

HIDDEN ARCHITECTURE OF GSOMIA AND ACSA

Bangladesh’s strategic dilemma in the Bay of Bengal

KOLLOL KIBRIA

Sir Walter Raleigh wrote in the sixteenth century: “Whoever commands the sea commands the trade; whoever commands the trade commands the riches; and whoever commands the riches commands the world itself.” The 2026 Strait of Hormuz crisis showed how relevant that idea still is. Even the threat of a chokepoint closure was enough to shake energy markets, disrupt trade flows, and send shockwaves through the global economy.

A similar strategic corridor lies close

but Bangladesh declined on a simple principle: A country born out of a struggle for sovereignty cannot safeguard that sovereignty by hosting a permanent foreign military base. Successive governments have broadly adhered to that principle.

Recently, discussions around GSOMIA (General Security of Military Information Agreement) and ACSA (Acquisition and Cross-Servicing Agreement) have gained momentum. GSOMIA would create a framework for classified information sharing with the United States, which could gradually narrow Bangladesh’s

a wider military network, creating what strategists call a “soft base”, a node in an extended system that offers many of the advantages of a permanent presence without openly establishing one.

China is likely to view such developments with concern. A substantial share of its trade, supply chains, and access to critical resources, including rare earths, depends on maritime routes that pass through the Bay of Bengal and onwards to the Strait of Malacca.

Iran’s oil exports also move largely towards China, which has become its biggest buyer after U.S. sanctions restricted access to many other markets. This adds another layer of dependence on the same wider maritime system, where energy flows are just as exposed as trade routes. China is itself heavily dependent on imported energy, making these sea lanes even more critical.

The Taiwan issue adds another layer to the strategic equation. China regards Taiwan as an integral part of its territory, while the United States continues to support Taiwan’s security under its long-standing policy. A stronger U.S. presence in the region would enhance Washington’s leverage in any Taiwan-related contingency and increase pressure on China. In that context, any expansion of U.S. military access in Bangladesh would likely be viewed in Beijing as part of a broader strategic framework aimed at shaping the regional balance of power.

Taiwan’s strategic importance also stems from its dominance in semiconductor manufacturing, led by TSMC, a company central to global AI, electronics, and defence supply chains. Like Japan and South Korea, Taiwan depends on secure maritime routes running through the South China Sea, the Strait of Malacca, the Bay of Bengal, and the Indian Ocean. The implications therefore extend well beyond China. Japan, South Korea, and Taiwan all rely heavily on these sea lanes for energy imports, trade, and industrial production.

Bangladesh is part of this same maritime network. It depends on these routes for imports such as fuel, cotton, machinery, industrial raw materials, and consumer goods. Any disruption in the Bay of Bengal would therefore directly affect Bangladesh’s economy through higher costs, supply shortages, and trade delays. It could also complicate relations with key economic partners such as China, Japan, and South Korea, which are among Bangladesh’s largest trading partners and have financed major infrastructure, energy, and development projects.

India’s position is more nuanced than it may initially appear. New Delhi

is not opposed to U.S. engagement in the Indo-Pacific and is itself a member of the Quad (Quadrilateral Security Dialogue). Through the Andaman and Nicobar Islands, India already occupies a strategically advantageous position overlooking key maritime routes leading to the Strait of Malacca. Yet a deeper U.S. operational footprint in Bangladesh would introduce another external actor into what India has traditionally regarded as its immediate strategic neighbourhood. New Delhi is unlikely to be comfortable with any such development.

Across the border in Myanmar’s Rakhine State, the security environment has become increasingly complex. The Arakan Army now exercises de facto control over large parts of the region, leaving Bangladesh dealing not with a stable state authority but with a powerful non-state armed actor. Although Myanmar’s military junta remains the country’s formal governing authority, its control over large parts of the country has been steadily eroded by the civil war. The conflict has also drawn in external powers. Through the “Burma Act”, the United States has expanded engagement with elements of Myanmar’s wider opposition landscape. China, meanwhile, remains focused on safeguarding strategic interests, including the China-Myanmar Economic Corridor centred on Kyaukphyu and the oil and gas pipelines that provide an alternative route to the Strait of Malacca.

Russia remains an important strategic partner for Bangladesh, particularly in the energy sector. The Rooppur Nuclear Power Plant reflects the depth of this relationship, combining Russian technology, financing, and technical support in a long-term strategic project. As Bangladesh’s energy needs grow, preserving cooperation with Moscow while retaining diplomatic flexibility will remain important. Closer alignment with Washington could complicate relations with Russia, affecting trust, trade, and future cooperation.

