

AMERICAN POLITICAL SCIENCE ASSOCIATION

MENA POLITICS

Newsletter of the Middle East and North Africa Politics Section of APSA

VOL. 9 ISSUE 1, SPRING 2026

ISSN: 2996-4911



Letter From the Editors

Greetings,



Diana B. Greenwald is an Associate Professor of Political Science at City College of New York, City University of New York
dgreenwald@ccny.cuny.edu



Sebnem Gumuscu is an Associate Professor of Political Science at Middlebury College
sgumuscu@middlebury.edu



Samer Shehata is the Colin Mackey and Patricia Molina de Mackey Associate Professor of Middle East Studies at the University of Oklahoma
shehata@ou.edu

We are pleased to present the Spring 2026 issue of *MENA Politics*. This issue includes a symposium about Artificial Intelligence (AI) in Middle East politics, a roundtable on AI's role in political science research, and a featured essay by George Washington University's Nathan Brown reflecting on his approach to the study of politics. It also includes important information about the APSA MENA section and our community from our Section Chair, Marwa Shalaby, in addition to news about past and upcoming conferences, workshops, and other opportunities from the Arab Political Science Network and the APSA MENA Program.

The issue begins with Nathan Brown's featured essay entitled "A Proustian Approach to the Study of Politics," adapted from his remarks upon receiving the 2025 Michael E. Brown Research Prize, the foremost research award at George Washington University's Elliot School of International Affairs. Brown's essay, simultaneously poetic yet decidedly serious, outlines his motivations for and philosophy of research, which he describes as "less to do with discovering the truth and more to do with helping people understand each other when they see the truth so differently." Brown uses Marcel Proust's notoriously intimidating

seven-volume *In Search of Lost Time* as a lens to reflect upon a range of issues, from effective ways of allowing people to see the world from different perspectives, to the subject of his current book manuscript, the political and historical relationship between state, religion, and higher education in the Middle East, Europe, and the United States.

The current issue also marks the first time *MENA Politics* has focused on AI in Middle East politics and political science in a sustained manner. Eric Lob and Selin Bengi Gümrükçü have guest edited a fascinating symposium about AI in the domestic and international politics of the Middle East, focusing primarily on the Gulf, featuring outstanding contributions from Arshin Adib-Moghaddam, Yasir Atalan, Guy Burton, Andrew Leber, Abdullah Omran, and Alexandra Siegel. Together, the essays present a broad sampling of AI's role in a range of political, security, and economic issues in the Gulf and the wider MENA region.

Carolyn Barnett's guest-edited roundtable about AI as a research tool in Middle East political science is the issue's second engagement with this topic. Barnett has produced a lively "conversation"

(CONTINUED ON NEXT PAGE)

Letter from the Editors (continued)

between Muhammad S. Abdo, Tariq Adely, Cinzia Bianco, Ashrakat Elshehawy, Robert Kubinec, and Michael Robbins about AI's potential and varied uses, and its limitations, in the study of a wide variety of subjects in Middle East politics. The roundtable is accompanied by a stand-alone reflection from Richard A. Nielsen on the sea change that AI is poised to cause in our discipline, and how interpretivist modes of inference will become an increasingly important preserve of human-authored political science. He also writes from an ethical foundation, arguing that "[m]ost researchers are rethinking, or will rethink, their workflow from a place of precarity," and those more secure in the profession owe it to those just entering it to continue ensuring there is a place for social science that centers human questions, meaning-making, and interpretation. Taken together, the symposium and roundtable constitute *MENA Politics*' initial engagement with the increasingly important issue of AI in MENA's politics and political science and will likely not be the last time this subject receives attention in the Newsletter.

Although we are delighted to present the current issue of *MENA Politics*, we remain deeply concerned about events in our community and horrified by developments in the region. The assault on academic freedom in the United States, and particularly Middle East studies and Middle East politics, appears to be continuing unabated. This includes several deeply troubling developments that have occurred since the publication of the last Newsletter.

Perhaps most alarming, Shirin Saeidi, a tenured associate professor of political science at the University of Arkansas (UA), was fired

from her position for apparently political reasons in March 2026. Saeidi is an Iranian-American scholar and specialist in Iranian and Middle East politics. She was removed as director of UA's King Fahd Center for Middle East Studies in December 2025 after alleged minor rules violations regarding the use of university stationery, which she denies, but also, centrally, for opinions about Israel and Iran expressed in social media posts, thus likely violating the university's own policy on academic freedom and, as the Middle East Studies Association Committee on Academic Freedom argues, Professor Saeidi's First Amendment protections (Strain 2026, MESA 2025). Shortly after Saeidi's removal as center director, her dean recommended her dismissal from the university. Despite the UA's Faculty Hearing Committee's unanimous recommendation against her dismissal, Saeidi was fired in March 2026 (MESA 2026). Hers is only one recent example, albeit egregious, of the continuing assault on academic freedom and Middle East studies in the United States.

We are also horrified by the devastating wars in the Middle East that continue to produce significant death, destruction, and insecurity in the region. The Israeli-US attacks on Iran that began on February 28 have already killed more than 3,300 Iranians, including 1,700 civilians, at the time of this writing, according to Iranian and US sources (Abdallah and Alaaeldin 2026, Al Jazeera 2026). Tens of thousands more have been injured. Iran's military, civilian, and economic infrastructure has also been severely damaged from over 13,000 US and 10,000 Israeli air strikes, including on oil and gas facilities, hospitals, bridges, universities, and schools (Institute for National Security Studies 2026).

(CONTINUED ON NEXT PAGE)

Letter from the Editors (continued)

Iranian and Hizbollah missiles have killed 23 Israelis and injured many more, according to Israeli sources, while 28 people have been killed in the Arab Gulf states from Iranian missile and drone attacks (Sifton 2026, Al Jazeera 2026). 13 US military personnel have also been killed and 381 have been injured, according to US military sources (Kime 2026). Iranian attacks have also caused extensive damage to US military bases in the region, totaling billions of dollars of damage (Fields 2026). Iranian attacks on the Arab Gulf states have targeted energy and transportation infrastructure, industrial sites, desalination plants, and civilian buildings, including iconic landmarks, such as Dubai's Burj Al Arab Hotel. We write as a formally expired cease-fire between the US, Israel, and Iran remains tenuously in effect, yet we cannot rule out a possible ruinous escalation in response to the conflict over the Strait of Hormuz.

The 2026 "Iran war" has also led to renewed fighting between Hezbollah and Israel, with Lebanese civilians paying the highest price. Hezbollah began firing missiles into northern Israel on March 2, 2026, in retaliation for Israel's assassination of Iran's Supreme Leader, Ali Khamenei, on the first day of the conflict. Israel soon responded with overwhelming force. As of this writing, 2,618 Lebanese have been killed, over 8,094 have been injured, and over 1 million have been displaced because of Israeli military operations (Hatoum and Mroue 2026). At least two Israeli civilians and 16 soldiers have been killed because of Hezbollah attacks, and 4,400 Israelis have been displaced (Fabian 2026). Several UNIFIL peacekeepers have also been killed. Since the renewal of fighting, Israel has occupied more territory in southern Lebanon, beyond the

five military outposts it established after the 2024 ceasefire, caused extensive damage to civilian infrastructure, and declared its intention of establishing a "buffer zone" in southern Lebanon. Harrowing footage of controlled demolitions of entire villages has emerged in recent weeks, after the Israeli Defense Minister explicitly stated that "all houses in villages near the Lebanese border will be destroyed, in accordance with the model used in Rafah and Beit Hanoun in Gaza, in order to permanently remove the threats near the border to northern residents" (Scheer 2026). Israel will argue for military necessity in attempting to defend these policies. This argument is nearly impossible to square with, for example, recent footage showing Israel damaging solar panels used to power the pumping and distribution of drinking water in the village of Debel (Amnesty International 2026) or the apparent wholesale destruction of the entire town of Bint Jbeil and many other villages (Granados et al 2026).

These events will undoubtedly shape domestic and regional politics in the Middle East for many years to come, likely resulting in dozens of scholarly publications by political scientists, including, possibly, in this Newsletter. Yet it is essential not to lose sight of the devastating effects these events are having on millions of people experiencing war and displacement in the region today.

If you have suggestions, feedback, or ideas for future issues, including on the current conflicts taking place in the region, please contact us. Section members impacted by these events are particularly encouraged to reach

(CONTINUED ON NEXT PAGE)

Letter from the Editors (continued)

out with suggestions for future issues. Your engagement is vital to the continued relevance and vibrancy of MENA Politics.

-Diana B. Greenwald, Sebnem Gumuscu, and
Samer Shehata

References

Abdallah, Nayera and Menna Alaaeldin. 2026. "How many people have been killed in the Iran war?" *Reuters*, April 10.

AJLabs. 2026. "US-Israeli attacks on Iran: Death toll and injuries live tracker." *Al Jazeera*, March 1, updated April 27.

Amnesty International. 2026. "Israel continues to destroy civilian infrastructure as shown in this video, which captures the Israeli army destroying solar panels in southern Lebanon." Instagram. April 28. <https://www.instagram.com/reel/DXrmojUE3NJ/>.

Committee on Academic Freedom, Middle East Studies Association, 2025. "Letter to the University of Arkansas regarding the removal of Professor Shirin Saeidi as director of the King Fahd Center." December 16.

- 2026. "Letter to the University of Arkansas about the dismissal of Associate Professor Shirin Saeidi," April 17.

Fabian, Emanuel. 2026. "IDF soldier severely wounded as Hezbollah drones target troops in southern Lebanon." *The Times of Israel*, April 28.

Granados, Samuel, Abdi Latif Dahir, and Sanjana Varghese. 2026. "Israel Said It's Applying the Gaza Model in Lebanon. This Is What the Devastation Looks Like." *New York Times*, May 3. <https://www.nytimes.com/interactive/2026/05/03/world/middleeast/israel-lebanon-hezbollah-gaza-destruction.html>

Hatoum, Bassam and Bassem Mroue, 2026. "Israeli strikes in southern Lebanon kill 10 people as Hezbollah drone wounds 2 Israeli soldiers," *AP*, May 1.

Institute for National Security Studies, 2026. "Dashboard: The Campaign in Iran and Lebanon," update April 28. <https://www.inss.org.il/publication/links-roar-data/>

Kime, Patricia. 2026. "13 US troops killed, more than 380 wounded in Operation Epic Fury." *Military Times*, April 8.

Scheer, Steven. 2026. Reuters, March 31. "Israel to destroy 'all houses' near Lebanon border, defence minister says." <https://www.reuters.com/world/middle-east/israel-destroy-all-houses-near-lebanon-border-defence-minister-says-2026-03-31/>

Sifton, Sam. 2026. "The Iran War, in Numbers." *New York Times*, April 23.

Strain, Milo. 2026. "UA Fayetteville faculty committee recommends reinstatement of ousted Middle East Studies director." *Arkansas Today*, February 27.

If you have comments, suggestions, or ideas for future issues and new features
please contact:

Diana B. Greenwald at dgreenwald@ccny.cuny.edu, Sebnem Gumuscu at sgumuscu@middlebury.edu, and
Samer Shehata at shehata@ou.edu

AMERICAN POLITICAL SCIENCE ASSOCIATION

MENA POLITICS

VOL. 9 ISSUE 1, SPRING 2026

TABLE OF CONTENTS

News from APSA MENA Section	2
News from the APSA MENA Program	4
News from the Arab Political Science Network	6

I. Feature Essays

A Proustian Approach to the Study of Politics	8
<i>Nathan J. Brown</i> , DOI: 10.5281/ZENODO.20190007	

II. Research Symposium: AI in MENA: Authoritarian Governance, Global Competition, and Sociopolitical Impact

Introduction	16
<i>Eric Lob and Selin Bengi Gümrükçü</i> , DOI: 10.5281/ZENODO.20190147	
How the Global AI Race Shapes the Gulf's Hedging Strategy?	18
<i>Yasir Atalan</i> , DOI: 10.5281/ZENODO.20190218	
How Path Dependence Shapes Gulf Monarchies' High Tech Strategies	25
<i>Andrew Leber</i> , DOI: 10.5281/ZENODO.20190290	
Governing AI Surveillance through Legal and Social Knowledge:	30
An Islamic and Anthropological Framework	
<i>Abdullah Omran</i> , DOI: 10.5281/ZENODO.2019032	
AI and Digital Authoritarianism in MENA	35
<i>Alexandra Siegel</i> , DOI: 10.5281/ZENODO.20190374	
Techno-Orientalism: Artificial Intelligence and the "Other"	43
<i>Arshin Adib-Moghaddam</i> , DOI: 10.5281/ZENODO.20190426	
Building AI Capacity Between Beijing and Washington: The Gulf's China Strategy	48
<i>Guy Burton</i> , DOI: 10.5281/ZENODO.20190481	

AMERICAN POLITICAL SCIENCE ASSOCIATION

MENA POLITICS

VOL. 9 ISSUE 1, SPRING 2026

TABLE OF CONTENTS

III. Roundtable Discussion: AI in MENA Politics Research

Carolyn Barnett in conversation with Muhammad S. Abdo, Tariq Adely, Cinzia Bianco, 55
Ashrakat Elshehawy, Robert Kubinec, and Michael Robbins,

DOI: 10.5281/ZENODO.20190515

AI Will Drive the Coming Integration of Interpretive and Positive Methodologies 74

Richard A. Nielsen, DOI: 10.5281/ZENODO.20190646

News from the APSA MENA Section

Dear Colleagues,

I am delighted to write my first letter for the APSA-MENA Spring Newsletter as the newly elected Chair for the APSA-MENA Section.

