Int. J. AI Adv. Comput. 1 (2026) 43-51

Contents lists available at AcademiansEdu

International Journal of AI and Advanced Computing

journal homepage:

Research Article

AI Multi-Agent Reinforcement Learning for Conflict Resolution and Forecasting in International Relations Policy

Miller

Independent Scholar

A R T I C L E I N F O

Keywords: Fake news detection, misinformation, machine learning, ensemble learning, hybrid classification, Logistic Regression, Random Forest, XGBoost, natural language processing, text classification, TF-IDF, SMOTE, class imbalance, model interpretability, SHAP, cross-domain generalization, adversarial robustness, social media analysis

A B S T R A C T

In an increasingly interconnected yet volatile world, the emergence of complex global challenges necessitates innovative approaches to conflict resolution and accurate forecasting of geopolitical developments. Traditional analytical methods in international relations often struggle to capture the dynamic and multifaceted interactions among nation-states and non-state actors, highlighting the urgent need for more sophisticated modelling tools that can account for strategic behaviour, emergent phenomena, and the nuanced interplay of diverse objectives. This paper explores the synergistic integration of Multi-Agent Reinforcement Learning (MARL) and Large Language Models (LLMs) as a novel algorithmic framework designed to advance the fields of proactive diplomacy and evidence-based policymaking. By leveraging the predictive capabilities of MARL within complex international relations simulations and combining them with the nuanced interpretive power of LLMs, this research proposes a comprehensive approach for analyzing geopolitical dynamics, simulating diplomatic negotiations, and optimizing strategic interventions. This integration facilitates a more holistic understanding of complex international phenomena, allowing for the analysis of emergent social outcomes from both macro-level trends and micro-level interactions, thereby illuminating the causal mechanisms underpinning international events and predicting the ramifications of various policy interventions.

The proposed framework enables the development of human-like agents capable of executing comprehensive multi-agent missions encompassing strategic planning, goal-oriented negotiation, and sophisticated social reasoning. These LLM-based agents can refine their strategies through self-play and memory augmentation, leading to continuous strategic evolution without direct human intervention and allowing for the rigorous evaluation of policy decisions in a simulated environment before real-world implementation. This approach not only enhances the quantitative assessment of geopolitical factors but also provides rich qualitative insights into individual-level social mechanisms, effectively bridging interpretability and predictability in international relations research. By offering a scalable and adaptable framework for understanding intricate international dynamics, these models empower policymakers to explore a multitude of scenarios and potential policy outcomes in a safe, simulated environment, thereby optimizing diplomatic initiatives for greater efficacy and mitigating unforeseen negative consequences. The continuous refinement of these models, incorporating lessons from real-world events and expert geopolitical analysis, ensures their sustained relevance and accuracy in an ever-evolving international landscape, ultimately contributing to the development of more robust, ethically sound, and effective strategies for fostering global peace and stability.

1. INTRODUCTION

The discipline of international relations grapples with the daunting task of understanding and managing a world characterized by deep interconnectedness, persistent volatility, and the emergence of complex global challenges that transcend national borders. From climate change and pandemics to economic instability and regional conflicts, the issues confronting policymakers today are inherently multi-faceted, involving a diverse array of actors with often conflicting objectives and strategic capabilities. Traditional analytical methods in political science and international relations, while providing valuable frameworks for understanding historical patterns and broad systemic dynamics, often struggle to capture the dynamic and multifaceted interactions among nation-states and non-state actors in a rigorous, predictive manner. These conventional approaches

https://doi.org/

Received 23 November 2025; Received in revised form 11 January 2026; Accepted 29 January 2026

Available online 27 February 2026

frequently rely on static models, simplified assumptions about rational behaviour, and linear projections that fail to account for the emergent phenomena, strategic adaptation, and feedback loops that characterize real-world geopolitical systems. This analytical gap highlights an urgent need for more sophisticated modelling tools capable of simulating the complex decision-making processes inherent in international relations, where diverse actors pursue often conflicting objectives within a dynamic and uncertain global landscape.

Multi-Agent Reinforcement Learning has emerged as a promising paradigm capable of addressing these analytical limitations. MARL extends conventional reinforcement learning to scenarios where multiple autonomous agents interact within a shared environment, each learning optimal policies through iterative trial and error, adapting their strategies based on the unfolding actions and reactions of other agents. This approach is particularly adept at simulating the intricate strategic interactions that define international relations, as it allows for the modelling of diverse actors with distinct goals, capabilities, and learning processes operating within a dynamic and interdependent system. The application of MARL in this domain enables the sophisticated modelling of strategic interactions, allowing researchers and policymakers to explore emergent behaviours, identify potential stable equilibria or pathways to cooperation even amidst divergent national interests, and rigorously test the potential outcomes of different policy interventions before they are deployed in the real world.

Complementing the strategic and predictive capabilities of MARL, Large Language Models offer an unprecedented capacity for nuanced interpretation and generation of human language, capturing the subtleties of diplomatic communication, cultural context, and strategic rhetoric that are central to international relations. The integration of LLMs with MARL creates a powerful synergy, endowing simulated agents with the ability not only to learn optimal strategies through interaction but also to communicate, negotiate, and reason in ways that mirror human diplomatic behaviour. This fusion allows for the creation of human-like agents capable of executing comprehensive multi-agent missions encompassing strategic planning, goal-oriented negotiation, and sophisticated social reasoning, thereby generating simulations that are both quantitatively rigorous and qualitatively rich. Such an integrated framework facilitates a more comprehensive understanding of complex international phenomena by enabling the analysis of emergent social outcomes from both macro-level trends and micro-level interactions, a dual analytical lens that is crucial for illuminating the causal mechanisms underpinning international events and for accurately predicting the ramifications of various policy interventions.

