SCIENTIFIC Publishing Limited

Journal of Intelligent Communication
https://ojs.ukscip.com/index.php/jic

Review

Hybrid-Threat Intelligence: A Critical Review of Semantic Integration Challenges and the Role of the HIPSTer Ontological Framework

R. Andrew Paskauskas [1,*], Evaldas Bružė [2], Giedrė Sabaliauskaitė [3], Raminta Matulytė [1] and Tomas Lavišius [1]

[1] Security Research Laboratory, Mykolas Romeris University, 08303 Vilnius, Lithuania
[2] Lithuanian Cybercrime Center of Excellence for Training, Research and Education (L3CE), 08303 Vilnius, Lithuania
[3] Faculty of Public Governance and Business, Mykolas Romeris University, 08303 Vilnius, Lithuania
[*] Correspondence: andrew@l3ce.eu

Received: 30 December 2025; **Revised:** 18 March 2026; **Accepted:** 24 March 2026; **Published:** 9 May 2026

Abstract: Contemporary hybrid threats employ coordinated campaigns across information, cyber, and physical domains, maintaining plausible deniability while exploiting institutional vulnerabilities. This review conducts a scoping analysis following the PRISMA ScR framework (Preferred Reporting Items for Systematic Reviews and Meta-Analyses—Extension for Scoping Reviews) to evaluate Open Source Intelligence (OSINT), Social Media Intelligence (SOCMINT), and Natural Language Processing (NLP) capabilities relevant to hybrid threat detection. We systematically assess these technologies against a 12-point operational requirements framework derived from documented Russian and Chinese military OSINT methodologies and influence operation tradecraft. The analysis incorporates high-Technology Readiness Level (TRL) European initiatives to ground capability assessments in operational experience. While individual analytical disciplines are technically advanced, current defensive systems remain siloed and lack the cross-domain reasoning necessary to correlate technical cyber indicators with coordinated narrative manipulation. Requirements for cross-platform correlation and adversarial adaptation show only prototype stage coverage. Our findings reveal a persistent "semantic gap": defensive systems collect extensive data but lack integrated semantic reasoning across domains and languages. To address this, we examine ontology-based approaches as architectural solutions, positioning the 'Hybrid Information Psychological Societal Threats handling system' (HIPSTer) framework as an illustrative case. HIPSTer specifically targets the multilingual nature of hybrid threats—particularly in Russian and Chinese contexts—achieving TRL-4 validation through high-efficiency semantic vectors and formal reasoning across diverse language benchmarks. Finally, the review analyzes how European regulations—including the General Data Protection Regulation (GDPR), AI Act, and Network and Information Systems Directive 2 (NIS2)—shape operational architectures through compliance by design imperatives. We conclude by outlining a prioritized research agenda to advance European hybrid threat detection toward operational maturity.

Keywords: Hybrid Threats; OSINT; SOCMINT; NLP; Cyber Threat Intelligence; Ontological Framework; HIPSTer

1. Introduction

Hybrid threats represent one of the most pressing challenges to contemporary security frameworks, combining conventional and unconventional means to exploit vulnerabilities across multiple domains simultaneously. Unlike traditional military confrontations or purely cyber operations, hybrid threats blend disinformation cam-

paigns, cyber-attacks, economic pressure, diplomatic manipulation, and kinetic actions into coordinated strategies designed to remain below the threshold of conventional armed conflict while achieving strategic objectives. The European Union's recognition of this challenge is reflected in sustained policy attention, from the 2016 Joint Framework on Countering Hybrid Threats through the recent Eighth Progress Report [1], which documents an escalating threat environment characterized by increasingly sophisticated foreign information manipulation and interference (FIMI) operations, critical infrastructure targeting, and the weaponization of migration flows. Russia's ongoing campaign against Ukraine exemplifies the contemporary hybrid threat paradigm, wherein military operations occur alongside massive disinformation efforts, cyber sabotage of energy infrastructure, and coordinated influence campaigns targeting Western democracies' political cohesion. The EU Parliament's October 2025 motion on hybrid provocations [2] underscores how these threats now extend from the information space into physical and military domains, demanding comprehensive detection and response capabilities.

The multi-domain character of hybrid threats requires intelligence capabilities that operate across diverse information environments and analytical methodologies. Open-Source Intelligence (OSINT) provides foundational monitoring of adversarial narratives, military movements, and strategic communications across publicly accessible sources ranging from state media to technical forums [3]. Social Media Intelligence (SOCMINT) extends this capability into the fluid, high-velocity environment of social platforms, where coordinated inauthentic behavior, bot networks, and influence operations unfold in real time across Twitter, Telegram, VKontakte, and emerging platforms [4]. Natural Language Processing (NLP) [5] and Large Language Models (LLMs) offer computational methods to process the massive data volumes these environments generate, enabling semantic analysis, entity extraction, sentiment assessment, and increasingly, the detection of AI-generated disinformation that now accounts for over 80% of observed social engineering attacks according to ENISA's 2025 Threat Landscape [6]. The convergence of OSINT, SOCMINT, and NLP is not merely additive but multiplicative, as effective hybrid-threat intelligence depends on correlating signals across information spaces to distinguish coordinated campaigns from organic discourse, technical hybrid methodologies from strategic hybrid threats, and state-sponsored operations from opportunistic actors.

Despite significant investment in OSINT platforms, SOCMINT tools, and NLP capabilities, the current landscape remains characterized by analytical silos that impede the integrated intelligence hybrid threats demand. Commercial platforms such as Palantir Gotham, Maltego, and Babel Street offer powerful capabilities within their respective domains but lack native mechanisms for cross-domain semantic reasoning. Law enforcement agencies and intelligence services frequently maintain separate workflows for open-source monitoring, social media analysis, and linguistic assessment, creating friction points that result in the loss of critical contextual signals. The proliferation of Large Language Models introduces both opportunities and challenges: while LLMs enable unprecedented scale in content analysis and generation, they simultaneously enable adversaries to produce synthetic media, deepfakes, and contextually sophisticated disinformation at minimal cost. Recent research demonstrates that LLM-generated misinformation exhibits linguistic patterns distinguishable from human-authored content [7], yet detection methods remain largely reactive rather than integrated into operational intelligence frameworks. The semantic gap between what these tools collect and what analysts require for actionable threat assessment represents a fundamental challenge that technical capability alone cannot resolve.

This review addresses the need for a comprehensive synthesis of hybrid threat intelligence capabilities, methodologies, and operational lessons learned from high-TRL European initiatives. While individual surveys have examined OSINT methodologies [8], SOCMINT detection techniques [9], or NLP applications in cybersecurity [10], no recent work provides an integrated analysis of how these domains must converge to address the hybrid threat challenge specifically. We position this review within the context of emerging European capabilities, including the STARLIGHT project's AI-enabled law enforcement tools [11], France's VIGINUM service for foreign digital interference detection [12], and Lithuania's NAAS platform for narrative analysis [13], examining how operational deployments reveal both technical achievements and persistent integration gaps. Our analysis distinguishes between hybrid methodologies (technical approaches combining multiple data sources or analytical techniques) and genuine hybrid threats (coordinated state or state-sponsored campaigns), a distinction that is often absent in the existing literature [14]. Furthermore, we incorporate the regulatory dimension, increasingly constraining operational practice, and examine how GDPR restrictions, AI Act transparency requirements, and data protection principles create compliance-by-design imperatives that shape operational threat intelligence architectures within the European

legal framework.

Following the review article methodology outlined in Section 1.1 below, the review proceeds through seven substantive sections. Section 2 examines adversarial operational tradecraft through Russian and Chinese military information operations, establishes the threat context, and derives operational requirements (Section 2.4) that defensive capabilities must address. Sections 3 through 5 provide critical analyses of OSINT, SOCMINT, and NLP/LLM capabilities, respectively, synthesizing academic literature with operational platform assessments to identify the current state-of-the-art and persistent gaps. Section 6 reviews high-TRL European initiatives that provide empirical evidence of integration challenges and operational lessons learned. Section 7 examines ontological approaches to cross-domain integration, positioning semantic reasoning frameworks as potential solutions to identified gaps. Section 8 synthesizes findings into a prioritized research agenda that accounts for technical, operational, and regulatory constraints. Our methodology relies exclusively on published literature, operational reporting, and publicly documented capabilities, adhering to the convention of a review article, which presents a critical synthesis without new empirical data. Close to 90 sources have been examined and span peer-reviewed academic publications, technical conference proceedings, EU policy documents, and operational system documentation, providing a comprehensive evidence base for our analysis and recommendations.

1.1. Review Article Methodology

This paper follows the PRISMA Extension for Scoping Reviews (PRISMA-ScR) framework [15], which provides structured reporting guidelines suitable for reviews that map the breadth of a research area and identify critical gaps rather than answer narrowly defined empirical questions. A scoping review approach was selected because hybrid-threat intelligence spans heterogeneous disciplines, cybersecurity, intelligence studies, natural language processing, regulatory analysis, and semantic-web technologies, whose diverse methodological traditions and publication venues preclude the standardized effect-size comparisons required for systematic reviews or meta-analyses.

1.1.1. Eligibility Criteria

Sources were included if they addressed one or more of the following: (a) OSINT, SOCMINT, or NLP capabilities applied to threat detection or intelligence analysis; (b) hybrid-threat characterization, including state-sponsored information operations, coordinated inauthentic behavior, or multi-domain campaign coordination; (c) ontological or semantic-integration approaches to security intelligence; or (d) European regulatory frameworks governing AI-enabled surveillance and threat-detection systems. Sources were excluded if they focused solely on kinetic military operations without an information or cyber component, or if they consisted of commercial marketing materials lacking independent evaluation or peer-reviewed validation. No language restrictions were imposed, consistent with the review's emphasis on multilingual capability assessment (R11).

1.1.2. Information Sources and Search Strategy

Literature searches were conducted across Scopus, Web of Science, and IEEE Xplore using combinations of terms including "hybrid threat," "OSINT," "SOCMINT," "social media intelligence," "disinformation detection," "ontology AND threat intelligence," "NLP AND cybersecurity," and "coordinated inauthentic behavior." These searches were supplemented by targeted retrieval of EU institutional documents (European Commission, ENISA, European Parliament), operational system documentation (NAAS, STARLIGHT, VIGINUM), and relevant legal instruments (GDPR, AI Act, NIS2 Directive, and applicable CJEU jurisprudence). Reference-list screening of included articles provided additional sources. The search period covered January 2018 to December 2025, with seminal earlier works included where foundational to the field.

1.1.3. Selection and Synthesis

Source screening and selection were performed by the lead author with verification by co-authors, following an iterative process in which initial broad retrieval was refined through relevance assessment against the 12-point operational-requirements framework established in Section 2.4. Given the heterogeneity of evidence types—spanning empirical studies, system documentation, policy analysis, and legal instruments—a narrative synthesis approach was employed rather than meta-analysis. Over 80 sources were retained in the final review, compris-

ing peer-reviewed academic publications, technical conference proceedings, EU policy documents, and operational system documentation.

2. Adversarial Contexts and Operational Tradecraft

Understanding defensive hybrid threat intelligence capabilities requires first examining the adversarial operational tradecraft these systems must detect and characterize. Russian and Chinese military information operations exemplify distinct yet increasingly sophisticated approaches to leveraging open sources and social media platforms for strategic advantage, establishing the threat baseline against which European detection frameworks operate.

Russia and China represent the most sophisticated and active state-level adversaries employing hybrid threat methodologies, combining technical capabilities, institutional resources, and operational frameworks that define the upper boundary of threats European defensive systems must address. Russian operations exemplify integration of cyber warfare, information operations, and proxy forces in coordinated campaigns (Ukraine, Baltic states), drawing from Soviet-era Maskirovka traditions of deception and disinformation updated for digital environments where information velocity overwhelms traditional counter-intelligence methods. Chinese approaches demonstrate long-term strategic positioning through coordinated media manipulation, economic leverage, and technical espionage, guided by doctrine, integrating psychological warfare, media warfare, and legal warfare into comprehensive influence operations. These two threat actors establish the operational benchmark: defensive systems capable of detecting Russian and Chinese hybrid operations will necessarily address threats from less capable actors, though the inverse does not hold. Both adversaries explicitly theorize and institutionalize these approaches within military doctrine, providing documented operational tradecraft against which European capabilities can be systematically assessed.

The European Union's institutional framing explicitly identifies Foreign Information Manipulation and Interference (FIMI) as a primary hybrid threat vector, characterized by coordinated campaigns employing inauthentic news articles, fabricated investigations, AI-generated deepfakes, and cross-platform narrative manipulation designed to interfere in elections, discredit public institutions, and erode democratic resilience [1]. Indeed, ENISA's 2025 Threat Landscape [6] documents that approximately a quarter of FIMI content focuses specifically on degrading the Union through negative narratives targeting France, Germany, and Poland, with accusations aimed at discrediting their governments, militaries, and intelligence services. These operations increasingly integrate information warfare with cyber-physical activities: investigative reporting, detailed pro-Russia groups using Telegram to recruit EU-based individuals for sabotage, vandalism, arson, and influence operations across the North Atlantic Treaty Organization (NATO) countries. This operational integration defines the threat baseline against which European detection capabilities must operate, justifying the focus on Russian and Chinese adversarial tradecraft examined throughout this review.

2.1. Russian Military OSINT and Information Warfare

Russian information warfare doctrine integrates military, intelligence, and psychological operations into a unified concept termed "information confrontation," wherein narrative control constitutes a domain of conflict co-equal with kinetic operations. This approach draws from Soviet-era *Maskirovka* traditions of deception, denial, and disinformation, updated for digital environments where information velocity and volume overwhelm traditional counter-intelligence methods [16]. The Russian General Staff's acknowledgement that contemporary conflicts co-occur across physical, information, and cognitive domains reflects a doctrinal sophistication that Western military establishments only recently began matching. Operations in Ukraine since 2014 demonstrate this integration in practice, with military actions coordinated alongside cyber-attacks on critical infrastructure, targeted disinformation campaigns on social platforms, and strategic communications through state-controlled media outlets such as RT and Sputnik to shape international perceptions while degrading Ukrainian command effectiveness [17].

VKontakte [18] and Telegram [19] emerge as primary platforms for Russian OSINT collection and influence operations, exploiting these services' widespread adoption in post-Soviet spaces and relatively permissive content moderation policies. VKontakte's 97 million monthly active users across Russia, Ukraine, Belarus, and Kazakhstan create a target-rich environment for monitoring military movements, identifying personnel through geotagged posts, and conducting network analysis to map organizational structures. Russian intelligence services systemati-

cally harvest open-source data from social platforms to supplement traditional collection, identifying military equipment transits, indicators of troop morale, and civil defense preparations through aggregated analysis of seemingly innocuous posts. Telegram's encrypted channels enable both collection and dissemination; Ukrainian defensive positions inadvertently revealed through soldier posts become targeting data, while Russian-operated channels coordinate information operations across language communities. The "grey zone" character of these activities—conducted through accounts with ambiguous attribution—complicates defensive responses, as distinguishing state operations from patriotic hackers or organic discourse requires correlation across multiple indicators.

Russian influence operations demonstrate increasing technical sophistication in content generation, network coordination, and narrative amplification. The Internet Research Agency's 2016 operations [20] targeting U.S. elections established templates subsequently refined: coordinated inauthentic behavior employs bot networks to amplify divisive content, authentic-appearing personas build credibility over months before deploying disinformation, and cross-platform coordination ensures narratives achieve saturation before fact-checking responses emerge. More recent campaigns show adaptation to platform countermeasures, using smaller and more resilient networks that evade detection algorithms focused on behavioral signatures. The integration of generative AI tools lowers operational costs while increasing output volume and linguistic sophistication, enabling campaigns to operate across multiple languages simultaneously with contextually appropriate messaging. The Portal Kombat network identified by France's VIGINUM [12] service in 2024 exemplifies this evolution: over 1,000 websites across European languages promoted coordinated pro-Russian narratives through culturally tailored content that exploited specific national debates. Defensive systems must therefore detect not merely technical indicators, such as bot behavior, but also semantic patterns that reveal coordinated narrative deployment across heterogeneous platforms and languages.

