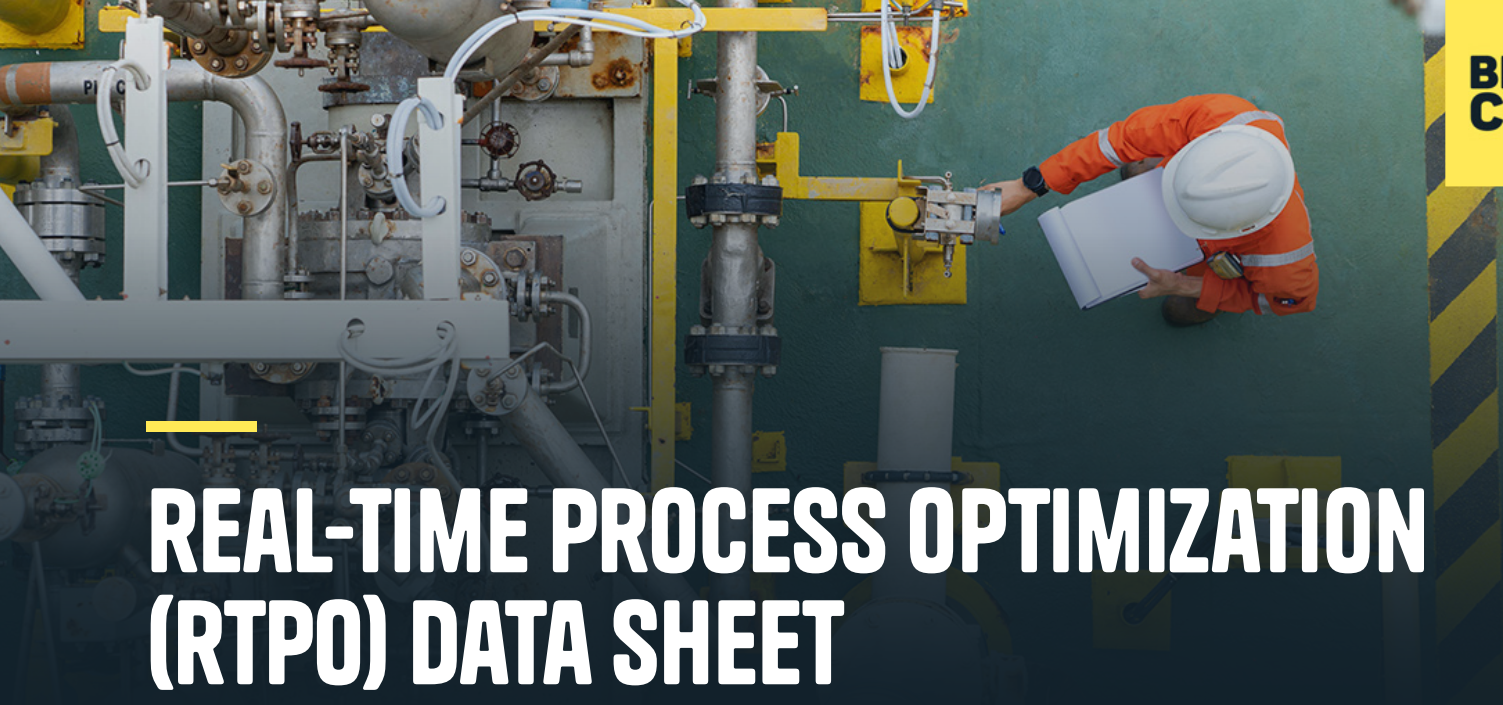

REAL-TIME PROCESS OPTIMIZATION (RTPO) DATA SHEET

THESIS

Most industrial AI fails not because the models are weak, but because the data was never structured to answer the question being asked. Braincube inverts this. Two proprietary technologies, the Product Clone and CrossRank AI, form a core, classical, stack that turns raw plant data into a decision-ready record of every product ever made. That foundation is what makes everything above it possible: real-time prescriptions through Real-time Process Optimization (RTPO), conversational access through Companion AI, and open, agent-driven interoperability through the Model Context Protocol (MCP).

The Foundation

Braincube, like many other software companies, has developed a core stack of traditional software applications called Product Clone and CrossRank. These form the backbone for the future of industrial automation - digital technologies that act as inputs for an array of AI-driven applications.