This paper will delve into the profound potential of integrating Multi-Agent Reinforcement Learning with Large Language Models to forge a new era of proactive diplomacy and evidence-based policymaking. It will explore how MARL frameworks can be leveraged to understand, predict, and potentially mitigate international conflicts, providing a novel perspective on global diplomacy and policy formulation. By simulating the strategic interactions and learning processes of various international actors, from sovereign states and international organizations to non-state actors and multinational corporations, MARL can illuminate pathways towards cooperation, defuse potential conflicts, and optimize policy interventions in an increasingly complex global arena. This capability offers a significant advantage over traditional analytical methods by providing a data-driven approach to understanding the underlying mechanisms of international relations, thereby enhancing the capacity for proactive policy intervention and fostering a more stable and peaceful global order. The subsequent sections of this paper will systematically explore the theoretical foundations, methodological frameworks, empirical findings, and profound implications of this innovative approach for the future of international relations and global governance.

2. LITERATURE REVIEW

The intersection of artificial intelligence and international relations has garnered increasing scholarly attention, with researchers exploring the potential of advanced computational methods to model, simulate, and forecast complex geopolitical dynamics. The literature reveals a growing recognition that traditional analytical frameworks, while valuable, often fall short in capturing the emergent, adaptive, and multi-layered nature of contemporary international interactions, thereby creating a fertile ground for the application of AI methodologies like Multi-Agent Reinforcement Learning and Large Language Models.

Early efforts to model international relations computationally often relied on system dynamics and game theory, providing valuable insights into strategic interactions but frequently simplifying the complexity of real-world actors and their learning processes. Integrated Assessment Models, for instance, have been traditionally employed to analyze long-term global challenges such as climate change, but they often overlook the intricacies of real-world negotiations by failing to simultaneously address multiple objectives and adaptive agents, thereby limiting their applicability in dynamic policy contexts. MARL, in contrast, offers a robust framework for such multi-objective, multi-agent scenarios, enabling a more nuanced understanding of strategic interactions in international diplomacy by allowing agents to learn and adapt their policies based on the unfolding consequences of their collective actions.

The application of MARL in international relations research has gained significant momentum, with studies demonstrating its versatility in capturing the varied nature of international interactions, from fully cooperative endeavours to intensely competitive confrontations. This methodological flexibility is evidenced by diverse applications ranging from optimizing cooperative systems to mastering complex, multi-player games like Diplomacy and Go, which require sophisticated strategic reasoning, negotiation, and the ability to anticipate opponents' moves. In the realm of international policy, models like RICE-N, a multi-region integrated assessment model enhanced with MARL, exemplify how this approach can simulate global climate and economic interactions to design and evaluate strategic outcomes for various negotiation frameworks, providing a powerful tool for exploring pathways to international climate cooperation.

The integration of Large Language Models into these simulation environments represents a recent and transformative advancement. LLMs endow simulated agents with the ability to process and generate natural language, enabling them to engage in realistic diplomatic communication, articulate nuanced positions, and adapt their rhetoric based on the evolving negotiation context. WarAgent, an LLM-powered multi-agent AI system, has demonstrated the capacity to simulate historical international conflicts such as World War I and II, offering novel perspectives on the triggers and conditions leading to war by allowing for dynamic "what-if" analyses that explore how alternative diplomatic interventions might have altered the course of history. Similarly, GovSim explores how societies of AI agents balance resource exploitation with sustainability, revealing the critical role of successful multi-agent communication for achieving cooperative outcomes in resource management dilemmas that mirror real-world challenges like managing shared water resources or combating overfishing.

The research further highlights the importance of communication protocols and reward structures in shaping the propensity for agents to engage in collaborative versus competitive behaviours. Studies have demonstrated that successful multi-agent communication is critical for achieving sustainable cooperation, particularly in contexts involving complex ethical considerations and long-term strategic planning. The ability of LLM-based agents to generate flexible, human-like behaviours allows for more realistic simulations, thereby enhancing the realism of policy evaluations and providing deeper insights into complex, real-world decision-making processes. Furthermore, these advanced simulations can explore the nuanced escalation risks inherent in military and diplomatic decision-making, particularly with the rise of autonomous AI agents in high-stakes contexts, offering a crucial platform for stress-testing crisis management protocols and identifying potential pathways to unintended conflict.

The literature also underscores the importance of interpretability and the quantification of uncertainty in these MARL models, ensuring their utility in real-world policy formulation, especially for forecasting policy or action pathways towards desired outcomes in global-scale, low-agent number scenarios. The ability to not only predict outcomes but also to understand the underlying reasons for those outcomes is paramount for building trust in AI-generated policy recommendations and for enabling human policymakers to exercise informed oversight. Moreover, the critical examination of emergent behaviours allows researchers to develop robust safeguards and ethical guidelines for the deployment of AI in international relations, ensuring that these powerful tools are used responsibly and in alignment with human values. The ongoing development of advanced AI systems, particularly those incorporating LLMs, offers new avenues for establishing credible commitments and mutual transparency among nations, which are vital for mitigating strategic uncertainty and fostering a more stable and cooperative international environment. This growing body of research collectively points towards a future where AI-driven simulations become an indispensable tool for diplomats, policymakers, and scholars, augmenting human judgment with the analytical power to explore complex strategic landscapes and craft more effective and resilient international policies.

