



Artificial intelligence in Acute Care Surgery: integrating new capabilities into a complex adaptive system

Philip S. Barie^{1,2}, Divya Kewalramani^{2,3,4}, Michael F. Rotondo⁵, Thomas M. Scalea^{6,7}, Mayur Narayan^{2,3,4}

Keywords:

Acute Care Surgery, artificial intelligence, complex adaptive systems, implementation, sociotechnical systems, surgical systems

Citation:

Barie PS, Kewalramani D, Rotondo MF, Scalea TM, Narayan M. Artificial intelligence in Acute Care Surgery: integrating new capabilities into a complex adaptive system. *Art Int Surg*. 2026;6:494-508. <https://dx.doi.org/10.20517/ais.2026.62>

Received: 8 Jul 2026

First Decision: 4 Aug 2026

Revised: 13 Sep 2026

Accepted: 21 Sep 2026

Published: 30 Sep 2026

Academic Editors:


Takeaki Ishizawa, Andrew Gumbs

Copy Editor:

Tong Wang

Production Editor:

Tong Wang



Abstract

Acute Care Surgery (ACS) emerged not as a single innovation, but as the progressive integration of operative discipline, intensive care, trauma systems, biologically informed resuscitation, and regionalized care infrastructure into a layered, complex, adaptive system. Artificial intelligence (AI) now represents the next major technological layer being introduced into this mature sociotechnical environment. This narrative review examines ACS from a systems perspective and argues that the success or failure of AI integration requires both algorithmic performance and how effectively these technologies couple with existing workflows, operational structures, feedback mechanisms, and clinical decision-making processes. Historical developments in ACS - including intensive care evolution, trauma system regionalization, damage control philosophy, and prior failures of reductionist biologic interventions - provide important lessons for understanding the opportunities and risks of AI deployment in high-acuity surgical care. Current AI

¹Division of Trauma, Burns, Acute and Critical Care, Department of Surgery, Weill Cornell Medicine, New York, NY 10065, USA.

²Artificial Intelligence for Comprehensive Care, Education, and Sustainability in Surgery (AiCCES) Consortium, Rutgers Robert Wood Johnson Medical School, New Brunswick, NJ 08901, USA.

³Division of Acute Care Surgery, Department of Surgery, Rutgers Robert Wood Johnson Medical School, New Brunswick, NJ 08901, USA.

⁴Rutgers Acute Care Surgery Laboratory, Rutgers Robert Wood Johnson Medical School, New Brunswick, NJ 08901, USA.

⁵Division of Acute Care Surgery, Department of Surgery, University of Rochester School of Medicine and Dentistry, Rochester, NY 14642, USA.

⁶Program in Trauma, University of Maryland Medical School, Baltimore, MD 21201, USA.

⁷R. Adams Cowley Shock Trauma Center, University of Maryland Medical System, Baltimore, MD 21201, USA.

Correspondence to: Dr. Divya Kewalramani, Division of Acute Care Surgery, Department of Surgery, Rutgers Robert Wood Johnson Medical School, New Brunswick, NJ 08901, USA. E-mail: divya.kewalramani@rutgers.edu

applications in ACS include augmented diagnostics, continuous physiologic surveillance, predictive analytics, operative video analysis, workflow optimization, and educational assessment. However, important issues remain, including alert fatigue, automation bias, implementation failure, inequitable model generalizability, and the risks of introducing reductionist computational paradigms into complex adaptive systems. A systems-oriented framework emphasizes that AI tools must be evaluated not as isolated products, but as components of broader clinical, organizational, and educational ecosystems. ACS, with its longstanding experience integrating layered technologies and coordinating multidisciplinary care under conditions of uncertainty, may represent one of the most informative environments in which to study responsible AI implementation in surgery.

INTRODUCTION

Acute Care Surgery (ACS) did not emerge as a single innovation, but rather as the consolidation of successive systems - operative discipline, sterile technique, structured intensive care, organized trauma networks, and biologically informed resuscitation - each layered atop the last and each only partially superseding its predecessor. ACS as practiced today is therefore best understood not as a specialty in the conventional sense but as a system of systems: a coordinated assembly of human expertise, infrastructure, protocols, and information flows working under time pressure with incomplete information, fluctuating physiology, and fragile (or broken) patients. In a complex adaptive system, interacting components respond to one another and to changing conditions, producing system behavior that cannot be predicted reliably from the properties of individual components alone. This framing matters because the next layer of innovation, artificial intelligence (AI), will be evaluated and either accommodated or rejected by exactly that pre-existing architecture.

The Special Issue of *Artificial Intelligence Surgery* to which this manuscript contributes asks how AI will reshape ACS. The question is generative, but it is also incomplete. AI cannot be considered apart from the system infrastructure into which it is introduced, any more than a pulmonary artery catheter or lung-protective ventilation could be considered apart from the intensive care units (ICUs) utilizing them. New technologies in ACS have historically succeeded not by replacing older systems but by interfacing with them - extending their reach, sharpening their measurements, or closing feedback loops previously left open. They have failed when introduced as solitary interventions into networks whose dynamics were not appreciated. The most instructive example remains the cycle of single-mediator anti-cytokine trials in sepsis, where biologically plausible interventions failed in heterogeneous critically ill populations because the host response was networked, redundant, and temporally dynamic^[1,2]. The reductionist gesture - isolate one variable, suppress it, expect linear benefit - did not survive contact with the system.

A systems perspective on AI in ACS therefore begins with three propositions. First, the value of any new tool is determined less by its standalone performance than by how it couples to the existing system. Second, the existing system in ACS is already complex and adaptive, with multiple interacting feedback loops, layered timescales, and limited slack capacity. Third, the failure modes of introducing AI into such a system are predictable from the historical record: noise overwhelming signal (low S:N), well-intentioned interventions producing iatrogenic load, and rigorous local validation that fails to generalize across institutional contexts. The argument that follows is that AI will succeed in ACS to the extent that it is treated as a new layer in a sociotechnical system^[3,4] rather than as a discrete clinical product. A recent scoping review of AI applications in ACS confirmed the breadth of activity in the field but also documented marked heterogeneity in validation, integration depth, and reporting standards across published work^[5], reinforcing the need for a systems-oriented framing.

