HARVARD LAW SCHOOL
NATIONAL SECURITY JOURNAL

ONLINE ESSAY

A New AI Strategy to Combat Domestic Terrorism and Violent Extremism

Jonathan Fischbach*

Recommended Citation

Jonathan Fischbach, *A New AI Strategy to Combat Domestic Terrorism and Violent Extremism*, HARV. NAT'L SEC. J. ONLINE (May 6, 2020), https://harvardnsj.org/wp-content/uploads/sites/13/2020/05/Fischbach_A-New-AI-Strategy.pdf.

* Attorney, U.S. Department of Defense. J.D., Cornell Law School, 2002; B.A., Princeton University, 1998. The positions expressed in this article do not necessarily reflect the views of any government agency or of the United States. This article was a winner of the Galileo Award in 2018, an annual competition sponsored by the Office of the Director of National Intelligence.

Table of Contents

INTRODUCTION: A REVEALING INVERSION

Data scientists utilize artificial intelligence (AI) in thousands of different contexts, ranging from analytics that design culinary masterpieces and identify illegal fishing, to algorithms that diagnose cancerous tumors, virtually compose symphonies, and predict vehicle failures.[1] Two communities within this expansive field, acting independently and without coordination, are currently experimenting with AI for the same narrow purpose—to determine whether machine-learning algorithms can discover patterns in demographic and behavioral data that identify actors likely to endanger innocent people. One group—the national security community—is tightly organized and well financed, operates at the federal level, attracts premier professional talent, and enjoys broad statutory access to data gathered in the United States. The other group—the social policy community—is decentralized and chronically underfunded, acts primarily at the state and local level, struggles to hire qualified data scientists, and lacks the authority and resources to establish access to data across localities. Yet the social policy community has already harnessed AI to achieve noteworthy policy outcomes, while the national security community still labors to understand its relevance and potential. Why?

The answer lies in the critical and underappreciated role of human beings in data collection. Intelligence agencies often acquire foreign intelligence as bits and snapshots of information that are bereft of context and rarely elicited directly by national security personnel. Social workers, however, use comprehensive intake tools to compile holistic histories and risk profiles of their subjects in order to optimize the delivery of appropriate services.[2] These case files are fertile terrain for machine-learning algorithms, which have unearthed striking and unforeseen causal relationships in numerous social service datasets and disciplines.

This track record suggests that the government cannot innovate its way to productive AI simply by developing more sophisticated software or by inventing better algorithms. To fully exploit machine technology, the national security community must reconceive its approach to data collection by reallocating resources from platforms that produce fragmentary collection to programs that leverage sources of data rich in biographic and narrative detail. But perhaps more significantly, this revelation underscores the vital importance of building bridges and establishing dialogue between professionals in the national security and domestic social policy realms. A natural first step would be a collaborative venture to determine how AI can help identify and ameliorate the conditions that induce domestic terrorism and violent extremism, emerging crises with both national security and social policy dimensions.

This article explores the hidden relevance of AI-enabled social welfare initiatives to national security programs. Part I recounts the efforts of one social welfare agency in Allegheny County, Pennsylvania to develop advanced analytics that predict violent acts against children before they occur. Part II surveys the use of AI in other social welfare programs to advance policies

---

[1] *See* Bernard Marr, *27 Incredible Examples of AI and Machine Learning in Practice*, FORBES (Apr. 30, 2018, 12:28 AM) https://www.forbes.com/sites/bernardmarr/2018/04/30/27-incredible-examples-of-ai-and-machine-learning-in-practice/#1ffff6347502 [https://perma.cc/83Q3-KFZ4].

[2] Paula Allen-Meares & Bruce A. Lane, *Social Work Practice: Integrating Qualitative and Quantitative Data Collection Techniques*, 35 SOC. WORK 452, 452–54 (1990).

that substantially overlap national security missions. Part III highlights critical distinctions between intelligence information and the social welfare data that drive AI successes in the social policy realm. Part IV proposes two strategies to overcome the inherent difficulties of using traditional intelligence collection to support AI capabilities. Part V envisions the fusion of social policy tradecraft with national security missions and resources to neutralize the threat of domestic terrorism and violent extremism.

I. A SOCIAL WELFARE TRIUMPH IN ALLEGHENY COUNTY

On June 30, 2011, firefighters rushed to the scene of a fire blazing from a third-floor apartment on East Pittsburgh-McKeesport Boulevard in Pittsburgh.[3] When the firefighters broke down the door, they found the body of seven-year-old KiDonn Pollard-Ford buried under a pile of clothes in his bedroom, where he sought refuge from the smoke. KiDonn's four-year-old brother KrisDon was lying under his bed, unconscious. He died in the hospital two days later. The Office of Children, Youth and Families (CYF) in Allegheny County, Pennsylvania previously received calls alleging that the children were neglected, but on each occasion, a case screener "screened the call out," or determined that the information did not warrant additional follow-up.

