

Hands-on



SAP Signavio – AI Agent Mining

AI Agent Mining with SAP Signavio -- Coffee Shop Scenario

Let's imagine you're Ella. On your way to class, you run by a coffee shop and fancy a breakfast. You enter the coffee shop, and an *Order Agent* takes your order:

AI Coffee Shop

New Order

Verbose: Off

You

Hi! I'd like to order a large latte with oat milk and a croissant please.

10:38:18 AM

Order Agent

Thank you for your order! Just to confirm, here's what you'd like:

- Large Latte with oat milk
- Croissant

Could you please provide your name for the order? Also, would you like anything else, or is this everything for today?

10:38:20 AM

My name is Ella. Oh yes, some caramel syrup for my latte, please.

Order

Menu & Inventory

Espresso 20 in stock	\$2.50 Coffee
Latte 15 in stock	\$4.00 Coffee
Cappuccino 18 in stock	\$3.75 Coffee
Americano 22 in stock	\$3.00 Coffee
Croissant 8 in stock	\$2.75 Pastry
Muffin 12 in stock	\$3.25 Pastry
Sandwich 5 in stock	\$6.50 Food

The *Order Agent* sums up your order and calculates the total cost:

 **Order Agent**
Thank you, Ella! Here's your order summary:

- Large Latte with oat milk and caramel syrup
- Croissant

Your current total is **\$8.50**.
Does everything look correct? If so, I'll check availability and get your order started!
10:41:24 AM

 **You**
Yes, please.
10:44:08 AM

As you confirm the order, the *Order Agent* hands over to the *Inventory Agent* checking the inventory for your order in the background. As all ordered ingredients are in stock, the *Inventory Agent* hands over to the *Barista Agent* preparing your order:

 **Order Agent**
10:44:08 AM

 **Inventory Agent**
10:44:10 AM

 **Barista Agent**
10:44:13 AM

 **Barista Agent**
Your order is being prepared, Ella!
Preparation Progress:

- Making 1x Large Latte (with oat milk and caramel syrup)
- Making 1x Croissant

All items have been prepared successfully!
Your order will be ready shortly—thank you for choosing us! If you have any other requests, just let me know.
10:44:14 AM

Soon after, you receive your latte and your croissant – yummy!

While this scenario looks fast and easy, a lot is happening in the background of the AI agents work. The step-by-step record of an AI agent's reasoning and action execution is called a *trace*. For instance, the total price needs to be calculated:

```
{
  "id": "ORD2073",
  "total": 8.5,
  "status": "pending",
  "customer": "Ella",
  "items": [
    {
      "name": "latte",
      "quantity": 1,
      "price": 5.75,
      "size": "large",
      "extras": [
        "oat milk",
        "caramel syrup"
      ]
    },
    {
      "name": "croissant",
      "quantity": 1,
      "price": 2.75,
      "size": null,
      "extras": []
    }
  ]
}
```

Furthermore, the stock is checked and updated:

```
{
  "report": {
    "order_id": "ORD5813",
    "all_available": true,
    "details": [
      {
        "name": "latte",
```

```

    "requested_quantity": 1,
    "available_quantity": 15,
    "status": "available"
  },
  {
    "name": "croissant",
    "requested_quantity": 1,
    "available_quantity": 8,
    "status": "available"
  }
],
"unavailable_items": []
}
...
}

```

Besides this, other scenarios can happen. There may be an item out of stock:


Order Agent

It looks like our menu offers:

- Muffin
- Sandwich
- Croissant
- Coffee drinks (espresso, latte, cappuccino, americano)

We do not have specific flavors (like blueberry) or types (like veggie) listed for muffins, bagels, or sandwiches. Would you like to proceed with:

- Sandwich (standard)
- Muffin (2)
- Skip the bagel (since it's not on the menu)?

Please confirm or let me know if you'd like to adjust your order!

9:49:55 PM

Or there could be a complaining scenario which may require the involvement of the *Customer Service Agent*:

 You

Excuse me, I ordered a cappuccino earlier and it was completely cold when I got it. I'm really not happy about this.

9:50:49 PM

 Order Agent

I'm so sorry to hear that your cappuccino was cold. I understand how disappointing that must be. I'd like to make this right for you.

Would you like a replacement drink or a refund? I can connect you with our customer service agent right away to resolve this issue.

Could you please provide the name your order was under, just to speed things up?

