



An AI That's Not Artificial at All

Messy and unpredictable situations—such as coordinating medical care during military conflicts—have exposed the limits of current design practices. Could a new methodology empower humans to innovate and solve problems?

Despite the fact that geopolitical tensions had been steadily increasing over the last week and the entire base was on high alert, the shock and violence of the missile attack on the airfield at 0230 was overwhelming. Trauma surgeon Colonel William Smith hastily entered the hospital and scrubbed in to prepare for the wave of casualties.

In 2017, the Defense Advanced Research Projects Agency (DARPA) ran a series of wargames with military logisticians, medical professionals, aviators, and planners. Building on a scenario like the one above, the goal of the wargames was to imagine new ways to evacuate and treat the wounded in future conflicts. As DARPA program managers, we worked through the profound logistical challenges of saving wounded service members across multiple military services and host nations. This experience led us to develop a new way of thinking about how to use artificial intelligence to help people work together in complex, multi-system, difficult-to-predict situations. Instead of focusing on what humans or machines are better at, we focused on the liminal spaces—gaps between systems and individuals, for example, or between the present and the future—with the aim of sewing together intelligences to empower humans to innovate and solve problems.

Amid a global pandemic that has revealed our inability to quickly bridge crucial gaps in knowledge and systems, this approach, which we call liminal design, seems particularly promising for its ability to mediate interactions between

people and coordinate their efforts around common goals. Virtual work, networks, and technology platforms have overturned long-held assumptions about how humans (and machines) can work together, especially in the last year and a half. Technology built with AI may soon be used to address emergent challenges such as fighting wildfires with input from interdisciplinary teams of firefighters, forest ecologists, and meteorologists, or, similarly, coordinating multi-institutional responses to public health or military emergencies. For all of their promise, however, these technologies and configurations challenge fundamental assumptions about decision making, autonomy, and scale in how we organize ourselves. As liminal design and other methods are developed to help orchestrate human-machine collaboration, their implications for governance will need close examination.

“Sir, we have 54 casualties inbound from the airfield. The runway is a mess and we aren’t getting medevac until the engineers can repair the craters. We need to stabilize the injured until we get transport,” reported the nurse as Colonel Smith prepared to wrap up his fourth procedure. It was going to be a long night.

The wargames, which were organized as part of a DARPA program called CASCADE (Complex Adaptive System Composition and Design Environment), forced planners to imagine a horrific future conflict with many unknowns. For

the last 20 years, military medical operations in Afghanistan and Iraq have been extraordinarily successful. Because opponents never impeded the US military's ability to fly and communicate in these conflicts, it has been possible to provide high-quality surgical care to severely injured soldiers within the "golden hour" of medical intervention (the first hour after a traumatic injury, when treatment is most effective), preventing death in most cases. Stabilized casualties were then evacuated rapidly by air to excellent hospitals in Europe and the United States to receive surgery.

However, in a great power conflict, the scale of engagement and the inability both to control the skies and ensure communications could lead to a volume and severity of casualties greater than any in living memory. The goal of the CASCADE wargames was to design a medical system that could support the military mission while minimizing the loss of life—and to do so without the ability to connect to, or even know, what might be happening in other, distant parts of the system.

We discovered that design paradigms we'd used in the past were inadequate in the chaotic settings we were exploring. We began with a user-centered design approach, which starts by considering the problem from the viewpoint of a single user, such as the surgeon Colonel Smith, and then prototypes a solution to the problem based on empathy and understanding; then we would use that feedback to converge on an optimal solution. But when we applied user-centered design in the CASCADE scenarios, we found that it failed to account for larger system-level impacts. It focused on the problem as defined, not as it might evolve. What's more, an excellent solution for one user, such as a surgeon, could even create new problems for other users such as pilots, blood supply managers, or logisticians at far-off hospitals.

To understand the problems better, our team interviewed trauma surgeons and nurses, most of whom were combat veterans from tactical units in Iraq and Afghanistan. We learned that the highest priority was time: How can we provide surgical intervention within the golden hour without being able to fly casualties back to better equipped hospitals? How can we amplify the ability of surgeons to cope with a surge of casualties an order of magnitude beyond anything they had experienced before?

