

Stealing a March

The Marine Corps can't change fast enough
without artificial intelligence

by LtCol Chris "Junior" Cannon, USMCR

In 1997 then-Commandant of the Marine Corps, Gen Charles Krulak, delivered his "Ne Cras" speech to the National Press Club in Washington, DC.¹ Gen Krulak recapped the Battle of the Teutoburg Forest during which the Roman pro counsel, Quintilius Varus, had his force of three legions, 30,000 men, ambushed and destroyed by an alliance of Germanic tribes. Varus, who had put down the same tribes three years prior, was heard muttering "ne cras, ne cras" (Latin for "not like yesterday") as he tried to withdraw his forces. The Germans had negated Roman first century combined arms—heavy cavalry and archery—by luring them into the marshes and forests. Weather also had an effect; Gen Krulak said that the Romans had failed because "their whole view was based on the promise or the premise that nothing would change." In contrast, the 31st Commandant stated, "To win in the 21st century the Corps must steal a march on change."

According to our 37th Commandant, Gen Robert B. Neller, our Corps is in danger of Romanesque complacency. In 2017, Gen Neller told Congress, "The Marine Corps is not organized, trained, equipped, or postured to meet the demands of the rapidly evolving future operating environment."² In a word, this future operating environment is the Pacific. To be more exact, the word is Indo-Pacific, since then-Secretary of Defense James N. Mattis renamed the combatant command in 2018.³ Gen David H. Berger took over as Commandant on 11 July 2019, issuing his *Commandant's Planning Guidance* (CPG) on the same day. Most recently serving as Deputy Commandant, Capabilities Development & Integration (DC, CD&I), Gen Berger oversaw force development

>LtCol Cannon is a Reservist with the MAGTF Staff Training Program and as a contractor supports AI/Machine Learning (ML) projects sponsored by the Office of Naval Research. The content of this article does not necessarily reflect the position or policy of the U.S. Government and no official endorsement should be inferred.

for the Marine Corps and the refining of expeditionary advanced base operations (EABO) concepts for the past year. EABs, which are mentioned 27 times in the CPG's 23 pages, are how the Marine Corps will fight composite warfare generally, and specifically in the Indo-Pacific. EAB requires forces designed to "stand-in" an advanced adversary's weapon engagement zones. As the Commandant clearly states, "If we expect to operate in an environment in which losing the hider-finder competition will result in attack by mass indirect fires, then we will train that way."

But change can be hard.

Learning how to fight an at risk, survivable force in the missile warfare era is a hard problem. This will get harder when the adversary optimally employs missiles using swarming tactics and machine learning from simulated salvos. This will get even harder when the adversary has hypersonic missiles. This becomes harder still when the hypersonic missiles can learn over their time of flight. The stakes are that either we learn—continuously—how to fight, hide, and win in an environment where stand-in assets are at constant risk, or that we create our own 21st century Teutoburg Forest, populated

with smoking holes of expeditionary bases distributed throughout the Western Philippine Sea. The Commandant emphasizes that significant change is required: "Effecting that change will be my top priority as your 38th Commandant."⁴

But change can be hard. To effect change, military Services typically create "institutional memories"⁵ by writing it into doctrine. The CPG has kick-started this process, stating,

We will conduct a comprehensive review of all doctrinal, reference, and warfighting publications to ensure that our doctrine, concepts, tactics, and procedures nest within and support composite warfare

and, "Unless specified within this document, all reference documents from previous Commandants are no longer authoritative." Regarding implementation, Gen Berger again references Gen Krulak, noting both that the 31st Commandant established the Marine Corps Warfighting Laboratory (MCWL) and that "the vehicle for change, in terms of wargaming in support of force design, will be the MCWL." The CPG further states, "We must invest robustly in wargaming, experimentation, and modeling & simulation (M&S) if we are to be a successful learning organization." Gen Krulak's main message in his "Ne Cras" speech was that we similarly needed to become a learning organization in order to steal this march.

But the Commandant's guidance on how to adapt to our "rapidly evolving

future operating environment” is not sufficient. On this, the CPG is clear:

Inevitably, EABO will evolve in implementation into a wide array of missions, with an equally wide assortment of force and capability combinations required to support them.

