

Center on International Cooperation

NYU | WAGNER

Turning Warnings into Action, and Action into Peace

BRANKA PANIC

Founder and Executive Director,
AI for Peace; NYU Center on
International Cooperation
Non-Resident Fellow
Lead author

ANGELICA LINDQVIST-MCGOWAN

Director of Operations and Outreach,
VIEWS, Department of Peace and Conflict
Research, Uppsala University
Contributing author

LUISA PORTUGAL

Senior Program Officer, Pathfinders
for Peaceful, Just and Inclusive
Societies, NYU Center on
International Cooperation
Contributing author



Reflections from Data for
Peace Conference 2026

September 2026

Table of Contents

1. Introduction	3
2. “Know Your Data:” Between Saturation and Scarcity	5
3. AI in Conflict Early Warning Early Action (EWEA)	9
4. Conflict Forecasting in the Age of AI	14
5. From Warning to Action: Closing the Field's Most Persistent Gap	18
6. Trust, Accountability, and Political Will	24
7. Nothing About Us Without Us:	28
Conclusion: What This Moment Demands	32
Endnotes	34

About the Data for Peace Conference

[Data for Peace Conference 2026](#) (15–17 June, Stockholm and online) brought together the research, technology, policy, and practitioner communities working on conflict prevention. Two in-person plenary days in Stockholm—attended by 150 participants on site—were followed by a virtual day of parallel tracks that drew more than 1500 registered participants. Session recordings from the three days have since generated over 1800 unique views, and counting. Formats ranged from panels and lightning talks to workshops, peer-review clinics, and scenario simulations—designed not only to showcase what is technically possible but to ask what is required for innovations to be applied responsibly and effectively to turn warnings into action and action into peace. The Conference was co-hosted by the [Department of Peace and Conflict Research at Uppsala University](#), the [NYU Center on International Cooperation](#), and [AI for Peace](#)—with support from the [Complex](#)

[Risk Analytics Fund \(CRAF'd\)](#), [Folke Bernadotte Academy \(FBA\)](#), and the [Finnish Ministry for Foreign Affairs](#).

The Conference generated the richest body of content the Data for Peace initiative has produced to date, and no summary can do justice to the full breadth of what was shared across both in-person and virtual days. This is not a comprehensive account but a thematic synthesis—an attempt to trace the threads that connected sessions across three days, rather than a reflection of any single session's full content. We begin by looking at where the field stands on peace and conflict data—how and whether AI is helping, where conflict modeling is headed, and whether the field is finally closing the gap between warning and action. We then turn to the harder questions: what trust, accountability, and political will are required to make warnings actionable—and who early warning systems ultimately serve and belong to.

For the full picture, we invite you to explore the recordings of individual sessions on our [YouTube channel](#), or complete conference days ([Day 1](#) and [Day 2](#)).

Acknowledgements

Editorial support and photography were provided by Symphony Chau.

About the Center on International Cooperation (CIC)

The Center on International Cooperation (CIC) is a nonprofit research center housed at New York University. For over three decades, CIC has been a leader in applied policy that links politics, security, justice, development, and humanitarian issues. CIC's mission is to strengthen cooperative approaches among national governments, international organizations, and the wider policy community to prevent crises and advance peace, justice, and inclusion. Learn more at cic.nyu.edu and @nyuCIC.

© New York University Center on International Cooperation, All Rights Reserved, 2026.

1. Introduction

Unprecedented Crisis, Unprecedented Capability

"Unprecedented challenges demand unprecedented cooperation—and for us, that cooperation is the Data for Peace community. It is you. That is why we brought together 57 sessions, 131+ speakers, and thousands of participants in the 2026 edition of Data for Peace. My invitation is simple: let these numbers be a starting point—to turn data into action, and action into peace."

—**Branka Panic**, Founder and Executive Director, AI for Peace

The picture the data paints is unambiguous, and it is grim. The world is experiencing its most severe deterioration in peacefulness since the end of the Second World War. The Uppsala Conflict Data Program (UCDP)¹ reports that interstate conflicts doubled for the second consecutive year in 2025, alongside a record of 65 state-based conflicts—the highest figure since 1946—and roughly 76,500 deaths from one-sided violence, a more than 400 percent year-on-year increase and the heaviest civilian toll since the 1994 Rwandan genocide.² The Global Peace Index marks a twelfth consecutive year of peace decline, with 119 countries less peaceful than in 2008 and geopolitical risk now exceeding Cold War-era levels.³ Military spending rose by 5.8 percent in 2025—the largest single-year increase in nearly two decades—while peacebuilding and peacekeeping together accounted for just 0.5 percent of that total. The human consequences are without precedent: over 120 million people forcibly displaced, some 300 million in humanitarian need, and a mounting toll measured in lives, livelihoods, and foregone development.

At the same time, our analytical and technological capabilities have never been greater. Technologies are developing faster than ever before—and, as it was echoed throughout the conference, never in the future will they move as slowly as they do today. The vocabulary itself tells the story: when the Data for Peace project began in 2019, we spoke about *emerging* technologies; today, we speak about *frontier* technologies, moving rapidly from theory to prototypes, pilots, and products. These frontier technologies are already transforming the wider environment in which peacebuilding takes place—information ecosystems, geopolitics, the diplomatic toolkit, and the very definition of power. The proof was

on this year's agenda: across three days, session after session showcased teams already embedding machine learning and generative artificial intelligence (AI) into their workstreams, from data collection and conflict analysis to early warnings and policy action. Yet at the institutional level—in governance frameworks, funding structures, and standard operating procedures—the field has barely begun to adapt. The result is the defining tension of our field: the gap between how fast technology evolves and how slowly governance and norms catch up.

