L6 Data Agents: From Individual Intelligence to Collective Autonomy

Dimas Cassimiro Nascimento[1], **Daliton da Silva**[1]

[1]Federal University of Agreste of Pernambuco (UFAPE)
Laboratory of Applied Data Intelligence (LIDA)
Garanhuns – PE – Brazil

{dimas.cassimiro,daliton.silva}@ufape.edu.br

Abstract. *We introduce Level 6 (L6) Data Agents, a framework in which collective autonomy emerges from structured interactions among agents under an adaptive algorithmic protocol. The framework incorporates a reputation-driven mechanism, where agents accumulate scores based on the quality and impact of their contributions, guiding selection and propagation. We also propose an iterative algorithm that combines evaluation, reputation updating, and adaptation, through which innovation emerges as a collective phenomenon. Additionally, a dynamic consensus threshold adapts to system performance, balancing novelty search and refinement and enabling robust and scalable collective autonomy.*

1. Introduction

Recent advances in large language models and data-centric AI have led to the emergence of data agents as a new paradigm for automating data science workflows, including pipeline construction, model development, and analytical reasoning [Alva and Pandey 2026, Sun et al. 2025a, Tang et al. 2025, Wang et al. 2025, Zhu et al. 2025]. These agents increasingly exhibit capabilities associated with autonomous data scientists [de Miranda and Campelo 2024], combining reasoning, tool usage, and interaction with complex data environments [Sun et al. 2026, Sun et al. 2025b, Zhou et al. 2025]. As a result, contemporary frameworks describe agents as autonomous entities capable of solving tasks, generating artifacts, and adapting their behavior.

In Figure 1, we illustrate the hierarchical spectrum of data agent autonomy from L0 to L5 [Zhu et al. 2025], showing the progressive increase in individual capabilities. Level 6 (L6) is highlighted separately to emphasize its distinctive feature: collective autonomy. Unlike previous levels, L6 agents interact, evaluate each other, and evolve through structured feedback, incorporating validation, reputation dynamics, and adaptive deployment to achieve system-level self-governance.

