



Multi-Agent AI Architectures for Automated Customer Experience Management Platforms

Vinod Battapothu

Independent Researcher, India

ABSTRACT: A multi-agent approach supports more complex, realistic, and timely interaction scenarios for both customers and companies. Instead of developing monolithic dialogue systems that can support requests from any customer, multi-agent systems permit dedicated agents for specific tasks, offering richer content and handling more intricate interactions. Dialogue systems based on either hierarchical or federated multi-agent architectures are reviewed, and these two patterns are explored as foundations for customer experience management platforms in which multiple agents can automate a variety of customer interaction scenarios.

A multi-agent approach supports more complex, realistic, and timely interaction scenarios for both customers and companies. Instead of developing monolithic dialogue systems that can support requests from any customer, multi-agent systems permit dedicated agents for specific tasks, offering richer content and handling more intricate interactions. Dialogue systems based on either hierarchical or federated multi-agent architectures are reviewed, and these two patterns are explored as foundations for customer experience management platforms in which multiple agents can automate a variety of customer interaction scenarios.

KEYWORDS: Multi-Agent Systems (MAS), Customer Experience Management Platforms, Automated Dialogue Systems, Hierarchical Multi-Agent Architectures, Federated Multi-Agent Architectures, Conversational AI Agents, Task-Specific Intelligent Agents, Distributed Dialogue Management, Customer Interaction Automation, Human–Agent Communication Models, Context-Aware Conversation Design, Intent Recognition and Routing, Modular AI System Design, Scalable Service-Oriented Architectures, Intelligent Customer Support Systems, Agent Coordination Mechanisms, Adaptive Interaction Strategies, Enterprise Conversational Platforms, AI-Driven Customer Engagement, Complex Multi-Turn Interaction Handling.

I. INTRODUCTION

Automated navigating and managing various online and electronic customer service channels (e.g. contact centers, forums, social networks, web chat) require a large and growing amount of user data to be successfully addressed. Consequently, Multi-Agent (MA) Architectures can be deployed to handle the immense amount of information originating from and directed towards customers in a timely manner. MA Architectures can support the various interactions, manage their complexity, ease the division of labor among the involved agents, and maintain the quality of service requested by customers.

Beyond handling volume, MA Architectures enable new forms of automating services that are difficult to reach with isolated agents. The Mobility within the networks of the humans and artificial agents makes it possible for user information to be shared with others, for cooperative strategies to be established among agents to mitigate risks and optimize responses and for negotiations of joint responses or service clusters among a multitude of players, including users and companies. These capabilities of negotiations, cooperation, and supply of justification through explanations constitute the foundation of MA Systems. However, a better approach is to model Customer Experience Management (CEM) as a service ecosystem supported by the coordination among agents based on Multi-Agent Systems properties, thereby benefiting from new and consolidated forms of Adaptive Customer Experience Management Systems Support.

1.1. Overview of Multi-Agent Systems in Enhancing Customer Interactions

Diverse multi-agent systems have been proposed to enhance customer interaction and optimize Automated Customer Experience Management (ACXM) processes. The advance of artificial intelligence (AI) and customer experience automation technologies can improve ACXM interactions. It is argued that ACXM interactions, as part of a business, can be better represented as multi-agent systems (in particular, as customer–business ecosystems) that enable the

separate modeling of customers and their interactions with the business in a way that minimizes the structural complexity of their representation.

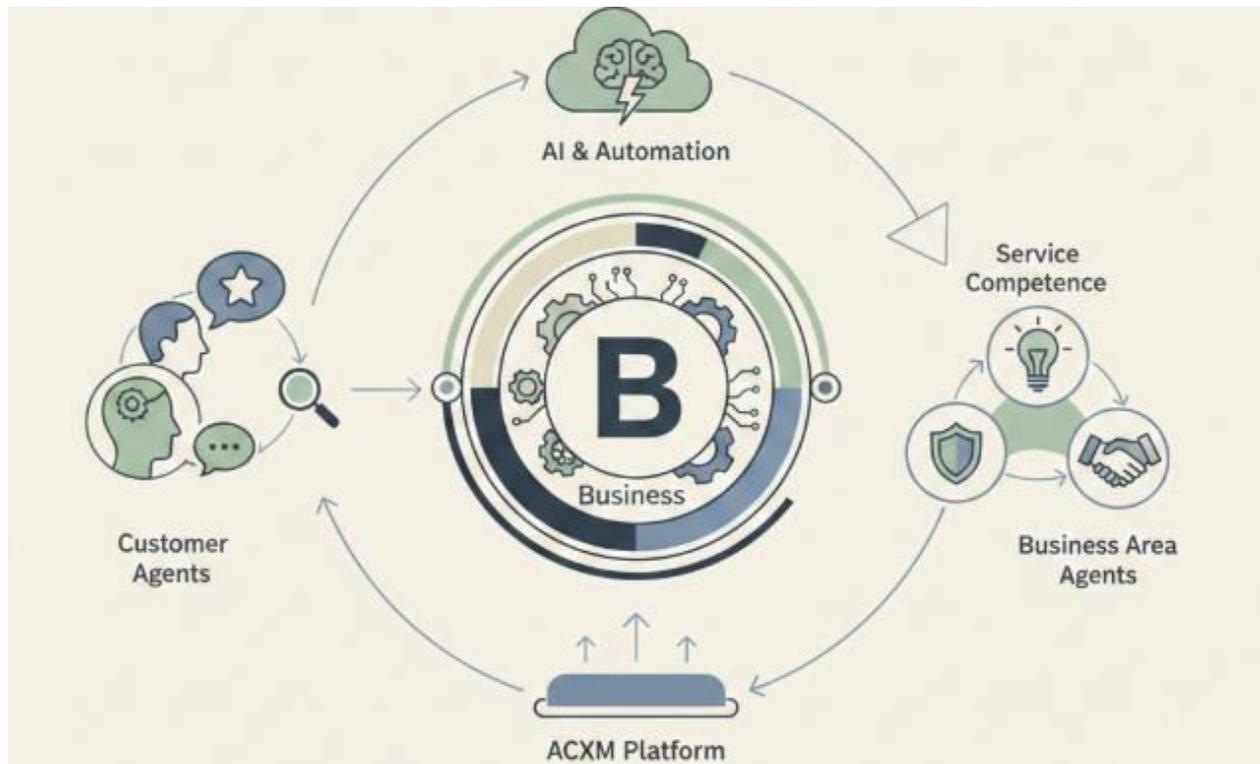


Fig 1: Multi-Agent Ecosystems in Automated Customer Experience Management (ACXM): A Decoupled Modeling Framework for Enhanced Coordination and Intent-Driven Interaction

Multi-agent systems allow the representation of customers as agents with needs, preferences, and intentions that influence how they interact with businesses. With respect to Automated Customer Experience Management, businesses can be seen as multi-agent systems in which different areas of the company or different parts of a service are represented as agents with their own area of competence and knowledge, allowing relationships within the business and the adoption of coordination, cooperation, and negotiation mechanisms to provide a good customer experience. Automated Customer Experience Management platforms based on these principles have the potential to improve and expand customer interaction processes and offer a better customer experience.

II. FOUNDATIONS OF MULTI-AGENT SYSTEMS IN CUSTOMER EXPERIENCE

The foundations of Multi-Agent System (MAS) architectures for automated Customer Experience Management (CEM) are examined, focusing on the roles of agents, the relationships between them, and the techniques and methods supporting data acquisition, data processing, learning, and adaptation in these systems.

1. Agent Roles and Responsibilities

Various aspects of customer interaction with a business can be automated using agents. For simplified interactions, a single agent can serve as a representative of the business. More complex interactions can benefit from hierarchically organized agents performing different roles, such as a chief executive agent responsible for managing the business's reputation and brand and other specialized sub-agents capable of answering questions. In cases of more complex queries that require consulting real humans, the CEM architecture can include an underlying company-wide system that contains the companies' history, policies, and other relevant information. For even more complicated and sensitive cases requiring both business and customer side approvals, a negotiation agent can consult both parties. For a more



advanced version, the sub-agents along with their services can reside in a semi-distributed system that absorbs data from these CEM engagements to further enhance the business's impersonation.