2.2. Chinese Military OSINT and Information Operations

Chinese military information operations reflect a distinct strategic culture and operational priorities while increasingly adopting techniques pioneered in Russian operations. The People's Liberation Army conceptualizes information warfare through the "Three Warfares" framework [21]: psychological warfare targeting adversary decision-maker perceptions, media warfare shaping domestic and international narratives, and legal warfare establishing favorable interpretations of international norms. This approach emphasizes long-term positioning over tactical opportunism, building information infrastructure and influence networks during peacetime to activate during crises. The Party-State-Military integration characteristic of Chinese governance enables coordination across ostensibly civilian entities—technology companies, educational institutions, diaspora organisations—that Western operational concepts struggle to categorise as intelligence activities.

Weibo, WeChat, and Douyin (TikTok's Chinese version) [22] serve as primary platforms for both domestic information control and international influence operations, though their roles differ substantially from Russia's platform usage [23]. In this context, Chinese authorities maintain comprehensive content moderation regimes that filter domestic discourse while tolerating—and occasionally amplifying—nationalism and anti-Western sentiment that serves state interests. For intelligence collection, these platforms also provide insights into Chinese military modernization, defense-industry developments, and elite political dynamics through analysis of patterns in officially permitted discussion. PLA Daily and other military publications provide overt strategic communications while technical forums and academic journals reveal capability developments, requiring specialized linguistic and cultural expertise to interpret correctly. The challenge for external OSINT analysts lies not in platform access—many Chinese platforms permit foreign observation—and more in linguistic barriers, cultural context requirements, and the risk of consuming deliberately misleading information placed for counterintelligence purposes.

Cross-platform coordination in Chinese information operations increasingly mirrors Russian sophistication while maintaining distinctive characteristics. Rather than employing bot networks extensively, Chinese operations often leverage genuine accounts—diaspora communities, nationalist volunteers, and coordinated authentic behavior by users responding to official guidance [24]. This approach produces operations more resilient to technical detection, as accounts exhibit authentic behavioral patterns while still advancing coordinated narratives. TikTok's algorithm-driven content distribution enables influence operations that exploit platform mechanisms rather than requiring overt coordination [25]; strategically placed content that resonates with target audiences receives organic amplification, effectively weaponizing the platform's commercial recommendation systems [26]. The language bar-

rier works bidirectionally: while Chinese operations targeting European audiences require translation capabilities (increasingly provided by LLMs), European monitoring of Chinese-language platforms remains limited by analyst availability and the inadequacy of machine translation for nuanced political discourse.

Recent research on Chinese information operations reveals tactical successes but strategic failures: while specific campaigns achieve temporary narrative dominance, sustained operations often generate backlash that damages Chinese soft power objectives. Opinion surveys across developed democracies show dramatic declines in favourable views of China, with South Korea experiencing a particularly sharp drop to 19% favourable by 2022 [27]. This paradox, effective tactical operations producing counterproductive strategic outcomes, suggests vulnerabilities that defensive systems might exploit through transparency operations that expose coordination and attribution. However, the growing use of AI-generated content, cross-platform narrative seeding, and exploitation of platform-recommendation algorithms indicates continued evolution toward more sophisticated and harder-to-detect operations.

2.3. Operational Challenges for Detection Systems

The adversarial tradecraft examined above establishes several requirements for practical defensive hybrid threat intelligence. First, detection systems must operate across multiple platforms simultaneously, as modern operations deliberately fragment activities to avoid single-platform detection. Second, semantic analysis capabilities must extend beyond keyword matching to identify coordinated narrative deployment across languages and cultural contexts. Third, attribution methodologies must account for deliberate ambiguity and multiple-tier operations where direct state involvement remains deniable. Fourth, temporal analysis matters: distinguishing long-term positioning from acute campaign activation requires historical baselines and anomaly detection. Finally, defensive capabilities must anticipate adversary adaptation to detection methods, as both Russian and Chinese operations demonstrate learning from platform countermeasures and adjusting tactics accordingly. These operational challenges—cross-platform fragmentation, semantic sophistication, deliberate ambiguity, temporal complexity, and adversarial adaptation—require formalization into specific technical requirements for systematic capability assessment, addressed in Section 2.4 below.

2.4. Requirements Framework for Defensive Capability Assessment

Building on the operational challenges identified above, this section formalizes the 12-point requirements framework used throughout the review to assess defensive hybrid-threat detection capabilities. The framework operationalizes the analytical criteria referenced in Section 1.1 and provides the structure through which all included sources were evaluated. The requirements draw substantially from documented Russian and Chinese information-operations tradecraft, which establishes the upper boundary of adversarial sophistication that European defensive systems must address.

The framework comprises two categories: detection-related requirements and operation-related requirements, reflecting the dual need to identify hybrid-threat activity and to integrate signals across domains, languages, and temporal patterns.

2.4.1. Detection-Related Requirements

- R1. Multi-source data harvesting: Automated collection from social platforms, news outlets, forums, and darknet sources matching the cross-domain information gathering demonstrated by Russian and Chinese operations.
- R2. Network structure mapping: Social network analysis revealing coordination patterns, influence hierarchies, and organizational relationships characteristic of state-sponsored troll operations and proxy networks.
- R3. Targeted campaign detection: Identification of coordinated disinformation, influence operations, and narrative manipulation across platforms as observed in Russian electoral interference and Chinese anti-Taiwan campaigns.
- R4. Strategic communications analysis: Monitoring state-sponsored media outlets (RT, Sputnik, CGTN), proxy organizations, and coordinated messaging architectures.
- R5. Toxic content identification: Detection of hate speech, extremism, incitement, and dehumanizing language indicative of psychological operations designed to polarize target societies.

2.4.2. Operation-Related Requirements

- R6. Cross-platform capability: Simultaneous monitoring across multiple platforms (Twitter/X, Facebook, Telegram, VKontakte, WeChat) to detect fragmented operations that Russian and Chinese actors deliberately distribute to avoid single-platform detection.
- R7. Semantic analysis beyond Context-aware understanding of coordinated narrative deployment across languages, cultural contexts, and coded communications that keyword-based systems fail to identify.
- R8. Attribution under ambiguity: Methodologies accounting for deliberate deniability, multi-tier operations (state → proxy → amplifier networks), and false-flag tactics characteristic of Russian gray-zone operations.
- R9. Temporal discrimination: Distinguishing long-term strategic positioning (years-long cultivation of proxy networks) from acute campaign activation (coordinated hashtag campaigns, targeted amplification).
- R10. Adversarial adaptation anticipation: Robustness to evolving adversary tactics as detection methods become known—critical given documented Russian and Chinese adaptation to platform moderation and detection systems.
- R11. Multilingual capability: Processing diverse European target languages (German, French, Polish, Baltic languages) and adversary languages (Russian, Mandarin) with comparable performance.
- R12. Operational cyber resilience: Robustness to cyberattacks, adversarial manipulation of detection systems, and infrastructure disruption given adversary capabilities in offensive cyber operations.

This requirements framework provides the evaluative structure for Sections 3–6 and underpins the synthesis presented in Section 8.

3. OSINT Capabilities and Integration Challenges

Open-Source Intelligence (OSINT) tools have become increasingly vital for detecting and analyzing hybrid threats from both state and non-state actors, leveraging vast amounts of publicly available data from sources including social media platforms, public databases, news outlets, technical forums, and geospatial information systems. These tools, particularly when combined with artificial intelligence and machine learning, can efficiently identify cyber threats, vulnerabilities, and sophisticated tactics such as botnet domain generation or malware concealed via steganography, often achieving high accuracy in near real-time processing environments. This section assesses OSINT capabilities against the requirements framework established in Section 2.4, with particular attention to multi-source data harvesting (R1), network structure mapping (R2), and strategic communications analysis (R4), including R6–8, and R11.

3.1. Capabilities

Contemporary OSINT research demonstrates impressive capabilities across multiple analytical domains. Suryotrisongko et al. [28] report the integration of explainable AI with OSINT through a random-forest model that achieves 96% accuracy in detecting botnet domain-generation algorithms while remaining robust against adversarial attacks. The authors note that the system's OSINT layer provides human-readable evidence that increases analyst trust in automated alerts, addressing a persistent challenge in AI-driven intelligence systems. Yadav, Kumar, and Singh [8] provide a comprehensive mapping of current OSINT data sources—from public databases to encrypted messaging platforms like Telegram—while reviewing AI and machine learning techniques deployed across national security, digital forensics, and cybercrime applications. Their analysis identifies the progression toward fully autonomous analytical models as the field's most pressing research frontier, though they acknowledge that significant challenges remain in achieving this goal.

The integration of large language models represents a notable advancement in OSINT automation. Yuan et al. [29] demonstrate how LLMs with external API access can autonomously decompose user queries and harvest fresh threat intelligence, with their chain-of-thought approach reportedly outperforming vanilla LLMs across four threat intelligence benchmarks. This capability addresses the challenge of maintaining current awareness across rapidly evolving threat landscapes. Similarly, Bizouarn, Abdulnabi, and Tan [30] report coupling clustering and machine learning with OSINT scraping to identify corporate infrastructure weaknesses, formatting heterogeneous open-source findings into actionable vulnerability reports that lower barriers for small organizations conducting continuous security assessments.

However, OSINT tools face notable limitations that become particularly evident in hybrid threat scenarios. Fauziyyah, Adrian, and Alam [31] examine image-embedded malware detection, revealing how current OSINT pipelines struggle with multimodal forensics: while VirusTotal's multiple scanning engines provide reasonable baseline detection, the authors report that detection rates drop sharply once payloads are concealed through symmetric-key steganography. This highlights fundamental gaps in current approaches to analyzing multimedia content that hybrid threat actors increasingly exploit. Nonum, Avwokuruaye, and Ezemonye [32] trace OSINT's evolution from military intelligence origins while cataloging privacy, legal, and ethical constraints that limit large-scale, AI-driven collection. Ivkova and Opirskyi [3] extend this analysis by examining information security risks exposed through OSINT collection itself, proposing granular data classification and leak-monitoring countermeasures to protect both personal and state assets.

Automation frameworks demonstrate both the sophistication of current tools and their operational constraints. Vacas, Medeiros, and Neves [33] developed IDSoSint, an automated ingestion system that converts 49 public data streams into real-time Snort intrusion detection rules, enabling the detection of botnet command-and-control traffic, brute-force attempts, and phishing without manual configuration. Shin and Jung [34] present a comparative analysis of four automation suites—TheHarvester, FOCA, Metagoofil, and Recon-NG—revealing complementary strengths in metadata mining versus email harvesting while identifying the need for orchestration frameworks that minimize analytical overlap and false positives.

The commercial OSINT landscape showcases advanced integration capabilities that academic research prototypes often lack. Babel Street [35] exemplifies state-of-the-art multilingual OSINT monitoring, conducting continuous searches across thousands of global sources in over 200 languages while providing adaptive text analytics, entity matching, and rich analytical perspectives, including geospatial mapping, temporal trend analysis, and sentiment assessment. The platform's strength lies in data integration—ingesting feeds from 30+ social media platforms alongside client databases and third-party sources into unified analytical environments with customizable dashboards featuring heat maps, network graphs, and collaborative workspaces. Babel Street's applications span large-event security monitoring and border threat detection, demonstrating operational deployment at scale.

WebIQ's Voyager suite [36] represents another high-capability commercial solution tailored for law enforcement and intelligence agencies. The platform excels in real-time monitoring across surface web, social media, Telegram channels, darknet forums, and chat boards, providing advanced filtering and automated alerting when relevant information emerges. Its Visual Media Analytics component integrates cutting-edge AI for semantic image search, audio transcription, perceptual hash matching, EXIF metadata extraction, and optical character recognition—capabilities that address the multimedia analysis gaps identified in academic research. WebIQ's modular architecture includes DarkCloud for dark web crawling and Atlas for specialized investigations, providing comprehensive coverage across overt and covert information spaces.

Traditional enterprise platforms like Palantir Gotham [37] and IBM i2 Analyst's Notebook [38] provide powerful capabilities for integrating diverse data sources into unified analytical frameworks with sophisticated visualization and machine-learning-driven predictive analytics. Palantir's strength lies in seamless multi-source data fusion, while IBM i2 specializes in social network analysis and pattern recognition for unraveling complex relationship networks. Open-source tools like Maltego [39] offer network mapping and entity relationship visualization capabilities. At the same time, programming libraries such as Python's Scrapy enable large-scale automated web crawling and data extraction at volumes far exceeding those achievable with manual collection methods. Data manipulation frameworks like Pandas and visualization libraries including Matplotlib provide essential infrastructure for processing and analyzing large OSINT datasets.

3.2. Critical Gaps in OSINT for Hybrid Threat Detection

Despite significant advances in OSINT automation and analytical sophistication, current approaches reveal three fundamental limitations when addressing hybrid threat scenarios.

3.2.1. Domain-Specific Analytical Silos

While individual studies and tools demonstrate sophisticated capabilities, each operates within narrow technical domains. Suryotrisongko et al. [28] focus exclusively on DNS-based botnet detection, Fauziyyah et al. [31] examine image-based steganography, and Bizouarn et al. [30] concentrate on infrastructure vulnerabilities. Com-

mercial platforms similarly specialize: Babel Street excels at multilingual information monitoring, WebIQ targets multimedia analysis, and Palantir focuses on data fusion and visualization. Hybrid threats, by definition, coordinate activities across multiple domains simultaneously—combining cyber operations, information campaigns, and physical activities—yet no current OSINT framework provides systematic cross-domain correlation capabilities that would enable analysts to connect botnet activity patterns with coordinated social media campaigns or link infrastructure vulnerabilities to broader influence operations.

3.2.2. Limited Contextual Integration

Current research demonstrates impressive automation in data collection and processing, particularly Yuan et al.'s [29] LLM-based orchestration and Vacas et al.'s [33] real-time feed integration. However, these approaches lack mechanisms for contextualizing threat indicators within broader campaign patterns characteristic of hybrid operations. OSINT tools can identify whether a domain generation algorithm is active or whether malware is present, but cannot systematically correlate these technical indicators with coordinated disinformation campaigns or geopolitical context to reveal hybrid threat attribution and intent. Commercial platforms aggregate and visualize vast amounts of data but leave higher-level synthesis—determining whether content represents part of a coordinated campaign and assessing its strategic purpose—to human analysts.

3.2.3. Absence of Multi-Modal Reasoning Frameworks

The literature demonstrates sophisticated analytical techniques within individual data modalities, yet no study provides frameworks for reasoning across the diverse information types that hybrid threats exploit. Hybrid operations typically combine technical cyber indicators, social media manipulation, geospatial intelligence, and traditional media sources in coordinated campaigns. Current OSINT research and commercial tools remain compartmentalized by data type and analytical technique: image analysis tools operate separately from text analytics, network mapping functions independently of sentiment analysis, and cyber threat indicators exist in isolation from information operations assessment. This fragmentation prevents systematic correlation of seemingly disparate indicators that collectively reveal hybrid threat campaigns.

The capabilities reviewed represent remarkable technical achievements in specialized domains. Yet, the fundamental challenge persists: hybrid threats exploit the seams between analytical disciplines, coordinating operations across domains that current OSINT approaches analyze in isolation. Effective hybrid threat detection requires not merely more sophisticated individual tools, but also semantic integration frameworks that enable systematic correlation of threat indicators across the information, cyber, and physical domains that contemporary adversaries strategically exploit. Section 7 provides a systematic assessment of these capabilities against the full requirements framework (R1–R12), revealing that while OSINT achieves strong coverage for multi-source harvesting (R1) and strategic communications monitoring (R4), the coverage is merely adequate in cross-platform correlation as critical gaps persist (R6), and is only partial for semantic reasoning (R7) and attribution methodologies (R8).