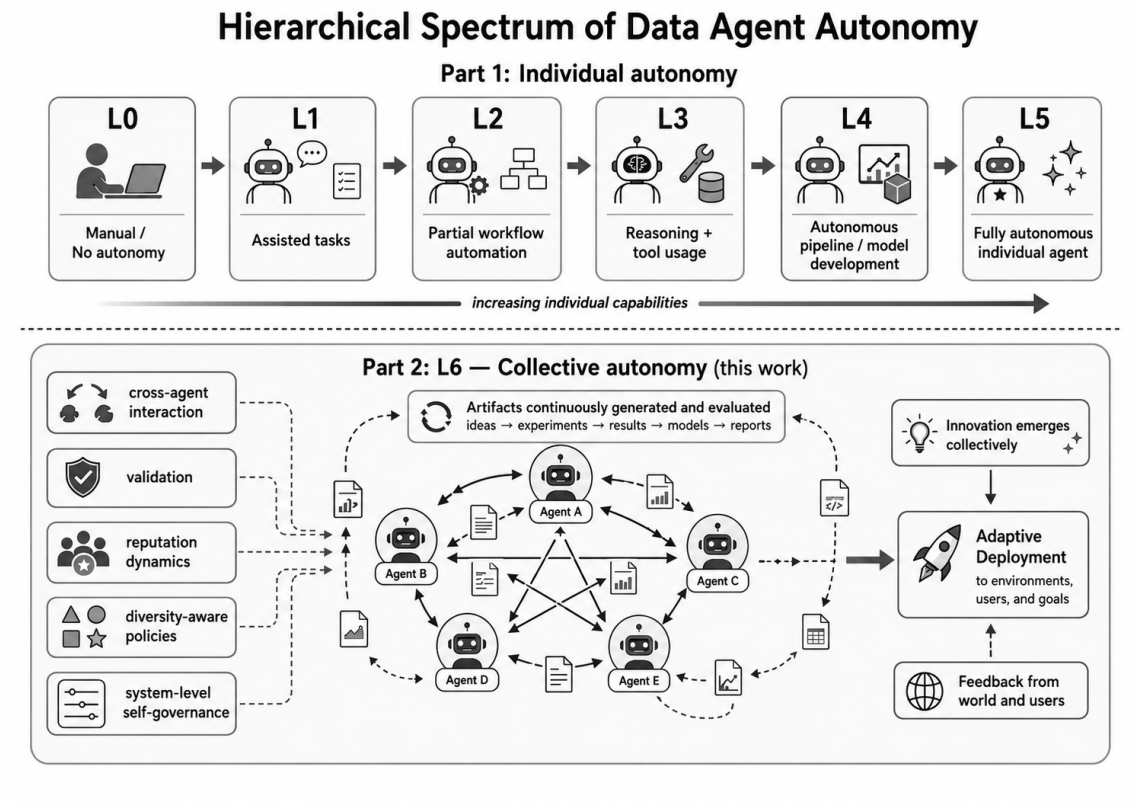


Figure 1. Hierarchical spectrum of data agent autonomy (L0–L5) and the emergence of L6 collective autonomy. While L0–L5 focus on increasing individual capabilities, L6 (this work) introduces cross-agent interactions, validation, and adaptive deployment as a collective process.

Despite these advances, existing approaches remain fundamentally centered on individual autonomy. Even at the highest levels of autonomy discussed in recent taxonomies, such as L5, agents are modeled as fully autonomous yet isolated decision-makers [Zhu et al. 2025]. This individualistic perspective becomes limiting in realistic settings where multiple agents simultaneously generate competing artifacts, requiring mechanisms for validation, coordination, and selection that go beyond single-agent capabilities. Without such mechanisms, systems risk deploying low-quality outputs or ignoring potentially valuable innovations due to the absence of shared evaluation and governance structures.

At the same time, recent work highlights the growing complexity of data ecosystems, where discovery, integration, and analysis are inherently collaborative and iterative processes [Fu et al. 2025, Piccialli et al. 2025]. These observations suggest that autonomy should not be viewed solely as a property of individual agents, but as an emergent property of interactions within a system of agents. In this work, we argue that autonomy must be extended beyond the individual level toward a collective dimension in which agents interact, evaluate each other, and evolve through structured feedback. In this perspective, the system itself becomes the locus of autonomy. Drawing inspiration from the hierarchical taxonomy introduced in [Zhu et al. 2025], we introduce Level 6 (L6) as a formal framework in which discovery and innovation are not isolated capabilities of individual agents, but are embedded in a collective process governed by validation, reputation,

and diversity-aware policies.

The proposed L6 framework models multi-agent systems as adaptive ecosystems in which artifacts are continuously generated and evaluated, enabling innovation to emerge as a collective phenomenon, shaped by interactions among agents rather than by isolated decisions. In this context, the contribution of this work is to establish the conceptual and formal foundations of L6, specifying the entities, variables, interaction mechanisms, and governance principles required for collective autonomy, thereby providing a basis for future implementation and empirical validation.

2. Related Work

Our work builds upon and extends the foundational taxonomy for data agents introduced by Zhu et al. [Zhu et al. 2025], who proposed a hierarchical framework from L0 to L5 that delineates progressive autonomy in data-related tasks. This taxonomy provides a systematic lens for classifying data agents based on their degree of independence, from manual operations to fully autonomous systems.

Prior surveys have examined the intersection of large language models and data-related tasks from various perspectives. Zhou et al. [Zhou et al. 2025] provide a comprehensive overview of LLM applications across the data lifecycle, covering areas such as configuration tuning, query optimization, and data cleaning. Tang et al. [Tang et al. 2025] focus specifically on LLM-based and agent-based approaches for data analysis, surveying systems that function as data analysts. Wang et al. [Wang et al. 2025] examine data science agents, analyzing their architectures and capabilities across different tasks. These perspectives are complemented by more recent published surveys, such as Sun et al. [Sun et al. 2025a], which review LLM-based agents for statistics and data science, and Alva and Pandey [Alva and Pandey 2026], who analyze agentic AI systems from an architectural and systems perspective.

Within the specific domain of autonomous data agents, several recent works have proposed architectures and frameworks that align with our L6 formulation. Sun et al. [Sun et al. 2025b] introduce a holistic architecture for orchestrating Data+AI ecosystems, while Sun et al. [Sun et al. 2026] propose a large-model-based data agent integrating multiple stages of the data lifecycle. Fu et al. [Fu et al. 2025] discuss the concept of autonomous data agents and their potential to transform smart data applications. From a broader perspective, Piccialli et al. [Piccialli et al. 2025] highlight the importance of interaction and coordination in autonomous multi-agent systems.

Beyond data-agent-specific surveys, the broader literature on LLM-based agents provides useful foundations for understanding the transition from individual autonomy to collective autonomy. Wang et al. [Wang et al. 2024] present a systematic survey of LLM-based autonomous agents, emphasizing key components such as profiling, memory, planning, and action. Similarly, Xi et al. [Xi et al. 2025] discuss the rise of LLM-based agents from the perspective of agent construction, application scenarios, and agent societies. These works are relevant to our proposal because L6 autonomy depends not only on task execution, but also on persistent interaction, adaptive behavior, and coordination among agents operating within a shared environment.

