



AI

IN

THE

IN A HOSPITAL'S INTENSIVE CARE UNIT (ICU), the sickest patients receive round-the-clock care as they lie in beds with their bodies connected to a bevy of surrounding machines. This advanced medical equipment is designed to keep an ailing person alive. Intravenous fluids drip into the bloodstream, while mechanical ventilators push air into the lungs. Sensors attached to the body track heart rate, blood pressure, and other vital signs, while bedside monitors graph the data in undulating lines. When the machines record measurements that are outside of normal parameters, beeps and alarms ring out to alert the medical staff to potential problems. • While this scene is laden with high tech, the technology isn't being used to best advantage. Each machine is monitoring a discrete part of the body, but the machines aren't working in concert. The rich streams of data aren't being captured or analyzed. And it's impossible for the ICU team—critical-care physicians, nurses, respiratory therapists, pharmacists, and other specialists—to keep watch at every patient's bedside. • The ICU of the future will make far better use of its machines and the continuous streams of

SPECTRUM.IEEE.ORG | NORTH AMERICAN | OCT 2018 | 31

IN THE INTENSIVE CARE
UNIT, ARTIFICIAL
INTELLIGENCE
CAN KEEP WATCH

BY BEHNOOD GHOLAMI,
WASSIM M. HADDAD
& JAMES M. BAILEY
ILLUSTRATIONS BY MCKIBILLO

ICU

data they generate. Monitors won't work in isolation, but instead will pool their information to present a comprehensive picture of the patient's health to doctors. And that information will also flow to artificial intelligence (AI) systems, which will autonomously adjust equipment settings to keep the patient in optimal condition.

At our company, Autonomous Healthcare, based in Hoboken, N.J., we're designing and building some of the first AI systems for the ICU. These technologies are intended to provide vigilant and nuanced care, as if an expert were at the patient's bedside every second, carefully calibrating treatment. Such systems could relieve the burden on the overtaxed staff in critical-care units. What's more, if the technology helps patients get out of the ICU sooner, it could bring down the skyrocketing costs of health care. We're focusing initially on hospitals in the United States, but our technology could be useful all around the world as populations age and the prevalence of chronic diseases grows.

The benefits could be huge. In the United States, ICUs are among the most expensive components of the health care system. About 55,000 patients are cared for in an ICU every day, with the typical daily cost ranging from US \$3,000 to \$10,000. The cumulative cost is more than \$80 billion per year.

As baby boomers reach old age, ICUs are becoming increasingly important. Today, more than half of ICU patients in the United States are over the age of 65—a demographic group that's expected to grow from 46 million in 2014 to 74 million by 2030. Similar trends in Europe and Asia make this a worldwide problem. To meet the growing demand for acute clinical care, ICUs will need to increase their capacity as well as their capabilities. Training more critical-care specialists is part of the solution—but so is automation. Far from replacing humans, AI systems could become part of the medical team, allowing doctors and nurses to deploy their skills when and where they're needed most.

IN ICUs TODAY, the data from the raft of bedside monitors is usually lost as the monitor screens refresh every few seconds. While some advanced ICUs are now trying to archive these measurements, they still struggle to mine the data for clinical insights.

A human doctor typically has neither the time nor the tools to make sense of the rapidly accumulating data. But an AI system does. It could also take actions based on the data, such as adjusting the machines involved in crucial ICU tasks. At Autonomous Healthcare, we're focusing first on AI systems that could manage a patient's ventilation and fluids. Mechanical ventilators come into play when a patient is sedated or suffers lung failure, a common ICU condition. And careful fluid management maintains the proper volume of blood flowing through a patient's circulatory system, therefore ensuring that all the tissues and organs get enough oxygen.

Our methodologies spring from an unlikely source: the aerospace industry. Two of us, Haddad and Gholami, are aerospace control engineers. We met at Georgia Tech's school of aero-

space engineering, where Haddad is a professor of dynamical systems and control and Gholami formerly worked as a doctoral researcher. Bailey joined the collaboration in the early 2000s when he was an associate professor of anesthesiology at the Emory University school of medicine. Haddad and Bailey first worked on control methods to automate anesthesia dosing and delivery in the operating room, which we tested in clinical studies at Emory University Hospital, in Atlanta, and Northeast Georgia Medical Center, in Gainesville, Ga. We then set our sights on more complex and broader control problems for the ICU. In 2013, Haddad and Gholami founded Autonomous Healthcare to commercialize our AI systems. Gholami is the company's CEO, Haddad is chief science advisor, and Bailey is chief medical officer.