4. SOCMINT for Hybrid Threat Detection

Social Media Intelligence (SOCMINT)—the systematic collection and analysis of data from platforms such as X/Twitter, Facebook, Telegram, and darknet forums—has become central to understanding and countering hybrid threats that blend cyber-attacks with large-scale influence operations. These tools have become increasingly effective in identifying and analyzing threats from both state and non-state actors, leveraging advances in artificial intelligence and machine learning to detect extremist propaganda, disinformation campaigns, and radical behavioral traits by analyzing vast amounts of unstructured social media data, including text, images, and network patterns, often in real time. This section evaluates SOCMINT capabilities against requirements established in Section 2.4, with particular focus on network structure mapping (R2), targeted campaign detection (R3), toxic content identification (R5), and cross-platform capability (R6).

AI-driven frameworks, including deep learning models and hybrid machine learning pipelines, have demonstrated high accuracy in identifying malicious posts, mapping events, and detecting automated bot accounts that spread misinformation or coordinate influence operations. However, analysis of the current literature reveals a significant conceptual divide between studies that explicitly address hybrid threats as coordinated, multi-domain

security phenomena and those that employ advanced "hybrid" methodologies to address traditional cybersecurity problems. This distinction proves critical for understanding both the capabilities and limitations of contemporary SOCMINT approaches.

4.1. SOCMINT Research Addressing Genuine Hybrid Threats

Only a subset of current SOCMINT research explicitly targets hybrid threats as coordinated campaigns spanning multiple operational domains. Mlinac [40] examines how AI systems themselves create hybrid threats through disinformation and influence operations in cyberspace, arguing that hybrid intelligence employed by social networks enables more effective exploitation of weaknesses in political and social systems, complicating counteraction and timely recognition of hybrid threats. This work emphasizes how adversaries leverage the sociotechnical properties of platforms themselves as weapons rather than merely as communication channels.

Dover [4] provides a critical assessment of Western SOCMINT capabilities in hybrid conflict scenarios, arguing that Western powers are systematically losing the information component of hybrid conflicts through adversaries' superior use of SOCMINT. Dover's analysis reveals that while SOCMINT provides valuable horizon-scanning capabilities for identifying emerging narratives and actor networks, it offers limited value for real-time warning during rapidly evolving hybrid conflicts—precisely when coordinated threat correlation becomes most critical. This assessment highlights a fundamental gap between SOCMINT's technical capabilities and its operational utility in actual hybrid threat scenarios.

Cárdenas et al. [41] address hybrid threats as non-traditional challenges to national security, proposing social media analysis as a detection tool for identifying radical behavioral traits and potential security threats. Their work emphasizes the integration of entity extraction, sentiment analysis, and content analytics to detect instability scenarios and radicalization, supporting timely alerts and crisis interpretation. However, their approach remains primarily focused within the social media domain rather than providing frameworks for correlating social media intelligence with cyber or physical threat indicators.

Mothe et al. [42] present instruments and tools designed to counter hybrid security threats that combine physical and cyber attacks, focusing specifically on identifying radical textual content across multiple languages and platforms. Their research acknowledges the multidomain nature of hybrid threats but focuses analytical capabilities on text-based content analysis rather than on cross-domain intelligence integration.

4.2. SOCMINT Research Employing Advanced Methodologies

The majority of current SOCMINT research focuses on applying sophisticated analytical techniques to specific detection challenges rather than explicitly addressing hybrid threat coordination. Ellaky et al. [9] report developing a hybrid deep learning architecture for social media bot detection based on BiGRU-LSTM and GloVe word embedding, achieving impressive performance metrics in identifying automated accounts. While bot detection is a crucial capability for identifying coordinated inauthentic behavior—a component of many hybrid influence operations—the research does not explicitly connect bot activity to broader patterns of hybrid threat campaigns.

Sangher et al. [43] describe LSTM and BERT-based transformer models for cyber threat intelligence, focusing on intent identification by analyzing exploitation of social media platforms from darknet forums. Their work demonstrates sophisticated natural language processing capabilities for understanding threat actor communications. Still, it operates within the bounded context of darknet intelligence rather than correlating these signals with surface-web influence operations or cyber-physical activities characteristic of hybrid campaigns.

Biagio et al. [44] present the MARPLE framework for social media threat intelligence, targeting terrorist threat detection through advanced analytics. Dragos et al. [14] undertake a comparative analysis of AI approaches for social media analysis, examining whether hybrid AI architectures suit hybrid threat detection—though their focus remains on technical methodology rather than operational hybrid threat scenarios. Bimyrzakyzy and Alimzhanova [45] employ social network analysis methods to identify cyber threats, while Arora et al. [46] use chatbots to predict cyber threats via sentiment analysis of social media. These studies showcase technical sophistication in specialized analytical tasks but do not explicitly address how their capabilities integrate into comprehensive hybrid threat detection frameworks.

4.3. Commercial SOCMINT Platforms

Commercial platforms demonstrate operational SOCMINT capabilities at scale, though with limitations similar to those observed in academic research.

Babel Street's multilingual monitoring across 200+ languages and real-time social media analysis from 30+ platforms provides comprehensive coverage of overt social media spaces [35]. The platform's sentiment analysis, network mapping, and collaborative workspaces support multi-analyst investigations of influence campaigns. However, these capabilities focus primarily on information collection and basic analytics rather than semantic reasoning about hybrid threat coordination across domains.

WebIQ's Voyager Suite [36], previously discussed in OSINT 3.1 Capabilities, extends SOCMINT capabilities to encrypted messaging platforms and darknet forums, providing law enforcement agencies with visibility across surface and dark-web social spaces. Here, its Visual Media Analytics component addresses multimedia content analysis—a capability gap identified in academic research—through AI-powered image search, audio transcription, and metadata extraction. The platform's real-time monitoring and automated alerting provide operational responsiveness. Nevertheless, these commercial solutions primarily aggregate and analyze social media data streams rather than systematically correlating social media intelligence with cyber technical indicators or physical threat assessments.

Specialized platforms focused on coordinated inauthentic behavior (CIB) represent the current analytical frontier for hybrid-threat Social Media Intelligence. Graphika [47], widely used in research collaborations with governments and major social-media platforms, excels at mapping influence-operation networks through advanced graph analysis, identifying coordination patterns associated with state-linked campaigns such as Russian IRA (Internet Research Agency) activity, Chinese state-aligned propaganda networks, and Iranian influence efforts. Its network-visualization and behavioral-pattern analysis support investigative-grade CIB identification.

Botometer [48] developed by Indiana University, provides research-standard bot-likelihood scoring through machine-learning analysis of account behavior, serving as a foundational tool for detecting automated amplification within influence campaigns. Together, these specialized platforms demonstrate that targeted network-analysis approaches deliver stronger performance for specific hybrid-threat components (CIB detection, bot identification) than general-purpose social-media monitoring tools. However, integration with cyber-technical and geospatial intelligence remains limited.

4.4. Critical Gaps in SOCMINT for Hybrid Threat Detection

Analysis of current SOCMINT research and operational capabilities reveals three fundamental limitations.

4.4.1. Conceptual Confusion between Hybrid Threats and Hybrid Methodologies

The literature demonstrates significant confusion between studies addressing genuine hybrid threats—coordinated campaigns spanning information, cyber, and physical domains—and those employing "hybrid" analytical methodologies for conventional cybersecurity problems. Only four studies [4,40–42] explicitly engage with hybrid threats as multi-domain security phenomena. The remaining research, despite achieving impressive technical sophistication in bot detection [9], darknet analysis [43], or extremist content identification [44], addresses isolated threat components rather than coordinated hybrid operations. This conceptual gap limits operational effectiveness: detecting bot networks or extremist content provides valuable intelligence, but understanding how these elements function as integrated components of broader hybrid campaigns requires analytical frameworks that current research does not provide.

4.4.2. Analytical Fragmentation across Threat Domains

Current SOCMINT research demonstrates strong performance within specific analytical niches but lacks frameworks for correlating threat intelligence across different social media contexts and platforms. Dover's critique [4] reinforces this limitation, noting that while SOCMINT provides horizon-scanning capabilities, it offers insufficient real-time warning capacity during hybrid conflicts. Sangher et al. [43] excel at analyzing darknet forums, Ellaky et al. [9] achieve sophisticated bot detection, yet no study provides systematic approaches for understanding how bot networks, extremist messaging, and darknet coordination function as integrated components of hybrid threat campaigns. Commercial platforms aggregate multi-platform data but leave synthesis to human analysts, creating

bottlenecks in threat assessment during time-sensitive scenarios.

4.4.3. Absence of Cross-Platform and Cross-Domain Intelligence Integration

The reviewed literature reveals minimal integration between SOCMINT and other intelligence disciplines, such as technical cyber intelligence or OSINT, beyond social media boundaries. Even studies explicitly addressing hybrid threats [41, 42] primarily operate within social media analytical boundaries rather than providing frameworks for systematically correlating social media intelligence with broader multi-domain threat indicators. Current approaches optimize specialized analytics within social media domains but cannot automatically connect social media influence campaigns with cyber infrastructure attacks, physical security incidents, or geospatial intelligence—the cross-domain correlation that defines effective hybrid threat detection. Commercial platforms like Babel Street and WebIQ, despite their sophisticated capabilities, similarly focus on social media monitoring rather than semantic integration with complementary intelligence sources.

The reviewed capabilities represent substantial progress in social media analysis and threat detection within bounded contexts. However, the fundamental challenge persists: hybrid threats orchestrate coordinated operations across multiple platforms and domains that current SOCMINT approaches analyze in isolation. Effective hybrid threat detection requires not merely more sophisticated social media analytics, but also semantic integration frameworks that enable systematic correlation between social media intelligence and cyber-physical threat indicators across the information environment that contemporary adversaries exploit strategically. Section 7 reveals this pattern systematically: while SOCMINT achieves moderate-to-strong coverage for network mapping (R2) and campaign detection within platforms (R3), critical gaps persist in cross-platform correlation (R6), semantic reasoning across domains (R7), and automated attribution (R8).

5. Natural Language Processing and Large Language Models

The integration of Natural Language Processing (NLP) with social media intelligence and open-source intelligence has undergone rapid evolution, transitioning from traditional text-mining approaches to sophisticated large-language model architectures capable of processing multimodal data streams in real time. This technological advancement has created new opportunities for automated threat detection, intelligence extraction, and information verification across diverse digital environments, from surface web social platforms to dark web marketplaces. Current research demonstrates significant progress in multilingual threat intelligence extraction, automated knowledge graph construction, natural language generation for intelligence summaries, and multimodal misinformation detection, leveraging transformer-based models, graph attention networks, retrieval-augmented generation, and agentic reasoning frameworks. This section evaluates NLP capabilities against requirements established in Section 2.4, with emphasis on toxic content detection (R5), semantic analysis beyond keywords (R7), multilingual capability (R11), and adversarial robustness (R10), including requirements R3, R6, and R8.

However, we acknowledge that a notable gap persists in the literature regarding specific detection and analysis of hybrid threats as coordinated campaigns. While current NLP research addresses isolated components of potential hybrid operations, such as misinformation detection, dark web monitoring, or social network analysis, few studies explicitly target the interconnected, multi-domain characteristics that define hybrid threat activities. This represents a critical research opportunity, as hybrid threats necessitate detection systems capable of simultaneously identifying patterns across different platforms, languages, and operational domains.

5.1. NLP Research Addressing Hybrid Threat Detection

Limited research explicitly frames NLP capabilities within hybrid threat contexts. Mlinac [40] examines how AI systems create hybrid threats through personalized disinformation and influence operations in cyberspace, focusing specifically on how natural language generation enables micro-targeted propaganda at scale. However, this work concentrates on how adversaries exploit NLP rather than how defenders can employ it for detection. Zapata Rozo et al. [5] develop an NLP-based framework to detect hostile social manipulation and extremist networks on social media platforms, employing entity extraction and relationship mapping to identify coordination patterns. While valuable, neither study addresses the fundamental challenge of detecting coordinated campaigns that span multiple domains simultaneously—combining disinformation, cyber operations, and physical activities as integrated threat

vectors characteristic of genuine hybrid operations.

5.2. Advanced NLP Capabilities for Threat Intelligence

Contemporary NLP research demonstrates sophisticated capabilities across specific intelligence domains. Wang et al. [49] introduce KnowCTI, a knowledge-based cyber threat intelligence system for entity and relation extraction that leverages cybersecurity ontologies to structure unstructured threat reports. The authors report that their approach enables systematic extraction of threat actor identities, techniques, and infrastructure from diverse intelligence sources, providing structured knowledge representations suitable for automated reasoning. This ontology-driven approach represents an advancement toward semantic integration, though KnowCTI focuses specifically on cyber threat intelligence rather than broader hybrid threat scenarios.

Perrina et al. [10] present AGIR, a natural language generation tool that converts STIX (Structured Threat Information Expression) graphs into human-readable cyber threat intelligence reports. The system demonstrates how NLP can automate intelligence production, transforming technical indicators into actionable narratives for diverse stakeholder audiences. However, AGIR operates within the bounded domain of cyber threat reporting, generating summaries from pre-structured data rather than performing cross-domain threat correlation or attribution.

Al-Yasiri et al. [50] propose an event extraction model that merges XLM-RoBERTa, BiGRU, and conditional random fields to process multilingual cyber threat intelligence feeds, addressing the challenge of processing threat intelligence across multiple languages. Their approach demonstrates the capability to extract structured threat events from unstructured text across diverse languages—a critical requirement given that hybrid threat operations frequently span linguistic boundaries. Pasupuleti [51] reports developing a hybrid NLP pipeline combining named entity recognition, sentiment analysis, and topic modeling to automate dark web threat intelligence, enabling systematic monitoring of underground forums for emerging threat discussions.

5.3. Misinformation Detection and Multimodal Analysis

Recent research addresses the challenge of detecting AI-generated misinformation and manipulated media. Chen and Shu [7] investigate the detection of LLM-generated misinformation, examining whether machine-generated propaganda exhibits distinguishable linguistic patterns from human-authored content. Their findings suggest that while current LLMs produce persuasive text, subtle distributional differences in language patterns may enable detection, though they caution that adversarial refinement of generation techniques could eliminate these signatures.

Huang et al. [52] conduct a comparative analysis of LLM-based misinformation detection strategies, evaluating approaches ranging from fine-tuned classifiers to zero-shot prompting of large language models for content verification. Their research reveals that while LLMs show promise for misinformation detection, performance varies significantly depending on prompt engineering, model architecture, and the sophistication of the misinformation. The authors note that detection systems must continuously evolve as generation techniques advance.

Qi et al. [53] present SNIFFER, a multimodal large language model for explainable out-of-context misinformation detection. SNIFFER analyzes both textual and visual content to identify instances of genuine media being repurposed with misleading context, a common disinformation tactic. The system's explainability component provides human-readable justifications for its assessments, addressing the transparency requirements critical for intelligence applications. However, SNIFFER focuses specifically on content authenticity rather than campaign-level coordination analysis.

Li et al. [54] introduce FactAgent, an LLM agent for fake news verification that decomposes verification tasks into tool-augmented subtasks. The system autonomously retrieves evidence, cross-references claims and synthesizes assessments through multi-step reasoning processes. This agentic approach demonstrates how LLMs can orchestrate complex analytical workflows. However, FactAgent primarily operates at the content level rather than identifying coordinated disinformation campaigns or correlating information operations with other hybrid threat indicators.

Marchiori et al. [55] examine LLM capabilities for Common Vulnerabilities and Exposures classification using hybrid Gemma-3 and MiniLM/XGBoost workflows, exploring whether large language models can automate vulnerability assessment. Their research indicates that while LLMs demonstrate promise for processing security advisories and extracting structured vulnerability information, current approaches require hybrid architectures combining neural and traditional machine learning components to achieve reliable performance.

5.4. Generative AI and Deepfake Implications

The proliferation of generative AI introduces dual challenges for hybrid threat intelligence. Current state-of-the-art generative models enable the creation of highly realistic synthetic text, images, audio, and video at minimal cost. ENISA's 2025 Threat Landscape report [6] indicates that AI-generated content now accounts for over 80% of observed social engineering attacks, representing a fundamental shift in the threat landscape. Adversaries can generate vast quantities of contextually appropriate propaganda across multiple languages, produce deepfake videos impersonating officials, or create synthetic personas with complete digital histories for influence operations.

