Real-Time Closed-Loop Optimization Delivers 43% Titer Improvement Over Historical Best

Authors:
Yossi Quint [1]*, Hunter Worssam [1], Christopher Horner[2]*, Mike Khor[2]

Affiliations:
1: Ark, 2: Culture Biosciences, *co-first authors

Highlights

By combining Ark's adaptive digital twin models with Culture Biosciences' API-native Stratyx™ bioreactor, the two teams delivered a 43% titer improvement over the historical best run showing what becomes possible when a model that adapts in real time meets a bioreactor built to respond to it.

The challenge: a fixed recipe cannot adapt to what it does not yet know

The standard approach to process development locks in a process recipe before the run begins. A recipe is designed, the bioreactor run starts, and the recipe is executed as planned. Two problems follow that no amount of upfront process development can fully solve.

First: biological variability is unavoidable. Media lot variation, seed culture state, inoculation density, and metabolic drift mean every run deviates from expectation. When the culture deviates from expectation, the static recipe has no mechanism to respond.

Second: the model becomes outdated the moment the run starts, which especially matters in R&D when you are deliberately exploring new parts of the design space where historical information is weakest. For example, a DOE-based model trained only on historical data may assume glucose consumption scales linearly with cell density, but if the culture shifts to a lactate-consuming phenotype mid-run, every subsequent feed decision is optimizing against a metabolic reality that no longer exists.

In both cases, there is no feedback loop. Any learning waits for the next experiment.

The solution: a digital twin that learns and adapts during the run

Ark and Culture Biosciences ran four conditions: one control and three digital twin conditions. The initial model trained on 33 historical runs from Culture Biosciences' lab. Each of the three digital twin conditions was seeded from a different optimization starting point identified by Ark's model prior to the run. The full deployment, from kickoff to inoculation, took three weeks.

During the run, online data was streamed continuously from Culture Biosciences' Stratyx bioreactor to Ark's platform, and offline samples were uploaded daily. At each update cycle:

- The model was retrained on the new data, updating its parameter estimates.
- A new optimization was run against the updated model to find the best possible go-forward recipe.
- Updated recipe transmitted back to the bioreactor.

What would otherwise be a one-time pre-run prediction became a continuously updating strategy that compounded across the duration of the run.

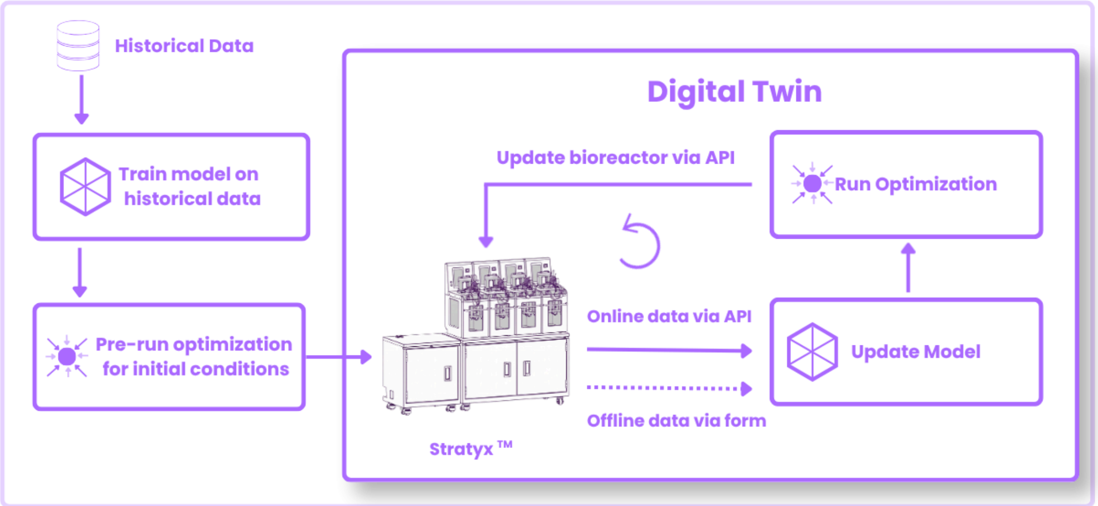

Early in the run, the culture behaved close to historical expectations and the digital twin held its original recommendations. Then, around day six, the culture entered territory the model had not seen before as expected when pushing into a new part of the design space. The digital twin responded: by day eight, it was prescribing meaningfully more feed (up to 82% difference) than the original recipe. A static recipe has no such mechanism.

The result: the top-performing condition (DT 3) produced the best titer outcome in the program's history.