Even as we increase our learning pace, change in the world is accelerating even faster. To adapt to this evolutionary pace, our warfighting and simulation activities need to fully implement modern analytical techniques of artificial intelligence (AI) and machine learning, which are featured prominently in the CPG. In the warfighting section of the CPG, the Commandant states,

We will make strategic investments in data science, machine learning, and artificial intelligence. ... We will leverage the investments other Services have made as a fast follower.

In order to become a fast follower in any field, a “first mover,” or original innovator, is required.⁶ The company with perhaps the longest AI history may be IBM, with Deep Blue and Watson projects. IBM invested \$1 billion to stand up just one Watson group focused on industrial research and development and fuel analytics.⁷ Watson was mainly developed for the marketing opportunity of beating Jeopardy champions in front of a live studio audience. IBM enjoyed similar gains when their Deep Blue AI beat grandmaster Gary Kasparov at chess in 1997.⁸ But leveraging Watson for applied problems has proven difficult. I worked briefly on a Watson instantiation for a government client in 2012. At the time, “training” Watson consisted of feeding it tens of thousands of question and answer pairs. It turned out that Watson actually had no intelligence, it just knew all of Wikipedia in the form of these question and answer pairs. The true technical challenges of Watson development consisted of harnessing a roomful of 2011 computing power to search all of Wikipedia in the time it took host Alex Trebek to read the Jeopardy answer.

Becoming a fast follower requires even more money or more expertise, to pick up on something that the first



We need to develop a culture of innovation, learning, and change in order to address the challenge of AI. (Photo by Matt Lyman.)

mover missed.⁹ Google caught up to AI innovation in 2014 with the affordable \$500 million acquisition of DeepMind, a London based lab.¹⁰ These types of acquisitions are sometimes described as acqui-hiring, or purchasing the company solely to acquire the talented employees working there in order to move them to similar projects. But even Google has been completely hands-off with DeepMind. DeepMind’s CEO, Dr. Demis Hassabis, has remained with the company since he founded it in 2010.

DeepMind’s story is instructive in what it takes to develop a warfighting concept that continually learns how to win. Throughout 2018, DeepMind developed an AI model called AlphaStar to play against professional Starcraft II players. Starcraft II is a realtime strategy war game first released in 2010. It is one of the longest running, most played video games in the world; Starcraft I debuted in 1998 when Gen Krulak was the Commandant. Here is how DeepMind describes how AlphaStar learns to play Starcraft II in industry terms:

More specifically, the neural network architecture applies a transformer torso to the units (similar to relational deep reinforcement learning), combined with a deep LSTM core, an auto-regressive policy head with a pointer network, and a centralised value baseline.¹¹

In the initial set of matchups, AlphaStar went 10-1 against the professionals, 5-0 against a lower rated pro, and 5-1 against a top player.¹² In July 2019, AlphaStar began playing anonymously on the competitive “ladder” in Europe, so a broader set of the AI’s effectiveness will be available soon.¹³

But AlphaStar is no overnight success. That buzzwordy soup of models was years in research, testing, and development. Even after this year’s long effort, the models supporting AlphaStar were not yet ready to fight. In the Marine Corps, we say that we must train how we fight. AI systems take this quite literally. After initial development, it took a week to “train” an AlphaStar agent on games involving human players. A week means that each agent runs 24/7 on 16 of Google’s proprietary Tensor Processing Units (TPUs). Next is another week or two of subsequent training, this time playing the best performing AlphaStar agents against themselves. Military training is expensive, but AI is in another class by itself. To make the pro beating version of AlphaStar, DeepMind trained 600 AI agents for 14 days on the TPUs. These are not cheap; the estimated training bill is \$28 million.¹⁴ To make an analogy, AlphaStar had a full preseason and regular season before taking on the pros. With our current wargaming and AI capability, employing EABO would be like taking

one spring practice and then advancing directly to the Super Bowl.

