

ODAI · OPTIMIZED DECISION-MAKING AI

ODAI Intellectual *Heritage*

What a judgment machine rests on

Heuristka's ODAI (Optimized Decision-making AI) links people and machines in a judgment chain that gives leaders decision advantage at every echelon. Human judgment, augmented by AI, for wicked problems.

“The first, the supreme, the most far-reaching act of judgment that the statesman and commander have to make is to establish by that test the kind of war on which they are embarking; neither mistaking it for, nor trying to turn it into, something that is alien to its nature.”

— CARL VON CLAUSEWITZ, *ON WAR*¹

SUMMARY

Every serious tradition that has studied judgment under uncertainty has reached part of the same answer. Strategists learned that war is a contest of wills fought through fog, and that the decisive act is to understand the kind of war one is in. The codebreakers of Bletchley Park learned to reason from fragments, and learned that a correct warning can fail for want of standing. Intelligence analysts learned that the gravest failures come not from missing information but from information that was present and not credited. Psychologists learned that human judgment fails in systematic, nameable ways, and that expert intuition is reliable only in environments that give honest feedback. Systems thinkers learned that the strongest levers on any system are its feedback and the visibility of its information, not its numbers. Students of military innovation learned that technology alone has never changed warfare. And students of AI safety, from Bletchley Park onward, learned that a machine’s capability can outrun the judgment that governs it, and that keeping the two together is a problem of design.

These traditions have lived in separate buildings. Strategists thought about power, psychologists about bias, designers about framing, analysts about uncertainty. Each produced insights; none produced a working discipline that a staff could use at the moment of decision. The raw materials existed for decades. The field did not.

This paper sets out the lineage behind ODAI, Heuristka’s Optimized Decision-making AI, and argues three things.

1. **ODAI rests on published, testable arguments, not on assertion.** Each element of its design traces to a body of work that a skeptical reader can check. The paper names the works and says what each contributes.
2. **The lineage is coherent.** Read together, these traditions describe one problem – how judgment fails, and how it can be protected – from several sides. The paper shows where they meet, where they disagree, and how the disagreements are used rather than hidden.
3. **What is new is the synthesis, and the instrument that makes it usable.** The ideas are, for the most part, well known. What no staff has had is the capacity to hold them all at once, for every actor, while the problem is moving. That capacity is what artificial intelligence now supplies, and it is why the discipline Heuristka names – *cognitive-heuristic intelligence*, the

systematic improvement of human judgment under uncertainty, complexity and adversarial adaptation – can now be practiced rather than only taught.

The lineage is not only historical. On 20 October 2026, in the National War College’s Kennan Lecture, the Vice Chairman of the U.S. Joint Chiefs of Staff, General Christopher J. Mahoney, drew on paired strategists from Sun Tzu and Clausewitz to George C. Marshall to tell officers preparing for joint command that speed is nothing without judgment, that they must understand the nature of the fight they have rather than the fight they want, and that candor is a duty. The Joint Staff published his remarks on the day Heuristka’s website went live. He identified a problem ODAI was created to fix, and he reached for the traditions this paper sets out (§12).

The paper takes the lineage in ten parts: Clausewitz; Boyd; Hinsley and Bletchley; intelligence analysis and the warning tradition; the science of judgment; systems, complexity and wicked problems; strategic theory; the triple trinity; military innovation and the organization; and AI safety. It closes with the practitioners whose questions work in the room, and with what ODAI unifies. Each section applies the same test – what does this tradition contribute to the judgment a staff must exercise, and what does it leave out? – and each ends with the gap the next tradition fills. The reader who wants the lineage at a glance will find it in the table in §13. The reader who wants to check it will find more than ninety notes. This paper, its companions and the films and podcast that go with them are published on the Heuristka Substack, insights.od-ai.ai.

01

The problem the lineage answers

In October 1962 the President of the United States was pressed by his military chiefs to strike the Soviet missile sites in Cuba and then invade. On the information the chiefs had, the advice was defensible. What they did not know was that ninety-eight tactical nuclear warheads were on the island, mated and mobile, held by commanders who would not wait for Moscow, and aimed at the beaches the invasion would use. The United States learned that thirty years later, from a Soviet general at a conference in Havana.²

The information that might have raised the question existed. American photo analysts had identified nuclear-capable short-range launchers on 25 October; the question of whether they carried nuclear warheads was not raised until after the crisis.³ The CIA had photographed the nuclear-storage vans at the Bejucal bunker three times and dismissed the site because it had only one fence, where an American site would have had two.⁴ The information was there. The orientation was wrong.

That is the problem every tradition in this paper has studied from its own side. It is not a shortage of information; the American government in 1962 had more than any government before it. It is the failure to place information inside an accurate picture of the situation before acting on it – to orient. The traditions name the failure differently: a mistaken judgment of the kind of war (Clausewitz); a broken

orientation (Boyd); a warning without standing (Hinsley); signals lost in noise (Wohlstetter); a bias with a name (Kahneman and Tversky); a wicked problem treated as a tame one (Rittel and Webber); a technology bought without the organization to use it (Murray and Millett). Each is right. None, alone, is enough to build a staff process on.

Nor is the problem one of scale alone, though scale is part of it. The crisis of 1962 involved three governments, three militaries and three peoples, and still the picture in the Cabinet Room was wrong. A small regional conflict of twenty states presents sixty such actors and some 1,770 primary interactions among them, each non-linear and each moving. The traditions that follow were, almost without exception, worked out by individuals thinking about individual minds or single institutions. The paper's last claim is that holding what they learned, for every actor, while the problem moves, is a task of a different order, and that it is the task a machine can now take on under human guidance.

The test this paper applies to each tradition is the same: what does it contribute to the judgment a staff must exercise, and what does it leave out?

02

Clausewitz: the kind of war, and the genius to see it

ODAI's first debt is to Clausewitz, for three ideas and a fourth that is treated later.

The kind of war. The sentence on this paper's first page is the end of an argument, not the beginning of one. "That test," in Clausewitz's phrase, is the political object: war is "merely the continuation of policy by other means," so the kind of war one is in is fixed by the purpose it serves and the means each side will commit to it, and the first act of judgment is to read that purpose correctly on both sides.⁵ He meant it as a test, not a slogan: neither mistaking the war for something it is not, nor trying to turn it into something alien to its nature. From this ODAI takes its first discipline, the distinction between war and warfare. War is the pursuit of a political objective sustained by political will, across every line of operation. Warfare is the activity along any one of them. A force can win every engagement and lose the war, and a staff that counts engagements will not see it happening.⁶ Every question ODAI puts to a plan is, in the end, a form of Clausewitz's first question: what kind of war is this, and does the action serve the purpose? It is still the first lesson the American military gives the officers it is preparing for joint command: in October 2026 the Vice Chairman of the U.S. Joint Chiefs of Staff paired Clausewitz with Eisenhower to tell them to understand the nature of the fight they have, not the fight they want (§12).

Friction and fog. Clausewitz separated two things that are often run together. Friction is the gap between the plan and its execution: "Everything in war is very simple, but the simplest thing is difficult," because countless small obstacles, none foreseeable, accumulate until performance falls short of intention.⁷ Fog is the uncertainty of the picture itself: war is "the realm of uncertainty; three quarters of the factors on which action in war is based are wrapped in a fog of greater or lesser uncertainty."⁸ Neither can be abolished, and Clausewitz did not pretend otherwise. Neither does ODAI. What a machine can do is draw

the fog off at the point of decision — thin it until the things that were always there can be seen — and keep a record of where friction has bent the plan, so that the next turn begins from what happened rather than from what was intended.

Genius. Clausewitz’s answer to fog was the commander’s genius, and especially *coup d’oeil*: “the quick recognition of a truth that the mind would ordinarily miss or would perceive only after long study and reflection.”⁹ Genius, as he describes it, is not a gift of birth alone. It is aptitude joined to what is learned and lived through, refined over a career until the commander sees what others cannot. That definition matters to ODAI for two reasons. It makes genius partly teachable, which is the premise of every war college. And it locates genius exactly where Boyd would later locate orientation: in a picture of the situation built from experience and continually revised. ODAI does not supply genius. It augments the conditions in which genius forms — the breadth of experience a mind can draw on, and the speed with which its picture is corrected — and it does so for the whole staff, not only for the rare officer in whom the gift arrives unaided.

Clausewitz’s trinity, the fourth idea, is treated with Thucydides and Boyd in §9.

What Clausewitz leaves out. He wrote for commanders and statesmen, not staffs, and his remedy for uncertainty was a quality of the individual mind. He had no account of how a staff should be organized to support judgment, no catalog of the systematic errors to which minds are prone, and no method for testing a commander’s picture against the evidence while there was still time to change it. Boyd supplied the first part of the answer: a theory of how that picture is built, and of why it is the center of everything else.

03

Boyd: orientation at the center

John Boyd is the most cited and least read theorist in military AI. The loop that carries his name — observe, orient, decide, act — is routinely drawn as four boxes in a circle and used to make one point: whoever cycles fastest wins. Boyd argued something different, and said so in plain terms. “Orientation is the *Schwerpunkt*. It shapes the way we interact with the environment — hence orientation shapes the way we observe, the way we decide, the way we act.” Without it, there is “no command and control worthy of the name.”¹⁰

The argument had a philosophical root. In the only essay Boyd ever wrote, “Destruction and Creation” (1976), he described thinking as a cycle of taking existing concepts apart and building new ones from the pieces, and he drew on Gödel, Heisenberg and the second law of thermodynamics to argue that any closed system of concepts, left to itself, drifts away from the reality it was built to describe. A mind or an organization that stops testing its picture against the world does not stay merely wrong; it grows more wrong, and more confident, together.¹¹ That is the mechanism behind Hinsley’s rule (§4), stated by a fighter pilot a generation before the formulation.

In his last sketch of the loop, orientation is the largest element, fed by genetic heritage, cultural traditions, previous experience and new information, and worked through analysis and synthesis. Arrows marked *implicit guidance and control* run from it directly to observation and to action.¹² Tempo, for Boyd, was a product of orientation, not a substitute for it. A force gets inside an adversary's loop by acting in ways his orientation cannot absorb, until his picture of the situation no longer matches the situation. Firing faster at an adversary whose orientation is intact does not get inside his loop. It spends ammunition at a higher rate. The scholarly reading of Boyd, above all Frans Osinga's, has made this case at length; the popular reading has not caught up.¹³ Some of the senior leadership has: in October 2026 the Vice Chairman of the U.S. Joint Chiefs of Staff invoked Boyd to warn future commanders that automated systems can validate bad information at speed (§12).

ODAI takes three things from Boyd and adds two.

- **Orientation as the center of gravity.** Everything in ODAI's design follows from it.
- **The moral, mental and physical planes,** ranked in that order, and the account of moral conflict as the magnification of menace, uncertainty and mistrust inside the opponent (§9).
- **Organic command and control.** Boyd's *Organic Design for Command and Control* argued for implicit communication and shared orientation over explicit orders. ODAI's intent refinement – guidance that travels up the chain as well as down – is that idea made into a staff function.