The European regulatory response to generative AI threats has advanced through targeted governance frameworks. The General-Purpose AI Code of Practice, endorsed by the Commission and AI Board in July 2025, provides voluntary compliance mechanisms for GPAI model providers, addressing safety, transparency, and copyright obligations that impact both adversarial content generation and defensive detection systems [56]. The Commission's Guidelines on Prohibited AI Practices (February 2025) establish clear boundaries for AI applications in threat intelligence contexts, particularly concerning practices involving harmful manipulation, social scoring, and biometric identification that could arise in surveillance-oriented threat detection deployments [57]. These frameworks create dual implications for hybrid threat intelligence: they constrain specific detection methodologies (e.g., indiscriminate social media profiling, automated behavioral scoring without transparency) while establishing standards for responsible AI deployment that operational systems must navigate, reinforcing the R12 compliance requirement for systems operating under European regulatory frameworks.

Detection of AI-generated content remains an active research challenge. While some studies report identifying distributional patterns in machine-generated text or visual artifacts in synthetic media, adversaries continuously refine generation techniques to eliminate these signatures. The asymmetry between generation and detection capabilities, where producing convincing synthetic content requires less sophistication than reliably detecting it, creates persistent advantages for threat actors employing generative AI in hybrid operations.

5.5. Critical Gaps in NLP for Hybrid Threat Detection

5.5.1. Absence of Cross-Domain Reasoning Frameworks

While individual studies demonstrate sophisticated capabilities within bounded domains—KnowCTI for cyber threat intelligence [49], AGIR for report generation [10], SNIFFER for multimodal fact-checking [53], and FactAgent for news verification [54]—no research provides frameworks for reasoning across the diverse information types that hybrid threats exploit simultaneously. Current approaches optimize individual analytical components but cannot systematically correlate linguistic patterns in social media propaganda with technical cyber threat indicators from dark web forums and geospatial intelligence from news sources. Hybrid operations characteristically combine these information streams in coordinated campaigns, yet current NLP research processes each data type through separate analytical pipelines without semantic integration, enabling cross-domain pattern recognition.

5.5.2. Limited Operational Explainability and Attribution

While recent work emphasizes explainable AI, particularly Qi et al.'s SNIFFER system, which provides human-readable justifications, current approaches primarily explain content-level assessments rather than campaign-level attribution. Operational hybrid threat intelligence requires not only identifying that specific content constitutes misinformation but also determining whether multiple pieces of content across platforms represent coordinated campaigns, attributing campaigns to specific threat actors based on tactical patterns and predicting likely campaign evolution. Current NLP systems excel at analyzing individual documents or posts but lack frameworks for reasoning about relationships between content pieces, campaign coordination indicators, and threat actor tradecraft. The explainability gap extends to temporal reasoning: while systems can assess whether a single claim is false, they cannot explain why a narrative emerged at a specific time, how it relates to concurrent cyber operations, or what strategic objectives it serves, insights critical for hybrid threat assessment.

5.5.3. Fragmentation of Detection Capabilities across Threat Components

Contemporary research addresses isolated components of hybrid information operations through specialized systems: Chen and Shu [7] detect LLM-generated text, Qi et al. [53] identify out-of-context media manipulation, Al-

Yasiri et al. [50] extract events from threat feeds and monitor dark web discussions [51]. However, hybrid threat actors orchestrate these tactics simultaneously, deploying LLM-generated propaganda alongside manipulated imagery, coordinating narratives across surface and dark web platforms, and synchronizing information operations with cyber-attacks. No current NLP framework provides a unified analytical environment in which these detection capabilities operate cohesively, automatically correlating detections across components to reveal coordinated campaigns. Analysts must manually integrate outputs from disparate systems, creating analytical bottlenecks and missing opportunities for correlation that automated semantic reasoning could identify.

The reviewed NLP capabilities represent remarkable technical achievements in natural language understanding, generation, and multimodal analysis. Performance improvements are substantial, with recent systems achieving F1 scores exceeding 90% for entity extraction and significant accuracy gains across various intelligence applications. However, these advances remain fundamentally siloed within specific analytical tasks. Effective hybrid threat detection requires not merely more sophisticated NLP components, but also architectural frameworks that enable these components to reason collectively about threat patterns spanning linguistic, visual, technical, and geospatial information sources across the multi-domain battlespace that contemporary hybrid threats characteristically exploit. Section 7 reveals the following pattern systematically: NLP achieves moderate coverage for identification of messaging strategies (R4), toxic content detection (R5), and threat intelligence across multiple languages (R11), but reveals persistent gaps in cross-domain reasoning (R7), campaign attribution (R8), and adversarial robustness (R10) that semantic integration frameworks must address.

6. High-TRL European Initiatives

Understanding hybrid threat intelligence capabilities requires examining not only academic research and conceptual frameworks, but also operational systems that have achieved sufficient technological maturity to operate in real-world security environments. This section analyses high-TRL European initiatives that demonstrate both the progress toward integrated hybrid threat detection and the persistent challenges that operational deployment reveals. The Technology Readiness Level scale, ranging from TRL-1 (basic principles observed) to TRL-9 (actual system proven in the operational environment), provides a valuable framework for assessing the maturity of European capabilities from research prototypes to deployable solutions. Assessment of these operational systems against requirements established in Section 2.4 reveals how technical maturity translates into coverage of multi-source harvesting (R1), campaign detection (R3), strategic communications monitoring (R4), and especially toxic content detection (R5) and multilingual capability (R11), while exposing persistent gaps in cross-domain correlation (R6–R8).

6.1. High-TRL Operational Systems

The European Union's investment in hybrid threat detection has produced several systems approaching or achieving operational readiness, with varying degrees of integration and scope. The Narrative Analysis and Alerting System (NAAS) [13], developed by the Lithuanian Military Academy (Coordinator), Mykolas Romeris University (Partner), and the Lithuanian Cybercrime Center of Excellence for Training, Research & Education [58] (L3CE, Project Manager), represents one of the most mature platforms currently available, having achieved TRL-8 status with operational prototype demonstration in real security environments. NAAS addresses the information domain of hybrid threats through comprehensive narrative monitoring across diverse sources, including RSS feeds, websites, Telegram channels, YouTube, and document repositories. The platform's architecture employs containerized microservices orchestrated by Apache Kafka, enabling scalable, real-time processing of multilingual content with automated language detection, entity extraction, sentiment analysis, and semantic clustering to identify coordinated narrative campaigns.

The system's analytical capabilities extend beyond simple keyword monitoring to include moral foundations assessment and source credibility evaluation, producing trend analysis and emerging narrative alerts through an intuitive dashboard interface designed for law enforcement analysts. Critically, NAAS demonstrates that sophisticated narrative intelligence can be delivered through modern, scalable architectures deployable within secure agency networks, addressing operational requirements for both performance and security. The platform's progression to TRL-8 within the Lithuanian operational context, with a clear pathway to full deployment (TRL-9) expected within

1–2 years, illustrates the successful translation of research capabilities into practical tools. Lithuania's broader cybersecurity ecosystem supports this development: the National Cyber Security Centre's 2024 report [59] documented 63% annual growth in cyber incidents, with social engineering attacks accounting for 59% of incidents, underscoring operational demand for platforms like NAAS. The annual Kibernetinis skydas OpEx 2024 exercises [60], involving 80 organizations providing critical services, demonstrated both the maturity of Lithuanian coordination mechanisms and the persistent need for automated intelligence tools capable of operating at the speed and scale of contemporary hybrid threats.

HIPSTer, discussed in detail in Section 7, should be understood as a research-stage off-shoot of NAAS rather than an operational platform. Whereas NAAS has progressed to TRL-8 with demonstrated performance in real security environments, HIPSTer remains at TRL-4 and explores semantic-integration methods not yet implemented in high-TRL systems. Its role in this review is therefore architectural and illustrative: it highlights how ontological approaches could complement or extend the capabilities of mature systems such as NAAS, rather than serving as an operational alternative.

The STARLIGHT project, funded under the EU's Horizon 2020 program, represents a complementary approach focused on AI-enabled misinformation detection across multiple modalities [11]. STARLIGHT has developed sophisticated component capabilities, including bot-detection algorithms to identify coordinated inauthentic behavior on social platforms, clickbait detection using natural language processing and information divergence measures, spam detection with entropy-based anomaly identification, and author-attribution techniques to distinguish propaganda sources from authentic users. The project's Meta-Detection Engine integrates advanced neural architectures, including recurrent neural networks, gated recurrent units, long short-term memory networks, and BERT transformers, to generate comprehensive credibility assessments of content. A particularly notable capability is the Defalsify-ai module, which addresses synthetic media threats through deepfake video detection, fake image and face identification, symbol manipulation recognition, and image geolocation verification. However, STARLIGHT's components remain at TRL 4–6, representing validated subsystems that require integration into unified operational platforms. The project exemplifies both the sophistication of current AI-driven detection capabilities and the integration challenge: powerful analytical modules that function effectively in controlled environments must be unified into cohesive systems capable of 24/7 operational deployment in complex, dynamic threat environments.

France's VIGINUM service, established to combat foreign digital interference, provides operational lessons from real-world campaign analysis. VIGINUM's examination of the Portal Kombat pro-Russian propaganda network identified 193 interconnected websites operating across multiple European languages, employing sophisticated tactics including mass content replication with localized adaptation, search engine optimization to maximize visibility, audience segmentation targeting distinct demographics, and infrastructure obfuscation through distributed hosting [12]. The analysis, conducted in late 2023, revealed network characteristics that automated systems must detect: coordinated timing of content publication across sites, linguistic patterns indicating machine translation or template-based generation, common infrastructure markers despite apparent organizational separation, and strategic narrative seeding timed to amplify existing societal tensions. VIGINUM's methodology combined technical forensics, IP address tracing, graphic comparisons, metadata analysis, and content analysis to map network structure and assess impact. While the network's direct traffic remained relatively modest, its polarizing content was amplified through social media sharing and search engine visibility, demonstrating how influence operations exploit platform recommendation algorithms to extend reach beyond direct audiences. VIGINUM's countermeasure recommendations span technical monitoring (AI-enhanced detection, network analysis automation), regulatory approaches (transparency requirements, platform accountability), and societal resilience (media literacy, public awareness campaigns), reflecting the necessarily multi-faceted response hybrid threats demand.

6.2. EU Coordination Mechanisms and Institutional Progress

The European Union's institutional response to hybrid threats has progressed from conceptual frameworks toward operational coordination, though significant integration challenges persist. The EU-NATO Hybrid Fusion Cell, established to enhance shared situational awareness, has achieved staff-to-staff information sharing and joint strategic communication efforts, with synchronized crisis response exercises and stress tests demonstrating improved coordination [1]. The November 2025 BlueOLEx [61] marked the first exercise following the adoption of the new EU Cyber Blueprint, clarifying crisis roles and responsibilities, simulating large-scale cyber incidents impact-

ing critical sectors across multiple member states. The exercise tested coordination between the EU-CyCLONe network of member states' cyber liaison authorities and the Commission, revealing both progress in procedural clarity and persistent challenges in real-time information exchange across organizational boundaries. The EU's External Action Service has established FIMI (Foreign Information Manipulation and Interference) infrastructure for centralized monitoring and alert sharing, though operational details remain limited in public documentation [62,63].

The Commission's November 2025 European Democracy Shield initiative [64] operationalizes counter-FIMI capabilities through a European Centre for Democratic Resilience (viz. Digital Services Act) enforcement against transnational information operations, and whole-of-society resilience measures, including media literacy programs and protection of civil society from abusive lawsuits and physical threats. This institutional framework positions democratic resilience as requiring both technical detection capabilities and societal psychological defenses against manipulation.

France's VIGINUM service represents national-level operational capability that could inform EU-wide approaches, while Lithuania's integration of NAAS within its national cybersecurity architecture demonstrates how advanced analytical platforms can support operational coordination. However, current coordination remains primarily organizational rather than technical: information sharing occurs through secure communications channels and coordinated procedures, but agencies primarily operate separate analytical platforms without unified semantic frameworks enabling automated correlation.

National-level initiatives beyond France and Lithuania demonstrate diverse approaches to operational hybrid threat capabilities. Germany's Cyber Security Strategy for Germany 2021 emphasizes security-by-design principles across critical infrastructure, establishes measurable objectives for public-private cooperation in threat detection, and requires AI systems deployed in security contexts to integrate cybersecurity foundations from inception rather than retrofit [65]. This strategic framework positions cybersecurity as a joint responsibility spanning government, industry, and research communities – an institutional model relevant for deploying hybrid threat intelligence across sectoral boundaries and thereby addressing operational cyber resilience through design-phase security integration (R12). The Netherlands' Data Protection Authority (Autoriteit Persoonsgegevens) provides complementary institutional innovation through its semi-annual Report on AI & Algorithms Netherlands (RAN), establishing proactive algorithmic risk supervision and multi-stakeholder cooperation frameworks for AI Act enforcement [66]. The Dutch model's emphasis on transparency, algorithmic literacy, and collaborative supervision between data protection and AI regulatory authorities offers operational lessons for deploying compliant threat intelligence systems across European jurisdictions and, thus, supporting attribution transparency (R8) and regulatory compliance architecture (R12).

The EU-NATO cooperation initiative, while advancing, highlights the following gap: staff exchanges and joint exercises improve human coordination, but the absence of shared technical platforms with common threat ontologies limits the speed and depth of collaborative analysis that hybrid threats increasingly demand.

6.3. Policy Enabling Framework and Implementation Challenges

The European policy landscape provides increasingly sophisticated frameworks for addressing hybrid threats, though implementation complexity creates operational friction. The Eighth Progress Report on the 2016 Joint Framework on Countering Hybrid Threats documents escalating threat environments characterized by sophisticated FIMI operations, critical infrastructure targeting, and weaponized migration flows [1]. The report details progress in validating EU Hybrid Rapid Response Teams, strengthening institutional cybersecurity, and integrating AI security considerations into economic risk assessments, while acknowledging persistent capability gaps. The European Parliament's October 2025 motion on hybrid provocations strongly condemns Russian military and hybrid activities, including drone incursions. It recognizes the evolution of the hybrid threat into the physical and military domains beyond purely informational or cyber operations [2]. The Commission's European Democracy Shield initiative, presented in November 2025, sets forth concrete measures to protect democratic pillars, including free elections, independent media, vibrant civil society, and strong institutions [64].

Moreover, assessments made by the European Union Agency for Cybersecurity (ENISA) provide operational context for these policy frameworks. The 2025 Threat Landscape report documents a maturing threat environment with rapid exploitation of vulnerabilities and growing complexity in adversary tracking, noting that AI-enabled phishing campaigns reportedly accounted for over 80% of observed social engineering attacks by early 2025 [6].

The agency's NIS360 2024 report assesses cybersecurity maturity and criticality across NIS2 Directive sectors, identifies areas for improvement, and facilitates progress monitoring [67]. The 2024 Report on the State of Cybersecurity in the Union, ENISA's first comprehensive assessment developed with the NIS Cooperation Group, recommends strengthening support for NIS2 implementation and developing horizontal EU policy frameworks to address supply chain security challenges facing both the public and private sectors [68]. The Council of the European Union's May 2024 Conclusions on the Future of Cybersecurity emphasize effective, non-fragmented implementation of the NIS2 and Cyber Resilience Act, while calling for greater SME support and explicitly acknowledging the dual benefits and challenges of AI and quantum computing for security [69].

6.4. Regulatory Constraints on Operational Deployment

The deployment of hybrid threat detection systems within the European Union is guided by increasingly sophisticated regulatory frameworks that fundamentally shape operational architectures. A comprehensive exemplary regulatory analysis, including a detailed treatment of national implementation through Lithuanian case law, is presented extensively in Bružė et al. [70]; this section summarises the key architectural implications for hybrid threat intelligence systems. In this context, the applicable regulatory landscape comprises three interconnected layers, with the protection of fundamental rights at its core.

The first concerns personal data protection and limits on data use. The General Data Protection Regulation (GDPR) [71] imposes substantive obligations, including, among others, data minimization, purpose limitation, storage limitation, the requirement to identify an appropriate legal basis for processing additional safeguards for special categories of data, and the obligation to conduct data protection impact assessments for high-risk processing operations. Where systems are deployed by law enforcement authorities for law-enforcement purposes, the Law Enforcement Directive (LED) [72] governs processing instead, a distinction with significant practical implications, as the LED allows Member States to specify certain national rules while still requiring fundamental rights safeguards. This dual regime may require differentiated access controls and tailored compliance configurations depending on whether the same technology is operated by civilian entities or law enforcement authorities.

