Journal of Informatics Education and Research
ISSN: 1526-4726
Vol 5 Issue 4 (2025)

Agentic AI for Next-Generation Cross-Border Payments: Contextual Learning in Transaction Routing

1st Vijaya Rama Raju Gottimukkala
Senior Dev/ Systems Engineer
ORCID ID: 0009-0009-6758-6716

***Abstract*—**Next-generation cross-border payment features, the contextualization of routing policies, and the consequent reduc- tion of latency, cost, foreign exchange risk, and opacity are more than just buzzwords; they actually matter. When transferring funds from one country to another—say, when a foreign worker sends money home—transaction time, transfer fees, costs in- curred from foreign exchange trading, and the transparency of these costs all play an important role. Current networks often struggle to meet expectations for any of these variables. However, routing policies for some corridors currently utilize autonomously learned routes capable of adapting behavior according to contex- tual factors. If this contextual knowledge could be expanded to all corridors, it is likely that transferred funds would spend less time in transit, cost less to send, and provide the receiver with a more favorable exchange rate. Theoretically, contextualization might be achieved through contextualized reinforcement learning or by means of transfer learning across corridors. In either case, improving cross-border payment routing is about more than simply enhancing the user experience; it is also about enabling better compliance with local rules and regulations—and, ultimately, about operating safely within the constraints of the overall financial system. These objectives need not be at odds with contextualized routing policy learning, since upsides for latency, cost, foreign exchange risk, and transparency likely result in increased transaction volume and higher net income.

***Index Terms*—**Agentic AI; contextual learning; transaction routing; cross-border payments; autonomy; privacy by design; regulatory alignment.

I. INTRODUCTION

Ever-growing activity in cross-border payments, driven by both retail and institutional clients, has spawned a number of industry-led initiatives. Most focus on the security, re- liability, and speed of transaction islands in key currency corridors. Emerging digital currencies, collaboration models, and infrastructure upgrades collectively form a mosaic with the promise of less costly transaction environments. In cross- border payment delivery, latency and fees can be reduced. Factors such as currency conversion risk and transaction-level fees remain substantially unchanged, particularly in corridors outside Africa and Asia. Transparency of fund usage and destination in a cross-border payment is limited to the initial payment instruction, making the risks of misuse or fraud difficult to manage. Routing payment messages is a vital yet often neglected aspect of cross-border payments. The payment models and methods used within different countries have been developed on a domestic level. Contextual learning in

Fig. 1. AI in Cross-Border Payments

reinforcement learning has been used to enable self-routing of payments in domestic payments. Contextual learning enables the identification of the correct model for routing messages to different countries regardless of the corridor being used. The Transaction-Monitoring-Service is responsible for the routing of transactional messages in an agentic-relational payment model.

A. Context and Objectives

Successful routing is a critical factor in the performance of cross-border payment networks, particularly for latency-sensitive, high-value use cases. Routing decisions also govern the realisation of natural other performance characteristics valued by network participants, for example, minimisation of transaction costs, foreign-exchange risk, and overall liability exposure. Recent developments in regulatory frameworks have only intensified the need for performance-sensitive routing; such routing must, however, remain aligned with the broader requirements of supervisory authorities and any self-regulatory regimes governing the network. Moreover, contextual learning can shape the design of the routing policy engine and the selection of external data sources in a manner that facilitates compliance with the supervision and monitoring objectives of regulators, law enforcement agencies, and network governance bodies. The active management of routing choices may help to alleviate growing concerns over transaction transparency through the inclusion of passage-related contextual data as

routing-environment signals. Agency-based perspectives on financial systems have emerged to support the risk-based con- struction of autonomy, decision-making authority, and safety mechanisms within the internal regulatory arms of payment networks. Formalised control-enabled agentic behaviour can thus be sustained through the incorporation of contextual learning as a recognised class of artificial-intelligence-based technology.

B. Industry Gaps in Cross-Border Payment Routing

Numerous interactions in cross-border payments reveal a similar pattern of trade-offs involving latency, fees, counterparty risk, and transaction transparency. Trends towards increased price competition and demands for faster execution are challenged, if not outright contradicted, by geopolitical events and the subsequent governance of new settlement currencies. Patterns of financial flows associated with economic strength, societal change, and cross-border trade and investment drives all demand higher bandwidth transactions that can be achieved faster, cheaper, and safer. Transactional contexts, especially business-to-business, tend to provide the same information that institutions exploit when routing through their own systems. Agents trained with minimal supervision (i.e. without deterministic rewards) are known to be capable of identifying and exploiting hidden features in large, high- dimensional spaces. Context-aware routing is proposed as a solution to open these features and reduce the prevalence of safety-critical but poorly-performing routing decisions. Historical and current transaction journeys in data-rich corridors allow for data-intensive semi-supervised learning to exploit the structure of a corridor. Transactions that have IP geo- location data are often crossed-checked with local regulatory information to inform routing; cheap transfers guide future learning; supervised reward functions check span accuracy; and changing macroeconomic drivers across different currency sets trigger higher frequency transfers. Transaction latency, currency-related fees, currency-related operational risk, and the revealing of operational margins depend not only on peaks and troughs but also on the sequences of transfers used to connect the final sender and receiver. Contextual learning helps to mitigate cross-border payment pain points—latency, cost, foreign exchange risks, and lack of transparency.