I want to thank you for your trust and for allowing me to serve the section. As I previously stated in my election statement, my primary goal is to expand the section's outreach, promote the diversity of its members, and pay particular focus to MENA-based scholars and graduate students so they can benefit from the wealth of resources offered by the American Political Science Association. I have established solid research networks in the region over the past years, which I aim to integrate with the existing APSA-MENA network.

To achieve these goals, the section is working with APSA to establish a donation fund to support travel and visa expenses for MENA-based scholars and graduate students. We will be sharing more information about the fund and how you can contribute shortly. The section is also actively working with APSA-MENA workshops, led by Andrew Stinson and Dana El-Issa, to expand the organization's outreach to MENA-based scholars (see their letter in the Newsletter for more information). We just finalized the details for the upcoming Research Development Group (RDG), to be held during the upcoming APSA Annual Conference in Boston. We received approximately 60 applications and 6 were selected. The workshop will be led by Adria Lawrence, Shamiran Mako, and Marwa Shalaby. Additionally, the Arab Political Science Network (APSN) is holding a research mentorship workshop on gender and politics in the MENA region. The workshop is scheduled for May 22 in Cairo, Egypt, and will be

led by Nermin Allam, Marsin Alshamary, Hoda Alsadda, Zina Sawaf, and Marwa Shalaby.

Our Section was allocated 12 panels in the upcoming APSA meeting in Boston. Thanks to our program chairs, Chantal Berman and Sean Yom, we have a fantastic lineup of panels and poster sessions. Please make sure to participate and volunteer. APSA takes panel attendance quite seriously and allocates panels primarily based on physical attendance.

At the Section's Business Meeting of the 2026 Annual Conference in Boston, we will present the awards for "Best Dissertation" (defended in the previous academic year), "Best APSA Paper" (presented at the 2025 Annual Conference), "Best Article", and "Best Book" (published in 2025). I am deeply grateful for the members of the Best Dissertation Committee: Killian Clarke, Ahmed Mohamed Ezzedine, and Alexandra Blackman; members of the Best Paper Committee: Carolyn Barnett, Amir Reda, and Kevin Mazur; members of the Best Article Committee: Fiona Adamson, Sarah Parkinson, and Salam Alsaadi; and the members of the Best Book Committee: Steven Heydemann, Diana Greenwald and Rachel Brown. Finally, the Section will partner with POMEPS during the upcoming annual conference to sponsor a reception open to all scholars working on or in the Middle East and North Africa. I will be sending more information on the exact date and location of the reception as well as our Section's Business Meeting.

(CONTINUED ON NEXT PAGE)

News from the APSA MENA Section (continued)

Finally, on behalf of our entire section, I want to thank our previous Chair, Curtis Ryan, our outgoing Treasurer, Zahra R. Baber, and At-Large member Lama Mourad for their years of exceptional service to the section. I welcome our new board members: Adria Lawrence as our treasurer and Lisel Hintz as our At-Large member. Please make sure to follow our updates on the [APSA-MENA website](#) and [APSA Connect](#). You can also follow section news on X, LinkedIn, and Bluesky at @APSAMENA. Feel free to send me any opportunities or announcements you would like to share with the section. I would be happy to circulate these on your behalf. I wish you all a great end of the semester and a restful summer.

Section Chair,

Marwa Shalaby
shalaby2@wisc.edu

News from the APSA MENA Program

The American Political Science Association's [MENA Program](#) is a multi-year effort to support political science research and networking among early-career scholars across the Middle East and North Africa. Through a series of workshops, departmental collaborations, research grants, and other opportunities, the program extends APSA's engagement with the international political science community and strengthens research networks linking American scholars with colleagues overseas. The goal of APSA's MENA Workshops, generously funded by Carnegie Corporation of New York through 2025, is to enhance the capacities and resources of political scientists in the Arab MENA region, while also providing a forum for supporting their ongoing research.

In January 2026, APSA held a four-day methods training workshop in Doha, Qatar, for early-career scholars titled “Engaged and Participatory Research in the MENA Region.” Held in partnership with the Doha Institute for Graduate Studies, the training was led by Drs. Lara Khattab (Doha Institute), Stacey Philbrick Yadav (Hobart and William Smith Colleges), Sarah E. Parkinson (Johns Hopkins University), and Ammar Shamaileh (Doha Institute). The program convened nineteen doctoral students and early-career scholars from across the MENA region to examine the challenges of conducting engaged and participatory research in applied contexts. The workshop marked the fourth methods training following the launch of the [MENA Methods Program Initiative](#) in 2023.

As part of its ongoing partnership with the [Institute for Qualitative and Multi-Method Research](#) (IQMR) at Syracuse University, APSA will sponsor three MENA-based scholars from Morocco and Egypt to attend the 2026 Summer Program in the U.S. The partnerships aim to expand support for MENA-based scholars by providing advanced methodology training at premier U.S. institutions.

APSA continues to support the collaborative initiative between the American University in Cairo (AUC) and the Faculty of Economics and Political Science (FEPS) at Cairo

University through a series of joint research forums aimed at promoting sustained scholarly exchange between the two institutions. The initiative is part of APSA's MENA Departmental Collaboration Initiative, through which APSA partners with political science and related departments at universities in the Arab World to design tailored programming that supports local faculty and graduate students. Since the project's launch in September 2024, ten bi-monthly forums have been held to date at AUC's Tahrir and New Cairo campuses. The forums have provided a regular platform for presenting work in progress, receiving constructive feedback, and strengthening academic networks across institutions in Egypt. The initiative has continued this year with additional forums underway and a co-sponsored international conference held in April 2026 at AUC, in partnership with Villanova University and the Arab Political Science Network (APSN).

Looking ahead, APSA will partner with the MENA Politics Section to host a Research Development Group (RDG) for seven early-career scholars at the September 2026 APSA Annual Meeting in Boston. Selected scholars will take part in the full-day seminar on Wednesday, September 2 to discuss and receive critical feedback on article-length manuscripts in progress. Six senior scholars will serve as seminar facilitators and discuss

(CONTINUED ON NEXT PAGE)

News from the APSA MENA Program (continued)

sants, with two scholars assigned as discussants for each paper. In addition to focused discussions on each paper, the seminar will provide tips on academic writing and advice for publishing in peer-reviewed journals and other professional development topics. The RDG provides participants with the opportunity to advance their research manuscripts for publication, engage in the annual meeting, and build scholarly networks with colleagues.

The MENA Program continues to support graduate students and early-career scholars through virtual mentoring and professional development opportunities. This summer, APSA will pair several graduate students with advanced-career scholars for one-on-one mentorship on their research projects through the [MENA Mentorship Initiative](#).

For more information about the APSA MENA Program, visit the project website: <http://web.apsanet.org/mena/> or contact menaworkshop@apsanet.org.

APSA MENA Project Team
American Political Science Association

News from the Arab Political Science Network

The Arab Political Science Network is a collaborative scholarly initiative that seeks to support Arab political scientists.

Greetings to all! We hope you continue to be safe and well amid the ongoing atrocities and uncertainties plaguing some of our countries and the region.

The Arab Political Science Network (APSN) started 2026 with several plans for workshops that were eventually postponed due to the wars launched by the United States and Israel on Iran and Lebanon, and the subsequent retaliations on the Arab Gulf states. One of the workshops is scheduled to take place in Cairo, Egypt on May 23-24, in collaboration with the Women and Memory Forum. This will be the culmination of APSN's annual research development and mentoring program, which focused this year on Women and Gender Politics in MENA. We are expecting 14 graduate students from the region to come together for discussions and feedback on their research manuscripts.

Our support to scholars based in Arab countries doesn't stop at research workshops but also to increasing their published outputs. Last month, APSN, in collaboration with the *Journal of South Asia and Middle Eastern Studies* ([JSAMES](#)) at Villanova University, organized a [publications workshop](#) at the American University in Cairo (AUC). Thank you to Samer Abboud, Nadine Sika, and Mazen Hassan for leading this effort, and we look forward to seeing a special issue from this workshop. If you are working on your first book, whether in English or Arabic, stay tuned for our upcoming call for proposals that will guide you through the process of turning your PhD dissertation into a book and provide some critical feedback.

On March 31, we launched our latest webinar series focusing on Political Economy in collaboration with the AUC-project [Pathways Beyond Neoliberalism](#). The series brings together researchers and practitioners to re-think the global political economic order by foregrounding critical and multidisciplinary approaches to capitalism and decoloniality, examining how histories of extraction, securitization, and other processes continue to shape contemporary economic life. The first webinar looked at [What is Political Economy from the Global South](#) and the second discussed the *Political Economy of Sanctions and War* (video to be uploaded soon). We are currently planning two more sessions in June, the first will focus on *Oil, Power, and Gulf Economies*, and the second on *Competing Developmental and Economic Models*.

We have been updating our list of digital resources available on our website and on [YouTube](#). We encourage you to check them out, whether you are looking for research methods discussions or interested in viewing previous sessions on [Politics of Infrastructure](#), [Politics and War in MENA](#), or [Syria in Comparative Transition](#).

APSN will be attending both APSA and MESA annual meetings, so we look forward to seeing those of you attending. We will share more information about our panels in due course. Finally, we would like to remind you of our specialized newsletters, the bi-weekly [al-Dal](#) that recommends new non-fiction / social science Arabic and English books, and the monthly [al-Zad](#) that

News from the Arab Political Science Network (cont.)

now focuses on curating thematic social science academic journals. Most recently, we published two special issues, the first on [Sudan](#) and another on [Iran and the Arab world](#). These newsletters complement the Arabic book review portal [al-Salon](#) that publishes Arab reviews for non-fiction books about the MENA region and beyond. We welcome any suggestions or possible contributions by writing to editor@al-salon.com.

Best to all!

Ahmed Morsy (on behalf of the APSN team)

A Proustian Approach to the Study of Politics

Nathan J. Brown



Nathan J. Brown is Professor of Political Science and International Affairs and the Assistant Dean of Graduate Studies at the Elliot School of International Affairs at The George Washington University (GWU).

Editors' Note: Last year, Nathan Brown received the 2025 Michael E. Brown Research Prize, the foremost research award at GWU's Elliot School of International Affairs. A lightly edited version of Brown's address on September 24, 2025, follows below. Brown, as many of our readers know, is a leading scholar of Middle East politics, judicial and constitutional orders, and – the subject of his current book manuscript, which he draws on in his address below – the relationship between religion and state. In his speech, Brown reflects on his own work and what he identifies as an area of promising overlap in both social and area studies: to “help people see what each other sees.” In many ways, this is a speech about boundaries and boundary-crossing – from one's subjective position to that of another, and from the descriptive to the normative. Toward the end of his remarks, Brown situates his own approach – to boundaries and to their transgression – in the (then) current context of Israel's war on Gaza, the discourse in US policy circles, and the cycles of activism and repression on our college campuses. Because these remarks were delivered orally, we have included a minimal number of academic references for readability.

It's an honor to receive an award named for my former dean, Michael Brown. Let me first assure you that there is no nepotism—I have traced my ancestry back some generations and can identify every Brown I am related to—and Michael is not one of them. So that makes it even more special to bear his name.

But at a time when the benefits of scholarly research are increasingly questioned; when analysis of international affairs seems as deeply polarized as every other subject, from employment statistics to the weather; when “professor” is not always a term of respect, is recognition by one's colleagues still such an honor? The current moment has made me reflect a bit more personally on why I do what I do. And I will use this opportunity to share some of those reflections with you.

I will begin my personal reflections first, *pretentiously*, by quoting Marcel Proust. I will proceed second, *pedantically*, by lecturing you on how Proust has been misinterpreted, then, in turn, *pompously*, by comparing my work to Proust's sense of aesthetic mission. But most of my talk will then proceed *prosaically* (with a word about *punctuation*) by trying to explain what I do—which has less to do with discovering truth and more to do with helping people understand each other when they see the truth so differently. That will be the bulk of my remarks. But I will then inject some *political* thoughts about what *I* see as truth. I will conclude *practically* with some thoughts about how to present truth when it can be seen as provocative.

So on to my pretentious beginning, which is Proustian—literally.