Several published works have also investigated mechanisms that enable agents to reason, act, and improve through feedback. ReAct [Yao et al. 2023] introduced a

paradigm in which language models interleave reasoning traces and actions, allowing agents to interact with external environments while maintaining interpretable decision processes. Reflexion [Shinn et al. 2023] extended this direction by showing how language agents can use verbal feedback and episodic memory to improve behavior across trials without updating model parameters. Voyager [Wang et al. 2023] further demonstrated open-ended lifelong learning through an agent that incrementally builds a reusable skill library from interaction with an environment. These works are not limited to data science, but they provide important conceptual foundations for L6 systems, especially regarding iterative improvement, feedback incorporation, and persistent adaptation.

Recent published systems also show how LLM agents can be specialized for data-intensive workflows. Data Interpreter [Hong et al. 2025] proposes an LLM-based agent for data science that uses hierarchical graph modeling and progressive verification to manage complex analytical tasks. This work is particularly related to our formulation because it highlights the need for decomposition, verification, and refinement in data science pipelines. In the broader data management context, Fernandez et al. [Fernandez et al. 2023] argue that large language models may transform data management tasks such as data integration, schema matching, query interfaces, and entity resolution. These contributions reinforce the idea that autonomous data agents should be evaluated not only as isolated assistants, but as components of larger data ecosystems in which correctness, reliability, and downstream utility are central concerns.

The multi-agent systems literature further motivates our emphasis on collective validation and governance. ChatDev [Qian et al. 2024] models software development as a communicative process among specialized agents responsible for design, coding, and testing. In turn, MetaGPT [Hong et al. 2024] introduces a multi-agent framework based on standardized operating procedures, assigning structured roles to agents to improve coordination and reduce inconsistencies in collaborative workflows. These works demonstrate the practical value of multi-agent collaboration, but they do not fully formalize how agents should validate each other, how artifact-level quality should affect long-term reputation, or how deployment decisions should adapt over time.

In this sense, our work complements existing research by shifting the focus from agent orchestration to collective autonomy. While prior systems emphasize task decomposition, tool use, role specialization, and conversational collaboration, our L6 formulation introduces an explicit mechanism for inter-agent validation, innovation-aware scoring, adaptive reputation, and consensus-based deployment. This allows the system to treat agents not merely as coordinated workers, but as members of a self-governing collective capable of evaluating artifacts, updating trust, and adapting deployment standards based on both peer consensus and post-deployment utility.

3. Proposed System Definition (L6)

We define a Level 6 system as $\mathcal{S}_{L6} = (\mathbb{A}, \mathcal{X}, V, R, \Pi, \phi, \delta^{(t)})$, where $\mathbb{A}$ is the set of agents and $\mathcal{X}$ is the space of artifacts, including pipelines, models, and analytical outputs. Here, Π denotes the deployment policy and $\delta^{(t)}$ the adaptive consensus threshold. In this system, autonomy is not localized within individual agents but emerges from the interaction among discovery, innovation, validation, and selection processes. The collective behavior induces a closed feedback loop: an artifact x is transformed via innovation into $\phi(x, \mathbf{c})$,

which is evaluated through V; evaluations update reputations R; reputations and innovation gains shape the deployment policy Π; and the deployed artifacts influence future tasks, data, and contextual signals. This loop governs the evolution of artifacts and agent behavior over time, with the threshold $\delta^{(t)}$ adapting to maintain alignment between system selectivity and collective capability. In this setting, collective autonomy refers to the capacity of the multi-agent system to evaluate, select, deploy, and adapt its behavior through system-level interactions without centralized control, distinguishing it from collective intelligence, which primarily concerns a group's ability to solve problems effectively. Individual autonomy, in turn, concerns what an agent can accomplish independently, whereas collective autonomy concerns what the system can govern and adapt through interactions among agents.

The framework formalizes autonomy as a two-dimensional construct: individual autonomy, captured by the L0–L5 hierarchy, and collective autonomy, which emerges at L6 from system-level interactions. In this setting, discovery, innovation, validation, selection, and evolution operate as integrated, feedback-driven processes. Thereby, innovation becomes a collective phenomenon shaped by continuous feedback, reputation-mediated evaluation, and adaptive governance, reframing the goal of autonomous data agents from isolated intelligence to coordinated, self-improving ecosystems.