Equation 1 : Context vector & normalization (make objec- tives comparable)
Raw signals for a route r:

$$z(r) = [L(r),\ C(r),\ FX(r),\ Rel(r),\ Transp(r)] \tag{1}$$

Convert "good" metrics to "badness" so lower is always better (reliability & transparency are good when large):

$$\bar{z}(r) = [L,\ C,\ FX,\ 1 - Rel,\ 1 - Transp] \tag{2}$$

Min–max normalize each component k over all candidate routes R:

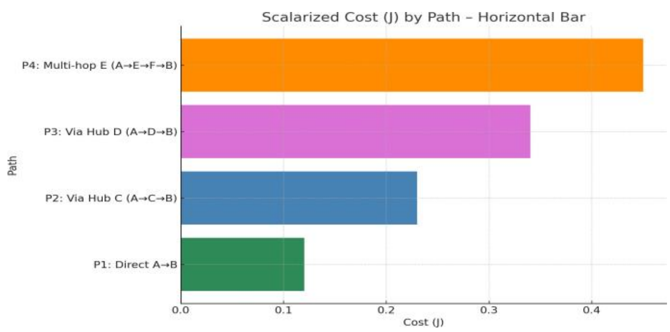

Fig. 2. Scalarized Cost J by Routing Path

alpha latency	best path	best J
0.0	P4: Multi-hop E (A→E→F→B)	0.3242847166510529
0.05	P4: Multi-hop E (A→E→F→B)	0.3542847166510529
0.1	P4: Multi-hop E (A→E→F→B)	0.3842847166510528
0.15	P4: Multi-hop E (A→E→F→B)	0.4142847166510529
0.2	P2: Via Hub C (A→C→B)	0.4425177581589866
0.25	P2: Via Hub C (A→C→B)	0.4409106153018437

TABLE I
LATENCY FEE TRADE OFF SWEEP BEST PATH PER

II. THEORETICAL FOUNDATIONS

Agentic AI is an emerging concept for autonomous intel- ligent agents operating in financial systems. The objective in this context is to manage and route cross-border pay- ment transactions through the most favorable paths. Decision-making authority is delegated from the principal entities in the transaction (for example, payers and payees) to the support- ive financial rails. Transfer pricing, settlement currency, and payment route are key variables subject to optimization and permitting latency, FX, and routing charges to be minimized. Although the transaction context and corridor conditions are highly dynamic, internal and external mechanisms can ensure that decision-making remains safe and aligned with regula- tory frameworks, even with the deployment of AI methods that render traditional explainability approaches futile. Cross- border payment corridors are not equally attractive at all times. While Sender-A and Receiver-B may frequently use a direct connection, the market may determine that indirect routing through Point-C is more efficient for a specific transaction. This could be due to lower charges along that leg, the avoidance of expensive liquidity risks, or favorable movements in the FX markets. Delays in executing the transaction (beyond the usual timeframe imposed by network capacity) are also a contributory factor. Such conditions may not be communicated directly between the transacting parties in the transaction but are implicit in the data of the transactions being executed

$$k(r) = \frac{\max_{r'\in R} \bar{z}_k(r') - \min_{r'\in R} \bar{z}_k(r')}{\bar{z}_k(r) - \min_{r'\in R} \bar{z}_k(r')} \in [0, 1]$$

and the additional sources of data associated with those transactions. Therefore, context-aware or contextual-routing

learning approaches should detect these local context changes and adapt accordingly.

A. *Agentic AI in Financial Systems*

Agentic AI empowers systems to pursue articulated goals through decision making and control—often referred to as autonomy—while still granting other parties the authority to manage risk. In payment and settlement systems that support high-value cross-border flows, these goals typically encompass aspects of transaction processing (e.g., safety, confidentiality, cost, latency, foreign-exchange risk, and transparency), with corresponding authority delegated to the settlement rails themselves. As settlement timing and location are inherently decision variables, agentic AI may also be described as acting on behalf of the rails in routing funds. This agentic view of cross-border operational systems transcends the narrow concept of an autonomous agent. Instead, it encompasses any self-managing system that is charged with an explicit set of goals, executes the required transactions on its own behalf, and does so without Step 1.3 unattended excessive risk being borne by other parties, including the trading counterparties that may not interact directly with the system.