3. METHODOLOGICAL FRAMEWORK

The methodological framework for applying Multi-Agent Reinforcement Learning integrated with Large Language Models to model international relations is built upon a multi-layered architecture that encompasses agent representation, environment design, learning algorithms, and validation protocols. This framework is designed to capture the fundamental complexity of international politics, where diverse actors with distinct goals, capabilities, and strategic cultures interact within a dynamic and uncertain global system, producing outcomes that are often emergent and difficult to predict using traditional analytical methods.

The foundation of this framework lies in the representation of key actors in international relations as autonomous agents, each endowed with distinct objectives, policy instruments, learning capabilities, and a sophisticated communication module powered by a Large Language Model. These agents are not merely rational utility-maximizers in the classical economic sense; rather, they are designed to exhibit a range of behaviours informed by historical patterns, cultural contexts, and strategic considerations. For instance, an agent representing a major power might be programmed with objectives related to national security, economic prosperity, and global influence, utilizing policy instruments such as diplomatic engagement, economic sanctions, military posturing, and treaty negotiations. An agent representing a smaller state might prioritize sovereignty, regional stability, and access to international aid, employing different strategic approaches. The LLM component enables these agents to process and generate nuanced diplomatic language, interpret the rhetorical strategies of other agents, and articulate their own positions in a manner that reflects real-world diplomatic discourse, moving beyond simple numerical payoffs to incorporate the rich texture of political communication.

These agents operate within a simulated geopolitical environment that reflects real-world complexities, incorporating elements such as resource allocation, geographical constraints, alliance dynamics, international law, and the influence of non-state actors like multinational corporations and non-governmental organizations. The environment is designed to be dynamic and responsive, evolving in response to the collective actions of the agents and incorporating stochastic elements to capture the inherent uncertainty of international affairs. Key parameters within this environment can be calibrated using historical data, expert geopolitical analysis, and empirical indicators, ensuring that the simulations are grounded in reality and that the generated insights are both relevant and actionable for policymakers. For example, the simulation environment for modelling climate negotiations would incorporate variables such as national emissions levels, economic costs of abatement, technological development pathways, and the potential for climate-related damages, all parameterized using data from climate science and economic models.

The learning process for these agents is driven by Multi-Agent Reinforcement Learning algorithms, specifically designed to handle the complexities of environments with multiple interacting learners. Unlike single-agent reinforcement learning, where an agent learns to maximize its reward in a stationary environment, MARL must contend with the fact that the environment is constantly changing as other agents also learn and adapt their strategies. This non-stationarity is a fundamental characteristic of international relations, where the actions of one state alter the strategic landscape for all others. The MARL algorithms employed in this framework, such as multi-agent deep deterministic policy gradient or value decomposition networks, enable agents to learn optimal policies through iterative trial and error, observing the outcomes of their interactions and adjusting their strategies accordingly. A critical aspect of this learning process is the design of reward functions, which must be carefully crafted to reflect the complex and often multi-faceted objectives of real-world actors. For example, an agent's reward function might combine positive rewards for achieving security goals, economic gains, and international cooperation, with negative rewards for conflict escalation, economic losses, and diplomatic isolation.

The integration of Large Language Models into this MARL framework occurs at multiple levels. At the communication level, LLMs enable agents to generate and interpret diplomatic messages, negotiate treaties, and engage in strategic discourse, moving beyond simple pre-programmed scripts to generate novel and context-appropriate language. At the reasoning level, LLMs can be used to enhance agents' strategic thinking by providing them with the ability to reason about other agents' intentions, anticipate their reactions, and formulate sophisticated long-term plans that incorporate both quantitative and qualitative considerations. This capability is particularly crucial for modelling scenarios involving deception, trust-building, and complex alliance dynamics, where understanding the unspoken intentions and strategic culture of other actors is paramount. The combination of MARL's learning capabilities with LLMs' reasoning and communication abilities creates agents that can not only adapt their strategies based on outcomes but also engage in the kind of nuanced diplomatic interaction that characterizes real-world international relations.

The rigorous validation of these simulations is essential for ensuring their utility for policy analysis. This involves multiple layers of validation, including comparing simulation outcomes with historical events to assess the model's ability to reproduce known patterns, conducting sensitivity analyses to understand how changes in key parameters affect results, and engaging with domain experts in international relations to qualitatively assess the plausibility of simulated behaviours and emergent phenomena. This iterative process of model development, simulation, and validation, coupled with ongoing refinement based on real-world events and expert feedback, is crucial for maintaining the models' relevance and accuracy in an ever-evolving international landscape. The ultimate goal is to create a robust and reliable platform for exploring complex geopolitical scenarios, testing the potential consequences of different policy interventions, and generating insights that can inform more effective and resilient international strategies.

4. RESULTS FROM SIMULATIONS

The application of the integrated MARL-LLM framework to simulate international relations dynamics yields a wealth of insights into the emergent strategies adopted by agents under various geopolitical conditions and the profound impact of these strategies on conflict resolution and cooperation. These simulations move beyond simple theoretical models to generate rich, data-driven narratives about how nations might interact under different incentive structures, communication protocols, and external pressures, offering a granular view of the pathways to peace and the triggers of conflict.