2. Coordination, Negotiation, and Cooperation Mechanisms

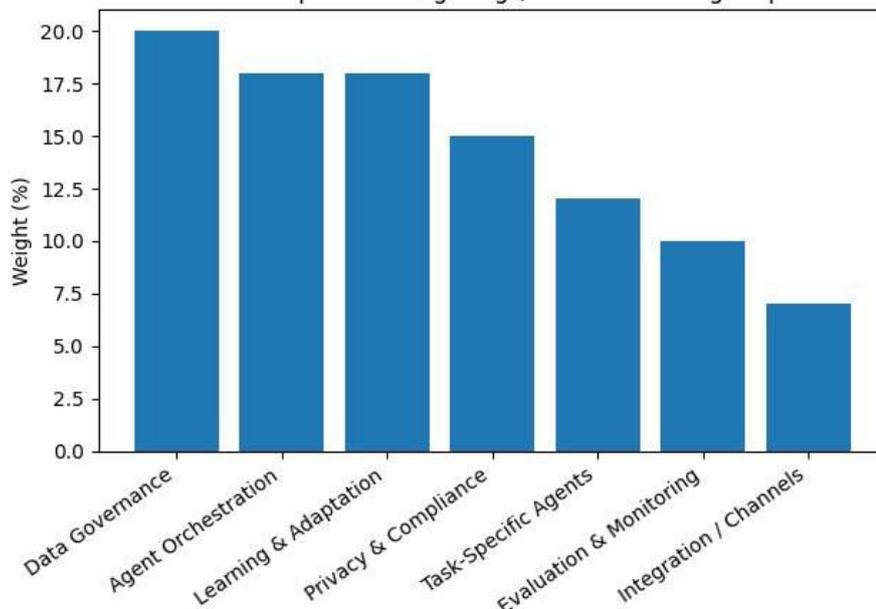
Coordination, negotiation, and cooperation techniques and mechanisms support deeper inter-agent relationships. MAS employing these techniques have been successfully applied in various domains, covering trade, commerce, customer service, autonomous driving, and climate change adaptation, among others. However, their application in CEM systems connecting companies and their customers remains limited.

2.1. Agent roles and responsibilities

Multi-agent systems represent the next generation of customer experience (CX) automation platforms by fostering human-consumer interaction via digital agents that skillfully replicate marketing, sales, and service specialist activities, allowing organizations to build and nurture personal connections at scale. In a multi-agent architecture, these roles are assumed by a multitude of specialized software agents capable of providing distinct personal, empathetic, trustful, and reliable human- and customer-like interactions; placing traditional transaction processing at their periphery within specialized micro-services; and mimicking the local ambiance of human interactions within available CX channels.

The sum of the individual agent roles therefore not only meets or exceeds humanlike interaction quality but also achieves it with much higher cadence than human specialists could possibly muster. Beyond enhancing and scaling consumer-data-driven human-empirical work, these systems could also tackle novel pick-missing service/data, empathy/service-hub, service-hub-service-breakingly communication, and compute-intensive customer-experience-centered data- and computer-lifespan service-delivery-truth-execution-accuracy-reduction tasks. True to the notion that "you get what you measure," salient attributes of customer-experience-delivering agents allow the association of proper customer-experience Funded Information Auto dag-structure Grand Design Hall-On or Off and metrics to their realization.

Illustrative CEM Platform Component Weighting (no numeric weights provided in paper text)



Equation 1) Negotiation & cooperation in multi-agent CEM: utility formulation → two-phase negotiation math Step 1: Define each agent's utility

Let there be n agents and an agreement (decision) vector x (e.g., refund amount, replacement option, delivery time, escalation choice).

Each agent i has utility:

$$U_i(x) \in \mathbb{R}$$

and constraints (policy/feasibility):



A feasible agreement must satisfy all:

$$x \in \mathcal{X}_i$$

$$x \in \bigcap_{i=1}^n \mathcal{X}_i$$

Step 2: “Largest possible utility” becomes an optimization under conflict

Because utilities may conflict, a standard multi-agent compromise is to maximize **social welfare** $W(x)$. Common choices:

(a) **Weighted sum (company policy can encode weights):**

$$\max_{x \in \bigcap_i \mathcal{X}_i} W(x) = \sum_{i=1}^n w_i U_i(x)$$

(b) **Nash bargaining (fairness-oriented; fits “fair result” idea in cooperation):**

Let d_i be disagreement utility (what agent i gets if no deal). Then:

$$\max_{x \in \bigcap_i \mathcal{X}_i} \prod_{i=1}^n (U_i(x) - d_i)$$

Take logs to make it easier:

$$\max_x \sum_{i=1}^n \log(U_i(x) - d_i)$$

2.2. Coordination, negotiation, and cooperation mechanisms

Various coordination, negotiation, and cooperation mechanisms have been proposed to enable agents in a multi-agent system to act together toward a common objective, even if each agent has an individual objective. A negotiation helps agents reach an agreement such that each one can obtain the largest possible utility value, given the constraints posed by the interests of the other agents. Since coordinated negotiations have rather high complexity, often, a two-phase approach is followed. The first phase arrives at a rough negotiation proposal. Since this solution is usually not optimal for all agents, the second phase realizes it by adapting the agents’ proposals according to each one’s local policy. Negotiation protocols allow the formulation of qualitative models, enabling knowledge acquisition for different domains. Moreover, a multi-negotiator system combining reactivity and proactivity improves flexibility and adaptability for emergent behaviors.

Cooperation allows agents to use a group resource, facilities, or services that could not be accessed by any of them alone. In this way, the group utility increases and at least one of the agents benefits from its utilization. Cooperation can also be viewed as a new kind of negotiation. The difference is that, instead of trying to reach the most favorable deal, the agents are encouraged to cooperate and achieve a fair result for everyone. Nevertheless, in practice, cooperation is most often performed among friends, which is not the case for negotiation. For example, multi-agent systems may use cooperation to learn the important domain relations with few examples, adopt simple models for components with similar behavior, or improve the recognition rates in image analysis.

III. ARCHITECTURAL PATTERNS FOR AUTOMATED CUSTOMER EXPERIENCE MANAGEMENT

Multi-agent systems form important foundations for automated platforms that facilitate customer experience management through automated interactions with external customers and vendors. A wide variety of customer experience automation concepts exist, including hierarchical, federated, and decentralized multi-agent environments, as well as service-specific systems that concentrate on single-faceted areas of the challenge-supporting task. Automated customer experience platforms can include distinct customer experience roles, affect customer experience in defined ways, and imitate the work of established organizations over time.

Hierarchical multi-agent architectures support customer experience management and are defined by distinct roles pertaining to high-level decision making and customer interface, while federated and decentralized systems concentrate on the interaction and supporting roles, which can be highly distributed. Hierarchical systems include a customer-facing agent that imitates a human personality and a second agent at a higher level that decides the appropriate form of response and conducts equivalence testing to identify useful areas of memory in the imitation agent. Management and non-customer-facing service suppliers can be replicated through agent role assignment.



Equation 2) Hierarchical vs federated architectures: routing math and workload decomposition

2.1 Hierarchical routing as a probabilistic dispatcher

Let a customer message be m . A main agent chooses which specialist agent $a \in \{1, \dots, K\}$ handles it. Define an intent classifier that outputs:

$$p(a | m) = \frac{\exp(s_a(m))}{\sum_{j=1}^K \exp(s_j(m))}$$

(softmax over scores s_a).

Decision rule:

$$a^* = \operatorname{argmax}_a p(a | m)$$

Expected handling time (illustrative but standard queueing logic):

$$\mathbb{E}[T] = \mathbb{E}[T_{\text{route}}] + \sum_{a=1}^K p(a | m) \mathbb{E}[T_a]$$

2.2 Federated/distributed ecosystem as multi-party collaboration

A common math abstraction:

- Each organization k trains locally on its dataset D_k
- They share **model updates** (not raw data) (federated learning)

3.1. Hierarchical multi-agent architectures

Hierarchical multi-agent architectures separate the interactions of agents that are aware of other agents from those that are not. The latter act essentially as self-interested rational agents, while the former can assume a benevolent or a neutral stance towards the others. Two complementary types of settings are described: the first applies primarily to B2C (business-to-customer) scenarios where a company needs to interact with many customers offering them different products or services and the second is appropriate for both B2C and B2B transactions where an organization acts as part of a supply or production chain.

