
AHMED HIMMI

On the Iran conflict & the Pentagon's Maven AI system

Who Chose When the Machine Recommended the Kill

*An essay on military AI, accountability,
and the dissolution of moral responsibility*

2026

The principle established at Nuremberg was precise and, at the time, revolutionary: following orders does not absolve the person who carries them out. The soldier who executed the command was responsible for the execution. The defense of obedience — *I was told to do this by someone with authority over me* — was rejected as morally and legally insufficient. The person who acted bore accountability for the action.

This principle now has a new problem. The orders are coming from machines.

Not entirely. Not in the crude sense of a robot issuing commands to a human soldier. What is actually happening is more nuanced and in some ways more difficult: a system called Maven, developed by the US Department of Defense and now operational in active conflict, processes satellite imagery, signals intelligence, and historical targeting data to recommend strikes. A human sits at a terminal, receives the recommendation, and decides whether to approve it. The human approves.

The question that the Nuremberg principle does not yet know how to answer is: who chose?

The history of warfare is partly a history of increasing distance between the person who decides and the person who dies.

The soldier with a sword made contact. They felt the resistance of the act. The archer was farther away — the arrow traveled, the outcome

arrived somewhere they could not fully see. The artillery gunner was farther still — the shell landed miles away, among people the gunner would never see and in circumstances they could not directly observe. The drone operator sits in a different country from the person on their screen, with a signal delay between intention and impact.

Each step in this progression accomplished something military — range, precision, safety for the operator. Each step also accomplished something psychological: it reduced the immediate human cost of killing for the person doing the killing. Not the moral cost. Not the long-term cost for the soldier who had to live with what they did. But the immediate, visceral, in-the-moment cost of ending a life.

That cost had a function. It was a deterrent — not on the individual level alone, but at the systemic level. Societies that go to war while feeling the weight of what war is go to war differently than societies that go to war from a clean distance. The psychological price of violence is part of what makes violence expensive enough to avoid. Remove the price, and the calculation changes.

Maven removes the price at a new level. It does not just create distance between the operator and the target. It creates distance between the operator and the judgment. The question used to be: can I do this? Now it is: should I approve what the system determined I should do?

Approval and choice are not the same thing.

This is the specific problem that AI recommendation in lethal decision-making creates, and it is worth being precise about. When you

choose something, you produce the judgment. You assess the information, weigh the considerations, arrive at a conclusion, and act on it. The cognition and the action belong to you. When you approve something, you ratify a judgment produced elsewhere. You confirm that you do not object. You lend your authorization to a conclusion you did not generate.

The distinction feels subtle. In practice it is enormous.

The person approving a machine's recommendation is operating under conditions that are structurally different from the person who chose. They are working with a confidence rating — a percentage, a score — that represents the machine's certainty, not their own. They are working under time pressure, because the target is time-sensitive and the system is optimized for speed. They are working inside an institutional context in which the system's recommendations have historically been reliable, which means that disapproving requires overcoming not just doubt but a track record that argues against doubt.

What does accountability look like for a person operating under these conditions? They technically chose. They pressed the button. No one forced them. But the judgment they ratified was not theirs, the certainty they acted on was the machine's, and the institutional weight of the system's reliability was pressing on them in the direction of approval.

This is not the defense of *I was following orders*. It is something for which we do not yet have precise language. It is closer to: I was following a recommendation that I had every institutional reason to trust and insufficient time and information to independently verify.

Nuremberg does not cover this. Neither does any existing framework for military accountability.

The people who built Maven were not trying to make war cheaper. They were trying to make it more precise.

The argument for AI targeting systems is substantively serious: human judgment under combat stress is error-prone. Cognitive load, emotional arousal, time pressure, incomplete information — these produce the mistakes that end in strikes on hospitals, wedding convoys, columns of refugees. A system trained on vast data, without fear or fatigue, without the specific distortions of a person who has been awake for forty hours in a war zone — that system might produce better judgments than the human it replaces.

This argument may even be correct, in the narrow domain it makes claims about. The machine may produce more precise targeting. Fewer unintended deaths. Reduced collateral damage in the specific, measurable sense.

But precision is not the only value at stake in the decision to use lethal force. There is also the question of who bears the weight of the decision. And a system optimized for precision is not optimized for weight-bearing. It does not carry anything. It has no face to confront in the mirror. It does not dream about the target it recommended. It does not wonder, in the years after, whether it was right.

The human who approved the recommendation does carry something. But less than they would have carried if they had generated the judgment themselves. The weight has been partially transferred to

the machine, and machines do not feel weight.

What this produces, across many decisions and many operators and many conflicts, is a systemic reduction in the psychological cost of lethal force at the institutional level. Individual operators may still feel profound weight. But the structural pressure — the recommendation, the confidence score, the time constraint, the institutional track record — is designed to reduce hesitation. And hesitation, in targeting decisions as in most human decisions, is where the weight of accountability actually lives.

There is a version of this concern that ends with: we should slow down, regulate, be careful, ensure meaningful human oversight.

This is the version that appears in policy documents, ethics guidelines, and the statements of AI companies that have signed various responsible use pledges. The concern is genuine. The policy response is insufficient — not because the people making it are cynical, but because the structure of the problem resists the solution.

You cannot put meaningful human oversight into a targeting decision that must be made in ninety seconds on the basis of data no human can process at the required speed. You cannot make the human genuinely responsible for a judgment they did not generate and cannot independently verify. You cannot restore the psychological weight of the decision without restoring the conditions that produced the weight — which means removing the machine, which means accepting the precision losses, which the military logic will not accept.

The accountability gap is not a bug to be patched. It is an emergent property of the system. It exists because the system was optimized for everything except the one thing that accountability requires: a single human mind confronting the full weight of an irreversible decision, with nowhere to distribute the responsibility.

Accountability is not only a legal concept. It is a psychological and civilizational one.

Societies that hold themselves accountable for the use of violence — where the people who authorize it can be clearly identified, where the decision to use force carries a weight that shapes future decisions — are different societies from the ones that distribute that accountability into systems complex enough that no single person can be said to have chosen.

The question that matters is not whether to use AI in warfare. That is largely settled, at least for the next decade, by the military logic that demands it. The question is what kind of people this produces. What kind of institutions. What kind of wars.

Maven is not the end of human decision-making in warfare. It is a point on a curve. The next system will be faster, more confident, more deeply embedded in the decision chain. The human in the loop will become less central, more of a formality — slower than the system, less precise, emotionally contaminated in ways the machine is not.

The curve does not end at a human pressing a button. It ends somewhere no one designed and everyone built together, one optimization at a time.

Who chose?

Not the machine. The machine has no stakes, no conscience, no future in which to live with what it recommended. The recommendation landed in a room and a human ratified it and a missile followed and people died and the machine moved on to the next target in the queue.

Not entirely the human either. The human approved a judgment they did not produce, under conditions designed to accelerate approval, inside a system whose track record argued against hesitation.

This is the specific horror of it — not that it is inhuman, but that it is post-human. It falls in the space between human and machine where our existing moral vocabulary does not reach. We have language for soldiers who chose. We have language for soldiers who obeyed. We do not yet have language for soldiers who approved a recommendation they had every structural reason not to question.

We will need it. The wars that produce the need are already being fought.

We all built this. The question of who chose belongs to all of us.