The framing has implications for how the field should evaluate, deploy, and regulate AI. A model with strong discrimination is a component, not a system. The same model embedded in differing workflows, escalation

pathways, staffing patterns, feedback structures, and rescue interventions will produce different outcomes. The relevant unit of analysis is the ensemble of practices that surrounds the tool. This manuscript traces how the existing ACS architecture came to look the way it does, where AI is being introduced into it, where the integration is most likely to fail, and what governance and educational structures will determine whether the next coupling proves durable.

Unlike reviews that catalog AI applications or compare algorithmic performance, the distinctive purpose of this review is to examine AI integration through the historical and systems architecture of ACS. Literature was identified through targeted searches of PubMed and reference lists of relevant publications, supplemented by the authors' knowledge of the ACS and surgical AI literature. Sources were selected for their relevance to the historical development of ACS, current and emerging AI applications, sociotechnical systems, and AI implementation. The review was narrative rather than systematic; formal study screening, evidence grading, and meta-analysis were not performed.

THE EXISTING SYSTEM: ACS AS A LAYERED, COMPLEX, ADAPTIVE ARCHITECTURE

The historical narrative of ACS is most often told as a sequence of advances. A systems reading reframes it as a sequence of couplings - each new capability requiring, and in turn enabling, the next layer of infrastructure.

The first such coupling was between operative discipline and infection prevention. Anesthesia made operative intervention possible^[6]; antisepsis and asepsis made it survivable^[7,8]. By the early twentieth century, Halsted had synthesized these principles into an operative philosophy of meticulous hemostasis, gentle tissue handling, and layered closure that allowed surgical ambition to expand without proportionate increases in physiologic insult^[9]. The system permitted larger operations but, in doing so, exposed a new problem: postoperative instability that could not be addressed at the operating table.

That problem required a second coupling between operative care and structured physiologic observation. Dandy's three-bed neurosurgical recovery unit at The Johns Hopkins Hospital in 1923 established the principle that, with vigilant surveillance, postoperative instability was dynamic and detectable^[10]. World War II-era shock wards extended the principle to active physiologic management^[11]. Ibsen's polio-era ventilatory care in 1950s Copenhagen and the maturation of postoperative recovery rooms produced the modern intensive care environment^[12], which in turn enabled survival of major operations and injuries that would previously have been terminal.

Survival exposed a third problem. As resuscitation, transfusion, and operative technique improved, patients began to survive the initial insult only to die later from progressive organ dysfunction. The recognition of multiple organ failure as the dominant cause of death in surgical ICUs reframed shock as a biologic rather than a purely hemodynamic problem^[13,14]. The cytokine era and its corollaries - gut barrier failure, ischemia-reperfusion injury, the compensatory anti-inflammatory response, and ultimately the persistent inflammation, immunosuppression, and catabolism syndrome (PIICS) - established that severe injury produced a sustained, coordinated transcriptomic response across immune, endothelial, autonomic, and metabolic networks^[15,16]. Investigators at multiple centers framed multiple organ dysfunction syndrome (MODS) as network instability within a complex adaptive system, susceptible to recalibration but not to single-target correction^[17,18].

A fourth coupling organized this biologic understanding at the level of geography. Regional trauma systems matched injury severity to operative and post-operative care capabilities, and verification programs through the American College of Surgeons Committee on Trauma standardized expectations across institutions^[19]. National registries - the National Trauma Data Bank and the Trauma Quality Improvement Program -

Table 1. Successive system couplings in the evolution of ACS

Existing capability	Failure mode revealed	New coupling introduced	System-level result
Operative intervention enabled by anesthesia and antisepsis	Postoperative instability and physiologic deterioration beyond the operating room	Surgery coupled with structured postoperative observation and intensive care ^[10-12]	Continuous physiologic surveillance and support of critically ill surgical patients
Advanced operative and resuscitative capabilities	Survival of the initial insult followed by delayed MODS	Hemodynamic resuscitation coupled with biologic and systems-level understanding of injury and critical illness ^[13-18]	Recognition of MODS, network physiology, and complex adaptive responses to injury
Institutional trauma expertise	Geographic mismatch between injury severity and available resources	Local trauma care coupled with regionalized trauma systems and verification programs ^[19-21]	Coordinated trauma networks and improved access to specialized care
Trauma surgery and surgical critical care	Fragmented emergency surgical coverage and workforce shortages	Trauma surgery, emergency general surgery, and surgical critical care integrated into the ACS model ^[22-24]	Unified service structure and sustainable workforce development
Data-rich digital health infrastructure	Information volume exceeding human cognitive capacity	Clinical care coupled with artificial intelligence and advanced analytics ^[32]	Potential for continuous data interpretation, decision support, and augmentation of clinical and operational performance

ACS: Acute Care Surgery; MODS: multiple organ dysfunction syndrome.

closed the loop by permitting risk-adjusted benchmarking and population-level feedback^[20,21].

The next coupling was administrative. Declining operative trauma volumes in many regions, expanding ICU responsibilities for trauma surgeons, and the persistent fragmentation of emergency general surgery produced a workforce mismatch^[22]. The ACS model consolidated trauma, emergency general surgery, and surgical critical care into a single service identity, and the formalization of a two-year fellowship through the American Association for the Surgery of Trauma created a sustainable training pipeline^[23,24].

The result is a discipline whose distinguishing feature is precisely its layered character. An acute care surgeon practices within an architecture that includes prehospital triage protocols, emergency department workflows, radiology pipelines, blood bank logistics, operating room (OR) scheduling, ICU infrastructure, infection control programs, registry-based quality improvement, and educational governance. Each layer was developed to address a specific failure mode of the prior layer, and each retains residual fragility. It is into this architecture that AI is being introduced [Table 1].

THE INFORMATION SUBSTRATE: FROM PAPER TO CONTINUOUS STREAMS

Before considering AI specifically, it is necessary to recognize the substrate on which it depends. The digital transformation of clinical information over the last quarter century has been at least as consequential to ACS as any single therapeutic advance, and it is the technology layer that most directly conditions what AI can and cannot do.

Earlier paper records were fragmented and vulnerable to transcription error, particularly in medication ordering. Computerized physician order entry and electronic prescribing standardized legible orders and enabled automated dose and interaction checking^[25]. Picture archiving and communication systems (PACS) eliminated film-based logistical bottlenecks and made imaging available simultaneously in the emergency department, OR, and ICU^[26]. Secure point-of-care medication dispensing created auditable chains linking prescribing, dispensing, and administration. Operative case logs, simulation-based assessments, and tele-ICU platforms extended digital connectivity into surgical education and into geographically distant or under-resourced hospitals^[27,28].