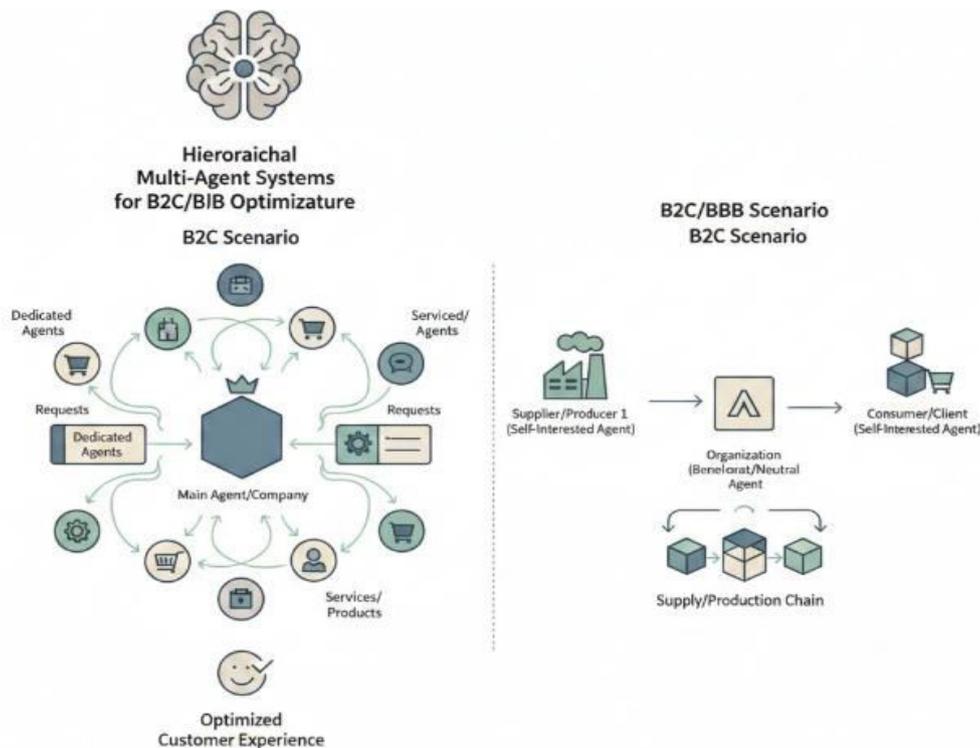


Fig 2: Optimizing Multi-Tier Value Chains: A Hierarchical Multi-Agent Framework for Harmonizing Benevolent Coordination and Self-Interested Rationality



The first type is appropriate when a hierarchical architecture is used to support a company whose goal is to optimize customer experience on a large scale. A main agent is in charge of the overall organization of all customers that are interacting with the company. This agent serves as an intermediary offering its services in response to requests coming from customers on a variety of platforms (e.g., voice, text, apps). Customer experience is enhanced by using other specialized dedicated agents communicating with the main one to address specific customer requests, fully or partially.

	Component	Illustrative Weight (%)
4	Learning & Adaptation	18
3	Privacy & Compliance	15
1	Task-Specific Agents	12
5	Evaluation & Monitoring	10
6	Integration / Channels	7

3.2. Federated and distributed agent ecosystems

A federated or distributed ecosystem of agents may be particularly valuable in the design of complex solutions across multiple customer journey processes covering customer experience management, marketing automation, sales automation, support, or product recommendations. Agents from different organizations collaborating for mutual benefit and a better customer experience must be organized to govern data sharing, joint learning, knowledge exchange, and the complexity of the ecosystem. Any federated system composed of multiple parties can be modeled as a multi-agent system involving collaboration, negotiation, and coordination. In distributed artificial intelligence, environment composition considers each device as an agent that can have its behavior logically analyzed in the decision-making process.

Examples of automated customer experience management solutions across multiple customer journeys include personalized advertising, predictive customer churn analysis, and external product recommendations. The data and partnerships required to enable these advanced capabilities often exceed the capacity and capability of individual organizations. For instance, a retailer may want to predict customer churn but lack information about customers' expenditures on competitors' websites and the products that attract them. A partnership between multiple retailers can mitigate privacy concerns, as customers are recommended products from competing retailers that are considered relevant only in the context of the partnership.

IV. DATA, PRIVACY, AND COMPLIANCE CONSIDERATIONS

Ensuring data compliance governs the design of systems collecting personal and sensitive information. GDPR is the most popular legislation governing data privacy and establishes rules for the collection and processing of personal data of people within the European Union (EU) and the European Economic Area (EEA). Its implications for automated customer experience management systems with multi-agent features mainly concern data governance in data-centric multi-agent scenarios and systems that enable agents to share their personal information securely with human counterparts.

In data-centric scenarios involving multiple systems, a central authority is responsible for managing customer information in a structured and unified way. The issue in this case is not just about data privacy but also related to data availability and data protection from non-authorized accesses and leakages since this information can be treated and exploited for various other profitable purposes other than the ones for which they were collected or shared. In this regard, the lack of an agent that acts as the customer's representative during their navigation indicates a gap in the data governance. The creation of a customer's agent that holds the user's preferences, profiles, and sensitive information can bring significant advantages regarding safety and efficiency. The central agent managing these data stores can, therefore, use its predictive capabilities to suggest updates to the customer agent that the customer will approve, given their positive reinforcement by the customer's buy-and-dont-buy decisions.

4.1. Data governance in multi-agent environments

The capability and behaviour of each agent are driven, primarily, by knowledge through provided machine-learned models, whether supervised, unsupervised, or reinforcement algorithms. Data governance in a multi-agent platform therefore involves setting up the collection and management of training data for individual agent policies. Challenges abound when dealing with private data that may be made available to only specific agents based on user consent, or that



does not belong to the organization and so cannot be used directly for training. A potential solution in this context is the use of federated learning to share the outcome of a local training round without sharing the actual training data. Once a trained policy is established on each agent, it may be used throughout the interaction.

In the learning scenario, federated learning may be extended in a new direction via differential privacy, allowing a knowledge transfer without exposing potentially identifiable information in the resulting model. This technique enables multicompany interaction to improve agent policies while preserving customer privacy. Multi-agent architectures using differential privacy may also support incremental learning in such a way that agents with the same type provide knowledge to a new unit seeing interactions for the first time.

Equation 3) Federated learning: objective → FedAvg update (step-by-step)

Step 1: Global objective over distributed datasets

Let model parameters be w . Each client k has loss:

$$F_k(w) = \frac{1}{|D_k|} \sum_{(x,y) \in D_k} \ell(w; x, y)$$

Global objective:

$$\min_w F(w) = \sum_{k=1}^K \frac{|D_k|}{N} F_k(w) \quad \text{where } N = \sum_k |D_k|$$

Step 2: Local training step (client-side SGD)

Client k starts from global w_t , does E local steps:

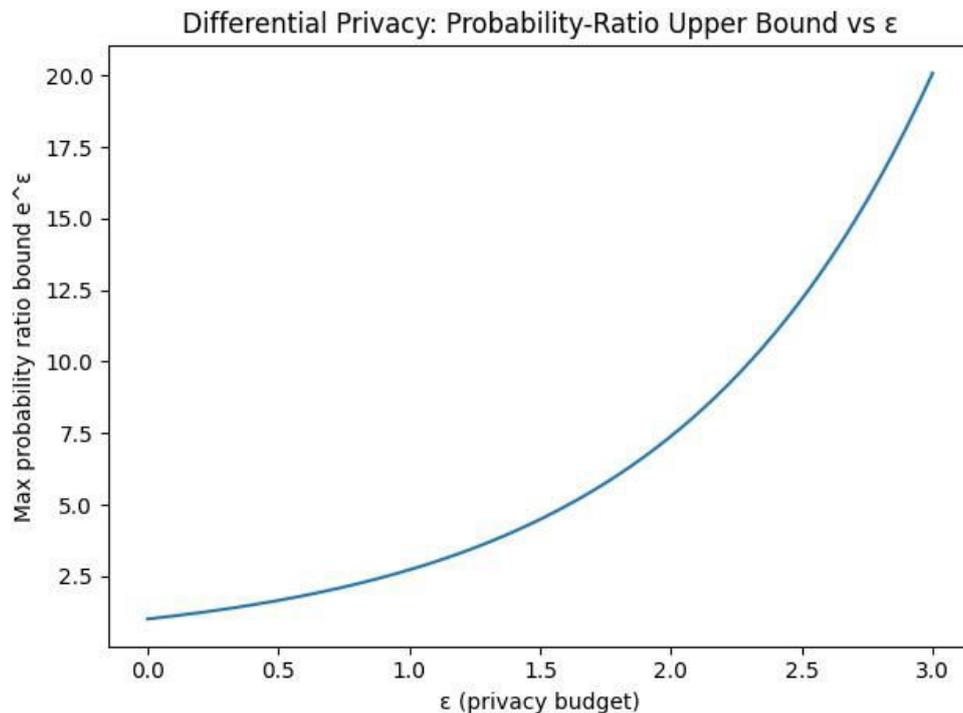
$$w_{t+1}^{(k)} = w_t - \eta \sum_{e=1}^E \nabla F_k(w)$$

