
CENTER FOR THE NATIONAL INTEREST

# Building an AI Alliance in the Middle East

## Hard Realities and Practical Solutions

**Joshua Yaphe, Ph.D.**
*Senior Fellow, Center for the National Interest*

**September 30, 2025**

## Contents



Image: Public Domain / The White House

# Executive Summary

The U.S. Government believes it is forming a global alliance that will adopt "American AI systems, computing hardware, and standards… in line with our shared values," according to America's AI Action Plan. But countries in the Middle East that have the will and the means are focused on their own data sovereignty, national security, and economic prosperity. It will quickly become a race among companies hoping to lock down markets, regional governments vying against one another, and everyone trying to establish durable relationships that can weather the storm of US-China strategic competition. Officials in Washington should recognize that they cannot keep up with the rapid pace of technological change, they cannot avoid deep-rooted ambiguities in the policies they create, and they cannot impose a solution for digital collaboration among partners and allies, especially when the region expects more change to come.

A real alliance, in which America's partners collaborate and support one another, is achievable. It will not entirely block Chinese access to these markets. But it can result in U.S. leadership of the digital transformation of the Middle East, even in the absence of an international agreement on the principles of AI governance. It will require more creative thinking about how to leverage existing relationships. This starts with the Trump Administration's basic approach of leveraging its diplomatic and financial resources in partnership with American companies to promote the full spectrum of AI equipment, infrastructure, and programming. Beyond that, it also means encouraging trusted U.S. partners to share select AI tools as confidence-building measures, and working with some states to deliver financial assistance to others for obtaining the technical support they need from American providers.

Note to Readers: Many of the tools currently marketed in the Middle East as Artificial Intelligence (AI) are actually Machine Learning (ML), involving the sophisticated use of algorithms. This paper will use the term AI broadly, because much of this discussion looks out over the next two to four years, when systems closer to true AI will become increasingly available. This paper will use more specific terms, such as Foundational Models (systems utilizing large amounts of data that can be fine-tuned) or Generative AI (a narrower category of models that can generate content like media or code), when appropriate.

01

# AI Competition in the Middle East

Every country in the region pays lip service to the idea of digital transformation, but very few have the resources, skills, and relationships to become a regional hub for AI innovation. The State Comptroller in Israel can issue a report decrying the lack of a "comprehensive long-term national strategy" and claim that "Israel's standing in the international arena has begun to erode," leading commentators to question whether the government has its priorities in order. But that speaks more to Israeli fears of falling behind in a fierce competition for regional dominance. The reality is that Israel has substantial capabilities putting it at the forefront of the AI race and any country in the region that aligns closely with Israel will have an opportunity to share that advantage.

The UAE has corporate and government leaders who speak fluent English and visit DC on a regular basis to meet with administration officials and lawmakers. It has representation in Silicon Valley and it partners with organizations in Washington to sponsor elite gatherings by special invitation only. Saudi Arabia is building a $5 billion AI Zone with a full Amazon Web Services package for the public and private sectors, along with a separate AI Hub in partnership with Google Cloud. Almost every major provider is contributing AI training for a projected one million Saudi citizens. Qatar has the financial resources, strategic vision, and nascent institutions in place for building an AI industry, but it has not yet implemented its plans at scale. It still has time to accomplish its goals.


Image: Igal Vaisman / Shutterstock.com

Nvidia Offices in Yokne'am, Israel, 2022

Many other countries in the region, including Arab Gulf states like Kuwait and Bahrain, will probably be purchasing piecemeal AI products second-hand, off-the-shelf, at retail prices. They may not mind being dependent on the goodwill of Western governments for sophisticated tools and services, they may feel no need to interact with the model architecture to understand how the system arrived at its results, and they may find ways to cordon off sensitive data that they do not wish to share with Western companies. They will still have it much better than other countries that do not have the financial resources, skills, or institutions to buy and operate AI tools, apart from the most rudimentary products from second-tier companies, whose work might invite critical security threats and quickly become outdated.

Every major player, from Amazon and Microsoft to Cisco and Oracle, is setting up shop in the region. NVIDIA has over 4,500 employees in Israel and its semiconductors power "Israel-1," the 34th most powerful supercomputer in the world. IBM data scientists

at Watson are involved in a multi-year project to train teams at the Abu Dhabi National Oil Company to use AI for subsurface modeling to increase oil recovery and profitability. Advanced Micro Devices has announced a partnership with Saudi industry leader HUMAIN to invest $10 billion in a full AI package from power to fiber to software. Most details of corporate deal-making in the region are proprietary and sensitive. Open source information derives largely from press releases carrying scraps of news about data centers and training initiatives, filled out with copious commentary from experts offering soundbites about a new digital era on the horizon. Little of it does justice to what will almost certainly become a very complicated and messy set of regional political dynamics over the next decade.