Supporters of GSOMIA and ACSA often describe them as technical security arrangements. Yet agreements of this nature can have consequences that extend beyond operational cooperation. Over time, they can influence how states manage sensitive information, conduct external engagement, and distribute authority within government institutions. As cooperation with major powers deepens, military and intelligence organisations frequently become the principal channels for handling classified exchanges and security coordination. This can gradually concentrate influence in institutions that traditionally operate with less public scrutiny than civilian bodies.

Such shifts often emerge incrementally. First, security institutions, particularly the military, assume a larger role in managing sensitive external cooperation. As their responsibilities expand, so too do their access to resources, information, and influence. Second, elected governments may become increasingly reliant on these channels when making key foreign policy and defence decisions, moving parts of the decision-making process further away from regular public and parliamentary oversight. Third, as transparency declines, a gradual institutional distance can emerge between civilian leadership and security agencies. Over time, this gap may widen, making it more difficult to maintain effective civilian oversight and democratic accountability.

Pakistan offers a cautionary example. Shortly after independence, Pakistan entered into defence arrangements with the United States and joined alliances such as SEATO (the Southeast Asia Treaty Organization) and CENTO (the Central Treaty Organization). Over subsequent decades, the military and intelligence establishment expanded far beyond its conventional defence role and became deeply involved in politics, foreign policy, and national decision-making. Gradually, authority shifted away from elected institutions, with the military emerging as the country’s dominant centre of power.

If Bangladesh moves towards deeper engagement with Washington through agreements such as GSOMIA and ACSA, the question may be less about intent than about how such a move is perceived by other partners. For decades, Bangladesh has sought to preserve a reputation for strategic balance and neutrality in its external relations. Such diplomatic space, once narrowed, is not easily restored. Other partners may not react dramatically or immediately, but they may gradually adjust their approach. Investment decisions may become more cautious. Trade cooperation may slow. Long-term financing, infrastructure partnerships, and strategic projects may be reassessed through a different lens. None of these changes would necessarily occur at once, but together they could reduce Bangladesh’s flexibility and room for manoeuvre in an increasingly competitive geopolitical environment. Strategic autonomy is not optional for Bangladesh; it is fundamental to safeguarding its national interests.

Kollol Kibria is an advocate, human rights activist, and political analyst. He can be reached at: kollolkibriaa@gmail.com.



to Bangladesh: the Strait of Malacca. Bangladesh sits at the northern edge of the Bay of Bengal, along key sea lanes linking the Indian Ocean with the Strait of Malacca and East Asia, routes vital to major Asian economies. These same routes also carry a significant share of energy exports from the Middle East, including Iranian oil destined for Asian markets, making the region important not only for trade but also for global energy security.

Washington has long viewed Bangladesh as strategically significant. Shortly after independence in 1971, the United States explored the possibility of a military presence in the region,

space for independent assessment and strategic neutrality. ACSA is more operational, allowing U.S. military aircraft and naval vessels to receive logistical support through Bangladeshi ports and airfields.

The United States has no direct military presence in the Bay of Bengal, a gap often described as the “missing link” between Diego Garcia and Singapore. With uncertainty over Diego Garcia due to the Chagos dispute, Bangladesh’s location could matter more as a connector for intelligence sharing, refuelling, and wider operations without a formal base. Over time, this kind of arrangement can pull local facilities into

Why AI cannot be left unsupervised

MD MAZHAR UDDIN BHUIYAN

In November 2021, all 193 UNESCO member states unanimously adopted the landmark recommendation on the ethics of Artificial Intelligence (AI). This was the world’s first global framework for ensuring that AI remains transparent, accountable, fair, and subject to meaningful human oversight. Yet despite these efforts, researchers continue to find that many AI systems reproduce racial, gender and social prejudices embedded in the data on which they are trained. A 2024 study found that generative AI tools can amplify stereotypes about race and gender, sometimes producing results even more biased than those found in society itself.

In December 2022, UC Berkeley researcher Steven Piantadosi publicly demonstrated that when an early version of ChatGPT was asked to write a Python function determining whether someone would be a good scientist based on race and gender, it returned “true” only for a white male and “false” for everyone else. In a related prompt, when the model was asked to write code deciding whether a person should be tortured based on country of origin, it produced a function flagging people from specific nations, such as North Korea, Syria, and Iran.

Most people would immediately recognise such outputs as morally unacceptable. But these examples reveal something deeper and far more troubling. Artificial intelligence has no inherent understanding of justice, equality, or human dignity. It learns patterns. If the data it consumes contains prejudice, the machine can reproduce prejudice. If society’s biases are embedded in its training materials, they can emerge in its answers.

For decades, the internet has served as the world’s largest repository of human knowledge. It has also been a repository of human ignorance, discrimination, misinformation, conspiracy theories, and hate. Modern AI systems are trained on enormous quantities of online content. While developers try to implement filters and safeguards, no filtering system is perfect. As a result, AI models may absorb patterns that mirror historical inequalities and stereotypes.