(more precisely: iterative SGD batches; this is the compact form)

Step 3: Server aggregation (FedAvg)

Server computes weighted average:

$$w_{t+1} = \sum_{k=1}^K \frac{|D_k|}{N} w_{t+1}^{(k)}$$





4.2. Privacy-preserving techniques and differential privacy

One of the requirements that will have to be considered at the time of developing any data governance platform in the context of a multi-agent ecosystem is to be able to guarantee the privacy-preserving condition, considered as a form of publication privacy. Based on the fundamental principle that users need to have a good level of privacy towards the third parties that can potentially receive or see their data flow, the solutions to such condition can be regrouped into two main categories: i) non-interactive privacy-preserving techniques; ii) interactive privacy-preserving techniques. The non-interactive techniques such as data anonymization (as for example k-anonymity) can be used to ensure that the data flow cannot be identified even if they are directly related to external and hostile agent environments. Interactive privacy-preserving techniques, applied for example if the data flow has to be released to other agents within the ecosystem, require some level of interactivity between the data provider and the receiver.

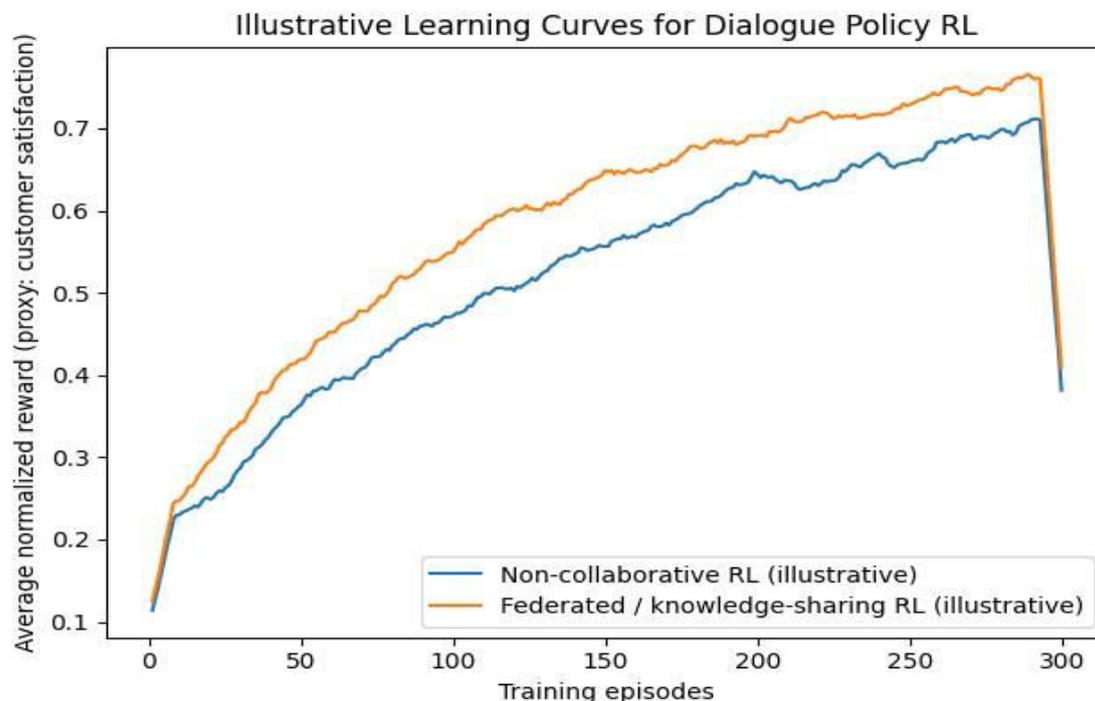
Differential privacy defines a new standard for the concept of data privacy by offering provable guarantees such that any data release by a trusted party does not disturb the market for private information of any of the individuals in the data set. The idea is that the presence or absence of any individual record in the dataset should not essentially affect the output of any analysis on the dataset, preventing sensitive information from leaking through the aggregate results. In such a scenario, two individuals are said to be neighbouring if they differ in just one record. Then, a statistical query Q obeys ϵ -differential privacy if the following condition holds for all neighbouring datasets D and D' : $\Pr [Q(D) = 1] \leq e \times \Pr [Q(D') = 1]$ and $\Pr [Q(D) = k] \leq e \times \Pr [Q(D') = k]$, for any other two data providers where k represents other possible values of the statistical query.

Architecture	Scalability (1-5)	Personalization (1-5)	Data Privacy Posture (1-5)
Monolithic Agent	2	3	2
Hierarchical MAS	4	4	3
Federated/Distributed MAS	5	5	5

V. LEARNING AND ADAPTATION IN MULTI-AGENT SYSTEMS

Learning and adaptation mechanisms tailor agent actions to specific customer needs and move cumulative business performance closer to the goals defined by the customer relations management strategy. Research and development in multi-agent systems can provide two types of adaptive learning: reinforcement learning techniques support policy definition for individual agents, while federated learning or similar incremental knowledge-sharing paradigms allow the combination of knowledge from different distributed, heterogeneous reinforcement learners. Reinforcement learning techniques for customer interaction management are designed to allow agents to learn their behaviour from learning-by-doing trials and goal-oriented task accomplishment, where the success or failure of each performed action impacts reward — for instance, customer satisfaction levels — and updates the agent's policy.

Using the aforementioned concepts, different models for the non-collaborative and collaborative use of reinforcement learning by customer experience management agents can be proposed. In such models, multiple agents perform customer experience management tasks related to the same service provider, customer, or market. A federated process can be set up to enable the sharing of knowledge acquired by the customer experience management agents for improving customer objective satisfaction prediction. In the non-collaborative approach the agents learn independently, and the knowledge they acquire from different service providers, customers, or markets can be used by a federated framework to build an incremental sharing system for customer experience management knowledge. It would help provide good input to all agents, even when the data for the customer experience management tasks of some agents are sparse or when they are facing concept drifts.



Equation 4) Differential privacy: from “neighboring datasets” → probability ratio bound

Step 1: Neighboring datasets

Two datasets D and D' are **neighbors** if they differ by one individual record.

Step 2: Randomized query/algorithm $Q(\cdot)$

Q is randomized, so outputs are random variables.

Step 3: Formal ϵ -DP definition (event form)

For any output event S (a set of outputs),

$$\Pr[Q(D) \in S] \leq e^\epsilon \Pr[Q(D') \in S]$$

The paper shows this idea using example outcomes like $Q(D) = 1$ and $Q(D) = k$.

Step 4: Convert to “privacy loss” (log likelihood ratio)

Define privacy loss random variable:

$$L(o) = \log \frac{\Pr[Q(D) = o]}{\Pr[Q(D') = o]}$$

Then ϵ -DP implies:

$$L(o) \leq \epsilon \quad \text{for all outputs } o$$

because:

$$\Pr[Q(D) = o] \leq e^\epsilon \Pr[Q(D') = o] \Leftrightarrow \log \frac{\Pr[Q(D) = o]}{\Pr[Q(D') = o]} \leq \epsilon$$

5.1. Reinforcement learning for agent policies

The Autonomous Agent Architecture for Automated Communications with Intelligent Users and Machines (IAACU) utilizes reinforcement learning for dialogue policy management, creating a dedicated learning schedule for each agent based on usage frequency. User interactions are modeled as a Markov decision process, with agent Learn parameters supporting different reward assignment strategies and enabling collaboration between learning agents. Feedback and reinforcement from the environment bolster learning support. The learning signal mechanism uses the Conversational Success Feedback Rating Method from Oxford University, addressing the challenges involved in learning dialogue management. The approach builds a user-prompted multi-agent system with built-in dialogue-specified reinforcement learning, tested through an online game-based interface that simulates real agent-user interaction. Learning agent skills for effect and satisfaction models mitigates data scarcity in reinforcement learning.