4. Problem Formulation

Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n\}$ be a set of agents, and let $\mathcal{X}$ denote the space of artifacts, including pipelines, models, analytical outputs, and other data-related products. Each agent operates over this space through two fundamental operators: discovery and innovation. An individual agent is defined as a functional entity $\mathcal{A}_i : (\tau_i, \mathcal{D}_i, \mathcal{E}_i, \mathcal{C}) \to \mathcal{X}$, where τ_i is a task description, $\mathcal{D}_i$ represents accessible data, $\mathcal{E}_i$ denotes the environment (e.g., data lakes, code interpreters), and $\mathcal{C}$ encapsulates contextual signals such as feedback from other agents or system-level state.

The discovery operator generates artifacts grounded in existing knowledge: $\text{Discover}_{\mathcal{A}_i} : (\tau_i, \mathcal{D}_i, \mathcal{E}_i) \to x_i$, where $x_i \in \mathcal{X}$. The innovation operator transforms an existing artifact into a new one by leveraging contextual signals: $\text{Innovate}_{\mathcal{A}_i} : \mathcal{X} \times \mathcal{C} \to \mathcal{X}$, so that $x_i' = \text{Innovate}_{\mathcal{A}_i}(x_i, \mathbf{c})$ for some $\mathbf{c} \in \mathcal{C}$. Innovation is modeled through a transformation function $\phi : \mathcal{X} \times \mathcal{C} \to \mathcal{X}$, yielding $x_i' = \phi(x_i, \mathbf{c})$. This formulation establishes a fundamental distinction: discovery produces artifacts from task and data, while innovation introduces transformations that generate novelty.

Problem Definition. The problem is to design a multi-agent system of *data agents* responsible for data-driven discovery, innovation, evaluation, and deployment, aiming to maximize artifact quality and utility while dynamically adapting governance thresholds.

5. Inter-Agent Validation

In L6 systems, artifacts must be evaluated collectively to ensure correctness, utility, and novelty. We define a validation function $V : \mathbb{A} \times \mathcal{X} \to [0, 1]$, where $V(\mathcal{A}_j, x_i)$ represents the evaluation assigned by agent $\mathcal{A}_j$ to artifact x_i. This evaluation extends naturally to innovated artifacts as $V(\mathcal{A}_j, x_i')$. Through this mechanism, the system evaluates both discovered and transformed artifacts, enabling a unified assessment that integrates quality and innovation. The consensus score for each artifact is computed as a weighted average

of other agents' evaluations, $\text{cons}(x) = \frac{\sum_{\mathcal{A}_j \neq \mathcal{A}_i} R_j^{(t)} V(\mathcal{A}_j, x)}{\sum_{\mathcal{A}_j \neq \mathcal{A}_i} R_j^{(t)}}$, where weights correspond to the agents' current reputations $R_j^{(t)}$ and the author of the artifact is excluded to avoid self-bias.

Innovation Quantification. To explicitly capture the contribution of innovation, we define the innovation gain as $\Delta_\phi(x_i) = \text{cons}(\phi(x_i, \mathbf{c})) - \text{cons}(x_i)$, where $\mathbf{c} \in \mathcal{C}$ denotes the contextual signals guiding the transformation, thereby reflecting its impact on the collective evaluation of an artifact. A positive value indicates improvement in consensus, while a negative value suggests degradation. This formulation isolates the marginal contribution of the transformation, decoupling innovation from the absolute quality of the artifact. Also, this signal is complemented by human feedback $F_\phi(x)$, capturing aspects of novelty and usefulness not fully reflected by peer evaluations. The composite artifact score $s(x)$ integrates consensus, human-assessed utility, and both innovation signals into a unified metric for reputation updates and deployment decisions, enabling the system to distinguish between incremental and transformative innovations.

Reputation Dynamics. Each agent is associated with a reputation score $R : \mathbb{A} \to \mathbb{R}^+$, which reflects its performance over time based on the quality of the artifacts it produces. Rather than explicitly decomposing reputation into separate components, we adopt an artifact-centric view in which each produced artifact contributes to the agent's evaluation through a composite score.

Each artifact x contributes a score $s(x) = \alpha \cdot \text{mean}(\mathbf{v}(x)) + \beta \cdot u + \gamma \cdot (\Delta_\phi(x) + F_\phi(x)) \cdot 2^{-1}$, combining the consensus of other agents' evaluations, human-assessed utility u, and innovation signals derived from both consensus ($\Delta_\phi(x)$) and human feedback (F_ϕ). We assume that $\alpha, \beta, \gamma \in [0, 1]$ are configurable parameters and that $\Delta_\phi(x)$ and $f_\phi = F_\phi(x)$ are defined on the same scale, which allows their direct aggregation. The agent's reputation at iteration t is computed as the average score over its generated artifacts, $\hat{R}_i^{(t)} = \frac{1}{|\mathcal{X}_i^{(t)}|} \sum_{x \in \mathcal{X}_i^{(t)}} s(x)$, and updated via an exponential moving average, $R_i^{(t+1)} = (1 - \lambda) R_i^{(t)} + \lambda \hat{R}_i^{(t)}$. This formulation ensures that agents are rewarded for producing high-quality artifacts, generating impactful innovations, and aligning with collective evaluations.

