitors daily performance across dozens or hundreds of wells. He or she strives to identify risks or improvement opportunities by manually identifying patterns in pressure, temperature and other production data. But what if we imagine a situation in which this engineer could access a tool that leverages the full production history, from not only the field in question but also

1. International Energy Agency 2024 World Energy Outlook (Net Zero Emissions Scenario) similar fields around the world? And if this was complemented with AI-enhanced data to identify anomalies and “bad actors” much earlier than possible for the human eye? It could then offer objective and condition-based advice, allowing the engineer to spend more time where he or she can create the most value for the company.

If implemented successfully at scale, BCG’s experience indicates that the **value potential of such solutions would be in the order of a 5-10% production uplift** combined with increased ultimate reservoir recovery and avoidance of bad OpEx and CapEx spend.

But before we dive deeper into concrete value examples, we will introduce BCG and Solution Seeker’s joint offering in this space: Production Compass.

Differences and complementarity of data driven vs legacy modeling approaches

Computational modeling has been used in oil and gas production optimization for many decades with large providers including Schlumberger, Halliburton, Petex, Kongsberg Digital and Aspentech each utilizing comprehensive and broadly adopted software suites, covering areas of the value chain both large and small.

A common starting point among legacy approaches when modeling the flow of hydrocarbons and particles is first-principle physical and mathematical laws, expressed through a set of partial differential equations. A typical example would be the Navier-Stokes equations, supported by one or more equations of state relating the properties of matter to a given set of physical conditions. Due to lack of analytical solutions and limitations in initial and boundary condition descriptions and computational power, these highly complex equation sets must be vastly simplified to make them feasible to solve. Examples of typical assumptions include reducing the equations from 3D to 2D or 1D; introducing closure relations that are empirically estimated and dependent on flow regime; only solving for parts of the fluid/particle mix, and leveraging numerical methods to integrate the equations over discrete steps in space and time.

While the first principles might describe the physics perfectly, the actual equation set that is solved is thus only an approximation, and one that relies heavily on human interpretation and intuition through the numerous assumptions made. Even with such limitations and uncertainty, it is important to point out that these models have served the industry well, and will continue to do so, in terms of guiding decisions in design and operation of oil and gas production systems.

But an AI or data-driven tool represents a complementary and fundamentally different approach.

Instead of starting with first principle laws, powerful machine learning algorithms are trained using all relevant operational and contextual data, only to be guided by physics when the data alone is insufficient. Leveraging similar AI and GenAI techniques, as in more well-known applications such as ChatGPT or DALL-E, these tools then deduce patterns in the data that can be used to estimate properties of the hydrocarbon flow, such as well flow rates in real time.

The main limitation of this approach is the availability and quality of the input data. However, in many real-life examples, it has been demonstrated that a data-driven approach already provides more accurate estimates than state-of-the-art legacy modeling approaches. This is perhaps not surprising given that AI algorithms are far better equipped to perform pattern recognition tasks based on data, even compared with the brilliant human brains that previously did the same job to simplify the equation sets in legacy modeling approaches.

For companies, this should not be a question of either one or the other. Given the value at stake and the fact that the applicability and reliability of either approach will vary with what is being modeled, the two approaches should be viewed as complementary rather than in competition. Data-driven methods should be welcomed as another source of information that can help production optimization, while also keeping in mind that turning new insights into value creation requires a lot more than simply implementing a new tool.

Production Compass — A Unique AI solution for production optimization at scale

While there are numerous AI applications in upstream, focusing on maximizing equipment uptime through condition based or preventive maintenance, very few AI offerings address the core problem of optimizing the flow of hydrocarbons. Production Compass does exactly that and it is a result of an exclusive partnership between BCG and our technology provider Solution Seeker. It combines unique AI production optimization technology (Solution Seeker) with change management and IT integration at scale (BCG).

Exhibit 1 illustrates the high-level data architecture for Production Compass. Most of the unique AI technology is found in the Smart data platform at the top. It is based on more than 15 years of research and co-development with leading oil and gas companies around the world, and the IP is protected by ten families of patents. Over this time, a one-of-a-kind global database with more than 500 billion operational data points has been built and is continuously being updated and used to train machine learning algorithms.