The triangle. ODAI draws Boyd's argument as a triangle rather than a circle: Orient at the center; Observe, Decide and Act at the points; feedback running both ways between Orient and every point. Boyd's own sketch returned feedback to Observe. ODAI routes it through Orient, where his argument says it does its work. The drawing is ODAI's; the logic is Boyd's.


The helix. A loop closes on itself. ODAI opens it into a helix, so that each turn builds on the last: Orient, Deliberate, Act/Audit, Iterate – whose initials are the system's name. In short form the third phase is called Act; the audit travels with the action rather than following it. The idea predates ODAI. In 2023 the Military Innovation Lab published an adaptation of Boyd's loop into a spiral built on dialectical reasoning: no end state, only hypotheses tested, refuted and refined into the next, the spiral standing for growth and the accumulation of wisdom.¹⁴ The helix is that spiral, engineered, with an audit trail that makes each turn's lessons available to the next. It is also "Destruction and Creation" made into a staff procedure: each turn takes the previous picture apart against new evidence and builds the next from what survives. And it has a biological reading, which Boyd's own sketch invites: he listed genetic heritage among orientation's inputs. The helix is a double helix. Its two strands, the human and the machine, are joined at every turn of the human-machine interface, and like DNA the structure carries what it has learned into what comes next: the audit trail is the staff's inheritance.



ODAI

BOYD'S OODA LOOP

Orientation at the center of judgment.



Orientation shapes observation, decision, and action—and is reshaped by feedback.

Heuristka's interpretation of Boyd's written work, which puts Orient at the center.

© 2026 Heuristka LLC · www.od-ai.ai · info@od-ai.ai · 99 Wall St, Suite 1772, New York, NY 10005

Figure 1. Boyd's OODA loop, drawn as a triangle: Heuristka's interpretation of Boyd's written work, which puts Orient at the center, with feedback running both ways between Orient and every point and, around it, the unfolding world. The drawing is ours; the logic is Boyd's.




Figure 2. The ODAI Helix: Boyd's loop opened so that each turn climbs. Orient, Deliberate, Act, Iterate, with the audit trail carrying each turn's lessons into the next. The full account is in the companion paper, *What Is ODAI?*

What Boyd leaves out. Boyd wrote about the individual mind and the organization as an organism. He did not supply a catalog of the specific ways orientation fails, a method for testing it, or a means of carrying orientation across a large staff. The intelligence and decision-science traditions supply the first two. The third had to wait for the machine. And before any of them, a small community in an English country house had already worked out, under the pressure of a war, what reasoning from fragments looks like when it is done well – and what happens when it is done well and ignored.

04

Hinsley and Bletchley: reasoning from fragments

ODAI's epistemic backbone comes from the tradition of Bletchley Park, and from the historian who recorded it, Sir Harry Hinsley.

Hinsley arrived at Bletchley in 1939, an undergraduate of twenty. He made his reputation in naval intelligence, reading the pattern of German radio traffic – who was transmitting, how much and from where – before the naval codes were broken. Young as he was, his colleagues came to call him “the Cardinal,” and his biographer attributes his gift to “an ability to sense that something unusual was afoot from the tiniest clues.”¹⁵

The gift was tested twice in 1940, and twice the warning failed to carry. In the spring, Hinsley's reading of the traffic told him that something was happening in the Baltic just before the German invasion of Norway; the warning went unheeded.¹⁶ Then, on 7 June, he warned the Admiralty that German heavy ships were moving out of the Baltic. No code had been broken; he was reading the pattern of the traffic, not its content, by a method the Admiralty regarded as untested, and it issued nothing. The next day *Scharnhorst* and *Gneisenau* sank the carrier *Glorious* and her two escorting destroyers, and more than 1,500 men died.¹⁷ The warning existed. It failed not on its evidence but on its standing: a new source, an untested method, an institution for which doing nothing was free.

The same method, once it had earned standing, changed the war. By 1941 Hinsley's reading of the German naval traffic was trusted enough to be acted on, and it was his analysis that pointed to the German weather ships in the North Atlantic as isolated vessels that must be carrying Enigma settings. Their capture that summer gave Bletchley the naval keys it had not been able to break by other means.¹⁸ The two episodes together are the whole lesson. A method of warning is only as useful as the standing it has when the warning comes, and standing is built in advance or not at all.

Two further lessons follow, and both are built into ODAI.

The epistemic discipline. Hinsley's later work as official historian of British intelligence in the Second World War showed, across several volumes, how strategic outcomes emerge from the interaction of information, interpretation and institutional behavior.¹⁹ ODAI takes three principles from it: *epistemic triangulation*, the disciplined separation of what is known from what is inferred from what is unknowable;

anomaly interpretation, the reading of weak signals and of deviations from the expected pattern, above all the anomaly that breaks the frame; and *blind-spot diagnosis*, the recognition that the gravest unknown unknowns come from refusing to update a mental model, not from a simple absence of data. ODAI's epistemic frame (§5) is that discipline formalized.

Hinsley's rule. In the formulation of the one of us who was his student: *the gravest strategic threats arise not from missing information but from refusing to update a mental model when the evidence contradicts it*. It is a formulation, not a quotation, and it is the organizing principle of the discipline Heuristka names. Its practical corollary is the *Glorious* lesson. A new source of warning must earn its standing before it is needed – run in shadow beside real planning and scored against outcomes – and when it is overridden, the override must carry a signature.

One practitioner belongs beside Hinsley here. Churchill insisted on reading a selection of Bletchley's raw decrypts himself, rather than only the summaries prepared for him.²⁰ It was not a method a staff could copy. It was a principal's refusal to let the most important material reach him only after others had decided what it meant – the instinct ODAI's provenance discipline tries to make routine.

What Bletchley leaves out. Bletchley was a small, extraordinary community with a single dominant source, and it relied on individuals of rare ability. Its discipline did not travel easily into large staffs, and it was never written down as a method while the war lasted; Hinsley's own history came three decades later. The task of turning that discipline into a profession, with rules that could be taught and judgments that could be checked, fell to the Americans after the war. ODAI is built so that the judgment on watch is not a matter of luck, and the American tradition is where the rules for that come from.

05

Intelligence analysis and the warning tradition

If Bletchley supplied the discipline of reasoning from fragments, the American intelligence tradition supplied its method, and its warning theorists supplied its most important findings.

Sherman Kent made intelligence analysis a profession: explicit, structured and accountable, with judgments stated so that they could be checked. His essay on words of estimative probability tried to make analysts say what “probable” meant in numbers, because a probability that cannot be read is no warning at all.²¹ Kent also set the limit that ODAI accepts as its own first principle. “Whatever the complexities of the puzzles we strive to solve, and whatever the sophisticated techniques we may use to collect the pieces and store them, there can never be a time when the thoughtful man can be supplanted as the intelligence device supreme.”²² He wrote that in 1965, about the computers of his day. ODAI is human-guided AI because Kent was right, and because nothing in the sixty years since has shown him wrong.

Richards Heuer built on Kent and on the heuristics-and-biases research of Kahneman and Tversky to explain how analysts' minds fail, and to offer structured remedies, above all the analysis of competing

hypotheses.²³ The chronology matters: Heuer drew on the psychologists, not the reverse. His contribution was to carry their findings into the analyst's chair and to show that structure could correct for them. His central move – set out every hypothesis, then look for the evidence that would *disconfirm* each, rather than the evidence that supports the favorite – is the opposite of how a busy staff naturally works, which is why it has to be built into the process rather than left to good intentions. ODAI's catalog of structured analytic techniques – more than a thousand of them – descends from that move.

Roberta Wohlstetter's study of Pearl Harbor established the distinction between signals and noise: the warning signals were present before the attack, embedded in a far larger volume of irrelevant information, and the failure lay in separating the one from the other.²⁴ More information does not fix that problem. It usually worsens it. Wohlstetter is the reason ODAI is designed to thin the fog rather than add to it.

Cynthia Grabo wrote the American intelligence community's handbook on strategic warning, and five of her findings are built into ODAI's epistemic frame.²⁵

- *Warning does not exist until it has been conveyed* – “to the policymaker, and ... he must know that he has been warned.” An assessment that never reaches the decision-maker is a stranded known.
- *The minority is usually right.* The most nearly correct judgments in crises have often come from a minority within each agency, so consensus is likeliest to be wrong exactly where being right matters.
- *Warning is cumulative, not current,* and can run out in a lull just before the blow. A dashboard of the latest reports hides the pattern a long memory shows.
- *Adduction:* intelligence asked to support a course already favored selects, without intent to distort, the facts that make it look desirable.²⁶
- *The fear of being wrong:* “the true professional will under-warn,” hedging until the recipient has to read between the lines.

Grabo also left a structured test for whether an actor will use force – five questions on commitment, attainability, capability, the exhaustion of other options, and risk – which ODAI's adversary model puts to the staff.

Barton Whaley's study of strategic deception reached the conclusion Grabo quoted with dismay: the deceiver is almost always successful, regardless of the sophistication of the victim.²⁷ ODAI's catalogs therefore treat deception as a standing condition of the environment, not an exception, and hold its techniques on all three sides – ours, the adversary's and third parties'.

The epistemic frame. Donald Rumsfeld's 2002 distinction among known knowns, known unknowns and unknown unknowns gave the analysis a popular vocabulary.²⁸ ODAI extends it to five categories, because two of the most dangerous states of knowledge are missing from the original: *false knowns*, held as certain and wrong (that Moscow controlled every weapon in Cuba), and *stranded knowns*, known to someone on our own side and never acting on the decision. A known is stranded by a seam, when it never reaches the decision; by doubt, when it arrives and is not credited, as Hinsley's warning was not; or by inconvenience, when it arrives and is not wanted. The last ODAI calls *convenient disbelief*: treating information that is

unwelcome as if it were unreliable. And among the unknown unknowns, ODAI separates those no one could have imagined from those a disciplined challenge could have raised. The ninety-eight warheads were imaginable. The launchers had been photographed.

What the intelligence tradition leaves out. It is a tradition of analysis for decision-makers, written mostly from the analyst's side of the table. It says less about the staff that must turn warning into a plan, or about the decision-maker who must act on it. Its remedies, from Kent's probability scale to Heuer's techniques, have mostly been taught in classrooms and applied beside the work rather than in it. And it took the failures of the mind largely as given. Why the mind fails as it does, and whether its shortcuts are only faults, was a question for the psychologists.

06

The science of judgment

The intelligence tradition knew that analysts erred. The psychologists showed how, and the disagreements among them turned out to be as useful as their findings.

Kahneman and Tversky. Their heuristics-and-biases program demonstrated that human judgment under uncertainty departs from the rules of probability in systematic, predictable ways: availability, representativeness, anchoring, framing, overconfidence and their many relatives.²⁹ The errors are not random noise around a correct answer. They have direction, and they recur – which is what makes them worth cataloging, because an error that recurs in a known direction can be watched for. That finding is the foundation of ODAI's cognitive bias catalog, which holds nearly a hundred entries, and growing, drawn from the general literature and then built out for the military and strategic world. Each entry is a working instrument with five parts: the bias, the distortion it produces, its strategic consequence, the human discipline that counters it, and the function ODAI performs to surface it.