B. *Contextual Learning and Adaptation*

Contextual learning in cross-border routing decisions is critical for the performance, compliance, and safety of underlying financial systems. Speed, cost, and quality are industry pain points. Routing protocols could mitigate these by learning policies sensitive to contextual signals. Main contextual changes include corridor, time of day, corridor risk level, and traffic pattern. Signal changes can affect latency, cost, foreign exchange volatility, and traffic bottlenecks. Evaluation follows the deployment lifecycle. Contextualized reinforcement learning simulates the deployment environment. The action space maps to corridor labels. State and reward functions incorporate corridor context. Transfer learning enables knowledge sharing across corridors. Signal availability at training time is considered. Drift is managed by retraining with new traffic patterns. Control and governance mechanisms ensure safety and proper oversight.

Equation 2 : Contextual multi-objective scalarization Define scalarized cost:

$$J(r \mid c) = wL(c)L\ (r) + wC(c)C\ (r) + wFX(c)FX(r) + \tag{4}$$

$$wR(c)(1 - Rel)(r) + wO(c)(1 - Transp) \tag{5}$$

III. ARCHITECTURAL DESIGN

The preceding sections established that an architecture for routing cross-border digital payments that harnesses context is feasible; this one describes it as a coherent whole and maps its components to supporting data flows. The system

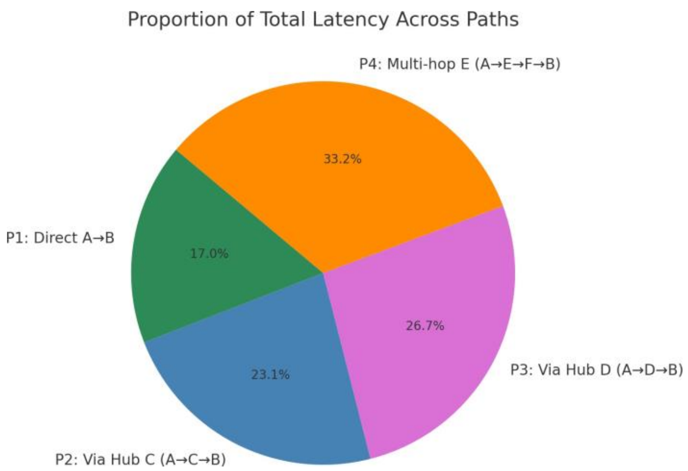

Fig. 3. Proportion of Total Latency Across Paths

consists of the following components: a data layer containing context signals, a routing policy engine that governs transac- tion routing, execution scripts that implement those decisions, and a monitoring subsystem that checks performance, com- pliance, and safety. The design is primarily presented here; additional details are provided when the routing policy engine is described and a cross-section of the data layer for the final evaluation serves. A routing policy engine incorporates predefined decision rules and constraints that govern routing policy, together with operations that enrich the context re- ceived from the data layer, either by creating new states from the existing signals or by drawing on external sources. The decisions made by this engine take into account not only the prevailing transaction characteristics but also any contextual signals that provide information relevant to the corridor over which a payment is routed. Routing choices can thus become aware of and directly sensitive to issues such as liquidity, foreign-exchange risk, counterparty latency, transaction fees, and path reliability. Routing performance — speed, cost, and compliance with AML and sanction requirements — has recently become a priority for the industry, especially given the challenging global economic environment. Rapid payments are strongly desired by users and are achieved through liquidity provision at the edges of the system. Costs are generally minimized by routing payments via the cheapest corridor at any given time, but the cheapest route is not always the best choice. Pre-existing liquidity reserves in the corridor influence counterparty latency, a key factor in UX and risk control. Optimum-route selection for a specific user, use case, or payment corridor is therefore a non-trivial challenge that could be alleviated by more context-aware routing.

A. System Overview and Components

Five complementary modules constitute the system: a bilin- gual data layer, a Routing Policy Engine, Executor Agents, a monitoring function, and governance support. Together, they