Consensus-Based Deployment with Adaptive Threshold. The system determines which artifacts to deploy through a consensus mechanism that aggregates evaluations weighted by reputation. An artifact x_i is selected for deployment if $\text{cons}(x_i) \geq \delta$ (a consensus threshold). Rather than remaining static, this threshold evolves over time based on system-wide performance. Let $\delta^{(t)}$ denote the threshold at iteration t. After each deployment cycle, the system computes the average consensus of deployed artifacts $\overline{C}^{(t)}$ and the average post-deployment utility $\overline{U}^{(t)}$ derived from human feedback or downstream task performance. The threshold is updated according to $\delta^{(t+1)} = \delta^{(t)} + \rho(\overline{C}^{(t)} - \delta^{(t)}) + \nu(\overline{U}^{(t)} - U_{\text{target}})$, where ρ and ν are learning rates, and U_{target} is a desired utility level. Only deployed artifacts $\mathcal{D}^{(t)}$ contribute to these averages, ensuring that the threshold adapts based on actual system performance. This mechanism allows the system to raise standards when performance is high and lower them when exploration is needed, ensuring that governance adapts to the system's evolving capabilities.

6. Consensus-Based Deployment with Adaptive Reputation

In this section, we present the L6-Consensus algorithm, an artifact-centric mechanism for managing deployments in a multi-agent system (Algorithm 1). At each time step, agents generate individual sets of artifacts (line 2), which are aggregated into a single candidate set (line 3). An empty set is initialized to hold deployed artifacts (line 4). Each artifact is then evaluated: the author is identified (line 6), votes from other agents are collected (line 7), and a consensus score is computed as a reputation-weighted average of these votes, excluding the author's own reputation to mitigate self-bias (line 8).

Deployment decisions are made based on the consensus score (line 9). Artifacts meeting or exceeding the threshold are deployed (line 10), at which point human utility and human innovation feedback are recorded (line 11), and the artifact is added to the deployed set (line 12). Artifacts below the threshold are not deployed, with human utility, human innovation feedback, and innovation gain all set to zero (line 14). A combined score integrating peer votes, human utility, innovation gain, and human innovation feedback is then calculated for each artifact (line 16), serving as the basis for reputation updates.

Reputation updates occur for each agent (line 18). If an agent produced artifacts, its estimated reputation is the average of their combined scores (line 20); otherwise, it remains unchanged (line 22). The reputation is updated via an exponential moving average (line 24), ensuring gradual changes over time. Finally, the deployment threshold is adapted for the next round (line 27).

7. Practical Considerations

The L6 system can be interpreted as a coupled stochastic dynamical system in which threshold adaptation, reputation updates, and deployment decisions evolve jointly. At each iteration, agents generate artifacts, evaluate the artifacts produced by others, receive feedback through deployment outcomes, and update their reputations accordingly. These components are not independent: changes in reputation affect future consensus scores, consensus scores determine which artifacts are deployed, and deployed artifacts provide the utility and innovation signals that influence both reputation and threshold adaptation. As a result, the system forms a closed feedback loop in which collective evaluation, individual performance, and governance criteria co-evolve over time. This feedback structure provides the operational basis for collective autonomy, which is evidenced when deployment, reputation, and governance decisions emerge from inter-agent interactions and adapt over time without being determined by a single agent. Its degree can be assessed through system-level indicators such as consensus quality, deployment utility, reputation adaptation, and threshold stability.