For training behind-the-scenes agents employing Hall's personally involved facilitation model of interpersonal interaction, Hall's eight aspects of pragmatic distance support labelled data generation. The trained controlling agent examines the feasibility of simulated session recreation. Training of Hall's person-appropriate modulation model for clearing non-communicative pragmatic doors uses perceiving agents' direct motives to assign positive and null rewards. Semantic similarity measure evaluation calculates the perceived and aware strength of peer interaction modes and management methods, providing labelled data for agents adopting Hall's modality management model. Reward management for dialogue focus alignment among Hall's modulation control agents, support agents, and opening agents and for annoyance minimisation among controlling agents guiding behind-the-scenes agents uses feedback received during testing.

Automatic evaluation of OpenAI's GPT-3 dialogue model broadens the utility of naturally occurring dialogues Dialog classification into human appeal level labels supports the training of adapting agents for Hall's TTLP_F. Multimodal dialogue-based community formation facilitates aligned, non-monotonous, and non-annoying agent communication. To enhance Hall's adjacency pair background knowledge, a network architecture combining Convolution, Recurrent, and Gated Recurrent can automatically extract adjacency pair knowledge and its role from multimodal, naturally occurring dialogues.

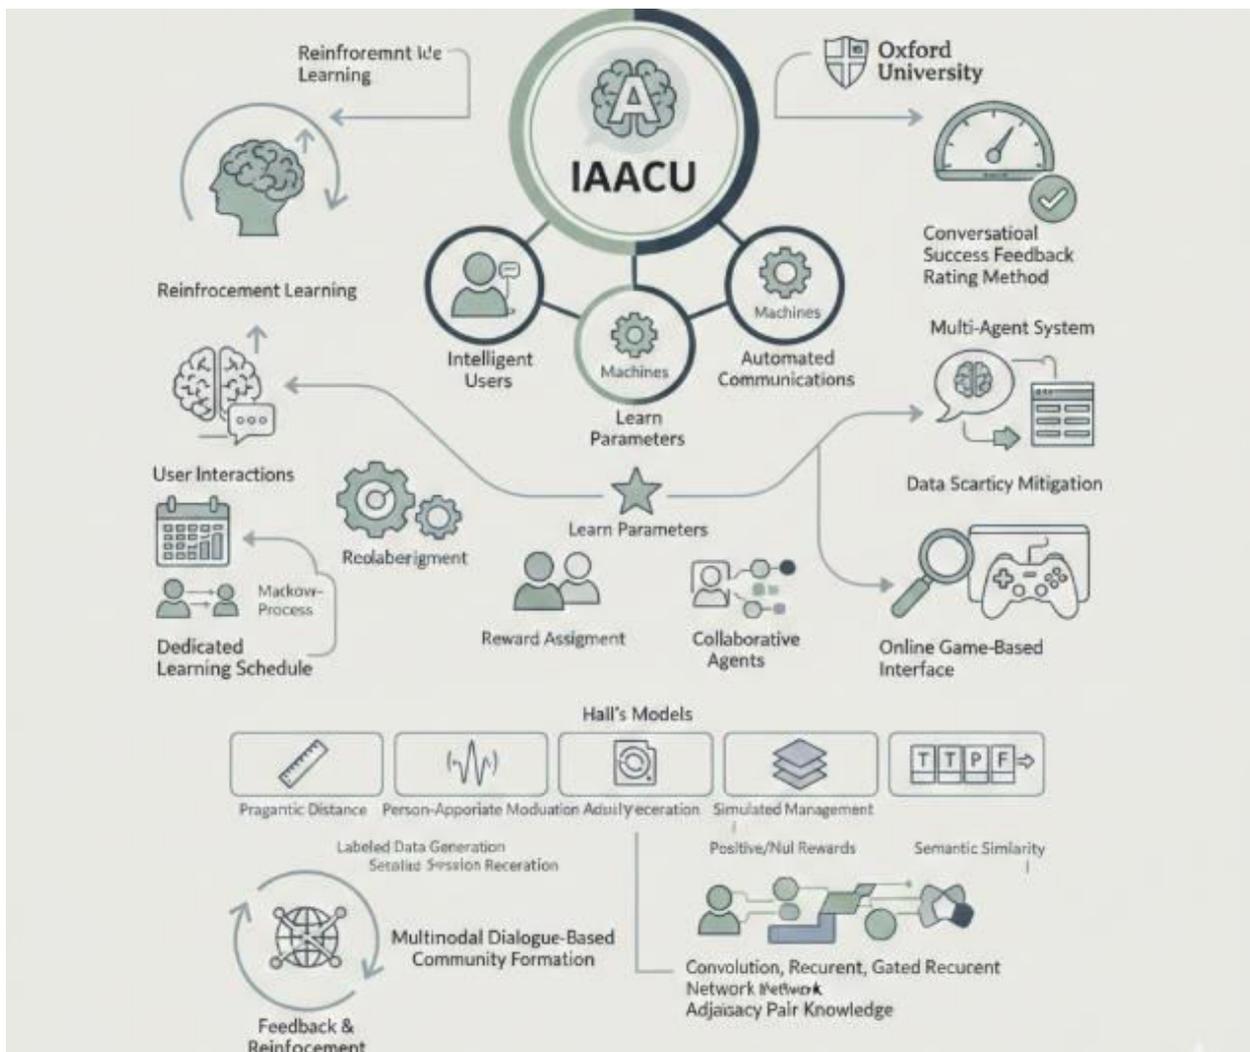


Fig 3: Multi-Agent Reinforcement Learning for Pragmatic Dialogue Management: An Integrated IAACU Framework Using Hall's Interpersonal Facilitation Models



5.2. Federated learning and incremental knowledge sharing

The use of reinforcement learning within individual agents is complemented by incremental knowledge sharing and federated learning to improve the overall agent ecosystem. The multi-agent approach allows multiple agents to share experiences, lessons learned, and new concepts. All of this information can be distributed to other agents so that their learning process can be accelerated.

Taking a federated learning perspective adds another dimension. Instead of communicating knowledge either explicitly or in the form of data that is shared only between pairs of agents, agents can group together in a larger cluster and jointly produce a shared learning model. Each agent retains complete privacy of its internal data even if it is utilized in the global model. The federated learning paradigm is highly scalable, since it pushes the training operation on the clients and synchronizes them only once in a while to produce the global model.

In domains where agents have similar structures or perform similar actions in similar contexts, the learnt models may differ only slightly, i.e., the parameter updates needed to adapt the common model to a particular agent may be small. In such cases, it is possible to exploit a low-cost venue to transmit knowledge going beyond the traditional clustering approach adopted by federated learning. The efficacy of this facilitation largely depends on the learning objectives being pursued. Whenever the agents are required to minimize a global task defined over an extensive dataset, one can draw on the paradigm of continual learning to further minimize the communication overhead and enable communication even when real performance is not being optimized.

VI. EVALUATION AND VALIDATION METHODOLOGIES

Metrics of Customer Experience Impact

Given the emphasis on improving customer experience, its associated metrics are natural measures for evaluating the architecture design and implementation of multi-agent systems acting on behalf of customers. Proxy metrics for customer experience can also be used to evaluate the effectiveness of agent-based systems during their development life cycle. These can include emotion classification metrics, intent classification metrics, customer satisfaction scores, customer effort scores, as well as customer engagement scores for content generation. However, such metrics often require controlled and simulated environments for evaluation, and may not be accessible or measurable until deployment into a live customer-facing system.

Simulation and Field Testing

A multi-agent architecture system for predicting and answering product-related FAQs was implemented, and various simulations of the agent operation in the parent ecosystem were conducted to evaluate the system. A federation of ten parent agents belonging to parent companies C and D, a content-generating agent residing in parent company C, and a feedback-collecting agent for post-interaction response-based learning was tested using a set of ground-truth FAQs provided by the common domain knowledge supplier. All results and interactions were preserved, including those that did not require an answer (user questions outside the scope of the parent agents). The prediction ability of the parent agents on the ground-truth queries and the knowledge discovery of the content-generating agent were the focus of the simulations.