This tragedy was followed by other episodes of children dying in Allegheny County after CYF call screeners dismissed allegations of abuse or neglect. Desperate to halt this trend, CYF contacted two social scientists exploring the use of AI to improve the response of social welfare agencies to allegations of mistreatment. The scientists exhaustively combed through all 76,894 allegations CYF screened between April 2010 and April 2014. For each allegation, the team created an entry composed of every piece of information the county knew about the family, comprising more than one hundred different data fields populated by eight separate databases. Each entry reflected whether children were re-victimized after CYF screened a phone call implicating their family.

The results of the study were jarring: CYF call screeners were screening *in* 48% of the lowest-risk families for additional investigation, and screening *out* 27% of the highest-risk families. The social scientists used the database to develop an algorithm—the "Allegheny Family Screening Tool"—that generated a machine assessment of the risk posed by each family accused of abuse or neglect. CYF deployed the algorithm in August 2016. Thereafter, when a screener fielded an allegation and made a preliminary decision to screen the call in or out, the final step was to click on an icon for the Allegheny Family Screening Tool to produce the algorithm's risk assessment.

After using the algorithm for sixteen months, Allegheny County reported that CYF call screeners accurately identified high-risk calls more frequently. The percentage of low-risk cases flagged for investigation dropped from nearly half of all investigations to one third, and audits revealed that screeners treated black and white families more consistently. A second round of technical modifications increased the algorithm's success rate at predicting bad outcomes from 78% to 90%.

---

[3] The story recounted in this Part is adapted from Dan Hurley, *Can an Algorithm Tell When Kids Are in Danger?*, N.Y. TIMES MAG. (Jan. 2, 2018), https://www.nytimes.com/2018/01/02/magazine/can-an-algorithm-tell-when-kids-are-in-danger.html [https://perma.cc/UX5K-P2SD].

* * *

At 3:50 p.m. on November 30, 2016, CYF received a call from a preschool teacher. Moments earlier a three-year-old child informed the teacher that her mother's boyfriend had "hurt their head and was bleeding and shaking on the floor and the bathtub." Local media outlets reported that the boyfriend had overdosed and died in the home.

The call screener searched CYF's database for records about the family. The database showed numerous allegations about the family dating back to 2008—substance abuse, inadequate hygiene, domestic violence, and inadequate food and medical care—but not a single allegation was substantiated. Drug use without more did not satisfy the minimum legal requirements to have a caseworker visit the home. To screen the call out and administratively close the file, the CYF screener had to predict the risk of future harm to the child. He typed in "Low risk." Prompted to assess the threat to the child's immediate safety, he chose "No safety threat."

The final step was to click the icon for the Allegheny Family Screening Tool. Three seconds later, the computer screen displayed a numeric scale ranging from 1 (lowest risk) to 20 (highest risk). The score for the child's family was 19. The screener reversed his initial decision, screened the call in, and recommended that a caseworker visit the home. The caseworker's investigation revealed that the girl and her two older siblings were in immediate danger, and they were removed from the home. As of 2018, all three children were thriving in placements with other relatives.

II. PREDICTIVE ANALYTICS IN THE SOCIAL SERVICES REALM

Allegheny County's breakthrough in marshalling AI to revamp its child welfare system is echoed in social policy achievements reported by other localities and social service domains. In 2016, the Oklahoma Department of Human Services built a machine-learning algorithm that analyzed child welfare data statewide to identify the factors most likely to predict child fatalities.[4] The algorithm revealed that fifteen data points—many of which would not set off alarm bells—correlate highly with subsequent child fatalities, including the presence of a child under the age of three, a lover in the home, young parents, and substance abuse.[5] Machine algorithms also illuminate the circumstances that case screeners weigh too heavily in the intake process, including the screener's own background and experiences, the gender and ethnicity of the child and their family, and the ease of attributing blame for a reported incident to the child (termed "blame ideology").[6]

These initiatives have produced tangible results. In Iowa, social scientists developed a predictive analytic using data from 6,832 families reported to child welfare services, focusing on

---

[4] Kathleen Hickey, *Saving Children, One Algorithm at a Time*, GCN (July 26, 2016), https://gcn.com/Articles/2016/07/26/child-welfare-analytics.aspx [https://perma.cc/UG7M-RBX4].
[5] *Id.*
[6] *See* Philip Gillingham, *Predictive Risk Modelling to Prevent Child Maltreatment and Other Adverse Outcomes for Services Users: Inside the 'Black Box' of Machine Learning*, 46 BRIT. J. SOC. WORK 1044, 1049 (2016).