The second layer concerns AI-specific governance. The AI Act [73] classifies systems by risk level, prohibits certain AI practices, and imposes substantial compliance requirements on high-risk systems, including human oversight, fundamental rights impact assessments, and technical documentation. Hybrid threat detection systems, depending on their use cases and system configuration, may fall within prohibited or high-risk categories under the AI Act and therefore require careful legal and technical assessment before deployment. The General-Purpose AI Code of Practice [56] and Guidelines on AI Practices [57] establish further boundaries relevant to both adversarial content generation and defensive detection methodologies.

The third layer concerns cybersecurity and operational resilience. For example, the NIS2 Directive [74] reinforces secure-by-design architectures, incident detection and reporting, and operational resilience requirements that apply where threat intelligence systems support essential entities.

In practice, these layers rarely apply in isolation, and design choices made to satisfy one layer often have direct implications for the others. For example, the GDPR and AI Act may apply simultaneously: lawful access to data is not sufficient if analytical models and their outputs must also meet AI Act governance requirements. This overlapping application creates internal tensions, as data minimisation under data protection law may conflict with the need for large, representative datasets under AI governance requirements. Furthermore, as the EU legislation usually includes exemptions, especially as related to systems used for national security purposes, the use cases must be carefully analysed to ensure the system is put in the exact regulatory dimension and does not fall between the gaps.

Recent jurisprudence of the Court of Justice of the European Union provides critical guidance for navigating these tensions. The SCHUFA judgment (C-634/21) [75] established that automated scoring systems constitute decisions based solely on automated processing under GDPR Article 22, requiring stricter transparency and human oversight when significantly affecting security assessments. The Dun & Bradstreet judgment (C-203/22) [76] clarifies that systems must provide meaningful information about the logic involved in automated assessments and that controllers cannot categorically invoke trade secret protections to deny transparency. These rulings reinforce that explainability is a legal prerequisite for operational AI-driven threat intelligence, directly supporting the R8 attribution requirement's emphasis on transparent, auditable methodologies.

Additional regulatory guidance further shapes these constraints. The European Data Protection Board's 2025

guidance on secure AI systems emphasizes data minimization, purpose limitation, and proportionality, requiring technical measures such as pseudonymization, automated deletion of identifiers once analytical utility is exhausted, and strict access controls for sensitive data. ENISA's Multilayer Framework for Good Cybersecurity Practices for AI [77] complements these obligations by outlining cybersecurity-by-design principles, adversarial-resilience requirements, and sector-specific controls that must be embedded into analytical pipelines. Together, these instruments reinforce that hybrid threat detection systems must integrate privacy, security, and governance safeguards directly into their architectures rather than treating compliance as a post-deployment add-on.

The regulatory environment creates clear trade-offs. Strong privacy protections and transparency requirements can constrain analytical depth and operational tempo, while purpose limitation can restrict data reuse for emerging threats. However, compliance is not insurmountable: well-designed systems can achieve operational effectiveness within regulatory frameworks through privacy-by-design methodologies, explainable reasoning architectures, and human-in-the-loop decision processes, approaches that enhance rather than merely constrain analytical capabilities. The compliance burden requires interdisciplinary expertise spanning legal, technical, and operational domains, reinforcing the R12 requirement for systems designed with regulatory compliance as an architectural foundation rather than a post-deployment obligation.

6.5. Emerging Research: The HIPSTer Ontological Framework

Against this backdrop of advancing capabilities and persistent gaps, research initiatives are exploring ontology-based approaches to hybrid threat detection that address identified integration challenges. L3CE's research contribution within the HIPSTer (Hybrid Information Psychological Societal Threats handling system) project focuses on developing semantic reasoning frameworks that enable automated correlation across the disparate data sources and analytical domains that hybrid threats exploit. The research, funded by various European Union and Lithuanian national programs, addresses a critical gap across operational systems: the absence of widely accepted frameworks for identifying and attributing hybrid threats that coordinate information operations, cyber-attacks, and physical activities into unified campaigns. L3CE/Security Research Laboratory's research scope is reflected in its publicly accessible HIPSTer GitHub repository, which extends to TRL 4 (technology validated in a laboratory environment). Strategically, the work focuses on ontological foundations, semantic-reasoning mechanisms, and knowledge-representation architectures capable of integrating OSINT, SOCMINT, NLP and technical cyber indicators within coherent analytical frameworks [78].

The HIPSTer approach distinguishes itself through structured semantic labeling of threat indicators across multiple dimensions including source credibility assessment, content type classification, information veracity evaluation, intent determination, and specific influence technique identification. This ontological structure enables reasoning that current platforms lack: rather than merely aggregating data and presenting visualizations, the framework can automatically assess whether observed patterns constitute coordinated campaigns, attribute activities to known threat actor tradecraft, and predict likely campaign evolution based on historical patterns encoded in the knowledge base. The research addresses cross-domain integration explicitly, providing mechanisms to correlate social media influence operations with cyber-technical indicators and geospatial intelligence, the multi-domain correlation that defines effective hybrid threat detection.

6.6. Persistent Integration Gaps Revealed by High-TRL Systems

Analysis of operational European initiatives reveals a consistent pattern: sophisticated capabilities within bounded domains paired with persistent challenges in cross-domain integration and semantic reasoning. NAAS demonstrates mature narrative-intelligence capabilities at TRL-8 but operates primarily within the information domain, without native integration of cyber-technical or physical threat indicators. STARLIGHT showcases impressive AI-driven detection for specific misinformation tactics, yet its component modules remain separate systems requiring integration rather than a unified analytical platform. Commercial tools provide powerful data aggregation and visualization, but leave semantic synthesis and threat attribution to human analysts, creating bottlenecks when information volume and velocity exceed human cognitive capacity. Coordination mechanisms have advanced procedurally through exercises and information-sharing protocols, yet remain fundamentally manual rather than technically integrated through shared semantic platforms.

The regulatory dimension adds further complexity. Frameworks such as the AI Act and GDPR provide nec-

essary governance, but their interaction creates implementation challenges that operational systems must navigate. Organizations deploying hybrid-threat detection capabilities face competing pressures: advancing analytical sophistication to keep pace with evolving threats while ensuring compliance with privacy, transparency, and accountability requirements. The solutions emerging from this tension will likely combine advanced AI capabilities with strong governance frameworks, privacy-preserving techniques, and explainable reasoning mechanisms. The progression from research prototypes (TRL 4–6) through operational demonstrations (TRL 7–8) toward fully deployed systems (TRL-9) across European initiatives suggests maturing capabilities. Yet the fundamental challenge persists: hybrid threats exploit the seams between analytical domains, and current approaches, however sophisticated within their respective niches, continue to analyze these domains largely in isolation.

This integration challenge provides the context for examining ontological approaches as potential solutions to the identified gaps. Bružė et al. [70] quantify this pattern, and the 12 requirements, derived substantially from documented Russian and Chinese military information-operations tradecraft, further reinforce the argument. High-TRL systems demonstrate moderate to strong coverage of domain-specific requirements (R1, R3, R4, R11), as reflected in empirical validations such as Kaunas University of Technology's toxic-content detection work on the adversarial DynaHate benchmark. While these results show operational viability for specific detection tasks (R5) and multilingual capability (R11) across English, Lithuanian, Russian, and Chinese, cross-domain integration requirements (R6, R7, R8) remain at prototype or research stages. This persistent gap underscores the need for semantic frameworks such as HIPSTer to bridge operational silos and support integrated hybrid-threat detection.

7. Ontological Integration and the Role of the HIPSTer Framework

The integration challenges identified across OSINT, SOCMINT, and NLP domains suggest fundamental architectural limitations in current approaches rather than merely incremental capability gaps. Hybrid-threat activity deliberately exploits connectivity between analytical disciplines, requiring defensive systems capable of reasoning across heterogeneous data types, threat indicators, behavioral patterns, and operational contexts. Addressing these challenges demands frameworks that support formal semantic structures rather than relying solely on statistical correlation or platform-specific analytics.

Ontology-based approaches, leveraging semantic-web technologies such as the Web Ontology Language (OWL) and the Resource Description Framework (RDF), provide architectural foundations for such integration by enabling machine-readable knowledge graphs that support automated inference, relationship traversal, and multi-domain reasoning. This approach has been extensively employed over the last five years in a series of papers on the architectural design of next-generation networks [79–81], which demonstrated how formal semantic representations can model complex systems involving assets, threats, vulnerabilities, controls, and compliance obligations [82]. These earlier studies established the theoretical basis for cross-domain threat modelling and showed that semantic-web technologies enable structured reasoning aligned with established cybersecurity standards.

The ontological approach adopted here is also grounded in a formal risk management methodology. Our prior work implemented the ISO 27005 risk-assessment model and ENISA's Threat Landscape taxonomy as a semantic ontology, demonstrating that the ISO 27005 structure, where threat agents exploit vulnerabilities to create risks affecting assets, maps naturally onto OWL/RDF representations [80,82]. This standards-aligned modeling supports automated reasoning, SHACL constraint validation, and machine-readable risk propagation, providing a rigorous foundation for hybrid-threat intelligence architectures.

It is important to clarify that the HIPSTer (Hybrid Information Psychological Societal Threats handling system) framework is not only an offshoot of NAAS discussed earlier in Section 6.1, but it has also been implemented on Stardog, which is both a graph database engine and a full semantic reasoning platform. While Stardog stores data as a knowledge graph [83], hence the initial "Create Database" step, it differs fundamentally from conventional property-graph systems in that it provides native support for OWL reasoning, SHACL validation, SPARQL query capability, SKOS taxonomies, and ontology-driven inference. During earlier ontology engineering work [79–81,83], multiple graph database platforms were evaluated, including European RDF engines. These systems offered efficient graph storage but lacked the integrated reasoning and validation capabilities required for standards-aligned modeling based on ISO 27005 and ENISA's methodology. Stardog was selected because it supports the full semantic-web technology stack needed for iterative ontology development, including repeated SHACL validation cycles and reasoning tests using HermiT, ELK, and Pellet within Protégé. This systematic evaluation provides the methodological

grounding for the semantic-integration approach adopted in HIPSTer.

The HIPSTer framework represents the evolution of this architecture for hybrid-threat contexts, incorporating Simple Knowledge Organization System (SKOS) taxonomies that formalise the cross-domain relationships and threat indicators identified in Sections 3–6. **Figure 1** introduces the core ontological structure underlying HIPSTer, illustrating how semantic classes and relationships enable automated reasoning across domains. This section demonstrates how semantic frameworks address the requirements for cross-domain integration (R6), semantic analysis beyond keywords (R7), and attribution under ambiguity (R8), while also contributing to targeted campaign detection (R3), toxic content identification (R5), and temporal discrimination (R9), which were revealed by persistent gaps across OSINT, SOCMINT, NLP, and high-TRL operational systems.

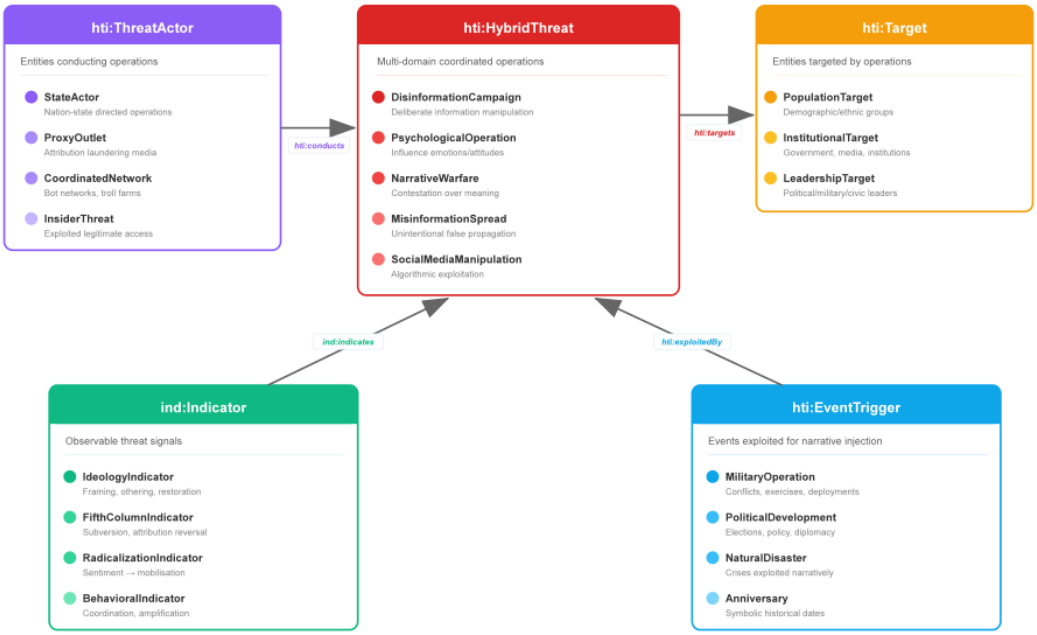

Figure 1. A schematic depicting the hybrid-threat ontological framework.

Note: Five principal classes, ThreatActor, HybridThreat, Target, Indicator, and EventTrigger, are linked through directional semantic relationships (hti:conducts, hti:targets, ind:indicates, hti:exploitedBy) that enable automated cross-domain reasoning from observable threat signals to campaign attribution. Subclasses capture the operational granularity required for hybrid threat characterisation, from state actors and coordinated networks through specific threat types to population, institutional, and leadership targets.

7.1. Architectural Overview

The HIPSTer framework applies ontological principles to hybrid threat contexts. The architecture combines formal OWL threat modeling with SKOS taxonomies that encode the operational vocabulary required for real-world threat characterization. As shown in **Figure 1**, the ontology's analytical foundation lies in its representation of hybrid-threat phenomena as interconnected semantic classes.

Where the OWL ontology provides the structural skeleton for threat reasoning, the complementary SKOS taxonomy populates this structure with the operational vocabulary required for real-world analysis. **Table 1** introduces HIPSTer's twelve functional capability domains, which form the basis of the SKOS taxonomy.

Table 1. HIPSTer's 12 primary capability domains are contained in the SKOS taxonomy.

Domain 1	Domain 2
Communication Metadata	Visual and Aesthetic Cues
Language and Phrasing	Cultural References and Symbolic Associations
Symbolism and Imagery	Social Media Metadata and Geolocation
Narrative and Storytelling	Content Structure and Framing Strategies
Community and Engagement Patterns	Dehumanization and Demonization
Behavioral Patterns and Activity Trends	Anomalies and Psychological Patterns

These domains are further expanded in **Figure 2**, which illustrates how more than 200 individual threat indicators are organized into a hierarchical classification spanning the full spectrum of hybrid-threat tradecraft, from communication metadata and linguistic patterns to behavioral signatures and psychological markers.

Figure 2. HIPSTer SKOS taxonomy: Twelve functional domains for hybrid-threat analysis.

Note: The taxonomy organises over 200 threat indicators into hierarchical domains spanning communication metadata, linguistic patterns, narrative techniques, behavioural signatures, and psychological markers. Each domain maps to formal ontological classes, enabling query-based analytical operations including quantifiable actor profiling, coordination detection through shared behavioural patterns, and multi-indicator threat assessment across domains.

To illustrate how these semantic structures operate in practice, **Figure 3** presents HIPSTer's end-to-end analytical workflow, showing how multi-domain data sources are transformed into integrated hybrid-threat intelligence outputs.