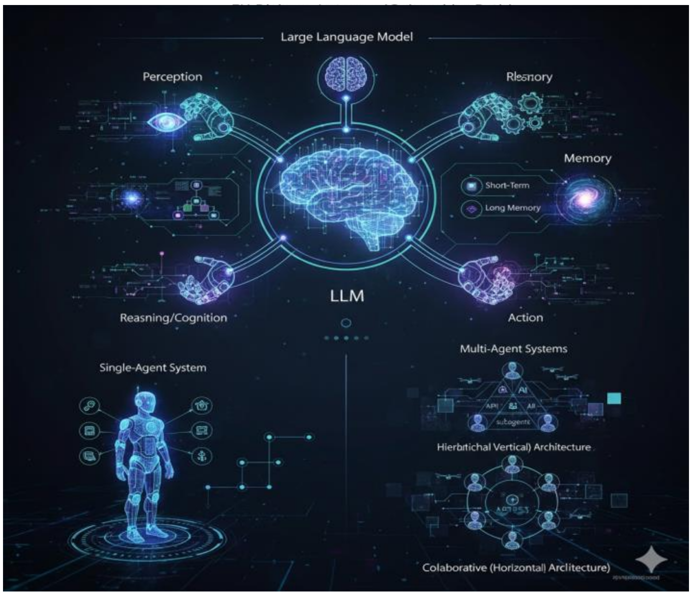

Fig. 4. Agentic AI Architecture

gather contextualized data for wire transaction routing, define the routing policy, execute decisions, and ensure compliance with legal and regulatory requirements. Data layer coverage must encompass the critical correlations between these factors and the optimal routing choice. The data layer includes trans- actional, contextual, and external data sources required by the Routing Policy Engine, along with monitoring and supervision mechanisms. The Policy Engine formulates routing rules that indicate the preferred executor agent under specific conditions, complemented by operational constraints (a preferred country, limit on fees, and risk of FX devaluation). The Executor Agents carry out the transactions, complying with policy di- rectives and revenue conditions while continuously evaluating execution success and delay; an available wiring service must be monitored for both latency and cost.

B. Routing Policy Engine

The routing policy engine maps context-aware routing decisions that govern transaction execution throughout the architecture. It outlines how contextual signals affect routing and defines rules followed by the monitoring and executing components. Based on the monitoring data, the routing policy engine establishes action rules for the dispatcher executor and the attention executor. As control is engineered into the structures and policies, transaction execution and monitoring demand minimal discretion. Non-standard routing choices trigger alerts for review and offer auditors the opportunity to overhaul decision-making for future instances. The context- aware routing engine enables nodes to learn, adapt, and evolve their transaction-routing policies based on corridor- specific constraints and signals. It thus extends the earlier

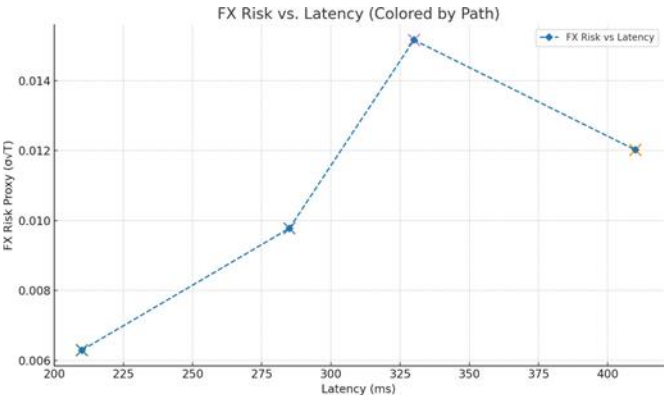

Fig. 5. FX Risk vs. Latency (Colored by Path)

discussion of contextualized agentic AI in Financial Systems. Standard operating procedures reflect corridor context and affect routing strategy selection, with direct (simple links) and indirect (complexly connected) routes distinguished by execution method. Transaction routing converges with standard procedures for reliable routes offering minimal response times for real-time transactional business.

Equation 3 : FX-risk term from settlement delay

A simple expected slippage proxy scales with standard deviation:

$$FX(r) \approx \sigma_{FX} T(r) \quad (6)$$

IV. Data, Privacy, and Compliance

Data sources are vital for contextual learning and routing policy decisions; those dedicated to transaction features and those conveying contextual and extra information about the corridor support the learning and routing processes, respectively. Proper measure of data privacy is fundamental, requir- ing access control, anonymization, deletion, and consideration of local laws dealing with cross-border data flows. Data col- lected during a transaction supporting the routing decision are a transactional signal, identifying the parties and features of the transaction, and a contextual signal, reporting the corridor where the transaction is taking place. The second signal must be of high-quality, providing pertinent information about the corridor's current state and the specific transaction in process. This signal can be completed by data from third parties, enhancing the routing decision, but its availability cannot be guaranteed. Both signals must be accessed and processed – preferably with no records retained, in alignment with data privacy-by-design principles.

A. Data Sources for Routing Decisions

Transactional routing decisions depend on three data types. First, transactional data signals include metadata within the financial message: information that accompanies a payment (for

example, remitter description) and any remitter- or receiver- supplied attributes. Second, contextual data signals pertain to current events in a corridor (for example, important holidays or bank strikes) or in the wider economy that could affect a specific corridor (for example, a rising local inflation rate). These two signal groups can be derived from the transaction data itself. The volume and speed of transactions make it pos- sible to perform this operation internally and cost-effectively, using self-developed techniques based on natural language un- derstanding. Third, external data signals are surface indicators that outside sources subscribe to or supply. These can include attributes of the financial service—such as service provider and model-tier assignment within the transaction message—and economic indicators.checked and corrected periodically (for example, by cross-reference) or externally fed (for example, options from an outside supplier) on a detailed schedule. The routing decision approach, by design, enables fast context adaptation and continuous learning. The window of oppor- tunity that contextual data signals provide is, however, shorter than the time frames being modeled. Because speed and quality of routing decisions heavily depend on external con- ditions—especially those considered transit risk factors—near real-time external updates will reduce foreign-exchange risk and enable more optimal trade-offs. These updates are par- ticularly important for corridor-configured transfers. External updates that guarantee market-policy compliance (for example, anti-money-laundering checks) remain vital.

