

TECHNO-REALISM AND AI-DRIVEN CYBER CONFLICT: A THEORETICAL ASSESSMENT OF TECHNOLOGY, INSTITUTIONS, AND STRATEGIC EFFECT

Noureen Akhtar^{*1}, Dr. Nazish Mahmood²

^{*1}PhD scholar at School of Politics and International Relations, Quaid-i-Azam University, Islamabad.

²Assistant Professor at School of Politics and International Relations, Quaid-i-Azam University, Islamabad.

^{*1}noureenakhtar_dss@live.com, ²nazish@qau.edu.pk

DOI: <https://doi.org/10.5281/zenodo.21818632>

Keywords

techno-realism, international relations theory, technology, constructivism, artificial intelligence, cyber conflict

Article History

Received: 15 May 2026

Accepted: 15 July 2026

Published: 31 July 2026

Copyright @Author

Corresponding Author: *

Noureen Akhtar

Abstract

International Relations theory has been unable to theorize technology fully, either because it is considered an exogenous input to material capability or because of its social saturation level and because its properties don't act independently but rather as a reflection of social meaning. This paper takes a purely theoretical look at techno-realism, a newish framework that posits the strategic impact of technology is channeled through the institutional and organisational identity of the state to which technology is applied or to which it belongs, as a possible fix for this problem. The paper reconstructs the key elements of the structural realism and offense-defense theory, the constructivism and the social-construction-of-technology tradition in science and technology studies, and critically analyzes the potential of techno-realism to provide a single coherent whole rather than merely a label for constructivist findings. It sees the conception of techno-realism, not as a total merger of realist and constructivist commitments, but as a justifiable, if less ambitious, theoretical division of labour, realist in its explanandum, both institutionally and partly ideationally responsive in its description of mechanism. The paper, building on this reading, outlines five theoretical research agendas for AI in cyber conflict in particular, outlining the questions that the applied literature in the field has not yet fully developed.

INTRODUCTION

There's been a longstanding and, in itself, unsolvable problem with technology in International Relations theory. Major paradigms recognize that technological change has an impact on international politics, but each one models technology as one of three ways that leave it largely under-theorized: as an exogenous event to which the system reacts, as a neutral tool whose outcomes stem directly from its physical characteristics, or as

a variable that is so subjectively socially constructed that its properties seem to have no effect whatsoever. The problem with power is a specific expression of a more general ambiguity that has long shaped the conceptualization of power in this discipline, as Barnett and Duvall (2005) maintain, between compulsory and material conceptions of power and more diffuse and structural and/or productive conceptions of power, which becomes an unresolved ambiguity in

a theory of the role of technology in power politics. Herrera (2006) refers to the theoretical space that is created as a result of this, the “uncaused cause”: technology is added to the mainstream in IR theorising as a force with consequences but without an explanation of what technology is, theoretically, or how its effects are actually produced.

This challenge has taken a new urgency in relation to the diffusion of Artificial Intelligence into military and security matters, for a reason specific to the use of AI and not to the applications of technology in general. Previous technologies, which were both materialistic and ideational, had been put pressure on IR theory, and each one of them was gradual enough to allow theory to stretch in practice incrementally rather than radically. Artificial intelligence is unique, for the purposes of this paper's argument, in that it compresses the very timescale that is typically the base of the incremental theoretical adjustment that typically has occurred when engaging with this kind of techno-issue, as the technology's impact on decision tempo discussed in great detail in Section five does apply equally to the scholarly and policy processes that have attempted to theorise and govern it as it does to the military and strategic processes that it directly confronts. A theory of technology as adequate to this moment can, in turn, only be those that effectively address and account for the underlying conflict between them—one that is not resolved by earlier, more incremental technological evolutions.

Recently, Sukma (2024) has suggested techno-realism as a possible account of this challenge, arguing that technology is a tool of power whose strategic impact is not directly inferable from the technical features of the technology, but rather depends on the institutional and organizational identity of the state that uses or absorbs it. The aim of this paper is a purely theoretical evaluation of the aforementioned proposal. It does not compare techno-realism to an empirical case, nor does it build and test the institutional-capacity model that I have developed elsewhere, and made operational

through the use of techno-realism. That kind of paper is not this one. Given this relationship, a disclosure is warranted at the beginning of this paper: This author has previously used techno-realism as the theoretical framework for applied research on the use of AI for cyber conflict in South Asia, which means this critical assessment is not completely objective. This prior commitment has been made explicit in the assessment that follows, and should be taken into account when reading the paper's critical engagement with the synthesis claim of techno-realism in Section four, as well as its conclusion that the synthesis claim requires qualification. The question addressed in this paper is more fundamental and limited than the applied work that follows, namely: Does techno-realism, on its own theoretical terms, solve the problem of technology that serves as its motivation and if so, what would a comprehensive enough theoretical treatment of AI in the context of cyber conflict specifically require of a theory seeking to do so?