Then, in the context of our wargames and models, we prototyped ways technology could strengthen the capabilities of forward surgical teams, which consist of a surgeon and nurses operating close to the geographic location of combat at high risk to their own safety, and shock trauma platoons, which are battlefield medical units with

limited ability to do surgery. We considered "autonomous combat care," for example, which uses AI and telemetry to stabilize patients by automatically monitoring symptoms of uncontrolled internal bleeding and administering anesthesia; this increases the number of patients each surgeon and nurse can care for.

Yet when we examined these interventions within the broader system, we found we were generating a game of "whack-a-mole"—creating new problems elsewhere in the system and across other classes of logistical support. More stabilized casualties required more airlift flights to advanced surgical care facilities, which consumed more fuel. Finite airlift capacity had to be traded off against the movement of other supplies, such as critical parts for tanks and aircraft.

"Ma'am, we only have three units left of AB negative blood, and we have two inbound that need three each. We also have a shortage of surgical instruments since the sterilizer went down yesterday. What should we do?" asked the medical supply technician. Major Julie Evans, who managed medical supply for the hospital, needed to come up with a solution quickly.

Shifting back from the surgeon and focusing on the bigger picture took us to a more system-centric design paradigm. Systems design requires first identifying the elementary parts of the system and then designing, building, integrating, and testing progressively larger assemblies of those parts to achieve optimal performance. We collected data from all the stakeholders, built models to obtain insights and identify leverage points, and then looked for interventions that had desired impacts.

In interviews, we learned about the critical role blood plays on the battlefield. Medical units only control medevac helicopters, which they will sometimes fly without patients just to get blood to those in need. This ad hoc logistics solution uses precious fuel and potentially slows the movement of critically injured patients to surgical units—another whack-a-mole cascade of unintended consequences that often results from such situationally necessary workarounds. Real-world knowledge like this, gleaned from interviews with surgeons, blood professionals, and airlift pilots, was essential to our wargames. We built a scenario involving competing priorities for airlift, supplies, and medical care backed by system simulation tools to explore choices and consequences.

Our analysis demonstrated that the existing system was neither robust nor flexible enough to respond in an

environment with limited airlift and communications. Maintaining consistent supplies of blood and intravenous fluids to the frontlines in response to surges in demand while also meeting other logistical constraints and moving patients to advanced surgical facilities in Europe or the United States was nearly impossible using conventional approaches. If the forward surgical team succeeded in quickly stabilizing patients, facilities capable of more substantive surgical procedures would soon begin to back up. Large hospitals had excess capacity because airlifts couldn't get patients to them, while medical units at the tactical edge, lacking communication, were unable to effectively evacuate patients or maintain supply.

Systems analysis has weaknesses, however. To make analysis manageable, causal inference and important nuances sometimes must be sacrificed, which can make it hard to tell which particular actions are helping or hurting on a system-wide level. In the simulation, we realized that life or death could be determined by the specific personnel capabilities, the precise nature of the injuries involved, and which service was running a given facility. For example, casualties in sea-based operations generally have more burns than ground war casualties do. Burns require intravenous fluids and skin grafts, not large quantities of blood. Simply flowing naval battle casualties to an Army or Air Force facility—a systems design-style solution—clogs those clinics with patients they are ill-equipped to treat. So on its own, as with user-centered design, systems design had disabling limitations. We realized we needed another type of design to combine individual knowledge and expertise into the broader context of the system; we needed an approach that was both flexible and scalable.

Major Evans started the logistics command and control (LogC2) app on her tablet and connected to a transient local network. She submitted requests for both the sterilizer repair parts and the units of AB negative blood. But given the lack of airlift, how was she going to get these vital supplies?

In the wargame, we realized we needed to focus on more than just maintaining operational speed and minimizing casualties. We needed to maximize options for individual users and increase the learning rate across the whole system. Providing individuals with more options and the autonomy to use them makes the rigid, monolithic systems of slow-moving bureaucracies and the technologies they use more adaptable to new situations and innovations. In Major Evans's case, if she were able to get a drone to fly blood to her unit, she

might be able to boost her dwindling blood supplies faster than it would take to fix the ruined airstrip. But thoughtful workarounds can only benefit the larger system if the know-how circulates throughout the enterprise and others can begin to help find the drones and arrange for the delivery. The larger system needs to effectively learn and adapt to the consequences of her changes or it will soon be caught in another cycle of cascading ad hoc responses to problems.