Exhibit 1 | Production compass high level data architecture

The Smart data platform takes in data, prepares and trains models based on it, and supports both standard and asset-specific applications

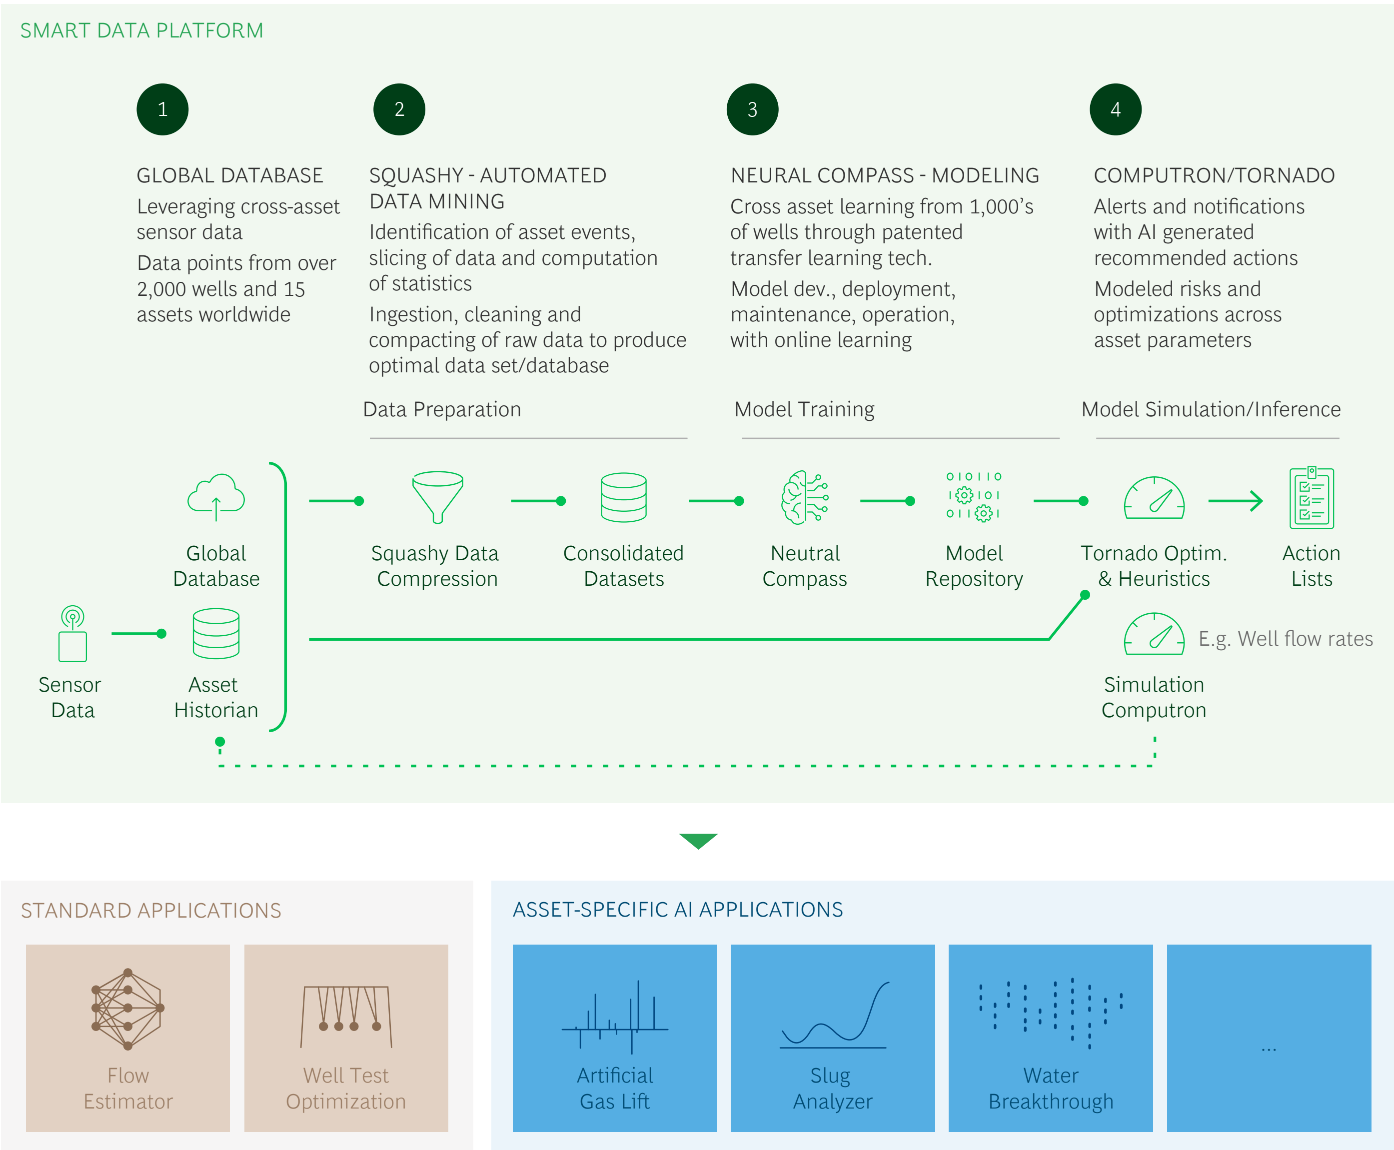

The Squashy automated data mining algorithm ingests, cleans and compacts the raw data (e.g. real time pressure and temperature time series) to produce optimal data sets for consumption by the AI algorithms. The data is then used to train the deep neural network algorithms in Neural Compass, leveraging among other things unique cross asset learning from thousands of wells in the global data base to extract insights relevant for the asset or well that is being modeled. In this way asset uniqueness is taken into account while minimizing the need for asset specific data to train the models. The technique is similar to what is applied in language models such as ChatGPT that are mostly trained on English texts, but despite differences in grammar and language rules, with a minimum of training can generate texts in other languages with surprisingly high quality. Finally, the main outputs from the Smart data platform are either estimated flow properties, such as well flow rates, that can also be fed back to the database or more derived properties such as alerts and notifications with AI-generated recommended actions.

In addition to the Smart data platform, there are two types of applications. First, the standard applications that are relevant across most assets and fields whether they are offshore vs onshore, conventional vs unconventional, large with many wells vs small with few wells, and across deployed IOR/EOR[2] technologies such as artificial lift or water injection. Examples here include the Flow Estimator that provides real-time and well-specific estimates of the produced oil, gas and water, and the Well Test Optimizer that supports both well test selection and execution.