→ Explore more: [Day 1 Opening Remarks](#)



2. “Know Your Data:” Between Saturation and Scarcity

A persistent theme ran across every session: **the imperative to know your data**. Yet even on this point, consensus proved elusive. Some participants argued that the field is already too saturated—that the volume of peace and conflict-relevant data now exceeds the capacity of most institutions to absorb it. Turning this saturation into opportunity rather than an obstacle, the United Nations Development Programme (UNDP)’s Crisis Risk Dashboard⁴ transforms a flood of open data into a centralized warehouse, enabling integrated analysis of environmental, political, and security risks. Building on the earlier versions of dashboards used for analytics, programing, advocacy and engagement within UNDP and at the national level, the tool has evolved into the Risk Anticipation Data Hub⁵ (RAD), piloted at the conference. The Hub was developed as a digital public good, providing an API to pull data directly into workflows, with navigation by topics, categories and other filters for a multidimensional understanding of risk.

→ Explore more: [Demo of the Risk Anticipation Data Hub](#)

Experts cautioned that **conflict data remains a representation, not a mirror**—shaped by what surfaces through reporting. The honest assessment is that the picture is uneven, depending on geography, indicator, and the populations represented—or excluded—in the underlying data.

→ Explore more: [AI-Assisted Conflict Event Coding: Fresh Insights and Hard Lessons from UCDP](#) and [Five Things Users Should Know About Conflict Event Data](#)

In some contexts, analysts are managing a plethora of signals drawn from social media, satellite imagery, event databases, and sentiment analysis. In others—frequently the most fragile, remote, or politically sensitive—coverage is minimal, and enduring gaps in granularity persist: the sub-level dynamics, gendered patterns of violence, and slow-accumulating effects that aggregate datasets fail to capture. Some events also remain nearly invisible regardless of context, due to factors such as reporting bias.

The gaps are as analytically significant as the data itself. The locations hardest to access, the forms of violence least reported, and the voices least represented are precisely the contexts in which predictive models perform worst. Office of the United Nations High Commissioner for Human Rights (OHCHR)'s Human Rights Data Exchange (HRDx)⁶ remote crisis-monitoring capability—deployed and tested during the latest Iran crisis—demonstrates one response to this access challenge. Processing roughly one million fragmented open-source signals daily across languages, tracking conflict events, classifying potential violations against international humanitarian and human rights law, and detecting connectivity disruptions that **are themselves early warning indicators**. The design principle is precise: the tool flags the question, never decides the answer—the system surfaces potential violations, the human officer reaches the conclusion.

→ Explore more: [Monitoring the Unreachable—From Early Warning to Early Action](#)

Another concern surfaced repeatedly across the three days: **the field's core datasets are at risk of disappearing**. Donor commitments are contracting, key institutional stakeholders are disengaging, and long-running collection efforts—some with decades of methodological continuity—are being defunded at precisely the moment demand for their outputs is highest. As more than one participant observed: *"without ACLED and UCDP data, this work would not exist,"* pointing to the field's dependency structure. Financing for data infrastructure must therefore be sustained, predictable, and insulated, to the extent possible, from the political volatility that affects the wider ecosystem.

"Something that we are doing is...licensing: moving away from the really grant-based model...towards licensing the things that [we're] producing as opposed to only relying on grant cycles."

—*Katayoun Kishi*, Head of Data Science, ACLED

But sustaining the field is not only a question of financing. **Equally important is connecting data across sectors and systems**, rather than continuing to develop methodologically rigorous but operationally isolated silos.

“It is true that we have more data than ever, and we have built many models, analyses, and platforms on top of that, but there is one challenge that remains in the core layer that feeds all those tools and that is that it is fragmented and siloed.”⁸

—*Antje Lehmann*, Fund Manager, Complex Risk Analytics Fund

Conflict data only becomes powerful when it is connected: when it reveals not just where violence is occurring but how it is evolving, how it is moving geographically, and how it is impacting populations—shaping displacement, generating protection risks, restricting access to health and education. No single dataset delivers that picture alone. A CRAF'd survey of more than one hundred ecosystem organizations confirmed the extent of the problem: **fragmentation, duplication, lack of standards, and persistent issues of data quality.**⁹

The Centre for Humanitarian Data¹⁰ found that only 75 percent of crisis data in major humanitarian contexts is up to date and timely—meaning some data points feeding operational decisions may already be old. It is precisely this structural deficit that explains CRAF'd's investment logic: not funding anticipatory action directly, but strengthening the underlying data infrastructure required to act earlier, faster, and more precisely.

→ Explore more: [Closing the Action Gap: From Signals to Systems for Anticipating the Impacts of Violence](#)

Alongside these structural challenges, the question is **how to produce data more sustainably, without eroding quality.** ACLED¹¹ already publishes twice weekly for every country in the world—and daily in active conflicts such as Iran—to keep pace with demand. Yet this is still achieved through purely human means: researchers manually visiting websites, watching videos, listening to radio, and translating what they find into structured data. For them, the turn toward AI has been driven by this unsustainable ratio of demand to capacity. But rather than relying on AI for event coding itself, this has meant using scrapers and text analysis scanning sources and flagging likely conflict-relevant reports, large language models (LLMs) transforming researchers' rough bullet points into polished event notes, and automated pipelines guiding researchers to signals rather than leaving them to find them manually. Some steps, however, remain difficult to automate—matching actors against established dictionaries, for instance. The goal with this approach is **targeted efficiency gains that preserve human judgment, not automation that replaces it.**¹²

Conflict Forecast¹³ has pursued a different logic altogether: in a declining funding environment, it has sought to build systems that are as automated as possible, requiring fewer people to run and monitor—though without sustained resources, more frequent data releases remain out of reach.

Sustaining data technically and financially is necessary but not sufficient. Data must also be sustained relationally—which means **making it relevant to the communities on the ground**. That is not a matter of handing them a dashboard; it requires connecting forecasts to operational meaning—how many people live in an area, where the nearest hospital is, how conflict is expected to move through that specific place. And it must be a two-way street: data should not only flow down to communities but up from them.