9:50:50 PM

 You

I'm Ella. I'd like to have a replacement drink (cappuccino).

9:51:45 PM

 Order Agent

9:51:46 PM

 Uses tool: Transfer To Customer Service Agent

Successfully transferred to customer_service_agent

9:51:46 PM

 Customer Service Agent

Thank you for letting me know, Ella, and I sincerely apologize for your experience. Here's what I'll do to resolve this:

- I'll arrange a replacement cappuccino for you right away.
- To make up for the inconvenience, I'd also like to offer you a **discount on your next order**.

Your satisfaction is very important to us. Would you like your replacement cappuccino prepared now, or at a specific time?

9:51:47 PM

In summary, different *AI agent traces* are produced due to varying scenarios (for example, an out-of-stock item or a customer complaint) and the corresponding actions and reasoning of each agent as it performs its tasks. However, from a business perspective, there most often is a desired, optimal process without deviations and bottlenecks. This is called *happy path* and it either is defined/refined before running a process analysis or after process analysis. The happy path in the AI coffee shop scenario is the following:




In the next steps, this activity will guide you through the configuration and analysis of the AI agent trace data generated by different scenarios that happened in the AI coffee shop.

Data Reference


Column header	Data type	Description
case_id	Case ID	This is your data's unique case identifier.
activity:id	Text	This is your data's unique event identifier.
time:timestamp	Date time	Represents the beginning of an event in the format YYYY-MM-DDTHH:mm:ss.SSS.
concept:name	Choice	Lists the general activity category.
concept:instance	Activity	Lists the event names in your data.
org:role	Choice	Differentiates the activity by the actor (user or agent).
ai:agent_name	Choice	Differentiates the events based on the executing agent's name.
ai:message	Text	Represents the text input from users.
ai:target_agent_name	Choice	Shows the agent to which a handover takes place.
time_finished	End timestamp	Represents the end of an event in the format YYYY-MM-DDTHH:mm:ss.SSS.
duration	Number	Shows the duration of an activity in milliseconds.
ai:model	Choice	Shows the chosen AI model.
ai:input_tokens	Number	Shows the number of AI input tokens.
ai:response_tokens	Number	Shows the number of AI response tokens.
ai:response_message	Text	Shows the response of the AI agent.
ai:tool_name	Choice	Tools that the AI agent used.

AI Coffee Shop Analysis in SAP Signavio


1. You can access SAP Signavio at <https://tinyurl.com/aspai2026> and log in with the user aspai.pohang@gmail.com (password: aspai.pohang)
2. Navigate to **Analysis Configuration** and open **ASPAI Agent Mining**
3. First, you individually create a copy of the [MASTER] Dashboard via “Save as” and choose a unique name for your copy..



4. You then enter your copy of the dashboard. By moving the *Edit* slider to the right, you can adjust the dashboard's name and its widgets or add new ones. On the bottom of the dashboard, you see the pages. In this example, there is one for the *Metrics* and one for the *Conformance*.