Gigerenzer. Gerd Gigerenzer argued against reading heuristics only as error. Simple rules of thumb, he showed, are often the right tool – “fast and frugal” – when they fit the structure of the environment in which they are used; the question is not whether a heuristic is biased but whether it is ecologically rational.³⁰ Nobody walking home at night computes the odds of a mugging in the dark alley ahead; the heuristic fires and they cross the street, and on most nights it is the right call. ODAI takes the correction seriously. Its catalog does not treat every shortcut as a fault. It asks whether the shortcut fits the problem in front of the staff, and the name *Heuristka*, from the same Greek root as *heuristic*, is a deliberate nod to the view that heuristics can be refined, taught and applied systematically rather than only suppressed. The disagreement between Gigerenzer and Kahneman is not resolved in the literature, and ODAI does not pretend to resolve it. It uses both: Kahneman's catalog to name the shortcut, Gigerenzer's test to ask whether, here, it is a shortcut or a mistake.

Klein, and the reconciliation. Gary Klein studied experts deciding under pressure – firefighters, commanders, nurses – and found that they rarely compare options. They recognize the situation, imagine

one course of action playing out, and act unless the simulation fails.³¹ Kahneman and Klein began on opposite sides of the question of when intuition can be trusted and ended in agreement: intuition is reliable only in environments regular enough to be learned, where the expert gets prompt and honest feedback on his judgments.³² That conclusion is a working principle inside ODAI. Strategy is rarely such an environment. Its feedback arrives late, ambiguous and contested, and the expert's intuition is therefore tested rather than trusted – not dismissed, but put to the question by the catalog and the techniques before it reaches the commander. Klein also gave the staff one of its most useful instruments, the pre-mortem: assume the plan has failed, and ask everyone in the room to write down why. It converts the reluctance to dissent into a task, and it is one of the techniques ODAI proposes when a plan has reached the point of looking settled.

Tetlock. Philip Tetlock's long study of expert political judgment found most forecasters barely better than chance, and the most confident big-idea thinkers among them, his "hedgehogs," often worse. His later work showed that judgment can be measured and improved: forecasters who think in probabilities, update in small steps and keep score become measurably better.³³ ODAI's audit trail is Tetlock's scorekeeping applied to the staff. Every assessment is recorded with its stated probability and scored against what happened, so that the staff learns how far to trust its own judgments and the machine's. It is also the answer to Kent's complaint: a probability that is written down and scored is one that can be read.

Janis. Irving Janis named groupthink: the tendency of a cohesive group under pressure to seek agreement at the expense of realistic appraisal, with its signs of self-censorship, illusory unanimity and the suppression of dissent.³⁴ His cases were drawn from American foreign policy, and the Bay of Pigs was among them – the same administration that, eighteen months later, handled the missile crisis with a deliberately adversarial process, the President absent from some meetings so that his advisers would argue. Grabo's minority-view principle is the intelligence corollary of Janis. ODAI answers both by keeping the disfavored hypothesis audible and by making the record, not the room, the arbiter of which judgments held.

Simon. Herbert Simon's bounded rationality is the theoretical root of the claim this paper's last sections depend on. Decision-makers, he showed, cannot optimize. They operate within limits of information, time and cognitive capacity, and so they satisfice – they search until they find an option good enough, and stop.³⁵ Organizations exist in large part because individual rationality is bounded; they divide the problem so that no one mind must hold all of it. Simon was also one of the founders of artificial intelligence, an author of one of its first working programs.³⁶ The two halves of his career meet in ODAI. The staff was the organizational answer to bounded rationality. The machine now offers a second answer: not to replace the bounded mind, but to widen the bounds within which it searches, so that the option it settles on is good enough against a far larger field than any staff could survey.

Allison. Graham Allison's three models of governmental decision – the rational actor, organizational process and governmental politics – remain the clearest account of how governments actually decide, and ODAI's adversary model uses all three.³⁷ His account of the 1962 crisis itself was first written two decades

before the record this paper relies on – the tactical warheads, the submarines, Bejucal – was declassified, and the revised edition could take in only part of it. It is cited here for its models, not its history.

The laboratory evidence on human and machine. Recent experiments add a finding none of these scholars could have tested. People working with a biased AI become more biased themselves, and are unaware of how far it moved them; working with an accurate AI, their judgment improves. Under pressure, many adopt the machine's answer rather than constructing their own.³⁸ The failure is miscalibrated trust in either direction, and the design of the system decides which way it goes. That is why ODAI's interface is continuous rather than a gate the human passes through once, and why it asks of its subject-matter experts the stance it calls creative skepticism: neither deferring to the machine nor dismissing it.

What the science of judgment leaves out. The psychologists described individual minds, mostly in controlled settings and mostly one decision at a time. They said less about adversaries, who exploit biases deliberately, or about institutions, which reward some errors and punish their correction. And they said almost nothing about the structure of the problem itself – why a strategic situation does not hold still while it is being judged. That is the subject of the systems tradition.

07

Systems, complexity and wicked problems

Meadows. Donella Meadows ranked the places to intervene in a system, from the weakest to the strongest. Adjusting its numbers – budgets, quotas, rates – is near the bottom. Near the top are the strength of its feedback loops, the structure of its information flows, the rules and goals that govern it, and the paradigm from which it arises.³⁹ Two consequences shape ODAI. The helix is built on feedback, because feedback is a stronger lever than any number a staff can change. And ODAI's principal means of changing behavior is the one Meadows ranked high and most reformers skip: making visible information that was invisible. Punishment secures compliance only while someone is watching. Incentives work only as well as they were designed. People act differently once they can see what they could not see before, without a new rule or a new threat. ODAI does not need to argue a commander into a correction. It only needs to show what was already there.

Senge. Peter Senge carried systems thinking into organizations, describing the learning organization: one that surfaces its mental models, holds a shared vision, and learns as a team rather than only as individuals.⁴⁰ Two of his instruments are close cousins of ODAI's. The first is the *ladder of inference*, which Senge took from Chris Argyris: the steps by which a mind moves from observable data, through selection and interpretation, to assumptions, conclusions and beliefs, each rung taken faster than the last and the lower rungs forgotten by the time the top is reached.⁴¹ Kent's separation of fact from inference from estimate is the same ladder, drawn from the analyst's side. ODAI asks every judgment to show its rungs. The second is the learning organization itself. The audit trail is Senge's learning organization given a

memory: the next long war is fought as one war, not as a succession of one-year wars, because what each rotation learned is on the record for the next.

Holland and the Santa Fe tradition. John Holland’s work on complex adaptive systems showed how coherent behavior emerges from many interacting agents following simple rules, how such systems adapt in ways no single agent intends, and why their behavior resists prediction even when every rule is known.⁴² Three of his points bear directly on strategic analysis. Agents in such a system do not merely interact; they anticipate one another, and change their behavior on the strength of what they expect, so that the system’s state includes everyone’s model of everyone else. Small changes propagate in ways that are not proportionate to their size. And the system’s regularities exist, but at the level of pattern rather than prediction – the kind of thing an experienced observer learns to recognize rather than compute. Strategic analysis is complex-adaptive-system modeling of this kind, with the difference that its agents are human, and a small regional conflict of twenty states, each with a people, a government and a military, presents some 1,770 primary interactions. No staff can hold them in view. That is the strongest single argument for a machine in the staff. It is also, by Holland’s own account, the strongest argument for why the machine must be guided by people who know which patterns matter, because the system will not yield to being computed.

Rittel and Webber. The design theorists Horst Rittel and Melvin Webber distinguished tame problems from wicked ones. A tame problem has a definitive formulation and a solution that can be checked. A wicked problem has neither: “there is no definitive formulation,” so that formulating the problem is the problem, and such problems are never solved; “at best they are only re-solved – over and over again.”⁴³ Every attempt at a solution changes the problem, and there is no way to test a solution without implementing it. Science solves tame problems, and so does AI. Strategy is wicked, and so nothing in it can be guaranteed: there is no perfect information, and the social sciences have no mathematics to settle what a government or a market will do. What a discipline can do is improve the odds. ODAI’s founders put the distinction in a sentence: the science of war is a tame challenge; the art of war is a wicked opportunity. ODAI holds the science and prompts the art.

Muller. Jerry Muller’s study of the tyranny of metrics supplies the caution that keeps the science from swallowing the art: not everything that is important is measurable, and much that is measurable is unimportant.⁴⁴ Worse, once a measure becomes a target, the people being measured begin to manage the measure rather than the thing it was meant to stand for. The American system is a master of measurement, and its characteristic failure is to measure what is easy – targets struck, sorties flown – and mistake it for what matters. ODAI is measured in effectiveness, not performance, and it asks of every metric what it measures, against what baseline, and whether it has become a target.

Rosling and Carr. The temperament the discipline asks for is neither optimism nor pessimism, which are emotional states, but what Hans Rosling called possibilism: seeing the world as it is, hoping and fearing only with reason.⁴⁵ Its older relative is E. H. Carr’s insistence, in *The Twenty Years’ Crisis*, that realism without ideas is barren and idealism without realism is dangerous – that analysis must begin with the world as it is and then imagine how it might be changed.⁴⁶ Carr’s narrow path between the two – realism

to understand the world as it is, enough idealism to envisage a better one that can realistically be reached – runs through the whole of ODAI’s method; one of us wrote his doctoral thesis on Carr’s work. ODAI’s method oscillates between the two: the catalogs establish what is, and the structured techniques open what might be. Carr is also the guide to how ODAI’s audit trail reads a decision afterward. A historian, he argued in *What Is History?*, needs imaginative understanding of the people he studies – of what they knew and why they acted – rather than moral judgment of them. The audit trail works the same way: it keeps the information available at the moment and the reason for the choice, stated at the time by the person or office that made it, so that it explains rather than exposes, and protects people from hindsight rather than presenting them for it.

What the systems tradition leaves out. It describes systems from outside, and often in peacetime. Its agents adapt, but they do not, as a rule, lie. It has less to say about an adversary who is deliberately trying to deceive the observer, to break the system being observed, and to read the observer’s own intentions in order to defeat them. That is the subject of strategic theory.

08

Strategic theory: interaction, perception and purpose

The strategic theorists of the nuclear age added what the systems tradition lacked: an opponent who thinks. Their work falls into three movements, and ODAI takes something from each.

Interaction: Schelling and Jervis. Thomas Schelling made strategy the study of interdependent decision: each side’s best move depends on what it expects the other to do, and much of the contest is conducted through signals, commitments and the manipulation of risk.⁴⁷ His most unsettling idea was “the threat that leaves something to chance” – the recognition that in a crisis neither side fully controls the escalation it is threatening, and that the loss of control is itself part of the bargaining. That is Clausewitz’s chance, residing in the military, read as an instrument of policy; and October 1962, with its unauthorized U-2 shoot-down and its submarine captain out of contact with Moscow, is the case Schelling’s readers have returned to ever since. ODAI’s adversary model and its war-gaming function are Schellingian. They test a plan not against a static target but against an opponent who is reading the plan and will respond, and they ask what the plan does to the risks neither side controls.