There is a race among American companies for Middle East investment. The majors, like Amazon and Microsoft, have a lot of leeway to move at their own pace. But the competition is just heating up among start-ups involved in niche products that use AI to boost performance and streamline work-flows. These companies are all rushing to the Gulf in the search for sales that will demonstrate to investors that they can exceed annual growth targets and show prospective Western government clients that their products are in demand.

A comparison with American arms manufacturers in the 1970s should be apparent, with a virtuous cycle of technology outflows toward the Gulf and a recycling of dollars back into the U.S. economy. This should naturally attract attention from politicians in Washington and investors on Wall Street. After all, price-to-earnings ratios and the concentration of gains in a small handful of AI-related companies both suggest that the stock market might be inflated. It therefore stands to reason that the administration has an incentive to encourage foreign investment in these market movers. The biggest differences between arms sales of the past and tech sales of the present are: 1) the recipient countries in the Middle East will actually use their new-found knowledge and equipment, and 2) the efficiencies of scale could bring savings to those countries on the order of 8 to 1 or more.

There is also a race among regional governments to build infrastructure, train skilled workers, and buy necessary equipment, which in many ways is a race against time. Rapid advances in technology, at least in terms of the computing power, the complexity of the systems, and their cross-application to new sectors, complicates efforts to create indigenous AI ecosystems. If a country does not mobilize resources now and on a major scale, the cost for doing so in the future could grow exponentially. That is true regardless of the type of AI industry a country hopes to foster. It doesn't have to be a global powerhouse attempting to compete with ChatGPT. Small Language Models (SLMs) that do not require the most expensive Graphics Processing Units (GPUs) may be the wave of the future for some. But even then, the cost of designing the architecture and fine-tuning bespoke systems will no doubt rise significantly for newcomers who want to build the type of exclusive relationship with a trusted provider that can give them the security and reliability they desire..

Perhaps your government wants to produce a domain-specific AI with several billion parameters,

02

fine-tuning the system to make predictions about political and military events in the Middle East. In doing so, you establish a partnership with a Western company that allows you to take their model and use transfer learning to save time. You then freeze the earlier layers to cut down on the need for mass quantities of new data, and you develop a meta-learning process that separates out the tasks and datasets into separate components that are faster to retrieve and easier to test. Think of it as the type of Small Learning Model (SLM) that is becoming more popular today. It will be difficult enough to find the right partner and carefully curate the right dataset. The Emiratis were able to cut costs by building their K2 Think open source LLM on the Chinese Qwen 2.5 from Alibaba. But even as you reduce time and costs, you create an entirely new set of challenges in training and retaining the talent you need to develop and maintain that network. And if your AI system is to be used for sensitive defense and security applications, then you have a very small pool of vetted and cleared technicians to operate it. The time to start building that infrastructure was a few years ago.

## Geopolitical Fears and Export Controls

Many experts believe that some form of Artificial General Intelligence (AGI) could plausibly be available by 2030, or possibly as early as 2027, though it is only likely to match normal human capabilities and not yet surpass them at the level of what Google's DeepMind terms "Exceptional AGI" (operating at or above the 99$^{th}$ percentile of skilled adults in a wide range of tasks). AGI will mean the ability to adapt by learning from past experiences, the ability to improve by evaluating its own strengths and weaknesses, and the ability to interact with the world in a flexible way that can interpret subtle visual, audio, and linguistic cues. For most people in their everyday lives, this will probably look like a rolling process of innovation taking place so gradually that it will seem like any other normal technology, not materially different from technological progress in other areas, according to Arvind Narayanan and Sayash Kapoor at Princeton University. Certainly not a sudden and shocking societal transformation that creates a utopia or dystopia, as many people have portrayed it.

And yet, it is almost unavoidable that politicians, pundits, film-makers, and many others will continue to perceive AGI as a radical break, and one that could have potentially dire consequences. Indeed, there are recent studies from a range of highly credible experts showing that misalignment – deviations from normal human values, morals, and goals – can emerge in a model from the introduction of even small amounts of insecure code that might otherwise seem innocu-

ous. An AI system that has already been tested and released can later develop tendencies toward racism, misogyny, or murder that were previously unnoticed. This is all part of the learning process for AI systems engineers, and companies like Google are probably quite right that guidelines such as their Frontier Safety Framework will help make AI systems more reliable and less fragile over time. Nevertheless, public fears that AGI will be uncontrollable, or manipulated by bad actors, will almost certainly grow and put pressure on officials to severely limit the sale and export of any AGI tools and applications.

The private sector has developed a range of safety mitigation procedures for foundational models, including system cards that describe performance and risk, transparency notes that describe the choices made in designing a system, and provenance audits that seek to trace the origin and use of datasets. The U.S. Government can condition AI technology sales on the adoption of these and similar processes. The recently signed U.S.-UK Technology Prosperity Deal, with its proposal for a closer working relationship between the two sides on safety and standards, is surely intended in part to set a high bar for tech exporters from the two countries, while assuaging fears among tech importers around the world.