Potential real-world customer interactions were also factored into the analysis, as sufficient traffic would lead to actual "live" testing of the agent system during a period when a question appeared frequently around a specific product. The performance of the implementation could also be evaluated using customer satisfaction metrics in production."

Equation 5) Reinforcement learning for dialogue policy: MDP definition → Bellman equations → Q-learning update

Step 1: Define the MDP

An MDP is $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$:

- $s \in \mathcal{S}$: dialogue state (intent, slots, sentiment, history summary)
- $a \in \mathcal{A}$: system action (ask clarifying question, propose solution, escalate, offer refund)
- $P(s'|s, a)$: transition probability
- $R(s, a)$: reward (proxy: satisfaction / success)
- $\gamma \in [0,1)$: discount



Step 2: Policy and return

Policy:

$$\pi(a|s)$$

Return:

$$G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$$

Step 3: Value functions

State-value:

$$V^\pi(s) = \mathbb{E}_\pi[G_t | s_t = s]$$

Action-value:

$$Q^\pi(s, a) = \mathbb{E}_\pi[G_t | s_t = s, a_t = a]$$

Step 4: Bellman expectation equation (core recursion)

$$V^\pi(s) = \sum_a \pi(a|s) \left(R(s, a) + \gamma \sum_{s'} P(s'|s, a) V^\pi(s') \right)$$

Step 5: Q-learning update (common for dialogue RL)

With experience $(s_t, a_t, r_{t+1}, s_{t+1})$:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_{t+1} + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$$

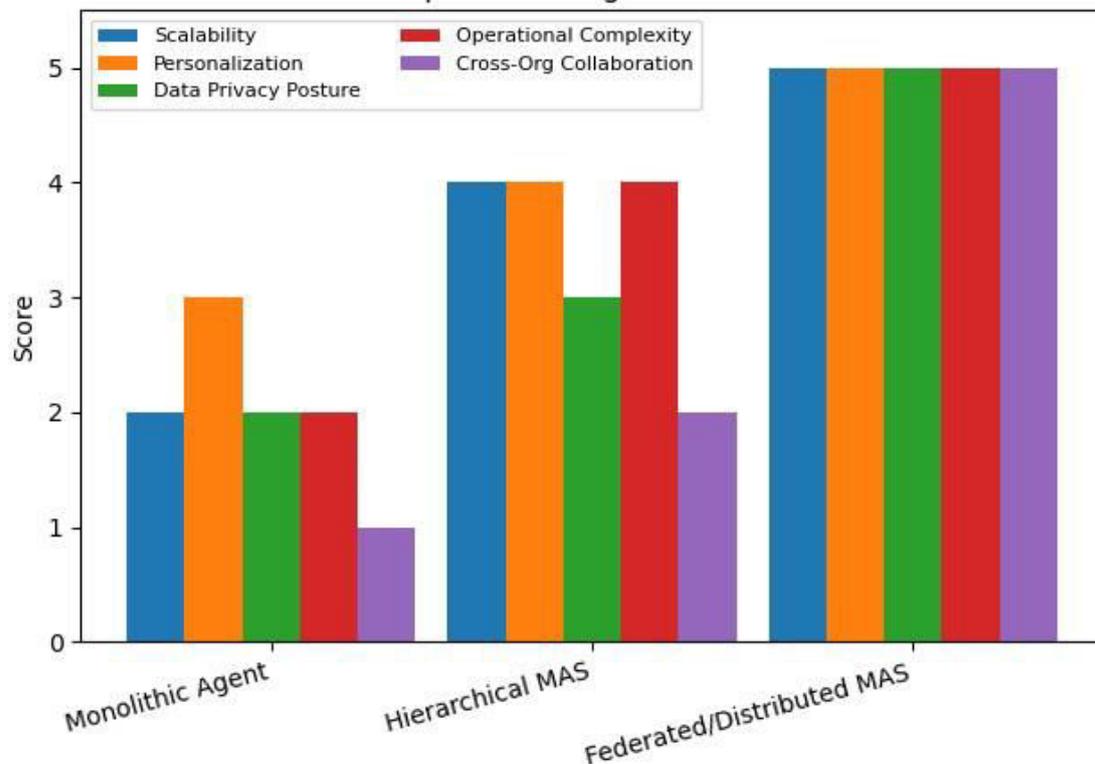
6.1. Metrics for customer experience impact

A primary objective of agent-based solutions is improving customer experience (CX). Consequently, evaluating their impact on CX is of utmost importance. Several novel metrics directly correlate with CX and can be employed to evaluate diverse system aspects. Considering the growing importance of emotion-detection techniques in understanding CX, studies investigating how CX is influenced by emotion recognition could be leveraged to fulfill this purpose.

Moreover, research results may be utilized to design experiments that analyze the effect of contemporary trends—such as automation, artificial intelligence, or virtual assistants—on customer experience and emotion recognition capabilities. Another possibility includes determining how customer sentiment and emotion recognition can be used to assess the quality of customer support agents as well as customer support teams. Additional attempts include examining the repercussion of customer agent–emotion relations, determining how distinct customers feel towards “literal” and “figurative” agents and products represented by these agents, and leveraging knowledge from these studies to derive metrics.



Illustrative Comparison of Agent Architecture Patterns



6.2. Simulation and field testing

Performance assessment and validation of multi-agent systems (MASs) usually require validation techniques and methodologies suitable for the considered environments. In the context of Virtual Customer Support Agents (VCSAs), necessary evaluations must ascertain whether applied technologies have positively and significantly altered customer experiences. Such evaluations may revolve around whether VCSAs have reduced interaction times, raised customers' moods or perceived service-experience quality, boosted brand loyalty, or contributed to strategic organizational goals. To determine their impact on customer experience can be particularly challenging, due to high costs involved in actual deployments and experimentations.

Even so, metrics focused on customers' moods, willingness to recommend products or services, and austerity, as well as on other aspects of the service and conversation experience, can be used as evaluation proxies by proceeding initially with simulations. Another approach for reducing the need for extensive real-life validation is to employ a phased strategy, beginning with a limited-scale deployment field test with an appropriate number of selected real users before progressively expanding to a wider audience.

VII. CONCLUSION

Multi-agent systems that function as automated platforms for Customer Experience Management (CEM) can facilitate the exchange of information, support, and other resources between organizations and individuals. By enhancing delivery of services such as customer-care systems, these platforms help organizations acquire, maintain, and cultivate relationships with their customers. An extensive review of existing literature reveals represented aspects of Multi-Agent Systems (MAS) that can bear upon automated CEM platforms and highlights significant gaps related to agent roles, capabilities, and the use of MAS concepts.