families that were either re-reported for alleged maltreatment or had a prior report substantiated.[7] Use of the analytic led case screeners to classify more families as low-risk, enabling the state to channel additional resources and supports to high-risk families.[8] In Broward County, Florida, social scientists used a predictive analytic to assess, for families reported to Florida's Office of Child Welfare, which social services and supports would most likely ensure that the families would not be reported a second time for alleged mistreatment.[9] The team projected that use of the algorithm could improve child welfare outcomes in Broward County by up to thirty percent through fewer inappropriate referrals and the enhanced use of "light touch services" in low-risk cases, such as periodic phone calls or visits, homemaker services, child care, and transportation.[10]

AI advances in the social policy realm are not limited to child welfare agencies. Leveraging existing administrative data, researchers at Cornell University, the University of Chicago, and Harvard University created a database of over one million bond court cases to develop an analytic that predicts whether defendants released on bail before trial will commit another criminal offense.[11] The scientists estimated that using the analytic to complement the intuition of criminal court judges would (1) reduce crimes committed by released defendants by up to 25%, without increasing the overall number of incarcerated individuals; and (2) lead judges to jail up to 42% fewer people without causing any increase in the crime rate.[12]

Public health experts in Illinois analyzed administrative data from the Illinois Department of Human Services to compile a dataset of 6,457 women who gave birth between July 2014 and May 2015.[13] Using this data the team developed an analytic to predict when a woman would experience an "adverse birth outcome," defined to include a preterm birth, low birth weight, death within the first year of life, or infant complications requiring admission to the neonatal intensive care unit.[14] The algorithm indicated that circumstances commonly assumed to increase the risk of adverse birth outcomes—such as mental illness, domestic violence, and prior incarceration—had little to no effect.[15] However, physical attributions and conditions, such as multiple pregnancies,

---

[7] Carol Coohey et al., *Actuarial Risk Assessment in Child Protective Services: Construction Methodology and Performance Criteria*, 35 CHILD. & YOUTH SERVS. REV. 151, 155 (2013).
[8] *See id*. at 160.
[9] *See* Ira M. Schwartz et al., *Predictive and Prescriptive Analytics, Machine Learning and Child Welfare Risk Assessment: The Broward County Experience*, 81 CHILD. & YOUTH SERVS. REV. 309, 317 (2017).
[10] *Id.* at 318–19.
[11] *See* Jon Kleinberg et al., *A Guide to Solving Social Problems with Machine Learning*, HARV. BUS. REV. (Dec. 8, 2016), https://hbr.org/2016/12/a-guide-to-solving-social-problems-with-machine-learning [https://perma.cc/QWY4-H93D].
[12] *Id.* In a similar study, researchers used machine learning algorithms to forecast the likelihood that a defendant arraigned on domestic violence charges would commit additional acts of domestic violence if released on bail. *See* Richard A. Berk et al., *Forecasting Domestic Violence: A Machine Learning Approach to Help Inform Arraignment Decisions*, 13 J. EMPIRICAL LEGAL STUD. 94, 94–95 (2016). The authors predict that use of the algorithm could cut the reoffending rate nearly in half over 24 months in the jurisdiction studied—that is, eliminate more than 2,000 post-arraignment arrests. *See id.* at 94.
[13] Ian Pan et al., *Machine Learning for Social Services: A Study of Prenatal Case Management in Illinois*, 107 AM. J. PUB. HEALTH 938, 938 (2017).
[14] *Id.* at 939.
[15] *Id.* at 941–42.

low pre-pregnancy weight, previous preterm birth, and maternal age of 40 years or older, had a surprisingly significant impact on health outcomes for newborns.[16]

III. CONTRASTING SOCIAL SERVICES DATA WITH FOREIGN INTELLIGENCE

A Mitre Corporation research team surveying the use of predictive analytics in child welfare was unequivocal in its assertion that "[d]ata is the most crucial part of a predictive analytics implementation; without useable data, no further work can be conducted."[17] Experts generally agree that predictive analytics require vast amounts of diverse data to expose patterns that inform the prediction of future events.[18] But data science initiatives in the Intelligence Community (IC) implicitly presume that any unit of information, aggregated in large quantities, can support the development of effective analytics. The reality is that foreign intelligence—the unit of information available to IC data scientists—is less useful to algorithms than are units of social welfare data because of the way foreign intelligence is collected. At root, this phenomenon is driven by discrepancies in *data context*. In three critical respects social-science data points acquire context through shared relationships to a common focal point—connective tissue that rarely binds foreign intelligence data points.