Product Clone

A Product Clone is not a time-series archive or an asset-based twin. It is a material-genealogy structure: one row per unit of product (a reel, a batch, a tonne), with every process variable that was in play time-aligned to that unit, from raw material to finished goods. Built per product rather than per machine, it answers the question control systems cannot: under exactly what conditions was this specific unit produced?

How Product Clones Are Different

Traditional industrial systems organize data around assets and timestamps. They answer questions such as: What happened on Line 3 at 10:04 AM?

Product Clones organize data around the product itself. They answer a different question: What conditions produced this specific batch, reel, or tonne?

This distinction is critical. Optimization requires understanding which conditions produced a specific outcome. By transforming time-series data into a product-centric record, Product Clones create the foundation for advanced analytics, RTPO, Companion AI, and agent-driven workflows.

How Product Clones Are Built

Product Clones are constructed by aggregating data from historians, MES systems, quality systems, and other operational sources. Braincube aligns process data to material flow, accounts for process delays and lagging, and creates a single product-centric record that captures the conditions under which every unit was produced.

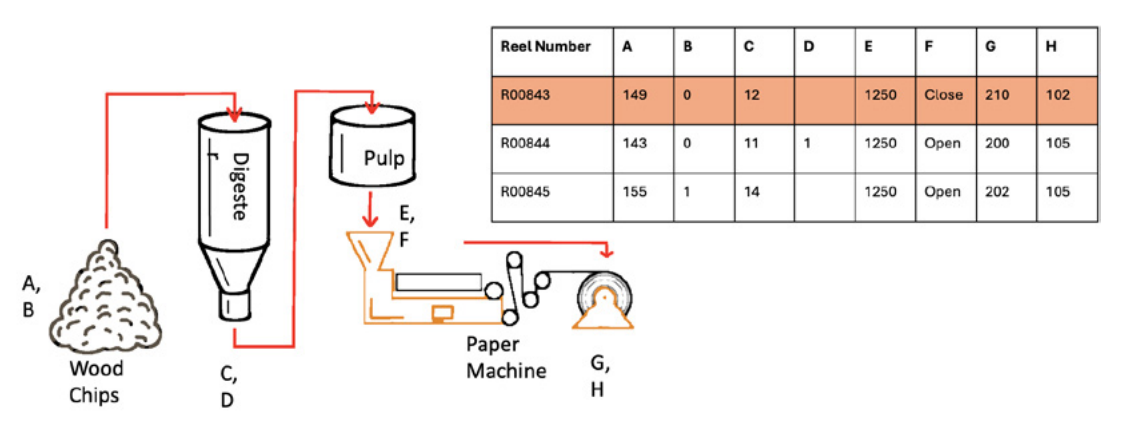

Figure 1: Product Clone transforming time-domain to product-domain

CrossRank AI

CrossRank AI measures the impact of every input variable on a chosen objective — yield, energy, throughput, scrap, quality. It is not correlation; it is a rule discriminator. The analyst defines a "good" range on the output, and for each variable CrossRank AI tests every possible interval to find the one that best separates good rows from bad, scored by a proprietary function balancing sensitivity, precision, and population balance. The result is a score and a golden operating range per variable.

PROPERTY	STANDARD CORRELATION	CROSSRANK AI
What is tested	Linear relationship between two variables	Whether a value range separates good from bad
Outliers / non-linearity	Sensitive / poorly handled	Robust / handled natively
Output	Coefficient + equation	Score + golden operating range
Harmful variables	Opposite sign, same magnitude	Explicit negative score

Table 1: CrossRank AI approach to sorting variables versus standard correlation algorithms

Two consequences: CrossRank AI does not search for lag (it consumes whatever pairing the clone provides, which is why clone construction is required), and it self-validates against randomly-categorized data — so a mis-built clone fails safe ("we find nothing") rather than producing false findings.

The Bridge: Real-Time Process Optimization

CrossRank AI tells engineers what has historically driven performance. Real-Time Process Optimization (RTPO) applies that intelligence continuously, helping operations adapt as conditions change and operate closer to optimal performance.

RTPO separates variables into contexts (non-controllable conditions) and settings (controllable levers). In Phase 1, it ranks variables by their impact on the desired outcome and filters the population to match current operating conditions. In Phase 2, it re-runs CrossRank AI on that filtered population to identify the operating ranges and setpoints most likely to achieve the desired result.