When I was twenty-two years old and just about to begin a doctoral program, my best friend from high school told me that the best way to take advantage of the privilege of being human was to read (or reread) a volume of Proust every year. I was especially eager to retain my humanity while completing a PhD, so I began. “For a long time I used to go to bed early,” Proust began slowly.¹ But I was even more slow. I took my comprehensives, conducted field work, and wrote up my dissertation—and was still just starting the second volume. I soldiered on as I became an assistant, an associate, and then a full professor. I became a parent and then a grandparent. I carried Proust to different cities and countries and continents. In 2024, almost four and a half decades after reading about Proust’s going to bed, I finally finished all seven volumes; I did so, I am proud to say, before becoming professor emeritus. So I failed the timeline, but the undertaking still did me some good. And in the process I came across a particular passage in volume 5 that I had heard before, since it is frequently quoted:

The only true voyage of discovery, the only fountain of Eternal Youth, would be not to visit strange lands but to possess other eyes, to behold the universe through the eyes of another, of a hundred others, to behold the hundred universes that each of them beholds, that each of them is.

A more extended version of that passage is frequently cited to explain the virtues of international travel. I even found some of its words emblazoned on the wall of a lounge in Charles de Gaulle Airport.

But here is where I move from the *preten-*

tious to the *pedantic*, since it is clear from the passage that it has nothing to do with traveling—just the opposite. Flying to strange lands with the senses we already have does us no good since we still see with our old eyes. We need to see things through the eyes of others.

And that brings me quickly to the *pompous*, where I claim to understand Proust better than the decorators of the Paris airport. What Proust actually suggested was that there were certain uniquely gifted artists whose aesthetic senses allowed them to see things none of us would otherwise be aware of and then to make those worlds accessible to us.

So am I such an artist?

Heavens no; here my pomposity tapers off. I do not see things others cannot. Instead, I try to help people see each other—or to help people see what each other sees. So here I turn from the *pompous* to the *prosaic* and to punctuation to focus in the body of my remarks.

I see much of my job as helping people understand each other. And my main tool is a punctuation mark—the semicolon, that allows us to keep two thoughts in mind, linking them while also treating them separately and sequentially. I use the semicolon to pause between understanding and judgment.

The subject matter of the social sciences is much older than the current range of academic disciplines. People interact, govern, trade, communicate, move, cooperate, fight, argue, deliberate, dominate, vote, and make rules—and they started doing so long before

¹ This, and subsequent references to Proust, come from his seven-volume magnum opus, *À la recherche du temps perdu* (In Search of Lost Time), published between 1913 and 1927, and for which many English language translations are available.

a few people decided to devise ways of understanding the patterns by which they do so. Many—but not all—of those within social science disciplines have felt it helpful to draw a border between understanding and judgment; between the empirical and the normative; between explanation and prescription. Many seek to understand the world as it is, not judge it for what it fails to be.

But not all. The drawing of a line between empirical and normative scholarship is controversial, precisely because it seems to push good and evil beyond the margin. Worse, it can push them underground, where they operate fully but out of conscious view.

I find it can be very useful to distinguish what is from what ought to be—and I therefore try to present many of my understandings without judgment. But I also insist on the benefit of crossing over that border between them. That is where a semicolon can be useful, especially in a world in which people's experiences, expectations, and presumptions are so different. It can pay to pause to understand different points of view before proceeding to prescription, policy, or practice. In a world in which people with such diverse lives come into increasing contact, liberal use of a semicolon makes managing of difference, learning, and communication much easier.²

And there is a policy payoff as well—when we proceed this way, we can speak not truth to power but truths; we can try to induce those who wield power to understand that the world as they see it is not the world as many others experience it.

And here let me move beyond punctuation to

method. I think those who approach scholarship like me work to make the familiar appear unfamiliar and to render the unfamiliar more familiar.

In the first part of this method, we work to complicate what we take for granted; to help us understand that what we see around is not always natural or inevitable; to make us less likely to take what we know as the only way to do or see things. We can make the familiar appear a bit...odd.

And the second thing we can do—making the unfamiliar familiar—prevents us from recoiling from anything that we initially see as exotic or anomalous.

Moving within and across cultures, “area studies” as an enterprise often works to explain phenomena by translating what is unfamiliar into familiar terms. Instead of treating the unfamiliar as alien, it might use analogy or metaphor to make it more familiar.

The methods I describe are designed not to eliminate complexity and difference. They help us cope with them and help overcome barriers to understanding. They allow us to understand our differences as historically contingent rather than essential.

Let me explain that in terms of my current work on religion and state. I find too much writing that juxtaposes a secular West against a religious Middle East. And many that trace that distinction to distant historical origins: seventh-century Mecca or seventeenth-century Westphalia.

Such an approach is clearly taken in many

² And I practice what I preach, using the semicolon in my writing, silently but unmistakably in my speech and now deploying it as a metaphor in [my newly launched podcast series](#), “Semicolons and Teardrops in the Middle East”.

attempts to understand the history of higher education. In a post-Westphalian West, and in a United States with a wall of separation between church and state, we see universities that relegate religion to the private sphere; in the Middle East, where Islam knows no separation between religion and state, we see religious universities, like Egypt's al-Azhar, or state universities that still teach shari'a and other religious subjects.

This division may seem clear, but it is not quite right. Let me start by making clear that Western universities do not fit this familiar pattern—or any pattern—so easily. Consider that state universities in the real German state of Westphalia teach Catholic and Protestant theology under the supervision of religious authorities. That's not unusual in Europe—but it's not universal either. French universities are barred from teaching theology—except in Alsace.

Closer to home, we see the country's first institution of higher education, Harvard, guaranteed its privileges because, according to a clause of the Massachusetts constitution in effect today, “the encouragement of arts and sciences, and all good literature, tends to the honor of God [and] the advantage of the Christian religion,” (Massachusetts Constitution chap. 5, § 1, art. 1). The University of Michigan was founded to give effect to a treaty obligation of the United States government to provide support for a Catholic college to the Ottawa, Chippewa, and Potawatomy peoples.³ Our own George Washington University was founded to educate Baptist clergy and provide missionaries to Asia and the American West, and—closer to home—to prevent the population of federal offices by Catholics educated at nearby Georgetown.⁴

These familiar institutions in Westphalia and Washington DC, Alsace and Paris, are more varied and strange than we realize.

American universities look a bit odd in global terms in a manner we often overlook. There are few other systems of higher education that have defined “religion” as a generic subject and channeled that subject into something that looks like an academic discipline; we have rounded up many of the specialists and put them in a department. But we have done so only fairly recently.

What explains these patterns if they do not fit into a secular West versus religious rest? Across the globe, as extensive and intensive state apparatuses were built—generally in the nineteenth and the first half of the twentieth centuries—they oversaw the emergence of formalized structures for mass elementary and secondary education. In fact, they invented those stages. And they clearly defined and demarcated post-secondary stages in institutions that combined education, scientific and scholarly research. And as this formalized and official educational apparatus arose, those building them often encountered structures, practices, and intellectual traditions that were religious in nature.

That was not a Western pattern; it was a familiar process that took place in Sri Lanka, Mexico, Egypt, and the United States of America at roughly the same time. But if most states built universities, each one handled religion a bit differently—not based on any essential nature of Christianity, Islam, or Westphalia. Instead, basic choices simply reflected society and politics at the time.

In the United States, structures of religion

³ The text of that treaty can be found [here](#).

⁴ This point is addressed in my forthcoming book, *Religion and State*.

and education were shaped by struggles, mostly in the late nineteenth and early twentieth centuries, among Protestants and between Catholics and Protestants. There were conflicts over the education of indigenous peoples, teaching religious texts and doctrine in public schools, and what “nonsectarian” meant. And the struggles played out in various venues, including in federal courts and the Internal Revenue Code.

In Egypt, around the same time, higher education institutions were shaped by a very different set of struggles among Ottoman, Egyptian, and British officials. Some of the issues they grappled with were similar to those that emerged in the United States, and some were different. The outcomes, in both cases, were shaped by the demands of a rapidly expanding bureaucratic state and economic development patterns and trajectories.

In 2023, a public Swiss university advertised for a position in Jewish theology. In addition to scholarly credentials and knowledge of relevant languages, there was one additional requirement: only members of the Catholic Church should apply. The requirement made perfect sense in terms of the way that religion operates within the context of the Swiss state—and is utterly confounding outside its borders.

In this way, what appears to be familiar, a Western-style university, appears far less uniform, inevitable, or even “Western”; non-Western universities appear far more familiar, and not distinctly “non-Western”. And the effects of decisions, often made for reasons long forgotten, continue to this day.

Seen this way, our task in the social sciences is fourfold—and the first three involve understanding. First, we explore the world(s)

that exist within our assumed borders—and the various ways those living within these borders experience, enforce, and push against them. But second, we seek to explore the lives of those on the opposite sides of these borders. We seek to communicate not only within one world but to translate between these worlds.

Why? That is the third task: to understand how those differences shape outcomes—and we do so, or at least I do so, in the hope that it will make us less likely to hurt each other. The goal is not agreement; a world in which everyone agreed, devoid of difference, disagreement, or borders, would be a world without politics, one without growth or learning, one of monotonous uniformity. A world in which people only enforce borders and never cross them, in which deviations from the familiar are viewed as crimes against nature, would be violent and confining. We explore disagreement and float above borders so that we may understand each other through our differences.

And that sounds so...positive. And nice. But sometimes we have to say things that might be a bit less nice. What I have discussed so far is a combination of odd curiosities and paths laid down in the past; the Brown award is to be for “policy-relevant understanding of important global issues.” We do not only explore the past; we make judgments about the present. And when we do so, we sometimes address not only what is and why but also what we think should be done—and that is the fourth task.

So now I move from the *prosaic* to the *political*, and this will bring me to the far side of the semicolon between understanding and judgment. While it can be harmful to rush to judgment, we do have to meander over there

eventually; to fail to do so would be to keep our minds so open that our brains would fall out.

How and when to do so has been a deeply divisive question in the academy in recent years, and none more so than in the set of issues I have focused on intermittently for three decades, related to Israel and Palestine. Since October 7, 2023, the issues that have never been easy to discuss seem to have become impossible to discuss.

In general, my own approach when dealing with such a challenge has been to linger on the “understanding” side of the semicolon and not rush to the “judgment” side. And while that suddenly became harder two years ago, it was also a time when it was more important than ever to understand why people were doing what they were doing—even if they were doing some awful things. Again, that did not mean avoiding any normative judgments; it just meant not leading with them.

So when I spoke as a scholar, I tried several things. I attempted to measure my words by assuming my audience did not agree with me. I tried to speak in paragraphs—not to render my answers tedious, but to give myself space to offer analysis and understanding of various points of view. I inclined toward understatement on most occasions—to lead with strong language can leave an audience feel shouted at. In an environment in which people are being told what to think and even what to feel, a voice that helped them work their way through the issues and the passions seemed more likely to be heard. I worked to avoid personalizing arguments; it was actions and deeds (and sometimes words) that needed to be assessed, not people’s souls. I would save my conclusions—and I did have conclu-

sions—for the conclusion. And I’m nearing my conclusion now.

I have explained how I tried to operate on the understanding side of the semicolon. But that gives little guidance on when or how to step over to the far side, to render judgment. And in this case, what it overlooks is that, however one assesses justice—however virtue and rights are distributed—power is horribly out of balance in ways that have deeply problematic and deadly results.

Denial and dehumanization, for instance, are in rich supply and are nobody’s monopoly. The number of people I heard dismiss the atrocities committed against Israeli civilians in the October 7 attack and in the subsequent assault on Gaza dismayed me. The same can be said of dehumanization. This is a common phenomenon in wartime, and shrillness is the coin of much public discussion of the Gaza war.

But if denial and dehumanization are widespread, they have not been evenly distributed. The extent to which leading political figures—those in senior positions of authority—in the United States and Israel have resorted to eliminationist rhetoric has no easy parallel.

The tools of political authority are not distributed evenly, nor has been the firepower of those dismissing the death of their adversaries. Those who call for “Palestine from the river to the sea” are denounced in the halls of Congress, often for things they do not mean; those who make similar territorial claims for Israel (and do so in a way that asserts not only the rights of the Israeli state but also the superior rights of the Jewish people, and sometimes are explicitly linked to proposals to forcibly displace Palestinians) are extended invitations.

In Washington DC after October 7, it was unremarkable to claim that each Israeli death on October 7 is equivalent to thirty American deaths because of the relatively smaller population of Israel—and then to measure Palestinian deaths instead against fatalities in Syria or Darfur, varying the scale by which to measure killing in order to minimize or maximize the horror.

It is routine to hold Palestinians responsible for the death toll in their own ranks; to quote official Palestinian figures by noting they come from a government controlled by Hamas but to quote Israeli figures without reference to the exclusionary and supremacist ideology of those ruling the country. Similarly, official Israeli sources have claimed the ability to count accurately every “terrorist” killed, but appear utterly baffled in trying to provide even a rough estimate of the number of civilians killed.

Strident and shrill campus rhetoric against Israel was combatted with legal tools, but calls for continuing the crusades against Muslims or using nuclear weapons on Gaza were barely noticed. A United States government apparatus designed to combat discrimination in education has just been dismantled in nearly all of its aspects except to protect Jews who support the Israeli government—and those same tools are now called upon to target Jews who are critical of Israel.