Robert Jervis showed how perception and misperception drive international politics: decision-makers see what their beliefs prepare them to see, draw lessons from the wrong analogies, and assume others understand their intentions as they do.⁴⁸ Mirror-imaging, the assumption that the adversary reasons as we do, is Jervis’s finding in its most dangerous form, and it is the first error ODAI’s adversary model exists to prevent. Jervis also set out the choice that ODAI’s strategic-inaction analysis is built to inform. Under what he called the deterrence model, restraint is read as weakness and invites aggression; under the spiral model, firmness is read as threat and provokes the very response it was meant to deter. Both are true of

some adversaries at some moments, and the whole question is which one applies now. A staff that asks only “can we strike it?” has assumed the deterrence model without noticing. ODAI puts the question.

Purpose under new conditions: Brodie and Kennan. Bernard Brodie saw at the start of the nuclear age that a new technology had changed the logic of strategy: “Thus far the chief purpose of our military establishment has been to win wars. From now on its chief purpose must be to avert them.”⁴⁹ His lesson for ODAI is methodological. When the nature of warfare changes, the logic of decision must change with it; new conditions need new theory, not retrofitted doctrine. A military that meets a new technology by making the old process faster has not absorbed Brodie’s point, whether the technology is the atomic bomb or the machine that now reads the battlefield.

George Kennan’s analysis of Soviet conduct showed what understanding an adversary on his own terms looks like: history, ideology, insecurity and the logic of a regime, read from the inside before any policy is proposed. The Long Telegram located the Kremlin’s conduct in a “traditional and instinctive Russian sense of insecurity” and in the uses Marxist ideology served for a regime that needed an external enemy – an account of motive, not of capability, and one that explained behavior the capability counts could not.⁵⁰ ODAI’s adversary model is built to that standard, within the limits every such model has, which it states.

What strategy is: Gray, Freedman and Nye. Colin Gray described strategy as a bridge between policy and military power, and insisted that culture, geography and politics are not context but substance; strategy is difficult, he argued, not because its principles are obscure but because it must be made to work through people, institutions and an enemy, all of them with wills of their own.⁵¹ Lawrence Freedman’s history of strategy treats it as the art of creating power, often through the stories actors tell about how a situation will unfold – scripts that guide action and that fail when the situation departs from them.⁵² Hinsley’s rule is the warning against holding to a script after the evidence has left it. Joseph Nye’s account of soft power, the capacity to attract rather than coerce, extends strategy to legitimacy and influence,⁵³ and so to the gray actors – third parties neither ally nor adversary – where ODAI’s analysis says most wars are decided and least planned for.

Sun Tzu. The oldest text in the lineage sets the highest standard: “supreme excellence consists in breaking the enemy’s resistance without fighting.”⁵⁴ It is also the definition of the threat the next section describes.

What strategic theory leaves out. It is a literature of insight for the individual strategist, and most of it treats the state as the actor. It does not tell a staff how to hold all of it in view, for every actor and for the parts inside each actor, while the problem is moving. The next section takes three of the oldest ideas in the lineage and shows how ODAI does exactly that.

09

The triple trinity

Strategic analysis rivals climate modeling in its non-linearity, with one difference: its variables are human. A small regional conflict of twenty states, each with a people, a government and a military, presents some 1,770 primary interaction vectors: sixty actors, each paired with every other. The number matters less than where its terms come from. The three actors inside each state, and the forces that move them, are Clausewitz's. What each actor wants is Thucydides'. The planes on which an action lands are Boyd's. These are three triads from three traditions, across twenty-four centuries, each built to describe the part of war that does not yield to counting. ODAI reads them together. We call the combination the triple trinity.

Where the thinking began

The idea was first set out in 2023, in the Military Innovation Lab's guide to building a creativity culture in a military staff, as a **double trinity**: Clausewitz and Thucydides read together as a lens for possibilist analysis.⁵⁵ Every actor in the international system, it argued, is made up of three parts – the people, the military and the government – subject, respectively, to passion, chance and reason, and each also subject to fear, honor and interest. Its working rule was the one that still matters most: never assume the people, the government and the military are aligned. Governments often fear their own people; militaries can feel threatened by either. What the people fear, honor or count as their interest may be far from the government's estimate, and the military's may differ from both. The gaps are where the strategic insight lies, and they move. A sudden coup in a stable state reveals tensions that conventional analysis had dismissed.

Boyd's three planes of conflict make the double trinity triple. They add what the first two lacked: an account of what an action does when it lands.

Clausewitz: the actors and their forces

Clausewitz closes the first chapter of *On War* with what he calls a *wunderliche Dreifaltigkeit*, which Howard and Paret translate as a paradoxical trinity. War is more than a chameleon that changes its character to suit each case. It is composed of three tendencies: primordial violence, hatred and enmity, which he treats as a blind natural force; the play of chance and probability, within which the creative spirit is free to roam; and war's subordination to policy, which makes it subject to reason alone.

He then places each tendency in a part of the state. "The first of these three aspects mainly concerns the people; the second the commander and his army; the third the government."⁵⁶

ACTOR	FORCE	NATURE
The people	Passion: primordial violence, hatred and enmity	A blind, non-rational force rooted in the population's emotions and its will to fight.

ACTOR	FORCE	NATURE
The military	Chance and probability: the play of creative spirit	The realm of uncertainty, luck and creative judgment, where the commander works through the fog of war.
The government	Reason: subordination to policy	The rational, instrumental force that keeps war subordinate to policy and directs violence toward a political purpose.

The three forces are connected. Clausewitz pictured the theorist's task as keeping an object suspended among three magnets, and warned that a theory which ignored any one of them, or fixed an arbitrary relationship among them, would be useless. A shift in one moves the others. A government's reason is pressed by the people's passion; an army's chance can overturn a government's calculation; a people's passion can be inflamed, or cooled, by what the army does and the government says. ODAI reads all three forces together, actor by actor, and treats their balance as something that moves with events.

October 1962 shows why.

- **What nearly decided the war was chance, residing in the military**, as Clausewitz said it would. Soviet officers in Cuba shot down a U-2 without Khrushchev's authority, and a submarine captain out of contact with Moscow ordered a nuclear torpedo readied.⁵⁷
- **What ended it was reason, reasserted in two governments**, by two men who had seen the last war and chose a blockade, a back channel and a secret compromise over the strike.

A model that read the governments alone would have seen neither the danger nor the escape.

Thucydides: what each actor wants

In 432 BCE, Athenian envoys at Sparta answered the charge of empire by naming the motives that had driven it: fear, honor and interest (1.75.3; 1.76.2). The triad has since been read as the realist's account of what moves states, and Richard Ned Lebow built a cultural theory of international relations on it.⁵⁸

Thucydides did not offer it as a law. It is one party's self-description in a quarrel, and the scholarship warns against making it universal:

- Honor in the *History* is a distinctively Athenian preoccupation, which Sparta does not share in the same way.
- The order of the three motives differs between the two passages.
- Recent textual work argues that honor in the first passage is a later insertion, added to make it match the second.⁵⁹

That contingency is what makes the triad usable. ODAI does not assume that every actor is moved by all three motives in the same proportion. It asks which motive carries most for this actor at this moment – economic survival in one state, national honor in another, fear of encirclement in a third – and which of the three actors within the state is most strongly moved. That is the work of the adversary model, and it is the guard against mirror-imaging the model exists to provide.

Thucydides also supplies the line that hybrid war has turned upside down. At Melos the Athenians told the islanders that “the strong do what they can and the weak suffer what they must.”⁶⁰ In the gray zone between war and peace the inversion is often closer to the truth: the weak do what they can, and the strong suffer what they must. A weaker actor with drones, proxies and a narrative can impose costs that a stronger one, bound by law, alliance and public opinion, cannot answer in kind. The adversary model asks, for every actor, which way round the Melian rule runs in this case.

Thucydides has no motive that corresponds to Clausewitz’s chance. He did not ignore chance: fortune runs through the *History*, nowhere more than at Pylos. But chance is not a reason anyone gives for acting. The gap is instructive, and we return to it below.

Boyd: where the effects land

Boyd’s briefing *Patterns of Conflict* set out to expose the nature of moral, mental and physical conflict, and ranked the three in that order of importance.⁶¹

- **Physical:** materiel, terrain and kinetic force.
- **Mental:** cognition, and the orientation that shapes observation, decision and action.
- **Moral:** the bonds that let an organization exist as one whole, and its will to go on.

Boyd broke the moral plane into a triad of its own. Moral conflict works by magnifying menace, uncertainty and mistrust inside the opponent, so that fear, anxiety and alienation overcome the courage, confidence and esprit that held the opponent together. Boyd argued that physical attrition alone is as likely to stiffen an enemy as to break him: a bomb does not sever a moral bond, and often tightens it.

What Boyd described as a mechanism now exists as a weapon. Cognitive warfare destroys an opponent’s resolve through narrative. Its aim is to convince the other side that it has no hope — that resistance will only make things worse — so that it gives up before it tries to resist. At its most complete, the target adopts the attacker’s political objectives as its own, and violence becomes unnecessary. It is Sun Tzu’s supreme excellence pursued by narrative instead of maneuver. Its modern form — the manipulation of social media, psychometric targeting, and content manufactured at volume for people and for the machines that feed them — works on exactly the plane Boyd ranked first. Its munitions are aimed at the moral bonds themselves: trust in institutions, in neighbors and in the shared facts on which a democracy argues. They magnify menace, uncertainty and mistrust inside a whole society, in peacetime as in war, degrade the norms that hold it together, and shift what it counts as acceptable behavior. The effect on any single vote or decision is hard to measure and much argued over; the corrosion is not the kind of effect that shows up in a single measurement. Boyd’s ranking was right, and it now cuts against us as well as for us. The moral plane is where the most consequential fighting of this century is being done, on both sides of the line between war and peace, and it is the plane a kill chain cannot see at all.

Boyd supplies what the other two do not: an operational vocabulary for the moral plane. Clausewitz tells the staff that passion exists and where it lives. Thucydides tells it what the passion is about. Boyd tells it what an action does to that passion once it lands — whether the action is a strike, a sanction or a lie

engineered to spread. That is why ODAI's catalogs hold the techniques of cognitive warfare and deception: blue, to see when they are aimed at us; red, to read the adversary's; gray, to understand the populations both are trying to move.

Reading the three together

	CLAUSEWITZ (ACTORS AND FORCES)	THUCYDIDES (MOTIVES)	BOYD (PLANES)
The will (least calculable)	The people: passion	Fear, honor	Moral
The contest of minds	The military: chance and creative genius	—	Mental
The calculable	The government: reason	Interest	Physical

The rows run from the least calculable to the most. The alignment is real but not exact, and the analysis lives in the misfits.