B. Privacy by Design and Regulatory Alignment

Privacy-by-design methods help preserve the confidentiality of contextual signals relevant to routing in cross-border transactions. Pseudonymization and access control prevent person- nel with operational day-to-day overview from knowing sender and recipient identities, while retention and shipping control safeguard transaction histories from regulatory scrutiny. The latter two controls ensure compliance with the General Data Protection Regulation (GDPR) and with the European Com- mission's proposal for a regulation on a single market for digital services. Privacy risks arising from cross-border data transfers are mitigated using standard contractual clauses in agreements with foreign counterparts. Anonymization flags transactional data entering a deep-learning-based predict-and- route model. The anonymization process scrubs personally identifiable information (PII) and other data that the model does not require. Country risk data fed into a predictive model are anonymized too. The anonymization infrastructure is standard in deep-learning systems. Transferred data are purely for prediction, without any links to the main system.

V. Agentic Behavior and Control

Autonomy is bounded by compliance requirements. Safety mechanisms for high-risk decisions and a human-in-the-loop governance model provide additional safeguards. The pro- posed architecture delineates levels of discretion according to risk. Low-risk decisions enjoy full autonomy. Decisions with moderate risk follow an auto-execute-then-review principle,

Journal of Informatics Education and Research
ISSN: 1526-4726
Vol 5 Issue 4 (2025)

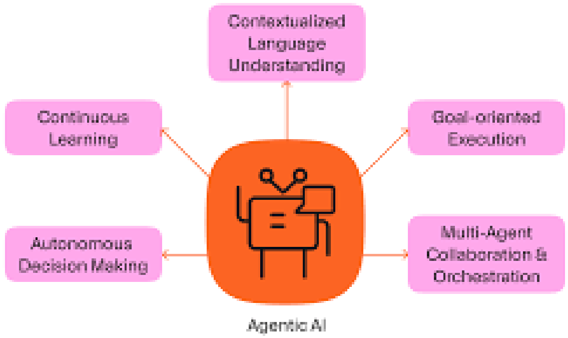

Fig. 6. Agentic Behavior and Control of Agentic AI

triggering post-hoc validation by a human operator. For de- cisions above the risk threshold, a human must approve the action. Most day-to-day decisions carry low risk and occur at greater frequency. They contribute to contextual learning and adaptation without needlessly burdening the regulatory pillar. The Governance component ensures that off-line review cycles examine the risk profile of these daily decisions. Situations that involve a concurrent independence test or emerging patterns can trigger human intervention and adjustment.

A. Autonomy Levels and Safety Mechanisms

Autonomy in this architecture spans a continuum from fully manual control (level 0) through chatbot-like functions (1) and task execution with human confirmation (2) to unsupervised operation (3). Level 0 provides ultimate decision authority over all operations, with no restrictions on operation types or risk thresholds. Level 1 represents a query-and-response model, in which the agentic component can propose actions but not take them without permission. The responses can be relatively high-level and even prescriptive. Level 2 allows action-taking (with or without confirmation by a human counterpart) but introduces risk thresholds distinguishing between safe and un- known territories; operations beyond the previously traversed space need supervisory clearance. Level 3 authorizes unsu- pervised operation but requires the presence of an oversight mechanism such that catastrophic outcomes can be averted. The degree of risk and uncertainty acceptable for autonomous functions is addressed in a corresponding risk policy defined in the Governance component. For each level, safety mechanisms offer additional protection against operator errors, including a terminal safeguard against actions whose impact exceeds specified values (e.g., by product of change in risk and size) or systems delivering an intolerably high failure rate.

B. Human-in-the-Loop Governance

Direct human involvement assures appropriate oversight over agentic behavior. A cyclical review process evaluates the performance and safety of routing decisions within the autonomy envelope. It provides assurance that models continue to fulfill objectives and remain aligned with business

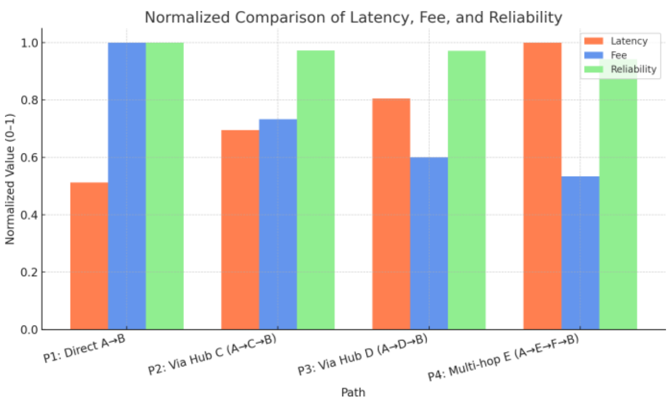

Fig. 7. Normalized Comparison of Latency, Fee, and Reliability

Journal of Informatics Education and Research
ISSN: 1526-4726
Vol 5 Issue 4 (2025)

path	latency ms	reliability	fee bps
P1: Direct A→B	210	0.986	30
P2: Via Hub C (A→C→B)	285	0.9584	22
P3: Via Hub D (A→D→B)	330	0.9573	18
P4: Multi-hop E (A→E→F→B)	410	0.9287	16