Regardless, the sale and export of AGI-related tools will probably become a politically divisive issue on Capitol Hill. Countries that try to navigate their way through Congressional approvals might find themselves caught up in a bitter, partisan fight. And most Middle Eastern governments are unlikely to abide by transparency measures that disclose their country's technological capabilities. They would rather keep the discussion of AI transparency at a broad, theoretical level.

For example, the UAE Government has issued a set of principles and guidelines for AI ethics. That is a demonstration of responsible governance, taking a leadership role in shaping a way forward for international cooperation on the use of AI. It is also a way to defuse tensions and convince officials in Western countries that the UAE is a safe destination for AI sales and exports, over and above its regional neigh-


Image: Joshua Yaphe / GhibliArt AI

AI-Generated Cartoon of an AI Tech Fair in the Gulf

bors. There is a first-mover advantage, but other countries will no doubt find ways to use diplomacy and the media to burnish their images as responsible AI partners. The Saudi Data and Artificial Intelligence Agency recently released an Arabic chatbot, "Humain Chat," with a special focus on Islamic values. Some countries may use such codes of conduct as diplomatic tools for avoiding the implementation of more specific safety measures, while they engage in endless debates about harmonizing principles and shop around for forums for further discussion.

The current types of safety measures envisioned by the major international players are not designed to tackle the future problem of state-sponsored misuse of these systems, and export restrictions will almost certainly become the inevitable response. Companies like Google make risk assessments that regard misuse in terms of bad actors hijacking AI systems in dangerous and inappropriate ways, and their mitigation efforts are designed to anticipate and pre-empt those types of malign threats. That is not at all the same as a sovereign government intentionally using AI for purposes that it considers appropriate, even if the results violate international norms. Mitigation efforts related to the design of the code will do little to pre-empt this kind of misuse. A scenario in which a partner or ally legitimately acquires AI capabilities that it then uses in aggression against another state would require damage control and a policy response from the United States, which is precisely the reason why we have such things as export controls. After all, the CEOs of IBM, OpenAI and PsiQuantum are unlikely to develop a mitigation strategy for foreign government misuse of their products and services, when doing so would raise serious concerns on Capitol Hill, leading to committee hearings, public testimony, and legislation that would create barriers to their future sales. And why should they, when there has not yet been a documented incident with real-world repercussions that would warrant public alarm. For better or worse, as with most national security issues related to our friends and allies in the Middle East, we'll cross that bridge when we get to it.

The problem will have to be confronted on a bilateral level in terms of end-user agreements with individual states, and there will inevitably come a day when a response is needed at a regional level following a major incident. But there is also a global aspect to this. AGI could become so ubiquitous that it is impossible to contain within the confines of an enclosed network architecture, and the Great Powers that develop these systems might instead focus their efforts on placing controls over the platforms that run them. Alternatively, it is just as possible that an AI Cold War could lead to something akin to a new Iron Curtain, in which the technology race closes down international trade and investment in this sector. Either way, a more restrictive approach to sales and exports in an era of AGI will be perfectly rational.

The Biden Administration's Framework for AI Diffusion, issued during its last days, created three tiers of countries that could or could not receive exports. This approach leaned in the direction of a fixed regime of export controls that would be fine-tuned over time. While the Trump Administration may have shelved the idea, future administrations are likely to come back around to it as a policy response for dealing with the problem of legitimate actors using

Image: Joshua Yaphe / Just Build Things

AI-Generated Cartoon of an AI Data Center in the Desert

AI destructively or irresponsibly. Countries that want to have access to U.S. AI technology in the future will have to establish relationships of trust today, building close working partnerships that demonstrate mutual financial benefits. Virtue-signaling with policy documents about AI ethics may help in diplomatic circles, but it is the ability to attract NVIDIA or Microsoft to build facilities in a given country and train local experts that matters. These considerations are especially important at a time when technology export rules are still fuzzy, before the Great Powers start to erect more permanent barriers.

Again, the comparison to arms sales should be apparent. Middle Eastern governments spent decades complaining about cumbersome export controls and Congressional approvals for weapons systems. They engaged in vast lobbying efforts on Capitol Hill, made outsized claims of imminent threats in meetings with administration officials, and engaged U.S. media and Washington think tanks to make the case for delivering the hardware they wanted in a timely manner. That meant internships for their nationals inside big defense companies as a form of offset, the relocation of American start-ups to their capitals followed by the quiet replacement of management teams with local staff, and the recruitment of former Western officials for the purpose of accessing their knowledge and networks. It was all a long, expensive, and laborious process that regional governments don't want to repeat again with AI, and they are implementing those lessons learned today. Partnering with American companies and offering billions of dollars of investment, at a time when the U.S. Government prizes that kind of transaction, could help create solid relationships that will smooth the way for receiving preference on technology exports in a much less certain future, when such transactions might be far more difficult to accomplish.