These developments were valuable in their own right. They were also, in retrospect, the prerequisite for everything that followed. Modern ICUs generate continuous, high-density physiologic data streams that no clinician can monitor comprehensively in real time^[29]. The pulmonary artery catheter, once the dominant symbol of invasive hemodynamic monitoring, has been substantially supplanted by pulse contour analysis, bedside ultrasound, and continuous noninvasive sensors. Registries permit risk-adjusted benchmarking at population scale. Trauma video review enables retrospective analysis of resuscitation performance. The substrate is now data-dense to a degree that the architects of the early ICU could not have anticipated.

Two important warnings emerged from this transition. The first was alert fatigue. The proliferation of automated electronic notifications - medication interactions, laboratory abnormalities, protocol reminders, sepsis screening flags - produced a phenomenon in which clinicians became habituated to frequent warnings and dismissed them increasingly^[30]. The S:N of clinical decision support became itself a determinant of its effectiveness, and in many implementations it became negative. The second warning was that automation at the data layer did not automatically translate into improved outcomes. A system can be data-rich and decision-poor. The intervening step - turning data into actionable judgment under time pressure - remained the constraint, and remains so today. AI is being introduced and directed precisely at this constraint.

AI AS A NEW LAYER IN A MATURE SOCIOTECHNICAL SYSTEM

The argument for AI in ACS is therefore not that it provides a new source of data; it does not. It is that machine-learning systems can ingest the existing data substrate and extract patterns at a scale and temporal resolution that humans cannot, and that they can do so continuously rather than episodically^[31,32]. Several application domains illustrate how this layer interfaces with the existing architecture. These applications span a broad spectrum of translational maturity, from retrospective or simulation-based studies to tools undergoing prospective evaluation or clinical deployment; their inclusion here illustrates points of interaction with the ACS system but does not imply equivalent levels of clinical validation or demonstrated efficacy [Figure 1].

Diagnostic augmentation is the most mature application. Convolutional neural networks for trauma imaging can flag hemorrhage, pneumothorax, or solid-organ injury within seconds of acquisition, prioritizing radiologist attention and supporting trauma team triage decisions^[33]. Importantly, these tools do not replace the existing PACS-based workflow; they are an inline layer that interprets the same images, in the same display environment, with the same downstream effect on operative and resuscitative decision-making. “Enhancements” succeed when integrated and fail when bolted on.

Risk prediction is the second domain. Traditional physiologic scoring systems - Acute Physiology and Chronic Health Evaluation (APACHE), Sequential Organ Failure Assessment (SOFA), Multiple Organ Dysfunction Score (MODS), the Trauma and Injury Severity Score (TRISS) - were extraordinary epistemic achievements that allowed organ dysfunction and injury severity to be quantified at the population level^[34-36]. Their structural limitation was that they depended on a small number of input variables, were largely static, and reflected the modeling conventions of their era. Machine-learning models trained on full electronic health record (EHR) and registry data have demonstrated improved discrimination for postoperative complications, sepsis onset, surgical site infection, and mortality, and have the additional property of updating predictions continuously as new data arrive^[37,38]. Whether this translates into outcome improvement depends not on the model’s discrimination in isolation, but on how its outputs couple to clinical action.

Surveillance and deterioration detection constitute a third domain. Continuous physiologic monitoring in the ICU generates patterns of variability - heart rate variability, respiratory rate dynamics, perfusion indices -




Figure 1. Systems architecture for AI integration in ACS. AI capabilities - diagnostic augmentation, predictive analytics, physiologic surveillance, operative guidance, operational management, and education and assessment - are introduced into the established ACS architecture of clinical teams and workflows, ED, OR, and ICU infrastructure, digital records and data streams, regional trauma networks, registries and quality-improvement systems, and training and governance. Adequate algorithmic performance is foundational, but translation into clinical benefit also depends on external and local validation, workflow integration, human oversight and accountability, management of alert burden and cognitive fit, equity and generalizability, training, and ongoing performance monitoring. Outcomes should be evaluated at the clinical and system levels, with monitoring results feeding back into tool and workflow redesign. ACS: Acute Care Surgery; AI: artificial intelligence; ED: emergency department; EHR: electronic health record; ICU: intensive care unit; OR: operating room.

whose nonlinear interactions are difficult for humans to interpret in aggregate but tractable for machine-learning models. Reduced heart rate variability, demonstrated decades ago in trauma populations, presaged the concept that autonomic regulation reflects network coherence, and that loss of coherence is itself prognostic^[39]. AI-enabled surveillance extends this insight by integrating multivariable streams into a single deterioration signal, so that detection precedes overt physiologic collapse rather than reacting to it. The clinical analog is the same shift that produced postoperative recovery rooms and shock wards a century earlier, scaled to the granularity of contemporary sensor networks.

Operative video analysis is a fourth domain, and an instructive one. Computer-vision systems trained on annotated operative video can recognize procedural phases, identify critical anatomy, and flag deviations from the expected operative sequence^[40]. The GoNoGoNet effort for laparoscopic cholecystectomy is an example of how such systems can be deployed across many institutions to support intraoperative decision making and education^[41,42]. The point is not that the system replaces the surgeon. It is that operative video, previously a teaching artifact reviewed retrospectively, becomes a real-time substrate for closed-loop feedback - an extension of the same principle that animated Theodor Billroth's systematic postoperative recordkeeping in the nineteenth century, now operating at frame rate. The same data substrate also supports noninvasive assessment of operative performance across institutions, enabling external benchmarking of intraoperative behavior in the same way that risk-adjusted outcome registries enabled external

Table 2. Current and emerging artificial intelligence applications across the ACS ecosystem