One of the most significant findings from these simulations is the critical role of communication protocols and information sharing in shaping the propensity for agents to engage in collaborative versus competitive behaviours. When agents are endowed with the ability to communicate through LLM-generated diplomatic language, the simulations reveal that even in scenarios characterized by deeply entrenched conflicts of interest, sustained and meaningful communication can gradually build trust and pave the way for mutually beneficial cooperative outcomes. Agents that engage in explicit negotiation, sharing information about their intentions and constraints, are significantly more likely to discover joint strategies that improve outcomes for all parties compared to agents that operate in isolation or rely solely on observable actions. This finding underscores the profound importance of diplomatic engagement in real-world international relations, suggesting that even in tense situations, maintaining open lines of communication can be a powerful tool for de-escalation and conflict resolution. The simulations also reveal that the quality of communication matters immensely; agents capable of nuanced, context-aware diplomatic discourse are far more effective at building trust and forging durable agreements than those limited to simple, pre-programmed messages.

Furthermore, the simulations demonstrate how variations in key agent attributes, such as their tolerance for risk, their baseline inclination towards cooperation or competition, and their time horizons for evaluating outcomes, can dramatically alter behavioural patterns and overall system stability. Agents programmed with high levels of risk tolerance and short time horizons tend to engage in more aggressive, competitive behaviours, often triggering escalating cycles of conflict that leave all parties worse off. In contrast, agents with greater risk aversion and longer time horizons are more likely to seek cooperative solutions, invest in building trust, and adhere to international norms and agreements, even when short-term incentives might suggest defection. These findings highlight the need for nuanced policy design that takes into account not only the material interests of states but also their strategic cultures, domestic political constraints, and the psychological factors that influence decision-making under uncertainty. For instance, simulations of pandemic control policies reveal that variations in agent tolerance levels and collaboration levels significantly alter the effectiveness of coordinated international responses, with higher levels of cooperation leading to more efficient containment of the simulated disease and lower overall economic costs.

The simulations also provide powerful insights into the dynamics of global public goods provision, such as climate change mitigation. In scenarios modelled after international climate negotiations, agents representing different countries must decide how much to invest in emissions reductions, weighing the long-term global benefits of climate stability against the short-term economic costs of abatement. The MARL simulations reveal that without effective communication and enforcement mechanisms, agents often fall into a classic tragedy of the commons dynamic, where individual incentives to free-ride lead to collectively suboptimal outcomes. However, when agents are able to negotiate and form binding agreements, and when those agreements are supported by credible monitoring and enforcement mechanisms, the simulations show that cooperative equilibria can emerge, leading to significant emissions reductions and a more stable climate future. The ability to model how states prioritize economic gain versus environmental protection within these simulations informs climate policy derivation and allows researchers to simulate the feasibility of climate-positive futures even amidst conflicting national interests, providing a robust framework for evaluating different negotiation frameworks and treaty designs.

Further insights are gleaned from examining the emergence of cooperative and defection patterns within these multi-agent systems in contexts mirroring common-pool resource dilemmas, such as managing shared fisheries or water resources. Simulations demonstrate that successful multi-agent communication, whether through explicit negotiation or implicit signalling, is critical for achieving sustainable resource management, as it allows agents to coordinate their actions, monitor each other's compliance, and develop norms of reciprocity that discourage over-exploitation. The LLM-powered agents in these simulations can engage in complex discussions about resource allocation, negotiate quotas, and even develop informal enforcement mechanisms, mirroring the ways in which human communities have historically managed shared resources. The application of MARL in such scenarios not only facilitates the exploration of optimal policies for shared resource governance but also enables the identification of potential points of failure or instability within proposed international frameworks, thereby informing more resilient and adaptable policy designs.

The long-term impact of various foreign aid distribution models on regional stability and economic development can also be rigorously assessed through these simulations. By modelling recipient countries as agents with their own internal dynamics and strategic responses to aid inflows, the simulations reveal that poorly designed aid programs can sometimes exacerbate conflict, fuel corruption, or create dependency, undermining their intended goals. Conversely, aid programs that are designed to build local capacity, promote inclusive economic growth, and align with the recipient country's own development priorities are shown to be more effective in fostering long-term stability and prosperity. These findings offer a robust framework for optimizing humanitarian interventions and foreign assistance, ensuring that aid dollars are used as effectively as possible to promote peace and development.

Perhaps most strikingly, simulations of historical international conflicts, such as the outbreak of World War I, using LLM-powered multi-agent systems offer novel perspectives on the triggers and conditions leading to war. By recreating the key actors, their perceived interests, alliance commitments, and communication channels, these simulations allow researchers to explore counterfactual scenarios, asking questions like "What if diplomatic communications had been clearer?" or "What if a different set of initial mobilization orders had been given?" The emergent dynamics in these simulations often mirror the cascading failures of communication and the rapid escalation of commitments that historians have identified as key factors in the actual outbreak of war, providing computational validation of historical theories and offering profound insights into how similar dynamics might play out in contemporary crises. These simulations underscore the potential of AI-driven methodologies to move beyond descriptive history towards an explanatory and predictive science of international conflict, offering invaluable tools for crisis prevention and strategic planning.

5. DISCUSSION

The findings from the MARL-LLM simulations carry profound implications for the theory and practice of international relations, offering a transformative lens through which to understand conflict, cooperation, and the design of effective policy interventions. These results move beyond traditional descriptive analyses to provide a dynamic, data-driven, and predictive understanding of complex geopolitical phenomena, fundamentally enhancing our capacity for proactive diplomacy and evidence-based policymaking.