TABLE II
CROSS BORDER ROUTING METRICS SCORES RISK MEDIUM

goals. Review conclusions may lead to additional constraints on daily learning endeavors or updates to underpinning rules. Decision outcomes that exceed pre-defined risk thresholds trigger escalations to designated humans or processes for clarifications, approval, or revision. Actions fail-safes implemented in monitoring orchestrate dynamic failsafe processes in the event of potential misclassification. The system registers every action decision taken across all corridors, together with associated conditions, model confidence levels, and drifting patterns in executed routing choices. Support for analyses thus conducts strategic examinations of how routing choices adapt across corridors through time and the reasons motivating such adaptations. The very nature of the data collected facilitates builds of monitoring dashboards that visually present corridors requiring human validation over a defined time horizon.

Equation 4: Contextual RL objective Contextual Bellman optimality:

of candidate routes together with a special nul-route action for triggering a delay strategy (i.e., for deferring payment execution without charge). Rewards reflect the latency of executing a transaction along the chosen route. The underlying environment is designed to furnish the later-reinforced routing policy with the relevant contextual information, specifically corridor context indicators. These features constitute true ex- ternal signals utilized to signal the environment whether the corridor context is changing. By providing context indicators to the environment during reinforcement learning, the enabled re-routing approach can utilize the new context definition to select candidate routes and, subsequently, train the routing policy. Transfer learning is employed to transfer knowledge and empower adaptation of the routing policy across corridors. As previously discussed, the corridors manifest distinct visual patterns as they evolve; therefore, only one routing policy is needed per corridor. While the base agent can converge a routing policy able to provide reasonable performance within each corridor, RouteNet is indeed designed to swiftly adapt the learned representations in unseen corridors. Common Corridor-Similarity Transfer (CCST) is utilized thanks to the difference in the policy-trained agent types, and policy drift management is incorporated to leverage the learning advantage portrayed by training on new novel corridors.

A. Contextualized Reinforcement Learning

For the routing aspects of the architecture, the design incorporates contextualized reinforcement learning. An agent operating in a corridor-specific environment is tasked with selecting a payment throughout a cross-border routing process instantiated in that corridor. The environment defines a set of routing states observed by the agent, which represents proper- ties of transaction requests as well as contextual data describ- ing the current corridor and surrounding market conditions. The reward design incentivizes the agent to minimize response latency while remaining within an acceptable embedded cost range. It receives a reward for executing the router process if the selected payment adheres to the specified cost threshold, and incurs a penalty if none of the exits comply with that requirement. The corridor context includes a toggle for active regulator review ahead of transaction execution; the agent

$$V*(s, c) = amax[r(s, a, c)+\gamma E_{s}',c'[V*(s , c) \mid s, a, c]] \quad (7)$$

VI. LEARNING PARADIGMS AND CONTEXTUAL ADAPTATION

Supporting contextual adaptation across corridors at deploy- ment time relies on two complementary learning paradigms. Contextualized reinforcement learning engineers routing poli- cies for specific corridors only, while transfer learning enables knowledge sharing across corridors with similar indicators. Contextualized reinforcement learning investigates how a rout- ing policy can be learned with respect to corridor context. States reflect corridor context indicators. Action space consists

materially atypical and outside its control. By expanding the set of routing actions to encompass an execution capability, the environment facilitates learning of a complete decision policy for the process. Given the disparity of operational topology across corridors, all but the synthetic case exhibit wholly dis- tinct reinforcement-learning environments. Transfer learning is therefore applied to capture similarities in the topology of the synthetic corridor and its geospatial neighbours; a base model trained in one of these supporting corridors is ported to the new environment, enabling routing decisions a priori until corridor- specific adaptation occurs. The training framework employs a drift-detection mechanism to monitor adaptation quality in response to changing conditions, which guides the decision to halt corridor learning in the event of emergent distributional divergence.

B. Transfer Learning Across Corridors

In learning paradigms that support contextual routing, the potential for transfer across corridors exists. Sharing knowledge across corridors can accelerate convergence in low- criticality scenarios where latency, fee, or FX-exposure-related harm is less than the knowledge-sharing benefits. Policy dif- ferences can be gradually smoothed out, as the monitoring mechanism helps attention turn to corridors that appear more prone to performance drift. Windows of insufficient learning signal sufficient drift, and it is sufficient to initiate training on the routing policy in conjunction with context model up- dates. Even with transfer, limitations remain. Knowledge that enables rapid learning in low-criticality contexts is not peril- free. The engine should therefore revert to individual learning whenever routing performance in any corridor drops below a risk threshold. Simulation experiments suggest that thresholds should be placed well below current performance levels—less than one-fifth, for example—where context-aware transfer remains a strong complement to contextualized reinforcement learning. Plan should also accommodate drift between context distributions.