Application domain	Representative AI applications	Primary data sources	Potential system-level benefit
Diagnostic augmentation	Automated detection of hemorrhage, pneumothorax, solid-organ injury, and other imaging findings ^[33]	Computed tomography, radiographs, ultrasound, and other imaging modalities	Faster diagnosis, prioritization of critical findings, and improved triage
Predictive analytics	Prediction of sepsis, surgical site infection, complications, mortality, and resource utilization ^[37,38]	Electronic health records, trauma registries, physiologic data streams	Earlier intervention and more individualized risk assessment
Physiologic surveillance	Detection of clinical deterioration and physiologic instability before overt decompensation ^[39]	Continuous bedside monitoring, wearable sensors, laboratory data	Earlier recognition of evolving critical illness and improved situational awareness
Operative guidance and assessment	Procedural phase recognition, critical anatomy identification, technical skill assessment, and video-based feedback ^[40-43]	Operative video and image data	Enhanced intraoperative decision support, education, and performance improvement
Closed-loop physiologic control	Adaptive management of mechanical ventilation, vasoactive medications, and organ support technologies ^[44]	Continuous physiologic monitoring integrated with therapeutic devices	More responsive physiologic optimization and reduced clinician workload
Operational management	Operating room scheduling, resource allocation, staffing optimization, and patient flow management ^[45]	Administrative, scheduling, and institutional operational datasets	Improved efficiency, throughput, and resource utilization
Education and competency assessment	Automated feedback generation, simulation assessment, individualized learning support, and competency tracking ^[46,47]	Simulation platforms, operative video, educational assessments	More scalable, objective, and personalized surgical education

ACS: Acute Care Surgery; AI: artificial intelligence.

benchmarking of perioperative results. Recent work distinguishing technical performance from operative decision-making in video-based assessment further underscores that intraoperative behavior is multidimensional, and that automated assessment systems will need to disentangle these dimensions to provide actionable feedback rather than a single composite score^[43].

Closed-loop physiologic control is a fifth domain. The intensive care pharmacopeia matured through the 1980s and 1990s toward titratable infusion therapy guided by invasive monitoring and serial laboratory assessment. Machine-learning integration with mechanical ventilators, vasoactive infusion systems, and renal replacement therapy promises continuous adaptive control rather than intermittent adjustment^[44]. These systems extend the historical trajectory of intensive care - from observation, to support, to deliberate physiologic shaping - and represent its logical next stage.

Operational and educational applications are a sixth domain. AI-assisted OR scheduling can incorporate predicted operative duration, acuity, and resource availability, addressing the persistent tension between emergent and elective demand for OR access^[45]. Natural-language processing can reduce the documentation burden that has become a substantial driver of clinician workload, generating draft notes from ambient capture for clinician review and endorsement^[46]. AI-enabled video review can provide individualized feedback to trainees, and machine-learning assessment of feedback quality itself has been shown to predict technical skill improvement in surgical learners^[47]. The latter is a particularly clean example of AI augmenting an existing system - surgical education - rather than supplanting it [Table 2].

INTERFACE FAILURES AND THE DISCIPLINE OF SYSTEMS THINKING

The historical record of ACS contains several lessons that bear directly on the integration of AI. Each is a variation on the same underlying point: complex adaptive systems do not yield to reductionist interventions, and well-engineered local successes do not automatically aggregate to system-level improvement. Documented experience with clinical AI deployment remains limited, particularly in ACS; consequently,

many of the most important implementation failures remain prospective risks rather than mature bodies of empirical evidence.

The first lesson comes from the anti-cytokine hypothesis. Through the 1990s, biologically plausible monoclonal antibodies and receptor antagonists were developed against endotoxin, tumor necrosis factor, interleukin-1, and other chemokines in the expectation that selective interruption of inflammatory mediators would attenuate organ dysfunction. Large trials failed repeatedly^[1,2]. The host response was nonlinear, biologically redundant, temporally dynamic, and variably expressed across patients. Suppression of one mediator did not reliably reverse established dysfunction. The analogous risk for AI is a generation of single-variable models that perform well in retrospective validation but fail to produce outcome benefit when deployed prospectively, because the clinical system into which they are introduced absorbs, dilutes, or redirects their effect. Predictive accuracy is necessary but not sufficient. What matters is the action loop to which the prediction couples. A sepsis alert that fires accurately but cannot be acted upon within the staffing structure of a community hospital, or a deterioration model whose alarm is delivered to a clinician with no decision-rights to escalate^[48], is a model that has been engineered without regard to the system that surrounds it.

The second lesson is alert fatigue. The history of clinical decision support is, in part, the history of well-intentioned signals becoming noise. AI-generated alerts, in the absence of end-user co-design, disciplined thresholding, and contextual relevance, will produce the same outcome. The technical fix is not better algorithms in isolation; it is a sociotechnical design process that treats alert burden as an outcome to be measured and minimized rather than an unavoidable cost of safety^[3,4]. The analog at the biologic level is the recognition that the inflammatory cascade is not solved by amplifying or suppressing any single signal; it is balanced by the compensatory anti-inflammatory response, and disequilibrium in either direction produces harm. Alert systems require an analogous regulatory loop.

The third lesson is the resuscitation correction. Early aggressive crystalloid administration, championed for decades as physiologically conservative, was eventually recognized as a contributor to tissue edema, endothelial disruption, dilutional coagulopathy, and pulmonary compromise, and was reframed as a biologic perturbation rather than a benign, salubrious undertaking^[49]. The principle generalizes. Any intervention - pharmacologic, mechanical, or computational - introduced into a complex adaptive system has non-monetary costs that may not be visible at the level of the immediate target parameter. AI deployments need surveillance for their second-order effects, which may include automation bias, deskilling of trainees who never learn (or attempt) to interpret unprocessed data, workflow distortion as clinicians reorient their attention toward the algorithm's outputs, and shifts in documentation behavior that change on what the next generation of models will be trained.

The fourth lesson is the failure of supranormal oxygen delivery. The pursuit of supranormal oxygen delivery targets, animated by the plausible idea that augmenting global oxygen transport to "repay" an oxygen "debt" incurred during shock and hypoperfusion would prevent or reverse organ failure, did not improve survival^[50]. Macro-circulatory optimization was insufficient because microcirculatory integrity and biologic regulation were the important determinants. The analog for AI is the temptation to optimize a single quantitative target - readmission rate, ICU length of stay, surgical site infection incidence - without attending to the regulatory loops that determine outcomes more broadly. Local optimization can degrade global performance; ACS has learned this lesson at considerable cost in adjacent domains.