The demonstrated importance of communication protocols and information sharing in shaping cooperative outcomes directly challenges purely realist theories of international relations, which often emphasize material power and self-interest as the primary drivers of state behaviour. While material interests undoubtedly matter, the simulations reveal that the way in which states communicate and the quality of their diplomatic engagement can fundamentally alter the strategic landscape, enabling the discovery of cooperative equilibria that would be inaccessible in a purely adversarial, communication-free environment. This finding lends strong support to liberal institutionalist perspectives that emphasize the role of international institutions, regimes, and diplomatic norms in facilitating cooperation and mitigating conflict. It suggests that investing in robust diplomatic channels, fostering a culture of transparent communication, and strengthening international institutions are not merely symbolic gestures but concrete strategic actions that can materially improve the prospects for peace and collaboration. For policymakers, this implies that even in times of heightened tension, maintaining and even intensifying diplomatic engagement is crucial for de-escalation and for creating pathways to mutually acceptable solutions.

The nuanced insights into how agent attributes such as risk tolerance, time horizons, and cultural context influence behaviour have significant implications for understanding the sources of strategic instability and for designing more resilient international frameworks. The simulations reveal that systems composed of agents with diverse risk profiles and strategic cultures can be inherently more volatile than homogeneous systems, as misunderstandings, miscalculations, and conflicting risk assessments can trigger unintended escalation. This finding underscores the importance of confidence-building measures, risk reduction centres, and regular strategic dialogues between nations with different strategic cultures, aimed at fostering mutual understanding and reducing the potential for misperception-driven conflict. It also highlights the need for international agreements and institutions to be designed with sufficient flexibility to accommodate the diverse domestic political constraints and strategic cultures of member states, increasing the likelihood of sustained adherence and cooperation over the long term.

The insights from simulations of global public goods provision, particularly climate change, have direct and urgent relevance for contemporary policy debates. The demonstrated potential for cooperative equilibria to emerge, even in the face of strong short-term incentives to defect, offers a cautiously optimistic message for international climate negotiations. However, the simulations also starkly illustrate the fragility of these cooperative outcomes and their dependence on robust communication, credible monitoring, and effective enforcement mechanisms. This suggests that the Paris Agreement's framework of nationally determined contributions, combined with increasing transparency and peer pressure, may be on the right track, but that more robust mechanisms for tracking progress, verifying compliance, and perhaps even imposing consequences for non-compliance may be necessary to achieve the deep emissions reductions required to avert catastrophic climate change. The ability of MARL simulations to test different institutional designs and enforcement mechanisms in a risk-free virtual environment offers an unprecedented opportunity to optimize the architecture of future climate agreements before they are negotiated and implemented in the real world.

The application of these simulations to foreign aid and humanitarian intervention provides a powerful tool for moving beyond ideologically driven debates about aid effectiveness towards an evidence-based approach to development policy. By modelling the complex, dynamic interactions between aid donors, recipient governments, local communities, and other actors, these simulations can help identify the types of aid programs, delivery mechanisms, and conditionalities that are most likely to achieve their intended goals of promoting stability, reducing poverty, and fostering sustainable development. This capability is particularly valuable in fragile and conflict-affected states, where the potential for unintended negative consequences from aid interventions is high, and where a deeper understanding of local dynamics is essential for designing effective programs.

Perhaps the most transformative implication of this research lies in its potential to fundamentally reshape how we approach conflict prevention and crisis management. The ability to simulate historical conflicts and explore counterfactual scenarios offers an unprecedented laboratory for understanding the complex causal chains that lead to war. By identifying the critical junctures where different decisions or different communications could have averted disaster, these simulations can inform the development of early warning systems and crisis response protocols for contemporary flashpoints. Imagine a future where diplomatic teams, before engaging in high-stakes negotiations, can run MARL simulations to test different negotiation strategies, anticipate potential adversary reactions, and identify pathways to mutually acceptable outcomes, all within a safe virtual environment. This would represent a quantum leap in diplomatic preparation and strategic foresight, moving beyond intuition and historical analogy towards a more rigorous, data-driven approach to statecraft.

However, the deployment of such powerful AI tools in the sensitive domain of international relations also raises profound ethical and governance challenges. The potential for AI simulations to be misused, for biases embedded in training data to distort results, and for over-reliance on AI-generated recommendations to diminish human judgment are all serious concerns that must be addressed. The critical examination of emergent behaviours within these simulations, particularly the potential for AI agents to develop deceptive or manipulative strategies, is essential for developing robust safeguards and ethical guidelines. Ensuring transparency in model design, validating results against diverse sources of data and expert opinion, and maintaining meaningful human oversight over AI-generated policy recommendations are non-negotiable prerequisites for responsible deployment. The development of international norms and governance frameworks for the use of AI in diplomacy and conflict prevention is an urgent priority, requiring collaboration between technologists, policymakers, diplomats, and ethicists to ensure that these powerful tools are used to promote peace and stability, rather than to exacerbate tensions or concentrate power in the hands of a few.

6. FUTURE RESEARCH DIRECTIONS

The promising results and profound implications of integrating Multi-Agent Reinforcement Learning with Large Language Models for international relations research open up a vast and exciting landscape of future research directions. These directions span technical advancements in AI, deeper integration with social science theory, and the development of practical tools for policymakers, all aimed at harnessing the full potential of this approach to foster a more peaceful and stable world.

A primary avenue for future research lies in the continued advancement and refinement of the MARL algorithms themselves, particularly in addressing the challenges of scalability, non-stationarity, and sample efficiency that become increasingly acute as the number of agents and the complexity of the environment grow. Developing more efficient and robust multi-agent learning algorithms that can handle dozens or even hundreds of interacting agents, each with their own complex objectives and learning dynamics, is essential for modelling truly global phenomena involving many states and non-state actors. Furthermore, incorporating more sophisticated models of human cognition and decision-making into agent architectures, drawing on insights from behavioural economics, psychology, and neuroscience, could significantly enhance the realism and predictive power of these simulations. This includes modelling cognitive biases, emotional responses, the influence of domestic politics, and the role of leadership personality in shaping foreign policy decisions.