A new design methodology

To build a system that was capable of encouraging individual innovation and system-wide learning, we came up with a new approach: liminal design. It employs four core concepts: abstraction, composition, mediation, and learning. Collectively, these ideas create the foundation for an “operating system” that works in an adaptive ecosystem, bridging the worlds of user-centered and system-centered design.

Abstraction

In the computer science context, *abstraction* entails removing extraneous details to make it easier to focus on the essence of a particular thing. A subway map, for example, doesn't pinpoint where stations are in physical space; instead, it depicts only information that is essential to a rider, such as lines, stops, and proximities to major landmarks. In liminal design, two powerful abstractions are functions and services. Functions are things that need to be done: in the case of medical logistics, providing blood or supplies. Services are the various ways in which those tasks could be accomplished. Framing these activities as services means that they can be provided by any number of sources when needed, and potentially from outside the existing system, resulting in a more flexible and resilient system.

Fortunately, the maintenance unit in the hospital still had enough 3D printer feedstock to make the sterilizer replacement part. The LogC2 system sent specifications and installation directions. Major Evans received notice that the sterilizer would be back up in the next hour.

Once Major Evans requested a sterilizer replacement part, the AI-enabled system could abstract her request by mapping it to a function (supply parts), which in turn could be provided by a capability as a service. In this case, the service was locally manufacturing the part with a 3D printer instead of shipping one from a logistics warehouse or depot.

In the context of our DARPA work, we developed abstractions for functions and services within the ecosystem but extended them to consider new combinations of other building blocks. For example, we studied this example of 3D printing specialized medical instruments or parts for sterilizers instead of using conventional supply chains. We realized that this choice was not about the details of materials manufacturing; it was actually about personnel and knowledge. The flexibility offered by printing medical instruments is offset by the burden of adding people with manufacturing expertise to medical teams. As a result, we needed to provide a way for AI to find options for the higher-level function of “source sterile instruments,” including making use of nearby 3D printers or even local contractors in the host nation. We needed a way to capture personnel skills, procedures, and supply chain information across disparate datasets and models—and then find effective combinations of them.

All of this is a daunting problem using traditional database engineering methods. One of the breakthroughs of our DARPA work was to translate this kind of abstract, cross-domain reasoning into datasets that AI systems can navigate. Using a form of metamathematics called category theory, we were able to map problems from one domain into another. This innovation, as much as any other, makes liminal AI possible.

Composition

The matching of things that might need to be done (functions) in a new situation with the different ways anyone or anything available in that situation might be able to accomplish those things (services) is what we call *composition*, the second core liminal design concept. In a technology platform such as Amazon’s marketplace, for example, the composition is the logic that matches buyers to distribution warehouses to sellers.

Composition is a powerful concept because it provides a way to change what is considered “inside” or “outside” of the system. While exploring options for medical professionals needing more blood supply, we used a technology platform to facilitate composition by balancing demand with the use of drones from a nonmedical unit to supply blood. This provided an unexpected service to satisfy a critical function while using system-level changes, such as combining Air Force and Army medical capabilities, to solve sudden problems, such as the shortage of blood.

Reliefswept over Major Evans’s face as the app on her tablet reported that an AI-guided drone was inbound with 10 units of AB negative blood from a nearby Army medical brigade. “Just in time,” she thought. Unknown to her was that the drone making the delivery wasn’t an Army or an Air Force asset, but instead one owned and operated by an automatically contracted host-nation support company.

Mediation

With the ability to translate knowledge from one domain into another, the simplest form of our third core concept becomes possible. *Mediation* is providing knowledge in context to facilitate a cooperative activity. Imagine an electrical engineer and a mechanical engineer talking about a robot design: the former uses circuit models while the latter uses mechanical drawings. Some combination of both types of knowledge is needed to make the design choice of where to put a motor or wiring harness.

In the blood delivery example, the mediation occurs between a medic and a provider of a service—logistics distribution—that is not normally part of a unit. Unmanned aircraft systems operation and airspace management are not medical functions, yet by using a mediation platform, one could obtain distribution as a service. Instead of explicitly supplying a unit with equipment and the know-how to operate it, having a liminal AI-enabled platform provides the flexibility to add capabilities wherever needed. This approach lets local problem solvers assemble the right solution in real time. In the medical setting, the approach is even more powerful when the system is “open,” allowing host-nation partners or units from other military services to contribute capabilities.