A fifth, less-often articulated lesson comes from damage control^[51]. Damage control surgery and damage control resuscitation succeeded not because either component was independently optimal but because their integration acknowledged that the patient was a system whose physiologic capacity to tolerate definitive

Table 3. Historical lessons from ACS relevant to artificial intelligence implementation

Illustrative ACS experience	Systems lesson learned	Analogous risk in artificial intelligence deployment
Failure of anti-cytokine therapies in sepsis and critical illness ^[1,2]	Complex adaptive systems rarely respond predictably to single-target interventions	High-performing predictive models may fail to improve outcomes if disconnected from clinical workflows and action pathways
Alert fatigue associated with clinical decision-support systems ^[3,4,30]	Excessive signaling degrades attention and reduces responsiveness to important alerts	Poorly designed AI notification systems may overwhelm clinicians and diminish trust in decision support
Recognition of harms associated with aggressive crystalloid resuscitation ^[49]	Interventions may produce important downstream effects beyond their intended targets	AI implementation may generate automation bias, workflow distortion, documentation changes, and trainee deskilling
Failure of surrogate endpoint optimization to improve outcomes ^[50]	Optimization of individual metrics does not necessarily improve overall system performance	Pursuit of isolated performance targets may degrade broader patient-, clinician-, or system-level outcomes
Success of staged damage control surgery and resuscitation ^[51]	Staged, adaptive approaches outperform rigid definitive strategies in environments characterized by uncertainty	AI systems should communicate uncertainty, evolving risk, and alternative trajectories rather than provide overly deterministic recommendations

ACS: Acute Care Surgery; AI: artificial intelligence.

repair had a finite envelope. The systems lesson is the value of staged, reversible interventions in environments of high uncertainty. AI systems that produce confident, one-off recommendations in conditions where the underlying physiology is shifting hour-to-hour are likely to misalign with how acute care surgeons actually reason. Systems that present probabilistic, time-resolved trajectories with explicit uncertainty are more likely to integrate with the discipline's native decision style.

A sociotechnical framing addresses these failure modes directly^[3,4,52]. The Salwei and Carayon work-system framework^[3], the McCradden intervention-ensemble concept^[53], and adjacent implementation-science approaches share a common insight: an AI tool is one element of a larger work system that includes people, tasks, technologies, physical environment, and organizational context. The tool integrates only if it fits the context, and the relevant unit of evaluation is the ensemble, not the model. For ACS, this means that the introduction of any AI system should be preceded by explicit workflow analysis, end-user engagement, articulation of the desired target state, prospective monitoring, and pre-specified de-implementation criteria. The same disciplines that distinguish a well-functioning trauma quality improvement program from a registry without consequence apply here.

The implementation-science maturity that ACS has developed in other domains is directly applicable [Table 3]. The Surviving Sepsis Campaign demonstrated that coordinated implementation could accelerate adoption of evidence-based care across diverse health systems and reduce mortality^[54,55]. The same disciplines - guideline and bundle development, performance measurement, audit and feedback, contextual adaptation - will be required to translate algorithmic performance into population benefit.

GOVERNANCE, EQUITY, AND WORKFORCE IN AN AI-INTEGRATED ACS

A systems perspective also clarifies that AI deployment in ACS is a governance problem, not only a technical one. Among several interdependent considerations, three merit particular attention [Table 4].

The first is data stewardship and validation to mitigate algorithmic bias. Models trained on data from large urban academic centers may underperform in rural, community, or international settings whose patient populations, comorbidity profiles, personnel, equipment, and workflow constraints differ from the training distribution. Trauma demographics, surgical infection epidemiology, and resource availability vary substantially across regions, and the consequences of generalization failure in acute care are immediate. Bias

Table 4. Systems-level requirements for responsible artificial intelligence integration in ACS

Domain	Key question	ACS analog	Emergent system capability
Data stewardship	How are data sourced, curated, governed, maintained, and authorized for AI development and deployment?	NTDB, TQIP, institutional registries, and quality-improvement databases ^[19-21]	Reliable, representative, and sustainable data ecosystems that support trustworthy AI systems
Validation	Does algorithm performance generalize across institutions, patient populations, and care environments?	Trauma center verification and external benchmarking ^[19-21]	Trustworthy performance across diverse clinical settings
Monitoring	How will performance be evaluated after deployment?	Registry-based quality improvement and continuous performance assessment ^[20,21]	Early detection of model drift, unintended consequences, and performance degradation
Implementation	How will AI outputs be incorporated into existing clinical workflows?	Sepsis bundle implementation and performance-improvement initiatives ^[54,55]	Reliable translation of predictive insights into clinical action
Education	Who is trained to use, interpret, and critically evaluate AI outputs?	ACS fellowship training and competency-based surgical education ^[58,59]	Appropriate trust calibration and effective human-AI collaboration
Accountability	Who retains responsibility for decisions informed by AI systems?	Morbidity and mortality review, peer review, and professional oversight	Transparent governance and preservation of professional trust
Equity	Does performance vary across populations, institutions, or resource settings?	Risk-adjusted outcomes assessment and benchmarking ^[20,21]	More equitable and safer deployment across diverse practice environments
Global implementation	Can systems developed in high-resource environments be adapted safely to other settings?	Regional trauma system development and global surgery partnerships ^[60]	Broader dissemination of expertise while minimizing implementation disparities

ACS: Acute Care Surgery; AI: artificial intelligence; NTDB: National Trauma Data Bank; TQIP: Trauma Quality Improvement Program.

auditing, external validation, and local recalibration are not optional adjuncts; they are part of safe deployment^[56]. The discipline's prior experience with risk-adjusted benchmarking through the National Trauma Data Bank and the Trauma Quality Improvement Program provides directly transferable infrastructure: the same registries that enabled identification of outlier institutions can support post-deployment surveillance of algorithmic performance across centers.

The second is transparency and trust. The history of clinical decision support shows that opaque recommendations are dismissed when they conflict with clinician judgment and over-trusted when they confirm it. AI systems in ACS will need to communicate not only their outputs but their confidence, the variables driving the output, and the conditions under which they should not be relied upon. This is particularly true in high-acuity environments where decisions are time-pressured, intervention is consequential, and reversibility is limited. The acute care surgeon at the bedside needs an interpretable signal, not a black-box recommendation, and the cognitive science of trust calibration is an underdeveloped element of most current deployment plans.