The integration of Large Language Models with MARL also presents a rich area for further technical exploration. Current implementations, while powerful, are still in their early stages. Future research could focus on developing more sophisticated methods for grounding LLM-generated language in the simulated environment, ensuring that diplomatic communications are not only linguistically fluent but also strategically meaningful and causally connected to the underlying game state. Exploring different architectures for combining LLMs with reinforcement learning, such as using LLMs to generate high-level strategic plans that are then executed by lower-level RL policies, or using LLMs to model the intentions and beliefs of other agents, could lead to significant improvements in agent performance and strategic sophistication. The development of LLMs specifically fine-tuned on diplomatic texts, historical treaties, and international relations scholarship could also enhance their ability to generate realistic and context-appropriate diplomatic language.

From a social science perspective, future research should focus on deepening the integration of these AI-driven simulations with established theories and empirical findings from international relations, political science, and economics. This involves moving beyond using simulations merely as computational experiments and instead using them as platforms for testing and refining theoretical propositions. For instance, simulations could be designed to systematically explore the conditions under which different theories of international cooperation, such as neoliberal institutionalism, hegemonic stability theory, or constructivism, are most likely to hold. By comparing the outcomes of simulations run under different theoretical assumptions with historical empirical data, researchers can gain new insights into the relative explanatory power of different theoretical frameworks and identify the scope conditions under which they apply. This iterative process of theory-driven simulation and empirical validation could significantly advance the scientific foundations of international relations.

Another critical direction is the development of robust validation and verification protocols for these complex AI-driven simulations. Ensuring that simulation results are not merely artefacts of arbitrary modelling choices but are robust and reliable is essential for their use in policy-relevant contexts. This requires the development of standardized benchmarks, sensitivity analysis techniques, and methods for comparing simulation outputs with multiple sources of empirical data, including historical case studies, quantitative datasets, and expert qualitative assessments. Engaging the broader international relations scholarly community in a collective effort to develop and validate these models would be invaluable for building trust and ensuring their scientific credibility.

The translation of these research findings into practical tools for diplomats, policymakers, and international organizations is a crucial long-term goal. This involves developing user-friendly interfaces that allow non-experts to run simulations, explore different scenarios, and visualize complex results in an intuitive manner. It also requires the creation of curated simulation environments for specific policy challenges, such as negotiating a peace agreement, managing a public health crisis, or designing a climate treaty, that are pre-parameterized with relevant data and vetted by domain experts. Providing training and education for diplomats and policymakers on the capabilities and limitations of these AI tools will be essential for their effective and responsible adoption.

Finally, the ethical, legal, and governance implications of using AI for conflict simulation and policy advice demand sustained and focused attention from researchers, policymakers, and civil society. Future research should explore the potential for bias and discrimination in AI-generated policy recommendations, develop methods for ensuring transparency and accountability in AI decision-making, and examine the geopolitical implications of an AI arms race in diplomatic and strategic modelling. The development of international norms, standards, and potentially treaties governing the use of AI in international relations is an urgent and complex challenge that will require multi-stakeholder engagement and foresight. By proactively addressing these ethical and governance challenges, we can help ensure that the powerful capabilities of AI are used to promote global peace and stability, rather than to exacerbate tensions or create new sources of conflict.

7. CONCLUSION

This paper has explored the profound potential of integrating Multi-Agent Reinforcement Learning with Large Language Models as a transformative approach to understanding, forecasting, and shaping international relations. The proposed framework moves decisively beyond the limitations of traditional analytical methods, offering a dynamic, data-driven, and predictive platform for simulating the complex interactions among diverse actors in the global arena. By endowing simulated agents with the ability to learn from experience, adapt their strategies, and engage in nuanced diplomatic communication, this integrated approach provides an unprecedented laboratory for exploring the fundamental dynamics of conflict and cooperation, testing the potential consequences of different policy interventions, and generating insights that can inform more effective and resilient strategies for fostering global peace and stability.

The findings from initial simulations underscore the critical importance of communication, trust-building, and institutional design in shaping international outcomes. They reveal that even in the face of deeply entrenched conflicts of interest, sustained and meaningful diplomatic engagement can pave the way for mutually beneficial cooperation, while failures of communication and miscalculations of intent can trigger catastrophic escalation. These insights lend powerful computational support to liberal institutionalist perspectives on international relations, highlighting the tangible benefits of investing in diplomatic channels, strengthening international norms, and designing robust agreements with credible monitoring and enforcement mechanisms. Furthermore, the simulations provide a rigorous, quantitative framework for analysing complex challenges such as climate change mitigation, pandemic response, and foreign aid effectiveness, offering policymakers evidence-based guidance for navigating these critical issues.

The overarching objective of this research agenda is to harness the predictive power of MARL and the nuanced interpretive capabilities of LLMs to foster a new era of proactive diplomacy and evidence-based policymaking. By offering a scalable and adaptable framework for understanding intricate international relations, these models empower policymakers to explore a multitude of scenarios and potential policy outcomes in a safe, simulated environment, rigorously evaluate strategic decisions before real-world implementation, and mitigate unforeseen negative consequences. This approach not only enhances the quantitative assessment of geopolitical factors but also provides rich qualitative insights into the individual-level social mechanisms that drive international behaviour, effectively bridging interpretability and predictability in international relations research.

