



# RAIDERS OF THE LOST CONTEXT


---

*The Existential Era of Enterprise AI*

**David Pontoppidan**


Head of Business AI, Nordics & Baltics | SAP

SAPSA Impuls | Gothenburg | 14 April 2026



# LLMs are getting better

AI's ability to complete long-horizon tasks doubles every 7 months



By 2028

# Agents Complete Tasks of a Full Day



So is speed why we need to transform?



Semiophore

# Data archeology **is a problem for meaning...** and management



## UNUSED POTENTIAL

**68%**

of available enterprise data goes unleveraged

*Seagate / IDC baseline*

**∞ dashboards**

**≈ 0 decisions**

Visibility becomes the product.  
Action stays scarce.

## THE EXTRACTION ERA

**\$34.8B**

global BI market  
2025

**32%**

of enterprise data  
actually put to work

**90%**

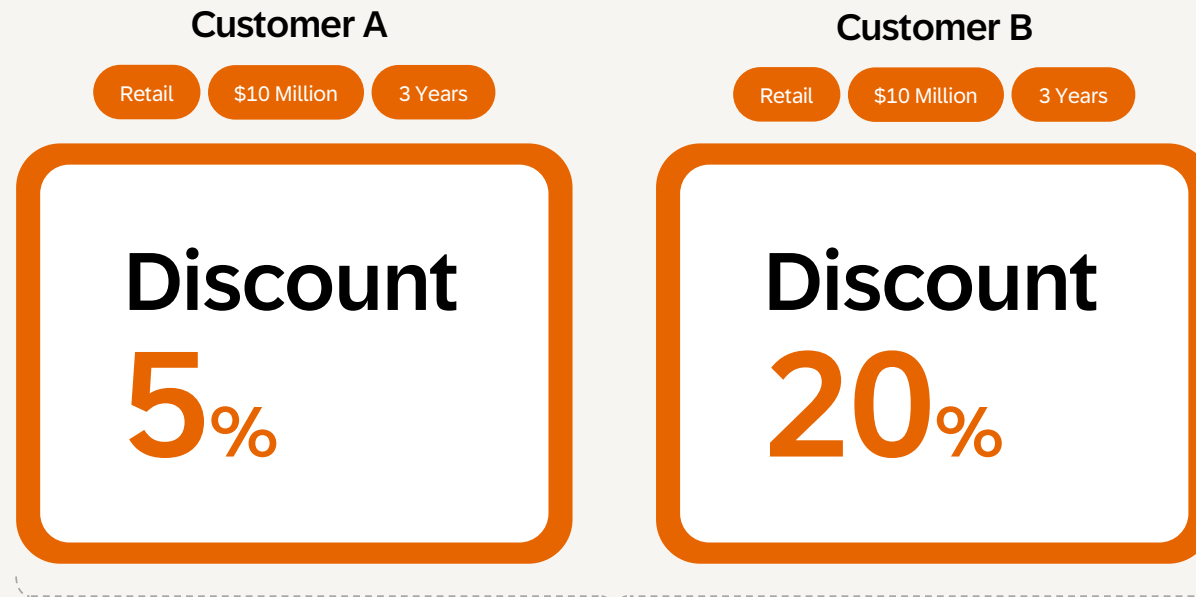
say data in the flow of work  
would help them perform better



# Data does not belong behind glass



# Systems of Record



How Should an Agent Know  
How to Choose Next Time?



# GAME OVER



You extracted the data but lost the context.



# Systems of Context

## Customer A

Retail

\$10 Million

3 Years

Discount  
5%

Past

- Exceptions
- Escalations
- Precedents

## Customer B

Retail

\$10 Million

3 Years


Discount  
20%

Past

- Exceptions
- Escalations
- Precedents

# Without context, we lose our foundation for truth

## LIVE BUSINESS OBJECT



### 20% discount request

Awaiting approval

|          |                     |
|----------|---------------------|
| Customer | Northwind Retail    |
| Approver | VP Sales            |
| Stage    | Deal desk 4 of 6    |
| Policy   | Retention exception |
| Quote    | Q-450021783         |



## EXTRACTION

*The numbers survive*


## DISPLAY OBJECT

### Discount exceptions > 15%

# 17

+2 vs last quarter

|           |        |
|-----------|--------|
| Trend     | Rising |
| Owner     | —      |
| Next step | —      |
| Policy    | —      |



## 01 RELATIONSHIPS

Customer, quote, approver.

## 02 INTENT

Why it exists. Which decision it supports.


## 03 PROCESS

Where it sits. What happens next.