→ Explore more: [Post-Mortem, Reset or Business as Usual? Conflict Early Warning in a Changing World](#)



3. AI in Conflict Early Warning Early Action (EWEA): Promise and Peril, Side by Side

If there was one technology that dominated the week, it was, predictably, AI. The room held genuine tension on the benefits and harms of AI—and that tension did not resolve itself. **What everyone agreed on is that AI is here to stay.**

Beyond that, the spectrum was wide: some argued it is already harming more than helping, distorting information environments and offloading judgment to systems no one fully understands; others pointed to cost savings, faster analysis, and reach into languages and contexts previously out of bounds. There was no conviction that the net impact of generative AI will be positive—but there was more optimism about its applications within the early warning field specifically, and a shared sense that demonstrating constructive uses matters.

Many projects showcased how LLMs and chatbots are now embedded into real workstreams. AI is proving capable at parsing massive volumes of qualitative information—extracting discrete claims, tagging them across categories, and organizing them cost-effectively at a scale no human team could match. As Living Signals¹⁴ demonstrated, a natural language processing (NLP)-driven infrastructure can capture real-time data from hundreds of local news sources and, when connected to tools like Claude and ChatGPT, expand the analytical lens well beyond violent conflict events.

→ Explore more: [From Conflict Events to Security 'Signals.' Using AI to Build a Search Engine for Peacebuilders](#)

"This isn't just a technical innovation but a political one, because we are expanding our search criteria beyond violent conflict events to address wider forms of security systematically, allowing us to better understand things like structural violence, slow emergency, gendered emergencies, and capturing indicators of positive peace."

—Sami Murphy, Co-Founder of Living Signals

→ Explore more: [From Conflict Events to Security 'Signals': Using AI to Build a Search Engine for Peacebuilders](#)



AI is already embedded in the data-production pipeline, as was mentioned above: building scrapers and automating source monitoring, processing raw articles into text features for forecasting models (on the premise that language itself carries predictive signals for armed conflict), extending collection into multimodal sources such as radio and television, unlocking reporting across multiple languages, and even developing tools like browser extensions that shield analysts from graphic and traumatic material.

→ Explore more: [Airwaves to Algorithms: Harnessing Radio Broadcasts to Track Conflict Narratives](#) and [AI-Assisted Conflict Event Coding: Fresh Insights and Hard Lessons from UCDP](#)

It is therefore facilitating broader data collection, at more local levels, from more diverse sources—and doing so faster. Other efforts use AI not to compete with hard security datasets on speed or granularity, but to surface a complementary layer of socio-political insight—tracking the gap between what mainstream media reports and what affected populations are actually saying about changing security conditions.

→ Explore more: [Reading the Room: Local-Language Narrative Monitoring as an Early Warning Tool in Conflict Settings](#)

The Institute for Economics and Peace¹⁵ (IEP) tested whether LLMs can replicate the expert judgment that underpins much of peace and conflict measurement—coding indicators like conflict escalation, intensity of internal conflict, and implementation of peace agreements. By forcing models to justify their scores level by level, the team made AI reasoning auditable rather than opaque. The results showed that when questions are well-defined, AI aligned closely with human experts; when concepts are more ambiguous, agreement broke down.

Four lessons emerged from their research:

- 1) AI delivers real efficiency gains, but experts remain essential for final judgment.
- 2) A model that performs well on one case may fail on others—reliability must be tested broadly.
- 3) How you phrase the prompt is itself a research choice.

- 4) A model that gives the same answer repeatedly is not necessarily giving the right one



"Why this approach is going to become increasingly common is that you can do this so quickly and cheaply. [Compared to] what we were paying for this data previously, [we are now] able to produce 20 years' worth of data for essentially 5% of the cost. So in terms of money, it's pennies on the dollar –and in terms of time, it's seconds on the hour."

—*Thomas Morgan*, Chief Research Officer, Institute for Economics and Peace

→ Explore more: [The LLM as Expert Conflict Coder: A Structured Adjudication Methodology and Its Limits](#)

Across the conference, sessions gave equal attention to the risks of applying AI in conflict early warning and early action: biased outputs, hallucinated facts, opaque decision pathways, and the dangerous erosion of analytical capacity when humans

As one speaker framed it:

—*Mert Can Yilmaz*, Research Engineer, Uppsala Conflict Data Program (UCDP), Department of Peace and Conflict Research, Uppsala University

A man with a beard, wearing a dark blazer over a light-colored shirt, stands on a stage. He is gesturing with his right hand towards a large screen behind him. The screen displays the title "Conflict data is not a census. It is a representation." in a serif font, with "representation" underlined. Below the title, a paragraph reads: "We see only what surfaces through reporting, and the gap is widest exactly where you'd least want it: the most violent, hardest-to-reach places." At the bottom of the screen, a citation is visible: "» sources: Lewis, A., Donati, D., Hays, M., Hoggan, S., & Thomas, M. (2020). The underreported death toll of war: A globalized measurement from a network." In the top right corner of the screen, a dark box contains the text: "» THE TAKEAWAY FOR THIS D... You are not tra... on a represent...". The stage is decorated with orange curtains, a black podium, several black speakers, and two orange armchairs. A small table with water bottles and glasses is next to the chairs. The audience's heads are visible in the foreground.

→ Explore more: [AI-Assisted Conflict Event Coding: Fresh Insights and Hard Lessons from UCDP](#)

4. Conflict Forecasting in the Age of AI: Contextualization and the Human Element

The central challenge in conflict forecasting is no longer simply whether models can predict risk, but how much of a complex and fast-changing reality they can capture—and how far those predictions can support decisions about what comes next. Across the forecasting sessions, speakers approached that problem from different directions: refining what aspects of armed conflicts models seek to capture, using AI to add context around their outputs, and expanding forecasting beyond violence itself to its human and economic consequences.

