



“Big Data” in the Intensive Care Unit Closing the Data Loop

The data generated in the process of medical care has historically not just been underused, it has been wasted. This was due in part to the difficulty of accessing, organizing, and using data entered on paper charts, but notable variability in clinical documentation methods and quality made the problem even more challenging. In the absence of a practical way to systematically capture, analyze, and integrate the information contained in the massive amount of data generated during patient care, medicine has remained a highly empirical process in which the disconnected application of individual experiences and subjective preferences continues to thwart continuous improvement and consistent delivery of best practices to all patients.

WHY A DATA-DRIVEN SYSTEM FOR CRITICAL CARE?

The pivotal studies in medical research have generally focused on the examination of the effect of a single drug, intervention, or diagnostic technique. And although essential, research at this level mutes the variation and interconnectedness that defines the modern day reality of medicine. This mode of research fails to capture systems issues such as important interactions between concomitantly applied therapies in a dynamic physiologic milieu. Hence, a high level of practice variability is inevitable as clinicians have been left with often conflicting and incomplete medical knowledge derived from a patient populace or clinical setting that may not reflect their own (1). Healthcare delivery has worked as well as it has to date because clinicians are bright, hard-working, and well-intentioned—not because of good system design or systematic data use.

Perhaps tempering enthusiasm for large-scale data archiving systems are the concerns of cost and efficiency. A recent study (2) suggests that the cost of implementing electronic health records is high, whereas the benefits remain ambiguous. This conclusion is understandable and highlights the fact that electronic health records are simply a necessary, albeit costly, first step in the process of reorganizing healthcare into a closed-loop system that can coherently and continuously generate and incorporate feedback to become better and more efficient.

The Institute of Medicine recently published “Best Care at Lower Cost: The Path to Continuously Learning Health Care in America” (3). This report states that “achieving higher quality care at lower cost will require fundamental commitments to the incentives, culture and leadership that foster continuous *learning*, as the lessons from research and each care experience are systematically captured, assessed and translated into reliable care.” To attain this vision, it is necessary to create the means to capture and archive individual clinical encounters to form a data substrate. Such a data substrate, if freely available, would create a means for clinicians and data scientists to address gaps and errors in knowledge, and support a version of *crowd sourcing for evidence creation* in clinical practice (4).

Author Contributions: All the authors participated in formulating the content and contributed to the writing of the manuscript.

L.A.C. and R.G.M. are supported by the National Institute of Biomedical Imaging and Bioengineering grant 2R01-001659.

Am J Respir Crit Care Med Vol 187, Iss. 11, pp 1157–1166, Jun 1, 2013
Internet address: www.atsjournals.org

The intensive care unit (ICU) presents an especially compelling case for clinical data analysis. The value of many treatments and interventions in the ICU is unproven, and high-quality data supporting or discouraging specific practices are embarrassingly sparse (5, 6). Guidelines developed to standardize practice are dependent on an evidence base that is surprisingly thin considering the copious data generated in the ICU. A knowledge gap of this magnitude is unacceptable for a medical discipline comprising 1% of U.S. gross domestic product (7), and for which ongoing demand is rising sharply (8). In a systematic review of multicenter randomized controlled trials evaluating the effect of ICU interventions (9), only one in seven studies showed benefit; the rest either had no measurable value or were found to be actually harmful. The purported reasons behind this perplexing observation are that the effects of interventions in the ICU are subject to the exceptional complexity of this environment and are particularly vulnerable to variation across patient subsets and clinical contexts. In fact, variations in human physiology are not nearly as problematic as the imposed variations, some inexplicable and even irrational, in local beliefs and practices (10). Some of this practice variability is due to lack of adherence to best practices, but the vast majority occurs simply because no evidence has been established for the issue in question (11). The traditional approach to evidence creation therefore needs to change and take advantage of the technical feasibility of creating complete, highly detailed critical care databases. These databases could motivate clinical investigations, support the development of clinical decision support tools, and permit the testing and perfecting of algorithms with the use of real-world data. The oncoming clinical use of “big” data sets such as genomics and proteomics will clearly require data management at this level. Finally, the unacceptable inefficiencies and wastes of the current system in conjunction with looming financial constraints demand a more thoughtfully engineered approach to healthcare delivery that continuously leverages technology to minimize both costs and nonproductive approaches while measuring and improving the clinical outcomes of individuals and populations.