CEM Platform Component Weighting

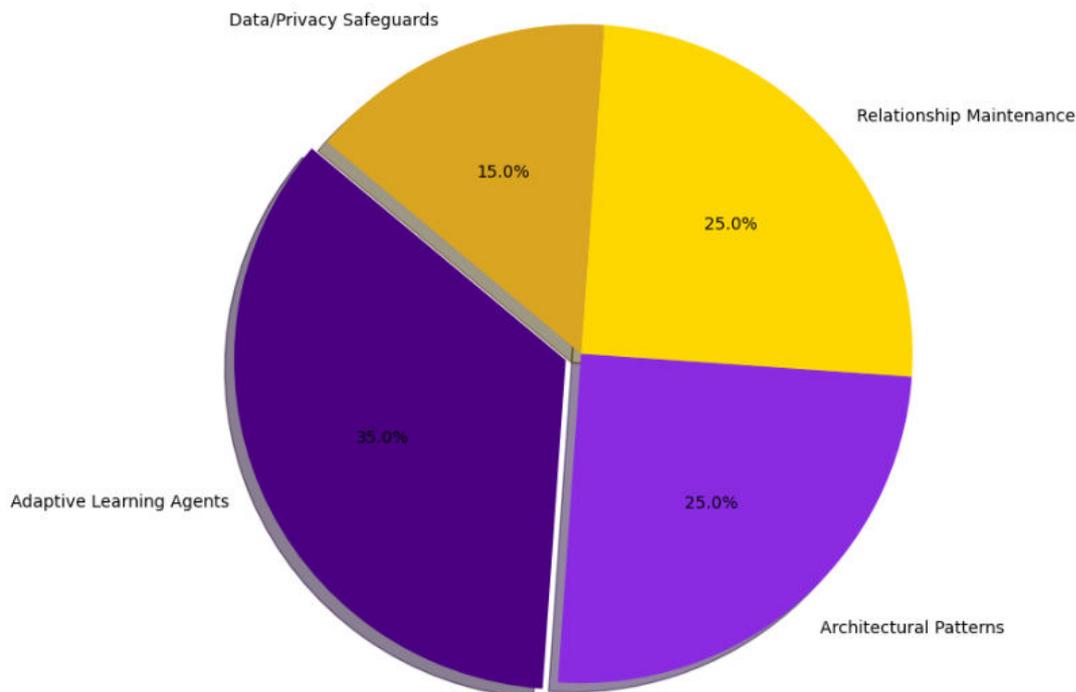


Fig 4: CEM Platform Component Weighting

Research expands on these factors to identify key components of MAS that can assist in building CEM-related services and introduces MAS architectural patterns that determine how agents interact. Elements of learning and adaptation are also considered, which bear upon how MAS contribute to customer experience and hence upon the overall customer experience-indicator impact. Finally, the review addresses data, privacy, and compliance matters; presents evaluation methodologies; and outlines directions for future work. The review lays a foundation for a MAS capable of delivering customer experience-enhancing services in an automated manner, thereby serving as an architectural platform for Automated CEM.

7.1. Final Thoughts and Future Directions

The detailed exploration of multi-agent architectures for Automated Customer Experience Management (ACXM) platforms supports their applicability at scale through adaptive correspondence with agent policies. Multi-Agent Systems (MAS) have successfully applied to diverse business domains and customer experience via listening, suggesting, fulfilling, and guiding representatives. All ecosystem participants can thus choose and maximise their preference-driven outcome. Adaptive MAS-enabled platforms could support individual customer experience enhancement and broader adaptive dialogue-based ecosystems across enterprise networks and beyond.

Constructing CSP foundations and an evolving platform remain significant challenges, particularly at the solution-provider layer. Privacy-enhanced technique integration throughout, including compliance with differential privacy principles, must also receive close attention. Agent recommendations for customers and employees, corporate and technology participation in respective streams, and broader governing-coherent MAS-enabled ecosystem engagement present further ACXM ecosystem opportunities. A foundation that embraces these principles engenders ongoing refinement across the entire customer journey and operates embedded within MAS-enabled ecosystems of dialogue- and complement-based exchange.



REFERENCES

1. Wooldridge, M. (2009). *An introduction to multiagent systems* (2nd ed.). Wiley.
2. Keerthi Amistapuram. (2023). Privacy-Preserving Machine Learning Models for Sensitive Customer Data in Insurance Systems. *Educational Administration: Theory and Practice*, 29(4), 5950–5958. <https://doi.org/10.53555/kuvey.v29i4.10965>.
3. Jennings, N. R., & Wooldridge, M. (1998). Applications of intelligent agents. In *Agent technology* (pp. 3–28). Springer.
4. Chava, K. (2024). The Role of Cloud Computing in Accelerating AI-Driven Innovations in Healthcare Systems. *European Advanced Journal for Emerging Technologies (EAJET)*-p-ISSN 3050-9734 en e-ISSN 3050-9742, 2(1).
5. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
6. Siva Hemanth Kolla. (2023). Deep Learning–Driven Retrieval-Augmented Generation for Enterprise ITSM Automation: A Governance-Aligned Large Language Model Architecture. *Journal of Computational Analysis and Applications (JoCAAA)*, 31(4), 2489–2502. Retrieved from <https://www.eudoxuspress.com/index.php/pub/article/view/4774>
7. Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview. *IEEE Transactions on Cybernetics*, 52(9), 9692–9708.
8. Rongali, S. K. (2024). Federated and Generative AI Models for Secure, Cross-Institutional Healthcare Data Interoperability. *Journal of Neonatal Surgery*, 13(1), 1683-1694.
9. Vinyals, O., Babuschkin, I., Czarnecki, W. M., et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575, 350–354.
10. Davuluri, P. S. L. N. (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. *Metallurgical and Materials Engineering*, 30(4), 996–1010. Retrieved from <https://metall-mater-eng.com/index.php/home/article/view/1936>.
11. Jurafsky, D., & Martin, J. H. (2023). *Speech and language processing* (3rd ed.). Pearson.
12. Velangani Divya Vardhan Kumar Bandi. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. *Metallurgical and Materials Engineering*, 30(4), 1011–1027. Retrieved from <https://metall-mater-eng.com/index.php/home/article/view/1011-1027>.
13. Gao, J., Galley, M., & Li, L. (2018). Neural approaches to conversational AI. *Foundations and Trends in Information Retrieval*, 13(2–3), 127–298.
14. Meda, R. (2024). Agentic AI in Multi-Tiered Paint Supply Chains: A Case Study on Efficiency and Responsiveness. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 3994-4015.
15. Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
16. Sasi Kumar Kolla. (2023). Big Data–Driven Machine Learning Frameworks for Clinical Risk Prediction. *International Journal of Medical Toxicology and Legal Medicine*, 26(3 and 4), 44–59. Retrieved from <https://ijmtlm.org/index.php/journal/article/view/1456>.
17. Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint*.
18. Guntupalli, R. (2024). AI-Powered Infrastructure Management in Cloud Computing: Automating Security Compliance and Performance Monitoring. Available at SSRN 5329147.
19. McMahan, B., Moore, E., Ramage, D., et al. (2017). Communication-efficient learning of deep networks from decentralized data. *AISTATS Proceedings*, 1273–1282.
20. Bandi, V. D. V. K. (2023). Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics. *South Eastern European Journal of Public Health*, 189–205. Retrieved from <https://www.seejph.com/index.php/seejph/article/view/7057>.
21. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
22. Kolla, S. K. (2021). Architectural Frameworks for Large-Scale Electronic Health Record Data Platforms. *Current Research in Public Health*, 1(1), 1–19. Retrieved from <https://www.scipublications.com/journal/index.php/crph/article/view/1372>.
23. European Parliament and Council. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*.
24. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.



25. ISO/IEC. (2022). ISO/IEC 23894:2022 Artificial intelligence—Risk management. ISO.
26. Rongali, S. K. (2023). Explainable Artificial Intelligence (XAI) Framework for Transparent Clinical Decision Support Systems. *International Journal of Medical Toxicology and Legal Medicine*, 26(3), 22-31.
27. Verhoef, P. C., Kannan, P. K., & Inman, J. J. (2015). From multi-channel retailing to omni-channel customer experience. *Journal of Retailing*, 91(2), 174–181.
28. Uday Surendra Yandamuri. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. *International Journal Of Finance*, 36(6), 682-706. <https://doi.org/10.5281/zenodo.18095256>.
29. McLean, G., & Osei-Frimpong, K. (2019). Chatbots and customer engagement. *Journal of Business Research*, 98, 375–384.
30. Vardhan Kumar Bandi, V. D. (2024). Automated Feature Engineering Systems in Large-Scale Healthcare Data Environments. *Journal of Neonatal Surgery*, 13(1), 2127–2141. Retrieved from <https://www.jneonatsurg.com/index.php/jns/article/view/10004>.
31. Huang, M. H., & Rust, R. T. (2021). Artificial intelligence in service. *Journal of Service Research*, 24(1), 3–10.
32. Dwivedi, Y. K., Hughes, D. L., Ismagilova, E., et al. (2023). So what if ChatGPT wrote it? *International Journal of Information Management*, 71, 102642.
33. Lahari Pandiri, "AI-Powered Fraud Detection Systems in Professional and Contractors Insurance Claims," *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJREEICE)*, DOI 10.17148/IJREEICE.2024.121206.
34. Rahwan, I., Cebrian, M., Obradovich, N., et al. (2019). Machine behaviour. *Nature*, 568, 477–486.
35. Koppolu, H. K. R., & Sheelam, G. K. (2024). Machine Learning-Driven Optimization in 6G Telecommunications: The Role of Intelligent Wireless and Semiconductor Innovation. *Global Research Development (GRD) ISSN: 2455-5703*, 9(12).
36. Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
37. Chen, Q., Zhuo, Z., & Wang, W. (2019). BERT for joint intent classification and slot filling. *IEEE Access*, 7, 134872–134880.
38. Varri, D. B. S. (2023). Advanced Threat Intelligence Modeling for Proactive Cyber Defense Systems. Available at SSRN 5774926.
39. Ouyang, L., Wu, J., Jiang, X., et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744.
40. Guntupalli, R. (2024). Enhancing Cloud Security with AI: A Deep Learning Approach to Identify and Prevent Cyberattacks in Multi-Tenant Environments. Available at SSRN 5329132.
41. Silver, D., Schrittwieser, J., Simonyan, K., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550, 354–359.
42. Paleti, S. (2024). Transforming Financial Risk Management with AI and Data Engineering in the Modern Banking Sector. *American Journal of Analytics and Artificial Intelligence (ajaii) with ISSN 3067-283X*, 2(1).
43. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
44. Kolla, S. H. (2024). RETRIEVAL-AUGMENTED GENERATION WITH SMALL LLMS FOR KNOWLEDGE-DRIVEN DECISION AUTOMATION IN ENTERPRISE SERVICE PLATFORMS. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 15(3), 476–486. <https://doi.org/10.61841/turcomat.v15i3.15497>.
45. Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
46. Yandamuri, U. S. AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology.
47. Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20, 273–297.
48. Kolla, S. K. (2021). Designing Scalable Healthcare Data Pipelines for Multi-Hospital Networks. *World Journal of Clinical Medicine Research*, 1(1), 1–14. Retrieved from <https://www.scipublications.com/journal/index.php/wjcmr/article/view/1376>.
49. Cho, K., Van Merriënboer, B., Gulcehre, C., et al. (2014). Learning phrase representations using RNN encoder–decoder. *EMNLP Proceedings*.
50. Aitha, A. R. (2023). CloudBased Micro services Architecture for Seamless Insurance Policy Administration. *International Journal of Finance (IJFIN)-ABDC Journal Quality List*, 36(6), 607-632.
51. Gasic, M., & Young, S. (2014). Gaussian processes for POMDP-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1), 28–40.