The paper is divided into six parts. The second section reviews the key existing paradigms of structural realism, the offense-defense theory, constructivism, and the social-construction-of-technology (SCOT) tradition of science and technology studies, each of which has addressed technology in its own different way, thus marking the theoretical space techno-realism positions itself to fill. The central theses of techno-realism are detailed in section three. In Section four, those claims are met with a critical examination of whether techno-realism can actually be considered a theoretical synthesis or is merely a way of bringing in constructivist insight within a realist framework, and whether it can overcome the internal tensions between constructivism and realism. The fifth section introduces the concept of artificial intelligence and cyber conflict, and asks what a techno-realist research agenda for this field would have to include that current treatments, such as those developed by this author and his previous applied work, have yet to be fully matured. In the last section, the limits of techno-

realism and the future direction of its further theoretical development are discussed before the paper concludes.

2- The Problem of Technology in International Relations Theory

2.1 Structural Realism: Technology as Capability

In structural realism, technology is treated as a substantive factor in relative capability; and the distribution of technology is one of the structural properties of the system (Waltz, 1979). This treatment is a continuation of the classical realism's earlier focus on power and capability as the main driver in international politics (Morgenthau, 1948), and a more recent one that explicitly focused on the offence side of the argument, with Mearsheimer's (2001) discussion of great-power competition, in which a state's technological and military capability is seen as the main determinant of its relative power position, largely independent of the institutional processes through which it is generated and managed. This is the virtue of parsimony; that technology does not need to be theorised as an entity in its own right since its impacts are incorporated into the capability distribution the theory already tracks. It has its costs, as Herrera (2006) puts it, for structural realism, since capability is not a function of the technology used to produce it, but of the outputs of the technology, it is hard to tell the difference between two technologies that happen to produce the same capability, but differ in the way they are organised, risks are distributed, or in the way in which their acquisition changes the states that acquire them. On this treatment, technology is real, but in this theory it is an inert aggregate on an underlying capability term that is not further broken down by the theory.

But it has a specific consequence and, in the interests of this paper, significant one, if applied to AI, the structural realism would equate AI-enabled capability with any other addition to a state's material powers, only in degree; it would be a different magnitude, a more or less precise or more or less scalable or more or less intelligent or

whatever, but ultimately the same capacity. A partial corrective to this realism, provided by Gilpin (1981), is the recognition that technological change is an important factor in the transition of one hegemony to another, and that differential rates of technological and economic growth across the different states is the primary mechanism that brings about the power shifts that restructure the international system periodically; but even this more dynamic treatment of the factor places the effect of technology in a substantial role in the contribution to differential rates of technological and economic growth, rather than in the institutional processes through which a given state translates technological potential into effective strategic capacity specifically. Both frameworks struggle to encompass Horowitz's (2010) general result that the institutional and financial capacity of the adapting state determines the diffusion of military innovations and consequently the strategic impact of the innovations on the state that acquires them, absent the exact type of institutional variable that structural realism's parsimony sought to avoid. This tension between a 'parsimonious' theory which cannot account for a well-supported empirical regularity, and a 'fuller' theory that does need those resources is the specific gap the following subsections trace across three further theoretical traditions, before Section three moves on to techno-realism itself attempting to close this gap.

2.2 Offense-Defense Theory: Technology as Structural Environment

Offense-defense theory is another strand of the realist tradition that is more technologically oriented, and asks whether a particular technology is more conducive to the attacker or the defender and what impact it has on the stability of the crisis (Jervis, 1978). This literature has focused on the specific properties of technology more than structural realism as such, and subsequent developments by Van Evera (1998) and Glaser and Kaufmann (1998) build on increasingly fine-

grained measures of the offense-defense balance that a particular technological configuration is believed to produce. But this literature also assumes that the strategic significance of technology can be deduced from its material characteristics, as soon as its characteristics are properly stated, again, largely without regard to the institutional environment in which they are placed. The more recent discovery by Garfinkel and Dafoe (2019) that the oscillation of offense-defense balance is not determined by properties of the technology itself, but depends on the context in which it is situated is a movement toward the limitation identified in this paper, without solving the theoretical problem, since it merely qualifies the oscillation of the balance between offense and defense in an empirical way.

The challenge of the offense-defense tradition to integrate AI clearly exemplifies this larger point. Unlike the classical cases of the machine gun, the tank, fortified trench, etc. as AI systems are traditionally classified as either offence or defence, the same underlying AI capability, automated detection, data fusion, autonomous response, could be used in defence to detect and neutralise an aggressor's attack faster than a human operator, or in offence to identify and exploit an aggressor's vulnerabilities faster than a human defender. The offense-defense ambiguity is not a mere side issue to the offense-defense theory but rather a direct attack on its foundational premise that a technology's offense-defense character can be discerned from its material properties without reference to the institutional and doctrinal options open to those who use it, specifying precisely the gap that the mediating claim of techno-realism aims to address.