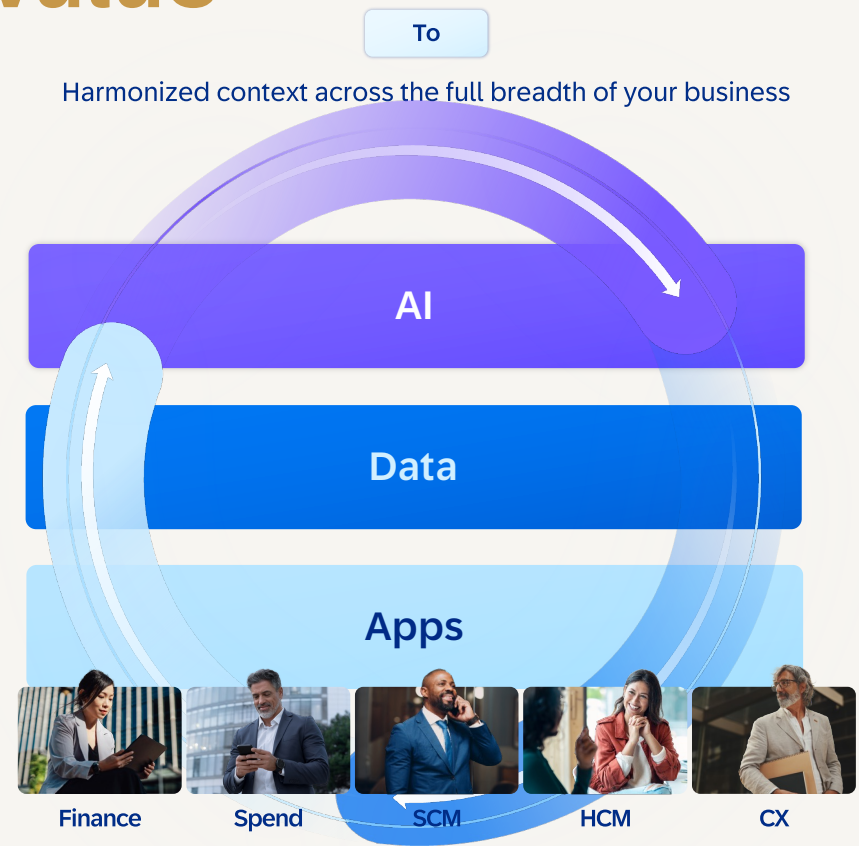
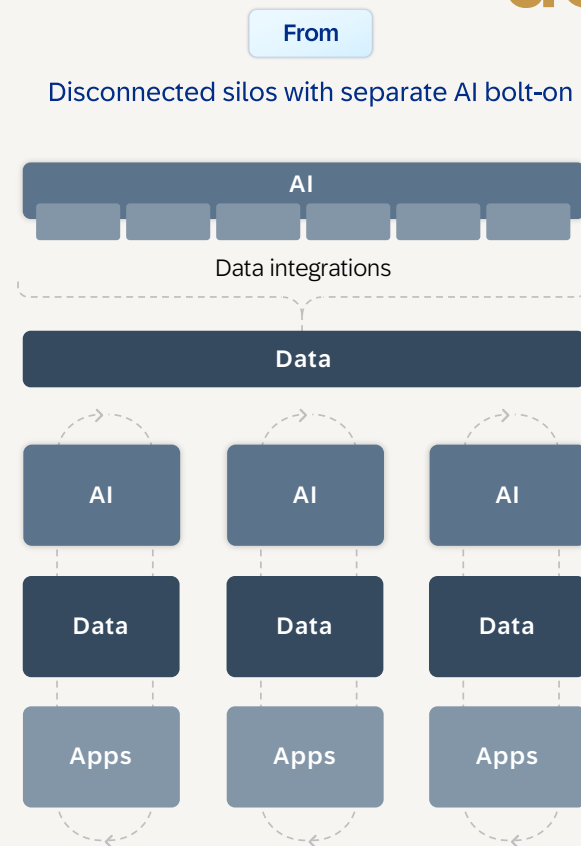
## 04 TRUST

Who may see, approve, or override.

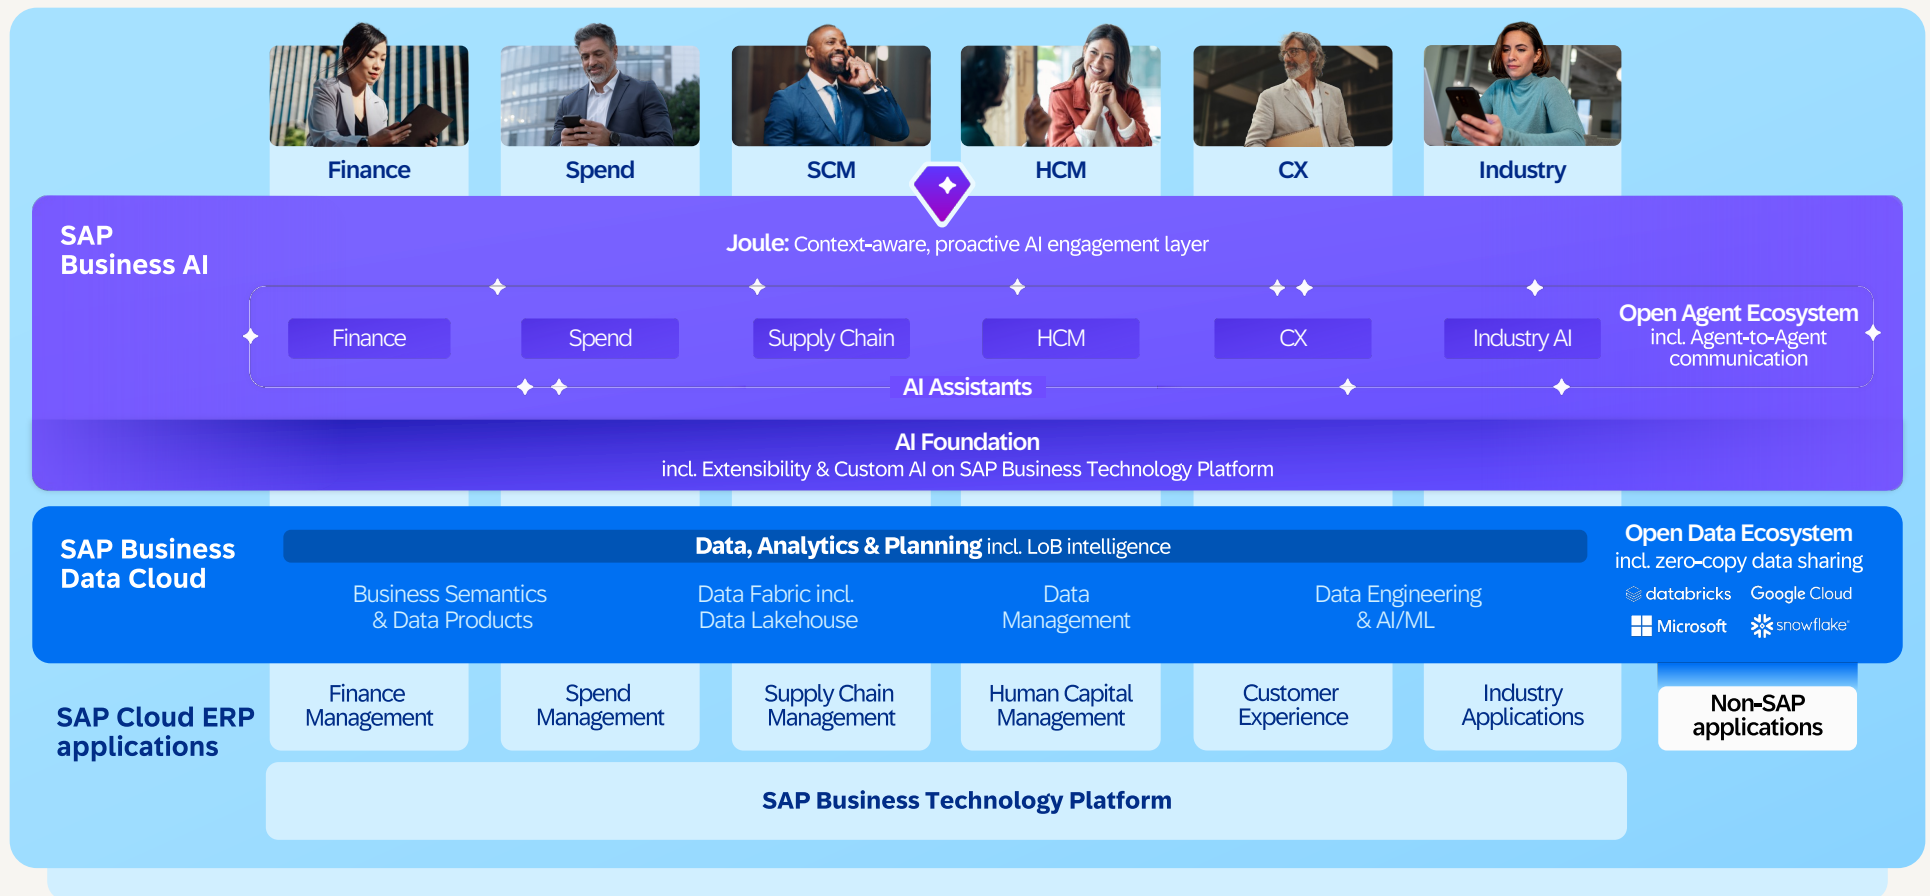
# Hard problems are often hard for a reason