The third is workforce and training. The acute care surgeon of the next decade will be expected to function alongside AI systems in diagnosis, surveillance, intraoperative guidance, and documentation. Training will need to include not only the use of these tools but the critical appraisal of their outputs, the recognition of their failure modes, and the ethical implications of their deployment^[4,56,57]. The Society of University Surgeons' Surgical Education Committee has articulated a position on the integration of AI into the training of medical students, residents, and fellows, proposing competencies and structured exposure across the surgical training continuum^[58]. Educational governance structures - the American Board of Surgery, the Accreditation Council for Graduate Medical Education, the Society of Critical Care Program Directors, and the Surgical Council on Resident Education - have a role to play in codifying these competencies, as does the American College of Surgeons through its national committees and organizations such as the Eastern

Association for the Surgery of Trauma and the Surgical Infection Society through guideline development. The integration is in progress; the framework is still being written. Current exposure to AI within graduate medical education remains variable and largely unstructured, indicating that ad hoc exposure will not be sufficient and that structured curricula are needed^[58]. Parallel developments in surgical simulation - where virtual reality, augmented reality, and AI-driven assessment are converging - provide a directly relevant model for how computational tools can be integrated into competency-based training without displacing the supervised operative experience that remains the core of surgical education^[59].

Equity considerations cut across all three [Table 4]. Tele-ICU and AI-augmented telemedicine platforms can extend acute care expertise into hospitals that lack on-site specialty coverage, partially mitigating regional disparities in access. They can also concentrate algorithmic risk in exactly those settings if validation has been performed only in better-resourced environments. The choice between these outcomes is not a property of the technology; it is a property of the deployment system. The global surgery dimension warrants particular attention: AI carries genuine potential to bridge resource gaps in low- and middle-income settings, but it also risks amplifying disparities when models trained on high-resource data are deployed without adaptation^[60]. Local context, infrastructure, and workforce considerations must shape adoption strategies rather than be retrofitted to them. A thermal imaging or video-analytic model validated at a high-volume academic center may behave unpredictably in a low-resource operating environment with different camera optics, lighting, and surgical workflow. Treating external validity as an empirical question, rather than an assumption, is a precondition for effective, responsible international deployment.

CONCLUSION

ACS at its core is a discipline of layered systems. Its history is that of successive couplings - operative discipline to infection control, operative care to physiologic observation, hemodynamic resuscitation to biologic insight, biologic insight to network thinking, and finally local capability to regional and national infrastructure. Each layer addressed a failure mode exposed by the prior one, and each retained vulnerability of its own. The most consequential advances were rarely standalone technologies. They were technologies whose value lay in how they interfaced with what already existed.

AI is best understood as the next layer, not as a discontinuity. Its potential contribution is real: continuous interpretation of data substrates that have outgrown human attentional capacity, individualized risk prediction beyond the limits of traditional scoring, real-time decision support coupled to operative video and physiologic streams, closed-loop physiologic control, and meaningful reduction of administrative burden. Its risks are equally real, and they are not principally technical. They are the risks of importing reductionist assumptions into a complex adaptive system, of optimizing local targets at the cost of global coherence, of generating alerts faster than clinicians can integrate them, and of deploying tools whose validation does not match the populations they will encounter. Future research should therefore move beyond measures of algorithmic performance to prospective evaluation of implementation endpoints, including effects on clinical decision-making, patient outcomes, workflow and cognitive burden, equity, generalizability across care environments, and unintended consequences.

The discipline that learned, painfully, that single-mediator anti-cytokine therapy could not reverse network dysregulation is well positioned to apply that lesson to AI. The same systems thinking that produced damage control resuscitation, lung-protective ventilation, regional trauma networks, and the ACS service model can be brought to bear on the integration of computational tools. The unit of analysis is the ensemble. The discipline is implementation. The historical record is the teacher. ACS has the institutional memory and the systems literacy to integrate AI, not as a replacement for its existing architecture but as its next coupling,

provided the integration is approached, as the discipline's previous transitions have been, with epistemic humility and a working knowledge of how systems actually fail. The next era of ACS will likely belong not to the most technologically aggressive systems, but to those most capable of integrating intelligence - human and artificial - into coherent, adaptive, and accountable care. If integrated thoughtfully, AI may become not a disrupter of ACS, but its next durable layer.

DECLARATIONS

Authors' contributions

Conceptualization, writing and editing: Barie PS, Kewalramani D

Critical review: Rotondo MF, Scalea TM, Narayan M

Availability of data and materials

Not applicable.

AI and AI-assisted tools statement

During the preparation of this manuscript, the AI tool ChatGPT (version 5, released 2025-08-07) was used solely for language editing. The tool did not influence the study design, data collection, analysis, interpretation, or the scientific content of the work. All authors take full responsibility for the accuracy, integrity, and final content of the manuscript.

Financial support and sponsorship

None.

Conflicts of interest

Kewalramani D is a consultant to PracticePilot and MedAstra. Narayan M is a consultant to Medcura, Gelectric, PracticePilot, and MedAstra, and serves pro bono on the board of directors of the SaveLIFE Foundation. Barie PS, Kewalramani D and Narayan M are Guest Editors of the Topic titled "Artificial Intelligence in Acute Care Surgery" of the journal *Artificial Intelligence Surgery*. Kewalramani D, Barie PS and Narayan M were not involved in any stage of the editorial process, including manuscript handling and decision-making. Rotondo MF and Scalea TM declared that there are no conflicts of interest.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Copyright

© The Author(s) 2026.

REFERENCES

1. Fink MP. Therapeutic options directed against platelet activating factor, eicosanoids and bradykinin in sepsis. *J Antimicrob Chemother.* 1998;41:81-94. [DOI PubMed](#)
2. Marshall JC. Such stuff as dreams are made on: mediator-directed therapy in sepsis. *Nat Rev Drug Discovery.* 2003;2:391-405. [DOI PubMed](#)
3. Salwei ME, Carayon P. A sociotechnical systems framework for the application of artificial intelligence in health care delivery. *J Cogn Eng Decis Mak.* 2022;16:194-206. [DOI PubMed PMC](#)
4. Mccradden MD, Anderson JA, Stephenson EA, et al. A research ethics framework for the clinical translation of healthcare machine learning. *Am J Bioeth.* 2022;22:8-22. [DOI PubMed](#)
5. Kewalramani D, Chattopadhyay K, Benton J, et al. A scoping review of artificial intelligence in acute care surgery: promise, pitfalls, and a path forward. *Int J Surg.* 2026;112:2214-24. [DOI PubMed](#)
6. Warren JC. Etherization, with surgical remarks. Boston, MA: William D. Ticknor; 1848. Available from: <https://wellcomecollection.org/works/axr94bb5>. [Last accessed on 23 Sep 2026].