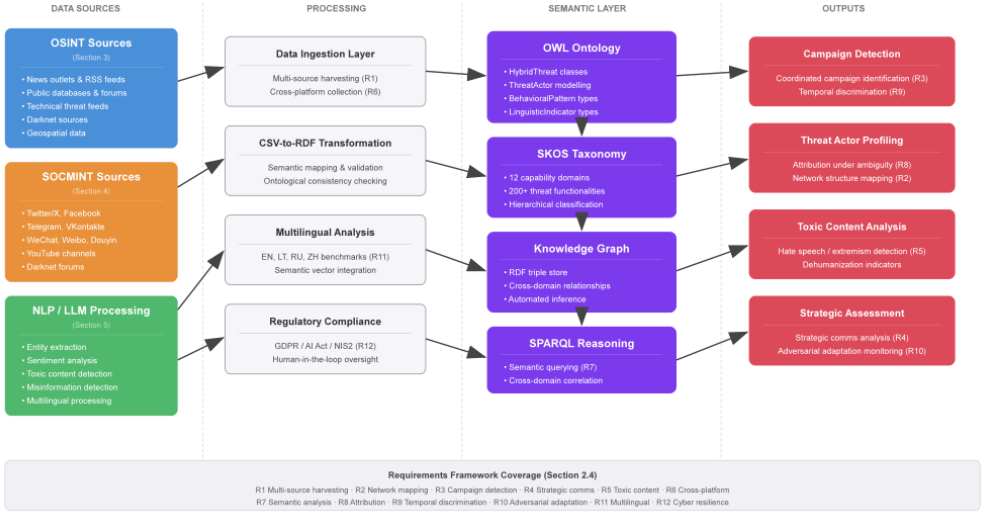

Figure 3. HIPSTer architectural workflow: Data flow from multi-domain sources to cross-domain threat intelligence.

Note: The pipeline progresses from OSINT, SOCMINT, and NLP data sources (left) through processing layers including data ingestion, CSV-to-RDF transformation, multilingual analysis, and regulatory compliance checking (centre-left) into the semantic integration layer comprising OWL ontology, SKOS taxonomy, knowledge graph, and SPARQL reasoning engine (centre-right-bottom), producing operational outputs including campaign detection, actor profiling, toxic-content analysis, and strategic assessment (right). Requirements framework references (R1–R12) indicate coverage points mapped to Section 2.4.

Together, **Figures 1–3** and **Table 1** demonstrate how HIPSTer operationalizes semantic integration across domains, addressing the architectural limitations identified in earlier sections. However, to systematically evaluate the framework's contribution, it is necessary to relate these semantic capabilities to the specific integration gaps and operational requirements identified in Sections 3–6.

Section 7.2, therefore, formalises this evaluation by using complementary analytical frameworks to assess both capability gaps and requirements coverage across the hybrid threat intelligence landscape.

7.2. Addressing Identified Integration Gaps

The preceding domain-specific reviews (Sections 3–5) and cross-domain integration analysis (Section 6) reveal systematic patterns in current hybrid threat intelligence capabilities: sophisticated analytics within bounded specializations, yet persistent challenges at the integration boundaries that sophisticated adversaries deliberately exploit. This section formalizes these observations through complementary analytical frameworks that assess both capability gaps and operational requirements coverage.

Two assessments structure this analysis.

First, a capability gap synthesis (**Table 2**) identifies critical limitations across OSINT, SOCMINT, and NLP/LLM domains, their operational consequences for detecting coordinated hybrid threats, and corresponding research priorities. This assessment reveals whether identified limitations represent performance deficits within existing system architectures (addressable through incremental improvement) or fundamental architectural constraints (requiring new paradigmatic approaches).

Table 2. Critical Capability Gaps and Research Priorities in Hybrid Threat Intelligence.

Domain	Critical Gap	Operational Consequence	Research Priority
OSINT (Section 3)	Domain-specific analytical silos	Cyber indicators, social media intelligence, and geospatial data were analyzed separately; missed campaign-level patterns	Semantic integration frameworks enabling automated cross-domain indicator correlation
OSINT	Limited contextual integration	Technical indicators (botnet activity, malware) are not systematically correlated with influence campaigns	Multi-modal reasoning architectures linking technical and informational threat indicators
OSINT	Absence of multi-modal reasoning	Tools process text, images, and signals independently; no unified analytical environment	Knowledge graph architectures supporting heterogeneous data type integration
SoCMINT (Section 4)	Conceptual confusion: hybrid threats vs. methodologies	Research addresses technical "hybrid" approaches without engaging multi-domain threat coordination	Formal threat taxonomies distinguishing genuine hybrid threats from hybrid analytical methods
SoCMINT	Analytical fragmentation across platforms	Bot detection, extremist content identification, and darknet monitoring operate independently	Cross-platform correlation mechanisms detecting coordinated campaigns across the surface/dark web
SoCMINT	Limited cross-domain integration	Social media intelligence is not systematically linked to cyber-technical or physical threat indicators	Semantic frameworks bridging SoCMINT with complementary intelligence disciplines
NLP/LLM (Section 5)	Absence of cross-domain reasoning	Content analysis, entity extraction, and sentiment assessment operate on isolated data streams	Ontological frameworks enabling semantic reasoning across diverse information types
NLP/LLM	Limited operational explainability	AI systems detect patterns but cannot explain campaign-level attribution or predict evolution	Explainable AI architectures providing human-interpretable reasoning chains for attribution
NLP/LLM	Fragmented detection capabilities	LLM-generated text detection, deepfake identification, and misinformation classification operate separately	Unified detection platforms integrating complementary AI capabilities with semantic oversight
Cross-Domain (Section 6)	Manual coordination mechanisms	Information sharing occurs through reports and communications rather than shared technical platforms	Collaborative intelligence platforms with common semantic frameworks, enabling automated correlation

Second, a requirements coverage evaluation (**Table 3**) systematically assesses how current capabilities across analytical domains address the operational requirements established in Section 2.4, requirements derived from documented Russian and Chinese hybrid threat tradecraft that European defensive systems must counter. This coverage assessment quantifies maturity levels across detection-focused requirements (multi-source harvesting, network mapping, toxic content identification) and operation-focused requirements (cross-platform integration, semantic reasoning, adversarial robustness, multilingual capability), revealing where capabilities are operational-ready and where significant development remains.

Table 3. Requirements Coverage across Hybrid Threat Detection Domains.

Requirement	OSINT (Section 3)	SoCMINT (Section 4)	NLP/LLM (Section 5)	HIPSTer Integration (Section 6–7)	Critical Gaps/Notes
R1: Multi-source data harvesting	*** Maltego, Babel Street, WebIQ Voyager provide comprehensive collection across platforms/forums	** Platform APIs enable data collection; limited darknet access	* Text-focused; limited multimedia harvesting	** Knowledge graph ingestion pipelines (CSV-to-RDF transformation)	Data provenance tracking standards; automated darknet monitoring
R2: Network structure mapping	*** IBM i2 Analyst's Notebook (social network analysis specialization) & Maltego (network mapping and entity visualization)	*** Graphika, Botometer provide advanced CIB detection and network analysis	* Entity extraction supports network construction	* Semantic network reasoning via ontological relationships	Cross-platform entity resolution; real-time network dynamics
R3: Targeted campaign detection	** Narrative tracking tools identify coordinated themes	*** CIB detection identifies coordinated inauthentic behavior campaigns	** Pattern recognition identifies narrative coordination	*** Ontological campaign modeling via HybridThreat class hierarchy	Automated attribution of campaigns to threat actors
R4: Strategic communications analysis	*** Media monitoring tools (multiple platforms) track state-sponsored outlets	** State media tracking across platforms	** Discourse analysis identifies messaging strategies	* Media-social correlation in early development	Real-time cross-source integration; propaganda technique classification
R5: Toxic content detection	* Limited automated toxic content identification	** Platform moderation tools; content policy enforcement	** detects hostile social manipulation and extremist networks on social media platforms	*** Kaunas University of Technology (KTU) validation benchmark	Multimodal toxicity (memes, images); contextual intent classification
R6: Cross-platform capability	** Manual aggregation across tools; no unified environment	* Platform-specific APIs and tools; limited cross-platform correlation	* Fragmented deployment across platforms	** Unified semantic framework via OWL/RDF knowledge graph; semantic reasoning; advanced query capabilities	Automated cross-platform entity resolution; real-time correlation
R7: Semantic analysis beyond keywords	* Predominantly keyword-based search; limited context understanding	* Basic sentiment analysis; limited semantic reasoning	* Semantic analysis gap in cross-domain reasoning	*** Sentence transformers, contextual understanding (KTU validation)	Cultural context understanding; coded language interpretation; sarcasm/irony detection
R8: Attribution under ambiguity	* Manual analyst attribution; limited automation	* Pattern-based actor profiling; requires human interpretation	* Limited campaign attribution capabilities	** TTP profiling via SKOS taxonomy (200+ functionalities); behavioral pattern matching	Automated attribution confidence scoring; multi-tier operation detection; false-flag identification
R9: Temporal discrimination	** Timeline visualization tools; historical analysis	** Trend analysis; temporal pattern detection	* Static models; limited temporal reasoning	* Campaign evolution modeling in development	Predictive analytics: distinguishing positioning vs. activation phases
R10: Adversarial adaptation anticipation	x Major gap; tools vulnerable to evasion	x Limited adversarial robustness; bot detection easily evaded	x adversarial adaptation requiring new detection paradigms	** DynaHate validation demonstrates tempered adversarial robustness	Adversarial training protocols; continual learning frameworks; red-team testing
R11: Multilingual capability	** Variable tool support; depends on platform	** Platform-dependent; major languages covered	** challenge of processing threat intelligence across multiple languages	*** KTU multilingual validation (Russian, Lithuanian, English); sentence transformers support 100+ languages	Low-resource language support; dialect handling; code-switching detection
R12: Operational cyber resilience	* Infrastructure-dependent; varies by deployment	* Platform security measures; limited hardening	* Model poisoning risks; adversarial example vulnerability	* Security considerations in the early design phase	Adversarial defences; secure model deployment; infrastructure hardening; fail-safe mechanisms

Note: Coverage ratings reflect both individual system capabilities and domain-wide maturity: *** (strong coverage, operationally validated), ** (adequate coverage, research-validated), * (partial coverage, prototype-stage), x (significant gap, research needed).

Together, these assessments provide empirical grounding for evaluating whether semantic integration frameworks, such as HIPSTer's ontological approach combining OWL threat modeling, SKOS taxonomy structures, and RDF knowledge graphs, represent incremental enhancements to existing systems or architectural prerequisites for effective hybrid threat detection. The analysis demonstrates that while individual analytical domains achieve moderate-to-strong maturity within their specializations, cross-domain integration requirements consistently show prototype-stage coverage or significant gaps, precisely the capabilities that ontological semantic reasoning was designed to address.

The following subsections present these assessments systematically: Section 7.3 synthesizes capability gaps and requirements coverage through the two tables. In contrast, Section 8.1 translates these findings into actionable research priorities for advancing European defensive capabilities.

7.3. Gap Synthesis and Research Priorities

Critical capability gaps identified across analytical domains, their operational consequences, and the corresponding research priorities are summarized in **Table 2**. The synthesis reveals a consistent pattern: current approaches optimize specialized analytics within domain boundaries but lack semantic integration frameworks that enable systematic correlation of threat indicators across the information, cyber, and physical domains that contemporary hybrid threats strategically exploit.

The research priorities emerging from a delineation of gaps in **Table 2** converge on common themes: semantic integration, multi-modal reasoning, cross-domain correlation, explainable attribution, and collaborative platforms with unified knowledge representation. These priorities align precisely with capabilities that ontological approaches – exemplified by HIPSTer's framework combining OWL threat modeling, SKOS taxonomy structures, and RDF knowledge graphs – were designed to provide. However, gap identification establishes research needs; requirements coverage assessment quantifies current capability maturity against operational demands.

Thus, **Table 3** complements this gap analysis by systematically evaluating coverage of operational requirements (Section 2.4) across analytical domains, quantifying the maturation gap between specialized capabilities and cross-domain operational needs that research priorities address in **Table 2**.

Indeed, the table systematically assesses requirements coverage across OSINT, SOCMINT, NLP/LLM, and HIPSTer integration, demonstrating that while individual capabilities achieve moderate-to-strong maturity within bounded scopes, systematic cross-domain correlation remains at prototype or research stages. The assessment reveals capability maturation patterns: detection-focused requirements (R1–R5) achieve adequate-to-strong coverage within specialized domains, multilingual capability (R11) demonstrates consistent strength across NLP and integration approaches, yet operation-focused requirements demanding cross-domain reasoning (R6–R8) and adversarial adaptation (R10) show persistent weaknesses. HIPSTer's ontological framework provides distinctively strong coverage of semantic integration requirements (R6, R7), where traditional approaches show gaps, though operational maturity (TRL-4) is required before deployment impact.

7.4. How Ontological Approaches Address Integration Gaps

The preceding gap analysis reveals that while individual analytical domains achieve moderate-to-strong maturity within their specialisations, cross-domain integration requirements consistently show prototype-stage coverage or significant gaps. Ontological frameworks address these architectural limitations through three fundamental capabilities.

First, unified semantic structures enable cross-domain correlation by integrating threat indicators spanning information operations, cyber-technical activities, and behavioural signatures within coherent knowledge representations. Rather than requiring manual analyst correlation between separate tools, explicit semantic relationships link behavioural patterns to content instances through actor entities, enabling automated detection of coordinated campaigns spanning multiple domains.

Second, formal ontological definitions provide contextual reasoning that surpasses keyword-based or purely statistical approaches. Observing linguistic indicators alone does not trigger threat classification, but co-occurrence with coordination patterns matching known threat actor tradecraft enables automated inference, the contextual assessment that **Tables 2** and **3** identify as persistently lacking across current platforms.

Third, multi-modal reasoning mechanisms represent diverse information types, technical cyber indicators, social media content, geospatial intelligence, linguistic features, as semantic entities within unified knowledge graphs, enabling queries that traverse cross-domain relationships to detect coordinated hybrid operations.

HIPSTer's current research scope extends to TRL-4 (technology validated in laboratory environment), with development and commercialization responsibility beyond TRL-4 resting with its enterprise partner, Novian (formerly ELSIS PRO), a Lithuanian small-to-medium enterprise partner responsible for advancing the research framework toward operational deployment [84]. While TRL-4 represents research validation rather than operational deployment, the framework demonstrates that formal ontological approaches can operationalize the semantic reasoning capabilities that current tools lack, establishing a viable architectural pathway from the integration gaps identified in this review toward the cross-domain correlation that effective hybrid threat detection demands.

Limitations of the Framework

Several limitations of the HIPSTer framework warrant acknowledgement. First, scalability beyond laboratory validation remains undemonstrated: the current TRL-4 implementation has been tested against curated datasets, and performance under operational data volumes, where real-time ingestion from multiple platforms generates orders-of-magnitude greater throughput, requires engineering maturation that falls within Novian's development remit. Second, the framework's analytical quality depends directly on the completeness and accuracy of the underlying knowledge graph; incomplete threat actor profiles, outdated taxonomic classifications, or gaps in the SKOS vocabulary would degrade reasoning outputs, necessitating continuous curation by domain experts. Third, while multilingual validation across English, Lithuanian, Russian, and Chinese benchmarks demonstrates cross-linguistic capability, coverage of other European target languages (German, French, Polish, Baltic languages beyond Lithuanian) remains untested. Fourth, operational deployment within European regulatory frameworks will require formal compliance assessment against GDPR data protection requirements and AI Act governance obligations – a process whose complexity Section 6.4 documents in detail. These limitations reinforce HIPSTer's current positioning as a research-validated architectural pathway rather than an operational capability.

7.5. Scope, Limitations, and Critical Assessment of the Landscape

The preceding sections have examined ontological integration from architectural, operational, and semantic perspectives. To contextualize these findings within the broader landscape of hybrid-threat intelligence, this subsection outlines the scope, assumptions, and methodological boundaries of the present review. Hybrid-threat intelligence spans heterogeneous disciplines and operational contexts, and any review of the field must balance breadth with analytical depth. The scope of this review is therefore intentionally delimited in several ways.

First, the 12-point operational requirements framework in Section 2.4 reflects documented Russian and Chinese information operations tradecraft, which establishes the upper boundary of adversarial sophistication relevant to European defensive systems. These requirements were derived from recurring operational patterns, cross-platform fragmentation, multilingual coordination, temporal activation, and adversarial adaptation, and not from the capabilities of any specific system. They serve as an evaluative structure for comparing OSINT, SOCMINT, and NLP approaches rather than as a prescriptive taxonomy.