Second, we have asset-specific applications where typically most of the value creation potential is found. These can address any production optimization challenge faced on the asset in question; they will usually require some additional tailoring, since the context of two different assets rarely is the same. More in-depth examples are given in the next section, and for example, could be applications that optimize artificial lift, manage slugging or instabilities in the flow, or provide early warnings and recommended actions when there is risk of water or gas breakthrough for specific wells.

The modular architecture of Production Compass allows for flexible integration with legacy systems through built in APIs. It can either be delivered as a standalone tool with its own user interface (see Exhibit 2 for a screenshot example), or web components or specific data can be integrated into any 3rd party user interface or dashboard. AI-enhanced data can also just be replicated to in-house databases for later consumption by legacy tools (e.g. well flow rates for history matching of reservoir models). In terms of delivery model, the most widespread is Software as a Service (SaaS), but other cloud platform or infrastructure options and on-premises solutions can also be catered for.

Production Compass is a solidly proven technology that has been implemented on 20+ assets and more than 2,000 wells around the world, continuously helping to solve a wide variety of production optimization challenges.

For more information on Production Compass and the underlying technology, please see the references in the “Further reading” section at the end of this article.

Example production optimization AI use cases

Exhibit 3 shows typical levers that can be used to increase production in a driver tree logic. Note that “Equipment uptime” is greyed out here. That is not to say it is not important, but in this article we focus more on optimizing production through addressing “Daily throughput efficiency”, “Maximum production potential” and “Reservoir recovery”. On the right side of the exhibit is a selection of AI use cases in which an offering such as Production Compass can add significant value. Next, we will look at four of these in more detail to illustrate different application areas. See also the factbox with information on the use of Production Compass at a South East Asia-based NOC.