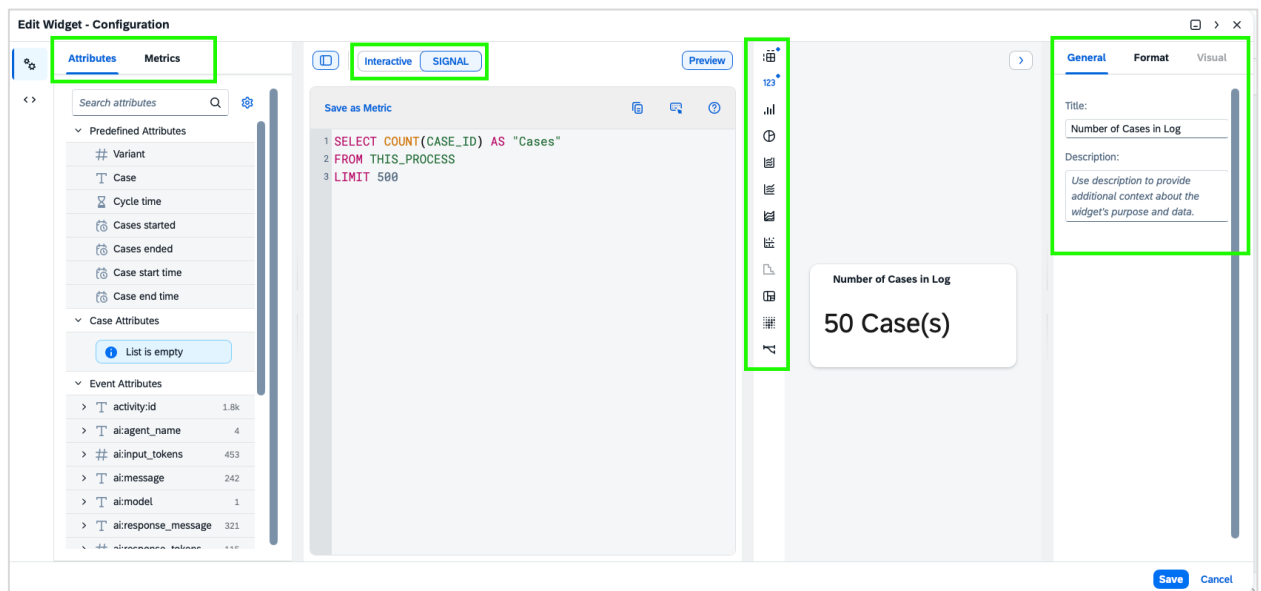


Additionally, you can access each single widget's configuration mode by clicking on the widget's three dot settings icon:




In there, you can see

- the *Attributes* and *Metrics* on the left side,
- the *Interactive drag & drop mode* or the *SIGNAL* code for the widget
- the *visualization options* such as table, value, bar charts, etc.
- and the *customization options* such as title, description, format, etc. on the right side.




In the following, the existing widgets are briefly introduced. For more extensive details on building and interpreting widgets, you can visit the [User Guide for SAP Signavio Process Intelligence](#).

- **Process discovery:** This widget automatically identifies, analyzes, and presents the process flows of the AI agents in the coffee shop scenario. In our example, when setting the *Activities* and *Paths* visualizations both to 100%, you can clearly see the most frequent path followed which is close to our happy path. However, there are several deviations where the process went other than expected.



- **Key metrics:** These widgets give you an overview on important numbers of your event log. It contains 50 cases and 1,774 events. Further, the average prompt-to-result-time is 16s 937ms, LLM calls on average take 1.21s, and the total LLM costs (calculated as sum of the consumed input and response tokens) is \$1.51. Also, we see that on average 2.9 agents are involved in a case.



- **AI agent and LLM information:** These three table widgets give you an overview on the used LLM (here gpt-4.1) the agents that are active in your log and the tools they used in the coffee shop scenario.

Utilized LLMs



LLM Model
gpt-4.1-2025-04-14

Agents in Log




Agent Name
Barista Agent
Order Agent
Inventory Agent
Customer Service Agent

Tools Used by Agents




Tool	used by
Estimate Prep Time	Barista Agent
Prepare Order	Barista Agent
Remake Order Item	Barista Agent
Offer Refund	Customer Service Agent
Offer Partial Refund	Customer Service Agent
Get Alternatives	Inventory Agent
Update Stock	Inventory Agent
Check Inventory	Inventory Agent
Process Order	Order Agent

- **Variant explorer:** Comparable to the process discovery, the variant explorer shows how the process actually runs by breaking it down to the different variants. Each variant represents a unique sequence of activities. Based on the selected aggregation (cases, events or average cycle time), it is helpful for detecting process deviations or bottlenecks, such as loops. For example, when comparing the first with the fifth most frequent variant, we see that in most cases the *Inventory Agent* does not update the stock after checking the inventory for the order.



Based on the selected variants, the widget additionally allows to automatically generate a BPMN model in the SAP Signavio Process Modeler via the button *Generate BPMN*.

- **Token overview:** Next, two table widgets give an overview of the total consumed tokens in respective costs, additionally separated by input and response. In addition, a pie chart shows how many of these tokens were consumed by which agents.



- **Conversation information:** The first table shows the input and response messages (= prompts) of the users, in our scenario the visitors (like Ella) who came to the AI coffee shop and ordered something. In the second table, the whole conversation history between users and AI agents is given for further reference.

case_id	Prompt
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_12	Hi there! Could I get a large espresso with caramel syrup and oat milk, please? And I'll take one muffin as well.
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_12	Sure! My name's Sam, and I'll go with a blueberry muffin, please.
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_12	No worries, the regular muffin is perfect! Everything else sounds good—thank you!
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_12	That's all for me, thanks so much! I'm all set. [END]
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_12	Thanks so much! Can't wait to enjoy it. Have a great day! [END]
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_13	Hi there! Could I please get two muffins and a small americano with caramel syrup? Thank you!
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_13	Sure! My name's Alex. For the muffins, could I get one blueberry and one chocolate chip if you have them, please? No extras, thanks!
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_13	That's totally fine, I'll go with two regular muffins. Thanks for checking! Everything else looks perfect.