- **The will row holds two of Thucydides' three motives.** This is his quiet verdict that states act mostly from drivers they later dress up as calculation.
- **The middle row has no motive, because chance is not chosen.** It is the row in which the staff's own orientation is tested.
- **The bottom row is the only one the kill chain was built to see.** Targets struck, sorties flown and munitions expended are all physical measures in the calculable row. The top two rows are where wars are decided, and where conventional models have least to say.

Read as three lists, the triads are a seminar exercise. Read as three axes, they form a structure a staff can run against every actor in the system. For each state, and for each of the three actors within it, the structure asks three questions:

1. What is the balance of passion, chance and reason, and which way is it moving?
2. Which motive — fear, honor or interest — is driving this actor now, and is it the same motive that drives the other two?
3. On which plane will a given action land, and what will it do to this actor's will?

The second question is the 2023 rule in working form. The answer for the people and the answer for the government are seldom the same, and the difference is often the most important finding the analysis produces.

Where it sits

The triple trinity lives in Orient, at the center of the ODAI triangle. Boyd drew orientation as fed by genetic heritage, cultural traditions, previous experience and new information. That is the qualitative material from which a picture of the situation is built, and the material most staffs leave to instinct.

The triple trinity gives that material structure without pretending to put numbers on it. ODAI does not score will, legitimacy or honor. It makes their configuration visible and asks the staff to weigh it, which is the translation layer the discipline exists to supply.

Feedback in the triangle runs both ways between Orient and every point, so the balance is read again on every pass. Clausewitz's magnets do not hold still, and neither does the model.

What the triple trinity leaves out. Nothing in it is new, and that is the point. Each triad is well known. What is new is holding them together, for every actor, while the problem is moving, which no staff has had the capacity to do – and which no staff will be organized to do until the staff itself changes. That is the subject of the next section.

 10

Military innovation and the organization

The traditions so far explain how judgment works and how it fails. None of them explains why an institution that knows all this still buys the machine and leaves the rest alone. The historians of military innovation do.

Murray and Millett. Williamson Murray and Allan Millett's study of military innovation between the world wars is the closest thing the field has to a controlled experiment. The major powers faced the same new technologies – the tank, the aircraft, the radio, the aircraft carrier, radar – in the same two decades, and drew very different results from them.⁶² The Germans and not the French made armored warfare out of the tank; the Americans and the Japanese, more than the British who invented it, made a strike weapon out of the aircraft carrier; the Marine Corps, almost alone, made amphibious assault a doctrine rather than a hope; the British, who were behind in most things, built the integrated air-defense system that won the summer of 1940, because they had a defined problem and organized radar, fighters and ground control around it. The finding that runs through the volume is that technology decided nothing. The forces that innovated successfully combined new technology with new organization and new doctrine, on a specific operational problem they had chosen to solve, in a culture willing to examine its own performance honestly – in exercises, war games and after-action study that were allowed to produce unwelcome results. Those that bought the machines without the rest got the machines and nothing else, and in 1940 the French had more and better tanks than the Germans.

Guderian and the sequence. The tank arrived on the Somme in September 1916, and it was not Blitzkrieg: the tank is a thing, and Blitzkrieg is what was done with it. It took until 1939 for an army to work out how to use it to decisive effect, and the officer credited with working it out was a signals officer, not a tank officer. Heinz Guderian's insight was that tanks and aircraft could be linked by the technology he knew, small mobile voice radio, so that a mass of armor and the aircraft above it could be commanded as one moving force.⁶³ He then had to fight the cavalry to get the technology adopted, and the Germans built a new organization around it, the panzer division. The organization found new tactics in the field.

After Poland, a campaign that was still superb, the army reviewed its own performance and retrained before France, and only after many iterations did what worked become doctrine. Even the name came afterward. ODAI's founders read the episode as four stages – technology, organization, tactics, and what works, written down (then called doctrine; ODAI calls it the innovation catalog) – multiplied by iteration; the ordering is theirs, not Murray and Millett's. Without new tactics there are no new operations, and without new operations no new strategy. AI is the twenty-first century's tank, and on its own it will do what the tank did on the Somme.

The staff that has not changed. The organization in question is older than any of the technologies. The staff system every modern headquarters runs was built by Berthier for Napoleon: a commander's intent converted into orders, orders descending through echelons, reports ascending, and the work divided into bureaux by function – the ancestors of the numbered staff sections on every joint staff today. The Prussians systematized it and added the education; the Americans adopted it; the reforms of 1986 made it joint.⁶⁴ Each reform changed one leg at a time – the organization, the education, what the staff ran on – and none could change the one thing that fixed the staff's capacity: the number of minds in the room and the hours in the day. Every advance in information made that ratio worse. It is the leg AI changes, and the reason the other two can now change with it.

From doctrine to the innovation catalog. A line widely attributed to General Mattis calls doctrine the last refuge of the unimaginative.⁶⁵ He was right about the misuse: doctrine hidden behind, treated as scripture, enforced against a world that has moved on. Done properly, the fourth stage is the record of what innovation has shown to work, so that the next army does not have to discover it again. ODAI calls it the innovation catalog: machine doctrine for the moment, constantly evaluated, updated and improved, and fed directly by every after-action review. The Ukrainians have written theirs continuously since 2022, on the ground, by finding what works and keeping it.

The war game as laboratory. Between the wars the Naval War College fought the war with Japan in its game rooms many times before it was fought once. Nimitz said afterward that the war with Japan had been gamed there so often and in so many ways that almost nothing in it came as a surprise – except the kamikaze.⁶⁶ The interwar games were used to speculate about the future, not to certify a narrative or a weapons program, and they were allowed to be lost. War-gaming since has too often been used to confirm what the sponsor already believed. In the Millennium Challenge of 2002, Paul Van Riper, the retired Marine general playing the adversary, sank much of the fleet in the opening days by means the script had not anticipated; the game was stopped, the fleet refloated and the script restored.⁶⁷ ODAI's war-gaming function exists to keep the adversary free to win, because a game the red side cannot win teaches nothing the sponsor did not already believe.

Operational design. The design tradition taught at the Army's School of Advanced Military Studies, and the practice of reframing that Van Riper championed in the Marine Corps, treat the commander's first task as understanding the problem rather than solving it, and expect the frame to be revised as the situation moves.⁶⁸ Design meets Rittel and Webber here. There is no end state in a wicked problem, only

the next state, and constant refinement, never perfection, is the essence of both operational design and ODAI's helix.

Adhocracy. Warren Bennis coined the word *adhocracy* in 1968 for organizations assembled around problems rather than functions, and Alvin Toffler made it famous two years later.⁶⁹ The idea suits the staff AI makes possible. The Napoleonic staff divides the work by function and echelon because, without a machine to hold the science of war, numbers of officers were the only way to hold it. With the machine, a planning cell can be small – assembled for the problem, at whatever echelon the problem sits, and disbanded when it is solved: a small staff with the expertise of a large one. The research on group size points the same way: beyond a handful of people, each addition costs a group more in coordination than it adds in thought.⁷⁰

Open mode and closed mode. John Cleese's lecture on creativity distinguished the open mode, relaxed, playful and tolerant of ambiguity, from the closed mode, purposeful and under pressure, and argued that good work needs both, in sequence: open to find the answer, closed to execute it.⁷¹ MILab's creativity-culture series built a staff method on the distinction, calling the second the action mode.⁷² Military staffs live almost entirely in the action mode. ODAI's Deliberate step is where the open mode is protected, and the triangle's feedback is how the staff returns to it when the action mode has stopped producing results.

The profession. Samuel Huntington defined a profession by three things: expertise, responsibility and corporateness.⁷³ The definition explains why ODAI's catalogs are curated rather than scraped. Expertise is what the profession knows and others do not; responsibility is to the society it serves; corporateness is the profession's own judgment of what counts as competent work. The catalogs are built by the profession, for the profession, and a machine that learned its war from the open internet would hold none of the three.

What the innovation literature leaves out. It explains why technology fails without organization and doctrine, but it was written about machines that did not think. It does not say what happens when the new technology is itself an instrument of judgment, or how to keep the human in charge of a tool that can reason faster than the staff it serves. That question has a tradition of its own, older than the laboratories now racing to answer it.

11

AI safety: capability and control

The question has been asked since the machine was invented, and it was first asked, like so much else in this paper, at Bletchley Park.

Turing and Good. Alan Turing, who led the attack on the German naval Enigma in Hut 8, whose decrypts Hinsley's naval section read, put the question in a lecture in Manchester in 1951. Once machines began to think, he argued, it would not take long for them to outstrip our feeble powers: "At some stage therefore we should have to expect the machines to take control."⁷⁴ His statistical assistant in Hut 8, I. J. Good, was

among the subtlest thinkers Bletchley produced. With Turing he worked out how the codebreakers should weigh fragmentary evidence, in a unit Turing called the ban, and after the war Good set the method out in *Probability and the Weighing of Evidence*.⁷⁵ It is the arithmetic beneath reasoning from fragments, and beneath what Tetlock would later find in the best forecasters: updating in small steps, each step the weight of a new piece of evidence. In 1965 Good gave Turing's warning its modern form. "The first ultraintelligent machine is the last invention that man need ever make," Good wrote, "provided that the machine is docile enough to tell us how to keep it under control."⁷⁶ The proviso is the whole of the AI-safety field in a clause. It also locates the answer where Bletchley located its own: not in the machine alone, but in the relationship between the machine and the people who must judge what it tells them.

Wiener. Norbert Wiener, the founder of cybernetics, supplied the warning that heads ODAI's AI-safety paper. If we use a machine "whose operation we cannot efficiently interfere" with once it has started, because its action is too fast and irrevocable for us to intervene, then we had better be quite sure that the purpose put into it is the purpose we really desire, "and not merely a colorful imitation of it."⁷⁷ Two ideas in that sentence run through ODAI's design. The first is that speed and irreversibility together set the price of a wrong purpose, which is why scrutiny in ODAI scales with the consequence of a decision and with how hard it is to reverse. The second is that the purpose written down is never quite the purpose intended – the gap later researchers would find, again and again, in machines that gamed the objective they were given. Machines, like metrics, are relentlessly literal (§7, on Muller).⁷⁸

The ironies of automation. The engineers and psychologists who studied automated control rooms, cockpits and plants reached a finding that the AI debate is only now rediscovering. Lisanne Bainbridge called it an irony: automation leaves the operator the tasks the designer could not automate, while taking away the practice that kept the operator able to do them, so that "the more advanced a control system is, so the more crucial may be the contribution of the human operator."⁷⁹ Raja Parasuraman and Victor Riley sorted the ways people go wrong with automation into misuse, disuse and abuse: trusting it too much, trusting it too little, and building it without regard for the people who must work beside it.⁸⁰ John Lee and Katrina See made the remedy a design goal: trust calibrated to what the system can actually do, neither more nor less.⁸¹ And Berkeley Dietvorst and his colleagues showed the other half of the problem: people lose confidence in an algorithm faster than in a human after seeing both make the same mistake, even when the algorithm is the better forecaster.⁸² The laboratory studies in §6 show the same failure in systems that talk back, and ODAI's answer, creative skepticism, is calibration made into a habit.