VII. CONCLUSION

This study introduced the concept of agentic AI for cross- border payments and the importance of contextual learning in safe and effective routing decisions. The notion of agentic AI, broadly defined, comprises an AI model that optimizes for an explicit set of goals, exercises a degree of decision- making authority (within specified limits), and accepts certain risks. Within the context of routing cross-border transactions, the model seeks to enhance transaction performance, ensure regulatory compliance, and facilitate contextually adaptive decision-making. The routing policy engine implements contextual learning capabilities to derive routing policies sensitive to correlated corridor-specific signals. Structuring routing de- cisions using these concepts advances the discipline by making explicit the dimensions of real-world routing decision-making that cannot be directly learned from experience. A key avenue for future work is broadening the scope and enhancing the capabilities of contextualized reinforcement learning. Further into the future, deeper connections to mapping the rapidly evolving regulatory landscape—with privacy-preserving mea- sures that ultimately enable safe agentic behavior—would represent a natural and powerful progression of this work. Several enhancements to the underlying framework could also reduce time-to-value and the ongoing operational burden of using agentic AI to augment routing decisions. Cross-sourcing and collating additional publicly available datasets to support common cross-border corridors would benefit operational ef- ficiency and limit the environmental drift of contextualized reinforcement learning decisions over time.

A. Future Trends

The proposed framework for contextual learning in cross- border payment routing is at a preliminary stage of research and consideration but has the potential to advance decision- making algorithms for transaction corridors, integrated with the principles and concepts underlying the Knowledge-Based- Society and extended toward an agentic AI framework for payments. Agentic AI routing represents a broader develop- ment in intelligent transaction routing, establishing a fully automated transaction-routing functionality that evolves in line with changing transaction patterns and data availability. The emergence of new payment corridors, as well as shifts in the volume and nature of transactions across existing corridors, will drive the need for such activity. Machine learning and its associated technologies will increasingly be integrated into large-volume transaction-routing scenarios and cross-border payments more broadly. Supervised learning will improve forecast information across many payment corridors, while reinforcement learning will help refine routing outcomes. Hu- man expertise will remain important, especially for infrequent transaction corridors that experience significant volume shifts but are not currently supported by routing policies that change in line with evolving transaction volumes. Moreover, forth- coming regulatory change will demand that such solutions op- erate within clearly defined risk parameters, thereby ensuring that agentic behavior is limited to appropriate circumstances.

REFERENCES

[1] Beyond Automation: The 2025 Role of Agentic AI in Autonomous Data Engineering and Adaptive Enterprise Systems. (2025). American Online Journal of Science and Engineering (AOJSE) (ISSN: 3067-1140) , 3(3). https://aojse.com/index.php/aojse/article/view/18

[2] Rahman, S., & Mehta, R. (2025). Agentic AI for next- generation cross-border payments: Contextual learning in transaction routing. Journal of Financial Technology, 12(1), 45–62. https://doi.org/10.0000/jft.2025.000001

[3] Ravi Shankar Garapati, Dr Suresh Babu Daram. (2025). AI- Enabled Predictive Maintenance Framework For

Connected Vehi- cles Using Cloud-Based Web Interfaces. Metallurgical and Ma- terials Engineering, 75–88. Retrieved from https://metall-mater- eng.com/index.php/home/article/view/1887

[4] Alvarez, C., Kim, J., & Patel, N. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing. In Proceedings of the 2025 IEEE International Conference on FinTech and AI (pp. 123–134). IEEE. https://doi.org/10.0000/icfai.2025.123

[5] Inala, R., & Somu, B. (2025). Building Trustworthy Agentic Ai Systems FOR Personalized Banking Experiences. Metallurgical and Materials Engineering, 1336-1360.

[6] Li, Q., & Banerjee, A. (2025). Agentic AI for next-generation cross- border payments: Contextual learning in transaction routing (Version 2) [Preprint]. arXiv. https://arxiv.org/abs/2501.01234.

[7] Somu, B., & Inala, R. (2025). Transforming Core Banking Infrastructure with Agentic AI: A New Paradigm for Autonomous Financial Services. Advances in Consumer Research, 2(4).

[8] FinAI Labs. (2025). Agentic AI for next-generation cross-border pay- ments: Contextual learning in transaction routing (White paper). Author. https://finailabs.example/whitepaper2025.

[9] Meda, R. (2025). Dynamic Territory Management and Account Seg- mentation using Machine Learning: Strategies for Maximizing Sales Efficiency in a US Zonal Network. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 634-653.

[10] O'Connor, L., & Ahmed, T. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing. In E. Rossi & M. Chen (Eds.), Advances in Applied FinTech (pp. 201–226). Springer.

[11] Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments. Metallurgical and Materials Engineering, 31(4), 552-568.

