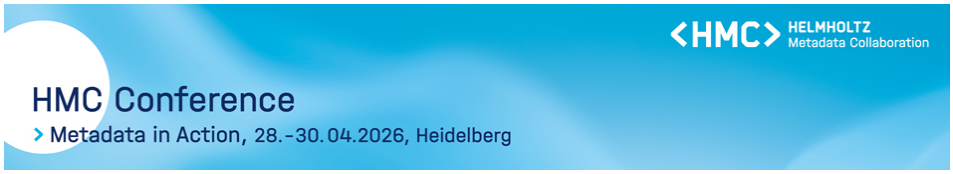

Beitrag ID: **74** Typ: **Talk**

The Agentic Automation Canvas: A Metadata Framework for Human-AI Task Delegation

Agentic AI systems—autonomous software driven by large language models (LLMs)—promise significant efficiency gains by performing tasks that traditionally required human judgment. However, their deployment fundamentally involves control inversion: humans must step back and allow the system to take command. The ease of building impressive prototypes with current LLMs creates a dangerous mismatch: stakeholders see quick demos and assume production-ready solutions are within reach, while the bulk of actual work—handling edge cases, ensuring reliability, integrating governance, and validating real-world performance—lies beyond the prototype. Without an explicit contract defining expectations before control inversion occurs, organizations face disillusionment when promised benefits fail to materialize.

We present the Agentic Automation Canvas (AAC), a structured metadata framework that captures the essential agreement between human stakeholders and agentic AI systems. By formalizing this as machine-readable metadata rather than traditional requirements documents, the AAC enables automated validation of stakeholder agreements, cross-system interoperability, and integration with institutional governance workflows. The canvas formalizes user requirements with quantified benefit expectations and balances them with developer feasibility assessments including model baseline capabilities, governance stages with assigned accountability, data access rights, and evaluation metrics for comparing outcomes against expectations. The AAC is implemented as an interactive web application (https://aac.slolab.ai) exporting versioned RO-Crate packages. Where possible, the schema maps to established vocabularies (Schema.org, PROV-O, DCAT, P-Plan, FRAPO, DUO); for agentic-specific concepts such as benefit metrics, baseline capabilities, and control inversion agreements, we introduce new terms under a registered https://w3id.org/aac/ namespace.

By requiring this contract before control inversion, the AAC bridges the gap between prototype enthusiasm and production reality. The resulting RO-Crate travels alongside the project as a machine-readable artifact designed to support governance, auditable decision-making, and benefit tracking throughout the collaboration lifecycle.

ONLY WORKSHOPS - Proposed interaction format

Alternative Track

1. From Harmonisation to Action(ability)

ONLY WORKSHOPS - Tentative audience

ONLY WORKSHOPS - Maximum number of participants

ONLY WORKSHOPS - Special technical requirements

Autor: LOBENTANZER, Sebastian (Helmholtz Munich)

Vortragende(r): LOBENTANZER, Sebastian (Helmholtz Munich)

Sitzung Einordnung: TALK SESSION

Track Klassifizierung: HMC Conference 2026 Track Topics: 4. Human-machine collaboration in (meta)data acquisition