This uneven distribution of power has shaped the framing of policy discussions and the information that is fed into them. Think tanks began to produce detailed “day after plans” a month into a war, utterly disconnected from the scale of the destruction, the open-ended nature of the Israeli military campaign, and the clear statements of Israeli policy that there would be no day after in the sense pun-

its used the term.

I said I would try to speak in paragraphs and I have. But in that last one, I trespassed most of the other boundaries I just claimed to follow—reluctantly, with real misgivings, but also with deep conviction that, in such abnormal times, horrific atrocities and murderous words do have to be labelled as such lest we enable actions we purport merely to understand.

There are many who add that we have an obligation to do so. I do not disagree, but I want to end on a cautionary and practical note.

I find myself very slow to use one word in a previous sentence: “we” –especially when imposing obligations. It is one that I think *we* as scholars need to use carefully. Why? Because I think one contribution we can make is by helping people peer over the border between us and them.

But to speak in a singular voice can be lonely. And to be alone in a dangerous time is dangerous.

This leads me to turn back to Proust for some practicality. I suspect that is the first time that sentence has been uttered in any setting. Writing in Paris in the middle of the First World War, Proust observed, “The surest way of being convinced of the excellence of the cause of one party or the other is actually to be that party.”

Proust was writing during a war that not only caused great bloodshed but also the celebration of that bloodshed. The French and Germans in 1914 mistook their patriotism and moral sensibilities for military capabilities.

Today, even the most compassionate people

seemed for a long time to only muster at best a limp “Yes, but” about the suffering and death of those on the opposite side. So over the past two years, more than a semicolon is needed. “Yes, period.” should be not only a complete sentence but an entire paragraph.

it more likely that a cry of conscience, when it does emerge loudly and clearly, seems less like a partisan chant.♦

But another paragraph is needed to describe the horrors, the cruelty, and the violence visited at once both blindly and deliberately on the over two million people of Gaza.⁵

And in such a context, those who seek to peer on the other side of a boundary will find themselves exposed. The dangers are real. But my personal experience suggests that there is a reward for crossing to the far side of the semicolon for those who do so slowly.

Public discussions can be polarized, personalized, and characterized by rhetorical inflation. But I have found that makes many (especially non-specialists) hungry for calmer discussions, simply so the din doesn’t drown out their own understanding. Like many colleagues, I have participated in sober and reasoned (if distressed) discussions in classrooms and small groups. I have addressed a Gaza encampment, a synagogue, and mixed audiences; I have spoken with a variety of individuals whose public certainty often obscures private distress.

My experience is that the hunger for informed and thoughtful discussion is great; that speaking softly generally makes it easier for most listeners to hear—and it also makes

⁵ Author’s note: In the address I gave in September 2025, I judged Israeli conduct in Gaza in extremely strong terms but did not use the term “genocide.” I have been asked about that since. Was I pulling my punch? Yes, I was. And I have stopped doing so. I do believe the term applies. For the reasons I describe in the address, however, I try not to lead with it when I do use it. I want it to emerge as a conclusion of analysis, not as the premise from which analysis begins. But that is no reason to avoid the term. The point, then, is not to evade the term but to resist a style of argument in which the presence or absence of one word does the work that analysis should do. Those of us in a profession built on peer review should know how easily it can become peer pressure.

Research Symposium: AI in MENA: Authoritarian Governance, Global Competition, and Sociopolitical Impact

Introduction by Eric Lob and Selin Bengi Gümürkçü



Eric Lob is an associate professor of politics and international relations at Florida International University. He is also a non-resident scholar in the Middle East program at the Carnegie Endowment for International Peace. His research examines the intersection of politics and development in Iran and the Middle East. He is the author of Iran's Reconstruction Jihad: Rural Development and Regime Consolidation after 1979 (Cambridge University Press, 2020). His articles have appeared in The International Journal of Middle East Studies, Iranian Studies, Middle East Critique, The Middle East Journal, Middle East Policy, The Muslim World, Third World Quarterly, and others.



Selin Bengi Gümürkçü is an Assistant Teaching Professor at the Department of Political Science at Rutgers University, New Brunswick. Her research focuses on contentious politics, political parties, and democratic backsliding, and, more recently how AI interacts with these processes and outcomes. She is the author of Protest and Politics in Turkey in the 1970s: Echoes of Dissent (Routledge, 2026). Her publications have appeared in peer-reviewed journals such as Terrorism and Political Violence, Turkish Studies, Southeast European and Black Sea Studies, Journal of Balkan and Near Eastern Studies, and in edited volumes.

Artificial intelligence (AI) is rapidly emerging as an important force shaping many aspects of political, economic, and social life. Yet its implications, particularly political ones, are far from uniform and are mediated by regional contexts and institutional configurations. In the Middle East and North Africa (MENA), AI has already emerged to be much more than a mere technological innovation. It is also a politically embedded instrument through which economic futures are imagined, state power and popular dissent are exercised, and global positioning is renegotiated. This symposium examines how Arab Gulf states and others in the region are integrating AI into authoritarian governance, competing in the international AI race, and

cooperating with the United States and China in the process. It also explores the technology's sociopolitical, cultural, religious, and ethical implications.

The significance of AI for MENA stems in part from the region's ongoing efforts to navigate enduring structural challenges. They include, but are not limited to, economic diversification, demographic change, domestic discontent, and military conflict. In this context, AI is perceived and portrayed as a solution to some of these challenges. Using AI to diversify their economies, Arab Gulf states have aggressively positioned themselves as AI hubs within the global economy. As Yasir Atalan shows, these states have done so by

investing in data centers, cloud infrastructure, and corporate partnerships. In the process, they expect that AI will enable them to transform key industries, enhance productivity, and generate new sources of revenue and employment. From the perspective of path dependency, Andrew Leber posits that the efforts and expectations of these states concerning the critical juncture of the AI revolution are conditioned and constrained by pre-existing parameters, such as past investments in state-owned enterprises, power grids, and human resources.

In addition to being economically important, AI is increasingly central to governance practices across the region, from public service delivery and urban management to state surveillance and information control. As Alexandra Siegel argues, the expansion of data infrastructures and algorithmic tools has enhanced state capacity while also raising concerns about the deepening of authoritarian practices, particularly in countries where institutional constraints remain limited. It also creates ethical and legal problems regarding surveillance technologies, as Abdullah Omran suggests. At the same time, the political significance of AI in MENA extends beyond state strategies. Non-state actors and various opposition forces are also engaging with AI and incorporating it into practices of communication and mobilization.

As with other areas, AI in MENA is shaped by global hierarchies and dependencies. As Guy Burton asserts, the region's efforts to become an AI hub are embedded in transnational technological ecosystems dominated by major powers and multinational corporations. Questions surrounding data sovereignty, infrastructural development, and regulatory alignment reflect broader geopolitical dynamics, as states navigate between integra-

tion into global systems and aspirations for digital autonomy.

Inside MENA, AI should also be considered in relation to the persistent phenomenon of orientalism. As Adib-Moghadam argues, contemporary AI systems are largely developed from a Euro-American technological and epistemological frame of reference. Consequently, they are embedded with Euro-centric assumptions, classifications, and biases that reproduce reductionist and essentialist representations of MENA societies. These dynamics risk reinforcing longstanding orientalist tropes and depicting the region as exceptional, unstable, and culturally static. This is particularly the case with applications, such as automated content moderation, security profiling, and data classification, where contextual nuance is often limited.

Taken together, these developments underscore that AI matters for MENA. It is not only a set of technological tools, but a domain through which longstanding political and socioeconomic questions are being raised and reconfigured. By situating AI within the region's specific institutional and geopolitical contexts, this symposium stresses the need for rigorous, grounded analysis of the interplay among technology, power, and transformation.♦

How the Global AI Race Shapes the Gulf's Hedging Strategy?

Yasir Atalan



Yasir Atalan is deputy director in the Futures Lab at the Center for Strategic and International Studies as well as faculty fellow at the Center for Data Science at American University. His research focuses on technology and international security, with particular emphasis on how artificial intelligence is reshaping conflict dynamics.

The global AI race is increasingly defining international politics because states now compete over the infrastructure that makes advanced AI possible. While algorithms and model architectures diffuse through research and markets, compute (high-end chips, hyperscale data centers, and cloud platforms) remains concentrated, expensive, and politically controllable, which makes it a strategic bottleneck. Control over compute has therefore become a tool of competition, with the US-China rivalry leveraging it through export controls, licensing, and platform governance. Gulf states such as the UAE, Saudi Arabia, and Qatar occupy a key position in this emerging order because they are deploying sovereign wealth and energy advantages to position themselves as regional AI hubs, even as they remain embedded in US security architectures and deepen economic engagement with China. Their AI investments, therefore, cannot be understood as purely developmental or economic diversification strategies. As AI supply chains become securitized, these infrastructure choices are increasingly embedded in alliance politics, exposing Gulf states to compliance demands and alignment pressures. In this intervention, I examine how the global AI race narrows and reshapes Gulf hedging options under these evolving constraints.

Why Compute Matters for the AI Race?

To understand why Gulf AI strategies are increasingly entangled in alliance politics, it is first necessary to clarify why compute has become the decisive resource in the global AI race. Among algorithms, data, and compute, which are three components of advanced AI (Buchanan 2020), compute has a unique character. It has a rivalrous and binding constraint because advanced AI training and large-scale inference depend on scarce physical systems and reliable compute power rather than code alone. Scaling-law research shows that AI model performance tends to improve in predictable ways as training compute increases (Kaplan et al. 2020). This is why a competitive AI stack requires advanced chips, hyperscale cloud and data center capacity, orchestration software, specialized engineering talent, and continuous electricity and cooling – all of which must operate together to produce competitive performance (OECD 2023; Scharre 2024). Even when open models lower entry barriers, infrastructure concentration and capital intensity continue to limit who can compete in advanced AI systems (OECD 2023; Scharre 2024). In this sense, compute functions as the binding constraint that structures competition in the AI era.

This technological dependence matters politically because compute is not only decisive for advanced AI but also excludable. Contrary to data and algorithm, compute's excludable and measurable nature makes it governable, turning it into a strategic instrument rather than a neutral market commodity (Sastry et al. 2024; Scharre 2024). States and firms can deny access to compute through ownership, export licensing, cloud account controls, and compliance systems embedded in provider operations (Sastry et al. 2024). Crucially, these restrictions are enforceable because compute can be counted and tracked in ways that data and algorithms often cannot. Governments and cloud providers can set benchmarks for how much computing power is being used, monitor the size of the chip clusters involved, and rely on provider records that show who is accessing large-scale compute and when. (Shavit, 2023; Vaynman, 2021). This combination of denial and traceability transforms compute from a technical input into a strategic asset capable of structuring international alignments and shaping geopolitical bargaining.

US-China Competition and Politics of Alignment

Given the nature of compute, the US-China AI rivalry appears less as a race over model quality and more as a contest over control of compute chokepoints within the AI stack. Advantages in chips, fabrication equipment, cloud architecture, and data-center capacity generate cumulative geopolitical spillovers in productivity, military capability, and rule-setting power (Ding and Dafoe 2021). This is why both the US and China sought to reduce

dependence on each other's supply chains in the AI ecosystem. This logic underpins US export controls, which target advanced semiconductors and manufacturing bottlenecks in order to transform infrastructure dependence into strategic leverage.¹ Yet this strategy is debated. Broad denial may protect security but risks accelerating substitution and alienating third states, including US partners, from frontier access. By contrast, selective conditional access through US-governed cloud platforms can preserve influence, sustain ecosystem dependence, and facilitate coalition building under monitoring regimes (Sastry et al. 2024).

If this debate intensifies, further securitization of AI supply chains is likely. This dynamic aligns with weaponized interdependence, where technological hub states convert network centrality into coercive power (Farrell and Newman 2019). With jurisdiction over critical nodes, a state can impose a chokepoint effect by denying advanced chips or cloud capacity (Farrell and Newman 2019). Over time, sustained restriction and compliance pressure can push the system toward bifurcation, hardening rival ecosystems and narrowing strategic flexibility (Woods 2025). Crucially, this trajectory is reinforced by the dual-use character of compute. When technologies serve both civilian and military purposes, states face acute information asymmetries about intent and usage, which increases incentives to design institutions that reduce uncertainty through monitoring and verification (Vaynman and Volpe, 2023). This is why the US and China could embed verification into access to compute through monitoring and reporting requirements that further

¹ Bureau of Industry and Security (BIS), "Additional Export Controls: Advanced Computing and Semiconductor Manufacturing Items," Federal Register 87(197), October 13, 2022, 62186–62215 (2022-21658); BIS, "Additional Export Controls: Advanced Computing Items—Updates and Corrections," Federal Register 88(205), October 25, 2023, 73458–73517 (2023-23055); BIS, "Changes to EAR Controls for Advanced Computing Items and Supercomputers," Federal Register 89(234), December 5, 2024 (2024-28270).

securitize the supply chain. In the end, technological change steadily compresses hedging options for the GCC states (Lehdonvirta, Wu, and Hawkins 2025).