2.3 Constructivism: Technology as Socially Constituted Meaning

The most comprehensive non-materialist approach to IR theory surveyed above is the constructivist approach, which argues that the meaning and consequences of material capability, including technical capabilities, are created

socially and subjectively, but not determined by the physical characteristics of the technology (Wendt, 1992; 1999). Adler's (1997) important description of constructivism as a middle-ground between rationalist and purely reflectivist approaches to international politics is relevant here since it implies that constructivism does not exclude material reality altogether, but rather that the significance of that material reality in political life is equally an issue of intersubjective understanding that is negotiated by actors. The edited volume by Katzenstein (1996) helped establish this logic in the more general constructivist literature on security studies, which seeks to show that security interests and threat perceptions themselves are culturally and historically contingent, and not directly dependent on material distributions of capability, which forms the foundation for considering technology's security implications as similarly identity- and norm-dependent. On this analysis, things are totally different depending on who possesses and understands the technology, and thus are radically different when interpreted otherwise. This is a memorable remark by Wendt (1995) on the difficulty of analyzing capabilities in terms of material force: Five hundred British nuclear weapons are less threatening than five North Korean nuclear weapons to the United States.

If the difficulty of constructivism, from the point of view of a theory of technology specifically, is that meaning is not "read directly off" material capability, then the difficulty of structural realism is that material characteristics appear to do no independent explanatory work of their own, but are read directly off "meaning". Inverted as it is, the difficulty this paper identifies is particularly relevant to the more radical variants of constructivism which do not have a meaningful role for material constraint, and it is these more interpretivist variants that are engaged by Guzzini's (2000) reconstruction of constructivism rather than the tradition as a whole, as techno-realism's engagement in Section four will demonstrate. A

committed constructivist cannot ignore the material properties of technologies, but the theoretical framework of constructivism is more rooted in the importance of ideas and shared understandings in the social construction of social facts and does not easily indicate how material properties feed into, limit or support the social construction of meaning in the first place; this was the subject of a significant amount of research in the science and technology studies literature, reviewed below.

But in order to be meaningful, we should be specific about the reason this is a hard challenge for AI, as opposed to a theoretical one. AI systems have material properties and computational requirements and data dependencies and error rates and latency characteristics and so on and so forth, that plausibly limit what social meanings can be plausibly attached to them; a system which is known to have a high failure rate can not plausibly be socially constructed as infallible without eventually running into the material fact of failure. A constructivist explanation that sees meaning as independent of and prior to these material constraints would have a hard time accounting for why some social constructions of the reliability of a particular AI system are more stable than others; indeed, norms and shared understandings, once created, can be undermined by material events that undermine them, as noted in literature on norm dynamics, but at other times and in other contexts (Finnemore & Sikkink, 1998). The framing set up in Section three, which is techno-realist, is developed in this paper in such a way as to directly deal with that dynamic; constructivism's architecture does not provide a straightforward supply of the material properties' independent causal role.

2.4 Social Construction of Technology: Symmetry Without a State

A third treatment has come from the tradition of social construction of technology in science and tech studies, but is not related to the international relations discipline; it has direct relevance to the

problem addressed in this paper. The basic foundational statement of the approach by Bijker, Hughes and Pinch (1987) is the principle of symmetry: technological artefacts and social groups are shaped by each other in an interactive process and the material properties of the artefact have to be taken as seriously as the social context it is introduced into, otherwise you may miss the significance of it. The social shaping of technology account by Mackenzie and Wajcman (1999), which is related to the present paper, offers a complementary point relevant to the argument put forward here: that the process of design of a technology is never fully closed or determined by its material properties, but instead only a part of it is closed and the rest realised through institutional and organisational choices, a point directly analogous to that of techno-realism's admission that the strategic effect of a technology is not directly derivable from its specifications but instead mediated by the institutional context in which it is used. On its own, this symmetry principle confronts the specific problem faced by each of the international relations traditions we have looked at so far, as it avoids structural realism's over-emphasis on material capabilities, and it avoids unqualified constructivism's over-emphasis on social meaning.

For the purposes of international relations theory specifically, the weakness of the social-construction-of-technology tradition is that it was developed to explain technological development and diffusion within and between firms, the professional community and the consumer public, but it is not a theory of the specific unit (the state) that remains the focus of international relations theory, namely the strategic, security-seeking actor in the anarchic international system. There is an additional theoretical challenge in implementing the symmetry principle to interstate strategic competition, because the main cases and the main theoretical postulates of the science and technology studies literature were constructed on a different unit of analysis. That's the very act of translation to which one can read the Herrera's

(2006) own project, a translation of a constructivist-compatible, techno-receptive

approach into the field of international relations theory in general, and techno-realism in particular.