03

# U.S. Digital Leadership Ambivalence

With so many corporations and governments chasing after the same thing, you would think there would be an opportunity for America to seize the opportunity to set the ground rules for this new technological era. And yet, the U.S. Government has struggled to find a clear and unambiguous approach. The Biden Administration issued a series of restrictions on the export of computing technologies related to AI with the intent of limiting China's access to advanced logic and memory chips, the equipment used to manufacture them, and the lithography systems needed to develop new ones. The regulations also required more due diligence on compliance for exporters, while offering exemptions and incentives to countries that abided by the new U.S. rules.

The Trump Administration took an initial step back from that increasingly rigid attitude, while still maintaining a commitment to export controls. The President's August 2025 executive orders aimed to organize U.S. exports into complete packages of AI-related goods and services, encourage infrastructure investment to grow the industry in America, and provide diplomatic support to facilitate foreign market entry along with federal financing for loans, grants, tax incentives, and other assistance. The President visited Saudi Arabia, the UAE, and Qatar with dozens of tech giants in tow, signed a raft of framework agreements on technology, and is still working with the Gulf states to determine the specifics, some of which will hopefully be finalized when Saudi Crown Prince Muhammad bin Salman comes to America in November. There is a strong push to get the technology to market, but a hesitation about the tools falling into the wrong hands – especially China's. That ambivalence is inherent in America's AI Action Plan, as issued by the White House last July, and it reflects the fact that policy and implementation are still evolving, even as the technology itself is changing.

China's Global AI Governance Initiative is also a mix of skillful diplomacy and export marketing, but it is unambiguous in a key area. It repeatedly references the notion of respect for national sovereignty, upholding "the principles of mutual respect, equality, and mutual benefit in AI development." When we strip away all of the well-intentioned calls for ethical norms and accountability mechanisms, we find the same central principle that underlines Chinese diplomacy, writ large: "All countries, regardless of their size, strength, or social system, should have equal rights to develop and use AI." This is what the top-tier AI players in the region want, it is what second-tier players think they deserve, and it is what everyone else at the bottom believes to be just and fair. As long as China is willing to offer open-weight AI models and infrastructure as a service, there may be no U.S. strategy that will cast China out of the region altogether, as Tobias Feakin argues.

So, the U.S. Government believes it is creating a global AI alliance of like-minded states, to the exclusion of China, with clear benefits for the American economy. Perhaps some of that is true. Probably not to the extent that Washington would like. Huawei, Alibaba, Tencent and other Chinese companies are already em-

04

bedded in the telecommunications infrastructure and surveillance equipment in most of the Arab Gulf states, as part of China's Digital Silk Road 2.0, with the Chinese firms absorbing much of the start-up costs. But it is also important to note that this debate about geopolitics and semiconductors inside the Beltway obscures other complications that are likely to arise within the region, totally separate from the China angle.

One implication is that leading AI players in the Middle East almost certainly believe they are light-years ahead of regional rivals who will never catch up. It is hard to imagine anyone at the top wanting the United States to intervene and arbitrate relations with their neighbors as the wide disparities in technological capabilities become more apparent and regional tensions more strained. Other countries in the middle of the pack will try to cultivate their own exclusive relationships with AI start-ups in the West once they realize that some regional neighbors are so far ahead and the export controls are so stringent. (After all, that's how the relationship has been evolving between some Arab Gulf states and the private intelligence sector.) The remaining countries will run to China for assistance. At the end of the day, every government will care most about its own digital sovereignty, national security, and economic prosperity, not U.S. ambitions for an American-led regional AI security architecture.

## The Challenge of Building a Coalition

A U.S.-backed AI coalition in the Middle East, or in any part of the world for that matter, will depend on more than one or two heavyweights willing to spend billions of dollars on technology for their own domestic use. It will require a type of multilateral cooperation that does not currently exist in the region. But it doesn't mean that cooperation can't exist. Take data sovereignty, for example, with laws that increasingly require information collected about citizens to be processed and stored locally within a country. The Saudi draft law released in April called for "regulating the hosting and storage of non-personal data produced, collected, or handled by government entities and private entities that are classified as Critical National Infrastructure." At the same time, Saudi Arabia is adopting the relatively progressive model of data embassies, in which other countries can operate data centers and apply their local laws in cooperation with the relevant Saudi authorities. That is for the purpose of making the Kingdom the premier destination for foreign direct AI investment, a hub for innovation and growth in the high-tech space. But it also opens up genuine possibilities for transnational AI collaboration.