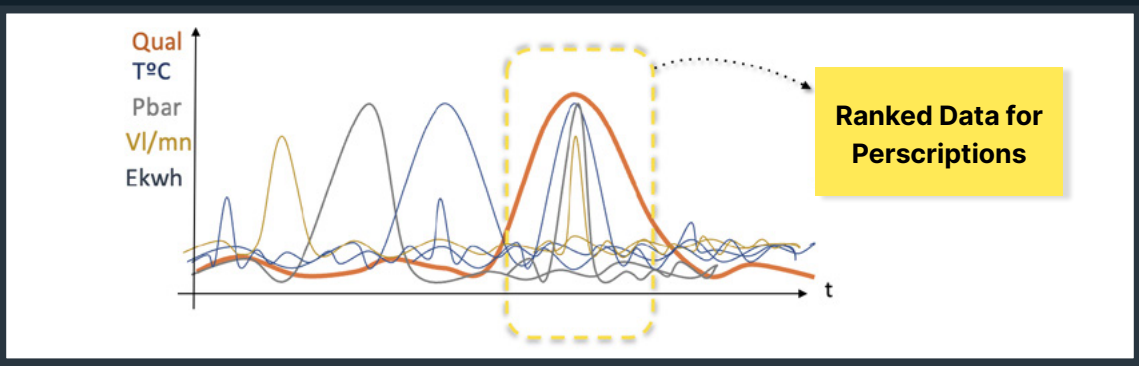

Figure 2: RTPO uses aligned and ranked data for set-point prescriptions

Operationalizing RTPO

RTPO recommendations are deployed through standardized implementation flows built using Node-RED. These flows connect Braincube intelligence with existing operational systems, allowing recommendations, alerts, and actions to be delivered directly into plant workflows.

This approach enables organizations to deploy repeatable optimization strategies across sites while maintaining compatibility with existing control systems and operating procedures.

The Future: Companion AI and MCP

The classical software stack produces decision-ready intelligence; the final question is access. The Companion AI is the conversational interface to the platform and RTPO, bringing grounded analysis and prescriptions into the flow of production for users who are not data scientists. The Model Context Protocol (MCP) is an open standard that exposes Product Clone data, CrossRank AI analyses, and RTPO outputs as structured, callable tools, so external AI agents and systems can reason over the plant's intelligence without bespoke integration. Where the RESTful API made data portable, MCP makes the intelligence portable — the connective tissue that lets the whole stack participate in a broader agentic ecosystem.

SEE HOW MUCH CAPACITY IS HIDDEN IN YOUR PROCESS

Schedule a performance assessment to see what
Real-Time Process Optimization could unlock in your operation.

Schedule a performance assessment

Glossary

Product Clone — Braincube's material-genealogy structure; one row per unit of product with all process variables time-aligned to it. The structural foundation of the stack.

Lagging / lag group — The time-shift aligning upstream variables to the product they influenced, computed per lag group via a backward outflow sum under plug-flow assumptions. Lag variance is not stored, so the aggregation period absorbs it.

Classical software — Computer system in which developers directly specify the rules and procedures that determine how inputs are processed into outputs, resulting in predictable and explainable behavior.

CrossRank AI — Braincube's engine scores how strongly each variable separates good outcomes from bad, returning a golden operating range. A robust, non-linear rule discriminator, not a correlation; it self-validates against random data.

Contexts / Conditions — Variables influencing a KPI that cannot be controlled in real time (ambient conditions, raw-material properties, recipe/SKU). RTPO filters the population to match the current context.

Settings — Controllable levers (temperatures, speeds, flows, dosing) for which RTPO recommends ranges and setpoints.

RTPO (Real-Time Process Optimization) — Continuous application of CrossRank against live data. Phase 1 ranks and filters; Phase 2 prescribes. Acts as an operator "GPS," moving setpoints rather than holding them.

AI Companion — The conversational interface to the platform and RTPO, bringing grounded analysis into the flow of production for non-specialists.

Model Context Protocol (MCP) — Open standard exposing clone data, CrossRank analyses, and RTPO outputs as callable tools so external AI agents can reason over the plant's intelligence.