The first problem concerns the level at which conflict is modeled. Country-level forecasts remain useful for cross-national comparison, but they can obscure the very dynamics practitioners need to understand. Two conflicts within the same state—even involving the same actors—can follow entirely different logics. ACLED's¹⁶ MOVE model (Movement of Violent Events) responds by mapping the distinct theaters of conflict within a country: how they shift over time, who the actors are, what events emerge, and how civilians are affected. The result is a "signature" for each conflict theatre, providing something closer to what practitioners might actually need: not simply an abstract national risk score, but a localized understanding of how a specific conflict may affect staff safety, operational continuity, and civilian protection.

Geography creates another problem. As the Joint Research Centre (JRC)¹⁷ emphasized, conflicts rarely stay contained within national borders. Many begin as internal disputes and then internationalize, while forecasting models and the data infrastructure beneath them are still often organized country by country. Capturing these dynamics requires not only better models, but cross-border data infrastructure most of the field has yet to build.

But even as forecasting models become better at capturing where and how conflict evolves, their outputs can only communicate so much on their own. A probability, risk score, or chart inevitably strips away much of the political and

operational context that gives it meaning. This is a different problem—and one where LLMs and agentic AI may have a more immediate role.

At ACLED, agentic AI is being explored as a layer on top of existing tools to contextualize their outputs, allowing users to interact with the data and draw operationally relevant meaning from it. At the JRC, a forthcoming chatbot, due by year's end, retrieves information through a retrieval-augmented generation (RAG) system and generates figures and maps on demand. It draws on sources like ACLED's event notes to give users under time pressure a richer understanding of the situation surrounding a forecast than a bar chart alone can provide.

→ Explore more: [Post-Mortem, Reset or Business as Usual? Conflict Early Warning in a Changing World](#)

While optimistic, speakers were honest about the technology's limits: LLMs will not, by themselves, make predictions more accurate, nor do they solve the explainability and interpretability problems that preoccupy forecasting researchers. Their real value is **contextualization**—situating a bare probability within a coherent model of the world drawn from the news cycle, internal documents, and other forms of information that help users understand what a forecast may actually mean. **The frontier here, in other words, is not better math; it is better meaning.**

This distinction matters as some of the hardest forecasting failures are not simply failures to detect risk. As one speaker noted, the onset of today's major wars—Gaza, Ukraine, Iran—is rarely just a quantitative problem. The harder problem is one of integrating a human analytical layer—the context-weighting judgment of experienced observers—into forecasts so they capture not only slow-moving trends but also the element of surprise.

But contextualizing a forecast is not the only way to make it more useful for action. The Violence and Impacts Early-Warning System (VIEWS)¹⁸ at Uppsala University and Peace Research Institute Oslo addresses a related bottleneck: translating forecasts of conflict-related fatalities into the humanitarian consequences affected populations may face if those forecasts materialize. In addition to estimating how many people are expected to die from armed conflict in a given time and place, VIEWS now forecasts impacts such as changes in gross domestic product (GDP) per capita and excess child mortality at a monthly resolution, up to 36 months ahead. The work responds to a persistent timing mismatch in early warning:



"On the one hand we have the conflict data which comes monthly and is updated frequently. But when it comes to understanding its consequences, we often face a delay problem, because social and economic indicators are typically reported annually. That means we observe conflicts quickly, but we observe its consequences slowly. Early warning systems therefore miss the economic and human impacts."

—**Borbála Farkas**, Research Engineer, Department of Peace and Conflict Research, Uppsala University

→ Explore more: [From Conflict Forecasts to Socio-Economic Impacts: Operational Use Cases and Challenges](#)

This represents a different route toward making forecasting more actionable. Where LLM-based tools seek to place more context **around** an existing prediction, impact forecasting seeks to **broaden what is predicted in the first place**. Both approaches respond to the same underlying gap between a model output and the decisions that early warning is ultimately meant to support.

But making forecasts more useful also raises other, more fundamental, questions that most AI conversations leave unspoken. **How much of the analytical and decision-making process should be delegated to the systems producing and interpreting them?** What happens when the systems and model access we treat as stable become unavailable—or, in a volatile geopolitical landscape, deliberately manipulated? Does routing sensitive information through commercial cloud infrastructure expose the very communities the work is meant to protect, making "do no harm" a design feature rather than an afterthought? And what happens when today's cheap token usage is no longer cheap? The conclusion was uncomfortable but clear: building sustainably means confronting the full geo-techno-political landscape now, before dependency hardens into vulnerability.

→ Explore more: [Post-Mortem, Reset or Business as Usual? Conflict Early Warning in a Changing World](#)

This, in turn, leads to an even larger question about the direction of the field. One session posed what may have been the most clarifying question of the conference: if resources were unlimited, **what would the ideal prevention model look like?** The difficulty participants had in answering it is itself diagnostic of how reactive rather than strategic the field has become.

→ Explore more: [If Data Fails, the Model Fails: Sustainability, Gaps, and Anticipatory Action](#)

A second question cut equally deep: **where is the VIEWS equivalent for peace?** The field has built sophisticated predictive infrastructure for forecasting violence, escalation, and now some of their downstream consequences. However, there is no comparable forecasting system for anticipating where and how peace is likely to consolidate, or projecting the conditions under which societies move toward stability rather than away from conflict.

Research and monitoring on positive peace exist—from the Positive Peace Index to a growing body of scholarship on resilience and social cohesion—but predictive systems built on that evidence base do not. The asymmetry is not accidental; it reflects a political economy in which crises mobilize resources and attention in ways that the promise of peace never has.

5. From Warning to Action: Closing the Field's Most Persistent Gap

The warning-to-response gap remains the central unresolved challenge—and the question that has followed this program since its inception (we cover this challenge in 2022 via a NYU Center on International Cooperation report).¹⁹ **Improvements in data availability, or forecasting accuracy do not, on their own, generate timely and context-sensitive action.** The problem is not access to information; it is turning information into decisions about where attention, resources, and action should go. Closing the gap requires, among other things, sustained engagement with decision-makers whose rotation through positions complicates long-term institutional partnerships, interfaces designed for operational use rather than analytical inspection, and products that translate model outputs into insight officials can act on within their own working timeframes.