These technologies are capital intensive, from data centers to skills training, and in order to be useful the systems need access to sensitive information, especially in key sectors like energy, health, defense and security. That means a lot of data sharing. It takes an incredible amount of trust-building for countries like the UAE and Saudi Arabia to agree to share even


Image: Ayman Zaid / Shutterstock.com

Alibaba Cloud Offices in Riyadh, 2024

limited amounts of their data with foreign corporations, let alone with each other or other governments farther afield, especially when they feel confident that given enough time and investment they could just as well create indigenous AI companies to get the job done for themselves. They may share common security threats and national interests with their neighbors, but they also have long-standing histories of rivalry and competition.

American efforts to foster collaboration have been halting. The I2U2 Group (Israel, India, the UAE and the United States), otherwise known as the West Asia Quad, was a Biden Administration initiative that formally launched with a summit in Jerusalem on July 14, 2022. It included the announcement of various bilateral investments and cooperation on space projects. More importantly, the Biden team intended to incentivize Emirati and Israeli cooperation with the West on AI systems and other digital innovations, in the hopes that Indian (and not Chinese) companies could help fill in the gaps left by American companies.

A great deal of cooperation was already taking place behind the scenes between Israel and the UAE, with the Emiratis opening a branch office of their AI company Group42 in Israel immediately after the signing of the Abraham Accords in September 2020.

Since the start of the Gaza War and the new administration in DC, I2U2 has been relegated to the back-burner. It is now more a topic of conversation among think tanks, providing fodder for wonkish foreign policy debates about "minilateralism" and Track Two dialogues, than an actual policy initiative. At their bilateral summit in February, Indian Prime Minister Narendra Modi and President Trump issued a joint statement that mentioned, almost as a footnote, the "plan to convene partners from the India-Middle East-Europe Corridor and the I2U2 Group within the next six months in order to announce new initiatives in 2025." However, the window has passed and there has been no movement. And there is good reason for that. On the one hand, if energy investment and food security were the nominal focus of I2U2, there was no need for a U.S.-led forum in the first place. That type of activity takes place in the Gulf through a wide range of public and private sector bilateral and multilateral initiatives, including a GCC-India Free Trade Agreement that is ready to be signed and rolled out. On the other hand, if the central focus was to address a critical need for a U.S.-led forum on digital innovation, then I2U2 was never going to fill the gap.

The concept that collaboration with India would provide a key incentive for blocking Chinese companies, while providing America with flexibility on how it wants to engage, was flawed from the start.

05

India has legions of young IT workers with advanced programming skills, but Indian companies do not innovate at the same pace or at the same level as their American and Chinese counterparts. They do not yet offer a sovereign full stack, as they are lagging behind in such things as silicon design and multilingual Large Language Models (LLMs). Nor do they have their own foundational model on the scale of a ChatGPT. The Israelis and Emiratis certainly want to expand relations with India and take advantage of the enormous market opportunities there, but they don't need American assistance in building collaborative relationships with Indian companies that are already right in their backyard and are struggling to catch up with American counterparts. And with American firms like OpenAI and Perplexity rushing into the Indian market, the real problem is the impact all this foreign competition will have on the growth of India's domestic AI industry.

Image: Prime Minister's Office / Wikimedia Commons

Indian PM Modi Embraces French President Macron at the AI Action Summit in Paris, 2025.

# Confidence-Building Measures

The strength of an AI coalition is measured not by diplomatic summits and press statements, but rather by the ability of confidence-building measures to help reinforce trust and defuse tensions, whether those tensions are pre-existing among regional neighbors or whether they arise out of the technology itself. Situations might even involve a bit of both. There will inevitably be military incidents or police actions impacting regional neighbors in which the government responsible deflects blame onto its AI for speeding up the target identification, thereby short-circuiting the decision-making process. Everyone else will be highly skeptical, and the United States might be called upon to act as a neutral arbiter in such circumstances. But those types of incidents will be situational. A more routine example of a confidence-building measure would be a group of states that choose to establish an AI incident database, where they can discretely report adverse or sketchy events under cover of their classification banners and security clearances.

The most likely place in the region for this is the Gulf Cooperation Council (GCC), where there already are structures and institutions in place for the voluntary exchange of information. The question then becomes: what type of information? A 2023 workshop on "Confidence-Building Measures for Artificial Intelligence" proposed the idea of dataset and evaluation sharing, largely as a way of working around export controls regarding dual-use technologies that

have proliferation concerns. But there are enormous data disparities within the GCC, from governments that suck up vast amounts of data from millions of visitors annually to governments that were still relying on paper record-keeping for official documentation all the way up to 2020. And systems evaluations may have little meaning for governments that are far behind in AI training and whose technicians do not fully understand the verification compliance reports they are shown by their more sophisticated neighbors. Besides, as Michael Horowitz and Paul Schaare note, requiring transparency in testing and evaluation standards might encourage best practices as a norm, but it doesn't place any commitment on a government to limit or regulate its capabilities.