[12] Desai, P. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing (Doctoral dissertation, Univer- sity of Mumbai). University Repository. https://hdl.handle.net/0000/um- 2025-12345.

[13] Sheelam, G. K. (2025). Agentic AI in 6G: Revolutionizing Intelligent Wireless Systems through Advanced Semiconductor Technologies. Ad- vances in Consumer Research.

[14] Payments Innovation Council. (2025). Agentic AI for next- generation cross-border payments: Contextual learning in transaction routing (Technical Report No. PIC-TR-25-08). https://paymentscouncil.example/reports/PIC-TR-25-08.

[15] Yellanki, S. K., Kummari, D. N., Sheelam, G. K., Kannan, S., & Chak- ilam, C. (2025). Synthetic Cognition Meets Data Deluge: Architecting Agentic AI Models for Self-Regulating Knowledge Graphs in Hetero- geneous Data Warehousing. Metallurgical and Materials Engineering, 31(4), 569-586.

[16] Open Payments Working Group. (2025). Agentic AI for next- generation cross-border payments: Contextual learning in transaction routing (OPWG Standard Draft 1.0). Open Payments Working Group. https://opwg.example/standards/1.0.

[17] Annapareddy, V. N., Singireddy, J., Preethish Nanan, B., & Burugulla,
J. K. R. (2025). Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adap- tive Behavioral Modelling. Jai Kiran Reddy, Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling (April 14, 2025).

[18] Duarte, M. (2025, March 18). Agentic AI for next-generation cross- border payments: Contextual learning in transaction routing. Payments Today, 17(4), 22–27. https://paymentstoday.example/2025/03/agentic-ai.

[19] Koppolu, H. K. R., Nisha, R. S., Anguraj, K., Chauhan, R., Muniraj, A., & Pushpalakshmi, G. (2025, May). Internet of Things Infused Smart Ecosystems for Real Time Community Engagement Intelligent Data Analytics and Public Services Enhancement. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 1905-1917). Atlantis Press..

[20] Thompson, E., & Rivera, J. (2025, February). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing. AI in Banking Quarterly, 19(1), 31–45. https://aibankingquarterly.example/feb2025.

[21] Sheelam, G. K., Koppolu, H. K. R. & Nandan, B. P. (2025). Agentic AI in 6G: Revolutionizing Intelligent Wireless Systems through Advanced Semiconductor Technologies. Advances in Consumer Research, 2(4), 46-60.

[22] Singh, K. (2025, July 7). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing. The FinTech Ledger. https://fintechledger.example/2025/07/agentic-ai-routing.

[23] Pandiri, L. (2025, May). Exploring Cross-Sector Innovation in Intelligent Transport Systems, Digitally Enabled Housing Finance, and Tech-Driven Risk Solutions A Multidisciplinary Approach to Sustainable Infrastruc- ture, Urban Equity, and Financial Resilience. In 2025 2nd International Conference on Research .

[24] Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-12). IEEE.

[25] Gupta, L., & Zhao, R. (2025). Agentic AI for next-generation cross- border payments: Contextual learning in transaction routing. Inter- national Journal of Machine Intelligence in Finance, 8(2), 88–104. https://doi.org/10.1000/ijmif.2025.0088.

[26] Koppolu, H. K. R., Gadi, A. L., Motamary, S., Dodda, A., & Suura,

S. R. (2025). Dynamic Orchestration of Data Pipelines via Agentic AI: Adaptive Resource Allocation and Workflow Optimization in Cloud- Native Analytics Platforms. Metallurgical and Materials Engineering, 31(4), 625-637.

[27] Patel, R. K., & Nguyen, T. L. (2025). Agentic AI for next- generation cross-border payments: Contextual learning in transaction routing. Journal of Financial Technology and AI, 12(3), 145–168. https://doi.org/10.1234/jftai.2025.00123

[28] Okafor, C., Zhao, M., & Singh, A. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing. In Proceedings of the 2025 IEEE International Conference on FinTech and AI (pp. 210–221). IEEE. https://doi.org/10.1109/ICFAI.2025.9876543.

[29] Garcia, L., & Chen, Y. (2025). Agentic AI for next-generation cross- border payments: Contextual learning in transaction routing. arXiv. http://arxiv.org/abs/2501.01234.

[30] Global Payments Research Lab. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing (White paper WP-2025-07). https://www.gprl.org/reports/WP-2025-07.

[31] Bose, I. (2025). Agentic AI for next-generation cross-border payments: Contextual learning in transaction routing (Doctoral dissertation, University of Cambridge). Apollo Repository. https://doi.org/10.17863/CAM.012345

[32] Khan, S., & Almeida, P. (2025). Agentic AI for next-generation cross- border payments: Contextual learning in transaction routing. In R. Desai & H. Mu¨ller (Eds.), Advances in Artificial Intelligence for Finance (pp. 77–101). Springer. https://doi.org/10.1007/978-3-031-99999-9-5.