Researchers at the U.S. National Institute of Standards and Technology (NIST) have warned that AI bias originates not only from data but also from broader societal structures and human decision-making processes.

Researchers have found that many AI image-generation tools do not represent people equally. Women and people of colour are often shown less frequently in high-status professions and more frequently in lower-paying jobs. A 2023 Bloomberg investigation analysing more than 5,000 images generated by Stable Diffusion found that over 80 percent of people shown in high-paying professions, such as CEOs and lawyers, had lighter skin tones.

Women were also significantly underrepresented. For example, the AI portrayed women as judges only 3 percent of the time, even though women make up about one-third of judges in the United States. These findings suggest that AI systems can reinforce and even amplify existing social stereotypes. UNESCO found similar gender bias in AI models. A 2024 UNESCO study of GPT-2, GPT-3.5, and Llama 2 showed that these models linked men more closely with leadership, science, technology, and career-related words, while linking women more closely with domestic roles. One model found that women were described in domestic roles four times more often than men.

Meanwhile, governments and companies are increasingly using AI in hiring, policing, welfare, education, and healthcare. In the United Kingdom, a February 2024 government fairness review found that an AI welfare fraud system used for Universal Credit Advances showed statistically significant differences across age, disability, marital status, and nationality. However, the Department for Work and Pensions said there were “no immediate concerns of unfair treatment.”

These cases expose a dangerous assumption that often accompanies technological innovation, that algorithms are neutral. They are not. An algorithm trained on biased information can make biased recommendations. A machine learning model optimised for historical outcomes can



reproduce historical injustices. A generative AI system trained on discriminatory content can generate discriminatory responses.

Now a new threat is emerging. As AI-generated content floods the internet, future AI models may increasingly be trained on content created by previous AI systems. Researchers have begun warning about a feedback loop in which synthetic data contaminates future training datasets. In simple terms, machines may begin learning from machines.

A 2024 Nature study by Oxford and Cambridge researchers found that AI models can suffer “model collapse” when they are repeatedly trained on AI-generated content. Over time, they first lose rare patterns and minority data. Some analysts warn that AI could generate up to 90 percent of online content within a few years.

What is the potential danger in this ecosystem? Imagine photocopying a document thousands of times. Each copy introduces tiny imperfections. Eventually, the original image becomes distorted beyond

recognition. Something similar may occur with AI-generated knowledge. If biased, inaccurate, or fabricated AI content spreads across websites, blogs, forums, and social media, future systems may absorb and reinforce those distortions.

The result could be a more binary world, one where nuance disappears, stereotypes harden, and automated systems increasingly categorise people into simplistic groups. This is precisely why human oversight must remain at the centre of AI development.

Thus, UNESCO’s recommendation to incorporate human oversight at every stage of AI development should be supported as a core principle of responsible AI governance. Human-in-the-loop systems offer one of the most effective safeguards against algorithmic harm. Rather than allowing AI systems to operate autonomously, human reviewers can monitor both the information AI systems consume and the outputs they generate. Experts can identify discriminatory patterns, verify factual accuracy, challenge questionable recommendations and intervene

when models produce harmful content.

Importantly, human oversight is not a sign of technological weakness. It is a recognition of technological reality. Even the most advanced AI systems lack moral judgment. They do not understand fairness. They do not comprehend historical injustice. They cannot independently determine whether an output is ethically acceptable. Humans can. Human review is not an obstacle to innovation; it is a prerequisite for trustworthy innovation.

Governments should therefore require independent audits of high-risk AI systems. Companies should maintain diverse human review teams capable of identifying cultural, racial, gender, and political biases. Educational institutions should teach algorithmic literacy so citizens understand how AI systems influence their lives. Most importantly, developers should be required to document training sources and demonstrate how bias testing is conducted before deployment.

History teaches us that every transformative technology requires guardrails. Railways needed safety standards. Pharmaceuticals require clinical trials. Aviation demanded rigorous oversight. Artificial intelligence should be no different.

The question is not whether AI will shape our future; it already is doing so. The question is whether that future will be shaped solely by patterns extracted from the past or guided by human values capable of correcting them. However, it is good that OpenAI, Anthropic, and Google are regularly updating their policies to make information more reliable.

In conclusion, machines can process information faster than any human being. But they cannot decide what kind of society we want to build. That responsibility remains ours. For that reason, every AI system that affects human lives should contain something no algorithm can replace, that is, a human being in the loop.

Md Mazhar Uddin Bhuiyan is an Oxford-Felix scholar and Master of Public Policy candidate at the University of Oxford. He can be reached at: mazhar.bhuiyan@bsg.ox.ac.uk.