A. *Context Through Horizontal Integration of Discrete Service Systems*

An important factor driving the success of AI in the social policy realm is the ability of social workers to trace the status and behavior of individuals across service systems. The Allegheny Family Screening Tool, for example, can access a rich trove of administrative data including medical records, criminal history information, records reflecting disability status, educational records, mental health information, and receipt of government benefits.[19] The data pool for the Broward County child welfare study consolidated 238,912 records from 85 datasets governing a diverse range of triggers for government services or interventions, including sexual abuse, physical abuse, human trafficking, alcohol and drug abuse, malnutrition, environmental hazards, criminal history, homelessness, housing assistance, counseling, child care, juvenile court proceedings, mental health services, adoption, family planning, respite care services, unemployment services, and transportation services.[20] Indeed, the Broward County team observed that efforts to develop a predictive analytic for child welfare outcomes "demonstrate[] the value and importance of integrating data from a variety of essential child welfare data sets in order to improve child welfare assessment instruments and improve child and family outcomes."[21] Similarly, the Illinois study on adverse birth outcomes leveraged a database that tracks participation in all Illinois government programs, capturing data ranging from biographic

---

[16] *Id*.
[17] CHRISTOPHER TEIXEIRA & MATTHEW BOYAS, MITRE CORP., PREDICTIVE ANALYTICS IN CHILD WELFARE: AN ASSESSMENT OF CURRENT EFFORTS, CHALLENGES AND OPPORTUNITIES 12 (2017), https://aspe.hhs.gov/system/files/pdf/257841/PACWAnAssessmentCurrentEffortsChallengesOpportunities.pdf [https://perma.cc/DC9T-F5D8].
[18] *See* Amir Gandomi & Murtaza Haider, *Beyond the Hype: Big Data Concepts, Methods and Analytics*, 35 INT'L J. INFO. MGMT. 137, 138–40 (2015).
[19] *See* Hurley, *supra* note 3.
[20] *See* Schwartz et al., *supra* note 9, at 312–15.
[21] *Id.* at 319.

information like age, pre-pregnancy weight, birth history, and education level; to behavioral and health indicators such as HIV status, mental illness, tobacco and drug use, homelessness, and child services involvement.[22]

The context that emerges from horizontally integrated data is invaluable to a predictive analytic. It enables algorithms to assemble and analyze a person's mosaic of interwoven attributes, behaviors, and events—for thousands of unique individuals. The value of machine learning is to reveal relationships between and among these data points that may elude the observation of human beings;[23] horizontally integrated data most closely approximates the layered factual circumstances a predictive analytic would be asked to evaluate.

Foreign intelligence is not similarly contextualized through horizontal integration of data across agencies and domains. Indeed, cultural and technical barriers preclude the IC from establishing databases that consolidate all the government's information about intelligence targets.[24] IC elements also gravitate by necessity toward more superficial modes of data collection in response to operational needs, resource constraints, and regulatory frameworks. The IC collection model prioritizes intelligence with obvious and immediate relevance to a human analyst; it does not strive for comprehensive coverage of a target's pattern of life, or cast a wide net for contextual data peripheral to the primary intelligence objective. While social welfare agencies apply intake procedures that produce a holistic profile of their subject, IC elements conceive subjects as two-dimensional personas caricatured through the lens of a specific threat.

B. *Context Through Longitudinal Data*

Predictive analytics rely on diverse subject matter, but also on data points collected over time, which allow algorithms to sequence significant events into personal timelines that enable the study of cause and effect. For example, the machine-learning algorithm developed to assess whether criminal defendants should be eligible for pre-trial release relied on a data pool that merged information about at least three non-contemporaneous events for each defendant: (1) the defendant's initial arrest or indictment; (2) whether the defendant was held for trial or released on bail; and (3) whether a defendant released before trial committed another criminal offense.[25] Similarly, analytics used to predict child welfare outcomes are grounded not only in investigation

---

[22] *See* Pan et al., *supra* note 13, at 939–42.

[23] *See* Hurley, *supra* note 3 ("What the screeners have is a lot of data, . . . but it's quite difficult to navigate and know which factors are most important. . . . [T]he human brain is not deft at harnessing and making sense of all that data.").

[24] *See* Marie-Helen Maras, *Overcoming the Intelligence-Sharing Paradox: Improving Information Sharing Through Change in Organizational Culture*, 36 COMP. STRATEGY 187, 190 (2017) ("A general culture of secrecy exists in the intelligence community. . . . The perceived risk of inadvertent disclosure of information serves as a barrier to better cooperation and sharing of intelligence."); NAT'L RESEARCH COUNCIL, INTELLIGENCE ANALYSIS FOR TOMORROW: ADVANCES FROM THE BEHAVIORAL AND SOCIAL SCIENCES 8 (2011) ("Broadly speaking, the nation's confederated intelligence system has produced specialization at the expense of integration and collaboration. The IC's inability to function as a unified team has been the subject of more than 40 major studies since the CIA's establishment in 1947." (citation omitted)); *see also* RICHARD A. BEST JR., CONG. RESEARCH SERV., R41848, INTELLIGENCE INFORMATION: NEED-TO-KNOW VS. NEED-TO-SHARE (2011), https://fas.org/sgp/crs/intel/R41848.pdf [https://perma.cc/4T4L-ED3R] (documenting challenges to post-9/11 efforts to establish common repositories of information that merge key intelligence produced by all IC elements).