# AI with the broadest context provides the deepest value



# Value is an architectural concept





# Your infra is no longer just your infra

## Which model?

GPT, Claude, Gemini, Mistral, Llama, Qwen ...

## Which framework?

LangGraph, CrewAI, Semantic Kernel, or none at all

## Which agents?

Custom-built, third-party, partner  
Single agent, multi-agent  
Agent scaffolding

## Which data?

SAP, Snowflake, Databricks, Big Query, Fabric  
Data lake, legacy  
Structured data, unstructured data

## Which governance?

Who approves? Who audits?  
Whose identity? Whose SLA?

# Your infra is no longer just your infra

## Which model?

GPT, Claude, Gemini, Mistral, Llama, Qwen ...

## Which framework?

LangGraph, CrewAI, Semantic Kernel, or none at all

## Which agents?

Custom-built, third-party, partner  
Single agent, multi-agent  
Agent scaffolding

## Which data?

SAP, Snowflake, Databricks, Big Query, Fabric  
Data lake, legacy  
Structured data, unstructured data

## Which governance?

Who approves? Who audits?  
Whose identity? Whose SLA?

### Amazon service was taken down by AI coding bot

Tech giant blames 'user error, not AI error' for incident in December involving its Kiro tool




Multiple Amazon employees told the FT that this was the second occasion in recent months in which one of the group's AI tools had been at the centre of a service disruption. © FT montage

Rafe Rosner-Uddin in San Francisco

### McKinsey rushes to fix AI system after hacker exposes flaws

Consultancy says it has found 'no evidence' that confidential client information was compromised



The hack is potentially embarrassing for McKinsey at a time when it is pitching for work advising blue-chip companies on how to use AI. © FT montage/Getty Images

### McKinsey, March 2026

An autonomous AI agent breached internal AI platform Lilli in two hours. No credentials needed. 22 API endpoints had no authentication. Full read-write access to 46 million messages and the system prompts governing AI behavior for 40,000 employees.

### Amazon, February 2026

Emergency internal meeting after AI-assisted code changes caused a string of outages including a six-hour retail site disruption. Safeguards for AI-assisted changes were not yet established.