How is aerospace similar to medicine? Both fields involve vast amounts of data that must be processed quickly to make decisions while lives hang in the balance, and both require that many tasks be performed simultaneously to keep things running smoothly. In particular, we see a role for feedback control technology in critical-care medicine. These technologies use algorithms and feedback to modify the behavior of engineered systems through sensing, computation, and actuation. They have become ubiquitous in the safety-critical systems of flight control and air traffic control.

However, there's a key difference between an airplane and a hospital patient. An airplane's design and control is based on well-established theories of mechanics and aerodynamics, whereas the human body involves highly complex biological systems that function and interact in ways we don't yet entirely understand.

Consider the management of mechanical ventilation. ICU patients may require this support because of direct trauma, lung infection, heart failure, or an inflammatory syndrome such as sepsis. The ventilator alternates between forcing air into the lungs and allowing the lungs to passively deflate. The device can be dialed up or down to do all of the work or to assist the patient's own efforts.

The interaction between human and machine is a subtle thing to manage. The human body has its own automatic mechanism to govern breathing, in which the nervous system triggers the diaphragm muscle to contract and pull downward on the lungs, thus initiating the intake of air. The ventilator must work with this innate drive; it should be synchronized with the patient's natural transitions between inhaling and exhaling, and it should match the natural air volume of the patient's breathing.

Unfortunately, mismatches between the patient's demand and the machine's delivery are all too common, which can cause a patient to "fight the ventilator." For example, a patient may naturally need more time to inhale, but the ventilator transitions to the exhalation prematurely. This and other synchronization problems with mechanical ventilation are associated with longer stints on the ventilator, longer stays in the ICU, and increased risk of death. Experts don't yet know why asynchrony has these detrimental effects, but patients clearly experience discomfort when trying to breathe out while the machine is




BREATHING EASIER CRITICALLY ILL PATIENTS who need help breathing are put on mechanical ventilators [1]. These machines push air into the lungs, but the rhythm can get out of sync with natural breathing patterns, causing patients to “fight the ventilator.” A smart control system could read airflow measurements [2] and identify different types of ventilator asynchrony [3] in real time via a machine-learning algorithm. In a fully autonomous system, an adaptive controller [4] would constantly adjust the ventilator’s airflow to keep it in sync with the patient. As a step toward the goal of full autonomy, a similar system could be used as a decision-support tool in the ICU, providing recommendations that respiratory therapists could use to make adjustments.

pushing air into their lungs, and their laboring muscles experience an additional workload. In ICUs in the United States, the share of patients on ventilators who experience severe asynchrony has been estimated to be between 12 and 43 percent.

The first step in addressing this problem is to detect it. Experienced respiratory therapists can identify different types of asynchrony if they continuously monitor the waveforms on a ventilator’s display screen indicating the pressure and flow. But in an ICU, one respiratory therapist typically oversees 10 or more patients and can’t possibly monitor all of them constantly.

At our company, we’ve designed a machine-learning framework that replicates that human expertise in detecting different types of asynchrony. To train our system, we used a data set of waveforms from patients on ventilators, in which each waveform had been evaluated by a panel of clinical experts. Our algorithm learned the signatures of different asynchrony types—such as a particular dip in the flow signal at a specific point in time. In our first assessments of the algorithm’s performance, we focused on what’s called cycling asynchrony, which is the most challenging type to detect.



FLUID MOVEMENTS

MOST ICU PATIENTS require infusion pumps and IVs [1] to drip fluid into their veins. Getting the fluid volume right is crucial: If levels are either too low or too high in the circulatory system, serious complications can arise. A smart control system could track real-time measurements [2] such as arterial blood pressure and the amount of blood pumped by the heart; the system could then feed the data into a physiological model [3] that represents how fluids move through the body's blood vessels and tissues. In a fully autonomous system, an adaptive controller [4] could continuously adjust fluid inputs to keep the patient stable. Initially, ICU physicians could use the technology as a decision-support system that provides recommendations.

Here the ventilator's initiation of the exhale doesn't match the patient's own exhalation. The accuracy of our algorithm in detecting cycling asynchrony in a new data set matched that of human experts.

We're now testing the algorithm at Northeast Georgia Medical Center's ICU to detect respiratory asynchrony in real patients and in real time. The technology has been incorporated into a clinical-decision support system, which is designed to help respiratory therapists assess a patient's needs. This framework can also provide researchers with a tool to better understand the underlying causes of asynchrony and its

impact on patients. Our long-term goal is to design mechanical ventilators that can automatically adjust their own settings in response to each patient's needs.

WHEN YOU PICTURE AN ICU, your mental image probably includes patients with plastic bags hanging from stands by their bedsides, fluids continually dripping into their veins through IVs. About 75 percent of patients require such fluid management at some point during their stay in the ICU.

However, calibrating the correct amount of fluid is far from an exact science. Just tracking a patient's fluid levels is a hard

task: No existing medical sensors can directly monitor fluid volume, so doctors rely on indirect indicators like blood pressure and urine volume. The amount of fluids that patients need depends on their illness and medications, among other things.