Second, the gap-analysis ratings in **Tables 2** and **3** are qualitative assessments grounded in the scoping review methodology described in Section 1.1. They synthesise evidence from peer-reviewed studies, operational system documentation, and high-TRL European initiatives. The ratings are not intended as definitive performance scores but as comparative indicators of maturity across domains, highlighting where capabilities remain at the prototype stage or are siloed.

Third, while ontology-based integration approaches are examined in Section 7, they are not presented as the only viable architectural solution. Alternative integration paradigms, including data lake architectures, graph database systems, federated learning, and emerging LLM-based cross-domain correlation, are recognised as active research areas. Ontologies are used here as an illustrative case because they directly address semantic-integration gaps identified in Sections 3–6 and because a companion technical publication provides empirical grounding for the HIPSTer framework [70]. The review does not claim ontologies are universally superior; rather, they exemplify one class of solutions aligned with the requirements framework.

Fourth, the review focuses on European capabilities, regulatory frameworks, and operational deployments. This regional emphasis reflects the EU's distinctive legal environment (GDPR, AI Act, NIS2) and the availability of high-TRL systems such as NAAS, STARLIGHT, and VIGINUM. Relevant work from other regions, including the United States, Israel, and Asia, was examined using the search strategies defined in the PRISMA ScR methodology (Section 1.1). However, the search returned several thousand potential entries, making a comprehensive global survey beyond the scope of this review.

Fifth, the literature base necessarily emphasizes recent publications (2023–2025) due to the rapid evolution of hybrid-threat methodologies and AI-enabled detection systems. Work on information warfare and influence operations is incorporated [21,23,27], providing essential conceptual grounding, but the review prioritizes contemporary evidence reflecting current adversarial capabilities.

Finally, several limitations should be noted. Economic considerations, particularly the feasibility of seman-

tic integration architectures for smaller states or organisations, remain underexplored in the available literature and warrant further study. Similarly, while Section 2 highlights adversarial adaptation, the accelerating role of AI-generated content (deepfakes, synthetic personas, LLM-generated propaganda) will further expand the threat landscape and require continued monitoring beyond the semantic ontological domain to consider the costs associated with the handling and analysis of vast amounts of data, as well as the costs involved in secure operations, including staffing, expertise, and organisational resources.

Validation of integrated systems also remains an open challenge: demonstrating superiority over siloed tools requires controlled experiments, shared benchmarks, and access to operational datasets that are often restricted. These limitations do not diminish the value of the present synthesis but highlight areas where further empirical research is needed.

8. Future Research Directions

This review has examined hybrid threat intelligence capabilities across three critical domains, Open-Source Intelligence, Social Media Intelligence, and Natural Language Processing with Large Language Models, revealing both substantial technical progress and persistent integration challenges. Analysis of high-TRL European initiatives demonstrates that operational systems achieve sophisticated performance within bounded domains while struggling with the cross-domain correlation that hybrid threats fundamentally demand. Assessment against the requirements framework established in Section 2.4 reveals a consistent pattern: individual domains achieve moderate-to-strong coverage for specialized capabilities, yet cross-domain integration requirements (R6–R8) remain systematically underdeveloped across research, commercial, and operational systems. The following synthesis identifies priority research directions that account for technical requirements, operational constraints, and regulatory imperatives shaping the European threat intelligence landscape.

8.1. Technology Readiness Level Positioning and Development Pathway

L3CE's research contribution focuses on ontological foundations, semantic reasoning mechanisms, and knowledge representation architectures, with research scope extending to TRL-4 (technology validated in a laboratory environment). This research scope encompasses ontology design validated through proof-of-concept implementations, SKOS taxonomy development covering identified functionality domains, semantic query development demonstrating analytical capabilities, and integration methodology establishing transformation pipelines from operational data formats to RDF representations. Development and commercialization responsibility beyond TRL-4 rests with Novian, as previously mentioned [84]. The development pathway targets progression to TRL-7 (system prototype demonstration in operational environment) through engineering maturation, user interface development for operational analysts, integration with existing law enforcement workflows, and pilot deployments with partner security agencies. The eventual pathway to TRL-9 (actual system proven in operational environment) extends 1–3 years following research completion, contingent on successful operational validation and user acceptance, positioning HIPSTer as bridging the gap between research capabilities (TRL 4–6) and operational needs revealed by high-TRL systems (TRL 7–9).

8.2. Prioritized Research Agenda

Based on identified gaps and European operational requirements, the following research priorities emerge in order of criticality and feasibility:

Priority 1—Semantic Integration Frameworks (1–3 years):

Developing ontology-based architectures that enable automated correlation of threat indicators across the OSINT, SOCMINT, and technical cyber intelligence domains is a foundational requirement. Research should advance formal knowledge representation methods, semantic reasoning algorithms for multi-domain pattern detection, and validation methodologies demonstrating improved detection accuracy over isolated analytical approaches. The HIPSTer ontological framework represents one approach, but broader research is needed to explore alternative semantic technologies, scalability optimization for operational data volumes, and interoperability standards enabling ecosystem development.

Priority 2—Explainable Attribution Mechanisms (2–4 years):

Operational threat intelligence requires not merely detecting coordinated campaigns but attributing them to specific threat actors with sufficient confidence to support decision-making. Research should develop attribution frameworks that combine semantic reasoning with machine learning, explainability architectures that provide human-interpretable evidence chains to support attribution assessments, and methodologies for quantifying confidence that enable analysts to assess attribution reliability. This research must address the inherent ambiguity in adversarial attribution while providing actionable intelligence under uncertainty.

Priority 3—Cross-Platform Coordination Detection (1–3 years):

Current SOCMINT approaches analyze individual platforms effectively but struggle to handle threat actors coordinating campaigns across Twitter, Telegram, VKontakte, darknet forums, and emerging platforms simultaneously. Research priorities include behavioral signature development identifying coordination patterns across heterogeneous platforms, temporal correlation algorithms detecting synchronized campaign activation, and privacy-preserving analysis methods enabling coordination detection without excessive personal data collection.

Priority 4—Regulatory-Compliant AI Architectures (immediate-ongoing):

Given European regulatory frameworks, research must advance AI architectures achieving high analytical performance while satisfying GDPR data protection requirements and AI Act governance obligations, cybersecurity obligations, and fundamental-rights safeguards. This includes privacy-by-design methodologies, differential privacy techniques that enable aggregate analysis without exposing individual data, and transparency mechanisms that satisfy explainability requirements without compromising operational security. Future work should also develop compliance-by-design methodologies and automated regulatory-verification tools, ensuring that hybrid threat detection systems integrate legal, technical, and operational safeguards from the earliest design stages.

Priority 5—Operational Validation and Human-AI Teaming (3–5 years):

Moving research prototypes to operational deployment requires validation methodologies assessing real-world performance, human-AI interaction research optimizing analyst workflows, and training frameworks ensuring operators can effectively utilize advanced systems. This includes developing evaluation metrics capturing operational utility beyond technical accuracy, conducting field trials with law enforcement and intelligence agencies, and iterative refinement based on operational feedback.

8.3. Regulatory Constraints and Compliance Imperatives

Regulatory constraints increasingly shape the design and deployment of hybrid-threat detection systems within the European Union, particularly as security-oriented AI capabilities intersect with stringent data-protection and algorithmic-governance requirements. The interaction between the GDPR, LED, the AI Act, and NIS2 creates a compliance environment in which technical innovation and regulatory alignment must be pursued simultaneously rather than sequentially. For hybrid-threat intelligence systems that process publicly available information containing personal data, monitor behavioural patterns, or generate automated assessments, these frameworks impose operational boundaries that directly influence system architecture, data-handling practices, and analytical workflows.

Within this landscape, the central challenge is not merely adhering to individual regulatory provisions but ensuring that cross-domain analytical systems remain compliant as they integrate OSINT, SOCMINT, NLP, and technical cyber indicators. High-risk classifications under the AI Act introduce requirements for transparency, human oversight, bias mitigation, and post-market monitoring, while GDPR principles such as data minimisation, purpose limitation, lawfulness, and proportionality constrain how threat-intelligence data can be collected, transformed, and retained. These constraints are particularly salient for systems designed to detect coordinated inauthentic behaviour, toxic content, or cross-platform influence operations, where analytical value often depends on correlating heterogeneous signals that may include personal or sensitive attributes.

The implication for future research is clear: compliance must become a design parameter rather than a post-hoc adjustment. Hybrid-threat detection systems will require architectures that embed regulatory safeguards into their semantic models, data-processing pipelines, and reasoning mechanisms. This includes developing compliance-by-design methodologies, automated verification tools capable of validating regulatory adherence across multilingual and multi-modal datasets, and reference architectures that demonstrate how high-performance analytical systems can operate within European legal constraints. Achieving this will require interdisciplinary teams that com-

bine legal expertise, technical capability, and operational understanding, a skillset combination that remains limited across most security organisations.

These considerations directly inform the research priorities identified in Section 8.2 and shape the trajectory of next-generation hybrid-threat intelligence systems. As regulatory frameworks continue to evolve, the ability to integrate analytical performance with demonstrable compliance will become a defining capability for both academic research and operational deployments. The following conclusions synthesize these findings and outline the broader implications for the development of hybrid-threat intelligence systems.

9. Conclusions

Hybrid threats represent adaptive, evolving security challenges that will continue to demand equally adaptive defensive responses. This review documents substantial European progress toward integrated threat intelligence capabilities, with high-TRL initiatives such as NAAS and STARLIGHT, and national platforms including VIGINUM, demonstrating the operational viability of advanced detection approaches. Yet the fundamental challenge persists: hybrid threats exploit coordination across domains that current systems still tend to analyse in isolation. Addressing this gap requires not only incremental improvements to existing tools but architectural innovation that enables semantic reasoning about multi-domain threat patterns.

While ontological approaches have historically remained at early research stages, the work-in-progress HIPSTer framework, developed through the L3CE/Security Research Laboratory (MRU) collaboration, demonstrates that these models can now reach TRL 4 (laboratory validation). By integrating high-efficiency semantic vectors with formal reasoning, this initiative has operationalised the detection of adversarial Chinese and Russian threat indicators with high accuracy and minimal computational overhead, as documented in the Data Availability Statement..

The regulatory dimension adds complexity but also drives innovation toward privacy-preserving, transparent, and accountable AI systems, capabilities that extend well beyond narrow compliance requirements. The research agenda outlined above provides a roadmap for advancing these capabilities. However, meaningful progress will depend on sustained investment, operational collaboration between research and practitioner communities, and policy frameworks that balance security imperatives with the protection of fundamental rights.

Ultimately, the hybrid-threat challenge is a systems-integration problem. Europe possesses sophisticated analytical components, but only through frameworks such as HIPSTer are we beginning to provide the semantic glue needed to bind them into coherent, cross-domain intelligence capabilities. Solving this integration challenge represents the central research frontier for European hybrid-threat intelligence.

Funding

This research was prepared as part of the project "Implementation of Mission-Based Science and Innovation Programs" (Project No. 02-002-P-0001), coordinated by the Security Research Laboratory, Mykolas Romeris University. The project is financed by the European Union through the Recovery and Resilience Facility (2021–2027 programming period, New Generation Lithuania plan), and co-funded by the state budget of the Republic of Lithuania administered via the Research Council of Lithuania (Lietuvos Mokslo Taryba). These combined mechanisms ensure alignment with both EU-level resilience and recovery objectives and Lithuania's national science and innovation priorities. The funding received for this project is restricted to supporting research and writing activities; it does not cover publication or printing fees.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The technical specifications of the HIPSTer ontological framework, including the illustrative query and multilingual validation results echoed in Section 6.6 and used to support the review's findings, are available as open-access research artifacts at https://github.com/SecOntologyLab/hipster-ontology. This resource is a contribution of the Security Research Laboratory, Mykolas Romeris University, and serves as a conceptual extension of the 5G hybrid threats preliminary ontology [79–81] developed by L3CE, available at https://purl.org/5g-hybrid-threats. Demonstration results for the multilingual toxic content detection methodologies for the DynaHate benchmark alluded to in Section 6.6 are based on the Kaunas University of Technology (KTU) hate speech detection pipeline, with technical details available at https://github.com/evavaic/KTU-Misijos-HIPSTer. No new primary empirical data were generated in this review article; all synthesized frameworks are derived from publicly available institutional reports and peer-reviewed literature.

Acknowledgments

The authors thank Edmundas Piesarskas (L3CE) for the comprehensive competitive analysis of OSINT tools and expert-based evaluation methodology that informed the SKOS taxonomy structure and HIPSTer's functional requirements framework. His contributions provided the empirical basis for the metadata schemas and the operational requirements framework documented in the data and docs folders of the companion GitHub repository [78].

Conflicts of Interest

The authors declare no conflict of interest.

References

1. European Commission. Eighth progress report on the implementation of the 2016 Joint framework on countering hybrid threats and the 2018 Joint communication on increasing resilience and bolstering capabilities to address hybrid threats. Available online: https://defence-industry-space.ec.europa.eu/system/files/2025-01/SWD_Annual-Progress-Report-2024.PDF (accessed on 26 November 2025).
2. European Parliament. Motion for a Resolution. Available online: https://www.europarl.europa.eu/doceo/document/B-10-2025-0437_EN.pdf (accessed on 26 November 2025).
3. Ivkova, V.S.; Opirskyi, I.R. Research of Existing OSINT Tools and Approaches in the Context of Personal and State Information Security. *Comput. Syst. Netw.* **2025**, *7*, 143–159. [CrossRef]
4. Dover, R. SOCMINT: A Shifting Balance of Opportunity. *Intell. Natl. Secur.* **2020**, *35*, 216–232. [CrossRef]
5. Zapata Rozo, A.; Díaz-López, D.; Pastor-Galindo, J.; et al. An NLP-Based Framework to Spot Extremist Networks in Social Media. *Complexity* **2024**, *2024*, 3380488. [CrossRef]
6. ENISA Threat Landscape 2025. Available online: https://www.enisa.europa.eu/sites/default/files/2025-10/ENISA%20Threat%20Landscape%202025_0.pdf (accessed on 26 November 2025).
7. Chen, C.; Shu, K. Can LLM-Generated Misinformation Be Detected? *arXiv preprint* **2024**, *arXiv.2309.13788*. [CrossRef]
8. Yadav, A.; Kumar, A.; Singh, V. Open-Source Intelligence: A Comprehensive Review of the Current State, Applications and Future Perspectives in Cyber Security. *Artif. Intell. Rev.* **2023**, *56*, 12407–12438. [CrossRef]
9. Ellaky, Z.; Benabbou, F.; Matrane, Y.; et al. A Hybrid Deep Learning Architecture for Social Media Bots Detection Based on BiGRU-LSTM and GloVe Word Embedding. *IEEE Access* **2024**, *12*, 100278–100294. [CrossRef]
10. Perrina, F.; Marchiori, F.; Conti, M.; et al. AGIR: Automating Cyber Threat Intelligence Reporting with Natural Language Generation. *arXiv preprint* **2023**, *arXiv.2310.02655*. [CrossRef]
11. European Commission. Sustainable Autonomy and Resilience for LEAs Using AI against High Priority Threats. Available online: https://cordis.europa.eu/project/id/101021797 (accessed on 26 November 2025).
12. Ministry of the Armed Forces. Disinformation, a Weapon of War. Available online: https://www.defense.gouv.fr/en/news/disinformation-weapon-war (accessed on 26 November 2025).
13. Mykolas Romeris University. Creation of Information Security and Information Threats' Detection, Analysis, Research and Education Ecosystem (NAAS). Available online: https://www.mruni.eu/en/creation-of-information-security-and-information-threats-detection-analysis-research-and-education-ecosystem-naas-nr-01-2-1-lvpa-v-835-03-000-nr-01-2-1-lvpa-v-835-03/ (accessed on 9 December 2025).
14. Dragos, V.; Forrester, B.; Rein, K. Is Hybrid AI Suited for Hybrid Threats? Insights from Social Media Analysis.