First, let us consider the previously mentioned estimation of well flow rates. This functionality alone does not lead directly to value creation, but well flow rates are the single most important input parameter to most production optimization decisions. Exhibit 4 displays results from a study across five mature upstream oil and gas companies comparing well flow rate estimates

2. Improved/Enhanced Oil Recovery

Exhibit 2 | Production compass example user interface

The system can be delivered as a standalone tool with its own user interface, as show here, or as web components or integrated into any third party user interface or dashboard

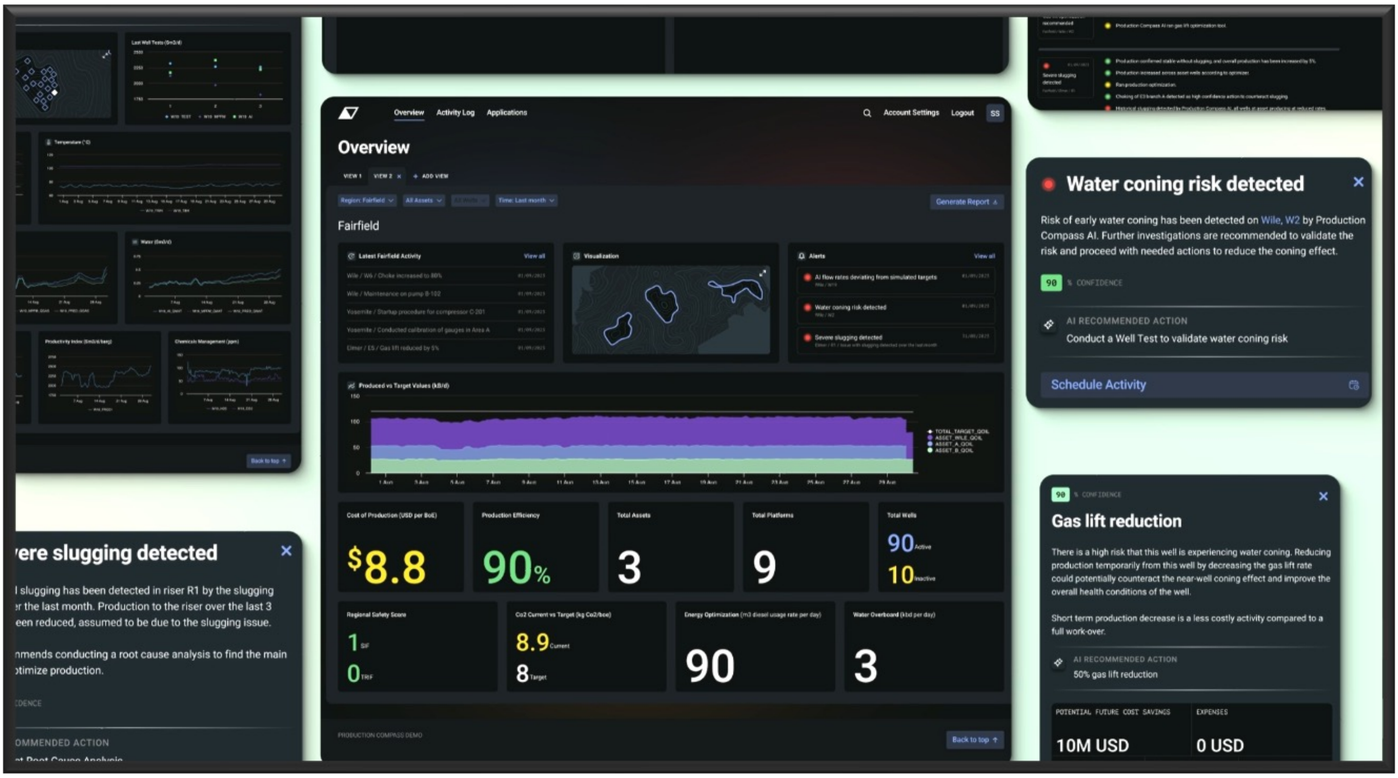

with results from well tests. In the graph on the left, we can see that without Production Compass, while using typical state-of-the-art multiphase flow meters and simulators, fewer than 40% of the estimates experienced an error of less than 20%. Many estimates also showed errors in the range of 50 to 100%. However, in the graph on the right we can see that with Production Compass, we get to more than 80% of the estimates experiencing an error of less than 5%. Such an improvement in accuracy of well flow rate estimate can enable significant value creation across all optimization levers on the left side of Exhibit 3.