7. Lister J. On the Antiseptic principle in the practice of surgery. *Br Med J*. 1867;2:246-8. DOI PubMed PMC
8. Semmelweis IP. Die Ätiologie, der Begriff und die Prophylaxis des Kindbettfiebers. Pest, Vienna, and Leipzig: C.A. Hartleben; 1861. Available from: <http://resolver.staatsbibliothek-berlin.de/SBB0001941E00000000>. [Last accessed on 23 Sep 2026].
9. Halsted WS. The training of the surgeon: the annual address in medicine delivered at Yale University, June 27, 1904. Baltimore: Johns Hopkins Press; 1904. pp. 267-75. Available from: <https://wellcomecollection.org/works/sap5h2ex>. [Last accessed on 23 Sep 2026].
10. Dandy WE. The brain. In: Lewis Practice of Surgery. Hagerstown, MD: WF Prior; 1932.
11. Beecher HK. Resuscitation and anesthesia for wounded men. Springfield, IL: Charles C Thomas; 1949. Available from: https://openlibrary.org/works/OL6587248W/Resuscitation_and_anesthesia_for_wounded_men. [Last accessed on 23 Sep 2026].
12. Ibsen B. The anæsthetist's viewpoint on the treatment of respiratory complications in poliomyelitis during the epidemic in Copenhagen, 1952. *Proc R Soc Med*. 1954;47:72-4. DOI PubMed PMC
13. Barie PS, Hydo LJ. Epidemiology of multiple organ dysfunction syndrome in critical surgical illness. *Surg Infect*. 2000;1:173-85. DOI PubMed
14. Baue AE. Multiple, progressive, or sequential systems failure: a syndrome of the 1970s. *Arch Surg*. 1975;110:779-81. DOI PubMed
15. Xiao W, Mindrinos MN, Seok J, et al. A genomic storm in critically injured humans. *J Exp Med*. 2011;208:2581-90. DOI PubMed PMC
16. Gentile LF, Cuenca AG, Efron PA, et al. Persistent inflammation and immunosuppression: a common syndrome and new horizon for surgical intensive care. *J Trauma Acute Care Surg*. 2012;72:1491-501. DOI PubMed PMC
17. Buchman TG. The community of the self. *Nature*. 2002;420:246-51. DOI PubMed
18. Marshall JC. Inflammation, coagulopathy, and the pathogenesis of multiple organ dysfunction syndrome. *Crit Care Med*. 2001;29:S99-106. DOI PubMed
19. American College of Surgeons Committee on Trauma. Resources for optimal care of the injured patient. Chicago: American College of Surgeons; 2014. Available from: <https://www.facs.org/quality-programs/trauma/quality/verification-review-and-consultation-program/standards/>. [Last accessed on 23 Sep 2026].
20. Fantus RJ, Fildes J. NTDB™ data points. Annual report 2005, dataset version 5.0. *Bull Am Coll Surg*. 2006;91:64. PubMed
21. Shafi S, Nathens AB, Cryer GH, et al. The trauma quality improvement program of the american college of surgeons committee on trauma. *J Am Coll Surg*. 2009;209:521-30.e1. DOI PubMed
22. Jurkovich GJ. Acute care surgery: The trauma surgeon's perspective. *Surgery*. 2007;141:293-6. DOI PubMed
23. Committee on Acute Care Surgery American Association for the Surgery of Trauma. The acute care surgery curriculum. *J Trauma*. 2007;62:553-6. DOI PubMed
24. Committee to Develop the Reorganized Specialty of Trauma, Surgical Critical Care, and Emergency Surgery. Acute care surgery: trauma, critical care, and emergency surgery. *J Trauma*. 2005;58:614-6. DOI PubMed
25. Bates DW, Leape LL, Cullen DJ, et al. Effect of computerized physician order entry and a team intervention on prevention of serious medication errors. *JAMA*. 1998;280:1311-6. DOI PubMed
26. Choplin RH, Boehme JM 2nd, Maynard CD. Picture archiving and communication systems: an overview. *Radiographics*. 1992;12:127-9. DOI PubMed
27. Lilly CM, Cody S, Zhao H, et al. ; University of Massachusetts Memorial Critical Care Operations Group. Hospital mortality, length of stay, and preventable complications among critically ill patients before and after tele-ICU reengineering of critical care processes. *JAMA*. 2011;305:2175-83. DOI PubMed
28. Kahn JM, Cicero BD, Wallace DJ, Iwashyna TJ. Adoption of ICU telemedicine in the United States. *Crit Care Med*. 2014;42:362-8. DOI PubMed PMC
29. Grogan EL, Norris PR, Speroff T, et al. Volatility: a new vital sign identified using a novel bedside monitoring strategy. *J Trauma*. 2005;58:7-14. DOI PubMed
30. van der Sijs H, Aarts J, Vulto A, Berg M. Overriding of drug safety alerts in computerized physician order entry. *J Am Med Inform Assoc*. 2006;13:138-47. DOI PubMed PMC
31. Loftus TJ, Tighe PJ, Filiberto AC, et al. Artificial intelligence and surgical decision-making. *JAMA Surg*. 2020;155:148-58. DOI PubMed PMC
32. Hashimoto DA, Rosman G, Rus D, Meireles OR. Artificial intelligence in surgery: promises and perils. *Ann Surg*. 2018;268:70-6. DOI PubMed PMC
33. Dreizin D, Staziaki PV, Khatri GD, et al. Artificial intelligence CAD tools in trauma imaging: a scoping review from the American Society of Emergency Radiology (ASER) AI/ML Expert Panel. *Emerg Radiol*. 2023;30:251-65. DOI PubMed PMC
34. Barie PS, Hydo LJ, Fischer E. Development of multiple organ dysfunction syndrome in critically ill patients with perforated viscus. Predictive value of APACHE severity scoring. *Arch Surg*. 1996;131:37-43. DOI PubMed
35. Talmor M, Hydo L, Barie PS. Relationship of systemic inflammatory response syndrome to organ dysfunction, length of stay, and mortality in critical surgical illness: effect of intensive care unit resuscitation. *Arch Surg*. 1999;134:81-7. DOI PubMed