One initiative launched at the conference this year addressed this gap directly: INFORM Warning,²⁰ a global system for forward-looking crisis risk analysis, explicitly designed not merely to produce better forecasts but to make them actionable. The system addresses the underserved 12-month time window, selecting the most relevant indicators, transforming them into standardized signals capturing future changes in risk, aggregating those signals into warning levels to prioritize countries, and publishing crisis alerts for the most critical emerging risks—translating forecasts into clear signals of where planning, funding, and action are needed. Its distinct value is prioritization—answering where action is needed:

→ Explore more: [Launch of INFORM Warning: Prioritizing Future Crises for Early Action](#) and download the [User Guide](#)





At the operational level, the Norwegian Refugee Council's CLEAR system²¹ addresses a complementary challenge—not which crises to prioritize but how fragmented signals reach responders in time to act. An open-source pipeline that ingests data from conflict monitors, media, displacement trackers, and community reports, CLEAR uses AI to cluster fragmented signals into coherent events and delivers actionable alerts to responders through a single open API—built on the premise that a warning's value lies not in its analytical sophistication but in whether it reaches someone positioned to act.

→ Explore more: [Lost in 90 Tabs: Closing the Last Mile Between Humanitarian Data and Field Decisions](#)

But both prioritization and delivery assume the signal is already in the system. There is a category of early warning that most conflict datasets are not built to detect—not because the data is unavailable, but because the systems were never designed to look for it. As one session highlighted: gender backlash and peace backlash move together—and gender backlash moves first. When women withdraw from public life, when their mobility shrinks, when networks weaken,

when economic exclusion deepens for female-headed households—these are not side effects of insecurity. They are its earliest visible form.



A deeper challenge concerns what early warning systems can see at all. In protracted conflicts, deterioration does not arrive as an event—it arrives as an accumulation, and accumulation is precisely what event-based systems are not designed to register. In Georgia, where conventional early warning never gained traction, UN Women built an alternative: a gender-responsive monitoring tool that reconceptualizes conflict-affected women not as informants but as primary knowledge producers. The tool rests on three premises: that human security manifests through small gendered disruptions long before formal systems register them; that data collection must be decentralized and validated by civil society on the ground; and that findings must align to decision-making cycles where action is possible.

→ Explore more: [Women Know First: Early Warning on Gender and Peace Backlash in Europe and Central Asia](#)



The conference also provided a deliberate deep-dive into **what closing the warning-to-action gap looks like from the government side of the table**. In a scenario exercise, participants were asked: if you had EUR 100 million to invest in data-driven conflict prevention, how would you allocate it across community engagement platforms, rapid response mechanisms, learning and evaluation, data collection and monitoring, decision support for policymakers, or forecasting and predictive analytics? The audience split almost evenly between community engagement, data collection, and rapid response. But the policymakers on the panel pushed back on the framing itself: the priority, they argued, is not choosing between tools but **building systems in which individual tools work together—and investing in partnerships** that reach beyond the early-warning community's own bubble, which are not fundamentally questions of money. The exchange surfaced a harder question: do decision-makers actually need more data? Or do they need less—three concise bullet points rather than a dashboard? That no consensus emerged is itself telling: the answer likely varies by institution, by crisis, and by the individual in the chair—which argues for designing decision support with users, not for them.

→ Explore more: [Early Warning in Action: Leveraging Data and Technology for Violence Prevention](#)

A further lesson emerged from the anticipatory action sessions: **warnings become actionable only when they are embedded in systems that predate the crisis.** The value of an early warning signal lies not in its precision alone but in its connection to pre-agreed triggers, so that no time is lost on ad-hoc decision-making or unclear governance structures. Yet even when these triggers materialize, institutions routinely hesitate. This could look like waiting for the field to converge on a shared reading of the situation, a process that consumes the very lead time the system was built to provide. And once settled, this process becomes resistant to contradicting evidence. Even where hesitation is overcome, current threshold models tend to activate when large numbers are projected to be affected, but what about contexts where only hundreds of people face extreme harm—not enough to trigger the system, but severe enough that inaction erodes the very trust on which future data collection depends? If communities provide information and no action follows, they will not provide it again. Thus, the threshold problem is not only analytical but relational—a failure of response today becomes a failure of data tomorrow.

→ Explore more: [Closing the Action Gap: From Signals to Systems for Anticipating the Impacts of Violence](#)



6. Trust, Accountability, and Political Will

In the transition from early warning to early action, **trust is a more decisive variable than technological capability**. This point was articulated, in varying formulations, across nearly every session: the analytical sophistication of a model is immaterial if the actors expected to operationalize its outputs do not trust the source, the methodology, or the institution conveying the findings. AI is now reshaping the conditions under which institutional trust is established and consequentially, also the speed and scale at which it can be eroded.

A revealing indicator emerged from the conference itself: across three days of programming, **political will was invoked more frequently than AI**. The evidence base on what works in conflict prevention is now substantial. The binding constraint is not analytical or methodological but implementational, and implementation depends on a resource—political will—that is presently in acute short supply. In a growing number of states, the political calculus is being recalibrated by active and escalating conflicts, producing a sustained reorientation toward defense expenditure and a corresponding contraction in peacebuilding commitments. This realignment is reshaping the operating environment of the entire field: budgets are tightening, donor priorities are shifting, and the institutional terms on which prevention work is conducted are being renegotiated, often unfavorably. No analytical instrument, however technically advanced, can substitute for an international system that has functionally deprioritized prevention as a strategic objective.

“For us, the question is not necessarily whether the signal is accurate, but it’s also about who trusts that signal. And with that trust who has authority to act on it and who has the resources to move.”