In a second example, we can consider the use case of drawdown optimization in connection with sand production. We found an oil producing asset in the Gulf of Mexico struggling with sand limiting production from three high volume subsea wells. Due to risk of damage to equipment and clogging of the wells and uncertainty in the contribution from each well, they were all conservatively choked back with an estimated production loss of 3 Mbbls/y[3]. By combining more accurate production allocation to the wells with granular correlations of sand contributions per well and identification of sand burst drivers, Production Compass could advise on more aggressive and de-averaged choke openings, remaining within the favored risk parameters. The resulting value potential was a 10% production uplift at the same overall sand production.

As our third example, an asset in the North Sea had struggled historically with slugging and flow instabilities in one of its risers. Conventional dynamic multiphase flow simulators such as OLGA and LedaFlow had been unable to recreate the slugging behavior. Production was thus choked back to avoid separator flooding and flow induced vibrations harmful to equipment. But just three months after Production Compass was deployed, it was able to pin down the root causes of the slugging by combining improved well flow rate estimates with systematic analysis of historical slugging data. The "worst actor" in this case was a branch of a multi-lateral well that produced excess water. By shutting this branch in and further opening the choke from other wells, the asset increased production by 5% with an annual cash flow impact of $120M.

3. Thousand barrels per year

Fourth, we can consider the case of optimizing ultimate recovery from a field. In this case the reservoir model is critical for informing overall drainage strategy as well as the optimal timing, location and spacing for any infill drilling campaigns. Maintaining this model through history matching is thus a core task for the subsurface discipline, and one of the most important input parameters for this calibration is the production allocation and fluid ratios per well. Needless to say, the error rates in well flow rate estimates discussed above quickly deteriorates the reservoir model when used in history matching. With a tool such as Production Compass, users can rapidly recalculate the full well production allocation history for a field with the improved accuracy it can bring to the reservoir model. It is difficult to put an exact dollar value on such a contribution, but it is well recognized among oil and gas professionals that having a reliable reservoir model is key to many production optimization decisions.

