

# PUBLIC HEALTH UNDER FIRE: AI FOR HEALTH IN CONFLICT ZONES

**Juntao Jiang**

College of Control Science and Engineering  
Zhejiang University  
Hangzhou, 310027, China  
juntaojiang@zju.edu.cn

## ABSTRACT

Conflict zones are characterized by persistent infrastructure disruptions, including prolonged power outages, network failures, and severe limitations in medical equipment, which fundamentally challenge conventional AI-for-healthcare systems. This work reviews the unique obstacles and deployment requirements of AI in such environments and discusses key technical strategies to address them. In particular, we emphasize a sustainable offline-first framework with dual-track MLLM deployment, the integration of portable and wearable AI-enabled sensing devices, fault-tolerant system design, and task-oriented evaluation, aiming to ensure resilient and practical medical support when traditional healthcare systems are partially or fully disrupted.

## 1 INTRODUCTION

Armed conflict remains one of the most severe challenges to global public health. Beyond the immediate toll of casualties, the disruption of health systems and the destruction of essential infrastructure mean that even basic healthcare becomes unattainable, and the health burden escalates across populations (Lancet, 2026). In such settings, otherwise preventable or treatable health conditions, such as trauma-related complications, infectious disease outbreaks, malnutrition, and poorly managed chronic illnesses, rapidly become leading causes of morbidity and mortality. Violence also destroys essential infrastructure for basic living and health, such as water supply and sanitation systems, producing long-term consequences for public health (ICRC, 2024). In Gaza, the health system has deteriorated to an unprecedented level, with most hospitals and clinics damaged or closed, severe shortages of essential supplies and workforce, and surging disease and malnutrition among displaced populations (Al Bakri et al., 2025; WHO, 2024; 2025). Persistent security risks further constrain the movement of health workers and medical supplies, and healthcare facilities and personnel are frequently exposed to violent threats (Makali et al., 2025). For example, in 2024, more than 3,600 attacks on healthcare services were recorded in conflict-affected areas (Mahase, 2025). As such incidents continue to rise, the World Health Organization has warned that attacks on healthcare facilities are becoming “the new reality,” and has emphasized the urgent need to prevent this from becoming the norm, as such violence severely hinders service delivery (Ghebreyesus, 2024).

In recent years, artificial intelligence has broad potential in public health, such as disease surveillance, decision-support systems, and optimization of resource allocation in relatively stable environments, providing new tools for public health practice (Fisher & Rosella, 2022). In such contexts, AI tools have been shown to improve efficiency, expand service coverage, and support data-informed decision-making (Fisher & Rosella, 2022; Panteli et al., 2025). However, applying these technologies directly in conflict-affected areas still faces fundamental challenges. Frequent power outages and network disruptions, shortages of computing resources and equipment, and fragmented governance structures make it difficult for AI systems to operate. In conflict-affected settings, health information systems are often fragmented and poorly coordinated (Ladadwa et al., 2024), thereby undermining the feasibility of data-driven technologies and real-time analytics, which require unified data collection, storage, and sharing. This reality highlights a critical contradiction in the use of AI for public health in conflict settings: a significant gap exists between potential value and actual usability. AI tools designed for conflict zones must assume from the outset that “unreliable infrastructure” is the norm rather than the

exception. Such systems should serve the practical needs of humanitarian organizations and local health authorities, providing decision support for limited personnel under conditions that prioritize safety and ethical compliance, rather than replacing human judgment. Based on this, this paper aims to propose an AI application approach for conflict-affected public health, highlighting the key technologies and design principles that should be prioritized in such environments. By repositioning AI as a tool to enhance system resilience and continuity, this study seeks to offer a feasible and responsible technological pathway for public health in conflict-affected areas.

## 2 APPLICABILITY OF MEDICAL AI IN CONFLICT SETTINGS

### 2.1 CHALLENGES OF AI FOR HEALTHCARE IN CONFLICT ZONES

War has a profound impact on the feasibility and effectiveness of AI applications in healthcare. The challenges of AI in conflict-affected healthcare can be grouped into four main categories, as summarized below; a detailed comparison is provided in Table 1.

