

Task Mining vs Process Mining vs Process Discovery

Three terms, often used interchangeably,
that mean three different things.

BY CARRV.AI

7 MIN READ · FEB 2026

CARRV.AI

ENTERPRISE INTELLIGENCE, CARVED AT THE SOURCE.

Three things, sold as one.

Three terms. Three categories. Three sets of vendors who would each like you to think they're synonyms — because each one wants its product to be all three.

They're not synonyms. They sit at different levels of the same stack, answer different questions, and produce different artifacts. Whether you need one, two, or all three depends entirely on what you're trying to do.

This is the comparison the vendor marketing pages won't write.

PROCESS MINING
Reads system logs.

TASK MINING
Aggregates user activity.

PROCESS DISCOVERY
Captures procedures.

Reads what systems remember.

ORIGIN

1999 · Eindhoven

INPUTS

Event logs from systems

ANSWERS

Where do cases get stuck?

The oldest category. Founded as an academic discipline by Wil van der Aalst at Eindhoven in the late 1990s, commercialized starting around 2011 (Celonis, Software AG, Fluxicon, Apromore), and now a real budget line in most Fortune 500 IT departments.

What it eats. Event logs from enterprise systems of record. SAP, Salesforce, ServiceNow, Oracle, Workday — anywhere a transaction leaves an audit trail. A typical input is a CSV with three columns: case ID, activity name, timestamp. Three columns and a few million rows.

What it does. Runs discovery algorithms — Alpha miner, Inductive miner, Heuristic miner — over the event log and reconstructs the implied process graph. Then layers conformance checking and performance analytics on top.

What it produces. A flow chart of how an order, a claim, or a ticket actually moved through your systems, along with bottleneck heatmaps, conformance gaps, and SLA breach analytics. Aggregated across thousands or millions of cases.

What it sees clearly. System-mediated work. Order-to-cash. Procure-to-pay. Incident-to-resolution. Anything that lives substantially inside one or two systems of record.

What it's blind to. Anything that happens outside the logs. If your AP clerks reconcile invoices in Excel, process mining never knows the Excel step existed. If an underwriter pulls a credit report in a browser, drops it into a PDF reader, scribbles a note, and then returns to the CRM — process mining sees only the CRM events. The gaps in the chain are invisible to it.

This is the dominant limitation of the category, and the one its vendors are quietest about. Real enterprise work happens substantially outside systems of record. Process mining cannot see that work.

Watches everyone, statistically.

ORIGIN

~2018 · commercial

INPUTS

User activity, many users

ANSWERS

Where does time go?

The middle child. Emerged commercially around 2018–2020 (NICE, Soroco, Kryon, UiPath Task Mining). Designed explicitly to fill the gap process mining can't see.

What it eats. User activity at the workstation. Mouse movements, clicks, keystrokes, window switches, application focus events — captured by a lightweight agent installed on the user's machine. Often paired with screenshots, sometimes with OCR over the screen content.

What it does. Records that stream of low-level events from many users, then uses machine learning to classify the activity into recognizable tasks (“processing an invoice,” “responding to a ticket”), cluster similar sequences, and produce statistical summaries.

What it produces. Time-and-motion analytics for desktop work. How many hours a week the AP team spends in Excel. How many windows the average rep switches between in a typical case. Which tasks consume the most user time across the department. Automation opportunity scores. Heatmaps.

What it sees clearly. The long tail of UI work that process mining misses. The Excel reconciliations. The copy-pastes between systems. The portals nobody documented. The cumulative cost of context switching.

What it's blind to. The procedure itself. Task mining aggregates. It tells you that across 200 AP clerks, the median invoice processing session involves 47 clicks across 6 applications. It does not tell you what those clicks are, in what order, with what inputs, leading to what outputs. It tells you where the time goes. It does not tell you how to do the work.

This is the category's central trade-off. You get coverage. You don't get procedure.

Captures one full procedure.

ORIGIN

~2010 · structured 2022+

INPUTS

Expert session, role chain

ANSWERS

How is this procedure done?

The category that's been here the longest under different names — earlier as screen-capture tooling, then as digital adoption platforms (WalkMe, Whatfix), and as standalone discovery tools (Epiplex500) for nearly two decades.

What changed recently is the structure. Capture is no longer just a recording, and the output is no longer just a help bubble. Modern process discovery captures structured, multi-actor procedures and generates downstream artifacts — simulations, SOPs, training material, AI training data — from a single source.