Exhibit 3 | Production optmization AI use cases

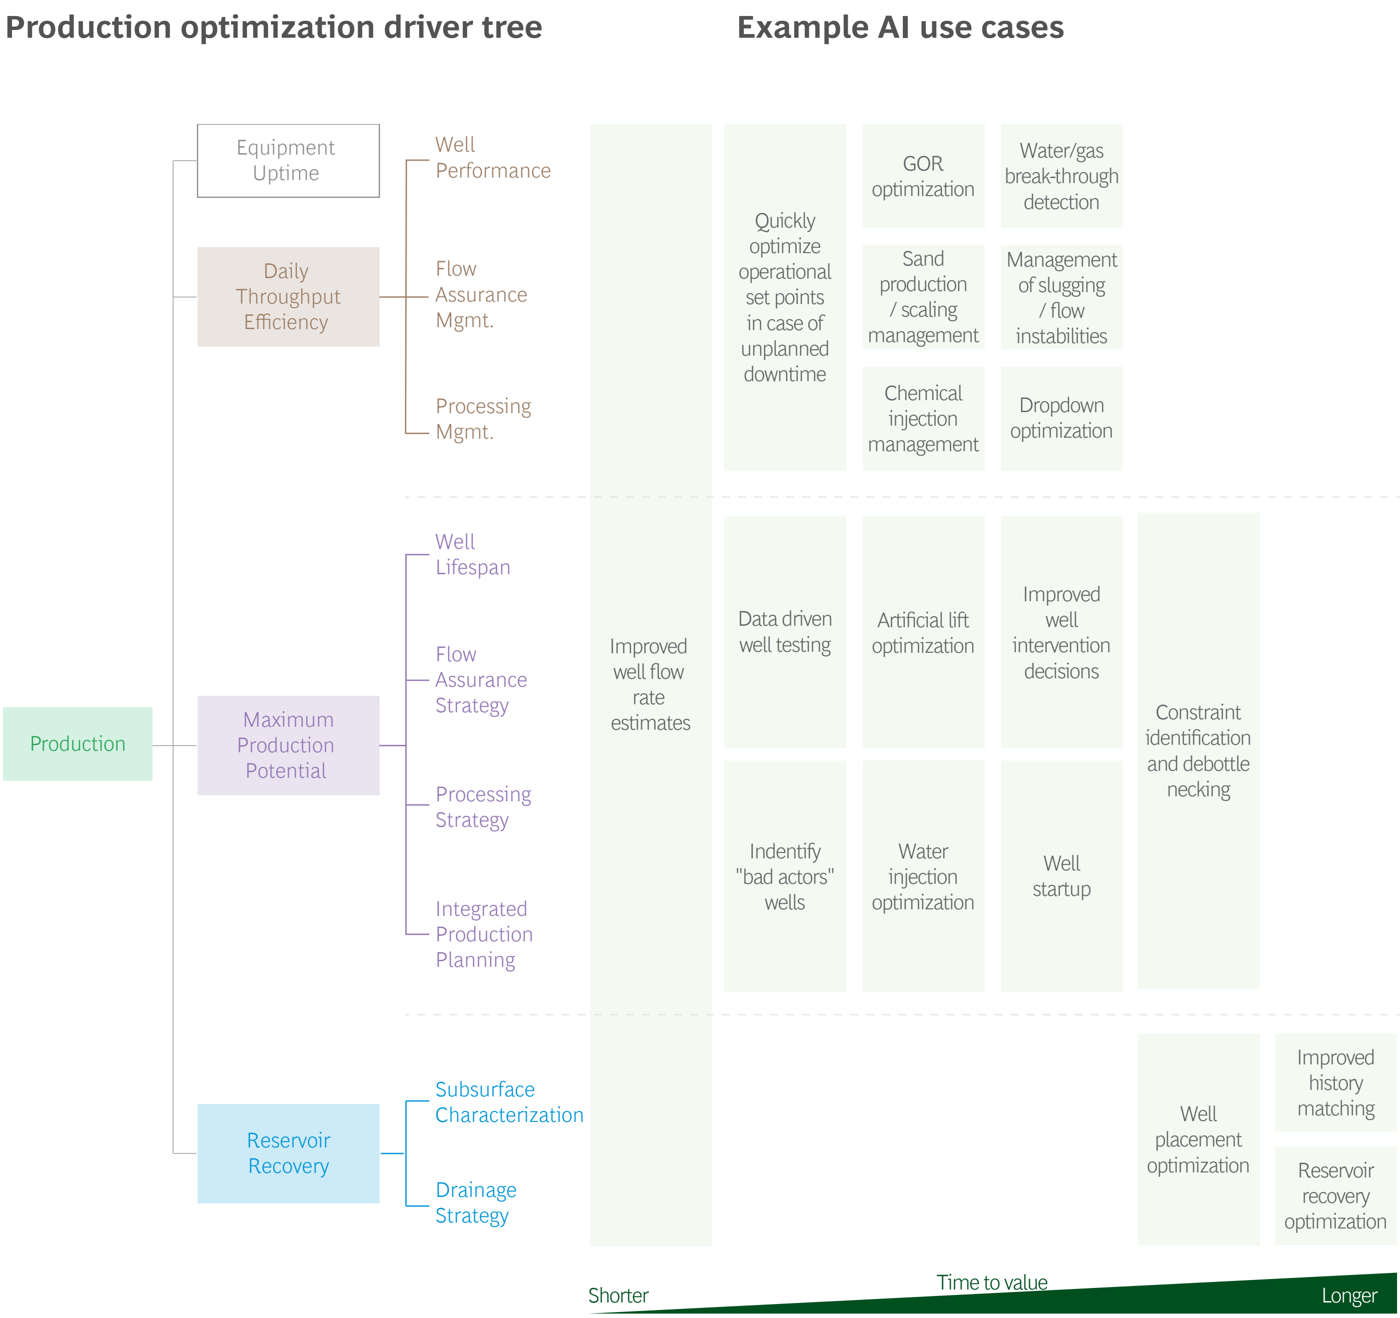

Exhibit 4 | Comparison of well flow rate estimates with results from well tests

Results from a study across five mature upstream oil and gas companies comparing well flow rate estimates with results from well tests