—*Jerome Galagade*, Head of Quality Assurance, M&E & Knowledge Management at Somali Resilience Program (SomReP), World Vision International

→ Explore more: [Closing the Action Gap: From Signals to Systems for Anticipating the Impacts of Violence](#)

In practice, this means that for policymakers to maintain trust in early warning systems, they need confidence that the underlying sources are credible and the logic is transparent. But **trust alone is insufficient without clarity on who is responsible for acting on the warnings**—a matter of coordination that becomes acutely difficult in dynamic situations where conditions shift faster than institutional mandates can adapt. Clear ownership of specific decisions at specific thresholds is essential if early warning is to translate into early action. But even ownership presupposes something more fundamental: decision space. A perfect forecast delivered to an actor who lacks the institutional mandate, political room, or operational capacity to respond is not a contribution—it is a transfer of moral burden without material consequence. The field must stop asking whether its products are good enough and start asking whether the recipients it has chosen are positioned to act. Where they are not, the answer is not better analysis—it is different partners. This demands a fundamental shift in engagement: away from delivering outputs and hoping they land, and toward sustained dialogue about what decisions need support, whether the analysis on offer is proximate to those decisions, and whether feedback loops exist to correct course when it is not.

→ Explore more: [Early Warning in Action: Leveraging Data and Technology for Violence Prevention](#)

For too long, the field's operating logic was linear: better data produces better models, better models produce more timely forecasts, and more timely forecasts produce action. Beneath this is an assumption that states are essentially rational—that if decision-makers understood how costly war is and how profitable peace is, they would choose accordingly. That assumption has not held. Motivations are shaped by political survival, domestic constituencies, and institutional incentives that do not always privilege human life. Understanding this requires the same analytical rigor the field brings to conflict dynamics but has rarely applied to

decision dynamics—what decision-makers actually want, what they can act on, and what constraints they face that are invisible from outside.

"The data doesn't make reality just by existing—it needs to be tied to choices, and those choices are made by someone. What the UN can usefully do is put that information in front of those decision-makers and make them understand that whichever way they choose to go, it comes with a price. If they choose to delay action, that's a price. If they choose not to act, that comes with a price as well."

—**Anayansi Lopez**, Head of Pillar - Human Rights, Civil Affairs, Peacebuilding, Guidance and Learning, Policy and Best Practice Service (PBPS), Division of Policy Evaluation and Training (DPET), Department of Peace Operations (DPO), United Nations

→ Explore more: [*The Data We Need to Invest in Peace*](#)

Once that ownership is established, it must be matched by **accountability**. When warnings are issued and no substantive response follows, the resulting inaction cannot be treated as a neutral outcome. Non-response is itself a form of decision. Where policymakers decline to act—whether because political will is lacking, because the issue is politically inconvenient, or because institutional attention has shifted elsewhere—those decisions should be systematically documented. The field would benefit from a sustained record of warnings issued, recipients notified, and responses taken or withheld. Such a ledger would be an infrastructure of institutional memory—a way to distinguish warnings that failed analytically from warnings that succeeded analytically but were not acted upon. Without it, the field cannot answer the most basic evaluative question about its own work: whether better warnings have, in fact, produced better outcomes.

One session brought particular clarity to the **political will question** in the domain where it may matter most: **atrocities prevention**. A review of nearly 600 studies by the Simon-Skjodt Center at the United States Holocaust Memorial Museum²² found that effectiveness depends not on the tool—naming and shaming, mediation, sanctions, peace operations—but on the conditions of its use: committed actors, international coordination, multiple tools deployed in combination, and application without political bias. The uncomfortable conclusion is not that the field lacks knowledge of what prevents atrocities but that it lacks the political will to implement what it already knows.

→ Explore more: [What Works in Atrocity Prevention?](#)

Closing this gap requires the same rigor applied to decision dynamics as to conflict dynamics: understanding what decision-makers actually want, what they can act on, and what constraints they face that remain invisible from outside.

→ Explore more: [The Data We Need to Invest in Peace](#)

Even where institutional political will is absent, **communities act anyway**—declaring neutral zones, establishing weapons-free areas, developing local protection strategies, as research from Indonesia and Colombia demonstrates. What they lack is not agency but accessible evidence: which strategies exist, which have worked, and under what conditions.

→ Explore more: [What Works in Atrocity Prevention?](#)

The investment case is straightforward: making that knowledge available enables bottom-up prevention that does not depend on top-down activation.

7. Nothing About Us Without Us: Communities as Partners, Not Data Points

The sharpest intervention of the conference came from those who have lived inside the crises these systems are built to address. The movement of troops, the spiking of prices, the shifting behavior of traders: these are not gaps in global coverage waiting to be filled by better algorithms, they are a fully functioning early warning system, already operational, already trusted, already acted upon by local communities.

“My uncle said no we cannot leave. The traders are still moving. The army barrack where we pass through while going home is not yet blocked. They have not seen ammunitions assembled in certain places there, they have not seen certain checkpoints and curfew. So, things are still ok. I can bet none of that data was available anywhere, in your Headquarters, not even in your country offices. But that was available locally on our ground.”

—John Jal Dak, Executive Director, Youth Social Advocacy Team (YSAT)

→ Explore more: [Closing the Action Gap: From Signals to Systems for Anticipating the Impacts of Violence](#)

Too many high-level early-warning architectures still call for action from capitals, designed by experts who have never set foot in the places they model. **Too many systems treat communities as victims to be counted and data resources to be extracted from—not as active actors with voice, agency, and their own analyses of what is happening to them.** Data is routinely collected from these communities, analyzed elsewhere, and rarely returned in a form that serves the people who generated it. The remark by one speaker that the conference featured far more data than people affected by it was itself a measure of how far the field still has to travel.

The conference surfaced another question: **who actually owns early warning?** The answer, from those operating closest to violence, was unequivocal—and it challenged the architecture most international systems take for granted. In the Democratic Republic of the Congo, communities form their own Local Protection Committees with support from the United Nations Organization Stabilization Mission in the Democratic Republic of the Congo (MONUSCO)—identifying conflict indicators, tracking armed group movements, and generating warnings through Community Alert Networks built on phone hotlines and solar-powered radios that reach places where no one has data or connectivity. In South Sudan, Ukraine, and the Philippines, Nonviolent Peaceforce²³ documented how civilian early warning systems co-exist alongside UN protection architectures, with communities increasingly using preventive tools—information-sharing and coordinated action—to stop violence before it occurs.

