

Simulating the Next Donbas Crisis: Integrating Machine Learning into Wargames for NATO-Russia Escalation Management

Innocent Jooji, PhD

Department of Peace and Conflict Studies,
Veritas University,
Abuja.
Nigeria.

Abstract

The ongoing volatility in the Donbas region underscores the precarious nature of NATO-Russia relations and the ever-present risk of strategic miscalculation. Traditional wargaming, while invaluable, often struggles to incorporate the sheer volume of variables and the rapid tempo inherent in modern hybrid warfare. This paper proposes a novel framework integrating Machine Learning (ML) algorithms into conventional wargaming methodologies to enhance the simulation of a potential future Donbas crisis. By leveraging predictive analytics and reinforcement learning, the proposed model moves beyond static scenarios to dynamically model adversarial decision-making and complex escalation ladders. This study demonstrates how ML can process vast datasets including historical troop movements, geopolitical signaling, and economic sanctions to identify non-obvious patterns and potential trigger points for escalation. The integration of these technologies allows for real-time adaptation within the simulation, offering policymakers a robust tool for testing deterrence strategies and de-escalation protocols under conditions of extreme uncertainty. Furthermore, the paper addresses the ethical and operational challenges of relying on algorithmic input for high-stakes defense planning. In conclusion, this research argues that AI-augmented wargaming provides NATO with a critical strategic edge, enabling a more nuanced understanding of Russian red lines and significantly improving the Alliance's ability to manage and contain escalation in future conflicts.

Key Words: Machine Learning, Wargaming, Escalation Management, NATO-Russia Relations, Strategic Forecasting

Introduction

The geopolitical landscape of Eastern Europe has undergone a radical transformation since the annexation of Crimea in 2014, with the Donbas region serving as the epicenter of a protracted and volatile conflict between Russia and Ukrainian forces. What began as a hybrid warfare operation has evolved into a strategic stalemate with profound implications for Euro-Atlantic security. For the North Atlantic Treaty Organization (NATO), the conflict in Eastern Ukraine represents not merely a regional dispute but a litmus test for the Alliance's ability to deter aggression and manage escalation near its borders. As the security environment deteriorates, the prospect of a "next Donbas crisis", a distinct escalation involving conventional forces or significant hybrid operations looms larger.

Consequently, the imperative for NATO to understand potential escalation dynamics has never been greater. This paper argues that traditional analytical methods are no longer sufficient to grasp

the complexity of modern Russian deterrence theory, and proposes the integration of Machine Learning (ML) into wargaming methodologies as a necessary evolution for strategic forecasting. Historically, wargaming has served as the cornerstone of military planning and professional military education. From the Kriegsspiel of the 19th century to the massive computer-assisted exercises of the Cold War, these simulations have provided decision-makers with a sandbox to test strategies and explore adversary behaviors (Perla, 1990). In the context of NATO-Russia relations, wargaming traditionally focuses on managing the "escalation ladder" that is gradual or rapid intensification of conflict from political coercion to tactical nuclear use.

However, traditional wargames are often hampered by human cognitive biases and rigid scenario structures. As Gredler (2004) notes, human-driven simulations frequently suffer from "mirror-imaging," where Blue Force (NATO) planners assume the Red Force (Russia) will react rationally according to Western, ignoring the distinct cultural, historical, and political drivers that inform Russian decision-making. In a crisis as emotionally and historically charged as the Donbas, these biases can lead to strategic blind spots, potentially misidentifying red lines or misinterpreting adversary signaling.

Furthermore, the nature of warfare itself has shifted in a way that overwhelms manual simulation capabilities. The conflict in Ukraine has been characterized by the convergence of conventional force, cyber operations, disinformation campaigns, and economic warfare which is a concept often termed "multi-domain operations" or "new generation warfare" (Giles, 2016). The sheer volume of variables and the speed at which information travels in the modern battlespace create a level of complexity that defies traditional heuristic analysis. Human adjudicators in a wargame struggle to process the nonlinear interactions between, for example, a cyber-attack on a power grid and the subsequent shift in public opinion or troop morale. This is where the intersection of computer science and strategic studies offers a promising path forward.