In Proceedings of the 2020 IEEE 23rd International Conference on Information Fusion, Rustenburg, South Africa, July 2020; pp. 1–7. [CrossRef]
15. Tricco, A.C.; Lillie, E.; Zarin, W.; et al. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Ann. Intern. Med.* **2018**, *169*, 467–473. [CrossRef]
16. Borisov, I. Maskirovka—The Art of Deception à La Russe. *RMT* **2024**, 192–207. [CrossRef]
17. Kalensky, J.; Osadchuk, R. How Ukraine Fights Russian Disinformation: Beehive vs Mammoth. Available online: https://www.hybridcoe.fi/wp-content/uploads/2024/01/20240124-Hybrid-CoE-Research-Report-11-How-UKR-fights-RUS-disinfo-WEB.pdf (accessed on 25 December 2025).
18. Kling, J.; Toepfl, F.; Jürgens, P. Entertainment Interspersed with Propaganda: How Non-Legacy-News Accounts Deliver Explicitly Political Content to Mass Audiences on Russia's Most Popular Social Network VK. *Inf. Commun. Soc.* **2025**, *28*, 1252–1269. [CrossRef]
19. Kruglova, L.A.; Shchepilova, G.G. Russian TV Channels and Social Media in the Transformation of the Media Field. *Online Media Glob. Commun.* **2025**, *4*, 371–386. [CrossRef]
20. Eady, G.; Paskhalis, T.; Zilinsky, J.; et al. Exposure to the Russian Internet Research Agency Foreign Influence Campaign on Twitter in the 2016 US Election and Its Relationship to Attitudes and Voting Behavior. *Nat. Commun.* **2023**, *14*, 62. [CrossRef]
21. Cochran, E.S. China's "Three Warfares": People's Liberation Army Influence Operations. *Int. Bull. Polit. Psychol.* **2020**, *20*.
22. OSINT. Advanced OSINT for China: SOCMINT on WeChat, Weibo, and More. Available online: https://www.osint.industries/post/advanced-osint-for-china-socmint-on-wechat-weibo-and-more (accessed on 27 November 2025).
23. Charon, P.; Jeangène Vilmer, J.-B. *Chinese Influence Operations: A Machiavellian Moment*; Institute for Strategic Research (IRSEM): Paris, France, 2021.
24. Wong, C. The Diaspora and China's Foreign Influence Activities. Available online: https://www.wilsoncenter.org/publication/diaspora-and-chinas-foreign-influence-activities (accessed on 27 November 2025).
25. Finkelstein, D.; Yanovsky, S.; Zucker, J.; et al. Information Manipulation on TikTok and Its Relation to American Users' Beliefs about China. *Front. Soc. Psychol.* **2025**, *2*, 1497434. [CrossRef]
26. Zhang, Y. Who Gets the Algorithm? The Bigger TikTok Danger. Available online: https://www.lawfaremedia.org/article/who-gets-the-algorithm-the-bigger-tiktok-danger (accessed on 27 November 2025).
27. Silver, L.; Huang, C.; Clancy, L. *Across 19 Countries, More People See the U.S. Than China Favorably—But More See China's Influence Growing*; Pew Research Center: Washington, DC, USA, 2022.
28. Suryotrisongko, H.; Musashi, Y.; Tsuneda, A.; et al. Robust Botnet DGA Detection: Blending XAI and OSINT for Cyber Threat Intelligence Sharing. *IEEE Access* **2022**, *10*, 34613–34624. [CrossRef]
29. Yuan, X.; Wang, J.; Zhao, H.; et al. Empowering LLMs with Toolkits: An Open-Source Intelligence Acquisition Method. *Future Internet* **2024**, *16*, 461. [CrossRef]
30. Bizouarn, K.M.; Abdulnabi, M.; Tan, J. OSINT and AI: A Powerful Combination for Company Vulnerability Detection. In Proceedings of the 2023 IEEE 21st Student Conference on Research and Development, Kuala Lumpur, Malaysia, 13 December 2023; pp. 246–250. [CrossRef]
31. Fauziyyah, A.K.; Adrian, R.; Alam, S. Analyzing Image Malware with OSINTs after Steganography Using Symmetric Key Algorithm. *SinkrOn* **2023**, *8*, 818–824. [CrossRef]
32. Nonum, E.O.; Avwokuruaye, O.; Ezemonye, T.M. Role of Open Source Intelligence (OSINT) in Cybersecurity and Threat Analysis. *Int. J. Latest Technol. Eng. Manag. Appl. Sci.* **2025**, *14*, 189–200. [CrossRef]
33. Vacas, I.; Medeiros, I.; Neves, N. Detecting Network Threats Using OSINT Knowledge-Based IDS. In Proceedings of the 2018 14th European Dependable Computing Conference, Iasi, Romania, 10–14 September 2018; pp. 128–135.
34. Shin, S.-M.; Jung, K.-H. A Comparative Study of OSINT Automation Tools. *Asia-Pac. J. Converg. Res. Interchange* **2024**, *10*, 1–13. [CrossRef]
35. Lynch, T.L.; Sulzer, M.A. New Towers of Babel: A Conceptual Argument for Digital Platforms as Unstable Linguistic Constructs. In *Literacies in the Platform Society*; Routledge: London, UK, 2025; pp. 21–39.
36. Puleri, J. Law Enforcement and Open Source Intelligence: Evolution, Technologies, and Privacy Issues. PhD Thesis, Utica College, Utica, NY, USA, 2021.
37. Galis, V.; Karlsson, B. A World of Palantir–Ontological Politics in the Danish Police's POL-INTEL. *Inf. Commun. Soc.* **2024**, *27*, 2438–2456. [CrossRef]
38. Xiao, P.; Xie, L.; Hang, F.; et al. Advanced Technique for Firmware Security Analysis through Heterogeneous Data Fusion and Knowledge Mapping. *PLoS ONE* **2025**, *20*, e0319660. [CrossRef]

39. Maltego. Available online: https://www.maltego.com (accessed on 27 November 2025).
40. Mlinac, N. Hybrid Intelligence as a Carrier of Disinformation and Hybrid Threats in Cyberspace. *Natl. Secur. Future* **2025**, *26*, 65–98.
41. Cárdenas, P.; Obara, B.; Theodoropoulos, G.; et al. Analysing Social Media as a Hybrid Tool to Detect and Interpret Likely Radical Behavioural Traits for National Security. In Proceedings of the 2019 IEEE International Conference on Big Data (Big Data), Los Angeles, CA, USA, 9–12 December 2019; pp. 4579–4588.
42. Mothe, J.; Ullah, M.Z.; Okon, G.; et al. Instruments and Tools to Identify Radical Textual Content. *Information* **2022**, *13*, 193. [CrossRef]
43. Sangher, K.S.; Singh, A.; Pandey, H.M. LSTM and BERT Based Transformers Models for Cyber Threat Intelligence for Intent Identification of Social Media Platforms Exploitation from Darknet Forums. *Int. J. Inf. Technol.* **2024**, *16*, 5277–5292. [CrossRef]
44. Biagio, M.S.; Simoncini, S.; La Mattina, E.; et al. MARPLE: A Framework for Social Media Threat Intelligence. In Proceedings of the 2024 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA), Victoria, Seychelles, 1–2 February 2024; pp. 1–6.
45. Bimyrzakyzy, A.; Alimzhanova, Z.M. Identifying Cyberthreats through Social Media Research. *Bull. Shakarim Univ. Tech. Sci.* **2024**, *3*, 42–49. [CrossRef]
46. Arora, A.; Arora, A.; McIntyre, J. Developing Chatbots for Cyber Security: Assessing Threats through Sentiment Analysis on Social Media. *Sustainability* **2023**, *15*, 13178. [CrossRef]
47. Ronzaud, L.; Carter, J.A.; Williams, T. *Summit Old, Summit New: Russia Linked Actors Leverage New and Old Tactics in Influence Operations Targeting Online Conversations about NATO Summit*; Graphika: New York, NY, USA, 2023.
48. Yang, K.-C.; Varol, O.; Davis, C.A.; et al. Arming the Public with Artificial Intelligence to Counter Social Bots. *Hum. Behav. Emerg. Technol.* **2019**, *1*, 48–61. [CrossRef]
49. Wang, G.; Liu, P.; Huang, J.; et al. KnowCTI: Knowledge-Based Cyber Threat Intelligence Entity and Relation Extraction. *Comput. Secur.* **2024**, *141*, 103824. [CrossRef]
50. Al-Yasiri, J.H.; Bin Zolkipli, M.F.; Farid, N.F.N.M.; et al. A Threat Intelligence Event Extraction Conceptual Model for Cyber Threat Intelligence Feeds. In Proceedings of the 2024 7th International Conference on Internet Applications, Protocols, and Services (NETAPPS), Kuala Lumpur, Malaysia, 6 November 2024; pp. 1–8.
51. Pasupuleti, M.K. Threat Intelligence Automation Using Natural Language Processing on Dark Web Data. *Int. J. Acad. Ind. Res. Innov.* **2025**, *5*, 399–411.
52. Huang, T.; Yi, J.; Yu, P.; et al. Unmasking Digital Falsehoods: A Comparative Analysis of LLM-Based Misinformation Detection Strategies. In Proceedings of the IEEE 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE), Shanghai, China, 21–23 March 2025; pp. 2470–2476.
53. Qi, P.; Yan, Z.; Hsu, W.; et al. SNIFFER: Multimodal Large Language Model for Explainable Out-of-Context Misinformation Detection. arXiv preprint **2024**, arXiv.2403.03170. [CrossRef]
54. Li, X.; Zhang, Y.; Malthouse, E.C. Large Language Model Agent for Fake News Detection. *arXiv preprint* **2024**, *arXiv.2405.01593*. [CrossRef]
55. Marchiori, F.; Donadel, D.; Conti, M. Can LLMs Classify CVEs? Investigating LLMs Capabilities in Computing CVSS Vectors. *arXiv preprint* **2025**, *arXiv.2504.10713*. [CrossRef]
56. European Commission. The General-Purpose AI Code of Practice. Available online: https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai (accessed on 5 July 2025).
57. European Commission. Commission publishes the Guidelines on prohibited artificial intelligence (AI) practices, as defined by the AI Act. Available online: https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-prohibited-artificial-intelligence-ai-practices-defined-ai-act (accessed on 9 December 2025).
58. L3CE, Lithuanian Cybercrime Center of Excellence for Training, Research & Education. Available online: https://www.l3ce.eu/en/about-l3ce/ (accessed on 9 December 2025).
59. National Cybersecurity State Report 2024. Available online: https://www.nksc.lt/doc/Nacionaline-kibernetinio-saugumo-ataskaita-2024.pdf (accessed on 28 November 2025).
60. National Cyber Security Exercises "Cyber Shield Opex 2024". Available online: https://www.nksc.lt/doc/KS2024_OPEX_Pratybu_ataskaita.pdf (accessed on 28 November 2025).
61. European Commission. Member States and Commission test collective cybersecurity crisis response. Available online: https://digital-strategy.ec.europa.eu/en/news/member-states-and-commission-test-collective-cybersecurity-crisis-response (accessed on 28 November 2025).
62. EU Disinfo Lab. A practical toolkit for detecting, assessing, and responding to Foreign Information Manipu-

lation and Interference (FIMI). Available online: https://www.disinfo.eu/publications/a-practical-toolkit-for-detecting-assessing-and-responding-to-fimi/ (accessed on 25 November 2025).
63. EEAS. Information Integrity and Countering Foreign Information Manipulation & Interference (FIMI). Available online: https://www.eeas.europa.eu/eeas/information-integrity-and-countering-foreign-information-manipulation-interference-fimi_en (accessed on 1 December 2025).
64. European Commission. European Democracy Shield and EU Strategy for Civil Society Pave the Way for Stronger and More Resilient Democracies. Available online: https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2660 (accessed on 26 November 2025).
65. Federal Ministry of the Interior, Building and Community. Cyber Security Strategy for Germany 2021. Available online: https://www.bmi.bund.de/SharedDocs/downloads/EN/themen/it-digital-policy/cyber-security-strategy-for-germany2021.pdf (accessed on 26 November 2025).
66. AP (Dutch Data Protection Authority). AI & Algorithmic Risks: Developments in the Netherlands. Available online: https://www.autoriteitpersoonsgegevens.nl/en/themes/algorithms-ai/ai-algorithmic-risks-developments-in-the-netherlands (accessed on 26 November 2025).
67. ENISA. ENISA NIS360 2024 Report: A Comprehensive Look at Cybersecurity Maturity and Criticality of NIS2 Sectors. Available online: https://www.enisa.europa.eu/news/enisa-nis360-2024-report (accessed on 26 November 2025).
68. European Union Agency for Cybersecurity (ENISA). *2024 Report on the State of Cybersecurity in the Union*; European Union Agency for Cybersecurity: Heraklion, Greece, 2024.
69. Council of the European Union. Council Conclusions on the Future of Cybersecurity: Implement and Protect Together. Available online: https://data.consilium.europa.eu/doc/document/ST-10133-2024-INIT/en/pdf (accessed on 26 November 2025).
70. Bružė, E.; Paskauskas, R.A.; Matulytė, R.; et al. Integration of Hybrid Threat Intelligence: The HiPSTer Ontological Method for Cross-Domain Correlation in Influence Operations. *Open Res. Eur.* **2026**, in press.
71. European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council. Available online: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02016R0679-20160504 (accessed on 15 December 2025).
72. European Union. Directive (EU) 2016/680 of the European Parliament and of the Council. Available online: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016L0680 (accessed on 17 December 2025).
73. European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council. Available online: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng (accessed on 17 December 2025).
74. European Union. Directive (EU) 2022/2555 of the European Parliament and of the Council. Available online: https://eur-lex.europa.eu/eli/dir/2022/2555/oj/eng (accessed on 17 December 2025).
75. European Union. Case C-634/21, SCHUFA Holding (Scoring). Available online: https://eur-lex.europa.eu/eli/C/2024/913/oj/eng?utm_source=copilot.com (accessed on 17 December 2025).
76. European Union. CJEU, 27 February 2025, CK v Magistrat der Stadt Wien and Dun & Bradstreet Austria GmbH, Case C-203/22. Available online: https://www.julia-project.eu/database/case-law/319 (accessed on 17 December 2025).
77. ENISA. Multilayer Framework for Good Cybersecurity Practices for AI. Available online: https://www.enisa.europa.eu/publications/multilayer-framework-for-good-cybersecurity-practices-for-ai (accessed on 7 December 2025).
78. GitHub. Hybrid-Threat Intelligence: The HIPSTer Ontological Framework. Available online: https://github.com/SecOntologyLab/hipster-ontology (accessed on 1 December 2025).
79. Paskauskas, R.A. A Preliminary Ontology for 5G Network Resilience: Hybrid Threats, Risk Reduction, Compliance. In Proceedings of the 2025 IEEE International Conference on Cyber Security and Resilience (CSR), Chania, Greece, 4–6 August 2025; pp. 490–497.
80. Paskauskas, R.A. Countering Hybrid Threats: Towards an Ontology for Securing 5G Networks. In *Computer and Communication Engineering*; Neri, F., Du, K.-L., San-Blas, A.-A., et al., Eds.; Springer Nature Switzerland: Cham, Switzerland, 2025; 2192, pp. 104–121.
81. Paskauskas, R.A. Decoding 5G Security: Toward a Hybrid Threat Ontology. *Open Res. Eur.* **2025**, *4*, 34. [CrossRef]
82. Paskauskas, R.A. ENISA: 5G Design and Architecture of Global Mobile Networks; Threats, Risks, Vulnerabilities; Cybersecurity Considerations. *Open Res. Eur.* **2022**, *2*, 125. [CrossRef]

83. Yadav, N.; Gopinathan, D. Semantic Exploring and Analysis on Visualization of Research Articles Based on Knowledge Graphs. In Proceedings of the 2023 Second International Conference on Informatics (ICI), Noida, India, 23–25 November 2023.
84. Novian. Novian's Consolidated Revenue Increased 2.4% in 2024 to EUR 38.9 Million. Available online: https://novian.io/news/novians-consolidated-revenue-increased-2-4-in-2025-to-eur-38-9-million/ (accessed on 28 November 2025).