Under mild stationarity assumptions, such as bounded evaluation scores, slowly varying artifact quality distributions, and stable feedback mechanisms, the system admits fixed points in expectation. In practical terms, this means that if the distribution of generated artifacts and the behavior of evaluators remain sufficiently stable, reputations and deployment thresholds tend to converge toward values that reflect the long-term contribution of each agent and the expected quality of deployable artifacts. However, convergence should not be interpreted as complete stabilization of behavior. Since agents may continue to generate novel artifacts and since innovation signals can perturb consensus and

Algorithm 1: L6-Consensus: Artifact-Centric Deployment with Adaptive Reputation

Input: Agents $\mathbb{A}$, reputations $\{R_i^{(t)}\}$, threshold $\delta^{(t)}$
Output: Updated reputations and deployed artifacts

for $t = 1, 2, \ldots$ **do**
  Generate $\mathcal{X}_i^{(t)}$ for each $\mathcal{A}_i$
  $\mathcal{X}^{(t)} \leftarrow \bigcup_{\mathcal{A}_i \in \mathbb{A}} \mathcal{X}_i^{(t)}$
  $\mathcal{D}^{(t)} \leftarrow \emptyset$ `// set of deployed artifacts`
  foreach $x \in \mathcal{X}^{(t)}$ **do**
    $i \leftarrow \text{author}(x)$
    $\mathbf{v}(x) \leftarrow \{V(\mathcal{A}_j, x) : \mathcal{A}_j \in \mathbb{A} \setminus \{\mathcal{A}_i\}\}$
    $\text{cons}(x) \leftarrow \frac{\sum_{\mathcal{A}_j \neq \mathcal{A}_i} R_j^{(t)} V(\mathcal{A}_j, x)}{\sum_{\mathcal{A}_j \neq \mathcal{A}_i} R_j^{(t)}}$
    if $cons(x) \geq \delta^{(t)}$ **then**
      Deploy x
      `// human utility and innovation assessed by humans`
      $u, f_\phi \leftarrow F_u(x), F_\phi(x)$
      $\mathcal{D}^{(t)} \leftarrow \mathcal{D}^{(t)} \cup \{x\}$
    else
      $u, f_\phi, \Delta_\phi(x) \leftarrow 0, 0, 0$
    end
    `// based on the proposed innovation quantification and reputation dynamics`
    $s(x) \leftarrow \alpha \cdot \text{mean}(\mathbf{v}(x)) + \beta \cdot u + \gamma \cdot (\Delta_\phi(x) + f_\phi) \cdot 2^{-1}$
  end
  foreach $\mathcal{A}_i \in \mathbb{A}$ **do**
    if $\mathcal{X}_i^{(t)} \neq \emptyset$ **then**
      $\hat{R}_i^{(t)} \leftarrow \frac{1}{|\mathcal{X}_i^{(t)}|} \sum_{x \in \mathcal{X}_i^{(t)}} s(x)$
    else
      $\hat{R}_i^{(t)} \leftarrow R_i^{(t)}$
    end
    $R_i^{(t+1)} \leftarrow (1 - \lambda) R_i^{(t)} + \lambda \hat{R}_i^{(t)}$
  end
  `// Compute the averages of` C `and` U `over the dataset at iteration` t
  Compute $\overline{C}^{(t)}$ and $\overline{U}^{(t)}$ over $\mathcal{D}^{(t)}$
  `// Update` δ `based on the average` C `and the difference between` U `and the target value`
  $\delta^{(t+1)} \leftarrow \delta^{(t)} + \rho(\overline{C}^{(t)} - \delta^{(t)}) + \nu(\overline{U}^{(t)} - U_{\text{target}})$
end

utility outcomes, the system preserves a degree of persistent exploration. This property is important in L6 settings, where the objective is not merely to identify consistently reliable agents, but also to maintain the capacity to discover and validate new solutions.

The interaction between innovation gains, consensus, and adaptive governance induces a feedback loop that balances exploitation and exploration. Exploitation occurs when artifacts produced by reputable agents receive high consensus scores and are more likely to be deployed. Exploration, in turn, is encouraged by the explicit inclusion of innovation gains and human-assessed novelty or usefulness in the artifact score. Thus, an agent is not rewarded solely for producing artifacts that conform to existing collective expectations, but may also improve its reputation by generating transformations that increase consensus or demonstrate practical value after deployment. This prevents the system from becoming excessively conservative and allows transformative artifacts to influence long-term reputation dynamics.

A central practical advantage of the proposed formulation is the explicit separation between artifact generation and artifact evaluation. Agents are responsible for producing artifacts, but the assessment of these artifacts is performed by other agents through the validation function. By excluding the author from validation and consensus computation, the mechanism reduces self-reinforcement, discourages inflated self-assessment, and promotes independent collective judgment. This design is particularly relevant in multi-agent environments where agents may specialize in different types of reasoning, data processing, modeling, or artifact generation. In such settings, collective validation allows the system to aggregate complementary perspectives while reducing dependence on a single evaluator.

Another important practical consideration concerns the role of reputation in consensus. Since votes are weighted by the current reputation of the evaluating agents, the system gives greater influence to agents with a stronger history of producing valuable artifacts. This creates an adaptive trust mechanism: agents that consistently contribute high-quality, useful, and innovative artifacts become more influential in future validation rounds. At the same time, because reputation is updated through an exponential moving average, sudden fluctuations are smoothed over time. This avoids excessive sensitivity to isolated artifacts while still allowing reputations to adapt when an agent's performance improves or deteriorates.