Getting the fluids right is particularly important for patients with sepsis, a life-threatening syndrome characterized by inflammation throughout the body. In these patients, the blood vessels dilate, thus reducing blood pressure, and fluid leaks from the tiniest vessels, the capillaries. As a result, less oxygen-carrying blood reaches the organs, which can cause organs to fail and patients to die. Doctors combat sepsis by dispensing drugs to boost blood pressure and pumping extra fluids into patients' circulatory systems.

It's important to add enough fluid, but not too much—an excess can cause complications such as pulmonary edema, a buildup of fluid in the lungs that can interfere with breathing. Studies have shown that fluid overload is associated with more days on mechanical ventilators, longer stays in the hospital, and higher rates of mortality.

Doctors therefore aim to maintain their patients' fluids at certain levels, which are based on models for an average patient. When the doctors come through the ICU on their rounds, they try to determine whether the patient is holding steady at the goal level by checking the mix of gases in the blood and monitoring blood pressure and urine output. Deciding when to add fluids and how much to add is highly subjective, and there's considerable debate about the best practices.

An AI system could do better. Rather than basing its decisions on goals established for an average patient, it could analyze a wide variety of physiological indicators for an individual patient in real time, and continuously dispense fluids according to that patient's specific needs.

At Autonomous Healthcare, we've developed a fully automated system that looks at indirect measurements of a patient's fluid levels (such as blood pressure and variation in the volume of blood pumped out by each heartbeat) and then feeds the data into a sophisticated physiological model. Our system uses these measurements to assess how fluids are moving between the body's blood vessels and tissues, constantly adjusting the parameters as new measurements come in. Our proprietary adaptive controller then adjusts the fluid-infusion settings accordingly.

One advantage of our technology is its attention to what control engineers call closed-loop system stability, which means that any perturbations to a normal state lead to only small and fleeting variations. Many engineering applications use control systems that guarantee closed-loop stability—when an airplane runs into powerful turbulence, for instance, the autopilot system compensates to keep the shaking to a minimum. However, most control systems for medical devices have no such guarantee. If doctors judge that a sepsis patient's fluid levels have dramatically dropped, they might push a large volume of fluid into the bloodstream, perhaps overcompensating.

We've already tested our automated fluid-management system in collaboration with William Muir, a veterinary anes-

thesiologist and cardiovascular physiologist. Working with dogs that were experiencing hemorrhages, we used our system to regulate their fluid infusions. Our system successfully kept the dogs in stable condition as measured by the volume of blood pumped with every heartbeat.

We need to do more testing in order to win regulatory approval for a fully automated fluid-management system for humans. As with our work on ventilator management, we can start by building a decision support system for the ICU. This "human in the loop" system will present information and recommendations to the clinician, who will then adjust the settings of the infusion pump accordingly.

LOOKING BEYOND VENTILATION and fluid management, other key aspects of patient care that could be automated include pain management and sedation. In the ICU of the future, we envision many such clinical operations being monitored, coordinated, and controlled by AI systems that assess each patient's physiological state and adjust equipment settings in real time.

To make this vision a reality, though, it won't be enough for engineers to produce reliable technology. We must also find our way through many regulatory barriers and institutional requirements at hospitals.

Clearly, regulators need to scrutinize any new autonomous medical system. We suggest that regulatory agencies make use of two testing frameworks commonly used in the automotive and aerospace industries. The first is *in silico* trials, which test an algorithm through computer simulations. These tests are useful only if the simulations are based on high-fidelity physiological models, but in certain applications this is already possible. For example, the U.S. Food and Drug Administration recently approved the use of *in silico* testing as a replacement for animal testing in efforts to develop an artificial pancreas for diabetics.

The second useful framework is hardware-in-the-loop testing, where hardware stands in for the object of interest, whether it's a jet engine or the human circulatory system. You can then test a device—an autonomous fluid pump, say—on the hardware platform, which will generate the same type of data you'd see on an actual patient's bedside monitor. These hardware-in-the-loop trials can show that the device performs well in real time and in the real world. Once these technologies have been proven with stand-ins for critically ill humans, testing can begin with real patients.

To bring these technologies into hospitals, the final step is to win the trust of the medical community. Medicine is a generally conservative environment—and for good reason. No one wants to make changes that might threaten the health of patients. Our approach is to prove our technologies in stages: We'll first commercialize decision-support systems to demonstrate their efficacy and benefits, and then move to truly autonomous systems. With the addition of AI, we believe ICUs can be smarter, safer, and healthier places. ■

▶ **POST YOUR COMMENTS** at <https://spectrum.ieee.org/smarticu1018>