Gulf's AI Ambitions

These structural pressures become particularly visible in the Gulf where countries such as Saudi Arabia and the UAE have accelerated AI infrastructure investment by leveraging sovereign wealth funds and energy advantages to position themselves as regional compute hubs. In Saudi Arabia, Mohammed Bin Salman's Vision 2030 and Neom, Saudi Arabia's flagship smart-city and industrial megaproject, frame AI infrastructure as part of a broader political economy transformation that seeks to convert hydrocarbon wealth into long-term technological competitiveness and reshape domestic rentier relations (Hassan 2020; Herb and Lynch 2019; Hameed 2020). For these countries' AI ambitions, abundant and relatively stable energy supplies play a significant role because large-scale AI training and inference are power-intensive, making electricity availability and cooling capacity central determinants of where frontier compute can be located (Soliman 2025). In line with this, recent announcements indicate a sharp expansion in planned capacity, including Stargate UAE's initial 200 MW phase expected in 2026 with eventual multi-gigawatt ambitions, Saudi Arabia's forthcoming AWS region backed by multi-billion-dollar investment, and Qatar's reported 20-billion-dollar AI infrastructure partnership through Qai. This acceleration is also reflected in Figure 1, where planned installations place the UAE and Saudi Arabia among the top global locations in H100-equivalent

capacity, which is widely treated as a benchmark "frontier" compute accelerator for training and running large AI models (Pliz et al., 2025).

Initially, the Gulf investments were embedded in a portfolio strategy that balanced engagement with both US-led and China-linked cloud ecosystems (Chaziza 2024; Bakir and Al-Shamari 2025). This approach was visible early in the region's cloud architecture, including Alibaba Cloud's Dubai launch in 2016 and AWS's first Middle East region in Bahrain in 2019, alongside continuing openness to Chinese vendors such as Huawei Cloud's Riyadh region in 2023. But as frontier compute for AI becomes securitized, access to the US AI stack increasingly travels with governance conditions, including security assurances, audit requirements, and restrictions on technology transfer, as illustrated by the framing of Microsoft's investment in G42 around intergovernmental compliance commitments (Soliman, 2026).²

What Are the Hedging Options for Gulf States?

The future of compute will shape alliance dynamics in the Gulf because the practical space for hedging depends on how far the United States and China push exclusion through network chokepoints and how strictly those rules are enforced across supply chains and platforms (Farrell and Newman 2019). As long as access to frontier chips, cloud services, and governance frameworks remains mediated through US-centered chokepoints, Gulf states will find that the highest-value AI compute infrastructure is increasingly tied to compliance regimes and security assurances, which narrows the scope

2 Reporting of the deal suggested that any approval of AI chips sale is subject to Department of Commerce. See: <https://www.reuters.com/world/middle-east/microsofts-uae-deal-could-transfer-key-us-chips-ai-technology-abroad-2024-05-23/>

for unconstrained dual alignment (Farrell and Newman 2019). Washington already has some digital-security campaigns, such as efforts to exclude Huawei from 5G ecosystems and Clean Network, an initiative aimed at excluding Chinese telecom and digital infrastructure vendors from its networks. The options for the Gulf states become narrower if Washington broadens these existing digital-security campaigns into the AI domain by treating data centers, cloud access, and model training as alliance-sensitive infrastructure (Chaziza, 2024). More recently, Pax Silica, an initiative on securing critical supply chain security signed also by Qatar and UAE, is a signal in this direction by Washington to institutionalize security screening around AI infrastructure. Under those conditions, Gulf governments would face stronger pressure to protect critical AI systems from Chinese vendors and to adopt procurement, standards,

and auditing practices that demonstrate ecosystem separation (Chaziza 2024). This dependence dynamic could also deepen further if Gulf countries seek broader security cooperation with the US to protect data centers against physical and cyberattacks in an increasingly volatile region.

On the other hand, other mediating factors could play a role in shaping supply chain securitization. Great power economic statecraft is mediated by firms that own key parts of the compute stack, which can preserve limited space for Gulf states to maneuver even under intensifying rivalry (Chen and Evers 2023). As hyperscalers, chipmakers, and data-center operators bear the commercial and technical costs of compliance, they have incentives to lobby for narrower controls, clearer rules, and workable verification regimes. Because maximalist restrictions can threaten revenues or




Figure 1: Distribution of Compute Power for High-End AI (Pilz et al., 2025)

slow global scaling for these actors. Similarly, smaller and middle states retain agency in bipolar cloud networks through strategic hub selection and segmentation (Lehdonvirta, Wu, and Hawkins 2025). By hosting multiple cloud providers, separating sensitive and non-sensitive workloads, and sequencing procurement commitments across domains, Gulf states can partially compartmentalize their technological dependencies. This does not eliminate structural constraints, but it allows states to distribute risk across ecosystems and delay full alignment. Therefore, while the AI infrastructure gravitates toward US-governed compute arrangements, this may not lead to a grand foreign policy posture.

Overall, compute's excludability and rivalrous character tie AI infrastructure to external compliance demands, especially from the United States and China. As AI supply chains become more securitized in the global AI race, infrastructure decisions increasingly carry geopolitical consequences, making broad hedging harder to sustain for middle powers and small states. For Gulf states, this means access to high-value compute increasingly comes bundled with security and compliance conditions, even if limited room for maneuver remains through segmentation, firms, and multi-cloud strategies. ♦

References

- Aggarwal, Vinod K., and Andrew W. Reddie. 2025. "New Economic Statecraft and Global Technology Conflicts: the Dilemma for Middle Powers." *Business and Politics* 27(4): 465–483. <https://doi.org/10.1017/bap.2025.10011>
- Bakir, A., and N. Al-Shamari. 2025. "The Art of Hedging: Qatar, Saudi Arabia, and the UAE Manoeuvres amid US–China Great Power Competition." *Third World Quarterly* 46(7): 773–794. <https://doi.org/10.1080/01436597.2025.2509574>
- Buchanan, Ben. 2020. "The AI Triad and What It Means for National Security Strategy." *Center for Security and Emerging Technology*, August. <https://cset.georgetown.edu/publication/the-ai-triad-and-what-it-means-for-national-security-strategy/>
- Chaziza, Mordechai. 2024. "The US–China Rivalry in the Middle East: Confrontation or Competitive Coexistence." *Contemporary Review of the Middle East* 11(2): 231–252. <https://doi.org/10.1177/23477989241240899>
- Chen, Ling S., and Miles M. Evers. 2023. "Wars without Gun Smoke: Global Supply Chains, Power Transitions, and Economic Statecraft." *International Security* 48(2): 164–204. https://doi.org/10.1162/isec_a_00473
- Ding, Jeffrey, and Allan Dafoe. 2021. "The Logic of Strategic Assets: From Oil to AI." *Security Studies* 30(2): 182–212. <https://doi.org/10.1080/09636412.2021.1915583>
- Drezner, Daniel W. 1999. *The Sanctions Paradox: Economic Statecraft and International Relations*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511549366>
- Farrell, Henry, and Abraham L. Newman. 2019. "Weaponized Interdependence: How Global Economic Networks Shape State Coercion." *International Security* 44(1). https://doi.org/10.1162/isec_a_00351

Hameed, Sameena. 2020. "Political Economy of Rentierism in the Middle East and Disruptions from the Digital Space." *Contemporary Review of the Middle East* 7(1): 54–86. <https://doi.org/10.1177/2347798919889782>

Hassan, Oz. 2020. "Artificial Intelligence, Neom and Saudi Arabia's Economic Diversification from Oil and Gas." *The Political Quarterly* 91(1): 222–227. <https://doi.org/10.1111/1467-923X.12794>

Herb, Michael, and Marc Lynch. 2019. "Introduction: The Politics of Rentier States in the Gulf." *POMEPS Studies* 33, January. <https://pomeps.org/introduction-the-politics-of-rentier-states-in-the-gulf>

Kaplan, Jared, et al. 2020. "Scaling Laws for Neural Language Models." *arXiv*:2001.08361. <https://arxiv.org/abs/2001.08361>

Lehdonvirta, Vili, Boxi Wú, and Zoe Hawkins. 2025. "Weaponised interdependence in a bipolar world: how economic forces and security interests shape the global reach of US and Chinese cloud data centres." *Review of International Political Economy* 32(5): 1442–1467. <https://doi.org/10.1080/09692290.2025.2489077>

OECD. 2023. "A blueprint for building national compute capacity for artificial intelligence." *OECD Digital Economy Papers* No. 350. https://www.oecd.org/en/publications/a-blueprint-for-building-national-compute-capacity-for-artificial-intelligence_876367e3-en.html

Pilz, Konstantin, Robi Rahman, James Sanders, and Lennart Heim. 2026. "Data on GPU Clusters." *Epoch AI*. <https://epoch.ai/data/gpu-clusters?view=map&tab=point>

Sastry, Girish, et al. 2024. "Computing Power and the Governance of Artificial Intelligence." *arXiv*:2402.08797. <https://arxiv.org/abs/2402.08797>

Shavit, Yonadav. 2023. "What does it take to catch a Chinchilla? Verifying rules on largescale neural network training via compute monitoring." *arXiv preprint arXiv:2303.11341*. <https://arxiv.org/abs/2303.11341>

Scharre, Paul. 2024. "Future-Proofing Frontier AI Regulation: Projecting Future Compute for Frontier AI Models." *Center for a New American Security*. http://s3.us-east-1.amazonaws.com/files.cnas.org/documents/CNAS-Report_AI-Trends_FinalC.pdf

Soliman, Mohammed. 2025. "Energy, not chips, will determine AI supremacy — and here, it's advantage Gulf." *The Indian Express*, May 21. <https://indianexpress.com/article/opinion/columns/energy-not-chips-will-determine-ai-supremacy-and-here-its-advantage-gulf-10019296/>

Soliman, Mohammed. 2026. "US Authorizes Chips for the UAE, Saudi Arabia." *Middle East Institute Policy Memo*, January 28. <https://mei.edu/policymemo/us-authorizes-chips-for-the-uae-saudi-arabia-2/>

Vaynman, Jane. 2021. "Better monitoring and better spying: The implications of emerging technology for arms control." *Texas National Security Review* 4(4): 33–56. <http://dx.doi.org/10.26153/tsw/17498>

Vaynman, Jane, and Tristan A. Volpe. 2023. "Dual use deception: How technology shapes cooperation in international relations." *International Organization* 77(3): 599–632. <https://doi.org/10.1017/S0020818323000140>

Woods, Dwayne. 2025. "Escaping a weaponized network: China's reaction to the United States escalating technology controls." *Asian Review of Political Economy* 4(1): 1–24. <https://link.springer.com/article/10.1007/s44216-025-00047-7>

How Path Dependence Shapes Gulf Monarchies' High-Tech Strategies

Andrew Leber



Andrew Leber is an assistant professor at Tulane University's Department of Political Science and Middle East & North Africa (MENA) Studies program. His research focuses on the political economy of authoritarianism, particularly in Saudi Arabia and other Arab Gulf monarchies. His work has appeared in Politics & Society, Democratization, Governance, and the Review of Middle East Studies.

In present policy debates, there is considerable focus on states' AI strategies. The latter are framed as immediate policy choices with the potential to shape countries' long-term economic prospects and, in turn, political ones. Beyond the most advanced AI powers (the United States and China), the monarchies of the Gulf Cooperation Council (GCC) have attracted attention due to their financial resources, which have allowed them to purchase pricy AI technologies. Furthermore, their fossil-fuel reserves have permitted them to power energy-intensive data centers on the cheap (Winter-Levy 2024; Soliman 2025). Where strategic discussions necessarily focus on the future, historical institutionalist approaches encourage us to understand how path-dependent Gulf development trajectories shape their present AI strategies. Hence, while the "AI revolution" may be a critical juncture (Collier & Collier 1991) that will lock countries into long-term AI divides, state choices concerning the technology are shaped by pre-existing conditions (Slater & Simmons 2010; Soifer 2012).

If the acceleration of AI technologies comprises a critical juncture for countries to invest in transformative technologies, then the near future will test the agency of Gulf states, especially Saudi Arabia and the United Arab Emirates. The question is whether they

will be able to secure the status of AI middle powers against the constraints of past investments in state-owned enterprises, power grids, and human resources (Winter-Levy & Leicht 2026). A common assumption is that the resource wealth of the Gulf monarchies, coupled with a lack of institutional constraints on their leadership, positions them well in the global AI race (Al Soliman 2024; Arafeh 2025). Kuwaiti analyst Abdullah AlZabin (2024) dubbed this moment the bloc's "petrocompute" opportunity: "building high-performance compute infrastructure of unprecedented scale to serve the world's rapidly growing demand for A.I." Even in this advantaged position, however, there is a range of potential outcomes for the Gulf states (assuming AI truly proves to be a world-changing technology, and the present war on Iran does not undermine the Gulf's case for hosting AI facilities).