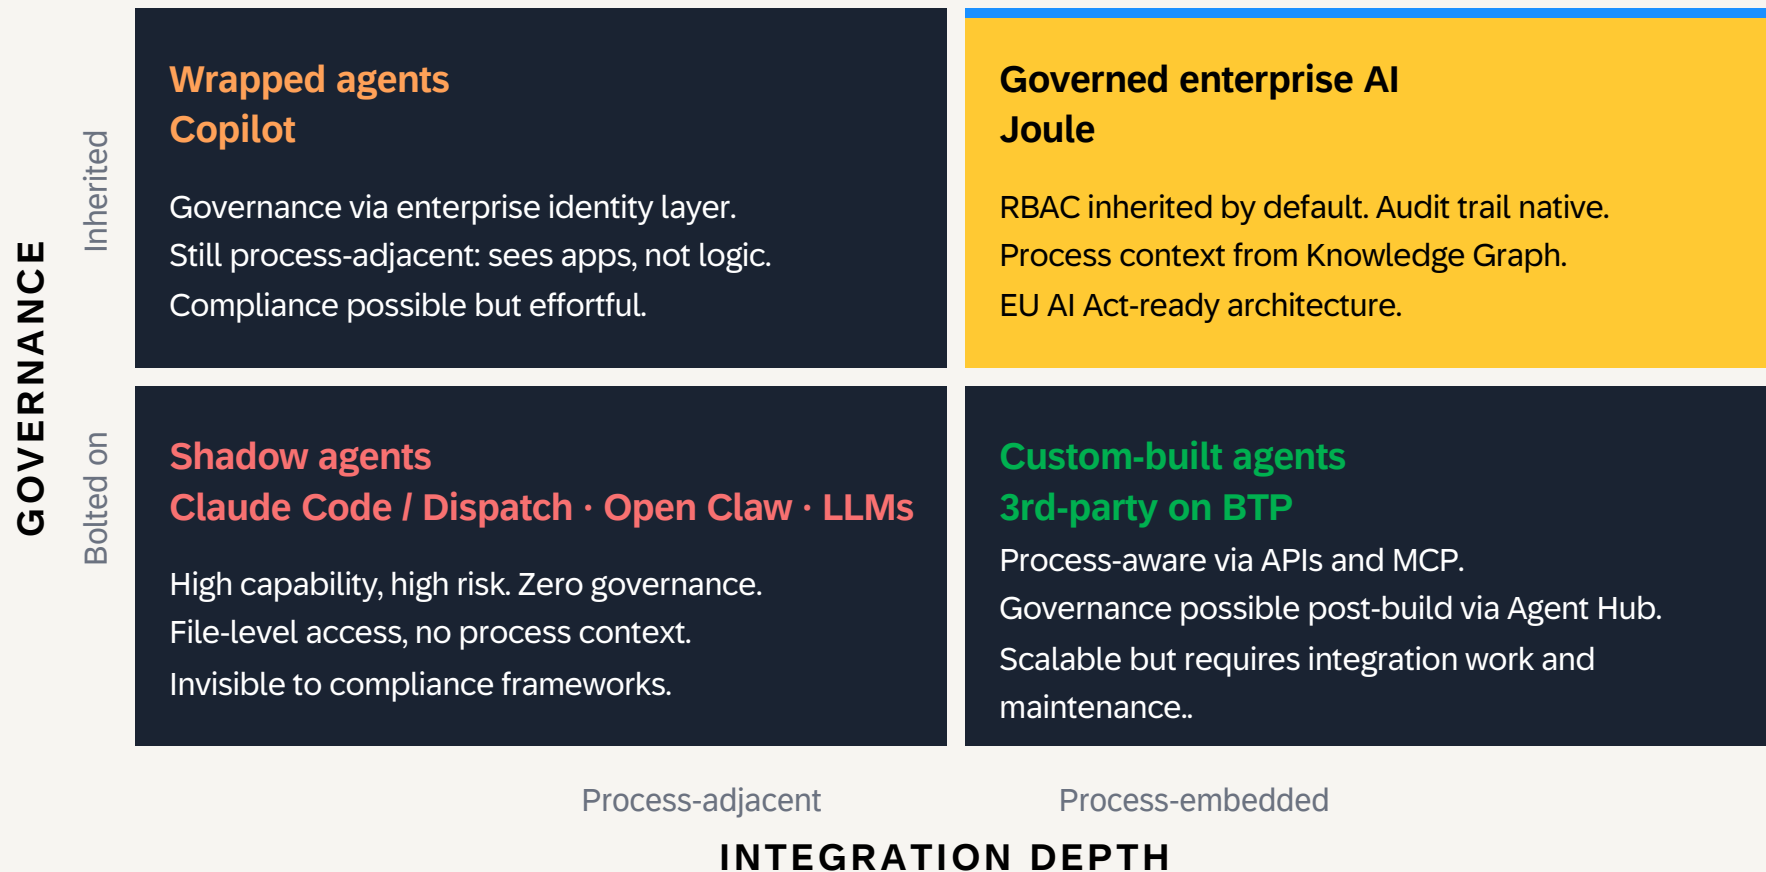
A detailed image of a bronze serpent, likely the Naga from the movie 'Indiana Jones and the Temple of Doom', coiled into a circle and eating its own tail. The serpent has a metallic, scaly texture and glowing orange eyes. A bright orange flame or lightning bolt is visible at the point where the head meets the tail. The text 'Agency eats truth' is written in a bold, yellow, sans-serif font across the center of the serpent's body.

Agency eats truth





# Where and how does context survive?





# What if you could keep the context embedded?



**Production & Operations  
Manager**



AI Assistant

Joule Agents

Joule Studio

Search

Q

Status:

Created X ; ▾

Issue Type:

No Filter X ▾

Delay Duration:

>=0 Hours

Order:

Material:

3 items

Created On:

Go

Adapt Filters (4)

Orders (11)

Change Dates and Quantity Edit Release Read Master Data Check Availability ▾ Related Apps ▾

| <input type="checkbox"/> | Order   | Material                                   | Open Quantity | Status  |  | Start                      | End                        | Progress of Operation                         | Issues | Created On         |
|--------------------------|---------|--|---------------|---------|--|----------------------------|----------------------------|---|--------|--------------------|
| <input type="checkbox"/> | 1003342 | MATS-E-BIKE<br>E-Bike - Finished           | 10 EA         | Created |  | Mon, Oct 13, 2025<br>07:00 | Mon, Oct 13, 2025<br>10:30 | Slag Processing (0010) <div><div></div></div> |        | Fri, Oct 3, 2025 > |
| <input type="checkbox"/> | 1003341 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 5 EA          | Created |  | Mon, Oct 13, 2025<br>11:00 | Mon, Oct 13, 2025<br>16:00 | Slag Processing (0010) <div><div></div></div> |        | Fri, Oct 3, 2025 > |
| <input type="checkbox"/> | 1003340 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 22 EA         | Created |  | Mon, Oct 13, 2025<br>15:00 | Mon, Oct 13, 2025<br>20:00 | Assembly (0010) <div><div></div></div>        |        | Mon, Oct 6, 2025 > |
| <input type="checkbox"/> | 1003345 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created |  | Mon, Oct 20, 2025<br>07:00 | Wed, Oct 22, 2025<br>11:30 | Slag Processing (0010) <div><div></div></div> |        | Mon, Oct 6, 2025 > |
| <input type="checkbox"/> | 1003347 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 15 EA         | Created |  | Mon, Oct 20, 2025<br>09:30 | Wed, Oct 22, 2025<br>13:30 | Slag Processing (0010) <div><div></div></div> |        | Tue, Oct 7, 2025 > |
| <input type="checkbox"/> | 1003348 | MATS-E-BIKE<br>E-Bike - Finished           | 15 EA         | Created |  | Mon, Oct 20, 2025<br>16:00 | Wed, Oct 22, 2025<br>14:00 | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |
| <input type="checkbox"/> | 1003350 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created |  | Thu, Oct 23, 2025<br>07:00 | Fri, Oct 24, 2025<br>10:30 | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |
| <input type="checkbox"/> | 1003352 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 15 EA         | Created |  | Thu, Oct 23, 2025<br>16:00 | Fri, Oct 24, 2025<br>11:30 | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |
| <input type="checkbox"/> | 1003358 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 20 EA         | Created |  | Thu, Oct 23, 2025<br>08:30 | Fri, Oct 24, 2025<br>09:30 | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |
| <input type="checkbox"/> | 1003361 | MATS-E-BIKE<br>E-Bike - Finished           | 15 EA         | Created |  | Mon, Oct 27, 2025<br>07:00 | Wed, Oct 29, 2025<br>9:30  | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |
| <input type="checkbox"/> | 1003361 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created |  | Mon, Oct 27, 2025<br>15:30 | Wed, Oct 29, 2025<br>17:00 | Slag Processing (0010) <div><div></div></div> |        | Wed, Oct 8, 2025 > |