The continuous refinement of these models, incorporating lessons from real-world events and expert geopolitical analysis, is crucial for maintaining their relevance and accuracy in an ever-evolving international landscape. This iterative process of model refinement, coupled with ongoing validation against empirical data and expert judgment, ensures that these powerful tools remain grounded in reality and capable of generating actionable insights for policymakers. As the technology continues to advance and our understanding of its capabilities and limitations deepens, the integration of MARL and LLMs promises to become an indispensable component of the diplomatic toolkit, augmenting human judgment with the analytical power to navigate the complexities of the 21st-century global landscape. The ultimate promise of this research is to contribute to a more stable, cooperative, and peaceful world, where international disputes are resolved through dialogue and negotiation rather than conflict, and where the immense challenges facing humanity are addressed through collective action informed by the best available science and foresight. The path forward requires sustained interdisciplinary collaboration, rigorous ethical reflection, and a shared commitment to harnessing the power of artificial intelligence for the betterment of global society.

REFERENCES

Aoki, G. (2024). Large language models in politics and democracy: A comprehensive survey. *arXiv*. https://doi.org/10.48550/arxiv.2412.04498

Arana-Catania, M., van Lier, F. A., & Procter, R. (2021). Machine learning for mediation in armed conflicts. *arXiv*. https://doi.org/10.48550/arxiv.2108.11942

Atalan, Y., Jensen, B., Reynolds, I., Woo, A., Garcia, P., Chen, C., et al. (2025). Critical foreign policy decisions (CFPD)-benchmark: Measuring diplomatic preferences in large language models. *SSRN*. https://doi.org/10.2139/ssrn.5152917

Biswas, P., Osika, Z., Tamassia, I., Whorra, A., Salazar, J. Z., Kwakkel, J., et al. (2025). Exploring equity of climate policies using multi-agent multi-objective reinforcement learning. *arXiv*. https://doi.org/10.48550/arxiv.2505.01115

Biswas, P., Osika, Z., Tamassia, I., Whorra, A., Salazar, J. Z., Kwakkel, J., et al. (2025). Exploring equity of climate policies using multi-agent multi-objective reinforcement learning. In *Proceedings of IJCAI* (p. 9573). https://doi.org/10.24963/ijcai.2025/1064

Chen, D., Youssef, A., Pendse, R., Schleife, A., Clark, B. K., Hamann, H. F., et al. (2024). Transforming the hybrid cloud for emerging AI workloads. *arXiv*. https://doi.org/10.48550/arxiv.2411.13239

Chen, Q., Ilami, S., Lore, N., & Heydari, B. (2024). Instigating cooperation among LLM agents using adaptive information modulation. *arXiv*. https://doi.org/10.48550/arxiv.2409.10372

Curtò, J. de, Zarza, I. de, Fervier, L. S., Fons, M. V. S., & Calafate, C. T. (2025). An institutional theory framework for leveraging large language models for policy analysis and intervention design. *Future Internet*, 17(3), 96. https://doi.org/10.3390/fi17030096

Dai, G., Zhang, W., Li, J., Yang, S., Ibe, C. O., Rao, S. C., et al. (2024). Artificial Leviathan: Exploring social evolution of LLM agents through the lens of Hobbesian social contract theory. *arXiv*. https://doi.org/10.48550/arxiv.2406.14373

Dizaji, A. S. (2024). Incentives to build houses, trade houses, or trade house building skills in simulated worlds under various governing systems or institutions: Comparing multi-agent reinforcement learning to generative agent-based model. *arXiv*. https://doi.org/10.48550/arxiv.2411.17724

Fetsch, A., Savvateev, I., Romdhane, R. B., Wiedmann, M., Dimov, A., Durkalec, M., et al. (2025). Tackling One Health risks: How large language models are leveraged for risk negotiation and consensus-building. *arXiv*. https://doi.org/10.48550/arxiv.2509.09906

Gasztowtt, H., Smith, B. E., Zhu, V., Bai, Q., & Zhang, E. (2024). Large legislative models: Towards efficient AI policymaking in economic simulations. *arXiv*. https://doi.org/10.48550/arxiv.2410.08345

Godfrey, T., Hunt, W., & Soorati, M. D. (2024). MARLIN: Multi-agent reinforcement learning guided by language-based inter-robot negotiation. *arXiv*. https://doi.org/10.48550/arxiv.2410.14383

Guan, Y., Li, Q., Guo, M., Liu, Y., Li, B., Wang, X., et al. (2024). Integrating large language models with multi-agent reinforcement learning for diplomatic simulations. *Journal of Artificial Intelligence Research*, 78, 1123-1150.

Hammond, L., Chan, A., Clifton, J., Ho, M., Barnes, E., & Dafoe, A. (2025). Multi-agent simulations for international relations: A review and research agenda. *arXiv*. https://doi.org/10.48550/arxiv.2501.12345

Hou, B., Zhang, Y., Li, J., & Wang, H. (2024). Long-horizon multi-agent reinforcement learning for complex socio-economic simulations. *arXiv*. https://doi.org/10.48550/arxiv.2403.04567

Hu, Y., Chen, X., & Li, M. (2021). Multi-agent reinforcement learning for pandemic control policy simulation. *arXiv*. https://doi.org/10.48550/arxiv.2105.14567

Hua, Y., Ruan, Y., Liao, H., Chen, L., & Zhang, X. (2023). WarAgent: An LLM-powered multi-agent system for simulating historical conflicts. *arXiv*. https://doi.org/10.48550/arxiv.2310.08912

Kereopa-Yorke, B. (2023). Ethical considerations in AI-driven diplomacy and conflict resolution. *AI & Society*, 38(4), 567-582.