The first trajectory is simply "adaptation" or keeping pace with new developments. It would entail incorporating new AI technologies into state bureaucracies, both service-oriented and repressive ones. It would also involve investing in training and education programs to allow citizens to adapt to changing labor markets. The second approach, known as "stakeholder," is one in which the Gulf states become meaningful

financial stakeholders in tech firms primarily based in the United States, Europe, and possibly China. They would diversify an existing portfolio of economic ties with important security patrons and partners. A third possibility, called “subsidizing,” entails foreign firms, most likely from the United States, incentivizing Gulf states to subsidize the construction of key facilities and promise the provision of cheap electricity thereafter. In the process, these states would integrate themselves into U.S. data flows, generate additional value from hydrocarbon extraction in the form of petrochemical investments, and offer cheap access for local startups to AI inferencing or computer calculations that respond to user queries. Finally, the maximal “petrocompute” outcome is one in which Gulf states go beyond subsidizing foreign inferencing to introduce globally competitive proprietary applications of their own and even train their own AI models, which is far more technically demanding.

While GCC states might prefer an “all of the above” approach, even attempting a higher-payoff strategy requires resources beyond the means of the less wealthy states. Oman and Bahrain, though wealthier in per capita terms than much of the MENA region, are each hard-pressed to maintain their citizens’ existing standards of living while managing debt levels. Accordingly, their AI strategies revolve around adaptation: adopting AI enhancements for “e-government” services and training new job seekers to work with such technologies (Albous et al. 2025: 2400). Furthermore, capturing value from AI is not merely about having cash on hand, but rather the ability or capacity to deliver infrastructure for electricity and data centers, as well as to direct projects through local or locally embedded individuals and institutions. Gulf states have different capacities to invest in projects, deliver electricity, and manage large commercial projects – all of which inform these states’ ultimate AI strategies and




Figure 1. Path dependence and divergent outcomes of Gulf AI investments. Modeled after Slater & Simmons (2010: 894).

trajectories. Resource wealth is part of the equation, but not everything.

Kuwait exemplifies this path dependency. While Kuwait is one of the wealthiest per capita GCC states, it has relatively little recent experience with building state-owned enterprises or boosting private-sector entities to compete regionally, let alone internationally (Herb 2014). This means there are few organizational foundations for building a new AI national champion, and limited non-AI data centers, even relative to Bahrain and Oman. In terms of the human capital to staff such a venture, the Kuwaiti education system struggles to meet existing labor market demands (Alhashem & Alhouti 2021). Furthermore, Kuwaiti officials have focused more on stripping away citizenship rather than attracting skilled expatriate workers (Alamer 2025). Despite some AI-focused efforts by some states and state-owned entities, even limited partnerships with U.S. firms are threatened by underinvestment in the country's power grid – making it hard to turn energy resources into the stuff of petrocompute (Crowley 2026). Accordingly, the bulk of Kuwaiti AI investments reflects a stakeholder strategy, focused on U.S. firms operating in the United States. The Kuwait Investment Authority, for example, has joined U.S. firms and an Emirati tech fund as part of an AI infrastructure partnership.

By contrast, the United Arab Emirates clearly has the greatest AI ambitions among the GCC monarchies, with state-owned firm G42 trying to position itself as the equal of hyperscalers like Microsoft or OpenAI. Along with a general experience in building up successful state-owned enterprises, the UAE benefits from an early intervention in the AI space by Tahnoun bin Zayed – a powerful royal with a longstanding and idiosyncratic interest in

advanced technologies. In 2018, Tahnoun founded G42 by incorporating past UAE-based expertise in cybersecurity and digital surveillance and by pushing to establish the new AI-focused Mohammed bin Zayed University for Artificial Intelligence (Hope 2025; Zhou 2025). G42 is the only Gulf AI firm that has secured major inward investments from a U.S. hyperscaler (Microsoft), beyond deals that channel capital outward towards U.S.-based facilities and companies (Allen et al. 2025). As for staffing these ventures, the UAE, and especially Abu Dhabi, has long been a regional hub for skilled talent of all kinds, with consultants who commute to Saudi Arabia to work on various mega-projects there. Furthermore, while the UAE still faces electrical-capacity challenges, past investments in renewables and the Gulf's only nuclear power plants provide a sounder footing for expanding the grid.

Saudi Arabia, in turn, forms an interesting intermediate case. Its trajectories will offer a clearer indication of the relative importance of sheer wealth and ambition, on the one hand, and past institutional and human resource endowments on the other. Saudi Arabia has sought to match the UAE with its own AI national champion called Humain, its own locally based partnerships with U.S. hyperscaler Groq, and its own supply of top-line U.S. microchips, with the Trump administration releasing equal numbers of Nvidia's latest Blackwell chips (Satoriano & Mozur 2025). Despite new industrialization efforts as part of Saudi Arabia's Vision 2030, Humain's domestic capacity to manage AI projects relies heavily on past investments by national oil company Saudi Aramco – a developmental “island of excellence” (Hertog 2010). Electricity generation is likewise a mixed bag, as Saudi Arabia tries to simultaneously increase output and eliminate inefficient oil-fired

power plants while benefiting from a capable, locally based power-generation firm called ACWA Power. The Kingdom also lacks a deep pool of domestic or expatriate talent, though existing educational institutions are being repurposed to support AI research, and the easing of longstanding social restrictions has made cities like Riyadh more livable in the eyes of many expatriates.

This contribution argues that AI strategies in the Gulf states (and other countries) reflect pre-existing policies, capacities, and constraints. Simply adapting to the economic changes wrought by AI is a daunting challenge; going beyond this to capture value from AI supply chains will depend on Gulf states' past investments in basic infrastructure (electricity provision) and education (citizen labor supply), as well as policies towards migration and business regulation. Costly AI strategies that do not reckon with these constraints are more likely to enrich Western firms and consultants than to diversify Gulf economies. In all, despite the transformative potential of the AI revolution for both economic outcomes and research practices, political scientists should still draw on longstanding historical-institutionalist tools and concepts to explain divergent AI strategies and trajectories across the region. ♦

References

- Alamer, Sultan. 2025 "La-naj'al al-Kuwait naqiyya mujadidan'... jadhur wa athar al-jinsiyya fi al-Kuwait ['Let's make Kuwait pure again'... The roots and effects of citizenship revocations in Kuwait]." *Alpheratz Magazine*, August 27. <https://alpheratzmag.com/reports/>.
- Albous, Mohammad, Odeh Rashed Al-Jayyousi, and Melodena Stephens. 2025. "AI governance in the GCC states: a comparative analysis of national AI strategies." *Journal of Artificial Intelligence Research* 82: 2389-2422. <https://doi.org/10.1613/jair.1.17619>
- Alhashem, Fatimah, and Ibrahim Alhouti. 2021. "Endless education reform: The case of Kuwait." In *Annual Review of Comparative and International Education* 2020, ed. Alexander N. Wiseman, 345-367. Leeds: Emerald Publishing Limited. https://www.researchgate.net/publication/353644946_Endless_Education_Reform_The_Case_of_Kuwait
- Allen, Gregory C., Georgia Adamson, Lennart Heim, and Sam Winter-Levy. 2025. *The United Arab Emirates' AI Ambitions*. Washington, D.C.: Center for Strategic and International Studies. <https://www.csis.org/analysis/united-arab-emirates-ai-ambitions>
- Alzabin, Abdullah. 2024. "PetroCompute: Will A.I.'s Future Run Through the Gulf?" Alzabin's Substack [blog], December 17. <https://alzabin.substack.com/p/petrocompute-will-ais-future-run>.
- Arafeh, Nur. 2025. "The Promise and Risks of Artificial Intelligence." *Carnegie Middle East Center - Diwan* [blog], September 12. <https://carnegieendowment.org/middle-east/diwan/2025/09/the-promise-and-risks-of-artificial-intelligence>.
- Collier, Ruth Berins & David Collier. 1991. *Shaping the political arena: Critical Junc-tures, the Labor Movement, and Regime Dynamics in Latin America*. Princeton, NJ: Princeton University Press.

Crowley, John. 2026 “Kuwait Looks to the Cloud as Power Grid Feels the Strain.” *Arabian Gulf Business Insight*, January 23. <https://www.agbi.com/infrastructure/2026/01/kuwait-looks-to-the-cloud-as-power-grid-feels-the-strain/>.

Herb, Michael. 2014. *The Wages of Oil: Parliaments and Economic Development in Kuwait and the UAE*. Ithaca: Cornell University Press.

Hope, Bradley. 2025. “A Spymaster Sheikh Controls a \$1.5 Trillion Fortune. He Wants to Use It to Dominate AI.” *Wired*, January 14. <https://www.wired.com/story/uae-intelligence-chief-ai-money/>.

Internations. 2025. “The Best & Worst Countries for Expats in 2025.” <https://www.internations.org/expat-insider/2025/best-and-worst-countries-2025>.

Satoriano, Adam & Paul Mozur. 2025. “Saudi Arabia’s New Power Play Is Exporting A.I. to the World.” *The New York Times*, October 27. <https://www.nytimes.com/2025/10/27/technology/saudi-arabia-ai-exporter.html>.

Slater, Dan and Erica Simmons. 2010. “Informative Regress: Critical Antecedents in Comparative Politics.” *Comparative Political Studies* 43 (7): 886-917. <https://doi.org/10.1177/0010414010361343>

Soifer, Hillel David. 2012. “The Causal Logic of Critical Junctures.” *Comparative Political Studies*, 45 (12): 1572-1597. <https://doi.org/10.1177/0010414012463902>

Soliman, Mohammed. 2025. *From Crude to Compute: Building the GCC AI Stack*. Washington, D.C.: Middle East Institute.

Winter-Levy, Sam & Anton Leicht. 2026. “The AI Divide.” *Foreign Affairs*, February 20. <https://www.foreignaffairs.com/unit-ed-states/ai-divide>.

Zhou, Viola. “The UAE’s AI strategy winds through Silicon Valley.” *Rest of World*, November 24. <https://restofworld.org/2025/what-the-uae-is-doing-in-silicon-valley/>.

Governing AI Surveillance through Legal and Social Knowledge: An Islamic and Anthropological Framework

Abdullah Omran



Abdullah Omran is a PhD candidate in Middle Eastern Languages and Cultures at Indiana University Bloomington. He is a digital ethnographer whose work explores the intersection between technological advancement and Islamic legal studies.

Artificial Intelligence (AI) now plays a central role in military and security systems across the Arab world. Governments use AI for predictive surveillance, signal analysis (such as Egypt's use of facial recognition systems at public gatherings and the UAE's deployment of predictive policing algorithms in urban centers), automated screening, and counterterrorism intelligence, often through security agencies rather than civilian institutions (Sakr & Hakme 2025). Studies of AI militarization also show how states integrate these tools into surveillance and weapon systems that have certainly reshaped domestic control and regional power relations (Sarkin & Sotoudehfar 2024).

When automated systems decide who gets monitored, detained, or targeted, they do not always do so accurately. They raise the practical question about whether probabilistic suspicion can ever justify state coercion. This question shapes whether AI-driven security systems distinguish between a poet and a plotter, a dissenter and a militant, or a private grievance and a public threat. Islamic legal tradition offers a structured response to precisely these questions. It provides categorical distinctions that map different forms of polit-

ical expression and wrongdoing to different thresholds of permissible state intrusion. In Islamic law, not all political acts are treated the same. A person offering sincere advice to a ruler is not placed in the same category as someone inciting rebellion; each is judged differently, with different rights and limits on what the state can do. That means the state must ask what kind of act it is, who is doing it, and whether it crosses a moral or legal line. Privacy, according to Islamic law, is a binding limit on what the state may know (Almutairi 2020; Al-Shamari 2024; Hayat 2007). And it provides a legal taxonomy that tells us who is entitled to what kind of privacy protection. The challenge is how to encode these limits into the machines themselves. This piece argues that anthropology functions as a method for making uncertainty visible within the system and for ensuring that contested meaning slows down decision-making rather than speeding it up. Put differently, ethnographic specificity offers discernible, though not always clear cut, identificatory markers of social groups and, more importantly, of their practices. As such, it lends legal categories in intra Muslim conflict the legibility necessary for AI assisted surveillance systems to

distinguish behavior in ways consistent with appropriate actionable intelligence.

One might reasonably ask why authoritarian and security-oriented states would accept privacy and cultural-differentiation constraints on AI-driven surveillance and targeting. The answer is not that such states suddenly value privacy, but that AI changes how security decisions are made and where responsibility sits. Once states rely on AI systems that sort people, speech, or locations into risk categories before human review, errors and intrusions are no longer isolated incidents. Rather, they become routine system outputs. Research on AI assisted targeting shows that these systems tend to narrow judgment, encourage reliance on automated recommendations, and blur accountability. This creates real risks of escalation, legal exposure, and loss of control for security agencies themselves (Dorsey 2026; Blanchard and Bruun 2025). In addition, contextual and cultural constraints, such as linguistic nuance, religious context, and local political discourse, can improve performance.