CRITICAL CARE DATABASES

Several commercial and noncommercial ICU databases have been developed, typically archiving patient demographics and aggregating information such as underlying disease, severity of illness, and unit- and hospital-specific information (e.g., length of stay, mortality, and readmission). The purpose of such databases is primarily to assess and compare the severity of ICU patient conditions and outcomes, as well as treatment costs, across participating ICUs on the basis of relatively few, highly selected pieces of information. For example, the noncommercial database collected by the Australian and New Zealand Intensive Care Society now contains such data from more than 900,000 ICU stays (12).

Among commercial ICU databases, APACHE Outcomes, created at Cerner by merging APACHE (13) with Project IMPACT (14), includes data from about 150,000 ICU stays since 2010. Although large, it contains incomplete physiologic and laboratory measurements, and lacks provider notes and

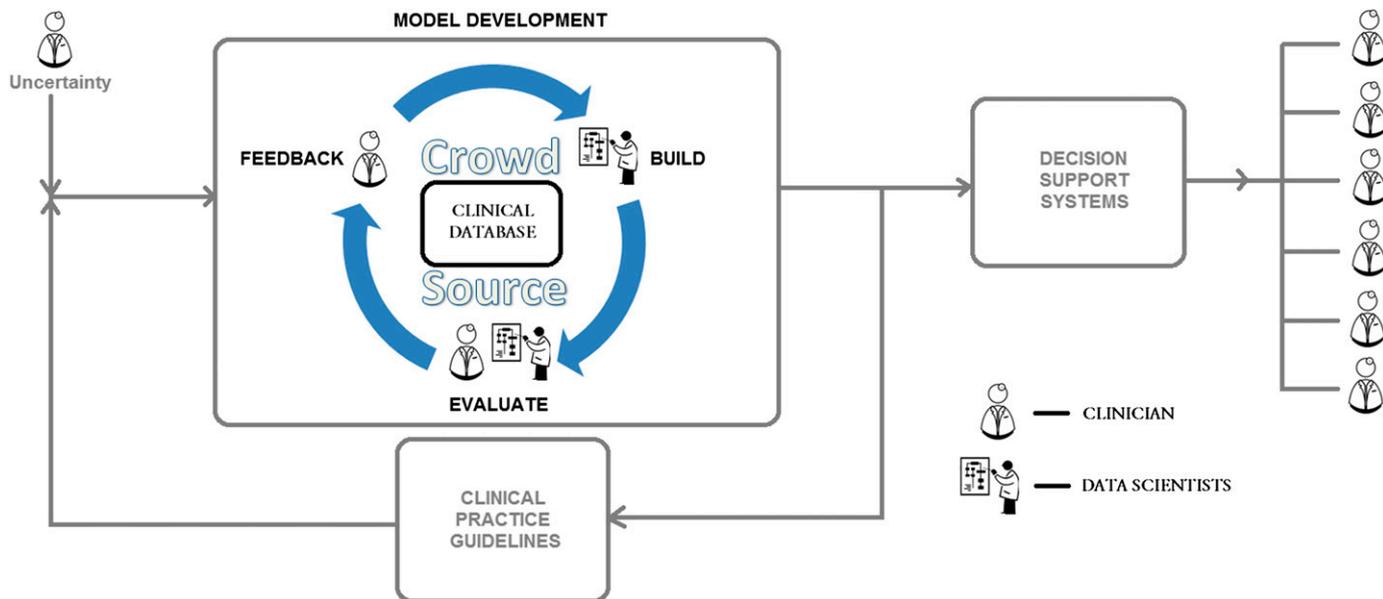


Figure 1. Data-driven learning system. Figure courtesy of Mohammad Ghassemi (Massachusetts Institute of Technology, Cambridge, MA) and Avraham Cooper (Harvard Medical School, Boston, MA).

physiologic waveform data. The commercial Philips eICU (15), a telemedicine intensive care support provider, archives data from participating ICUs. Philips eICU is estimated to maintain a database of over 1.5 million ICU stays, and is adding 400,000 patient records per year from over 180 subscribing hospitals in the country. As in the other projects above, it does not archive waveform data; provider notes are captured if they are entered into the software. This tightly controlled database is made available to selected researchers via the eICU Research Institute (16).