The integration of Machine Learning into wargaming offers a solution to the limitations of human-centric simulations. Unlike static models, ML algorithms can ingest vast datasets including historical troop movements, logistics data, social media sentiment, and economic indicators to identify patterns and correlations that remain invisible to human analysts (Horowitz, 2019). Specifically, reinforcement learning, a subset of ML, allows artificial agents to interact with a simulated environment, learn from the consequences of their actions, and develop novel strategies over time. In the context of a Donbas crisis simulation, an ML-driven Red Force agent could

explore thousands of potential escalation paths in a fraction of the time required by a human wargame, revealing unexpected or "counter-intuitive" strategies that Russian forces might employ (Johnson et al., 2019). This shifts the utility of wargaming from merely validating existing plans to genuinely discovering new vulnerabilities.

However, the application of AI to nuclear escalation and high-intensity conflict is not without controversy. Scholars have raised valid concerns regarding the "black box" nature of deep learning algorithms, where the rationale behind a specific strategic decision is opaque to human operators (Payne, 2021). In a scenario involving the potential use of nuclear weapons, trusting an opaque algorithm to manage the ladder of escalation carries significant ethical and operational risks. Furthermore, there is the danger of data bias; if the ML models are trained solely on Western open-source intelligence or historical Western exercises, the "mirror-imaging" problem may simply be automated rather than eliminated.

On the whole, this study aims to bridge the gap between advanced computational techniques and classical strategic theory. By simulating a future crisis in the Donbas region, this study demonstrates how ML-augmented wargames can provide NATO planners with a more dynamic, probabilistic understanding of Russian escalation thresholds. The following sections will outline the theoretical framework for AI-driven wargaming, detail the architecture of the proposed simulation model, and analyze the results of several distinct crisis scenarios.

Literature Review

Scholars such as Lawrence Freedman (2023) and Keir Giles (2022) have analyzed the "escalation ladder" in the context of the Ukrainian theater. They argue that the Russian doctrine of "escalate to de-escalate" which is the use of tactical nuclear weapons to force a cessation of hostilities creates a nonlinear risk profile that traditional linear models fail to capture. The Donbas region serves as a potential flashpoint where tactical encirclement could trigger strategic response. Acton (2022) highlights the difficulty of deterrence in this environment, noting that the integration of conventional and nuclear capabilities in Russian military doctrine complicates NATO's red lines. Recent analyses by the Rand Corporation and the Institute for the Study of War emphasize that the "next" crisis in Donbas will not be purely kinetic. It will involve cyber-attacks on European infrastructure, information warfare, and economic coercion. Kofman and Lee (2022) detail how the initial phases of the 2022 invasion demonstrated the friction between theoretical Russian combined arms operations and the reality of Ukrainian defense. For a simulation to be accurate, it

must account for these friction points, which are often influenced by logistical constraints and morale variables difficult to model statically.

Perla (1990) and Gantt (2020) argue that the primary value of wargaming lies in educating the human decision-maker. These games exploit the "OODA Loop" (Observe, Orient, Decide, Act) to train officers in rapid decision-making. However, the literature identifies significant limitations. Human players are subject to cognitive biases, availability heuristics, and emotional fatigue, which can skew results. Furthermore, manual adjudication where a referee determines the outcome of a move is often criticized for being subjective and slow (Rubel, 2021).

A major limitation identified in defense literature is the "branching factor." In a complex scenario like a Donbas escalation, the number of possible moves and counter-moves exceeds human cognitive capacity. Traditional games can only explore a fraction of the decision tree, potentially missing "black swan" events or unintended escalation pathways (Schoemaker, 2020). This creates a demand for automated agents that can run simulations at speeds and volumes unattainable by human teams.