“Sir, we’re going to need to bring the network down and go emissions controlled. We have positive ID on inbound drones that will hit anything emitting radio comms,” reported the intelligence chief to the Command Operations Center’s operations officer. The operations officer looked at Major Mike O’Connell, the logistics chief: “Do the forward units have any way to sustain operations?” Major O’Connell, like all good logisticians, was already ahead of this one.

Mediation becomes particularly important when the conditions of the system change. In the context of modern military conflict, communications are constantly attacked and jammed or used sparingly, if at all, to prevent adversaries from targeting them. One solution to this challenge is to design a hardened, fully integrated system to support centralized command and efficiency. But experience in the pandemic has

revealed that even the most sophisticated control systems were brittle and failed to adapt.

A more resilient approach is to anticipate that the system will be fragmented by communications availability and local teams will need a way to assemble the capabilities needed to accomplish their goals. Yet if the sub-teams were never intended to work together, how can they do so? In Major O’Connell’s case, how might a forward unit combine capabilities from the host nation and nearby units to accomplish the mission if their radios don’t work together and they don’t have any history of collaboration? Our team found that we can use AI to essentially stitch together different components that weren’t intended to work together, whether that involves different units, radio systems, or procedures.

Major O’Connell and his staff had already used a feature of the LogC2 app called the “digital twin” to simulate and game out responses to the scenario they were now experiencing in real life. He knew that in the absence of communications, the forward logistics and medical teams could combine and adapt their capabilities for spikes in demand while drones in the air would act as a virtual warehouse for supplies under constant demand, such as food. More importantly, the combat engineers would still be getting supplies to patch the craters in the airfield via unmanned ground vehicles.

Learning

To adapt to a changing and unpredictable world, a system needs a mechanism first for sensing and interpreting the environment, then for continually updating and honing its responses. Thus *learning* is our final core liminal design concept. Generally, we’ve used feedback loops to facilitate learning within the system, but we’ve also worked to facilitate human learning as well. One way to do that is through training and simulation, as with the “digital twin” feature used by Major O’Connell. The digital twin includes a detailed model based on data obtained from the real system. It can show users how local innovations affect system-level goals by providing feedback on the probable consequences of decisions or changing policies. Running such simulations before a crisis occurs lets decisionmakers rehearse different approaches and priorities and helps strengthen their intuitions about the possible ramifications of their choices. As the wargame demonstrated, such feedback also helps decisionmakers identify when to make big operational changes, such as identifying the conditions that signal a need to shift from the efficient yet vulnerable hub-and-spoke distribution system to a more resilient, dispersed one enabled by drones.

Gaps are everywhere

As powerful as the methodology is in the military operations context, we believe that liminal design could be effectively applied to almost any problem that spans individual insight and previously disparate systems. As a case study in applying liminal design in an altogether different domain, two of us (Main and Russell) led the development of an innovation platform called Polyplexus, which focuses on the challenge of catalyzing innovation by dynamically combining far-flung ideas, perspectives, and capabilities.

The Polyplexus platform allows a diverse community of more than 3,500 researchers, inventors, and citizen scientists to collaborate, share insights, and develop new ideas. So far, they have worked on scientific challenges that include understanding how the vagus nerve functions, repairing the social safety net, and generating local magnetic fields using semiconductors.

Here liminal design principles offer something traditional user-focused or system-focused platforms are unable to, being hemmed in by system boundaries. Instead, liminal design centers on the cross-disciplinary sweet spots where disruptive ideas and technologies are found. Equal parts serious game, social network, research database, and discussion forum, Polyplexus provides a framework for participants, or “plexors,” to make conjectures by applying the liminal design concepts of abstraction and composition and combining evidence with personal experience, knowledge, intuition, and creativity.

For example, a micropublication (a brief, peer-reviewed research result) from the materials science literature might show that a design tool coupled to a 3D printing technique can be used to make a material that has tunable mechanical and thermal properties. This insight could be combined with a micropublication from the electric aircraft and drone community noting an unmet need for materials with specific combinations of weight, porosity, and thermal expansion for higher power-density batteries. The conjecture micropublication would combine these to hypothesize that the 3D-printing technique could be used for the battery problem, and would need additional micropublications to support the hypothesis that the technique is appropriate to the specific sorts of materials and constraints of the battery application. The platform provides a set of rules that define evidence and manage submissions to make the ideas more shareable.