Over the past decade, the Laboratory of Computational Physiology at the Massachusetts Institute of Technology, Beth Israel Deaconess Medical Center, and Philips Healthcare, with support from the National Institute of Biomedical Imaging and Bioinformatics, have partnered to build and maintain the Multi-parameter Intelligent Monitoring in Intensive Care (MIMIC) database (17). This public-access database, which now holds clinical data from over 40,000 stays in Beth Israel Deaconess Medical Center ICUs, has been meticulously deidentified and is freely shared online with the research community via PhysioNet (18). It is an unparalleled research resource; nearly 600 researchers have free access to the clinical data under data use agreements. This community includes investigators from more than 32 countries and is growing by over 50% per year. In addition, thousands of investigators, educators, and students have used the waveform data, which is freely available to all. The MIMIC database is also unique in capturing highly granular structured data including minute-by-minute changes in physiologic signals as well as time-stamped treatments with dosages. Such granularity enables modeling of the individual dynamic response to a physiologic insult or clinical intervention, leading to improved risk-benefit calculation and outcome prediction (19).

CREATING DATA-DRIVEN TOOLS

Predictive Modeling, Prognostication, and Outcomes

The MIMIC database has allowed our group to develop predictive models with actionable outputs that potentially lead to measurable improvements in process and/or outcome. Such models could support appropriate early triage regarding level of care and monitoring,

as well as the allotment of costly resources such as specialist-requiring interventions and/or technologies. For example, these tools could assist emergency departments if limitations in ICU resources lead to regionalization of critical care (20). Ongoing examples of our current investigations include the following: prediction of which patients with hypotension will respond to fluid resuscitation and which ones will proceed to develop multiorgan failure; Markov modeling to determine the proper duration for a trial of aggressive ICU care among high-risk patients (21); and fuzzy modeling to predict whether gastrointestinal bleeding will stop with conservative treatment alone or requires an endoscopic or surgical intervention. Artificial intelligence methods have also been used with the MIMIC database to predict whether a laboratory test is significantly changed from the last determination by modeling the treatments and the physiologic response during the interim period among patients who presented with gastrointestinal bleeding (22). The goal is to reduce unnecessary testing, which contributes to patient discomfort, use of staff time, iatrogenic anemia, increased laboratory costs, and medical errors that result from false-positive results.

In the research setting, a data-dependent process known as “decision analysis” systematically analyzes complex decisions and quantifies the expected benefits versus harms of different treatment options (23). Although this approach has had little traction among clinicians, it has the potential to be helpful in prognostication and treatment personalization. In the ICU, patients receive technologically advanced interventions that can sustain life in the face of severe organ dysfunction(s). However, these treatments are not without cost. And although the financial costs to society are not trivial, of significantly more importance are the costs to the individual patient. The prolongation of intensive care carries the major downside of increased suffering on the part of patients and their family members, and should always be performed with the greatest thought and care, especially if the interventions are unlikely to lead to outcomes consistent with patient preferences.

The development of point-of-care decision support tools based on MIMIC is currently in progress. These are intended to provide reliable prognostication based on a database cohort

comparable for the variables that influence relevant outcomes. Outcomes of concern will not be limited to mortality or length of stay, but will be extended to include likelihood of discharge to and anticipated duration of stay in a skilled nursing facility, and probable need for repeated hospitalizations and/or procedures, for example, hemodialysis. Such prognostications will become more robust as MIMIC is expanded to include data from the Massachusetts All Payer Claims Database (24).