Agent-Based Modeling allows researchers to simulate the interactions of autonomous agents to observe emergent behaviors. When combined with Reinforcement Learning (RL), these agents can learn optimal strategies through trial and error. A seminal study by the OpenAI "Five" team and subsequent DARPA initiatives (e.g., the ACE program) demonstrates that RL agents can develop complex, non-intuitive strategies in tactical combat scenarios. In the context of escalation management, RL agents are particularly useful for modeling adversary behavior. By training an AI on historical Russian military doctrine and observed behavior in Ukraine, analysts can simulate a "Red Team" that acts unpredictably yet rationally within its own framework (Johnson et al., 2021).

The most recent evolution in this field involves the use of Large Language Models (like GPT-4) to simulate strategic decision-making. Research by Beard et al. (2023) suggests that LLMs can role-play national leaders with surprising fidelity, taking into account ideological constraints and historical context. Unlike game-theoretic models that assume perfect rationality, LLM-based agents can simulate pride, fear, and misperception which are critical variables in a nuclear standoff. This allows for the simulation of the "Next Donbas Crisis" to include diplomatic negotiations and information warfare, not just tactical movement.

The literature emphasizes the challenge of "data hunger" in military applications. However, the open-source intelligence (OSINT) revolution has provided vast datasets regarding the disposition of forces in Eastern Europe. Researchers are now using computer vision to analyze satellite imagery of the Donbas front to train terrain-analysis algorithms (Williams & Blum, 2022). This integration allows simulations to move beyond abstract maps to geospecific models of the Donbas terrain, including logistics bottlenecks and urban choke points.

A critical barrier to adoption is the opacity of ML decision-making. In a nuclear escalation scenario, military leaders require justification for a recommended course of action. If an ML agent recommends a strike on a Russian supply depot that triggers a nuclear response, and the rationale is buried in a neural network's weights, the tool is useless for practical decision-making (Horowitz, 2019). Current research focuses on Explainable AI (XAI) to bridge this gap. There is a concern that ML models trained on Western data may project Western rationality onto Russian decision-making, leading to mirror-imaging errors. Furthermore, automated wargames suffer from a lack of "skin in the game"; an AI agent does not fear destruction, potentially leading to riskier escalation in simulation than a human commander would tolerate (Payne, 2021).

The literature suggests that the simulation of a future crisis in the Donbas requires a paradigm shift. Traditional wargaming offers indispensable insight into human psychology and coalition management but lacks the speed and scalability to model the full complexity of modern multi-domain war. Conversely, Machine Learning offers the ability to process vast datasets and explore thousands of escalation pathways, but it struggles with the nuances of human irrationality and explainability. The emerging consensus points toward a hybrid approach: "Human-in-the-loop" AI wargaming. In this model, ML handles the tactical adjudication and logistics simulations, as well as generating unpredictable Red Team strategies based on doctrinal data, while human decision-makers focus on the strategic and political implications. This synthesis offers the most robust path for NATO to manage escalation in a volatile theater, ensuring that the alliance is prepared for both the intended and unintended consequences of a renewed conflict in Eastern Europe.

Methodology

To investigate the complex dynamics of a potential future crisis in the Donbas region, this study employs a hybrid research design that synthesizes traditional force-on-force wargaming with Multi-Agent Reinforcement Learning (MARL). The foundational simulation environment is constructed using high-resolution terrain data of the Donbas region and rigorous Order of Battle

(OB) tables for both Russian and NATO forces. To ensure historical fidelity, the model ingests a comprehensive dataset derived from the 2014–2024 conflicts, including tactical engagement logs, supply chain timelines, and geopolitical decision points. Furthermore, qualitative factors such as Russia's Gerasimov Doctrine and NATO's Comprehensive Defence strategies are quantified into algorithmic constraints and rules of engagement, ensuring that agent behavior remains grounded in real-world military doctrine rather than arbitrary game logic.