[25] *See* Kleinberg et al., *supra* note 11.

narratives and antecedent events archived in government databases, but also in post-allegation criminal, child welfare, and open-source data that map a family's trajectory after a reported incident.[26]

The IC may occasionally construct detailed chronologies for high-profile targets, but these efforts are the exception, not the rule. Foreign intelligence that lacks longitudinal context is not regarded as inherently insufficient or incomplete, and analysts routinely supply informed inferences to frame and contextualize the snapshots of activity revealed in intelligence collection. Thus intelligence reports about a terrorist or cyber hacker would generally eschew a lengthy treatment of the actor's childhood, and accentuate the present circumstances that create danger or suggest a threat-mitigation strategy—such as the individual's unique capabilities, geographic location, access to resources, and personal network. Consequently, the sample of foreign intelligence data points enriched by historical facts is miniscule compared to the volume of longitudinal data in the administrative datasets that shape predictive analytics in the social services realm.

C. *Context Through Documentation of Interventions*

Both social policy and national security officials administer a variety of interventions to disrupt potential threats. At the state and local level, social sector responses range from passive observation to engagement through counseling and social services, from support through government benefits and assistance to arrest and prosecution. Strategies to disrupt national security threats include intelligence gathering, criminal prosecution, covert action, or kinetic operations.

When the government at any level responds to individuals who pose a threat, there is obvious AI value in combining the subject's demographic and behavioral information with data that chart the nature, implementation, and aftermath of these interventions. Data scientists exploit the availability of this information in social services data to develop analytics that not only predict risk but also recommend interventions to neutralize the specific risk identified.[27] For example, the predictive analytic developed to improve child protective service investigations in Broward County revealed not only that certain low-risk cases were systematically over-referred to local agencies and the juvenile courts, but also that giving families "too much"—i.e., delivering social services that were unnecessarily intensive—was more harmful than not serving the family at all.[28]

There is no parallel movement within the national security community to consolidate and leverage data on efforts to disrupt national security threats. IC elements have limited visibility into the strategies or operations used by other agencies to combat threat actors of general concern,[29] and even within agencies there are no uniform standards for documenting an intervention or rating

---

[26] *See* TEIXEIRA & BOYAS, *supra* note 17, at 13.
[27] *See* Sun-Woo Choo, *Predictive Analytics, Recommender Systems, and Machine Learning: The Power of Data for Child Welfare*, ABT ASSOCIATES (May 14, 2018), https://www.abtassociates.com/insights/perspectives-blog/predictive-analytics-recommender-systems-and-machine-learning-the-power [https://perma.cc/7Y3S-346Z] ("[Machine learning] can be used in child welfare in two ways: to make predictions (e.g., about the scale of a child's endangerment), or to make recommendations (e.g., what services and referrals can be provided to parents to maximize child well-being).").
[28] *See* Schwartz et al., *supra* note 9, at 318.
[29] *See* sources cited *supra* note 24.

its effectiveness. IC elements may occasionally synchronize intelligence gathering with separate operational activity to reveal how a target was impacted by an intervention, or to capture any subsequent adjustment of an adversary's strategies, behaviors, and tradecraft. But these data samples are too small to support analytics that could recommend specific national security interventions for newly identified threats.

IV. APPLYING ROBUST ANALYTICS TO NATIONAL SECURITY PROBLEMS

IC elements have invested substantial resources in AI, and these initiatives have yielded some impressive results. The preceding discussion does not imply otherwise. But in the current threat landscape, AI will ultimately be measured by its ability (1) to predict whether newly encountered individuals pose a threat to our national security; (2) to reveal with precision the geographic, demographic, political, and socio-economic conditions likely to foment threats against the United States; and (3) to prescribe the specific strategies and interventions best suited to disrupt identified threats. Data scientists have proven that contextualized data enable machine-learning algorithms to supply these answers in the social policy realm. The question is whether these successes translate to national security challenges.

Obviously, the government cannot collect contextualized foreign intelligence through the overt and transparent intake procedures that social service organizations use to gather information. However, two creative strategies could posture the national security community to achieve the same AI breakthroughs that data scientists have delivered for social welfare agencies.

A. *Leveraging Domestic Administrative Data to Decode Foreign Threat Activity*

Algorithms mining administrative datasets have correlated demographic and behavioral data with event outcomes to produce non-intuitive insights. Significantly, the same contextualized data compiled by domestic social service agencies could revolutionize the study of foreign threat actors. The foreign activities we characterize as national security threats—e.g., terrorism, nefarious cyber activity, weapons proliferation, and transnational organized crime—are manifestations of the same underlying conditions and behaviors that engender adverse social and criminal justice outcomes in the United States.[30] By developing a more sophisticated understanding of the common attributes that link foreign threat activity and negative domestic outcomes, social scientists and psychologists could leverage the troves of social welfare and law enforcement data maintained by federal, state, and local governments to develop robust analytics that predict national security threats. Instead of relying exclusively on foreign intelligence leads to locate specific overseas actors who threaten America's security, algorithms trained and refined by domestic social welfare data could reason inductively to identify the behaviors, locations, conditions, and circumstances overseas that are material to our national security. Mining these results could generate novel leads that improve the exploitation of existing intelligence collection; broadening the perspective of IC analysts evaluating national security risk in the same way that the Allegheny Family Screening Tool improved the assessments of CYF call screeners.