The existential-risk tradition. In 2003 Martin Rees, the Astronomer Royal, argued that this century's technologies put civilization at risk from error as well as from malice. In 2012 he co-founded the Centre for the Study of Existential Risk at Cambridge with the philosopher Huw Price and Jaan Tallinn, a founding engineer of Skype, and in 2014 Tallinn co-founded the Future of Life Institute. Nick Bostrom defined the control problem: how to keep a system more capable than its makers pointed at what they actually want. Stuart Russell gave it the formulation closest to ODAI's. The danger, he argued, lies in the standard model of AI itself, a machine optimizing a fixed objective, because the objective written down is never quite the one intended; his remedy is machines that are uncertain about what people want, and

therefore defer to them, ask before acting and allow themselves to be switched off.⁸³ From this tradition ODAI takes its central insight, that control is a design problem, decided by what a system is built to pursue and by whether people remain able to correct it.

From restraint to engineering. The tradition divides on the remedy. Some of its members ask the frontier laboratories to pause or to pace themselves. ODAI's founders doubt that restraint can carry the weight. Restraint rarely holds when breaking it pays – the naval limits agreed in Washington in 1922 lasted until Japan gave notice in 1934, and lapsed at the end of 1936 – and where restraint has held, it held because it was verified and enforced.⁸⁴ The governance that works in other dangerous fields is engineered and inspected: aviation's safety investigation, candid because it is privileged, and the regulator with power to ground an aircraft.⁸⁵

The rules now being written for AI point the same way. Europe's AI Act requires that the people overseeing high-risk systems be able to remain aware of automation bias; the U.S. Defense Department requires that autonomous weapon systems be designed to allow appropriate levels of human judgment over the use of force; and civil society's Pro-Human AI Declaration asks that people keep the authority and capacity to understand, guide and override the systems they use.⁸⁶ Each states a requirement. None says how judgment is to be kept whole at the moment of decision, when the machine is fluent, fast and usually right. And as Edsger Dijkstra observed of software, testing can show the presence of a flaw, never its absence. Testing alone cannot show a system safe; it has to be kept correctable.⁸⁷

What ODAI takes. Read together, these traditions describe the layer where ODAI works. Russell's deferential machine is a design principle for the model; ODAI supplies the other half, at the layer around the model where people and machines decide together. That is the difference between a human in the loop and human-guided AI. A human in the loop is a position in the machine's process, and it erodes as the machine grows more capable, because the person at the checkpoint defers more – Bainbridge's irony, at the scale of AI. Human-guided AI is an epistemology of human cognitive dominance over the machine: a theory of how knowledge is formed, tested and owned, in which the person, not the model, holds authority over what counts as known. The person frames the problem and the machine tests the frame; the machine questions the person and the person questions the machine; and the person decides what the machine's evidence is worth, and records why.

Wiener's two warnings become two disciplines: scrutiny scaled to consequence and irreversibility, and a purpose stated before the machine is set to work and tested against what it then does. A warning cannot be cleared with a click, and the record keeps each decision as it was made, on the information of the time, with the reason given by the person or office that made it. Good's proviso is met not by a docile machine but by people who remain able to judge it. ODAI's AI-safety paper, *Judgment Built In*, develops the argument.⁸⁸

What the AI-safety tradition leaves out. It has studied the machine – its objectives, its capabilities, whether it can be controlled – far more than the judgment it is meant to serve. The human-factors school studied the operator at a console; it did not study a staff deciding a wicked problem against an adversary who adapts. The existential-risk tradition asks how to keep the machine from doing harm; it says less

about how the people using it can be made better at their own work. Neither offers a theory of the decision itself, of how judgment fails in a room and how it can be protected there, which is what the rest of this paper supplies. The practitioners in the next section show what the human brings to that room.

12

The practitioners

Theory earns its place in a staff by producing questions that work in the room. A handful of practitioners have left questions of that kind.

Powell. Colin Powell's instruction to every intelligence briefer is a whole method in one line: tell me what you know, tell me what you don't know, then tell me what you think – and always distinguish which is which.⁸⁹ It is Kent's separation of fact, inference and estimate, and Argyris's ladder, made into a habit of command. ODAI puts the instruction to every briefer, human or machine.

Rumsfeld. Rumsfeld's known knowns, known unknowns and unknown unknowns (§5) became a joke and remain a method. The hardest category is the third, because finding it requires imagination, and the 9/11 Commission's verdict on the system was that imagination was exactly what it had run short of.⁹⁰ Imagination can be prompted. A machine shaped by people who have sat in those rooms can put in front of a planner the question no one thought to ask, the case no one remembered and the consequence no one traced.

Churchill. Churchill's habit of reading Bletchley's raw decrypts himself (§4) had costs as well as benefits; his commanders did not always welcome a prime minister reading their intelligence and pressing them on what it meant. Its principle is ODAI's all the same: the decision-maker should be able to see the evidence under the judgment, and the judgment should be able to stand up to his questions.

Eisenhower. Eisenhower's remark that "plans are worthless, but planning is everything" explains why a staff process matters more than its product.⁹¹ The plan is overtaken on contact. What survives is the understanding the planning built, and the staff's ability to replan from it. ODAI's audit trail keeps that understanding after the plan is gone.

Rochefort and Layton. In May 1942 the codebreakers at Pearl Harbor could read only fragments of Japanese naval traffic. Joseph Rochefort, who ran the station, and Edwin Layton, Nimitz's fleet intelligence officer, reasoned from the fragments to an explicit call – Midway the objective, the first week of June the date, four carriers the force – against the view in Washington that the blow would fall elsewhere and later. To settle the argument they had Midway report in the clear that its water plant had failed, and watched the Japanese signal that "AF" was short of water. Layton gave Nimitz the date, the bearing and the hour; Nimitz positioned his three carriers on the strength of it, and the Japanese lost four.⁹² The intelligence was never complete. The orientation was forward-looking, the hypothesis was tested, and the probability call was made rather than hedged into uselessness. It is the warning tradition, Hinsley's rule and Boyd's

orientation working together, in one of the decisive weeks of the century — and it is the case that shows what Hinsley’s method does when the warning has standing. Nimitz had decided to trust his intelligence officers before the fragments arrived.

Mahoney. The practitioners’ questions are still being put, and by the people whose job it is to put them. On 20 October 2026, in the National War College’s Kennan Lecture, General Christopher J. Mahoney, Vice Chairman of the U.S. Joint Chiefs of Staff, gave officers preparing for joint command four insights, each drawn from a pair of strategists.⁹³ The first was that speed is nothing without judgment. He invoked Boyd to warn that automated systems can validate bad information at speed, deepen cognitive bias and take in an adversary’s deception, and he reminded his audience that algorithms are not accountable; commanders are. “Speed does not connote understanding,” he said. “Do not ever outsource your judgment ... Because if you’re making the wrong decisions, you’re just losing faster.” The second paired Clausewitz with Eisenhower: understand the nature of the fight you have, not the fight you want. The fourth drew on Marshall: candor is a duty, and execution an obligation. “Tell the truth as you see it. Not what the boss necessarily wants to hear, not what makes the briefing go a little smoother.” Boyd’s orientation, Clausewitz’s first act of judgment and the briefer’s duty to say what the room does not want to hear are ideas this paper traces to their sources; here they were put to the next generation of joint commanders by the second-ranking officer in the armed forces. The Joint Staff published the lecture on the day Heuristka’s website went live. He identified a problem ODAI was created to fix: how to keep that judgment whole when machines can process faster than any staff can understand. That it was the Kennan Lecture is fitting. Kennan’s standard for understanding an adversary from the inside (§8) is the one ODAI’s adversary model is built to. The correspondence between the lecture and ODAI’s argument runs almost point for point, and much of it was in print sixteen days earlier, in *Smart War, Fast War; What Is ODAI?* sets the three side by side in an appendix.

What the practitioners leave out. They had the questions and, at their best, the judgment to use them. They did not have the capacity to put every question to every actor while the problem moved. Their successes depended on rare individuals in the right chair at the right moment. A discipline cannot depend on that.

13

What ODAI unifies

Each tradition in this paper contributes a part of ODAI’s design, and each leaves something out that another supplies.

TRADITION	WHAT IT CONTRIBUTES	WHERE IT LIVES IN ODAI
Clausewitz	The kind of war; friction and fog; genius; the trinity	The first question of Orient; the four layers; the triple trinity

TRADITION	WHAT IT CONTRIBUTES	WHERE IT LIVES IN ODAI
Boyd	Orientation as the center; destruction and creation; the moral plane	The ODAI triangle and the helix: Orient, Deliberate, Act/Audit, Iterate
Hinsley and Bletchley	Reasoning from fragments; the anomaly that breaks the frame; warning given standing	The epistemic frame; Hinsley's rule; a new source of warning given standing before it is needed, and every override signed
Kent, Heuer, Wohlstetter, Grabo, Whaley	Fact, inference and estimate; analysis of competing hypotheses; signal and noise; intent; deception	The uncertainty-mapping engine; Grabo's five questions; the deception catalog; the minority view kept audible
Kahneman and Tversky, Gigerenzer, Klein, Tetlock, Janis, Simon	Named biases; ecological rationality; tested intuition; the pre-mortem; calibration; groupthink; bounded rationality	The cognitive bias catalog; the decision-discipline engine; the scored audit trail
Meadows, Senge, Holland, Rittel and Webber, Muller, Rosling, Carr	Feedback and visibility; the ladder of inference and the learning organization; emergence; wicked problems; the limits of metrics; possibilism	The helix's feedback; the audit trail as memory; measures of effectiveness over measures of performance
Schelling, Jervis, Brodie, Kennan, Gray, Freedman, Nye, Sun Tzu	Interdependent decision; the deterrence and spiral models; new conditions, new theory; the adversary on his own terms; legitimacy	The adversary model; blue, red and gray; the strategic-inaction engine
Clausewitz, Thucydides and Boyd together	Actors and forces, motives, planes	The triple trinity, in Orient
Murray and Millett, Guderian, design, Bennis and Toffler, Cleese, Huntington	Technology, organization and doctrine; the sequence (ours): technology, organization, tactics, the innovation catalog; framing; the staff built around the problem; the open mode; the profession	The case for reform; ODAI as the platform for the new staff; the human-guided interface; curated catalogs
Turing and Good, Wiener, Bainbridge, Parasuraman and Riley, Lee and See, Dietvorst, Rees, Bostrom, Russell, Dijkstra	The purpose put into the machine; control as a design problem; the ironies of automation; calibrated trust; the deferential machine	Human-guided AI, an epistemology of human cognitive dominance; scrutiny scaled to consequence; flags that cannot be clicked away; the record
Powell, Rumsfeld, Churchill, Eisenhower, Rochefort and Layton	Questions that work in the room	Every briefing ODAI supports

Where the traditions disagree, and what ODAI does with it. The summary promised that the disagreements would be used rather than hidden, and the paper has kept a list. Kahneman and Gigerenzer disagree about whether a heuristic is a fault; ODAI names the shortcut with the one and tests its fit with

the other. Kahneman and Klein disagreed about intuition and settled on a condition; ODAI applies the condition and finds strategy mostly fails it. Thucydides has no motive for Clausewitz's chance; ODAI keeps the empty cell, because it marks the row in which the staff's own orientation is tested. Boyd drew a circle; ODAI draws a triangle and says whose drawing is whose. Allison's models are kept and his history is set aside. Mattis's line on doctrine is quoted as a caution and then turned around. Jervis's two models of the adversary contradict each other, and the contradiction is the question ODAI asks. Part of the existential-risk tradition would pause the frontier; ODAI doubts that restraint which rewards whoever breaks it can hold, and engineers judgment in wherever AI is already used. Russell builds deference into the model; ODAI builds human cognitive dominance into the layer around it, and treats the two as complements. A lineage in which everyone agreed would be a lineage no skeptic should trust.