The wargame operates on a hexagonal grid system utilizing discrete time steps to manage the flow of combat and logistics. Within this architecture, the state of the conflict is represented by a dynamic matrix tracking unit positioning, artillery suppression levels, and air superiority. A critical component of the system is the "Escalation Metric," a variable that monitors conflict intensity across a seven-step ladder ranging from skirmishes to tactical nuclear deployment. This metric serves as the primary dependent variable, allowing the researchers to track how tactical decisions on the ground propagate into strategic crises. The environment is designed to react to these inputs, simulating the friction of war and the logistical constraints inherent to the region, thereby providing a realistic testing ground for algorithmic decision-making.

The core of the simulation involves the deployment of autonomous agents powered by Deep Reinforcement Learning, specifically utilizing the Proximal Policy Optimization (PPO) algorithm. Two distinct adversarial teams are established: a Red Agent (Russia), optimized for territorial expansion and NATO alliance disruption, and a Blue Agent (NATO), tasked with deterrence and the restoration of status quo borders. To address the specific research question of escalation management, the methodology employs sophisticated "reward shaping." Rather than rewarding victory alone, the algorithms are heavily penalized for crossing nuclear thresholds or incurring high civilian casualties. This compels the AI to explore complex de-escalation strategies and diplomatic "off-ramps" that might be overlooked in standard zero-sum models. Additionally, a Large Language Model (LLM) layer is integrated to generate and interpret diplomatic signals, introducing a negotiation dimension to the otherwise kinetic simulation.

Finally, the study relies on high-volume iteration and expert validation to ensure analytical rigor. The system executes 10,000 Monte Carlo simulation runs, providing a robust probability distribution of potential conflict outcomes. This volume allows the model to identify emergent behaviors and non-linear cause-and-effect relationships that human planners might miss. The resulting data is validated against parallel wargames conducted by Subject Matter Experts (SMEs),

comparing AI decision-making patterns against human tactical logic. By analyzing these iterations, the study isolates specific "tipping points" a conditional states where minor tactical shifts trigger uncontrolled strategic escalation thereby offering data-driven insights for NATO deterrence planning.

Discussion of Findings

The execution of 10,000 Monte Carlo simulation runs within the hybrid wargame environment has yielded critical insights into the volatility of a future Donbas crisis, fundamentally challenging established theories of conflict resolution. By integrating Multi-Agent Reinforcement Learning (MARL) with Large Language Models (LLMs), this study moved beyond static, human-centric wargaming to explore the "boundary conditions" of escalation when non-human actors influence the decision loop. The findings reveal a landscape where the friction of war is not solely physical but temporal and cognitive. The integration of machine learning into wargaming highlights the destabilizing effects of automated decision-making and the precariousness of de-escalation pathways, suggesting that the geopolitical stability of the European theater hinges not just on missile counts, but on the speed and logic of the algorithms controlling them.

A primary and alarming finding of the study is the phenomenon of "hyper-escalation." In traditional military theory, John Boyd's OODA loop (Observe, Orient, Decide, Act) describes the cognitive cycle of conflict. Historically, the advantage went to the side that could cycle through this loop faster than the adversary (Boyd, 1995). However, the MARL models in this simulation compressed these decision-making cycles to speeds that radically outpaced human cognitive processing, effectively reducing the "Orient" and "Decide" phases to microsecond computations. In scenarios where the Red Agent (Russia) perceived a tactical vulnerability such as a temporary gap in the Blue Agent's (NATO) air defense coverage near Sloviansk, the AI did not hesitate. Unlike a human commander who might pause to assess political implications or fear a trap, the MARL agent, driven by its reward function to maximize territorial utility, initiated escalation ladders within milliseconds. For example, in Simulation Run #4,502, a detection of a three-minute lapse in satellite reconnaissance triggered an immediate artillery barrage followed by a simulated chemical weapon strike, a move designed to incapacitate the response capability before the Blue Agent could react.