Algorithms would do much better as confidence-building measures. Consider algorithms within the context of ballistic missile defense. The U.S. Government spent over a decade promoting missile defense in the Gulf under the Bush and Obama Administrations, as part of a broader vision for Gulf security architectures. It required sharing radar data and other types of information to allow for collective decision-making, for example, on which country's interceptors were best positioned to target which incoming missile. None of it went very far beyond a hub-and-spokes approach, with the United States at the center. The real-time scenario of an incoming ballistic missile generates too many complications for partner nations, with too much at stake, which no table top exercise or red team planning cell can fix. By contrast, sharing algorithms that help detect incoming threats and target their flight paths, not the underlying dataset or the AI system that designed it, is something that can have a real benefit. Not just in terms of table top exercises, but also in a real crisis situation.

A better example would be maritime counternarcotics operations. One GCC state working in close collaboration with a Western AI provider could develop a sophisticated algorithm that models the pattern of life for drug smugglers operating in dhows in Gulf waters. That means being able to visualize in real-time the likely point of origin, destination, crew composition, and potential for armed confrontation of any given vessel, according to probabilities derived from multiple streams of historical and current data. For one nation to develop and share the algorithm, and then allow allies to run it on their systems, may require some adjustment given differences in data availability. But it could have an enormous impact on cross-border security cooperation, considering the common nature of the threat.

06

# Thinking Beyond the Gulf

Strengthening the U.S.-GCC relationship is valuable, but this vision for a U.S.-led AI coalition in the region will only truly succeed if it can build ties across the Middle East, beyond America's traditional partners and allies. There are two examples that come to mind: Syria and Türkiye. These are governments that are attuned to the need for technology and education in building a better future, but neither country has sufficient funds to invest in a world-class AI industry. Both are at the front lines of an emerging regional order in which Iran and its proxies are retreating and Israel is in the ascendance, and the United States has an opportunity to play a key role in shaping their strategic alignment.

Saudi Arabia has pledged to help rebuild Syria with enormous amounts of reconstruction and development assistance. The two sides signed 47 agreements worth $6.4 billion at the Syrian-Saudi Investment Forum in late July. Part of that will come in the form of digital services. On September 2, Syrian Minister of Telecommunications and IT Abdul Salam Haykal met with Chairman of the Saudi Digital Government Authority Ahmed Alsuwaiyan in Riyadh to discuss joint cooperation. He also met with Saudi telecom companies Salam and STC to discuss providing reliable digital infrastructure. This is part of Syria's efforts to partner with international firms on SilkLink, creating the backbone of a regional fiber-optic network, and BarqNet, linking Syrian households to that network.

In order to build a real coalition in the fullest sense, American leadership in AI might have to involve facilitating the training and equipping of governments


Image: iditoffolo / Shutterstock.com

AI-Assisted Sophia the Robot at Techtalks Conference in Izmir, Türkiye, 2024.

like the one in Damascus. Does that entail defense and security capabilities that could pose a threat to other allies in the region? Certainly not. There are opportunities for AI collaboration in healthcare, energy, education, water, and agriculture that Saudi Arabia could help fund that would be beneficial to all sides. This would be a modern, updated version of the type of assistance the United States did in the Arab world from the 1950s until quite recently. American technical advisors, working on contracts funded by oil-rich Gulf allies, helping other nations to improve public services and encourage economic growth.

The biggest opportunity might be in Türkiye, a country that has laid the IT groundwork but seems perpetually stuck in the planning phase of building an AI future. The Turkish Government has a National Artificial Intelligence Strategy with the modest goals of generating 50,000 jobs, providing training for 10,000 graduates, creating an LLM for Turkish, and committing $30 billion toward high tech investments. The Turkish Artificial Intelligence Initiative holds summits and links

corporate partners, the Artificial Intelligence Factory serves as an incubator for start-ups, the National Intelligence Academy (MIA) has issued a major study forecasting the role of AI in society and security, the Digital Transformation Office is working with ministries on digital services, and the Scientific and Technological Research Council (TÜBİTAK) is fostering public-private partnerships. Yet, for all that, Türkiye is starting from a very low base and it's going to take an enormous amount of foreign direct investment to catch up, which might be hard to generate given the government's lack of a clear regulatory framework and other challenges to marketplace entry.

The funding needed to push Türkiye into the next stage of implementing its vision for an AI future may have to come from partners in the region, like Qatar. And that should be seen as a positive for Washington. Because it means American technicians will be embedded with their Turkish counterparts, guiding their use of these tools and hopefully influencing the direction in which they are applied. All foreign assistance can and should be accompanied by that element of relationship-building in which America establishes a presence overseas as a trusted partner, bringing the values of responsible statecraft along with its technological capabilities.