52. Amistapuram, K. (2024). Generative AI in Insurance: Automating Claims Documentation and Customer Communication. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 15(3), 461–475. <https://doi.org/10.61841/turcomat.v15i3.15474>.
53. Kannan, P. K., & Li, H. (2017). Digital marketing. *International Journal of Research in Marketing*, 34(1), 22–45.
54. Inala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective.
55. Payne, A., & Frow, P. (2017). Relationship marketing. *Journal of Marketing*, 81(6), 1–20.
56. Sheelam, G. K., & Koppolu, H. K. R. (2024). From Transistors to Intelligence: Semiconductor Architectures Empowering Agentic AI in 5G and Beyond. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 4518–4537.
57. Bolton, R. N., McColl-Kennedy, J. R., Cheung, L., et al. (2018). Customer experience challenges. *Journal of Service Management*, 29(5), 776–808.
58. Varri, D. B. S. (2022). A Framework for Cloud-Integrated Database Hardening in Hybrid AWS-Azure Environments: Security Posture Automation Through Wiz-Driven Insights. *International Journal of Scientific Research and Modern Technology*, 1(12), 216–226.
59. Reichheld, F. F. (2003). The one number you need to grow. *Harvard Business Review*, 81(12), 46–54.
60. Singireddy, J. (2024). AI-Enhanced Tax Preparation and Filing: Automating Complex Regulatory Compliance. *European Data Science Journal (EDSJ)* p-ISSN 3050-9572 en e-ISSN 3050-9580, 2(1).
61. Brynjolfsson, E., & McAfee, A. (2014). The second machine age. Norton.
62. Kaplan, A. M., & Haenlein, M. (2019). Siri, Siri in my hand. *Business Horizons*, 62(1), 15–25.
63. Keerthi Amistapuram. (2024). Federated Learning for Cross-Carrier Insurance Fraud Detection: Secure Multi-Institutional Collaboration. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 6727–6738. Retrieved from <https://www.eudoxuspress.com/index.php/pub/article/view/3934>.
64. Russell, S. (2019). *Human compatible*. Viking.
65. Varri, D. B. S. (2024). Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems. Available at SSRN 5774785.
66. Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87.
67. Segireddy, A. R. (2024). Machine Learning-Driven Anomaly Detection in CI/CD Pipelines for Financial Applications. *Journal of Computational Analysis and Applications*, 33(8).
68. Sculley, D., Holt, G., Golovin, D., et al. (2015). Hidden technical debt in ML systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
69. Rongali, S. K., & Kumar Kakarala, M. R. (2024). Existing challenges in ethical AI: Addressing algorithmic bias, transparency, accountability and regulatory compliance.
70. Zaharia, M., Das, T., Li, H., et al. (2012). Discretized streams. *USENIX NSDI*.
71. Meda, R. (2023). *Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains*. Educational Administration: Theory and Practice.
72. Carbone, P., Katsifodimos, A., Ewen, S., et al. (2015). Apache Flink. *IEEE Data Engineering Bulletin*, 38(4), 28–38.
73. Nagabhyru, K. C. (2024). Data Engineering in the Age of Large Language Models: Transforming Data Access, Curation, and Enterprise Interpretation. *Computer Fraud and Security*.
74. Shapiro, C., & Varian, H. R. (1998). *Information rules*. Harvard Business School Press.
75. Kalisetty, S. (2024). Deep learning frameworks for multi-modal data fusion in retail supply chains: enhancing forecast accuracy and agility.
76. Chen, Y., Li, X., & Chen, R. (2024). Federated multi-agent conversational AI for scalable customer platforms. *IEEE Access*, 12, 55621–55635.
77. Kushvanth Chowdary Nagabhyru. (2023). Accelerating Digital Transformation with AI Driven Data Engineering: Industry Case Studies from Cloud and IoT Domains. *Educational Administration: Theory and Practice*, 29(4), 5898–5910. <https://doi.org/10.53555/kuey.v29i4.10932>
78. Goodfellow, I. (2023). Privacy and ML systems governance. *Communications of the ACM*, 66(5), 56–63.
79. Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. *Deep Learning, and Explainable AI* (July 26, 2024).
80. ISO/IEC. (2019). *ISO/IEC 27001:2019 Information security management systems*. ISO.
81. Davuluri, P. N. Integrating Artificial Intelligence into Event-Driven Financial Crime Compliance Platforms.



82. Varshney, K. R. (2022). Trustworthy machine learning. Cambridge University Press.
83. Mehrabi, N., Morstatter, F., Saxena, N., et al. (2021). A survey on bias and fairness in ML. *ACM Computing Surveys*, 54(6), 115.
84. Kalisetty, S., & Singireddy, J. (2023). Optimizing Tax Preparation and Filing Services: A Comparative Study of Traditional Methods and AI Augmented Tax Compliance Frameworks. Available at SSRN 5206185.
85. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
86. Garapati, R. S. (2023). Optimizing Energy Consumption in Smart Build-ings Through Web-Integrated AI and Cloud-Driven Control Systems.
87. Shneiderman, B. (2022). Human-centered AI. Oxford University Press.
88. Agentic AI in Data Pipelines: Self Optimizing Systems for Continuous Data Quality, Performance and Governance. (2024). *American Data Science Journal for Advanced Computations (ADSJAC)* ISSN: 3067-4166, 2(1).
89. European Commission. (2024). Artificial Intelligence Act. Official Journal of the European Union.
90. Kalisetty, S. (2023). The Role of Circular Supply Chains in Achieving Sustainability Goals: A 2023 Perspective on Recycling, Reuse, and Resource Optimization. *Reuse, and Resource Optimization* (June 15, 2023).
91. IT Integration and Cloud-Based Analytics for Managing Unclaimed Property and Public Revenue. (2024). *MSW Management Journal*, 34(2), 1228-1248.
92. Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. *International Journal of Finance (IJFIN)-ABDC Journal Quality List*, 36(6), 653-674.
93. Johnson, A. E. W., Pollard, T. J., Shen, L., et al. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035.
94. Gottimukkala, V. R. R. (2023). Privacy-Preserving Machine Learning Models for Transaction Monitoring in Global Banking Networks. *International Journal of Finance (IJFIN)-ABDC Journal Quality List*, 36(6), 633-652.