Table 1
How Four Traditions Treat Technology's Strategic Effect

Tradition	Locus of Technology's Effect	Central Limitation	Key Source
Structural realism	Material capability, read directly from technical properties	Cannot distinguish technologies producing equivalent capability via different institutional processes	Waltz (1979)
Offense-defense theory	Technology's offensive or defensive material character	Assumes offense-defense character is fixed independent of institutional/doctrinal context	Jervis (1978); Van Evera (1998)
Constructivism	Intersubjective meaning constituted through social practice	Material properties risk becoming theoretically inert relative to social meaning	Wendt (1992, 1999)
Social construction of technology (STS)	Symmetric co-shaping of artefact and social group	Developed for firms/publics, not for the security-seeking state as unit of analysis	Bijker et al. (1987)
Techno-realism	Institutional/organisational capacity mediating material capability	Coherence of the mediation claim as genuinely realist is contested (Section 4)	Sukma (2024)

Note. Summarises Section 2's survey; techno-realism's own claim, reconstructed in Section 3, is included for comparison.

Techno-Realism Reconstructed

Sukma's (2024) techno-realism makes three claims, all of which, along with his techno-realism, are an attempt to solve the problem he has surveyed in a deductive manner in Section two. The first is continuist; technology operates as a tool of power, and is in continuity with the classical realist

paradigm's focus on capacity, survival and the importance of material conditions determining state action in anarchy. The claim puts techno-realism in a position of realism rather than in opposition to it, a difference from the more radical challenge to realist ontology in the approach of constructivism.

The second claim is mediating: the strategic impact of technology is not apparent from technical specifications, but needs to be mediated by the institutional and organisational profile of the state that uses it and/or absorbs it. It is this claim that makes the paradigm's distinctive theoretical work, and in this area techno-realism diverges from the structural-realist and offense-defense readings that are the focus of Section two, both of which, in differing ways, permit a more direct reading of strategic effect from the material properties of a technology. The practice of regarding an institutional variable as a third variable between a presumed structural or material cause and an observed effect is not new in international relations theory: the literature on international institutions, for example, treats international institutions as an intervening variable that connects the underlying distribution of power to the types of cooperation or conflict that this distribution fosters (Krasner, 1982). The novelty of techno-realism is that it uses a variable not different from between two variables (the technology and the strategic impact) but analogous to it (a state's domestic institutional capacity specifically), as the variable to be explained.

The third claim is synthetic: by not placing institutional mediation in a constructivist framework, but rather a realist one, it does not give up realism's parsimony and the centrality of the state as the unit of analysis, and it does introduce the institutional and organisational sensitivity that the constructivist and science-and-technology-studies traditions bring to bear on technology specifically. Section four now critically engages with this synthetic ambition, realist commitments and institutional sensitivities to the mediated effects of technology.

3.1 Techno-Realism Against Herrera's Prior Synthesis

Situating techno-realism in the context of Herrera's (2006) earlier attempt helps to understand the particular formulation of techno-realism in Sukma's (2024) work. Whereas techno-

realism starts with realism, Herrera starts with an account of technology that is compatible with constructivist approaches, and asks what is the interaction between the two? How is it possible to reconcile an account of technology as a socially embedded system with the traditional interest in institutional and structural change in international relations theory? Central instead of techno-realism is Herrera's idea of interaction capacity, the systemic property that defines how densely and how quickly units in the international system can interact with one another, the level at which the effect of technology is felt, as opposed to the individual state's capacity to absorb or wield a given technology.

The difference of point and level of analysis is not just one of style. Herrera's systemic focus is appropriate for accounting for large-scale, long-run transformations of the sort his case studies examine, as well as the railroad's and the atomic bomb's impacts on the overall interaction capacity of the international system, but not so immediately so for explaining why two particular states, confronted by what is essentially the same techno-structural pressure in the same systemic environment, might nonetheless react very differently, with respect to the stability of their bilateral relationship, to that pressure, a comparative, dyad-level question to which the applied work of techno-realism has been honed. The two syntheses are thus complementary rather than competing: Herrera's account is more likely to account for the long-term structural transformation of the properties of the system, while the institutional-mediation claim of techno-realism is more likely to account for a difference in response to a given techno-structural pressure between otherwise equally situated states within a given systemic context. Either alone, synthesis has not yet considered the insights of the theory of international relations. A full theory of technology would probably have had to include insights from both levels of analysis.

Assessing Techno-Realism's Theoretical Coherence

Three questions arise concerning techno-realism's purported synthesis, each of which will be examined here, with the aim of constructive critical intermingling rather than uncritical adoption or dismissal.

4.1 Does Institutional Mediation Import Constructivist Commitments Without Resolving Their Tension with Realism?

The most pointed question techno-realism must answer is, is the mediating claim of the following nature, that institutional/organisational identity mediates the strategic effect of technology, doing real work, or is it a constructivist insight re-translated into realist language without a tension between the two logics being resolved. It's more than a semantics issue. The traditional definition of the state as a unit that is "only" functionally different, differentiated from other states by their relative capability, not by their institutional nature, is problematic when it is said that institutional identity is "doing a lot of theoretical work" that determines the effectiveness of the state. The theory has progressed, at least, a considerable distance from the functional-similarity assumption that classically and structurally realist theories share (Waltz, 1979).