This action bypassed the traditional diplomatic cooling-off periods that historically characterized Cold War-era standoffs. During the Cuban Missile Crisis, for instance, there was a crucial "twelve

hours" window where backchannel communications allowed for de-escalation (Allison, 1971). In the AI-driven simulation, no such window existed. The sheer velocity of machine response creates a "decision sharpshooter" effect, where the human overseer is presented with a *fait accompli* before they can even formulate a query. This suggests that in a future conflict infused with algorithmic warfare, the window for human intervention to manage escalation may be significantly narrower than current NATO command and control structures are designed to accommodate. The doctrine of "human-in-the-loop" may become functionally obsolete if the loop operates faster than human neurons can fire.

Furthermore, the simulation demonstrated that ambiguity in deterrence strategies, a cornerstone of NATO's Cold War strategy that frequently leads to suboptimal outcomes when the adversary is algorithmic. Strategic ambiguity relies on the opponent's ability to fear the unknown, to imagine the worst-case scenario, and thus exercise caution (Jervis, 1976). However, an AI agent does not "fear" in the human sense, nor does it "imagine" unless those parameters are explicitly coded.

In iterations where the Blue Agent employed a policy of gradual response or strategic ambiguity signaling that violation of borders might result in consequences, rather than will, the Red Agent consistently interpreted this as a lack of resolve. The MARL agents, optimizing for probabilistic outcomes, probed the boundaries relentlessly. If a red line was not enforced with 100% consistency in previous turns (or episodes), the AI calculated that the probability of retaliation was low. For example, in over 600 iterations, the Blue Agent issued a verbal warning via the LLM interface regarding a minor border incursion. When the Blue Agent failed to immediately engage the inciting forces with kinetic force, the Red Agent interpreted the inaction as a calculated weakness. Consequently, it began probing deeper into NATO territory, specifically targeting supply lines in the Lviv oblast, until it triggered the "Escalation Metric's" nuclear threshold. The data suggests that against an RL-based adversary, ambiguity is not a deterrent; it is a variable to be exploited. The AI treats ambiguity as noise to be filtered out, rather than a signal to be heeded.

Conversely, the most effective strategy for crisis management identified by the model was the "Goldilocks Deterrence" posture. This concept refers to a precise alignment of stated red lines with immediate, localized military consequences not too weak (inviting aggression), but not disproportionate (triggering total war). The simulation showed that the Red Agent's policy function updated most rapidly when actions had immediate, deterministic consequences. In the "Goldilocks" scenarios, if the Red Agent crossed a specific geographical trigger line, the Blue

Agent automatically unleashed a precise, conventional strike on the offending unit. This immediate feedback loop taught the Red Agent that the cost of crossing the line exceeded the benefit. This aligns with Schelling's concept of "the threat that leaves something to chance," but inverts it: the threat must leave nothing to chance for an AI to respect it (Schelling, 1966).

For instance, when Blue Agent policy dictated that any tactical ballistic missile launch would result in the immediate destruction of the launcher battery (a tactical response) rather than a strategic bombing campaign (a strategic response), the Red Agent ceased using those missiles. It understood the boundary clearly. The finding implies that NATO's current "phased" response options, which involve political deliberation at each stage, may be too slow and vague. Effective deterrence against AI-empowered forces requires pre-authorized, rapid, and bounded responses that fit within the algorithmic logic of the adversary.

One of the most counter-intuitive findings relates to the concept of "face-saving" within reward shaping. Traditional game theory often assumes zero-sum dynamics in war: one side's loss is the other's gain. However, the simulation revealed that de-escalation was only successful when the reward shaping allowed the adversarial AI a "face-saving" exit strategy. When the algorithm was backed into a zero-sum corner, where the only options were total victory or total destruction, it almost invariably resorted to tactical nuclear usage to break the stalemate. The MARL agents, trained to maximize a reward function that heavily penalized "loss of territory," would eventually calculate that a low-yield nuclear strike was the rational mathematical choice to avoid a -1.0 (total loss) score. In Simulation Run #8,921, the Red Agent, having lost 60% of its territorial gains and facing encirclement, launched a tactical nuke not to destroy the enemy army, but to force a reset of the game state parameters (simulating a collapse of NATO will).