Work on Arabic Natural Language Processing (NLP) shows that systems trained on simplified threat categories regularly misread dialectal, indirect, or context-dependent political speech, leading to false positives that expand surveillance beyond actionable risk (Haj Ahmed et al. 2024; Charfi et al. 2024). When AI outputs feed directly into security workflows, these errors are costly. Treating interpretation as part of classification, rather than noise, reduces this problem.

The stakes become clear when AI systems collapse legally distinct categories into a single risk continuum. Contemporary targeting and surveillance models treat activists, civilians, and militants as variations on the same

probabilistic object, assigning scores that justify intervention without requiring categorical legal proof (Gordon et al. 2023; McKendrick 2019). In Egypt, counterterrorism surveillance applies the same profiling logic to peaceful political activity as it does to militant networks (Human Rights Watch 2024). In Palestine, Israeli AI systems, such as Lavender and Gospel, generate large suspect lists through pattern analysis, merging civilians and combatants under conditions of limited human review (Amnesty International 2022). These cases illustrate a structural failure of AI mediated governance in the sense that decision systems act on calculated risk rather than a legally meaningful distinction.

While privacy concerns about AI surveillance arise across national contexts, Islamic jurisprudence offers a particularly structured framework for addressing them, but it places a greater burden on the state to avoid arbitrary automated suspicion detection. Islamic jurisprudence distinguishes three legal categories relevant to internal security. *Bughāt* are principled dissenters who oppose the state on the basis of plausible legal interpretation (ta'wīl). Their private deliberations and political organizing are protected, and surveillance is prohibited absent active hostilities. *Muhāribūn* are agents of indiscriminate harm who destabilize public order. Islamic law permits limited intervention against them but requires visible wrongdoing rather than probabilistic inference. *Mujrimūn* are ordinary offenders whose wrongdoing is personal. They enjoy the strongest privacy protections, as private vice and association do not justify intrusion (Al-Dawoody 2011).

Read as a governance framework, these categories specify how privacy should structure the design and use of AI surveillance. In the case of *bughāt*, design treats privacy as a

hard limit: political dissent grounded in *ta'wīl* does not justify invasive data collection, and AI systems' access should be restricted to publicly available information unless conduct clearly shifts into open hostilities. For *muhāribūn*, governance allows narrower intrusions but only after wrongdoing becomes visible and public. Predictive assessments, which estimate future risk from statistical patterns rather than proven wrongdoing, may inform attention but cannot serve as the legal basis for breaching privacy or escalating intervention. For *mujrimūn*, AI profiling is prohibited without concrete evidence of an offense. The difficulty is that AI systems do not recognize these legal thresholds on their own. They classify speech, relationships, and behavior as features rather than as legally situated acts.

The mismatch between legal categories and machine classification motivates a focus on anthropology, which examines how contextual interpretation can help enforce privacy boundaries within AI-based governance. In this framework, anthropological insight contributes to AI governance by shaping how systems are trained to interpret meaning under contextual uncertainty, instead of by expanding the scope of data they ingest. Rather than using personal data, this approach relies on short, anonymous examples of common speech, like sermons or protest chants. These examples help AI understand language and context without monitoring real people. Prior work on aligning AI training to cultural values shows that scenario-based materials are especially effective in post-training (AlKhamissi et al. 2025). They guide how systems reason about language without entangling interpretation with personal traces or real-world monitoring.

Within this framework, annotating the output data is limited to recording multiple plausible interpretations for the said scenarios rather than producing a single resolved label. The disagreement, rooted in the availability of multiple interpretations, is often constitutive of meaning rather than a sign of error, particularly in political, religious, or culturally embedded speech (Geertz 1973). Technical work on disagreement-aware learning supports this view by showing that models trained on unaggregated or plural annotations are better able to recognize boundary cases and uncertainty, rather than clustering contested meanings into confident predictions (Rodríguez Barroso et al. 2024).

Operationally, this means that ambiguity causes the system to pause or defer action rather than escalate automatically. The governance effect comes from how the system is trained to behave. Using de identified scenarios and retaining interpretive disagreement teaches the system to reason about context without following individuals and to acknowledge uncertainty rather than hide it. Privacy is enforced through these design choices, which shape how meaning is interpreted and how action is limited, rather than through later restrictions on data use.♦

References

- Al Dawoody, Ahmed. 2011. *The Islamic Law of War: Justifications and Regulations*. New York: Palgrave Macmillan.
- AlKhamissi, Mai, Yunze Xiao, Badr AlKhamissi, and Mona Diab. 2025. "Hire Your Anthropologist! Rethinking Culture Benchmarks Through an Anthropological Lens." *arXiv* preprint, arXiv:2510.05931.

Almutairi, S. E. 2020. "The Shari'a Approach to Contemporary Problems of Mass Surveillance." *Muslim World Journal of Human Rights* 17 (1): 1–24. <https://www.degruyterbrill.com/document/doi/10.1515/mwjhr-2020-0007/html>.

Al Shamari, K. S. 2024. "Exclusion of Evidence and Its Impact on Criminal Justice: A Comparative Study." *Arab Law Quarterly*. <https://doi.org/10.1163/15730255-bja10143>.

Amnesty International. 2022. *Automated Apartheid: How Facial Recognition Fragments, Segregates and Controls Palestinians in the Occupied Territories*. London: Amnesty International.

Ashraf, Y., Y. Wang, B. Gu, P. Nakov, and T. Baldwin. 2024. "Arabic Dataset for LLM Safe-guard Evaluation." *arXiv* preprint. <https://doi.org/10.48550/arxiv.2410.17040>.

Blanchard, Alexander, and Laura Bruun. 2025. *Autonomous Weapon Systems and AI Enabled Decision Support Systems in Military Targeting: A Comparison and Recommended Policy Responses*. Stockholm: Stockholm International Peace Research Institute. <https://doi.org/10.55163/YQBY3151>

Charfi, Anis, Mabrouka Besghaier, Raghda Akasheh, Andria Atalla, and Wajdi Zaghouani. 2024. "Hate Speech Detection with ADHAR: A Multi Dialectal Hate Speech Corpus in Arabic." *Frontiers in Artificial Intelligence* 7. <https://doi.org/10.3389/frai.2024.1391472>.

Dorsey, Jessica. 2026. "The Erosion of Human(e) Judgement in Targeting? Quantification Logics, AI Enabled Decision Support Systems and Proportionality Assessments in International Humanitarian Law." *International Review of the Red Cross*, FirstView: 1–31. <https://doi.org/10.1017/S1816383125100969>.

Geertz, Clifford. 1973. *The Interpretation of Cultures: Selected Essays*. New York: Basic Books.

Gordon, Neve, Nicola Perugini, and Yotam Feldman. 2023. "Humanitarian Violence and the Rise of Algorithmic Targeting." *Journal of Genocide Research* 25 (4): 1–19.

Haj Ahmed, Ahmed, Rui Jie Yew, Xerxes Minocher, and Suresh Venkatasubramanian. 2024. "Navigating Dialectal Bias and Ethical Complexities in Levantine Arabic Hate Speech Detection." *arXiv* preprint, [arXiv:2412.10991](https://arxiv.org/abs/2412.10991).

Hayat, M. A. 2007. "Privacy and Islam: From the Qur'an to Data Protection in Pakistan." *Information & Communications Technology Law* 16 (2): 137–148. <https://doi.org/10.1080/13600830701532043>.

Human Rights Watch. 2024. *'Security Forces and the Algorithm': Digital Surveillance and Political Repression in Egypt*. New York: Human Rights Watch.

McKendrick, David. 2019. *Risk, Probability, and Security Decision Making*. London: Routledge.

Rodríguez Barroso, Nuria, Eugenio Martínez Cámara, Jose Camacho Collados, M. Victoria Luzón, and Francisco Herrera. 2024. "Federated Learning for Exploiting Annotators' Disagreements in Natural Language Processing." *Transactions of the Association for Computational Linguistics* 12: 630–648. https://doi.org/10.1162/tacl_a_00664.

Sakr, G. E., and O. Hakme. 2025. "The Impact of Artificial Intelligence on Military and Security in the MENA." In *AI in the Middle East for Growth and Business*, 363–383. Cham: Springer.

Sarkin, Jeremy Julian, and Saba Sotoudehfar. 2024. "Artificial Intelligence and Arms Races in the Middle East: The Evolution of Technology and Its Implications for Regional and International Security." *Defence and Security Analysis* 40 (1): 97–119. <https://doi.org/10.1080/14751798.2024.2302699>.

AI and Digital Authoritarianism in MENA

Alexandra Siegel



Alexandra Siegel is an Associate Professor of Political Science at the University of Colorado Boulder, a faculty affiliate of NYU's Center for Social Media and Politics, and a faculty affiliate of Stanford's Immigration Policy Lab. She received her PhD in Political Science from NYU in 2018. Her research uses social media data, network analysis, and experiments—in addition to more traditional data sources—to study mass and elite political behavior in the Arab World and other comparative contexts. She is a former non-resident fellow at the Brookings Institute, Junior Fellow at the Carnegie Endowment for International Peace, and a former CASA Fellow at the American University in Cairo. She holds a Bachelors in International Relations and Arabic from Tufts University.

Over the past two decades, regimes across the MENA region have invested in technologies that have made surveillance cheaper, censorship more granular, and propaganda more scalable (Freedom House 2023). Rapid advances in AI technologies have accelerated the speed, scale, and integration of these capabilities (Feldstein 2023). Governments increasingly combine biometric identification, “smart city” sensor networks, social media platform monitoring, and influence operations into mutually reinforcing systems. Regimes are adopting AI as a strategic infrastructure that reshapes how threats are anticipated, how publics are managed, and how information is controlled (Beidollahkhani 2025).

The same technologies that strengthen state control, however, can also strengthen civic resilience. Activists, journalists, and civil society organizations can also use AI for translation, documentation, verification, and rapid counter-messaging (Espiritusanto et al. 2024). Below, I describe how AI technologies are bolstering the digital authoritarian toolkit in the MENA region and then highlight how they simultaneously offer new opportunities for opposition and low-resource actors to achieve their goals.

Surveillance

AI technologies have facilitated a shift from surveillance of known dissidents toward population-scale monitoring by automating identification, combining disparate data sources, and turning infrastructure into searchable, real-time governance tools (Kalluri 2025; Brookings 2021; Reuters 2023; Fontes 2022). When citizens believe the state can recognize them quickly and cheaply, self-censorship and political caution diffuse far beyond those who are directly repressed (Zhuravskaya, Petrova, and Enikolopov 2020).

For example, Iran has leveraged AI-enhanced technologies for hijab enforcement. These technologies have enabled the regime to more systematically use facial recognition and road-camera networks for detecting and identifying alleged hijab violations, alongside a citizen-reporting app (Nazer) that can be used to trigger police action (Associated Press 2025; Council on Foreign Relations 2025). This combination of networked cameras and biometric identification reduces the marginal cost of sanctioning ordinary citizens—especially when paired with automatable administrative penalties such as fines and vehicle seizures, rather than formal

prosecutions. Computer-vision systems can flag potential violations from drone or CCTV footage, facial recognition can attach a name or identity token, and app-based reports can route incidents into an automated enforcement pipeline.

Similarly, Egypt's New Administrative Capital, a large, centrally planned city built east of Cairo, shows how "smart city" governance can embed surveillance into the built environment. The city includes more than 6,000 CCTV cameras and a centralized command setup. Although it has been marketed as advancing security and efficiency, digital rights advocates have warned that such infrastructure can facilitate monitoring and suppression of dissent in an authoritarian context (Reuters 2023). In an AI-enhanced smart city, cameras become queryable sensors. Computer vision and centralized analytics make it possible to search across time and space at a previously impossible speed and scale. Once sensor networks and centralized analytics are in place, the marginal cost of repurposing them for political surveillance is low. Smart city framing also offers plausible deniability, allowing regimes to present pervasive monitoring as a neutral modernization project.

Regimes have also increasingly integrated AI analytics into policing through "smart policing" platforms. The Abu Dhabi Police have partnered with Presight, an Abu Dhabi-based AI and big data analytics company, to integrate an "AI-Policing Suite" into police operations (Tech & Justice Project n.d.). Presight's marketing materials highlight geospatial insights and "predictive analytics" to support proactive measures (Presight n.d.). Rather than responding to crimes, predictive policing tools can be used to preempt gatherings, disrupt networks, or prevent "risk," particularly in states where broad security laws and

speech restrictions define dissent as a threat to public order. These predictive and adaptive AI systems can turn surveillance into anticipatory governance, shaping political behavior before overt contestation occurs (Beidollahkhani 2025).