→ Explore more: [Who Owns Early Warning? Unpacking the Intersection of International Response and Local Knowledge](#)

In Burkina Faso, Mali, Niger, and Guinea-Bissau, Build Up²⁴ has gone further: communities are embedded from the first step, co-designing questionnaires rather than receiving them, co-analyzing data rather than merely providing it, and retaining autonomy over how findings are used. In South Sudan, Benin, and Burkina Faso, the Sentinel Project²⁵ is developing a framework where communities provide the insight that national and international reporting structurally cannot reach—small-scale violence in remote areas that never makes it beyond the local information ecosystem, and indicators that are simply unobservable from outside, while receiving the broader situational awareness that makes their own preventive actions more effective—closing the information loop in both directions.

→ Explore more: [Seeing Small Signs: Developing a Community-Focused Approach to Violence Early Warning](#)

The common thread is a **shift from communities as data sources to communities as co-owners of the questions asked, the analysis produced, and the decisions that follow.**

But ownership without protection is extraction by another name. The principle voiced across these sessions was direct: **do not collect data you cannot protect, and do not collect data you cannot act on.** Where there is no realistic prospect of response, gathering information only transfers risk to those providing it—particularly when that data implicates governments as perpetrators of violence against their own people. When warnings consistently produce no action,

communities stop reporting altogether, leaving institutions that depend on them progressively blinder. The compounding logic is simple: extraction erodes trust, eroded trust produces silence, and silence degrades the very systems that claimed to serve those communities in the first place. The harder question the conference raised but could not resolve is **whether the greater challenge today lies in producing better warnings or in building institutions willing to act on them.**

→ Explore more: [Who Owns Early Warning? Unpacking the Intersection of International Response and Local Knowledge](#)

Building institutions capable of action starts with building systems worthy of trust. That means investing in languages and translation as core infrastructure, not as an afterthought. It means finally operationalizing a principle this field invokes but has yet to embed in its own design: **nothing about us without us.** Systems built without the communities they claim to serve do not remain neutral—they eventually work against them, whether by design or as the unintended consequence of exclusion hardened into architecture. And it means confronting the ethics of our own supply chains—from the data annotators whose invisible labor underpins our models to the field researchers whose safety we rarely discuss in the same rooms where we celebrate the insights they make possible. If our systems depend on people we never name and never protect, we have not built ethical infrastructure; we have built extractive ones.

→ Explore more: [Data Collection in the Context of Insecurity: Challenges, Limitations and Innovations in the Sahel](#) and [Nothing About Us Without Us: Ten Years of a Community-Owned Early Warning System in Mali \(2016-2025\)](#)

The same holds for the spaces where the field convenes: ensuring that those closest to conflict can be present alongside those who analyze it is a shared responsibility—and one that becomes possible when organizers and funders work together to treat Global South and conflict-affected participation as core to the design of these spaces, rather than as an aspiration to be met when budgets allow.

The conference also brought a distinct lens rarely foregrounded in early warning discussions: **the city as a unit of action.** Much of the field's architecture operates at the national or international level—producing country-level risk scores, engaging state actors, seeking political will from capitals. But violence is experienced locally, and in an increasingly urbanized world, cities are where prevention either succeeds or fails in practice. Unlike national governments constrained by geopolitical complexity, municipal authorities often possess what the field most struggles to find elsewhere: decision space, operational proximity, and the mandate to act. In

the session by Peace in Our Cities,²⁶ panelists examined how urban governments are integrating data and AI into violence prevention—not as a smart-city exercise in optimization, but as a governance challenge requiring inclusivity, participation, and ethical boundaries. The tensions are real: digital divides, surveillance risks, algorithmic bias, and the temptation to prioritize efficiency over trust. But so is the opportunity. Cities are the scale at which early warning and early action can converge within a single institution, and where the gap between data and decision is shortest. The conference pushed beyond "human in the loop" to a harder standard: "community in the loop". A single oversight mechanism cannot carry the accountability that AI-informed decisions demand. What the Peace in Our Cities session demonstrated—residents as governing partners in technology decisions, not beneficiaries of them—is not an aspiration. It is a design requirement.

"AI can process enormous amounts of data quickly and translate complexity into something usable, helping us understand city dynamics and get closer to the root causes of violence. It also lets us zoom in and understand the hyper-local dynamics of violence."

—*Luisa Portugal*, Program Officer, NYU Center on International Cooperation

→ Explore more: [Human-Centered, Technology-Empowered Violence Prevention in Cities](#)

Conclusion: What This Moment Demands

This article captures only a fraction of what the conference produced. Sessions ranged from climate-security modeling and earth observation to nuclear diplomacy datasets; from displacement forecasting and pastoral mobility to macroeconomic approaches to prevention; from data collection under insecurity in the Sahel to live simulations on Syria and Somalia; from demining and arms-risk indicators to new databases on mediation and peace processes. The full breadth—over sixty sessions across three days—is available through the [conference recordings](#). But across that breadth, a set of conclusions converged.

What emerged from Stockholm was not a single recommendation but a set of interdependent conditions: data infrastructure must be funded as a public good, not a grant-dependent project; AI must enhance human judgment, not replace the obligation to exercise it; forecasting models must move beyond national scores toward contextualized, localized meaning; early warning systems must be embedded in pre-agreed triggers and community relationships that make action possible before the crisis, not after the dashboard turns red; trust, once broken by extraction without return, cannot be rebuilt by better visualization; the participation of Global South and conflict-affected communities in the field's own convenings must be enabled as deliberately as the datasets and models produced about them; and gender-sensitive, locally grounded, and ethically accountable approaches are not cross-cutting add-ons but foundations.