The best answer that techno-realism can muster is that mediation through institutional capacity does not run counter to realism's core principles as long as institutional capacity is defined as an organisational capability to transform raw material or technological capacity into strategic capacity rather than as an ideational or identity variable. Not an ideational construct, as constructivism does in its main theory, but a property, a property of a state that can be, in principle, observed, measured and shown to be either present or absent, at least in an organisational sense, but in a real sense, as well: the ability to coordinate, integrate, respond and endure is a tangible, if organisationally rather than materially located, property of a state, it is present or not, and it can, in principle, be observed and measured in ways that are continuous with the way realism treats

other sorts of capability. Techno-realism, then, is innovative not in that it brings in a constructivist ontology in full, but in that it extends realism's understanding of capability to encompass computational elements as well as physical and technological elements, but also organisational and institutional capacity. If this answer is in the affirmative, a further question must be asked, namely whether, once this capacity is described with the required level of granularity, a workable theory can be described without invoking the kind of shared understanding and social practices that constructivism centralizes in its theory.

This is the question that is raised by another consideration, this time from within the realist tradition. The idea of what constitutes capability is never a consistent school, however, and the inclusion of unit-level and domestic-institutional factors, state strength, elite cohesion, and the capacity of societies to extract them from the system does not necessarily imply a rejection of the central tenets of realism (Rose, 1998). Neoclassical realism, which allows for the inclusion of domestic institutional variation without compromising realism, has an analogue in the greater diversity of the realist tradition: the use of institutional capacity as a variable intervening between a specific class of systemic pressure (in this case, cyber pressure due to artificial intelligence), relaxes, if it does not eliminate, the concern this subsection has raised. If the need to ask whether a tradition so diverse as to accommodate the ideational and unit-level variables of this sort is, in any meaningfully realist sense, realist has been a familiar criticism in the literature of neoclassical realism, it is equally applicable to techno-realism, and this paper does not solve a debate that the neoclassical realism literature has not settled; it only notes that techno-realism's own move is not unique, but shares company in a longer history of a debate regarding the boundaries of realism.

4.2 Can Institutional Capacity Be Specified Without Constructivist Content?

The theory of institutional capacity needed for the kind of strategic effects that techno-realism requires should go beyond simply stating that institutions matter to explaining which characteristics of institutions lead to the amplification or absorption of a given technology's effects. Each of these, coordination, integration, responsiveness and resilience, is a plausible candidate, but each, on inspection, requires a degree of common understanding by the actors of the situation, shared enough to act together, not at cross-purposes; shared enough to interpret an issue at the technical level and draw conclusions that are relevant to strategic judgements; shared enough to set a tempo of action for a given situation. All of these assumptions are in some sense ideational, rather than just material, and thus a fully specified understanding of institutional capacity is not without some of the same type of shared-understanding variable constructivism puts at the heart of its understanding of social action (Wendt, 1999).

The objection of this paper is that this is not, properly understood, a fatal objection against the techno-realist paradigm, although it does require it to be more explicit than Sukma's (2024) original formulation of it is on the matter of the specific sense in which it is realist despite containing ideational content of this kind. It is the most defensible, in that it is realist with respect to its explanandum, keeps the classical realist concern with capability, survival and strategic effect as the outcomes to be explained, and uses theoretical resources of the kind developed by Bijker et al. (1987) and Herrera (2006), which are compatible with constructivism and with science and technology studies - to explain how those outcomes are actually produced. In this reading, techno-realism is not a unification or reconciliation of realism and constructivism, but an accommodation of the theoretical labour that involves using realist categories to define what is to be explained, while allowing for the more ecumenical, institutionally and ideationally conscious account of what makes it. It is a lesser

claim than a complete theoretical synthesis, but, this paper suggests, it is defensible and useful, at the very least, if it is explicitly formulated as such, and not claimed to be a "seamless fusion" that is not really a fusion at all.

4.3 Is Techno-Realism Falsifiable?

A third question that might be asked is whether the techno-realism that is defined in this way, makes predictions that are sufficiently different from its competitors that they can be subjected to evidence rather than being made to fit the results of what is observed. This is a serious worry about any theory whose central mediating variable is multidimensional and partly ideational in content, like institutional capacity, and, as established in Section 4.2, can be specified in a fairly broad variety of ways, so as to make it fit almost any observed profile of stability or instability. What is defensible and what this author's own applied work has tried to operationalise would be to identify in advance the dimensions of institutional capacity and its observable measures, and explicitly those measures that would contradict the theory's predictions of the types and levels of instability, such as a materially and technologically similar pair of states that have similar outcomes of stability, but their institutional capacity clusters are convergent rather than divergent, which the theory's own logic cannot explain. The ability of techno-realism to operationalise this claim remains an empirical matter, but for this purely theoretical paper to state the theoretical requirement is clear: the claim of institutional mediation must be pre-specified in a sufficiently precise and falsifiable way or the paradigm stands a strong chance of being adduced to unfalsifiability that has been historically leveled at looser versions of constructivism.