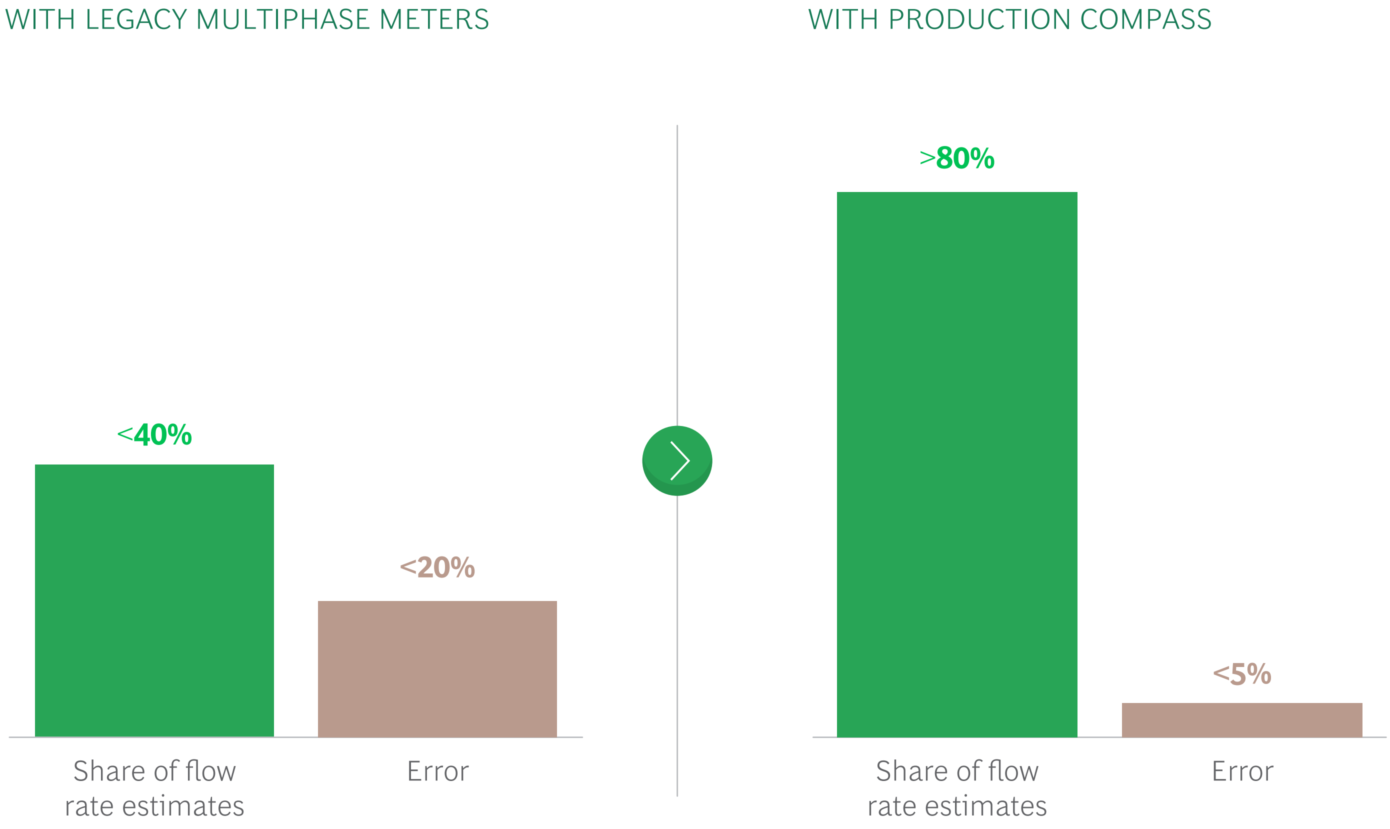

Production Compass at SEA NOC

An NOC in South East Asia has been running Production Compass on four assets since 2019. Standard applications such as the Flow Estimator and Well Test Optimizer are in use, proving their value at more than 1,500 wells (mostly in shallow water with dry wellheads, and with complex field settings with compartmentalization and completions across multiple hydrocarbon zones). In addition, asset-specific applications have been deployed to address specific production optimization challenges.

One particular challenge for the company was related to water breakthrough negatively impacting well potential and overall reservoir recovery as well as leading to costly well intervention operations. With the high number of wells, plus a lack of resolution and accuracy from sensors, it was difficult to manually detect soft signals in the data and therefore to implement mitigating actions in time. In many instances, early warnings of the water breakthrough in the sensor data were not even visible to the human eye in hindsight.