36. Barie PS, Hydo LJ, Shou J, Eachempati SR. Decreasing magnitude of multiple organ dysfunction syndrome despite increasingly severe critical surgical illness: A 17-year longitudinal study. *J Trauma*. 2008;65:1227-35. [DOI PubMed](#)
37. Bertsimas D, Dunn J, Velmahos GC, Kaafarani HMA. Surgical risk is not linear: derivation and validation of a novel, user-friendly, and machine-learning-based predictive optimal trees in emergency surgery risk (POTTER) calculator. *Ann Surg*. 2018;268:574-83. [DOI PubMed](#)
38. Brennan M, Puri S, Ozrazgat-Baslanti T, et al. Comparing clinical judgment with the MySurgeryRisk algorithm for preoperative risk assessment: a pilot usability study. *Surgery*. 2019;165:1035-45. [DOI PubMed PMC](#)
39. Morris JA Jr, Norris PR, Ozdas A, et al. Reduced heart rate variability: an indicator of cardiac uncoupling and diminished physiologic reserve in 1,425 trauma patients. *J Trauma*. 2006;60:1165-73. [DOI PubMed](#)
40. Mascagni P, Vardazaryan A, Alapatt D, et al. Artificial intelligence for surgical safety: automatic assessment of the critical view of safety in laparoscopic cholecystectomy using deep learning. *Ann Surg*. 2022;275:955-61. [DOI PubMed](#)
41. Madani A, Namazi B, Altieri MS, et al. Artificial intelligence for intraoperative guidance: using semantic segmentation to identify surgical anatomy during laparoscopic cholecystectomy. *Ann Surg*. 2022;276:363-9. [DOI PubMed PMC](#)
42. Laplante S, Namazi B, Kiani P, et al. Validation of an artificial intelligence platform for the guidance of safe laparoscopic cholecystectomy. *Surg Endosc*. 2022;37:2260-8. [DOI PubMed](#)
43. Kewalramani D, Grunhut J, Jopling J, Narayan M. Distinguishing signal from perception: what video-based assessment reveals about surgical decision-making and technical performance. *J Am Coll Surg*. 2026;242:1478-80. [DOI PubMed](#)
44. Komorowski M, Celi LA, Badawi O, Gordon AC, Faisal AA. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat Med*. 2018;24:1716-20. [DOI PubMed](#)
45. Bellini V, Guzzon M, Bigliardi B, Mordonini M, Filippelli S, Bignami E. Artificial intelligence: a new tool in operating room management. Role of machine learning models in operating room optimization. *J Med Syst*. 2019;44:20. [DOI PubMed](#)
46. Sezgin E, Sirrianni J, Linwood SL. Operationalizing and implementing pretrained, large artificial intelligence linguistic models in the us health care system: outlook of generative pretrained transformer 3 (GPT-3) as a service model. *JMIR Med Inform*. 2022;10:e32875. [DOI PubMed PMC](#)
47. Kewalramani D, Roman DS, Lagos SA, et al. Quality matters: artificial intelligence-based assessment of feedback quality predicts technical skill improvement. *Surgery*. 2025;187:109625. [DOI PubMed](#)
48. Liu SI, Shikar M, Gante E, et al. Improving communication and response to clinical deterioration to increase patient safety in the intensive care unit. *Crit Care Nurse*. 2022;42:33-43. [DOI PubMed](#)
49. Cotton BA, Guy JS, Morris JA Jr, Abumrad NN. The cellular, metabolic, and systemic consequences of aggressive fluid resuscitation strategies. *Shock*. 2006;26:115-21. [DOI PubMed](#)
50. Gattinoni L, Brazzi L, Pelosi P, et al. A trial of goal-oriented hemodynamic therapy in critically ill patients. *N Engl J Med*. 1995;333:1025-32. [DOI PubMed](#)
51. Rotondo MF, Schwab CW, McGonigal MD, et al. 'Damage control': an approach for improved survival in exsanguinating penetrating abdominal injury. *J Trauma*. 1993;35:375-83. [DOI PubMed](#)
52. Sittig DF, Singh H. A new sociotechnical model for studying health information technology in complex adaptive healthcare systems. *Qual Saf Health Care*. 2010;19:i68-74. [DOI PubMed PMC](#)
53. Corrêa NK, Galvão C, Santos JW, et al. Worldwide AI ethics: a review of 200 guidelines and recommendations for AI governance. *Patterns*. 2023;4:100857. [DOI PubMed PMC](#)
54. Rhodes A, Evans LE, Alhazzani W, et al. Surviving sepsis campaign: international guidelines for management of sepsis and septic shock: 2016. *Crit Care Med*. 2017;45:486-552. [DOI PubMed](#)
55. Levy MM, Rhodes A, Phillips GS, et al. Surviving sepsis campaign: association between performance metrics and outcomes in a 7.5-year study. *Crit Care Med*. 2015;43:3-12. [DOI PubMed](#)
56. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366:447-53. [DOI PubMed](#)
57. Char DS, Shah NH, Magnus D. Implementing machine learning in health care - addressing ethical challenges. *N Engl J Med*. 2018;378:981-3. [DOI PubMed PMC](#)
58. Kewalramani D, Jawa RS, Martin CA, et al. Position statement from the society of University surgeons, surgical education committee: Artificial intelligence in surgical training for medical students, residents, and fellows. *Surgery*. 2026;190:109849. [DOI PubMed](#)
59. Riddle EW, Kewalramani D, Narayan M, Jones DB. Surgical simulation: virtual reality to artificial intelligence. *Curr Probl Surg*. 2024;61:101625. [DOI PubMed](#)
60. Valli AD, Tingle SJ, Kazerooni S, et al. Artificial intelligence in surgical care within low-income and middle-income countries: a scoping review of development, validation, and deployment. *eClinicalMedicine*. 2026;94:103836. [DOI PubMed PMC](#)

Disclaimer/Publisher's Note: All statements, opinions, and data contained in this publication are solely those of the individual author(s) and contributor(s) and do not necessarily reflect those of OAE and/or the editor(s). OAE and/or the editor(s) disclaim any responsibility for harm to persons or property resulting from the use of any ideas, methods, instructions, or products mentioned in the content.



© The Author(s) 2026. Open Access This article is licensed under a Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, sharing, adaptation, distribution and reproduction in any medium or format, for any purpose, even commercially, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.