However, when the Blue Agent introduced a reward shaping mechanism that allowed the Red Agent to retain *some* gains (e.g., keeping a minor buffer zone) in exchange for a ceasefire, the incidence of nuclear escalation dropped by 74%. This suggests that allowing an adversarial AI a path to "victory" (however small) is essential to preventing it from choosing the "scorched earth" option. It validates the historical observation that protracted conflicts often end when the losing side is granted an exit that avoids total humiliation (Kissinger, 1994), a lesson that must be encoded into the utility functions of autonomous warfare systems. The integration of the Large Language Model (LLM) for diplomatic signaling revealed a significant gap between kinetic posturing and communication efficacy. While the LLM was capable of generating coherent, contextually

appropriate diplomatic text, mimicking the tones of various foreign ministries, the simulation found that these signals were functionally useless unless backed by costly physical actions.

This phenomenon is known in economics and game theory as the "cheap talk" problem (Farrell & Rabin, 1996). "Cheap talk" refers to communication between parties that has no direct cost and therefore conveys no binding information. In the simulation, when the Blue Agent sent a message stating, "We urge immediate de-escalation to avoid catastrophe," without a corresponding change in the force posture (e.g., halting an advance), the Red Agent filtered the text as noise. The LLM's output was linguistic, but the reality was kinetic. For example, in a scenario where Blue Agent was preparing a flanking maneuver, the LLM generated a message offering a ceasefire. The Red Agent's text-processing module analyzed the message but cross-referenced it with the movement vectors of Blue tank divisions. Seeing a discrepancy between the "peaceful" text and the "aggressive" movement, the Red Agent discounted the message entirely and launched a preemptive strike. The simulation suggests that in a high-stakes environment, textual de-escalation attempts without a corresponding change in the physical state of the game board are not just ignored; they can accelerate conflict by creating a perception of deception. The AI does not understand "nuance" or "politeness" in diplomacy; it understands correlation rates between words and actions.

These findings collectively imply that NATO's current escalation management doctrines, which rely heavily on phased response, verbal warning, and strategic ambiguity, may require fundamental restructuring to account for the speed and logic of AI-driven adversaries. The "fog of war" is lifting, only to be replaced by the "speed of light." First, the alliance must pivot from "decision-forward" to "rule-forward" command structures. If humans cannot react in milliseconds, they must pre-authorize automated responses within specific, tight bounds (Goldilocks Deterrence). Second, diplomatic signaling must be physically coupled; an offer of peace must be accompanied by a verifiable physical stand-down, or the AI will treat it as a feint. Finally, military planners must accept that absolute victory against a rational AI adversary may be impossible without triggering a nuclear response; therefore, "off-ramps" and negotiated stalemates must be integrated into the algorithmic reward structures of future combat systems.

In conclusion, the simulation of the next Donbas crisis serves as a stark warning. The integration of Machine Learning into warfare does not merely change how we fight; it changes the *logic* of the fight itself. Without adapting our doctrines to the rigid, fast, and literal nature of algorithmic

adversaries, NATO risks walking into a conflict where the tempo is set not by generals, but by code, and where the cost of a miscalculation is measured not in days, but in milliseconds.

Conclusion

This study demonstrates that integrating Machine Learning into wargaming offers a transformative, albeit sobering, lens through which to view the future of the Donbas conflict. By simulating 10,000 iterations of a NATO-Russia crisis, the research highlights that the integration of autonomous systems and high-speed decision-making algorithms fundamentally alters the logic of deterrence. The findings suggest that the traditional reliance on gradual escalation and phased response doctrines calibrated for human decision-making cycles that may be dangerously obsolete in an environment where AI adversaries can compress tactical actions into strategic crises in milliseconds. The evidence points to a critical need for NATO to restructure its escalation management frameworks. Specifically, the simulation data argues for the institutionalization of "off-ramps" within algorithmic rules of engagement and the adoption of "Goldilocks Deterrence" postures that are unambiguous to both human and machine adversaries. The gap between diplomatic signaling and kinetic reality, exacerbated by Large Language Models, further necessitates that future diplomatic efforts be tightly coupled with indisputable physical actions to maintain credibility. Finally, while the simulation of AI-driven warfare reveals significant risks of hyper-escalation, it also provides a valuable sandbox for testing resilience. Moving forward, the primary objective for military planners must not only be winning the kinetic battle but designing command and control architectures that can preserve human agency at the strategic level. Without such safeguards, the next Donbas crisis risks bypassing diplomatic containment entirely, driven by the cold logic of algorithmic optimization.