However, by working closely together with the company subject matter experts and leveraging the improved well flow rate estimates as well as other available historic sensor data, an AI-based water breakthrough risk indicator was established. Based on pattern recognition in the data leading up to prior events, the indicator is now able to provide engineers with early warning of wells at high risk of water breakthrough. At one of these assets, the latest version of the AI algorithm detected 93% of water breakthroughs. It also sends notifications to the relevant engineers so that they can take appropriate actions in time. These notifications enable a condition-based workflow, reducing the time spent analyzing water breakthrough data by 97.5%.

What it takes — and where to start

There are numerous examples of the oil and gas industry successfully being at the frontier of technology development and adoption. Examples include the engineering achievements required to realize large offshore and subsea developments, and seismic interpretation that has pushed the boundaries of supercomputing and big data analytics for decades.

However, when it comes to the digital revolution over the past decade, there is a clear perception in the industry that investments have far exceeded the returns.

Root causes are multiple; some of them may be linked to large legacy oil and gas organizations being less well-suited to agility, or to cross-disciplinary collaboration and efforts to fundamentally rethink their ways of working. But one thing is for sure: **success requires a lot more than simply the implementation of tools**.

In fact, at BCG we use a rule of thumb that applies to all digitally-enabled change across industries: the rule of 10, 20 and 70. It states that to ensure success, roughly 10% of your effort must go into developing front end applications and algorithms, a further 20% is related to the underlying data and IT platform — but 70% of all effort must go into changing the way people work and how decisions are made.

Driving this type of change is at the heart of what we do at BCG, and our Production Compass offering is much more than just a tool. In essence, our methodology is based on four steps:

1. Identify the biggest untapped production optimization value pools for your company. Set ambitious, yet realistic, top-down targets. Ensure that business leaders get on board with the vision and are committed to the change.
2. Examine relevant production optimization workflows from end to end and identify frictions keeping them from functioning optimally today. Such frictions could both be related to data, tools and insights as well as adjustments to how decisions are made, information flow and collaboration, team constellations and other operating model elements.
3. Rethink workflows, leveraging AI solutions, and build an action and value realization plan. Include milestones both for leading KPIs (e.g. major activities completed, production efficiency improved, maximum production potential increased) and value realization KPIs (e.g. production increased, cost reduced).
4. Execute, monitor and course-correct. Implement solutions and changes in workflow, communicate changes, train workers and relentlessly track progress and make adjustments as needed to drive all the way to value delivered.

In terms of where to start, this naturally depends on the nature of your asset portfolio and the specific production optimization challenges your company may be facing. However, our advice is often to start with a single asset or a small cluster as a pilot, and apply the methodology outlined above in close collaboration with the asset manager and their team. Decide on one or a few use cases to start with and then confidently scale up, in terms of increasing functionality, new use cases and rolling out across the asset portfolio and geographies.

Given the industry context, combined with the ultra-rapid development seen in AI and GenAI, there is no doubt that this technology will disrupt production optimization as well as other parts of the upstream value chain in the years to come. Companies that want to stay ahead of the curve will need to act boldly and embrace this opportunity at scale.

Further reading

Video introducing Production Compass value proposition.

Video with teaser version of sanitized Production Compass demo.

Video elaborating how a Production Compass Virtual Flow Meter is set up, maintained and delivered as an AI service.

BCG Publication: *AI Adoption in Energy Should Focus on Agility, Not Algorithms.* This publication highlights that achieving a lasting competitive advantage with AI requires companies to ensure quality data, promote a culture of innovation, and quickly upskill workers.

White paper: *A deep latent variable model for semi-supervised multi-unit soft sensing in industrial processes.* This paper elaborates on a modeling strategy that is relevant for quantifying confidence in sensors.

Scientific paper: *Multi-unit soft sensing permits few-shot learning.* This paper shows that transfer learning from data between wells enables good predictions also for wells with very little data.

Scientific paper: *Multi-task learning for virtual flow metering.* This paper addresses the same method as item 3 above (transfer learning between well sensor data) but applies it to virtual flow metering.


bcg.com