Unraveling Complexity and Variability

Databases like MIMIC that include detailed clinical information provide researchers an opportunity to accumulate safety and efficacy evidence, to discover patient subpopulations that experience important variances in efficacy or unanticipated delayed adverse effects, and to uncover interactions between and among simultaneous treatments as drugs become used in wider, more diverse patient populations than those possible during premarket approval clinical studies. Such an active national surveillance system would allow drugs to be monitored longitudinally over their entire market life, providing the Food and Drug Administration timely access to new information for evaluating a drug's benefit–risk profile. In a previous article, we described a pharmacovigilance surveillance system that will extend the Adverse Event Reporting System of the Food and Drug Administration (25). As further examples, there are ongoing analyses using MIMIC to examine the effect of long-term use of selective serotonin reuptake inhibitors (26) and proton pump inhibitors on outcomes of critical illness. A similar approach could be gainfully applied to epidemiological issues as well.

Clinical databases such as MIMIC represent an opportunity to study clinical areas where practice variation exists as a result of either lack of or conflicting medical knowledge. In a previous article, we used MIMIC to explore practice variation and health outcomes in critically ill patients admitted for or subsequently developing hypotension in an ICU (27). The decision to administer intravenous fluids or vasoactive agents, and the volume or dose chosen, largely depends on clinician preference, local practice patterns, and unsystematic process-related factors at the time of the hypotensive event (28). Clinician decision making in the absence of strong guidelines and evidence from randomized controlled trials is highly driven by prior training and experience (29) and results in significant variability in care quality.

Our study that analyzes discrete hypotensive episodes (27) is a first step in systematically examining the management of hypotension in the ICU, and suggests several new approaches for investigation and standardization of care processes. Importantly, the results also support the need to personalize treatment by providing a patient- and context-specific optimal blood pressure target range. This can be accomplished by investigating previous admissions for the patient, and others in the same age range with similar comorbidities and presentation in the database. A target blood pressure level can then be objectively determined by mapping blood pressure to a measure of tissue perfusion. This target level setting could subsequently be fine-tuned and fully contextualized by adjusting for illness severity, treatments received, and phase of disease process.

Practice variation is unavoidable when medical knowledge is lacking or conflicting. The analysis of care variability and the performance of comparative effectiveness studies in these situations require the understanding (and modeling) of the reasoning behind the administration or avoidance of an intervention, including provider and institution bias. This may be partially addressed by including the providers and the institutions as covariates in the modeling.

Bias (which adds noise to the models) might also be offset by augmentation of the signal amplitude by increasing the sample size and the number of institutions (30). The association between a treatment process with a treatment outcome, even if confounded by the noise of provider and institution bias, should be detectable with a sufficiently scaled-up data set. Finally, we redirect the reader to the multisociety statement published by the American College of Chest Physicians, American Thoracic Society, and Society of Critical Care Medicine that addresses the role of clinical research, including retrospective database analysis, in the practice of critical care medicine (1).

CHANGE AND CHALLENGE

Although useful and important, the MIMIC database is currently limited by its restriction to a single institution. Plans are underway to scale up the project to include data from intensive care units in the United Kingdom and France. An international database affords numerous benefits apart from augmenting true signals as discussed above. The practice variation across ICUs that can be examined is much richer. Cross-validation of models across institutions will determine which findings are institution-specific and which are generalizable. Most importantly, knowledge discovery is accelerated exponentially if more investigators participate in clinical data mining.

Clinicians at the front line of care should be at the core of this dynamic learning system, fully supported by engineers to collaborate on the daily translation of questions into strategies for database interrogation, modeling, and analysis. This learning system will engender a medical culture in which clinicians and engineers work together in a mutually supportive environment where cross-specialty communication is not only possible but intrinsic and continuous. Our vision is for the development of a care system consisting of “clinical informatics without walls” (Figure 1), in which the creation of evidence and clinical decision support tools is initiated, updated, honed, and enhanced by crowd sourcing. In this collaborative medical culture, knowledge generation would become *routine and fully integrated* into the clinical workflow. This system would use individual data to benefit the care of populations and population data to benefit the care of individuals. Medicine should long ago have incorporated this kind of data-driven approach, but unfortunately, for technical (including privacy and security), financial (including data ownership), and cultural reasons, this has not yet occurred. Work of this nature has the potential to identify and scale best practice, and generate huge dividends in human health and the rational use of healthcare resources.

Author disclosures are available with the text of this article at www.atsjournals.org.