What it eats. The same raw signals as task mining — clicks, keystrokes, screens, application events — but treated very differently.

What it does. Captures one user, one full task, one complete session, and reassembles those signals into a structured, ordered, end-to-end representation of that single procedure. Not statistics over many users. The procedure itself, for one user (or one role chain, when the process passes through several hands), with all its steps preserved, its decision branches marked, its actor handoffs traced.

A deliverable, not a dashboard.

What it produces. A structured artifact that can be replayed as an interactive simulation, generated into a written SOP, exported as a SCORM training module, rendered as a video walkthrough, or fed as training data to an AI agent. Not analytics. A deliverable.

What it sees clearly. The actual procedure. End to end. With the screens. In the order they happened. Including the decision points and the handoffs across roles.

What it's blind to. Breadth. Process discovery captures fidelity over scale. You don't deploy it across 200 users to find automation opportunities. You deploy it on the people who know the process to capture the procedure with structure. Different question, different tool.

SIMULATION

Replayable, click-through.

DOCUMENTS

SOPs, video, SCORM.

AI TRAINING

Structured procedure data.

FROM A SINGLE CAPTURED SOURCE.

What an honest matrix looks like.

If a procurement team forced all three categories onto one slide, the table their vendors won't print:

	PROCESS MINING	TASK MINING	PROCESS DISCOVERY
INPUT	Event logs from systems of record	User activity across many users	Expert session (or role chain)
GRANULARITY	Cases across thousands	Aggregated tasks across users	One procedure, fully specified
OUTPUT	Flow graphs + analytics	Statistics + heatmaps	Structured walkthroughs
BEST QUESTION	Where do cases get stuck?	Where does time go?	How is the procedure done?
BLIND SPOT	Anything outside systems	One discrete procedure	Cross-org scale
BUYER	Ops excellence, IT	Automation CoE, RPA team	L&D, knowledge ops
GOES TO	Dashboards	Bot backlogs	Training, sims, AI

Not because the table is wrong, but because it pins each product to a single column when the marketing teams want all three.

When to use which.

The categories are not redundant. They are complementary at different altitudes.

PROCESS MINING when

you need to understand flows across a large population of cases inside systems of record. Where are the bottlenecks? What's the conformance rate? Where do SLAs break? Process mining will answer these well and is mature.

TASK MINING when

you need to find automation opportunities across many users doing similar UI work. Which repetitive tasks are eating aggregate time? Where would a bot pay back? Where is unrecognized work happening outside monitored systems? Task mining will surface these.

PROCESS DISCOVERY when

you need to capture how a specific procedure is actually done — for training new hires, documenting it for compliance, building a simulation, or feeding it to an AI agent. Where do the experts click? What are the decision points? What happens at the handoff between roles? Process discovery is the answer.

A mature enterprise probably runs all three at different times for different reasons. The mistake is assuming they overlap so much that one can substitute for the others.

Why the categories blur.

Three reasons.

01 Adjacent telemetry.

Task mining and process discovery use overlapping data sources. If you have an agent capturing clicks for task mining, the temptation to claim discovery capability with the same agent is irresistible. Marketing departments rarely resist temptation.

02 TAM expansion.

Every vendor in this space would prefer to compete in a \$20B category instead of a \$2B one. The cleanest way to enlarge your category is to claim adjacent categories are the same as yours. Process mining vendors now claim task mining. Task mining vendors now claim discovery. Discovery platforms claim mining. The vocabulary becomes pollution.

03 Genuine technical overlap.

All three are converging on similar substrates — agents, OCR, ML classification, structured event capture. A platform that does one well often has the components to do another. The question is whether it does it well, or just plausibly enough for a demo.

If you only have ten seconds.

PROCESS MINING

reads system logs and tells you about flow at scale.

TASK MINING

watches screens and tells you about time across users.

PROCESS DISCOVERY

captures one full procedure and produces the artifact you actually train, simulate, or automate from.

Three categories. Three layers. Three answers.

Not synonyms.

If anyone tries to sell you all three from a single platform, ask them which column of the matrix they came from. The honest ones will tell you.

CARRV.AI

ENTERPRISE INTELLIGENCE, CARVED AT THE SOURCE.



Task Mining vs Process Mining vs Process Discovery

FOUNDATIONS · FEB 2026

CARRV.AI / RESOURCES