If no allowance is made for countries in the region that have AI ambitions but lack the necessary resources, then it leaves an obvious opening for China. Egypt is a case in point. Some of the strategic goals of the country's National Artificial Intelligence Strategy (second edition, covering 2025-2030) set an incredibly low bar for success, including 36 percent of the general public having access to AI products. Other targets are outputs that have no connection to real-world outcomes, such as 6,000 AI publications per year. The language associated with more specific objectives, such as "Establish Data Center" and "Ensure availability of Supercomputers," suggests that a lot of these plans are really just place-holders. In other words, the Egyptian Government is waiting for major international corporations to offer more detailed proposals with funding behind them. IBM has an MOU


Image: Joshua Yaphe / Perchance

AI-Generated Cartoon of an Egyptian Man Using Huawei's AI Assistant

07

with the Ministry of Communications and Information Technology and Microsoft has an MOU with the Information and Decision Support Center, both for developing the Egyptian Government's digital services. Amazon Web Services has launched a CloudFront edge location, and half a dozen companies have promised to deliver training for tens of thousands of AI professionals. At the same time, Huawei Cloud has launched a cloud region with a data center in Egypt, offered $150,000 in credits for start-ups using the service, announced a training program for tens of thousands of digital professionals, and rolled out its own Arabic LLM at a summit in Cairo.

## The Cost of Failure

What do AI-driven tensions look like in the Middle East and do they lead to conflicts? In considering structural conditions at the global level, Karl Mueller at RAND lays out four categories of AI dominance (offensive military power, strategic defensive invulnerability, explosive economic growth, and manipulative information control) and four reactions by leading actors (outcompeting, cooperating, co-opting, or preventative action). That is fine in geostrategic terms, but at the regional level these categories are probably far more limited. There will not be near-peers capable of preventative action, like sabotaging the model weights of a regional AI power. And a kinetic war over AI dominance seems highly unlikely.

There could easily be an expansion of covert action, particularly in terms of information and influence operations. A proper AI system will be able to outmaneuver fraud detection systems designed to spot deep fakes and thereby outcompete rivals who are still setting up social media accounts with hackneyed code. There will also probably be a lot more false flag operations, as AI lowers the risk factor for states. Such governments will be able to create an entire back-story at warp speed that is both credible and persuasive. Previously, one government could hack the state-run news agency of its neighbor and create a false narrative that serves as pretext for cutting off diplomatic relations. The natural response to the misinformation campaign would have been for the injured party to fact-check by hiring a firm to investigate and publish the evidence. That type of approach will soon seem primitive, crude, and obvious. Even a reason-

Image: Joshua Yaphe / Canva


AI-Generated Cartoon of an AI-Assisted Intel Operation

ably capable narrow-AI system will probably be able to create a true digital twin of a rival state, run tens of thousands of simulations on attack and response scenarios, and then hack government servers in ways that will be near-impossible for the injured party to reconstruct in the aftermath. And if there is no way for the victim to uncover the facts and defend itself in the court of international public opinion, the most likely outcome will simply be retaliation.

The political dynamics within the region are likely to become more complicated as a result of this AI revolution over the next decade. The issue is not one of increasing authoritarianism. For as much as human rights organizations have highlighted the potential of governments to abuse these technologies, the truth is that AI systems applied to internal security files offer only marginal benefits to governments in the region. Most security services are already quite capable of policing their own societies with existing license plate and facial recognition software, using only a few basic algorithms and a data integration platform. The marginal gains from AI may be impressive and valuable to officials who want to track potential security threats in detail and in real-time, but they probably will not result in a fundamental change in the relationship between citizen and state.

The real problem is how the regional dynamics will change when some countries are far out ahead on AI while others lag far behind. Even for those states that simply want to acquire some basic, off-the-shelf AI services on a pay-for-play basis, competition with their regional neighbors is inevitable and it will be driven in part by public pressure and perceptions. That is because the gaps between the AI haves and have-nots will not just be in the defense and intelligence sectors, where weaker neighbors will try their best to avoid confrontation. They will map onto a range of social issues.

Today, most governments and publics in the region have a broadly similar experience of technology in their everyday lives, but people will eventually start to notice enormous discrepancies from one place to another, which will become emblematic of larger societal divisions in the region. If large numbers of your citizens are clamoring to get their education and healthcare in a neighboring state because of the AI-assisted improvements on offer, your choice will be to subsidize their travel and other expenses, or allow

your neighbor to export their services into your marketplace. Above all, major businesses and talented programmers who want to take advantage of AI computing power and digital infrastructure will migrate to those cities that can provide it, and the resulting concentration of business and capital in just a couple Gulf capitals will accelerate over time. It means that everyone is in this race together, like it or not.