---

[30] *See, e.g.*, N. Veerasamy, *A High-Level Conceptual Framework of Cyber-Terrorism*, 8 J. INFO. WARFARE, no. 1, 2009, at 43, 46 ("[T]he distinction between cyber terror and cyber crime . . . does not lie in the mechanics of the event, but rather in the intent that drove the person's actions." (citation omitted)).

Predictive analytics that apply lessons learned from social policy data to inform foreign intelligence activities pose obvious risks to privacy and civil liberties. But these concerns can be addressed through transparent processes and careful regulation. Since the AI value of social policy data is divorced from the identities of the people indexed in social service databases, the requirements governing this program would compel the government to strip all personally identifiable information from administrative records used to develop algorithms for national security purposes. The regulations would further prohibit federal government personnel from viewing, querying, or otherwise accessing these data pools. Only machine tools vetted and ratified by a cleared oversight panel representing government, private sector, and civil liberties interests would be approved to process the data. The panel would also sample and audit results transmitted by the machine tools to ensure that personal information is appropriately protected. Finally, the panel would advance key privacy and civil liberties objectives by working collaboratively with national security officials to assess whether use of these analytics expands or narrows the population of U.S. persons and other individuals targeted as subjects of interest.

B. *Promoting AI as a New Organizing Concept for Foreign Intelligence Collection*

Traditionally, the IC has structured intelligence activity around frameworks that optimize human data processing. Its collection apparatus is partitioned into areas of specialization demarcated by geographic region (e.g., South America, Africa), threat category (e.g., counterterrorism, cyber), specific target, method of collection, etc. This modular infrastructure may contribute certain efficiencies and promote subject matter expertise among human analysts, but it impedes the IC from consolidating data across topics and modalities into data pools that emulate the rich administrative datasets supporting predictive analytics in the social services realm.

These pools of contextualized social services data are a boon to the IC because they model the conditions that would prime AI for success in the national security domain. To that end, the examples outlined above suggest that the IC has the resources and capabilities to generate datasets capable of supporting effective predictive analytics. The missing ingredient is a commitment by national security officials to establish AI as a new center of gravity in the IC that can orient data integration efforts, tasking requirements, and future intelligence gathering around a core mission of creating datasets primarily designed to support machine analytics.

Data scientists could lay the groundwork for these data pools in three phases. First, the IC should begin consolidating existing records on national security threat actors into a single repository, prioritizing records with numerous data points and rich narrative detail. National Security Presidential Memorandum 7 (NSPM 7) already directs federal agencies to develop technical architecture "to advance the integration, sharing, and use of identity attributes and associated derogatory information" for individuals falling into five enumerated threat categories.[31] Though NSPM 7 contemplates the creation of separate databases for each threat category, national security officials should direct the executive agents for each database to harmonize the structure, formatting, and functionality of all NSPM 7 architecture. This will ensure that identity records

[31] Memorandum on Integration, Sharing, and Use of National Security Threat Actor Information to Protect Americans, 2017 DAILY COMP. PRES. DOC. 722, at 2 (Oct. 4, 2017), https://www.govinfo.gov/content/pkg/DCPD-201700722/pdf/DCPD-201700722.pdf [https://perma.cc/WNE7-92QM].

initially archived in threat-specific databases can flow downstream into data pools that merge all-source and all-threat intelligence to support predictive analytics.

Second, the IC should direct agencies to coordinate responses to common intelligence taskings with an eye toward corroborating and enhancing identity information gathered by other agencies. Horizontally integrated social services records could provide a model for these efforts. By reverse-engineering the anatomy of these records, data scientists can establish common standards for identity records that agencies can build cooperatively for use in IC data pools. Where, for example, the IC has acquired extensive human intelligence (HUMNIT) on a group of priority threat actors, respective agency partners should be prompted to determine whether signals intelligence (SIGINT), geospatial intelligence (GEOINT), or open-source intelligence (OSINT) capabilities could augment that knowledge to enrich identity records until they meet the threshold for inclusion in the foreign intelligence data pool.[32]

Third, the IC should consider pursuing new sources of information that may have limited utility to a human analyst, but unique value to a machine algorithm. As just one example, foreign countries with advanced social welfare systems (and threat actors within their borders) possess vast stores of social security and social welfare data that algorithms have successfully mined to advance other AI initiatives.[33] The IC could leverage relationships with foreign partners to either gain direct access to this data under appropriate conditions, or launch joint ventures with foreign partners to develop advanced analytics supported by a multi-national social welfare dataset.