The other test, besides pre-specification, is whether the predictions of the theory are distinguishable in practice from those of other theories. The predictions of structural realism are that stability outcomes should follow the material and technological capability of a technology, those

of offense-defense theory are that stability outcomes should follow the offense-defense character of a technology, and those of techno-realism are that stability outcomes should follow the institutional capacity of a technology, independent of, or at least not solely determined by, either. If a test of discrimination, then, requires a case, or set of cases, in which these three predictions differ, e.g., a materially weaker state with a lower off/defavour techno-likelihood than its competitor but a higher institutional capacity, a case that, according to techno-realism, is unstable, but a genuinely conducive one to the very contrary. This paper does not "retroactively" select the cases because they have come to the author's attention as fitting the theory's prediction, but seeks to identify them in advance for the purpose of providing an "avoiding confirmation bias" kind of theoretical comparison, which the qualitative-methods literature on theory-testing case studies specifies as essential, and in which techno-realism's applied literature has begun to meet the needs outlined by the kind of pre-specified falsification criteria and critical-case selection which this author has tried to provide for.

A Theoretical Research Agenda for AI and Cyber Conflict

Artificial intelligence in cyber conflict is a demanding test case for techno-realism specifically, because it combines exactly the properties that make the paradigm's mediating claim both most necessary and most difficult to specify precisely. At the same time, AI-based systems are very material—they are built with specific hardware, use specific amounts of data and computing power, generate specific measurable outputs—and are very contextually dependent—the same basic model or system can be used by different organizations with varying amounts of oversight, verification and integration into strategic decision making processes. A fully developed techno-realist research programme for this domain would require at least three things

that do not yet have a complete and precise theoretical account in applied research, such as that of this author.

First, the agenda demands a more exact description of what is special about AI, in contrast to past military technologies, as a hard case for the materialist approaches that are surveyed in Section two. The theory that this author has identified as most distinctive of AI is that it compresses the very timeframe during which institutional mediation must take place and consequently places unprecedented demands on only those capacities of institutions which the theory has identified as being most crucial, namely, capacities for tempo-compression, attribution-erosion and signalling-degradation. This is a specific claim regarding the relationship of AI to the theory's mediating variable and a full-fledged research program needs to be formulated in these exact terms. Other technologies that have imposed institutional demands on the states which acquired them, most obviously nuclear weapons, imposed such demands over a longer institutional-adjustment timescale, over decades, or even more, before their full destabilising potential was realised in a live crisis. AI's tempo-compressing property eliminates this adjustment runway, that is, the exact meaning in which the paper argues this technology poses a more difficult theoretical case than the technologies on which techno-realism's mediating claim was originally built but to which it has until now been applied without this distinction being made explicit.

Second, the agenda needs to be clear on whether institutional capacity is a technology-neutral variable or whether it needs a different or more refined specification when it comes to AI, one that could not have applied to the other technologies in their respective eras, such as nuclear weapons and precision-guided munitions. The tentative position for this paper is that generally speaking, the four dimensions are technology neutral, but that the tempo-compressing property of AI makes responsiveness a relative dimension of importance for AI that it was not for the previous technologies.

Reliability is a fact that produces a testable, comparative prediction: the responsiveness should have a greater discriminating capacity between stable and unstable outcomes in AI-driven cases than in comparable cases with earlier, less tempo-compressing technologies.

Third, the agenda of techno-realism needs to take a stance on the role of institutional mediation in the wider diffusion of AI capability among (state and non-state) actors which undermines the conceptual logic of international order regardless of shifts in relative material power relations (Kello 2017). Along with this, a second question arises, namely whether a similar mechanism between AI capability and the state's strategic effect is also present at the international system level, where the state's control over AI has been completely lost, though this paper does not attempt to answer it, and it is a significant extension of the theory that remains unanswered.

Fourth, the linkage between the institutional-mediation thesis of techno-realism and the classical security-dilemma logic found in much of the cyber-deterrence literature (Buchanan, 2017; Jervis, 1978) is also part of the agenda. If institutional capacity shapes not only how a state can accurately distinguish its own defensive AI-equipped preparations from an adversary's offensive ones, but also the severity of the underlying security dilemma AI-driven capability creates in the first place, then institutional capacity could moderate deterrence stability not only directly, as has been suggested in the past, but also the severity of the security dilemma generally that AI-driven capability is likely to foster. The present state of the proposed direction for future theoretical development is sketchy and is not fully developed because to develop it adequately would be beyond the scope of the present paper, and would require an engagement with the literature of the security dilemma that is certainly as deep as it is possible to go within the confines of this paper.

A fifth factor is related to all of the four above and is the unit of analysis in which institutional

mediation should be described: This paper's applied analysis and the research direction outlined above, predominantly reflects the state-centred unit of analysis realism typically takes. Artificial intelligence makes this choice of unit even more complex than earlier military technologies: AI systems are often designed, trained and maintained by private companies, which may or may not have a formal relationship to a given state's security apparatus, and thus the institutional capacity that will be relevant to a given AI system's strategic impact can be located in part outside the state's own security apparatus. The interaction between state-level institutional mediation and firm-level or public-private institutional arrangements would require a much more complex techno-realist account than this paper provides in this applied context, and that the mediating variable must take place at the state level; this latter assumption, however, has been largely accepted to make this paper tractable, and a more complex techno-realist account would be necessary to specify where else that the mediation might occur.