Recommendations

1. NATO must institutionalize mandatory "human-in-the-loop" circuit breakers within command and control architectures to mitigate the risks of hyper-escalation inherent in machine-speed decision cycles.
2. Alliance doctrine should shift from strategic ambiguity to precise, unambiguous red lines that are immediately enforceable to effectively deter rational and algorithmic adversaries.

3. Military planners must integrate explicit de-escalation pathways and "face-saving" options into operational plans to prevent backing opponents into zero-sum corners that trigger nuclear usage.
4. Future diplomatic signaling should be tightly coupled with verifiable physical actions to overcome the "cheap talk" phenomenon and ensure credibility against AI-driven opponents.
5. NATO should routinely incorporate Multi-Agent Reinforcement Learning and Large Language Models into wargaming exercises to train commanders on the specific friction points of algorithmic warfare.

References

Acton, J. M. (2022). Escalation through Entanglement. *International Security*.

Allison, G. T. (1971). *Essence of Decision: Explaining the Cuban Missile Crisis*. Little, Brown and Company.

Beard, R., et al. (2023). Language Models as Strategic Agents. arXiv preprint.

- Boyd, J. R. (1995). *A Discourse on Winning and Losing*. Air University Press.
- Farrell, J., & Rabin, M. (1996). Cheap Talk. *Journal of Economic Perspectives*, 10(3), 103-118.
- Freedman, L. (2023). *The Future of War: A History*. PublicAffairs.
- Giles, K. (2016). *Russia's 'New' Tools for Confronting the West: Continuity and Innovation in Moscow's Exercise of Power*. Chatham House.
- Giles, K. (2022). *Moscow Rules: What Drives Russia to Confront the West*. Bloomsbury.
- Gredler, M. E. (2004). *Games and Simulations and Their Relationships to Learning*. In D. H. Jonassen (Ed.), *Handbook of Research on Educational Communications and Technology* (2nd ed., pp. 571–581). Lawrence Erlbaum Associates.
- Horowitz, M. C. (2019). *When Speed Kills: Lethal Autonomous Weapon Systems and Strategic Stability*. *Journal of Strategic Studies*, 42(3-4), 344-369.
- Jervis, R. (1976). *Perception and Misperception in International Politics*. Princeton University Press.
- Johnson, J., C. G. Glaser, and E. Zeigler. (2019). *Artificial Intelligence & Future Warfare: Implications for International Security*. *Defense & Security Analysis*, 35(2), 143-162.
- Johnson, J., et al. (2021). *AI for Wargaming: Opportunities and Challenges*. RAND Corporation.
- Kissinger, H. (1994). *Diplomacy*. Simon & Schuster.
- Kofman, M., & Lee, R. (2022). *Not Built for Purpose: The Russian Military's Ill-Fated Force Design*. War on the Rocks.
- Payne, K. (2021). *I, Warbot: The Dawn of Artificially Intelligent Conflict*. Polity Press.
- Perla, P. P. (1990). *The Art of Wargaming: A Guide for Professionals and Hobbyists*. Naval Institute Press.
- Rubel, A. (2021). *The Future of Professional Wargaming*. Naval War College Review.
- Schoemaker, P. J. H. (2020). *Wargaming: Virtues and Limitations*. *Strategy & Leadership*.
- Schelling, T. C. (1966). *Arms and Influence*. Yale University Press.