Search

Q

Status:

Created ×

▾

Issue Type:

No Filter ×

▾

Delay Duration:

>=0 Hours

Order:

🔗

Material:

3 items

🔗

Created On:

📅

Go

Adapt Filters (4)

⬆

★

Orders (11)

Change Dates and Quantity

Edit

Release

Read Master Data

Check Availability ▾

Related Apps

⚙️

🔗 ▾

| <input type="checkbox"/> | Order   | Material                                   | Open Quantity | Status  |   | Start                      | End                        | Progress of Operation  | Issues    | Created On         |
|--------------------------|---------|--|---------------|---------|---|----------------------------|----------------------------|------------------------|-----------|--------------------|
| <input type="checkbox"/> | 1003342 | MATS-E-BIKE<br>E-Bike - Finished           | 10 EA         | Created | 🔄 | Mon, Oct 13, 2025<br>07:00 | Mon, Oct 13, 2025<br>10:30 | Slag Processing (0010) | 🕒 📦 📄 📊 🚚 | Fri, Oct 3, 2025 ▸ |
| <input type="checkbox"/> | 1003341 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 5 EA          | Created | 🔄 | Mon, Oct 13, 2025<br>11:00 | Mon, Oct 13, 2025<br>16:00 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003340 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 22 EA         | Created | 🔄 | Mon, Oct 13, 2025<br>15:00 | Mon, Oct 13, 2025<br>20:00 | Assembly (0010)        | 🕒         |                    |
| <input type="checkbox"/> | 1003345 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created | 🔄 | Mon, Oct 20, 2025<br>07:00 | Wed, Oct 22, 2025<br>11:30 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003347 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 15 EA         | Created | 🔄 | Mon, Oct 20, 2025<br>09:30 | Wed, Oct 22, 2025<br>13:30 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003348 | MATS-E-BIKE<br>E-Bike - Finished           | 15 EA         | Created | 🔄 | Mon, Oct 20, 2025<br>16:00 | Wed, Oct 22, 2025<br>14:00 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003350 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created | 🔄 | Thu, Oct 23, 2025<br>07:00 | Fri, Oct 24, 2025<br>10:30 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003352 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 15 EA         | Created | 🔄 | Thu, Oct 23, 2025<br>16:00 | Fri, Oct 24, 2025<br>11:30 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003358 | MATS-ROAD-BIKE<br>Road Bike - Finished     | 20 EA         | Created | 🔄 | Thu, Oct 23, 2025<br>08:30 | Fri, Oct 24, 2025<br>09:30 | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003361 | MATS-E-BIKE<br>E-Bike - Finished           | 15 EA         | Created | 🔄 | Mon, Oct 27, 2025<br>07:00 | Wed, Oct 29, 2025<br>9:30  | Slag Processing (0010) | 🕒         |                    |
| <input type="checkbox"/> | 1003361 | MATS-GRAVEL-BIKE<br>Gravel Bike - Finished | 10 EA         | Created | 🔄 | Mon, Oct 27, 2025<br>15:30 | Wed, Oct 29, 2025<br>17:00 | Slag Processing (0010) | 🕒         |                    |

Joule

Bike\_Road\_Tire 1 of production order 1003341. The planned start date could stay the same when using the suggested parts.

Alternative Part for component Bike\_Road\_Tire 1 (BRT151)

Alternative Details

Name:  
Bike\_Road\_Tire 3 (BRT153)