- **Infrastructure constraints.** In conflict settings, unstable power and fuel shortages (Dafallah et al., 2023) often prevent medical equipment from operating and make cloud computing inaccessible, while damaged communication infrastructure leads to network outages (Alkhalil et al., 2024) that undermine remote consultations and cloud-based AI. Advanced imaging devices may be destroyed or unrecoverable (Suji et al., 2024), forcing reliance on portable, low-power equipment, which makes traditional AI systems difficult to operate.
- **Data and information flow disruption.** Data continuity and quality deteriorate as record-keeping becomes chaotic and electronic health systems fail. Imaging and clinical data may be missing or unusable, making AI inference unreliable even if models remain available. For example, EMR crashes in Gaza created gaps and delays in documentation, which directly compromised data quality and continuity (Dafallah et al., 2023).
- **Human and organizational constraints.** Conflict environments cause healthcare worker shortages (Dafallah et al., 2023; Ali et al., 2025), exhaustion, and breakdown of specialist collaboration. AI can no longer serve as a narrow specialist tool but must support broader decision-making across the care pathway, increasing functional demands and risk.
- **Security and governance risks.** Healthcare systems in conflict zones face physical destruction, information attacks, and supply-chain disruption. Devices may be damaged or evacuated, and maintenance becomes difficult due to a lack of support and safety concerns. These risks reduce AI feasibility and raise ethical and operational challenges.

### 2.2 CONFLICT ZONES NEED AI FOR HEALTHCARE

Despite the substantial technical, infrastructural, and ethical challenges that constrain conventional AI deployment in conflict-affected settings, these environments may in fact exhibit a heightened need for appropriately designed healthcare AI. Armed conflict simultaneously erodes the medical workforce, collapses specialist-based care, and generates repeated mass casualty incidents that overwhelm human decision-making capacity, while also producing widespread psychological trauma (Ahmed et al., 2024; Hemmeda et al., 2025; Amsalem et al., 2025) in the near absence of mental health professionals. Under such conditions, narrowly scoped, task-specific AI systems become inadequate; instead, generalist, workflow-level AI, such as multimodal large language models (MLLMs), can help preserve a minimal standard of care by supporting end-to-end clinical decision-making, rapid training of inexperienced personnel, stabilization of triage and resource allocation under extreme overload, and low-threshold psychological support when no alternatives exist. In this sense, the very fragility and scarcity that complicate AI deployment in conflict zones also make resilient, safety-aware AI tools potentially indispensable for maintaining healthcare continuity.

## 3 SUSTAINABLE OFFLINE-FIRST FRAMEWORK OF AI FOR HEALTHCARE IN CONFLICT ZONES

**Framework Overview.** In conflict-affected areas such as Gaza, where prolonged blockade and infrastructure damage are the norm, power and internet outages are not isolated incidents but persistent

Table 1: Factors affecting AI applications in normal vs. conflict settings

Factor Category	Normal Environment	Conflict Environment	Impact on AI applications
<b>Power and Energy</b>	Stable electricity supply	Frequent outages and fuel shortages	Servers unavailable; requires low-power and offline AI
<b>Network and Communication</b>	Stable 4G/5G/fiber coverage	Base stations destroyed, network outage, extremely high latency	Remote consultation and cloud inference/data synchronization are infeasible
<b>Hardware Availability</b>	Advanced imaging and monitoring devices widely available	Devices damaged, maintenance difficult, spare parts unavailable	Reliance on portable devices and lightweight AI models
<b>Data Availability and Quality</b>	Structured electronic health records, complete imaging and real-time monitoring data	Missing data, paper records, incomplete or low-quality imaging, interrupted transmission	AI inference inputs are incomplete and noisy; requires missing data handling, and support for low-quality inputs
<b>Personnel and Organizational Structure</b>	Sufficient healthcare workers and clear specialization	Personnel casualties, extreme shift burden, organizational disorder	AI must provide broader decision support and automation; needs generalist capability and full-process assistance
<b>Security and Threats</b>	Data security norms and relatively safe environment	Devices destroyed, data stolen, information warfare	AI systems must be attack-resistant, privacy-preserving, and fault-tolerant
<b>Maintainability</b>	Engineering support and sustainable maintenance	Lack of technicians, hazardous environment	AI models need rapid deployment, self-recovery, and low maintenance overhead