V. CASE STUDY: DEVELOPING AN ANALYTIC TO PREDICT EXTREMIST ACTS

In October 2016, the President updated the nation's strategic implementation plan to address violent extremism in the United States, observing:

> In many ways, the threat of violent extremism today is more challenging than ever before. Violent extremists have demonstrated an ability to entice people to travel great distances, to direct attacks, and to inspire others to act from afar. They have utilized the Internet and other technologies, specifically social media platforms, as a means to reach a greater number of people in more places, tailor messages to appeal to different audiences, and reach out to potential recruits individually.[34]

Unlike other national security threats, the government cannot combat the domestic influence of violent extremists by strengthening America's borders or by targeting hostile actors

---

[32] For a brief introduction to these four intelligence collection disciplines, see *What Is Intelligence?*, OFF. OF THE DIRECTOR OF NAT'L INTELLIGENCE, https://www.dni.gov/index.php/what-we-do/what-is-intelligence [https://perma.cc/NL3U-UZ6P].

[33] *See, e.g.*, Longbing Cao, *Social Security and Social Welfare Data Mining: An Overview*, 42 IEEE TRANSACTIONS ON SYS., MAN & CYBERNETICS—PART C 837, 837 (2012).

[34] EXEC. OFFICE OF THE PRESIDENT OF THE U.S., STRATEGIC IMPLEMENTATION PLAN FOR EMPOWERING LOCAL PARTNERS TO PREVENT VIOLENT EXTREMISM IN THE UNITED STATES 1 (2016), https://www.dhs.gov/sites/default/files/publications/2016_strategic_implementation_plan_empowering_local_partners_prev.pdf [https://perma.cc/JST7-87JQ].

and networks overseas. Lacking an obvious countermeasure, the government has witnessed homegrown violent extremism proliferate as individuals radicalized in the United States conducted lethal, ideologically motivated attacks to promote the ideologies and agendas of foreign and domestic terrorist groups. Just within the last twenty months:

- On December 6, 2019, Mohammed Saeed Alshamrani, an aviation student from Saudi Arabia, killed 3 people and injured 8 others in an attack at Naval Air Station Pensacola in Pensacola, Florida.[35] The Department of Justice concluded that the attack was an act of terrorism motivated by jihadist ideology.[36]
- On August 3, 2019, Patrick Crucius killed 22 people and injured 26 others in a mass shooting at a Walmart store in El Paso, Texas.[37] In a manifesto published prior to the attack, Crusius stated that the attack was "a response to the Hispanic invasion of Texas."[38]
- On October 27, 2018, Robert Bower—an anti-Semite and proponent of white nationalism—killed 11 people and injured 6 others at the Tree of Life Synagogue in Pittsburgh, Pennsylvania.[39]

These acts of violent extremism bedevil national security and law enforcement officials because they are notoriously hard to predict. According to the National Consortium for the Study of Terrorism and Responses to Terrorism: "There are *no known* pathways, definite set of risk factors, or reliable predictors that would indicate who is likely to commit violent acts driven by extremism."[40] The diffuse nature of this threat is punctuated in the President's first national strategy for preventing violent extremism: "Individuals from a broad array of communities and walks of life in the United States have been radicalized to support or commit acts of ideologically-inspired violence."[41] "[They] come from different socioeconomic backgrounds, ethnic and

---

[35] William P. Barr, Attorney Gen., U.S. Dep't of Justice, Announcement of the Findings of the Criminal Investigation into the December 2019 Shooting at Pensacola Naval Air Station (Jan. 13, 2020), https://www.justice.gov/opa/speech/attorney-general-william-p-barr-announces-findings-criminal-investigation-december-2019 [https://perma.cc/6SGT-CUZC].

[36] *Id.*

[37] Vanessa Romo, *El Paso Walmart Shooting Suspect Pleads Not Guilty*, NPR (Oct. 10, 2019, 4:31 PM), https://www.npr.org/2019/10/10/769013051/el-paso-walmart-shooting-suspect-pleads-not-guilty [https://perma.cc/J34N-GXSM].

[38] Nicholas Bogel-Burroughs, *'I'm the Shooter': El Paso Suspect Confessed to Targeting Mexicans, Police Say*, N.Y. TIMES (Aug. 9, 2019), https://www.nytimes.com/2019/08/09/us/el-paso-suspect-confession.html [https://perma.cc/G34J-UUYX].

[39] Campbell Robertson et al., *11 Killed in Synagogue Massacre; Suspect Charged with 29 Counts*, N.Y. TIMES (Oct. 27, 2018), https://www.nytimes.com/2018/10/27/us/active-shooter-pittsburgh-synagogue-shooting.html [https://perma.cc/WS5Q-GMJQ].