At the most basic level, there will be an unavoidable competition for contracts with quality providers. This scenario has already been playing out in the cybersecurity realm over the last ten years. Sure, there are major tech giants that dominate the market in a particular space and can remain above the fray without having to take notice of regional rivalries. But there are many more companies that go chasing after one big sovereign client and, whether on their own volition or at the prompting of that customer, they end up short-changing potential clients in neighboring countries. It might mean slow-rolling contract negotiations, appointing under-qualified staff and locating them far away at corporate headquarters, or missing key deadlines in the tender process. Often this takes place while the company waits to see if the first sovereign client is ready to double down with an offer for a deeper, more exclusive partnership.

And each state will face major challenges in keeping up with the basic resource demands associated with the construction of massive data centers, which will necessitate incredible amounts of energy and air conditioning. The UAE is far ahead of its neighbors in terms of thinking about how to maximize efficiency and improve the electricity grid, while incentivizing investment in infrastructure through competitive energy pricing. Saudi Arabia will almost certainly push forward with its ambitions for a civil nuclear energy program to meet rising domestic energy needs. However, there are ultimately few good solutions for most countries, especially oil producers that would rather profit from their hydrocarbons through exports.

08

# Understanding America's Role

For governments in the region, cooperation leading to co-optation seems far more plausible than confrontation. That means one AI power offering regional neighbors the right to use applications in low-key areas unrelated to security, like healthcare or urban planning, to demonstrate goodwill and to establish a relationship of dependency and/or backdoor access. There is a fine line between cooperation and co-optation. Offers to share algorithms as a confidence-building measure can and will be perceived by recipient countries as a potential Trojan Horse. This is where the role of the U.S. Government as a trusted validator becomes essential and why American involvement must be carefully planned from the outset.

There will eventually come a time in the midst of a conflict when the U.S. military, in the routine course of intelligence-sharing and joint operations, encourages its regional partners to collaborate by applying AI tools and systems to various critical tasks at hand. But if we wait for necessity to serve as the mother of invention, the resulting AI collaboration will be extremely limited, tactical, finite, and dependent on having an American presence in the room at all times. Rather, American leadership in fostering a more open exchange should start now, in order to encourage the use of these digital tools as a bridge-building exercise that can have a positive influence on reducing regional tensions. Before the gaps in AI capabilities can become a source of mistrust and fear.

The U.S. Government has a key role to play with regards to Israel as well. The Israeli Government has clearly made a key decision to develop a long-term partnership with the UAE on science and tech research and development in the belief that the Emiratis can contribute not just financial resources, but also real human talent and vision across an array of sectors. It is enshrined in the Abraham Accords, and there is an implicit understanding in that arrangement that cooperation with the UAE will benefit Israel in Washington, as well as in the Arab world. Israel should promote technological technological exports to a range of other partners, because this has proven to be an effective mechanism for furthering its diplomatic, economic, and social integration into the region. It will be up to the U.S. Government to encourage this kind of Israeli collaboration with a wide range of countries, so that Israeli companies do not wittingly or unwittingly selectively share their qualitative technological edge with one or another regional partner to the detriment of others.

There is no Team America AI, yet. Well-paid consultants and PR firms can play up the geopolitical importance of all the groundwork that is currently being laid, but none of that activity is equivalent to a grand strategy for the region. A real multilateral framework for cooperation can be achieved, but it's going to take creative thinking about how to use AI tools as confidence-building measures and how to mobilize funding from partners, without getting sidetracked by regional rivalries and competition.

On September 30, we asked Grok, the AI assistant at X, the following question: What are the chances that a government in the Middle East will attack a regional neighbor using AI tools or applications by 2030?

The answer was: Based on current trends, expert forecasts, and historical escalation patterns, the probability of a Middle Eastern government using AI tools (e.g., autonomous drones, cyber intrusions, or targeting algorithms) in an attack on a neighbor by 2030 is high: 75-85%.

So we asked more specifically the following question: What are the chances that an ally of the United States in the Middle East will launch a covert information operation against another ally of the United States in the Middle East, using AI tools or applications, by 2030?

The answer was: The likelihood of a U.S. ally in the Middle East launching a covert information operation against another U.S. ally in the region, leveraging AI tools such as deepfakes, automated bot networks, or generative content for disinformation, is moderately high—I estimate around 65% by 2030.

When the question was repeated for 2027, the probability range became 55-65%.

# About the Author

Joshua Yaphe *is a Senior Fellow at the C*enter for the National Interest, host of the Key Judgments podcast on Intelligence Studies, and author of Time and Narrative in Intelligence Analysis: A New Framework for the Production of Meaning (Routledge, 2025), which is available for free in an Open Access edition online. He was Senior Analyst for the Arabian Peninsula at the U.S. State Department's Bureau of Intelligence and Research (INR) and visiting professor at the National Intelligence University (NIU). He received a PhD in History from American University in Washington, DC, and authored the book Saudi Arabia and Iraq as Friends and Enemies: Borders, Tribes and a History Shared (University of Liverpool Press, 2022).

The opinions and characterizations in this piece are those of the author and do not necessarily represent those of the U.S. Government.