Case	User Query	Agent Response
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_15	Sure! My name is Alex, and for the muffins, could I please get one blueberry and one chocolate chip? Thanks so much! That's perfect, thank you all so much! Have a great day! [END]	Thank you, Alex! Here's your order sum You're very welcome, Alex! **Status:** Or
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_16	Hi, could I please get a sandwich and a hot chocolate? I'll have a turkey sandwich, please, and a medium hot chocolate. Thanks! That's totally fine, I'll go with the chef's choice for the sandwich. And I'll have a latte instead of the hot chocolate, please. Thanks! No extras for the latte, and nothing else to add—I'm all set. Thanks for your help! [END] Great, thank you so much! I'll just wait for my order. [END] Thank you! Looking forward to it. [END]	Thank you for your order! To ensure even It looks like the options available are a ge Here's your order summary: **Sandwich **Inventory Check:** - **Sandwich** You're welcome! Your order is being prep You're most welcome! Everything is in sto
db8616b6-5ba1-4aa0-b08a-dcbaa033479a_sim_17	Hi! I'd like a small espresso with an extra shot, please. That's everything for me! Sure, it's Alex. Thank you! Great, thank you so much! Have a wonderful day. [END]	Thank you for your order! Here's what I've **Inventory Status for Order ORD1184 (A Thank you, Alex! **Order complete.** Hav

- **Process conformance:** On the bottom, you can switch dashboard pages to *Conformance*. Process conformance analysis is one of the key pillars of process mining. After you copied the dashboard, you need to re-establish the mapping to the corresponding process model to allow process conformance checking via its “Settings”.

[MASTER] AI coffee shop (copy)
Save as
Auto-insights

Settings

Process Conformance


Please link a model in order to use this widget

Map activities to process model


In order to use Conformance, Throughput or Diagram widgets, select a BPMN model to link with the dashboard.

ASP AI


[ASP AI] ASP AI



Inside the activity mapping screen, select “Auto-map”.



- It enables comparison between the as-is process with the to-be process and vice versa. Compared to the similar *Variant Explorer* that is mainly used for comparing the different process variants, the process conformance widget enables comparison and the identification of executing frequency (*Hotspots* view) and average activity sequence of a variant (*Variant path* view). For instance, in the first variant, we see that compared to our happy path (to-be process) the agents skip the steps of calculating the total of an order, updating the stock, and estimating the preparation time.



Note that our happy path BPMN model is already linked to the dashboard in the file **Agent Mining with Signavio_Coffee Shop_Dashboard.json** you imported.


Widget Creation: You can check if the “Edit” mode is turned on




Click the “Create Widget” button



Select “Text to Widget” icon



You can try any widget creation based on the data in the dashboard. For example “Average Duration by AI Agent”. You can click “Create” if you like this widget.



Analysis Questions

Question 1



The Inventory Agent in the Activity 2 exercises skips the stock update step in several process variants. What business risk does this create for the coffee shop? Using your process mining analysis, identify which variant shows this deviation most frequently and suggest at least two ways to address it - one in the agent configuration and one in the process design.

Question 2



Your process mining analysis shows the average time from user prompt to result is approximately 16 seconds across 50 cases. A competing coffee shop uses human staff and averages 90 seconds from order to confirmation. Using the event log metrics you analyzed, make a case for or against replacing the human-staffed process with the AI agent system. Consider cost, speed, error rate, and customer experience.

Question 3



Process conformance checking compares the actual process (as-is) against an intended reference model (to-be). In this activity, the happy path served as the to-be model. What assumptions are built into defining that happy path? Who should be involved in defining the to-be model in a real enterprise deployment, and what happens when the actual process diverges significantly from the reference model?

Question 6



The coffee shop event log includes the full text of customer messages in the ai:message column. In a real deployment, what privacy and data protection considerations would apply to storing and mining this data? How would you redesign the log to reduce privacy risk while preserving analytical value?