conditions. Therefore, AI systems designed for civilian healthcare and humanitarian organizations must adopt an offline-first operating paradigm from the outset: 1) To cope with prolonged power disruptions, the system should prioritize low-power, modular edge computing units, compressing MLLMs to a scale that can run on NPU or ARM devices, and incorporating energy-aware inference mechanisms that dynamically reduce inference frequency or preserve only essential functions when battery levels are low, thereby ensuring basic intelligent support even in extreme isolation. 2) The system should be integrated with mobile devices to acquire multimodal inputs through mobile phone cameras, handheld or wearable devices (e.g., portable ultrasound, handheld X-ray, vital-sign sensors), and text-based entry, while maintaining information continuity under limited or absent public connectivity through local autonomy and intermittent synchronization: all patient data, triage records, and model outputs should be stored locally in encrypted form to avoid dependence on remote servers; when brief network windows occur, delay-tolerant networking (DTN) can be used to transmit only essential anonymized statistics and updates; and different healthcare sites can exchange critical status information via short-range wireless communication or physical media to avoid “information islands.” 3) Unlike consumer-grade AI, systems for humanitarian institutions should follow institution-level configuration and governance: updates should be managed centrally by hospitals or coordination centers, permissions should be separated by role, and structured audit logs should be generated to support ethical and compliance review, ensuring that AI functions as a controlled humanitarian infrastructure rather than an unconstrained autonomous technology.

**Dual-track Deployment of MLLMs.** In conflict settings, the value of MLLMs lies in their multimodal capabilities: they can interpret and analyze diverse inputs such as imaging, trauma photos, and monitoring data, while integrating text-based medical history and symptom descriptions for comprehensive assessment. They can answer clinical questions, provide diagnostic and treatment

guidance, support triage and on-site emergency decisions, and also help rapidly train inexperienced healthcare workers and offer psychological support, making them a “all-in-one medical assistant” when resources are scarce and specialist systems collapse. To remain usable in conflict-affected environments, deployment should follow a dual-track architecture: when network connectivity is available, the system may connect to cloud APIs or centrally hosted models to access higher-capacity inference; when offline, it must fall back to a fully functional local system. The offline track is the primary focus, but the online option is retained for enhanced performance whenever feasible. Deploying MLLMs under on-device track should prioritize: 1) Lightweight models that can run directly on mobile devices, edge hardware, or ordinary PCs with fast inference, ensuring usability when infrastructure is degraded. For mobile deployment, models with  $\leq 2\text{B}$  parameters are generally the most reliable choice, while ordinary PCs can often support larger models (e.g., 4B–7B) depending on available memory and acceleration; Knowledge distillation and quantization techniques can be used to compress the model to a size that can run on small devices; 2) open-source models (e.g., LLaVA-series (Li et al., 2024; Shu et al., 2024; Zhang et al., 2025; Xu et al., 2025), Qwen-VL series (Wang et al., 2024; Bai et al., 2025; Yang et al., 2025)) to avoid dependence on cloud APIs that become inaccessible during network outages; and 3) generalist capability, since multiple specialized models cannot realistically run on a single device in the field. Therefore, the chosen MLLM must cover a wide range of tasks, including medical image interpretation, clinical Q&A, medication and procedural guidance, and psychological support.

**Portable and Wearable AI-Enabled Devices.** In conflict settings, wearable or handheld imaging devices should prioritize completing essential clinical tasks with a minimum viable hardware configuration: a portable ultrasound or handheld X-ray/digital radiography module, a camera for wound and trauma photography, and vital-sign sensors (heart rate, blood pressure, temperature). The device must also include sufficient local storage, battery capacity, and offline compute to function during power outages and network loss. An embedded lightweight MLLM can then analyze images and monitoring data in real time, provide triage levels, urgent care priorities, basic diagnostic cues, and on-site treatment guidance, and answer clinicians’ key questions, ensuring core medical decision support remains available even when connectivity and specialist systems collapse.