The use of human feedback only after deployment establishes a clear causal pipeline from collective validation to real-world utility. Before deployment, artifacts are assessed through inter-agent consensus. After deployment, external feedback provides evidence about their practical usefulness and innovation value. This distinction is important because peer consensus and real-world utility may capture different aspects of quality. An artifact may be theoretically promising but less useful in practice, or it may receive moderate consensus while proving valuable after deployment. By incorporating post-deployment feedback into the composite score, the system can correct purely internal assessments and align reputation updates with observed utility.

The adaptive threshold also plays an important governance role. A fixed deployment threshold may be too rigid in dynamic environments: if set too high, it may prevent useful exploration; if set too low, it may allow low-quality artifacts to be deployed. The

proposed update rule adjusts the threshold according to the average consensus of deployed artifacts and their post-deployment utility. When deployed artifacts consistently achieve high consensus and strong utility, the system can raise its deployment standards. Conversely, when utility falls below the desired target or when exploration becomes necessary, the threshold can be relaxed. This mechanism allows the system to adapt its governance policy to the observed performance of the multi-agent collective.

In practical implementations, the parameters α, β, and γ should be selected according to the intended governance profile of the system. A higher value of α emphasizes peer consensus and favors artifacts that are broadly accepted by the agent collective. Alternatively, a higher value of β gives more importance to human-assessed utility and is appropriate when deployment outcomes are the primary source of validation. In turn, a higher value of γ increases the role of innovation, encouraging agents to propose artifacts that transform existing solutions rather than merely reproduce high-consensus patterns. Similarly, the learning rate λ controls how quickly reputations react to new evidence, while ρ and ν regulate the sensitivity of the deployment threshold to consensus and utility signals.

Care must also be taken when interpreting innovation scores. The innovation gain $\Delta_\phi(x)$ captures the marginal improvement produced by a transformation in terms of consensus, while $F_\phi(x)$ reflects human-perceived novelty or usefulness. These two signals are complementary, but they may not always agree. A transformation may increase peer consensus without being perceived as novel by humans, or it may be considered novel and useful despite receiving limited initial support from other agents. For this reason, the composite score combines both signals rather than relying on a single measure of innovation. This design makes the system more robust to cases in which internal validation and external judgment diverge.

Finally, the proposed mechanism supports a scalable and modular implementation of L6 collective autonomy. Artifact generation, inter-agent validation, human feedback collection, reputation updating, and threshold adaptation can be implemented as separate modules connected through a shared artifact registry. This modularity allows the system to incorporate different types of agents, validation functions, feedback sources, and deployment policies without changing the core logic of the framework. Overall, the formulation provides a practical basis for governing autonomous multi-agent systems in which quality control, innovation assessment, reputation dynamics, and deployment decisions are integrated into a unified adaptive process.

8. Conclusions and Future Works

We introduced a formal framework for L6 data agents that integrates discovery, innovation, and human-assessed utility into a collective system governed by validation, reputation, and adaptive deployment. Artifacts are evaluated through a reputation-weighted consensus mechanism, and each agent's contributions are summarized in a composite score $s(x)$ that combines consensus, utility, and innovation gain. By coupling artifact-level evaluation with agent-level reputation dynamics, the proposed framework provides a structured way to regulate deployment decisions while continuously updating the influence of each agent according to its observed contributions. This perspective extends the existing autonomy hierarchy by adding a collective dimension, redefining data agent au-

tonomy as a transition from individual mastery to coordinated, self-governing intelligence that actively balances utility, reputation, and innovation.

The proposed formulation also emphasizes that autonomy at L6 is not merely a matter of increasing the capabilities of isolated agents, but of designing mechanisms through which agents can evaluate one another, incorporate external feedback, and adapt collective standards over time. In this sense, consensus, reputation, and adaptive thresholds act as governance components that support both quality control and exploration. This allows the system to reward reliable contributions while preserving space for novel artifacts whose value may only become evident after deployment.

Future work includes empirical validation of the L6 framework, exploring alternative composite scoring functions, adaptive learning rates, mechanisms for agent diversity, safeguards against reputation concentration or excessive consensus bias, and robustness to adversarial behavior and unreliable feedback.