Bike Road Tire Quantity:  
2

Supplier:  
Continental

Measure:  
700c x 23mm

Choose and release

Cancel

Open in App

👍 🗨️ ⓘ

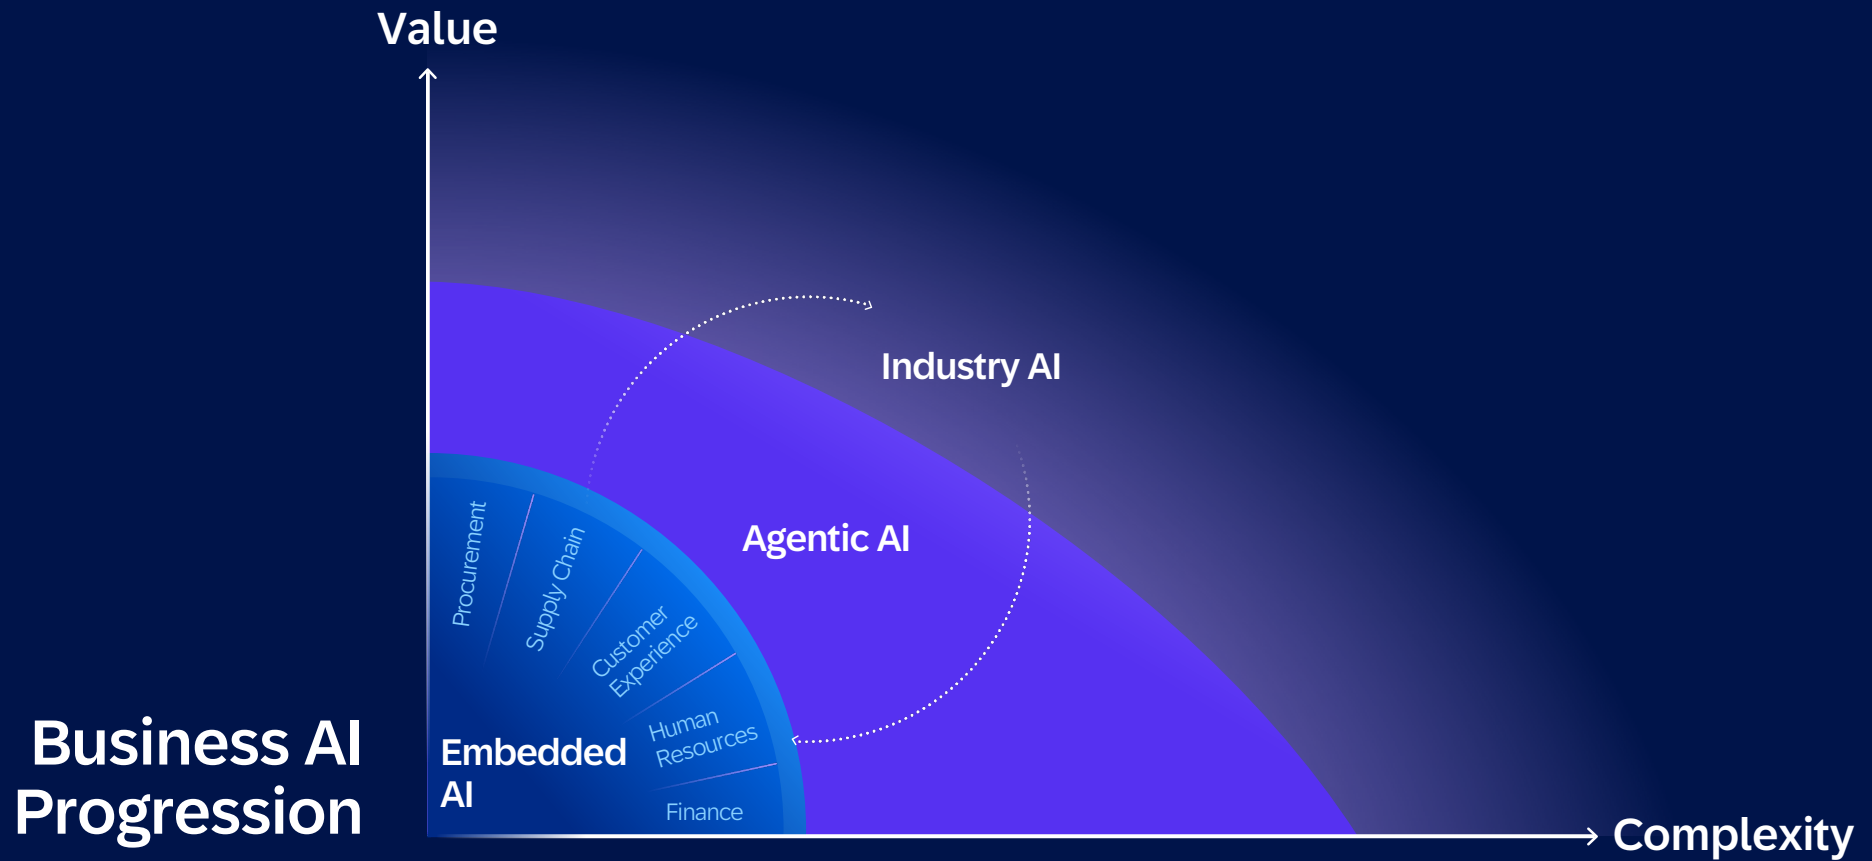
Message Joule...

🔍

Joule uses AI. Verify results.




# What if you could deepen the context













# ConTextTab: A Semantics-Aware Tabular In-Context Learner

Marco Spinaci<sup>\*1</sup> Marek Polewczyk<sup>\*2</sup> Maximilian Schambach<sup>\*2</sup> Sam Thelin<sup>2</sup>  
<sup>1</sup>SAP France <sup>2</sup>SAP SE  
{firstname.lastname}@sap.com

## Abstract

Tabular in-context learning (ICL) has recently achieved state-of-the-art (SOTA) performance on several tabular prediction tasks. Previously restricted to classification problems on small tables, recent advances such as TabPFN [18] and TabICL [30] have extended its use to larger datasets. Although current table-native ICL architectures are architecturally efficient and well-adapted to tabular data structures, their exclusive training on synthetic data limits their ability to fully leverage the rich semantics and world knowledge contained in real-world tabular data. At the other end of the spectrum, tabular ICL models based on pretrained large language models such as TabuLa-8B [12] integrate deep semantic understanding and world knowledge but are only able to make use of a small amount of context due to inherent architectural limitations. With the aim to combine the best of both these worlds, we introduce **ConTextTab**, integrating semantic understanding and alignment into a table-native ICL framework. By employing specialized embeddings for different data modalities and by training on large-scale real-world tabular data, our model is competitive with SOTA across a broad set of benchmarks while setting a new standard on the semantically rich CARTE benchmark. Code and model checkpoints are available at <https://github.com/marcospinaci/contexttab>.