To pursue this fivefold agenda by identifying the unique relationship between AI and institutional mediation, asking whether the dimensions of institutional capacity are technology neutral or AI specific, to extend the theory's logic to systems-level diffusion, and to examine its implications on the severity of a security dilemma, and to dive into the appropriate unit of analysis in light of the public-private development process of AI, would constitute theoretical groundwork that a mature account of techno-realism in the context of AI and cyber conflict would need, before and apart from further empirical testing it might subsequently undergo.

Limitations and Future Development

This paper ends the theoretical evaluation with a candor that is appropriate for the constructive but critical tone of the evaluation in Section four of the paper. Techno-realism, as defined here, is a less ambitious theoretical success than a complete

resolution of the realist/ideational dichotomy in international relations theory; it is, rather, a mutually viable division of labour in theory, with realist categories for defining the outcome the theory should be able to explain and with resources from the constructivist tradition and the science-and-technology-studies tradition for defining the mechanism producing the outcome. It is a real contribution, the problem of technology's undertheorisation, which Herrera (2006) points to, is real, and few existing frameworks take up the challenge of describing it in a parsimonious, institutional-specific way as in the combination of parsimony and institutional specificity that characterizes techno-realism, but it is not the seamless synthesis that the techno-realistic paradigm itself sometimes implies and which remains unqualified without the more modest description of a division of labour.

The theory's predictions should be directly testable, as Section 4.3 of this paper requires, across a range of technologies, and not just artificial intelligence, in order for the theory to be falsifiable; and across a range of technologies, the more you compress the tempo, the more the theory's mediating logic should be important, which is a comparison this paper has suggested but not tested, especially responsiveness.

There is one caveat with the assessment created in this paper. In this paper no attempt has been made to illustrate the points with empirical examples because the paper is meant to have a purely theoretical character. Its arguments about the falsifiability of techno-realism and its research program for artificial intelligence are thus elaborated as theoretical statements in this paper and not as experiments in falsification that have been reported. It is true that there is an initial empirical application, in that the authors' published work used a techno-realist institutional-capacity model to test it in the cyber competition between India and Pakistan (Akhtar & Mahmood, 2026). That application should be considered to be a first step, not a final step, in the empirical testing that the theoretical assessment of this paper

calls for. The Section 4.3 requirement of discriminating and cross-case comparison is not fulfilled by just a single dyadic case, but is essential across a range of technologies and state pairs to establish the framework's falsifiability.

Conclusion

This paper is simply a conceptual analysis of the concept of techno-realism as a means to understand the role of technology in international relations in general and artificial intelligence in cyber conflict in particular. It has claimed that techno-realism addresses a real and persistent problem, that is the lack of a theory of technology in the dominant paradigms that each clears up one side of the tension between the materiality of technology and its socially mediated strategic meaning, while permanently shutting down the other. This paper, it has been argued, is a resolution of techno-realism, rather than a single combination of realism and constructivism, realist in its explanandum, and institutionally, but to some extent also ideationally, attentive in its account of the mechanism that produces strategic effect.

This reframing has another explicit precedent to be noted. However, not all that advanced research which has been productive in the study of world politics follows the techno-realist model of paradigmatic synthesis but is carried out by a pragmatic mixing of elements of different traditions to solve a certain, finite problem, without being presented as a solution for the deeper metatheoretical debate represented by these traditions. Judging by this lens, techno-realism's division-of-labour structure is not an incomplete synthesis but well-precedented example of the kind of problem-focused theorising which Sil and Katzenstein would defend: exactly the kind of eclectic theorising, in other words. The opposite argument to the one offered by Copeland (2000) on this point comes from the far side: from the argument that the best way to move forward for both traditions is not to say one is right and the other wrong, but to carefully articulate the

scope restrictions within which one is more successful than the other, and within which the other is more successful than the one.

This reframing is not meant as a devaluation of the contribution of techno-realism, but rather as a clarification which the further development of the paradigm will be better served by an explicit adoption of this paradigm. A theory which tries to be a synthesis invites just the sort of criticism that Section four has developed; namely, it has imported ideational content without having gone far enough to resolve the problem in question. A theory which tries to be a disciplined division of labour between explanandum and mechanism can incorporate ideational content without appearing to breach its basis. This paper argues that the continued theoretical viability of techno-realism is less tied to the stronger synthesis claim than to the claim's actual theoretical content being described more exactly to be able to survive this paper's critique.

To what extent this division of labour will turn out to meet the particular theoretical requirements the AI poses on it, this paper has left as a research question, not as a finding, is a question that will have to be answered by further theoretical and, in time, empirical work. What this paper has sought to provide is a more limited but, it argues, essential foundation for the research it calls upon: a clear definition of what techno-realism asserts - a more careful investigation of whether it holds together as a set of assertions on its own terms - and a delineation of what needs to be added to a more comprehensive theoretical account of AI-inflected cyber conflict. Techno-realism, on the assessment developed here, is a genuine and useful contribution to a problem international relations theory has long left unresolved, it is not, on that same assessment, a finished theory, and treating it as one would foreclose exactly the kind of critical theoretical development this paper has sought to encourage.