References

Alva, L. R. and Pandey, B. (2026). Agentic ai systems in the age of generative models: Architectures, cloud scalability, and real-world applications. *Artificial Intelligence Review*, 59:88.

de Miranda, B. A. and Campelo, C. E. C. (2024). How effective is an llm-based data analysis automation tool? a case study with chatgpt's data analyst. In *Simpósio Brasileiro de Banco de Dados (SBBD)*, pages 287–299. SBC.

Fernandez, R. C., Elmore, A. J., Franklin, M. J., Krishnan, S., and Tan, C. (2023). How large language models will disrupt data management. *Proceedings of the VLDB Endowment*, 16(11):3302–3309.

Fu, Y., Wang, D., Ying, W., Zhang, X., Liu, H., and Pei, J. (2025). Autonomous data agents: A new opportunity for smart data. *arXiv preprint arXiv:2509.18710*.

Hong, S., Lin, Y., Liu, B., Liu, B., Wu, B., Zhang, C., Li, D., Chen, J., Zhang, J., Wang, J., Zhang, L., Zhang, L., Yang, M., Zhuge, M., Guo, T., Zhou, T., Tao, W., Tang, R., Lu, X., Zheng, X., Liang, X., Fei, Y., Cheng, Y., Ni, Y., Gou, Z., Xu, Z., Luo, Y., and Wu, C. (2025). Data interpreter: An llm agent for data science. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19796–19821, Vienna, Austria. Association for Computational Linguistics.

Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Zhang, C., Wang, J., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., and Schmidhuber, J. (2024). Metagpt: Meta programming for a multi-agent collaborative framework. In *Proceedings of the Twelfth International Conference on Learning Representations*.

Piccialli, F., Chiaro, D., Sarwar, S., Cerciello, D., Qi, P., and Mele, V. (2025). Agentai: A comprehensive survey on autonomous agents in distributed ai for industry 4.0. *Expert Systems with Applications*, 291:128404.

Qian, C., Liu, W., Liu, H., Chen, N., Dang, Y., Li, J., Yang, C., Chen, W., Su, Y., Cong, X., Xu, J., Li, D., Liu, Z., and Sun, M. (2024). Chatdev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for*

Computational Linguistics (Volume 1: Long Papers), pages 15174–15186, Bangkok, Thailand. Association for Computational Linguistics.

Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., and Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652.

Sun, M., Han, R., Jiang, B., Qi, H., Sun, D., Yuan, Y., and Huang, J. (2025a). A survey on large language model-based agents for statistics and data science. *The American Statistician*, pages 331–344.

Sun, M., Han, R., Jiang, B., Qi, H., Sun, D., Yuan, Y., and Huang, J. (2026). Lambda: A large model based data agent. *Journal of the American Statistical Association*, 121(553):1–13.

Sun, Z., Wang, J., Zhao, X., Wang, J., and Li, G. (2025b). Data agent: A holistic architecture for orchestrating data-ai ecosystems. *arXiv preprint arXiv:2507.01599*.

Tang, Z., Wang, W., Zhou, Z., Jiao, Y., Xu, B., Niu, B., Zhou, X., Li, G., He, Y., Zhou, W., Song, Y., Tan, C., Wang, B., He, C., Wang, X., and Wu, F. (2025). Llm/agent-as-data-analyst: A survey. *arXiv preprint arXiv:2509.23988*.

Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L., and Anandkumar, A. (2023). Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., and Wen, J.-R. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.

Wang, P., Yu, Y., Chen, K., Zhan, X., and Wang, H. (2025). Large language model-based data science agent: A survey. *arXiv preprint arXiv:2508.02744*.

Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Cheng, W., Zhang, Q., Qin, W., Zheng, Y., Qiu, X., Huang, X., and Gui, T. (2025). The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68:121101.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. (2023). React: Synergizing reasoning and acting in language models. In *Proceedings of the Eleventh International Conference on Learning Representations*.

Zhou, X., He, J., Zhou, W., Chen, H., Tang, Z., Zhao, H., Tong, X., Li, G., Chen, Y., Zhou, J., Sun, Z., Hui, B., Wang, S., He, C., Liu, Z., Zhou, J., and Wu, F. (2025). A survey of llm $\times$ data. *arXiv preprint arXiv:2505.18458*.

Zhu, Y., Wang, L., Yang, C., Lin, X., Li, B., Zhou, W., Liu, X., Peng, Z., Luo, T., Li, Y., Chai, C., Chen, C., Di, S., Fan, J., Sun, J., Tang, N., Tsung, F., Wang, J., Wu, C., Xu, Y., Zhang, S., Zhang, Y., Zhou, X., Li, G., and Luo, Y. (2025). A survey of data agents: Emerging paradigm or overstated hype? *arXiv preprint arXiv:2510.23587*.