LEO ANTHONY CELI, M.D., M.S., M.P.H.

ROGER G. MARK, M.D., Ph.D.

*Institute for Medical Engineering and Science
Massachusetts Institute of Technology
Cambridge, Massachusetts*

and
*Department of Medicine
Beth Israel Deaconess Medical Center
Boston, Massachusetts*

DAVID J. STONE, M.D.

*Departments of Anesthesiology and Neurological Surgery
and the Center for Wireless Health
University of Virginia
Charlottesville, Virginia*

ROBERT A. MONTGOMERY

*Harvard Medical School
Boston, Massachusetts*

References

1. Tonelli MR, Curtis JR, Guntupalli KK, Rubenfeld GD, Arroliga AC, Brochard L, Douglas IS, Gutterman DD, Hall JR, Kavanagh BP, *et al.*; ACCP/ATS/SCCM Working Group. An official multi-society statement: the role of clinical research results in the practice of critical care medicine. *Am J Respir Crit Care Med* 2012;185:1117–1124.
2. Kellerman AL, Jones SS. What it will take to achieve the as-yet-unfulfilled promises of health information technology. *Health Aff (Millwood)* 2013; 32:63–68.
3. Institute of Medicine, Committee on the Learning Health Care System in America. Best care at lower cost: the path to continuously learning health care in America. Washington, DC: The National Academies Press; 2012.
4. Lakhani KR, Boudreau KJ, Loh PR, Backstrom L, Baldwin C, Lonstein E, Lydon M, MacCormack A, Arnaut RA, Guinan EC. Prize-based contests can provide solutions to computational biology problems. *Nat Biotechnol* 2013;31:108–111.
5. Vincent JL, Singer M. Critical care: advances and future perspectives. *Lancet* 2010;376:1354–1361.
6. Vincent JL. Is the current management of severe sepsis and septic shock really evidence based? *PLoS Med* 2006;3:e346.
7. Halpern NA, Bettles L, Greenstein R. Federal and nationwide intensive care units and healthcare costs: 1986-1992. *Crit Care Med* 1994;22:2001–2007.
8. Angus DC, Kelley MA, Schmitz RJ, White A, Popovich J Jr; Committee on Manpower for Pulmonary and Critical Care Societies (COMPACCS). Caring for the critically ill patient. Current and projected workforce requirements for care of the critically ill and patients with pulmonary disease: can we meet the requirements of an aging population? *JAMA* 2000;284:2762–2770.
9. Ospina-Tascón GA, Büchele GL, Vincent JL. Multicenter, randomized, controlled trials evaluating mortality in intensive care: doomed to fail? *Crit Care Med* 2008;36:1311–1322.
10. Walkey AJ, Wiener RS. Risk factors for underuse of lung-protective ventilation in acute lung injury. *J Crit Care* 2012;27:323.e1–e9.
11. James BC, Savitz LA. How Intermountain trimmed health care costs through robust quality improvement efforts. *Health Aff (Millwood)* 2011;30:1185–1191.
12. Stowa PJ, Hart GK, Higlett T, George C, Herkes R, McWilliam D, Bellomo R; ANZICS Database Management Committee. Development and implementation of a high-quality clinical database: the Australian and New Zealand Intensive Care Society Adult Patient Database. *J Crit Care* 2006;21:133–141.
13. Zimmerman JE, Kramer AA, McNair DS, Malila FM, Shaffer VL. Intensive care unit length of stay: Benchmarking based on Acute Physiology and Chronic Health Evaluation (APACHE) IV. *Crit Care Med* 2006;34:2517–2529.
14. Cook SF, Visscher WA, Hobbs CL, Williams RL; Project IMPACT Clinical Implementation Committee. Project IMPACT: results from a pilot validity study of a new observational database. *Crit Care Med* 2002;30:2765–2770.
15. eICU Program Solution. Baltimore, MD: Koninklijke Philips Electronics N.V.; 2012 [accessed 2013 Feb 10]. Available from: <http://eicu.mediaroom.com/index.php?s=28701>
16. McShea M, Holl R, Badawi O, Riker RR, Silfen E. The eICU research institute - a collaboration between industry, health-care providers, and academia. *IEEE Eng Med Biol Mag* 2010;29:18–25.