## 1 Introduction

Tables with information stored in a structured format are a dominant data format in many



# RELATIONAL TRANSFORMER: TOWARD ZERO-SHOT FOUNDATION MODELS FOR RELATIONAL DATA

Rishabh Ranjan<sup>01\*</sup>, Valter Hudovertnik<sup>0</sup>, Mark Znidar<sup>0</sup>, Charilaos Kanatsoulis<sup>0</sup>,  
Roshan Upendra<sup>1</sup>, Mahmoud Mohammadi<sup>1</sup>, Joe Meyer<sup>1</sup>, Tom Palczewski<sup>1</sup>,  
Carlos Guestrin<sup>0</sup>, Jure Leskovec<sup>0</sup>  
<sup>0</sup>Stanford University, <sup>1</sup>SAP Labs LLC  
{ranjanr, guestrin, jure}@stanford.edu

## ABSTRACT

Pretrained transformers readily adapt to new sequence modeling tasks via zero-shot prompting, but relational domains still lack architectures that transfer across datasets and tasks. The core challenge is the diversity of relational data, with varying heterogeneous schemas, graph structures and functional dependencies. In this paper, we present the *Relational Transformer (RT)* architecture, which can be pretrained on diverse relational databases and directly applied to unseen datasets and tasks without task- or dataset-specific fine-tuning, or retrieval of in-context examples. RT (i) tokenizes cells with table/column metadata, (ii) is pretrained via masked token prediction, and (iii) utilizes a novel *Relational Attention* mechanism over columns, rows, and primary-foreign key links. Pretrained on RelBench datasets spanning tasks such as churn and sales forecasting, RT attains strong zero-shot performance, averaging 93% of fully supervised AUROC on binary classification tasks with a single forward pass of a 22M parameter model, as opposed to 84% for a 27B LLM. Fine-tuning yields state-of-the-art results with high sample efficiency. Our experiments show that RT’s zero-shot transfer harnesses task-table context, relational attention patterns and schema semantics. Overall, RT provides a practical path toward foundation models for relational data.

## 1 INTRODUCTION

Foundation models [3, 48] have emerged as a powerful paradigm in natural language processing (NLP) [7] and computer



# SAP-RPT-1



['ræpɪd]-[wʌn]



# Are you simulating context – or aware of it?

Old World – 1h

New World – 3 sec

```
# Install AutoGluon and its dependencies (might take a few minutes)
pip install autogluon==1.5.0
pip install autogluon.tabular[all],test
restart_python

Looking in indexes: https://pypi.org/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/lnworkbench-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/cdi-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/lnworkbench-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/proxy-deploy-releases-by-periscope-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/lnworkbench-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/cdi-pypi/simple, https://[REDACTED]:[REDACTED]common.repositories.cloud.sap/artifactory/api/pypi/proxy-deploy-releases-hyperspace-pypi/simple
Collecting autogluon==1.5.0
  Downloading autogluon-1.5.0-py3-none-any.whl.metadata (12 kB)
Collecting autogluon.core==1.5.0 (from autogluon.core[all]==1.5.0->autogluon==1.5.0)
  Downloading autogluon_core-1.5.0-py3-none-any.whl.metadata (13 kB)
Collecting autogluon.features==1.5.0 (from autogluon==1.5.0)
  Downloading autogluon_features-1.5.0-py3-none-any.whl.metadata (12 kB)
Collecting autogluon.tabular==1.5.0 (from autogluon.tabular[all]==1.5.0->autogluon==1.5.0)
  Downloading autogluon_tabular-1.5.0-py3-none-any.whl.metadata (16 kB)
Collecting autogluon.timeseries==1.5.0 (from autogluon==1.5.0)
  Downloading autogluon_timeseries-1.5.0-py3-none-any.whl.metadata (13 kB)
Collecting autogluon.multimodal==1.5.0 (from autogluon==1.5.0)
  Downloading autogluon_multimodal-1.5.0-py3-none-any.whl.metadata (13 kB)
Collecting autogluon.timeseries==1.5.0 (from autogluon.timeseries[all]==1.5.0->autogluon==1.5.0)
  Downloading autogluon_timeseries-1.5.0-py3-none-any.whl.metadata (13 kB)
Requirement already satisfied: numpy<1.17,>=1.5.4 in /databricks/python3/lib/python3.12/site-packages (from autogluon.core==1.5.0->autogluon.core[all]==1.5.0->autogluon==1.5.0) (12.1.3)
Requirement already satisfied: scipy<1.17,>=1.5.4 in /databricks/python3/lib/python3.12/site-packages (from autogluon.core==1.5.0->autogluon.core[all]==1.5.0->autogluon==1.5.0) (1.15.1)
Requirement already satisfied: scikit-learn==1.8.0 in /databricks/ovthn3/lib/ovthn3.12/site-packages (from autogluon.core==1.5.0->autogluon.core[all]==1.5.0->autogluon==1.5.0) (1.8.0)
```

**1. Constants and Configuration**

Next, we define our sample size, test set size, set PLANT as our prediction target, and identify lesser columns that wouldn't be available at prediction time.

```
POST /api/v1/predict
{
  "prediction_config": {
    "task_type": "classification",
    "task_name": "PLANT",
    "model_name": "PLANT",
    "task_type": "classification"
  },
  "task_type": "classification"
}

{
  "prediction": "PLANT",
  "task_type": "classification"
}
```

rpt\_context\_for\_playground.csv ⓘ

[Start Tour](#)
[Reset](#)
[Export](#)

Rows (2015)

[Set Predict](#)