**Fault Tolerance.** In conflict settings, fault tolerance should be designed around three core principles: 1) graceful degradation, where the system automatically scales down to essential functions when power, compute, or network resources drop, ensuring minimum service continuity rather than complete failure; 2) local redundancy, with key data and models stored in multiple local copies and synchronized across nearby devices when connectivity permits, preventing single-point failures from disabling the system; and 3) checkpointing and automatic recovery, where the system periodically saves its operational state and can rapidly restore the latest checkpoint after a reboot or crash.

**Development-Phase Evaluation** During the development phase, evaluation is not aimed at achieving state-of-the-art performance, but at assessing practical usability under conditions that deliberately simulate the real difficulties of conflict-affected environments rather than ideal laboratory settings. In addition to standard accuracy measures, it should adopt task-oriented metrics such as task completion rate, time-to-decision, and system failure rate to better reflect clinical utility in emergency scenarios. System robustness is systematically evaluated by emulating prolonged network outages, severe power and compute constraints, and degraded or incomplete input data, with particular emphasis on whether the system can maintain essential medical functions, provide conservative and interpretable outputs, and degrade gracefully instead of failing abruptly under extreme conditions.

## 4 CONCLUSION

This paper reviews the key challenges and requirements of applying AI to healthcare in conflict zones, and proposes technical directions that can substantially improve the resilience and practicality of medical AI under extreme infrastructure disruption. While continuing technical breakthroughs can further advance the deployment of AI-assisted care in conflict settings, the sustained operation of healthcare systems ultimately still depends on the involvement of the UN and humanitarian organizations, and on warring parties’ adherence to international law. We hope that, with the help of these technologies, people in conflict zones can still receive basic medical care and protection.

## ETHICS STATEMENT

We have ensured that our study follows ethical standards, with no direct involvement of human subjects and no foreseeable risk of harm. We disclose no conflicts of interest or sponsorship. This work is motivated by the urgent need to support healthcare delivery in conflict-affected regions, where civilians face substantial barriers to basic medical services. We focus on technical and system-level contributions, while recognizing the broader humanitarian context in which these challenges arise. Not discussing an explicit stance in the paper does not imply indifference to the atrocities involved. The authors hold clear moral judgments regarding these events; the absence of such discussion reflects respect for academic ethics and cross-cultural readability.

## THE USE OF LARGE LANGUAGE MODELS

We used large language models to refine the writing and support the literature review, and we conducted a careful verification of the content, taking full responsibility for all material presented in this work.

## REFERENCES

- Syed Hassan Ahmed, Aabia Zakai, Maha Zahid, Muhammad Youshay Jawad, Rui Fu, and Michael Chaiton. Prevalence of post-traumatic stress disorder and depressive symptoms among civilians residing in armed conflict-affected regions: a systematic review and meta-analysis. *General Psychiatry*, 37(3):e101438, 2024.
- Deema Al Bakri, Yousef Khader, Rana Khatib, Fawzi Al-Hammouri, Mousa Aabed, Yehia Abed, Yousef Zureikat, Omar Lattouf, Ahmad Shatat, Abdullatif Husseini, et al. The war on gaza and its impact on public health: challenges and pathways to recovery. *Frontiers in Public Health*, 13: 1664850, 2025.
- Shayan Ali, Bilal Irfan, Wesam Abdeljaber, Elias Nasser, Muaaz Wajahath, Mosab Nasser, and Khaled J Saleh. Digital health in humanitarian crises: A case study of gaza and the west bank. *Digital Health*, 11:20552076251365010, 2025.
- Munzer Alkhalil, Aula Abbara, Caroline Grangier, and Abdulkarim Ekzayez. Ai in conflict zones: the potential to revitalise healthcare in syria and beyond. *BMJ Global Health*, 9(8):e015755, 2024.
- Doron Amsalem, Shilat Haim-Nachum, Amit Lazarov, Yossi Levi-Belz, John C Markowitz, Maja Bergman, Alexandra Klein Rafaeli, Lisa A Brenner, Nitsa Nacasch, Milton L Wainberg, et al. The effects of war-related experiences on mental health symptoms of individuals living in conflict zones: a longitudinal study. *Scientific reports*, 15(1):889, 2025.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Alaa Dafallah, Osman KO Elmahi, Maisoon Elbukhari Ibrahim, Rania Elfatih Elsheikh, and Karl Blanchet. Destruction, disruption and disaster: Sudan's health system amidst armed conflict. *Conflict and health*, 17(1):43, 2023.
- Stacey Fisher and Laura C Rosella. Priorities for successful use of artificial intelligence by public health organizations: a literature review. *BMC Public Health*, 22(1):2146, 2022.
- Tedros Adhanom Ghebreyesus. Attacks on health are becoming the new reality; we must stop this becoming the norm. <https://www.who.int/news-room/commentaries/detail/attacks-on-health-are-becoming-the-new-reality--we-must-stop-this-becoming-the-norm>, 2024.
- Lina Hemmeda, Alaa S Ahmed, Gasm Alseed Abdelmonim Gasm Alseed Fadlalmoula, Doaa Rabeie Hassan, Mohammed Haydar Awad, Gotof Salim, Abeer Osman Mukhtar Elabid, Alaa Azhary mohammed Ali, Asmaa O Algasim, Ethar Hamad, et al. Displacement-related stressors in a sudanese war-affected community; identifying the impact of war exposure and ongoing stressors on trauma symptom severity: a national multi-center cross-sectional study. *BMC Public Health*, 25(1):1325, 2025.