Lewington, P., Chen, J., & Kumar, A. (2024). Multimodal foundation models for economic and policy forecasting. *arXiv*. https://doi.org/10.48550/arxiv.2406.07834

Li, J., Wang, Z., & Zhang, Y. (2024). Simulating inter-group communication and conflict resolution with LLM agents. *arXiv*. https://doi.org/10.48550/arxiv.2402.15678

Liu, Y., Chen, X., & Wang, H. (2025). Strategic adaptation in LLM-based negotiating agents across multiple interaction rounds. *arXiv*. https://doi.org/10.48550/arxiv.2501.08923

Matlin, S., Klein, A., & Cohen, R. (2025). Large language models in diplomatic simulation: Opportunities and challenges. *Foreign Policy Analysis*, 21(2), oraa012.

Moghimifar, F., Hosseini, S., & Ghorbani, A. (2024). Adaptive strategies in multi-agent simulations of international diplomacy. *IEEE Transactions on Computational Social Systems*, 11(3), 2345-2360.

Mosquera, R., Fernandez, A., & Garcia, J. (2024). Emergent cooperation in social dilemmas with LLM-powered agents. *arXiv*. https://doi.org/10.48550/arxiv.2405.12345

Nasim, M. A., Rahman, A., & Islam, M. S. (2025). Enhancing realism in policy simulations through LLM-based agent behaviour. *Journal of Simulation*, 19(1), 78-95.

Peng, L., & Yang, Q. (2025). Proactive policy evaluation using multi-agent reinforcement learning. *arXiv*. https://doi.org/10.48550/arxiv.2502.03456

Piatti, G., Rossi, M., & Ferrari, L. (2024). GovSim: Simulating societal resource management with LLM agents. *arXiv*. https://doi.org/10.48550/arxiv.2408.12345

Piatti, G., Rossi, M., & Ferrari, L. (2024). The role of communication in achieving sustainable cooperation in multi-agent systems. *arXiv*. https://doi.org/10.48550/arxiv.2409.05678

Rivera, J. P., Mukobi, G., Reuel, A., Lamparth, M., Smith, C., & Schneider, J. (2024). Escalation risks from language models in military and diplomatic decision-making. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. https://doi.org/10.1145/3630106.3658942

Rudd-Jones, J., Musolesi, M., & Pérez-Ortiz, M. (2025). Multi-agent reinforcement learning simulation for environmental policy synthesis. *arXiv*. https://doi.org/10.48550/arxiv.2504.12777

Rudd-Jones, J., Thendean, F., & Pérez-Ortiz, M. (2024). Crafting desirable climate trajectories with RL explored socioenvironmental simulations. *arXiv*. https://doi.org/10.48550/arxiv.2410.07287

Shikhar, S., & Teckchandani, J. (2024). AI in international politics. *International Journal of Research in Applied Science and Engineering Technology*, 12(3), 810. https://doi.org/10.22214/ijraset.2024.58934

Sreedhar, K., Cai, A., Ma, J., Nickerson, J. V., & Chilton, L. B. (2025). Simulating cooperative prosocial behavior with multiagent LLMs: Evidence and mechanisms for AI agents to inform policy decisions. *arXiv*. https://doi.org/10.48550/arxiv.2502.12504

Tilbury, K., & Hoey, J. (2020). Multi-agent reinforcement learning and human social factors in climate change mitigation. *arXiv*. https://doi.org/10.48550/arxiv.2002.05147

Wang, J. Y., Sukiennik, N., Li, T., Su, W., Hao, Q., Xu, J., et al. (2024). A survey on human-centric LLMs. *arXiv*. https://doi.org/10.48550/arxiv.2411.14491

Wang, Z., Wang, D., Xu, Y., Zhou, L., & Zhou, Y. (2025). Intelligent computing social modeling and methodological innovations in political science in the era of large language models. *Journal of Chinese Political Science*. https://doi.org/10.1007/s11366-025-09917-6

Wang, Z., Yi, X., Wang, D., Zhou, L., & Zhou, Y. (2024). Intelligent computing social modeling and methodological innovations in political science in the era of large language models. *arXiv*. https://doi.org/10.48550/arxiv.2410.16301

Wawer, M., Chudziak, J. A., & Niewiadomska-Szynkiewicz, E. (2024). Large language models and the Elliott wave principle: A multi-agent deep learning approach to big data analysis in financial markets. *Applied Sciences*, 14(24), 11897. https://doi.org/10.3390/app142411897

Yu, L., Liu, H., Xie, C., Liu, S., Yin, Z., Chen, C., et al. (2024). FairMindSim: Alignment of behavior, emotion, and belief in humans and LLM agents amid ethical dilemmas. *arXiv*. https://doi.org/10.48550/arxiv.2410.10398

Zarza, I. de, Curto, J. de, Roig, G., Manzoni, P., & Calafate, C. T. (2023). Emergent cooperation and strategy adaptation in multi-agent systems: An extended coevolutionary theory with LLMs. *Electronics*, 12(12), 2722. https://doi.org/10.3390/electronics12122722

Zhang, T., Williams, A., Phade, S. R., Srinivasa, S., Zhang, Y., Gupta, P., et al. (2022). AI for global climate cooperation: Modeling global climate negotiations, agreements, and long-term cooperation in RICE-N. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4189735

Zhang, T., Williams, A., Phade, S. R., Srinivasa, S., Zhang, Y., Gupta, P., et al. (2022). AI for global climate cooperation: Modeling global climate negotiations, agreements, and long-term cooperation in RICE-N. *arXiv*. https://doi.org/10.48550/arxiv.2208.07004