The four layers. The traditions also sort by altitude. At the top is the world system, volatile, uncertain, complex and ambiguous, which the systems thinkers and the strategic theorists describe. Below it is the statesman, who sets the purpose; Clausewitz and Brodie speak to him. Below that is the commander, who turns purpose into design; Boyd, the design tradition and the practitioners speak to him. At the bottom is contact, where the design meets reality, and where Hinsley's fragments and Grabo's indicators first appear. ODAI reads all four layers on every pass, because a failure at any one of them propagates to the others.

The three engines. The lineage converges on three things ODAI does. *Uncertainty mapping* separates the knowns and the known unknowns from the generators of unknown unknowns, and searches for the false knowns and the stranded knowns – the knowledge that is in the building but has not reached the decision. *Decision discipline* tests the coherence of political aims, timing against speed, and action against inaction, putting the staff's reasoning through the catalogs and the structured techniques before it reaches the commander. *Strategic-inaction analysis* closes the blind spot in every conventional framework by assessing the consequences of the null option: whether to act at all, and what not acting would cost. All three are human-guided. The machine proposes; the profession decides. ODAI does not compute answers. It computes the conditions under which answers can be trusted.

Why now. The traditions in this paper have been available for decades, some for centuries. They were never joined into a working discipline, because no staff could hold them all at once, for every actor, while the problem moved. Simon explained why: the mind is bounded, and the staff was the organizational answer to the bound. Artificial intelligence offers a second answer. It can hold the science of war – the catalogs of what is known, the record of what has happened, the techniques for testing a judgment – and put it in front of the people who must practice the art. That is the division ODAI is built on: the science of war at every planner's fingertips through the machine; the art of war done by humans, iterating with one another up and down the chain. Murray and Millett supply the warning that goes with the opportunity. A technology adopted without the organization to use it, and the tactics that organization finds and the record of what works, has never changed warfare, and AI will be no exception.

The discipline. Heuristka calls the discipline cognitive-heuristic intelligence: the systematic improvement of human judgment under uncertainty, complexity and adversarial adaptation. Its name joins the Greek *heuriskein*, to discover, to uncover what is hidden, with *-ka*, for the craftsman. It joins two lineages, which

the founders carry: the epistemic tradition of Bletchley Park, passed from Hinsley to his student, and the practitioner tradition of military operations. Its creed begins with the name: *Uncover what is hidden. Augment judgment. The commander decides.*

The judgment chain is old. Clausewitz described its first link, Hinsley lived its hardest test, and Boyd drew its center. What is new is an instrument that lets a staff practice the whole of it at once.

NOTES

Notes

1. Carl von Clausewitz, *On War*, ed. and trans. Michael Howard and Peter Paret (Princeton: Princeton University Press, 1976), bk. 1, chap. 1, §27, pp. 88–89.
2. Gen. Anatoly Gribkov's disclosure at the Havana conference of January 1992 is recorded in James G. Blight, Bruce J. Allyn and David A. Welch, *Cuba on the Brink: Castro, the Missile Crisis, and the Soviet Collapse* (New York: Pantheon, 1993). The count of ninety-eight (eighty cruise-missile warheads, twelve for the Luna rockets and six bombs) follows the Soviet records cited there and by Dobbs; National Security Archive, Electronic Briefing Book 397 (16 October 2012), notes that U.S. policymakers learned of the Luna warheads only in 1992.
3. National Security Archive, EBB 397: U.S. photo analysts identified FROG (Luna) launchers on 25 October 1962; a DIA report of 29 October noted them "without speculating on the possibility of tactical nuclear deployments," and a memorandum of 2 November records Gen. Maxwell Taylor and DCI John McCone considering the possibility that nuclear warheads had been deployed with them.
4. Michael Dobbs, *One Minute to Midnight: Kennedy, Khrushchev, and Castro on the Brink of Nuclear War* (New York: Knopf, 2008), on the Bejucal reconnaissance passes and the CIA's assessment.
5. Clausewitz, *On War*, bk. 1, chap. 1, §24, p. 87.
6. *On War and Warfare: Definitions and Consequences* (National Security Desk, NSD/DEF-2026-001, 2026).
7. Clausewitz, *On War*, bk. 1, chap. 7, p. 119.
8. Clausewitz, *On War*, bk. 1, chap. 3, p. 101.
9. Clausewitz, *On War*, bk. 1, chap. 3, p. 102.
10. John R. Boyd, "Organic Design for Command and Control" (briefing, May 1987), slides 16 and 25.
11. John R. Boyd, "Destruction and Creation" (essay, 3 September 1976), in *A Discourse on Winning and Losing* (Maxwell AFB: Air University Press, 2018).
12. John R. Boyd, "The Essence of Winning and Losing" (briefing, June 1995).
13. Frans P. B. Osinga, *Science, Strategy and War: The Strategic Theory of John Boyd* (London: Routledge, 2007).
14. *How to Build a Creativity Culture II* (National Security Desk, Military Innovation Lab, 12 June 2023).
15. Richard Langhorne, "Francis Harry Hinsley, 1918–1998," *Proceedings of the British Academy* 120 (2003): 263–74, at 263.
16. Langhorne, "Francis Harry Hinsley," 263.
17. F. H. Hinsley et al., *British Intelligence in the Second World War*, vol. 1 (London: HMSO, 1979); Bletchley Park Trust, record for F. H. Hinsley. Hinsley's own account of the episode is on video at www.od-ai.ai.
18. Hinsley et al., *British Intelligence in the Second World War*, vol. 1, on the captures of the weather ships *München* (May 1941) and *Lauenburg* (June 1941); F. H. Hinsley and Alan Stripp, eds., *Codebreakers: The Inside Story of Bletchley Park* (Oxford: Oxford University Press, 1993).
19. F. H. Hinsley et al., *British Intelligence in the Second World War*, 5 vols. (London: HMSO, 1979–90).

20. David Stafford, *Churchill and Secret Service* (London: John Murray, 1997); Hinsley et al., *British Intelligence in the Second World War*, vol. 1.
21. Sherman Kent, “Words of Estimative Probability,” *Studies in Intelligence* 8, no. 4 (1964).
22. Sherman Kent, *Strategic Intelligence for American World Policy*, preface to the 1965 edition, p. xviii, as quoted in Jack Davis, “Sherman Kent and the Profession of Intelligence Analysis,” CIA Sherman Kent Center, *Occasional Papers* 1, no. 5 (2002).
23. Richards J. Heuer Jr., *Psychology of Intelligence Analysis* (Washington, DC: CIA Center for the Study of Intelligence, 1999).
24. Roberta Wohlstetter, *Pearl Harbor: Warning and Decision* (Stanford: Stanford University Press, 1962).
25. Cynthia M. Grabo, *Anticipating Surprise: Analysis for Strategic Warning*, ed. Jan Goldman (Washington, DC: Joint Military Intelligence College, 2002): warning conveyed and the minority view, chap. 1; the five questions on intent, chap. 5; the lull, chap. 6; cumulative warning and “the true professional will under-warn,” chap. 9.
26. Thomas G. Belden, “Indications, Warning, and Crisis Operations,” *International Studies Quarterly* 21, no. 1 (1977): 181–98.
27. Barton Whaley, *Stratagem: Deception and Surprise in War* (Cambridge, MA: MIT Center for International Studies, 1969; Boston: Artech House, 2007), as quoted in Grabo, *Anticipating Surprise*, chap. 7.
28. Donald H. Rumsfeld, Department of Defense news briefing, 12 February 2002.
29. Amos Tversky and Daniel Kahneman, “Judgment under Uncertainty: Heuristics and Biases,” *Science* 185, no. 4157 (1974): 1124–31; Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).
30. Gerd Gigerenzer, Peter M. Todd and the ABC Research Group, *Simple Heuristics That Make Us Smart* (New York: Oxford University Press, 1999).
31. Gary Klein, *Sources of Power: How People Make Decisions* (Cambridge, MA: MIT Press, 1998); Klein, “Performing a Project Premortem,” *Harvard Business Review*, September 2007.
32. Daniel Kahneman and Gary Klein, “Conditions for Intuitive Expertise: A Failure to Disagree,” *American Psychologist* 64, no. 6 (2009): 515–26.
33. Philip E. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton: Princeton University Press, 2005); Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown, 2015).
34. Irving L. Janis, *Victims of Groupthink* (Boston: Houghton Mifflin, 1972).
35. Herbert A. Simon, *Administrative Behavior* (New York: Macmillan, 1947); Simon, “A Behavioral Model of Rational Choice,” *Quarterly Journal of Economics* 69, no. 1 (1955): 99–118.
36. Allen Newell and Herbert A. Simon, “The Logic Theory Machine,” *IRE Transactions on Information Theory* 2, no. 3 (1956): 61–79. The program was written with J. C. Shaw.
37. Graham T. Allison, *Essence of Decision: Explaining the Cuban Missile Crisis* (Boston: Little, Brown, 1971); 2nd ed., with Philip Zelikow (New York: Longman, 1999).
38. Moshe Glickman and Tali Sharot, “How Human–AI Feedback Loops Alter Human Perceptual, Emotional and Social Judgements,” *Nature Human Behaviour* 9 (2025): 345–59; *AI Is Not the Problem* (National Security Desk, MIL/AI-2026-002, 2026), §VII, and the studies cited there.
39. Donella H. Meadows, *Leverage Points: Places to Intervene in a System* (Hartland, VT: The Sustainability Institute, 1999); Meadows, *Thinking in Systems: A Primer* (White River Junction, VT: Chelsea Green, 2008).
40. Peter M. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization* (New York: Doubleday, 1990).
41. Chris Argyris, *Overcoming Organizational Defenses* (Boston: Allyn and Bacon, 1990); Peter M. Senge et al., *The Fifth Discipline Fieldbook* (New York: Doubleday, 1994), where the ladder of inference is set out for practitioners.
42. John H. Holland, *Hidden Order: How Adaptation Builds Complexity* (Reading, MA: Addison-Wesley, 1995).
43. Horst W. J. Rittel and Melvin M. Webber, “Dilemmas in a General Theory of Planning,” *Policy Sciences* 4, no. 2 (1973): 155–69.
44. Jerry Z. Muller, *The Tyranny of Metrics* (Princeton: Princeton University Press, 2018).
45. Hans Rosling, with Ola Rosling and Anna Rosling Rönnlund, *Factfulness* (New York: Flatiron, 2018).