✖

K

or

Shift

Click

+ Add Row

+ Add Column

Columns ▾

| <input type="checkbox"/> | SALESDOCUMENT ↑↓ | SALESDOCUMENTTYPE ↑↓ | SALESORGANIZATION ↑↓ | DISTRIBUTIONCHANNEL ↑↓ | ORGANIZATIONDIVISION ↑↓ | BILLINGCOMPANYCODE ↑↓ | TRANSACTIONCURRENCY ↑↓ | CREATIONDATE ↑↓ | CRE  | + |
|--------------------------|------------------|----------------------|----------------------|------------------------|-------------------------|-----------------------|------------------------|-----------------|------|---|
| <input type="checkbox"/> | 2231837          | TA                   | 4200                 | 10                     | 10                      | DK20                  | EUR                    | 2019-04-24      | 14:0 |   |
| <input type="checkbox"/> | 2223397          | ZMUN                 | 2500                 | 10                     | 10                      | CN10                  | CNY                    | 2019-04-09      | 08:4 |   |
| <input type="checkbox"/> | 2419778          | ZMUN                 | 1200                 | 10                     | 10                      | TR10                  | TRY                    | 2020-04-16      | 09:0 |   |
| <input type="checkbox"/> | 2066126          | ZMUN                 | 10                   | 10                     | 10                      | D011                  | EUR                    | 2018-06-04      | 11:4 |   |
| <input type="checkbox"/> | 2459116          | ZMUN                 | 400                  | 10                     | 10                      | UK10                  | EUR                    | 2020-06-29      | 17:3 |   |
| <input type="checkbox"/> | 2230823          | ZMUN                 | 10                   | 10                     | 10                      | D011                  | USD                    | 2019-04-23      | 11:5 |   |
| <input type="checkbox"/> | 2347601          | ZMUN                 | 4000                 | 10                     | 10                      | MX10                  | USD                    | 2019-11-29      | 19:5 |   |
| <input type="checkbox"/> | 2084783          | ZMUN                 | 10                   | 10                     | 10                      | D011                  | EUR                    | 2018-07-05      | 14:1 |   |
| <input type="checkbox"/> | 2031648          | ZICP                 | 10                   | 10                     | 10                      | D011                  | EUR                    | 2018-04-04      | 03:2 |   |
| <input type="checkbox"/> | 2073287          | ZMUN                 | 2200                 | 10                     | 10                      | D300                  | EUR                    | 2018-06-14      | 16:4 |   |
| <input type="checkbox"/> | 2014413          | ZMUN                 | 10                   | 10                     | 10                      | D011                  | EUR                    | 2018-03-01      | 09:0 |   |
| <input type="checkbox"/> | 2456011          | TA                   | 10                   | 10                     | 10                      | D011                  | EUR                    | 2020-06-24      | 08:3 |   |
| <input type="checkbox"/> | 2244467          | ZMUN                 | 10                   | 10                     | 10                      | D011                  | EUR                    | 2019-05-17      | 08:5 |   |
| <input type="checkbox"/> | 2143721          | ZMUN                 | 700                  | 10                     | 10                      | PL10                  | PLN                    | 2018-10-31      | 12:4 |   |

Showing 1 to 100 of 2015 rows

[Previous](#)

Page 1

of 21

[Next](#)
[Show 100 ▾](#)
[Predict](#)



## AI Agent Landscape

Save as ▾

⊗ Type: **AI Agent** ⊗ AND Quality Seal: **Approved** ⊗ AND Business Capabilities: **Customer Service OR Finance OR Procurement** ⊕

Add an Initiative ⊗

JUN

JUL

AUG

SEP

OCT

NOV

View

AI Agents ▾



Show suggested Joule Agents



Show suggested Applications

Layout



Sort



Level

1 ▾



100%

**A** AI Agents ☐ N/A ☒ SAP ☐ Non-SAP ☐ Custom ☐ Not in Use **A** Application ☐ Not in use ☒ In use**B** Customer Service

Catalog Optimization Agent

✓ SAP Intelligent...

✓ Call Center Mg...

✓ SAP Claims Mg...

Shopping Assistant Agent

✓ SAP Commerce

Customer Sentiment Agent

✓ Mailsnake

Self-Service Agent

Quote Creation Agent

✓ SAP Intelligent...

✓ SAP S/4HANA...

Response Drafting Agent

✓ Call Center Mg...

Knowledge Creation Agent

✓ SAP S/4HANA...

SAP SuccessFactors

Case Classification Agent

✓ SAP Claims Mg...

✓ SAP Intelligent...

Proactive Outreach Agent

✓ Mailsnake

✓ Call Center Mg...

Q&amp;A Agent

✓ SAP Commerce

**B** Finance

Accounts Receivable Agent

✓ SAP FSCM

Vendor Risk Agent

✓ LedgerLink

✓ SAP S/4HANA...

✓ SAP Analytics...

Dispute Resolution Agent

✓ SAP S/4HANA...

SAP Risk and Assu...

Compliance Monitoring Agent

✓ SAP GRC

✓ SAP S/4HANA...

✓ Bookend

✓ SAP Analytics...

Invoice Processing Agent

✓ Bookend

✓ InvoSync AI

✓ SAP S/4HANA...

SAP Global Trade...

✓ SAP S/4HANA...

**B** Procurement

Contract Intelligence Agent

✓ SAP Ariba

Supplier Risk &amp; Compliance Agent

✓ SAP Ariba

✓ SAP Fieldglass

Expense Report Validation Agent

✓ SAP Concur

✓ SAP S/4HANA...

Spend Insights Agent

✓ BuyWise

✓ SAP Ariba

Sourcing Agent

✓ SAP Ariba

✓ ServiceNow

✓ SAP Analytics...

✓ SAP Ariba

Sourcing Intelligence Agent

✓ SAP Ariba

✓ SAP Fieldglass

Meeting Location Planner Agent


✓ SAP Concur

✓ SAP Concur

✓ SAP S/4HANA...

# What about agents in the real world

**AGIBOT**



**ANYbotics**




**BOOSTER  
ROBOTICS**



**BostonDynamics**




**GALBOT**



**HUMANOID**



**UNITREE**



**NEURA  
ROBOTICS**





**Martur Fompak x Humanoid, EWM Pick & Place**  
2026 Q1 on-site pilot

**SAPPHIRE PREVIEW**



# The Nordics should be leading this journey

## TRUST CULTURES

01

High-trust societies understand governance is not a friction.

Good governance is the foundation for truth.

## DATA MATURITY

02

Early ERP use and digitisation means more data is still preserved here.

That is a real advantage when AI depends on context.

## COORDINATION INSTINCT

03

The Nordic model is coordination and consensus on truth.

That is exactly the muscle AI needs.

# What we have today

Available

# 2,400

Joule Skills



Available

# 350+

Embedded AI  
use cases

Available

# 40+

Joule Agents

Discovery Center

Your hub to discover  
SAP Business AI use  
cases and agents





# The map forward

**Open your  
gates**

**Joule in every LoB**

**Preserve your  
language**

**AI Operating Model on BDC**

**Build for the  
future**

**Design for AI on RISE**



*"Every civilisation has a bit of the **Athenian** and of the **Spartan** within it.*

*So do you and I.  
We decide who we let out."*

---

-- Johan Norberg, Peak Human

*Will you join the expedition?*



**David Pontoppidan**  
david.pontoppidan@sap.com