But beneath these institutional conditions ran a thread easy to forget in a week of dashboards and benchmarks: behind every spike in violence is a community that, most probably, already knew what was coming. **Storytelling, lived experience, and local knowledge are what turn analytics into understanding and early action.** A line on a chart is a starting point, not an answer.

The field now faces a choice that no model can make for it: whether to keep refining instruments of observation separately, or to connect them into systems where better data, responsible AI, political commitment, and community ownership work as one. The conference showed that the pieces exist.

What remains is the architecture to hold them together—and the willingness to listen before we measure. As reflected by one of co-organizers of this year’s conference:

“That architecture will not be built by any one institution or discipline. It depends on the relationships between those who collect and interpret data, those who develop new tools, those who make policy, those who act in crises, and—most importantly—those who live with the consequences. The work ahead is therefore as much about sustaining and deepening this community as it is about advancing the technologies and systems discussed throughout the conference.”

—Angelica Lindqvist-McGowan, Director of Operations and Outreach, Violence & Impacts Early-Warning System (IEWS), Department of Peace and Conflict Research, Uppsala University

That spirit of collective commitment is what Data for Peace is ultimately built on. The task now is to carry it forward—continuing to connect the people, knowledge, tools, and institutions needed to turn warning into action. We hope the conversation continues well beyond Stockholm, and we invite the community to stay engaged, keep building together, and join us again in Oslo next year.

Endnotes

-
- ¹ Webpage, “The Uppsala Conflict Data Program,” *UCDP*, n.d., <https://ucdp.uu.se/>.
- ² Shawn Davies, Therése Pettersson, Magnus Öberg, “Organized violence 1989–2025, and violent political protests,” *Journal of Peace Research* 63, no. 4 (July 2026): 705–723, <https://academic.oup.com/jpr/article/63/4/705/8703754>.
- ³ <https://www.economicsandpeace.org/wp-content/uploads/2026/06/Global-Peace-Index-2026-Report.pdf>.
- ⁴ Webpage, “Crisis Risk Dashboard,” *UNDP*, n.d., <https://data.undp.org/products/crisis-risk-dashboard>.
- ⁵ Webpage, “UNDP Global Risk Anticipation Data Hub,” *CRAF’d*, n.d., <https://crafd.io/projects/undp-datahub>.
- ⁶ Webpage, “Home,” OHCHR, n.d., https://www.ohchr.org/en/ohchr_homepage.
- ⁷ Data for Peace, “Post-Mortem, Reset or Business as Usual? Conflict Early Warning in a Changing World,” June 24, 2026, YouTube video, <https://www.youtube.com/watch?v=WUI6bDAPQ3o&t=1s>.
- ⁸ Data for Peace, “Closing the Action Gap: From Signals to Systems for Anticipating the Impacts of Violence,” June 24, 2026, YouTube video, <https://www.youtube.com/watch?v=4azWcrvY4w0&t=2s>.
- ⁹ Complex Risk Analytics Fund, “CRAF’d Ecosystem Strategy: How Our Pooled Funding for Crisis Data Builds, Protects, and Nurtures the Global Crisis Data Ecosystem” (New York: United Nations, 2025), https://crafd.io/s/2025_CRAFdEcosystemStrategy_Digital.pdf.
- ¹⁰ Webpage, “Home,” *Centre for Humanitarian Data*, n.d., <https://centre.humdata.org/>.
- ¹¹ Webpage, “Home,” *ACLED Data*, n.d., <https://acleddata.com/>.
- ¹² Data for Peace, “Five Things Users Should Know About Conflict Event Data,” June 24, 2026, YouTube video, <https://youtu.be/127jv2npKWM>.
- ¹³ Webpage, “Home,” *Conflict Forecast*, n.d., <https://conflictforecast.org/>.
- ¹⁴ Webpage, “Home,” *Living Signals*, n.d., <https://www.livingsignals.io/>.
- ¹⁵ Webpage, “Home,” *Institute for Economics and Peace*, n.d., <https://www.economicsandpeace.org/>.
- ¹⁶ Webpage, “Home,” *ACLED Data*, n.d., <https://acleddata.com/>.
- ¹⁷ Webpage, “Joint Research Centre,” *European Commission*, n.d., https://commission.europa.eu/about/departments-and-executive-agencies/joint-research-centre_en.
- ¹⁸ Webpage, “Home,” *VIEWS*, n.d., <https://viewsforecasting.org/>.
- ¹⁹ Branka Panic, “Can Emerging Technologies Lead a Revival of Conflict Early Warning/Early Action? Lessons from the Field,” *NYU Center on International Cooperation*, April 6, 2022, <https://cic.nyu.edu/resources/can-emerging-technologies-lead-a-revival-of-conflict-early-warning-early-action-lessons-from-the-field/>.
- ²⁰ Webpage, “INFORM Warning,” *European Commission*, n.d., <https://drmkc.jrc.ec.europa.eu/inform-index/INFORM-Warning>.
- ²¹ Webpage, “CLEAR Initiative,” *Norwegian Refugee Council*, n.d., <https://clearinitiative.io/>.
- ²² Webpage, “Tools for Atrocity Prevention,” *United States Holocaust Memorial Museum*, n.d., <https://preventiontools.ushmm.org/>.
- ²³ Webpage, “Home,” *Nonviolent Peaceforce*, n.d., <https://nonviolentpeaceforce.org/>.
- ²⁴ Webpage, “Home,” *Build Up*, n.d., <https://howtobuildup.org/>.
- ²⁵ Webpage, “Home,” *Sentinel Project*, n.d., <https://thesentinelproject.org/>.
- ²⁶ Webpage, “Home,” *Peace in Our Cities*, n.d., <https://peaceinourcities.org/>.

Center on International Cooperation

 NYU | WAGNER

105 East 17th Street

New York, NY 10003-2170

cic.nyu.edu | [@nyuCIC](https://twitter.com/nyuCIC)