Our Marine Corps is not going to acqui-hire a world class AI capability. We would not foot the bill just for training. The Marine Corps is not manned, trained, or equipped to be a fast follower in technological pursuits, particularly AI and data science. It is not hard to see why, as the CPG states:

We do not currently collect the data we need systematically, we lack the processes and technology to make sense of the data we do collect, and we do not leverage the data we have to identify the decision space in manning, training, and equipping the force.

But it is not too late to be a slow follower. StarCraft II is a challenging, complex decision game that has emerged by consensus as a “grand challenge” for AI research. But it is not nearly as complex as actual warfare. The secret sauce to AlphaStar and most production AI systems is what is known as supervised learning: drawing knowledge from observing what humans actually do. In this regard, the Marine Corps has a deep reservoir of knowledge which can benefit AI. The Commandant has stated as much. Immediately after the well-intentioned “fast follower” sentence referring AI and machine learning, the CPG states

These tools will empower our existing analytical community to leverage the advanced education investments the Marine Corps is making in the 88XX community.” Further solutions the CPG prescribes include: deliberate [S]ervice-level O-6 [colonel] and O-5 [lieutenant colonel] talent management, permanent manning adjustments, fiscal programming, and the temporary allocation of highly-qualified manpower from the MCU student population, are all elements of a likely solution for proper resourcing of this critical effort.

To implement AI for concept development and operational use, the Marine Corps needs another shift in how we invest in that technical talent. In 2011, I participated with the planning teams in Quantico that developed the

Commandant’s Career-Level Education Board and the Commandant’s Professional Intermediate level Education Board.¹⁵ Although most of the selectees were destined for resident PME, these boards were designed to send “the best and fully qualified officers” for graduate-level education, even if they did not want it. When other Services send their members for 88XX-level graduate education, they gain a new *primary* MOS, not an additional one. For the critical technical expertise needed to establish a continuing development of data science, AI, and wargaming for operational problems, the Marine Corps needs to follow suit.

Gen Krulak recognized this fact in his CPG, also issued the first day that he was Commandant, stating,

We have limited resources and, therefore, we must get the maximum return on every dollar spent, every Marine assigned to instructional duties, and every minute of instruction. This will require constancy of purpose, stability among faculty and trainers, and an understanding of education and learning theory.

The graduate education “problem” has been studied repeatedly, with different points of emphasis and with different recommendations. However, one thing holds constant: the only way to achieve “maximum return on every dollar spent” is to assign technical Special Education Program graduates to their new 88XX

The CPG states, “All of our investments in data science, machine learning, and artificial intelligence are designed to unleash the incredible talent of the individual Marine.” Gen Krulak’s CPG similarly stated,

It is my intent that we reach the stage where Marines come to work and spend part of each day talking about warfighting: learning to think, making decisions, and being exposed to tactical and operational issues.

These two statements, made 24 years apart, are inextricably linked. The latter exposes how Gen Krulak envisioned our Corps would “steal a march on global change.” Gen Krulak sought to institutionalize innovation by seeking it from “all of our Marines and civilians from private to general and GS-1 to SES [senior executive service],” an act that would continue a culture of learning, a culture of innovation, a culture of change. But Gen Krulak was ahead of his time. With current technology, we can make tactical decision games, models and simulations, tools for MOS-specific problems, and even virtual reality in a collaborative Marine environment that collects the data for training AI models and other data science purposes. But this cannot happen unless we make structural changes for this system to be built to modern standards and by Marines, for Marines.

Stealing a march on change will take more than just organizational stabil-

To implement AI for concept development and operational use, the Marine Corps needs another shift in how we invest in that technical talent.

MOS permanently. This is an acqui-hire within our own ranks. This means telling some supremely talented mid-career Marines that they will be doing AI research under stop-move orders (if not stop-loss). Needs of the Marine Corps. This would be true of each billet that the Commandant suggests altering in order to enable his top priority of changing Marine Corps force design.

ity. It will also take directive action between activities that goes further than what the CPG suggests. The document makes references to both AI/ML and wargaming but does not directly link the two. This linkage is necessary. It is interesting that CMC states (emphasis in the original) that “we will build a Wargaming Center on the Marine Corps University (MCU) campus”



We need to look at what the other Services are doing in AI and data science technology. (Photo by PFC Samuel Ellis.)