Learning in Polyplexus is captured by the participants. As each plexor shares information and

ideas or participates in a moderated discussion, the result is a promising and purposeful interaction that benefits all participants, as well as society at large.

Engineering the next wave of innovation

These glimpses of the liminal approach to designing adaptive ecosystems are provocative, but there are challenges that need to be overcome before it will be a part of AI designers' toolboxes. The first challenge is developing and defining a vocabulary and the equivalent building blocks and representations for liminal design. We also recognize that although we have adapted powerful ideas from computer science, these may not work for all problems; liminal design needs to incorporate insights from other fields. For example, a key aspect of liminal design is understanding how specific human individuals, their interrelationships, and their collective goals and motives may influence outcomes, which will require integrating cognitive and social science concepts.

Furthermore, the idea of an operating system for a distributed collective intelligence hints at a provocative theory of computation that extends beyond machines to groups of people whose individual and collective abilities are enhanced by a machine platform. This model will require concepts from disciplines as disparate as game theory, behavioral economics, cognitive psychology, and the theory underpinning distributed systems such as cloud services. Taken together, these insights can catalyze a form of human-machine symbiosis, stitching together diverse intelligences and giving people more options in turbulent times.

Aspects of this are beginning to take shape in academic research. The computer science concept of aggregate programming can monitor crowds across thousands of devices, although it has yet to capture human insights. Alternatively, work by the computer scientist Michael Bernstein and collaborators on "flash organizations" provides glimpses of synergistically using AI mediation to blend human and machine intelligence to coordinate complex tasks. Fully unlocking the potential of these ideas will require joint efforts among computer scientists, economists, and social scientists—people who rarely collaborate. But if they did, they might begin to address the core challenge here: capturing the human dimension of these open, adaptive ecosystems and exposing deep and fundamental questions about the nature of human experience, intelligence, and effective collaboration. Such efforts may have a profound impact not only on computer science and AI research but on the systems of governance that underpin how groups of people are able to accomplish collective, everyday goals.

The pandemic has ruthlessly exposed the brittleness of many of the systems supporting society. Yet it has also inspired a wave of collaborative innovation to address those systemic shortcomings. These emerging ecosystems of problem solvers will want to collaborate, share knowledge, and coordinate the use of resources in ways that will catalyze development of entirely new systems for innovation. By focusing on the gaps between system insights and human ingenuity, we are optimistic that a liminal design approach will provide new tools to meet society's most pressing challenges.

John Paschkewitz is a partner and associate director at Boston Consulting Group. He was a DARPA program manager from 2015 to 2021, initiating and transitioning multiple programs on systems and technologies to improve adaptability, resilience, and competitive advantage. Bart Russell is a program manager in DARPA's Defense Sciences Office, focused on understanding the variability of human cognitive and social behavior to empower decisionmakers, improve analytics, and generate autonomous and AI systems that enable human adaptability.

John Main is a research and development veteran with experience in academia, government research management (including DARPA), and industrial product and market development. The authors thank Sean Winkler and Brent Thomas for their contributions to the medical wargames and analysis.

RECOMMENDED READING

- Thomas Both, "Human-Centered, Systems-Minded Design," *Stanford Social Innovation Review* (Mar. 9, 2018).
- John D. Foley, Spencer Breiner, Eswaran Subrahmanian Eswaran, and John M. Dusel, "Operads for complex system design specification, analysis and synthesis," *Proceedings of the Royal Society A* 477, no. 2250 (2021): 20210099.
- Brent Thomas, *Preparing for the Future of Combat Casualty Care: Opportunities to Refine the Military Health System's Alignment with the National Defense Strategy* (Santa Monica, CA: RAND Corporation, RR-A713-1, 2021).
- Brent Thomas, Katherine Anania, Anthony DeCicco, and John A. Hamm, *Toward Resiliency in the Joint Blood Supply Chain* (Santa Monica, CA: RAND Corporation, RR-2482, 2018).
- Melissa A. Valentine, Daniela Retelny, Alexandra To, Negar Rahamti, Tulsee Doshi, and Michael S. Bernstein, "Flash Organizations: Crowdsourcing Complex Work by Structuring Crowds as Organizations," *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017).