REFERENCES

- Akhtar, N., & Mahmood, N. (2026). Machine-speed deterrence: Institutional capacity as the moderating variable in AI-driven cyber crises. *International Journal of Emerging Multidisciplinaries: Social Science*, 5(2), 236-258.
<https://doi.org/10.54938/ijemdss.2026.05.2.676>
- Adler, E. (1997). Seizing the middle ground: Constructivism in world politics. *European Journal of International Relations*, 3(3), 319-363.
<https://doi.org/10.1177/1354066197003003003>
- Barnett, M., & Duvall, R. (2005). Power in international politics. *International Organization*, 59(1), 39-75.
<https://doi.org/10.1017/S0020818305050010>
- Bijker, W. E., Hughes, T. P., & Pinch, T. (Eds.). (1987). *The social construction of technological systems: New directions in the sociology and history of technology*. MIT Press.
- Buchanan, B. (2017). *The cybersecurity dilemma: Hacking, trust and fear between nations*. Oxford University Press.
- Copeland, D. C. (2000). The constructivist challenge to structural realism: A review essay. *International Security*, 25(2), 187-212.
<https://doi.org/10.1162/016228800560515>
- Finnemore, M., & Sikkink, K. (1998). International norm dynamics and political change. *International Organization*, 52(4), 887-917.
<https://doi.org/10.1162/002081898550789>
- Garfinkel, B., & Dafoe, A. (2019). How does the offense-defense balance scale? *Journal of Strategic Studies*, 42(6), 736-763.
<https://doi.org/10.1080/01402390.2019.1631810>

- George, A. L., & Bennett, A. (2005). Case studies and theory development in the social sciences. MIT Press.
- Gilpin, R. (1981). War and change in world politics. Cambridge University Press.
- Glaser, C. L., & Kaufmann, C. (1998). What is the offense-defense balance and can we measure it? *International Security*, 22(4), 44-82. <https://doi.org/10.1162/isec.22.4.44>
- Guzzini, S. (2000). A reconstruction of constructivism in international relations. *European Journal of International Relations*, 6(2), 147-182. <https://doi.org/10.1177/1354066100006002001>
- Herrera, G. L. (2006). Technology and international transformation: The railroad, the atom bomb, and the politics of technological change. State University of New York Press.
- Horowitz, M. C. (2010). The diffusion of military power: Causes and consequences for international politics. Princeton University Press.
- Jervis, R. (1978). Cooperation under the security dilemma. *World Politics*, 30(2), 167-214. <https://doi.org/10.2307/2009958>
- Katzenstein, P. J. (Ed.). (1996). The culture of national security: Norms and identity in world politics. Columbia University Press.
- Kello, L. (2017). The virtual weapon and international order. Yale University Press.
- Krasner, S. D. (1982). Structural causes and regime consequences: Regimes as intervening variables. *International Organization*, 36(2), 185-205. <https://doi.org/10.1017/S0020818300018920>
- Legro, J. W., & Moravcsik, A. (1999). Is anybody still a realist? *International Security*, 24(2), 5-55. <https://doi.org/10.1162/016228899560130>
- MacKenzie, D., & Wajcman, J. (Eds.). (1999). The social shaping of technology (2nd ed.). Open University Press.
- Mearsheimer, J. J. (2001). The tragedy of great power politics. W. W. Norton & Company.
- Morgenthau, H. J. (1948). Politics among nations: The struggle for power and peace. Alfred A. Knopf.
- Rose, G. (1998). Neoclassical realism and theories of foreign policy. *World Politics*, 51(1), 144-172. <https://doi.org/10.1017/S0043887100007814>
- Sil, R., & Katzenstein, P. J. (2010). Analytic eclecticism in the study of world politics: Reconfiguring problems and mechanisms across research traditions. *Perspectives on Politics*, 8(2), 411-431. <https://doi.org/10.1017/S1537592710001179>
- Sukma, I. M. (2024). Techno-realism: Navigating new challenges in the contemporary role of technology in politics. *Security and Defence Quarterly*, 46(2), 24-46.
- Van Evera, S. (1998). Offense, defense, and the causes of war. *International Security*, 22(4), 5-43. <https://doi.org/10.1162/isec.22.4.5>
- Waltz, K. N. (1979). Theory of international politics. Addison-Wesley.
- Wendt, A. (1992). Anarchy is what states make of it: The social construction of power politics. *International Organization*, 46(2), 391-425. <https://doi.org/10.1017/S0020818300027764>
- Wendt, A. (1995). Constructing international politics. *International Security*, 20(1), 71-81. <https://doi.org/10.2307/2539217>
- Wendt, A. (1999). Social theory of international politics. Cambridge University Press.