at nearly the same instant that DOD is about to award a department wide cloud computing contract.¹⁶ Historical wargaming efforts are not designed to operate on this scale, but AI/ML is. Wargames focused on “talent management, predictive maintenance, logistics, intelligence, and training,” challenges referenced in the CPG, will not yield intelligent answers with tens or hundreds or thousands of simulation runs. Each AlphaStar agent learned to play StarCraft by experiencing up to 200 years of realtime play. That means millions of games played per agent. B.H. Liddell Hart once wrote, “There is no excuse for anyone who is not illiterate if he is less than three thousand years old in mind.”¹⁷ If we do not follow the actual industry standards of fast followers, then our AI models will be functionally illiterate.

Gen Krulak ended his “Ne Cras” speech in 1997 by recalling that the Latin saying had a duel meaning. Five years after Varus’s defeat, Caesar Augustus led an expedition back to Germany to recover the three lost eagles. After finding the German tribes united, Augustus too said, “Ne cras,” which also means “a warning for tomorrow.” The Roman Empire did not fall because of Varus’s intransigence, the eagles were subsequently recovered on three sepa-

rate expeditions, and the Empire continued to thrive for another 400 years. Our current Commandant is clearly focused on the future, with a commitment to a warfighting concept that has a chance to succeed in the Western Philippine Sea and anywhere else in the world. We have been warned, now it is time to march.

Notes

1. National Press Club, “Marine Corps Future,” C-SPAN, (October 1997), available at <https://www.c-span.org/>.
2. U.S. Senate, Statement by Gen Robert B. Neller, Commandant United [sic] States Marine Corps Before the Senate Committee on Appropriations Subcommittee on Defense, (Washington, DC: May 2017).
3. Ryan Browne, “U.S. Rebrands Pacific Command amid Tensions with China,” CNN, (May 2018), available at <https://www.cnn.com/>.
4. Gen David H. Berger, *38th Commandant’s Planning Guidance*, (Washington, DC: July 2019).
5. John A. Nagl, *Learning to Eat Soup with a Knife: British and American Army Counterinsurgency Learning during the Malayan Emergency and the Vietnam War*, (Santa Barbara, CA: Praeger, 2002).

6. Scott D. Anthony, “First Mover or Fast Follower?,” *Harvard Business Review*, (June 2012), available at <https://hbr.org/>.

7. Staff, “IBM Watson Group Unveils Cloud-Delivered Watson Service to Transform Industrial R&D, Visualize Big Data Insights and Fuel Analytics Exploration,” IBM, (January 2014), available at <https://www-03.ibm.com/>.

8. William Saletan, “Chess Bump, The Triumphant Teamwork of Humans and Computers,” *Slate*, (May 2007), available at <https://slate.com/>.

9. Ava Seave, “Fast Followers Not First Movers Are the Real Winners,” *Forbes*, (October 2014), <https://www.forbes.com/>.

10. Catherine Shu, “Google Acquire Artificial Intelligence Startup DeepMind For More Than \$500M,” *TechCrunch*, (January 2014), available at <https://techcrunch.com/>.

11. The Alpha Star Team, “AlphaStar: Mastering the Real-Time Strategy Game StarCraft II,” *DeepMind*, (January 2019), available at <https://deepmind.com/>.

12. Tom Simonite, “DeepMind Beats Pros at Starcraft in Another Triumph for Bots,” *Wired*, (January 2019), available at <https://www.wired.com/>.

13. Staff, “DeepMind Research on Ladder,” *Blizzard Entertainment*, (July 2019), available at <https://news.blizzard.com/>.

14. Information for “Cloud TPU Pricing” is available at <https://cloud.google.com>.

15. Headquarters Marine Corps, *MARADMIN 488/11, FY12 Commandant’s Career-Level Education Board*, (Washington, DC: August 2011).

16. Ron Miller, “President throws latest wrench in \$10B JEDI cloud contract selection process,” *TechCrunch*, (August 2019), available at <https://techcrunch.com/>.

17. B.H. Liddell Hart, “Why Don’t We Learn from History?” (Crows Nest, AUS: George Allen & Unwin, 1944).