[40] NAT'L CONSORTIUM FOR THE STUDY OF TERRORISM & RESPONSES TO TERRORISM, SUPPORTING A MULTIDISCIPLINARY APPROACH TO ADDRESSING VIOLENT EXTREMISM: WHAT ROLE CAN EDUCATION PROFESSIONALS PLAY? 1 (2015), https://www.start.umd.edu/pubs/START_LessonsLearnedfromMentalHealthAndEducation_EducatorSummary_Oct2015.pdf [https://perma.cc/S6MC-2F6M].

[41] PRESIDENT OF THE U.S., EMPOWERING LOCAL PARTNERS TO PREVENT VIOLENT EXTREMISM IN THE UNITED STATES 2 (2011), https://www.dhs.gov/sites/default/files/publications/empowering_local_partners.pdf [https://perma.cc/6SEY-ZGJ2].

religious communities, and areas of the country, making it difficult to predict where violent extremist narratives will resonate."[42]

Demystifying the recent outbreak of domestic terrorism and violent extremism—and discerning the antecedents that telegraph future attacks—is the quintessential policy predicament that could be unlocked by an AI solution. These extremist acts appear to be a function of three factors: (1) individual attributes and predispositions; (2) community conditions; and (3) the influence of extremist ideologues. All three dynamics are expressed and captured in data currently compiled by government agencies or available in open-source collection. But analysts have struggled to uncover patterns in this data that explain the aforementioned attacks, highlight existing high-risk conditions, and recommend services and interventions to neutralize these risks.

By merging relevant information into the type of contextualized data pools described above, the IC could give machine-learning algorithms the opportunity to identify patterns in extremist behavior that elude human observation. Ideally these data pools would be populated by information from local, state, federal, and non-government sources. State and local administrative data from social service agencies could contribute person-level information reflecting biographic, behavioral, and criminal justice indicators, as well as a wealth of community-level demographic information. IC elements would enrich these data pools with intelligence information on known extremist threat actors, prioritizing inputs that provide visibility into domestic infrastructure, messaging tradecraft, Internet activity, and pattern-of-life. Government open-source analysts, in collaboration with private partners, could supplement these databases with social media data that researchers have previously aggregated and mined to conduct detailed impact analyses of social media accounts promoting violent extremism.[43] Establishing comprehensive data coverage across a broad universe of individuals and communities exposed to extremist influences may enable machine-learning algorithms to engineer analytics that can accurately predict when extremist expression will tip into lethal and destructive attacks.

This collaboration would require national security and social policy professionals to engage and trust one another—a dynamic conspicuously absent from government policymaking in recent decades. The COVID-19 pandemic underscores the importance of integrated, whole-of-government strategies to combat existential threats; an approach that will serve the nation well when it can renew in earnest the campaign to eradicate domestic terrorism and violent extremism.

CONCLUSION

The national security community tends to regard itself as a favored son among government and commercial sectors—privileged by abundant resources, influence with policymakers, broad access to data, and top-to-bottom talent in its workforce. Buoyed by this sense of exceptionalism, national security officials have been slow to recognize that intelligence agencies and defense elements are uniquely *disadvantaged* in the race to evolve core capabilities through AI innovation.

---

[42] *Id.* at 1.

[43] *See* Tamar Mitts, Countering Violent Extremism: Do Community Engagement Efforts Reduce Extremist Rhetoric on Social Media? 11–13 (Apr. 6, 2017) (unpublished manuscript), http://tamarmitts.com/wp-content/uploads/2016/06/Mitts_community_engagement_v3.pdf [https://perma.cc/MU66-8XC4] (describing use of Twitter data on Islamic State supporters in the United States to measure the online expression of extremist ideology).

Numerous public and commercial entities have enjoyed a more natural transition to advanced analytics, and have discovered that machines can process the same data that humans review to automate manual functions, anticipate problems before they materialize, and recommend novel solutions to stubborn challenges.

Security agencies have fallen behind this curve. Fragmented intelligence—acquired through covert means and subject to strict collection and retention rules—is an inherently poor fit for AI. The result is a low ceiling for even the most advanced machines and algorithms attempting to spin data straw into rich and revealing patterns.

These realities demand new intelligence programs customized for machines. As a government we should reassess how traditional intelligence collection is resourced and prioritized at a time when the most valuable national security data for machines increasingly resides in open sources and mediums. As a society we should reexamine the questionable assumption—codified in current rules—that privacy interests are impacted equally by human and machine review of the same information. Advanced analytics cannot be optimized in data environments constrained by collection and processing rules designed for human analysts. To strike the appropriate balance between civil liberties and national security we should explore regulatory frameworks that advance a more contemporary notion of privacy by prescribing unique and segregable roles for humans and machines in modern information domains.