Feed strategy: digital twin vs. static recipe for top performing condition

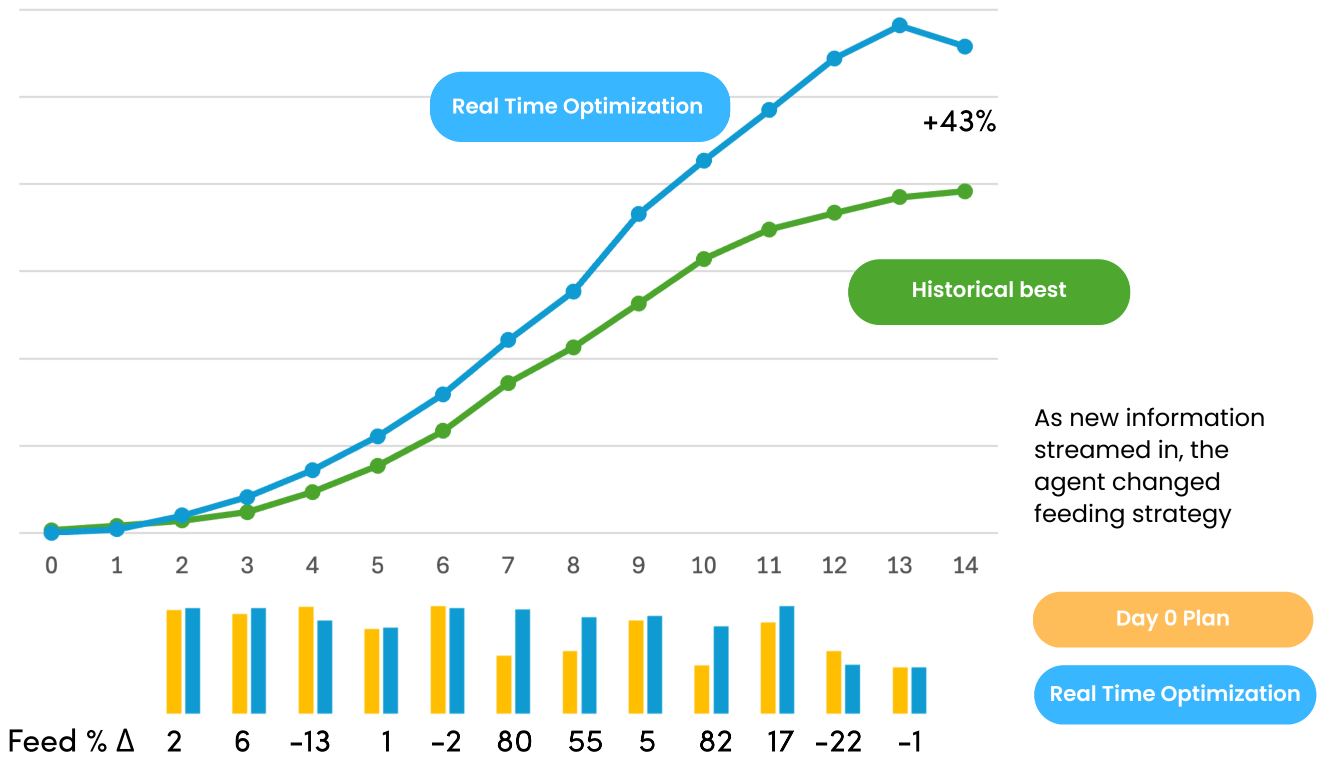

Run	vs. Best
DT 3 ★	43%
DT 2	17%
Control	21%†
DT 1	—

† Control improvement can be attributed to a light-shielding process correction applied to all conditions, independent of the digital twin.

How the digital twin worked: real-time modeling and API-native infrastructure

Ark's digital twin platform was integrated directly with Culture Biosciences' Stratyx bioreactor system via Application Programing Interface (API). Process data flowed from Culture Biosciences' platform into Ark's models in near-real-time. Updated process recommendations flowed back.

Ark was used to train a base model on 33 historical runs from Culture Biosciences using a hybrid architecture: a biologically informed cell metabolism model embedded within a multiphysics cell culture simulator. During the run, online data was streamed continuously and offline samples uploaded daily. At each cycle, the model retrained on the new data, ran a fresh optimization, and transmitted an updated process recipe. What would otherwise be a one-time pre-run prediction became a continuously updating strategy that compounded across the duration of the run.

Culture Biosciences' Stratyx bioreactor streams real-time process data externally and provides full API integration over all system parameters - feeding, agitation, pH, and more. Ark's recommendations were transmitted directly to the system at each optimization cycle, eliminating the need for manual transcription while preserving the established control framework. Because the bioreactor is API-native and the optimization cycle runs automatically, decisions based on model recommendations can be made around the clock. Cells don't sleep, and now neither does the model.

What this means for your process

This collaboration demonstrated three capabilities with implications for any organization running cell culture:

- Intra-run optimization delivers measurable benefits. The feed strategy for the top run diverged meaningfully from the static recipe with digital twin prescribing meaningfully more feed in the critical mid-run growth window, a shift the static recipe would never have made. That difference was the margin between a good run and the best run in program history.
- A digitally-native bioreactor removes the last bottleneck to closed-loop process optimization. Full API integration and real-time trend data is the prerequisite for autonomous optimization. Culture Biosciences' architecture demonstrated that this integration can be achieved without custom middleware or brittle workarounds.
- Digital twins compound in value. Models improve as training data grows. The current results were achieved on 33 runs; future iterations will benefit from the data generated by this study and every subsequent run, making the system progressively more accurate and responsive. There is also a major opportunity for adaptive DOE to improve the efficiency and reduce the number of runs needed for data training.

Ark-Culture Partnership

Ark and Culture Biosciences are partnering to bring closed-loop process optimization to process development teams. Together, the two platforms provide everything required to move from static, recipe-driven runs to adaptive, data-driven bioprocessing. Companies interested in leveraging this joint capability are encouraged to reach out to explore what a pilot collaboration could look like for their process.

To learn more about Ark visit https://www.ark-biotech.com/

Special Thanks

Ark: Robert Seares, Nick Austin, Raj Balaji-Wright, Jaron Thompson, Aidan O'Donnell, Daisie Ogawaand, and Jordan Lim.
Culture Biosciences: Babu Sivaraman, Jonathan Meed, Ramesh Veeramalla, and Savitri Beaver