46. E. H. Carr, *The Twenty Years' Crisis, 1919–1939* (London: Macmillan, 1939); Carr, *What Is History?* (London: Macmillan, 1961).
47. Thomas C. Schelling, *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960), chap. 8, “The Threat That Leaves Something to Chance”; Schelling, *Arms and Influence* (New Haven: Yale University Press, 1966).
48. Robert Jervis, *Perception and Misperception in International Politics* (Princeton: Princeton University Press, 1976), chap. 3 on the deterrence and spiral models.
49. Bernard Brodie, ed., *The Absolute Weapon: Atomic Power and World Order* (New York: Harcourt, Brace, 1946), 76.
50. George F. Kennan, telegram from Moscow, 22 February 1946 (the “Long Telegram”), part 2; “X” [Kennan], “The Sources of Soviet Conduct,” *Foreign Affairs* 25, no. 4 (July 1947).
51. Colin S. Gray, *Modern Strategy* (Oxford: Oxford University Press, 1999); Gray, *The Strategy Bridge: Theory for Practice* (Oxford: Oxford University Press, 2010).
52. Lawrence Freedman, *Strategy: A History* (New York: Oxford University Press, 2013).
53. Joseph S. Nye Jr., *Soft Power: The Means to Success in World Politics* (New York: PublicAffairs, 2004).
54. Sun Tzu, *The Art of War*, trans. Lionel Giles (London: Luzac, 1910), chap. 3.
55. *How to Build a Creativity Culture I* (National Security Desk, Military Innovation Lab, 19 May 2023), “Double Trinity Lens in Possibilism”; also *Overview: All-in-One Introduction to MILab* (2023).
56. Clausewitz, *On War*, bk. 1, chap. 1, §28, p. 89.
57. National Security Archive, *The Submarines of October*, Electronic Briefing Book 75 (2002); Dobbs, *One Minute to Midnight*.
58. Thucydides, *History of the Peloponnesian War* 1.75.3 and 1.76.2; Robert B. Strassler, ed., *The Landmark Thucydides* (New York: Free Press, 1996); Richard Ned Lebow, *A Cultural Theory of International Relations* (Cambridge: Cambridge University Press, 2008).
59. Daniel P. Tompkins, “‘Honor, Fear and Profit’: Non-Universal Terms in Thucydides,” *Classical Association of the Middle West and South*, 2008; Jeffrey S. Rusten, “Honour, Fear, and Benefit—In That Order: The Interpolation of Τιμή in 1.75.3 of Thucydides’ Speech of the Athenians at Sparta,” *Histos* 18 (2024): 126–39.
60. Thucydides 5.89, in the Crawley translation as revised in Strassler, *The Landmark Thucydides*.
61. John R. Boyd, “Patterns of Conflict” (briefing, December 1986), in *A Discourse on Winning and Losing* (Maxwell AFB: Air University Press, 2018); Osinga, *Science, Strategy and War*.
62. Williamson Murray and Allan R. Millett, eds., *Military Innovation in the Interwar Period* (Cambridge: Cambridge University Press, 1996), especially Murray on armored warfare, Geoffrey Till on the aircraft carrier, Millett on amphibious warfare, Alan Beyerchen on radar, and Murray’s concluding chapter.
63. Heinz Guderian, *Achtung – Panzer!* (1937), trans. Christopher Duffy (London: Arms and Armour, 1992); Guderian, *Panzer Leader* (London: Michael Joseph, 1952). On the German army’s review of its own performance after Poland, Williamson Murray, “The German Response to Victory in Poland: A Case Study in Professionalism,” *Armed Forces & Society* 7, no. 2 (1981): 285–98. On the term, Karl-Heinz Frieser, *The Blitzkrieg Legend: The 1940 Campaign in the West* (Annapolis: Naval Institute Press, 2005), 4–6.
64. Dallas D. Irvine, “The Origin of Capital Staffs,” *Journal of Modern History* 10, no. 2 (1938): 161–79; J. D. Hittle, *The Military Staff: Its History and Development*, 3rd ed. (Harrisburg, PA: Stackpole, 1961); Goldwater–Nichols Department of Defense Reorganization Act of 1986, Pub. L. 99-433. The argument is developed in *What Is ODAI?* (Heuristka, 2026), §8.
65. The line is widely attributed to Gen. James Mattis. On his view of study and doctrine, Jim Mattis and Bing West, *Call Sign Chaos: Learning to Lead* (New York: Random House, 2019).
66. Chester W. Nimitz, address to the Naval War College, Newport, 10 October 1960.
67. Micah Zenko, “Millennium Challenge: The Real Story of a Corrupted Military Exercise and Its Legacy,” *War on the Rocks*, 5 November 2015; Zenko, *Red Team: How to Succeed by Thinking Like the Enemy* (New York: Basic Books, 2015).
68. Headquarters, Department of the Army, ATP 5-0.1, *Army Design Methodology* (Washington, DC, July 2015); Joint Chiefs of Staff, Joint Publication 5-0, *Joint Planning*.

69. Warren G. Bennis and Philip E. Slater, *The Temporary Society* (New York: Harper & Row, 1968); Alvin Toffler, *Future Shock* (New York: Random House, 1970).
70. Marcia W. Blenko, Michael C. Mankins and Paul Rogers, *Decide & Deliver: Five Steps to Breakthrough Performance in Your Organization* (Boston: Harvard Business Review Press, 2010); J. Richard Hackman, *Leading Teams* (Boston: Harvard Business School Press, 2002).
71. John Cleese, lecture on creativity in management, Grosvenor House, London, 23 January 1991; Cleese, *Creativity: A Short and Cheerful Guide* (London: Hutchinson, 2020).
72. *How to Build a Creativity Culture*, I–II, and the accompanying syllabus (National Security Desk, Military Innovation Lab, 2023).
73. Samuel P. Huntington, *The Soldier and the State* (Cambridge, MA: Harvard University Press, 1957), chap. 1.
74. Alan Turing, “Intelligent Machinery, A Heretical Theory” (lecture, Manchester, c. 1951), first printed in Sara Turing, *Alan M. Turing* (Cambridge: Heffer, 1959); reprinted in *Philosophia Mathematica* 4, no. 3 (1996): 256–60.
75. I. J. Good, *Probability and the Weighing of Evidence* (London: Charles Griffin, 1950); Good, “A. M. Turing’s Statistical Work in World War II,” *Biometrika* 66, no. 2 (1979): 393–96.
76. I. J. Good, “Speculations Concerning the First Ultrainelligent Machine,” *Advances in Computers* 6 (1965): 31–88, at 33. On Good in Hut 8 with Turing, Jack Copeland, ed., *The Essential Turing* (Oxford: Oxford University Press, 2004).
77. Norbert Wiener, “Some Moral and Technical Consequences of Automation,” *Science* 131, no. 3410 (6 May 1960): 1355–58.
78. Jack Clark and Dario Amodei, “Faulty Reward Functions in the Wild,” OpenAI, 21 December 2016; Victoria Krakovna et al., “Specification Gaming: The Flip Side of AI Ingenuity,” DeepMind, 21 April 2020; METR, “Recent Frontier Models Are Reward Hacking,” 5 June 2025.
79. Lisanne Bainbridge, “Ironies of Automation,” *Automatica* 19, no. 6 (1983): 775–79, at 775.
80. Raja Parasuraman and Victor Riley, “Humans and Automation: Use, Misuse, Disuse, Abuse,” *Human Factors* 39, no. 2 (1997): 230–53.
81. John D. Lee and Katrina A. See, “Trust in Automation: Designing for Appropriate Reliance,” *Human Factors* 46, no. 1 (2004): 50–80.
82. Berkeley J. Dietvorst, Joseph P. Simmons and Cade Massey, “Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err,” *Journal of Experimental Psychology: General* 144, no. 1 (2015): 114–26.
83. Martin Rees, *Our Final Century* (London: Heinemann, 2003); Centre for the Study of Existential Risk, University of Cambridge, www.cser.ac.uk; Future of Life Institute, *Asilomar AI Principles* (2017); Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014); Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* (New York: Viking, 2019).
84. Treaty for the Limitation of Naval Armament, signed at Washington, 6 February 1922; Japan gave notice of termination on 29 December 1934, and the treaty’s limits lapsed on 31 December 1936. The argument is developed in *Judgment Built In* (Heuristka, 2026), §1.
85. *Machin v. Zuckert*, 316 F.2d 336 (D.C. Cir. 1963), on the privilege that protects military aviation’s safety investigations.
86. Regulation (EU) 2024/1689 (the AI Act), Article 14(4)(b); U.S. Department of Defense, Directive 3000.09, *Autonomy in Weapon Systems* (25 January 2023), §1.2.a; Future of Life Institute, *Pro-Human AI Declaration* (March 2026), futureoflife.org/pro-human-ai-declaration.
87. Edsger W. Dijkstra, *Notes on Structured Programming*, EWD249 (2nd ed., April 1970).
88. *Judgment Built In: AI Is Not the Problem. Capability without Judgment Is* (Heuristka, 2026).
89. Colin Powell, *It Worked for Me: In Life and Leadership* (New York: Harper, 2012).
90. National Commission on Terrorist Attacks upon the United States, *The 9/11 Commission Report* (New York: Norton, 2004), chap. 11.
91. Dwight D. Eisenhower, remarks at the National Defense Executive Reserve Conference, 14 November 1957.

92. Edwin T. Layton, with Roger Pineau and John Costello, *“And I Was There”: Pearl Harbor and Midway – Breaking the Secrets* (New York: Morrow, 1985).
93. Maj. Joshua J. Pena, USMC, “Vice Chairman Urges Future Commanders to Prioritize People over Platforms at NDU Address,” Joint Chiefs of Staff, 21 October 2026, on the Kennan Lecture at the National War College, National Defense University, 20 October 2026. The quotations are as the Joint Staff reported them; the four insights are its summary.

Published on the Heuristka Substack, insights.od-ai.ai.

© 2026 Heuristka LLC · www.od-ai.ai · info@od-ai.ai · 99 Wall St, Suite 1772, New York, NY 10005

Heuristka™, ODAI™ and the ODAI Helix™ are trademarks of Heuristka LLC.



A new category of AI. *For judgment.*

Human judgment, augmented by AI, for wicked problems.

REQUEST A BRIEFING info@od-ai.ai

PAPERS, FILMS AND THE PODCAST insights.od-ai.ai