- ICRC. Urban warfare and violence. <https://www.icrc.org/zh/law-and-policy/urban-warfare-and-violence>, 2024.
- Reem Ladadwa, Mahmoud Hariri, Muhammed Mansur Alatras, Yasir Elferruh, Abdulhakim Ramadan, Mahmoud Dowah, Yahya Mohammad Bawaneh, Wassel Aljerk, Preeti Patel, Abdulkarim Ekzayez, et al. Health information management systems and practices in conflict-affected settings: the case of northwest syria. *Globalization and Health*, 20(1):45, 2024.
- The Lancet. No health without peace. *The Lancet*, 2026.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
- Elisabeth Mahase. Conflict zones: Assaults on healthcare reached “new levels of horror” in 2024, finds damning report, 2025.
- Samuel Lwamushi Makali, Okechi Ogueji, Leonie Tax, and Larissa Fast. What works and who knows? the importance of community and health worker perspectives in defining solutions to attacks on healthcare. *Conflict and Health*, 19(1):46, 2025.
- Dimitra Panteli, Keyrellous Adib, Stefan Buttigieg, Francisco Goiana-da Silva, Katharina Ladewig, Natasha Azzopardi-Muscat, Josep Figueras, David Novillo-Ortiz, and Martin McKee. Artificial intelligence in public health: promises, challenges, and an agenda for policy makers and public health institutions. *The Lancet Public Health*, 10(5):e428–e432, 2025.
- Fangxun Shu, Yue Liao, Le Zhuo, Chenning Xu, Lei Zhang, Guanghao Zhang, Haonan Shi, Long Chen, Tao Zhong, Wanggui He, et al. Llava-mod: Making llava tiny via moe knowledge distillation. *arXiv preprint arXiv:2408.15881*, 2024.
- Trisha Suji, Richard Sullivan, and Gemma Bowsher. Radiology in conflict: scoping review. *Conflict and Health*, 18(1):8, 2024.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- WHO. Children’s lives threatened by rising malnutrition in the gaza strip. <https://www.who.int/news/item/19-02-2024-children-s-lives-threatened-by-rising-malnutrition-in-the-gaza-strip>, 2024.
- WHO. Health system at breaking point as hostilities further intensify in gaza, who warns. <https://www.who.int/news/item/22-05-2025-health-system-at-breaking-point-as-hostilities-further-intensify--who-warns>, 2025.
- Guowei Xu, Peng Jin, Ziang Wu, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let vision language models reason step-by-step. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2087–2098, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Shaolei Zhang, Qingkai Fang, Zhe Yang, and Yang Feng. Llava-mini: Efficient image and video large multimodal models with one vision token. *arXiv preprint arXiv:2501.03895*, 2025.