17. Saeed M, Villarreal M, Reisner AT, Clifford G, Lehman LW, Moody G, Heldt T, Kyaw TH, Moody B, Mark RG. Multiparameter Intelligent Monitoring in Intensive Care II: a public-access intensive care unit database. *Crit Care Med* 2011;39:952–960.
18. MIMIC II databases [updated 2012 Oct 4; accessed 2013 Feb 10]. Available from: <http://physionet.org/mimic2>
19. Mayaud L, Lai PS, Clifford GD, Tarassenko L, Celi LA, Annane D. Dynamic data during hypotensive episode improves mortality predictions among patients with sepsis and hypotension. *Crit Care Med* 2013;41:954–962.
20. Kelley MA, Angus D, Chalfin DB, Crandall ED, Ingbar D, Johanson W, Medina J, Sessler CN, Vender JS. The critical care crisis in the United States: a report from the profession. *Chest* 2004;125:1514–1517.
21. Lai PS, Shrimel MG, Berket BS, Scott DJ, Celi LA, Hunink M. Using Markov models to obtain the optimal duration for a trial of intensive care in patients with active cancer and septic shock [abstract]. *Am J Respir Crit Care Med* 2012;185:A3998.
22. Cismondi F, Celi LA, Fialho A, Vieira SM, Reti SR, Sousa JM, Finkelstein SN. Reducing ICU blood draws with artificial intelligence. *Int J Med Inform* 2013;82:345–358.
23. Weinstein M. Clinical decision analysis. Philadelphia: W. B. Saunders Co.; 1980.
24. Center for Health Information and Analysis. All-payer claims database. Center for Health Information and Analysis, Commonwealth of Massachusetts; 2013 [accessed 2013 Feb 10]. Available from: <http://www.mass.gov/chia/researcher/health-care-delivery/hcf-data-resources/apcd/>
25. Moses CM, Celi LA, Marshall J. Pharmacovigilance: an active surveillance system to proactively identify risks for adverse events. *Popul Health Manag* (In press)
26. Berg KM, Ghassemi M, Marshall J, Donnino MW, Marshall J, Celi L. Pre-admission use of selective serotonin reuptake inhibitors is associated with ICU mortality [abstract]. *Am J Respir Crit Care Med* 2012; 185:A6840.
27. Lee J, Kothari R, Ladapo JA, Scott DJ, Celi LA. Interrogating a clinical database to study treatment of hypotension in the critically ill. *BMJ Open* 2012;2:e000916.
28. Takala J. Should we target blood pressure in sepsis? *Crit Care Med* 2010; 38(10, Suppl):S613–S619.
29. Angus DC. Caring for the critically ill patient: challenges and opportunities. *JAMA* 2007;298:456–458.
30. Silvers N. The signal and the noise. New York: The Penguin Press; 2012.

Copyright © 2013 by the American Thoracic Society
DOI: 10.1164/rccm.201212-2311ED

Maternal Distress and Childhood Wheeze: Mechanisms and Context



The article by Wright and colleagues in this issue of the *Journal* (pp. 1186–1193) represents yet another important contribution to our understanding on how early-life stress can program the development of persistent wheeze in children (1). Based on five daily salivary cortisol collections over 3 days in 261 pregnant women, they reported a twofold increase in wheeze at 2 years of age among children born to women with higher bedtime cortisol levels. This association was independent of maternal overweight in pregnancy (and observed lower morning cortisol levels); maternal overweight on its own increased the likelihood of childhood wheeze more than threefold. Finally,

they found an additive association between overweight and cortisol profiles, such that wheeze at 2 years of age was more likely in obese pregnant women with abnormal cortisol profiles. As discussed by the authors in their current publication and in more detail by Wright in a previous review article (2), maternal cortisol influences the development of structure and function in the fetal lung and can affect cytokine production. Both maternal overweight and hypothalamic–pituitary–adrenal axis disruption have the capacity to contribute to oxidative stress, altering the delicate prooxidant–antioxidant balance of the placenta and the tight regulation of cortisol transfer to the fetus (3, 4). Adjustment