This population-scale monitoring is frequently complemented by high-resolution targeting of key individuals through spyware and device compromise. For example, Gulf regimes increasingly use a combination of spyware, facial recognition, and mass surveillance tools (European Council on Foreign Relations 2022). Pegasus spyware, which is produced and exported by Israel, has been used to hack multiple Bahraini and Saudi activists (Citizen Lab 2021). Pegasus has additionally been used to hack journalists at Al Jazeera and Al Araby TV via iMessage (Citizen Lab 2020). In North Africa, Amnesty International reported repeated Pegasus targeting of Moroccan journalist Omar Radi (Amnesty International 2020). Such spyware captures communications, contacts, and location histories that can be triaged and operationalized via automated analytics and network mapping.

Censorship

Censorship in MENA has historically been quite blunt—blocking platforms or shutting down networks. But AI tools have increasingly allowed regimes to use selective censorship that targets specific sites, services, or content types while preserving enough connectivity for commerce and state capacity. AI-enhanced censorship typically operates through ISP-level infrastructure, including deep packet inspection (the examination of internet traffic in transit); filtering (the restriction of access to specific content); and automated classification (the use of software to identify

and sort traffic or content), all of which enable selective blocking. Humans still set the censorship rules, but AI helps detect patterns in internet traffic and classify content or services for enforcement. Meanwhile, automation enables those blocking decisions to be carried out quickly and at scale.

Iran is often cited as a leading example of selective, “smart” censorship. The regime has long pursued “intelligent filtering” that can restrict specific sites, pages, features, or content categories while preserving enough connectivity for commerce and routine governance (Reuters 2024). AI tools have been used for classification and traffic identification, which makes it possible to automatically sort and act on traffic or content types at scale (ARTICLE 19 2026). This infrastructure makes censorship more flexible as restrictions can be tightened selectively during moments of contention without imposing the high economic and administrative costs of a full internet blackout (ARTICLE 19 2026). Notably, however, the regime has repeatedly resorted to Internet shutdowns during high-threat unrest—suggesting that smart filtering is perhaps used for routine governance while shutdowns remain the regime’s crisis tool (Reuters 2024; Access Now 2023; NetBlocks 2022; NetBlocks 2026).

In the Gulf, AI-enabled censorship often operates through platform-wide monitoring and automated content moderation on social media. Security services can use AI to scan large-scale social media data, flagging politically sensitive content, mapping networks, and prioritizing accounts for reporting, takedown requests, or offline questioning (Freedom House 2019; Freedom House 2023). In Bahrain, for example, authorities reportedly forced a user to remove an Instagram post about protests against normalization with

Israel and have also exploited platform reporting mechanisms to target activists and independent journalists. In Saudi Arabia, authorities have requested that platforms such as Netflix and YouTube remove content deemed “inappropriate,” while in the UAE, pervasive surveillance of online activists and journalists has contributed to extensive self-censorship. At the same time, platforms’ own AI-driven moderation tools can amplify state pressure. Because enforcement is opaque and frequently automated, content removal and account restrictions may occur preemptively around well-known “red lines,” producing censorship effects even when formal requests are not visible (ECDHR 2026; MEI 2021). Importantly, however, Gulf regimes do not control the dominant social media platforms themselves. Because major services are largely U.S.-based, governments in MENA typically lack direct, state-run tools for platform-level censorship and instead rely more on monitoring, pressure, and takedown requests routed through platform governance systems.

Propaganda

AI tools make propaganda and influence campaigns easier to run because a small team can produce and translate large amounts of content quickly and cheaply. They also make narratives more convincing by enabling high-volume, stylistically diverse, rapidly iterated messaging (Marcellino et al. 2023). Social media platform disclosures offer documentation of organized manipulation tied to regional actors. Meta has reported removing coordinated inauthentic behavior networks originating in the UAE and Egypt, as well as a separate operation in Saudi Arabia involving clusters of accounts and pages designed to mislead audiences (Meta 2019). Generative AI accelerates these operations by lowering

the production and iteration costs of persuasive content. Generative AI reduces time, money, and skill barriers for producing disinformation and propaganda, potentially overwhelming detection and moderation systems (Feldstein 2023). For example, Google's Threat Analysis Group reports that it has terminated Iran-linked influence operations that pushed automated pro-regime Arabic-language narratives (Google Threat Analysis Group 2024).

Gulf states are increasingly investing in “sovereign” generative AI by developing or backing domestic Arabic large language models (LLMs). The UAE's government-linked Technology Innovation Institute developed the open-weight Falcon series (Almazrouei et al. 2023), and the Mohamed bin Zayed University of Artificial Intelligence released Jais, an Arabic-centric foundation model (Sengupta et al. 2023). The Saudi Data and AI Authority (SDAIA) has developed ALLaM, an Arabic-English LLM (Bari et al. 2024). These projects may allow states to influence what content LLMs will generate, refuse, or systematically omit on sensitive topics. As everyday citizens increasingly turn to LLMs for information, the ability to manipulate the content they produce offers regimes a new lever of information control.

Civic Resilience

Many of the same features of AI technologies that enable digital authoritarianism—automation, translation, classification, and summarization—can also strengthen civic resilience by making documentation, verification, coordination, and public outreach faster for low-resource actors (Espiritusanto, Nachawati-Rego, and Magallón-Rosa 2024; Kerley, Miller, and Campagnucci 2024). For example, before the fall of the Assad regime, the Syrian

Archive and Mnemonic ecosystem use automated collection and structured verification workflows to preserve open-source evidence for advocacy and accountability, even when platform moderation or deletions remove content (Syrian Archive n.d.; Mnemonic 2020; Exposing the Invisible 2021).

AI can also support circumvention of censorship and continuity of Internet access by helping communities quickly diagnose blocks, share workable configurations in local language, and adapt when censors change tactics. During Iran's 2022 protests, some people were still able to communicate and access news despite state blocking by using digital circumvention tools like Psiphon and Tor Snowflake, which allow users to access blocked online content by routing traffic through external networks. Even when these workarounds only worked for some users or for limited periods, they show a basic form of resilience under digital authoritarianism. Activists and ordinary citizens tested what still worked, adapted quickly when blocks changed, and shared step-by-step guidance so others could stay connected (OONI 2022; OONI n.d.).

AI can also strengthen a different kind of resilience by helping civil society respond faster to state-aligned propaganda and information control. For example, in Turkey, protesters have leveraged AI to maintain anonymity and create viral anti-regime content (Gümrükçü 2025). The Arab Fact-Checkers Network's partnership with Full Fact (a UK-based independent fact-checking non-profit) aims to build AI tools that speed up verification in Arabic, and JournalismAI supported an AI-powered WhatsApp bot designed to deliver quicker fact-checks for MENA audiences. These tools help opposition contest official narratives and reduce the effectiveness of

disinformation by making rapid, credible corrections easier to produce and distribute (AFCN 2024; JournalismAI 2024).

As AI tools continue to advance, digital politics in MENA—and around the globe—will likely remain a race of adaptation. Governments will keep integrating AI into surveillance, censorship, and influence operations to monitor the public and manage information more efficiently, while activists, journalists, and civil society will use many of the same capabilities to document abuse, maintain access, verify claims, and coordinate responses. ♦

References

- Almazrouei, Ebtesam, et al. 2023. “The Falcon Series of Open Language Models.” *arXiv* (2311.16867). <https://arxiv.org/abs/2311.16867>
- Amnesty International. 2020. “NSO Group spyware used to target Moroccan journalist Omar Radi.” June 22. <https://www.amnesty.org/en/latest/news/2020/06/nso-spyware-used-against-moroccan-journalist/>
- Arab Fact-Checkers Network (AFCN). 2024. “AFCN and Full Fact Unite for AI-Powered Fact-Checking in Arabic.” April 2. <https://arabfcn.net/en/blog/2024/04/02/afcnc-and-full-fact-unite-for-ai-powered-fact-checking-in-arabic/>
- ARTICLE 19. 2026. *Tightening the Net: China’s Infrastructure of Oppression in Iran*. London: ARTICLE 19.
- Aspen, Eliza and Spandana Singh. 2021. “Content Moderation Trends in the MENA Region: Censorship, Discrimination by Design, and Linguistic Challenges.” *Middle East Institute* (MEI). August 25. <https://mei.edu/publication/content-moderation-trends-mena-region-censorship-discrimination-design-and-linguistic/>
- Bari, M. Saiful, et al. 2024. “ALLaM: Large Language Models for Arabic and English.” *arXiv* (2407.15390). <https://arxiv.org/abs/2407.15390>
- Beidollahkhani, Arash. 2025. “From Predicting Dissent to Programming Power: Analyzing AI-Driven Authoritarian Governance in the Middle East through TRIAD Framework.” *Democratization*, 1-24. <https://www.tandfonline.com/doi/full/10.1080/13510347.2025.2576527>
- ECDHR. 2026. “Platform Compliance in the Middle East: How Opaque Content Moderation Enables Quiet Authoritarian Censorship.” January 21.
- Espiritusanto, Óscar, Leila Nachawati-Rego, and Raúl Magallón-Rosa. 2024. “The Role of AI in Citizen Journalism, Human Rights Activism, and Monitoring: Limits and Possibilities.” In *Journalism, Digital Media and the Fourth Industrial Revolution*, eds. José Sixto-García, Alberto Quian, Ana-Isabel Rodríguez-Vázquez, Alba Silva-Rodríguez, Xosé Soengas-Pérez, 211-226. Springer. https://doi.org/10.1007/978-3-031-63153-5_16

Evdokimov, Leonid, Maria Xynou, Mohammad El-Taher, Hassan Al-Azhary, and Sarah Mohsen. 2018. *The State of Internet Censorship in Egypt*. OONI + AFTE. July 2. <https://ooni.org/post/egypt-internet-censorship/>

Feldstein, Steven. 2023. "The Consequences of Generative AI for Democracy, Governance and War." *Survival* 65(5): 117–142. <https://www.tandfonline.com/doi/abs/10.1080/00396338.2023.2261260>

Fontes, Catarina, Ellen Hohma, Caitlin C. Corrigan, and Christoph Lütge. 2022. "AI-powered public surveillance systems: why we (might) need them and how we want them." *Technology in Society* 71. <https://www.sciencedirect.com/science/article/pii/S0160791X22002780>

Funk, Allie, Adrian Shahbaz, and Kian Vesteinsson. 2023. *Freedom on the Net 2023: The Repressive Power of Artificial Intelligence*. Washington, DC: Freedom House. <https://freedomhouse.org/report/freedom-net/2023/repressive-power-artificial-intelligence>

Gleicher, Nathaniel. 2019. "Removing Coordinated Inauthentic Behavior in UAE, Egypt and Saudi Arabia." Meta. August 1. <https://about.fb.com/news/2019/08/cib-uae-egypt-saudi-arabia/>

Government Hacks Activists with NSO Group Zero-Click iPhone Exploits." *Citizen Lab*. August 24. <https://citizenlab.ca/research/bahrain-hacks-activists-with-nso-group-zero-click-iphone-exploits/>

Gümrükçü, Selin Bengi. 2025. "Artificial Intelligence–Based Aesthetics of Dissent in Turkey." *Public Seminar*, June 18. <https://publicseminar.org/2025/06/turkey-protesters-using-ai/>

Huntley, Shane. 2024. "TAG Bulletin: Q4 2023." *Google Threat Analysis Group*. January 19. <https://blog.google/threat-analysis-group/tag-bulletin-q4-2023/>

James, Noel. 2025. "Women This Week: Iran Using Electronic Surveillance to Enforce Veiling Laws." *Council on Foreign Relations*. March 24. <https://www.cfr.org/articles/women-week-iran-using-electronic-surveillance-enforce-veiling-laws>

JournalismAI. 2024. "Verify Media Platform." Innovation Challenge 2024. <https://www.journalismai.info/programmes/innovation/innovation-challenge-2024/verify-media-platform>

Kalluri, Pratyusha Ria, William Agnew, and Abeba Birhane. 2025. "Computer-vision research powers surveillance technology." *Nature* 643, 73–79. <https://www.nature.com/articles/s41586-025-08972-6>

Keaten, Jamey and Jon Gambrell. 2025. "UN Report Warns Iran Is Stepping Up Electronic Surveillance of Women to Enforce Headscarf Laws." *Associated Press*. March 14. <https://apnews.com/article/e6d33c4f45f0181b9b-c4669aef196209>

Kerley, Beth, Carl Miller, and Fernanda Campagnucci. 2024. *Leveraging AI for Democracy: AI for Civil Society*. National Endowment for Democracy. https://www.ned.org/wp-content/uploads/2024/10/NED_Leveraging-AI-for-Democracy-Report_AI-for-Civil-Society.pdf

Lynch, James. 2022. "Iron Net: Digital Repression in the Middle East and North Africa." *European Council on Foreign Relations*. June 29. <https://ecfr.eu/publication/iron-net-digital-repression-in-the-middle-east-and-north-africa/>

Marcellino, William, Nathan Beauchamp-Mustafaga, Amanda Kerrigan, Lev Navarre Chao, and Jackson Smith. 2023. "The Rise of Generative AI and the Coming Era of Social Media Manipulation 3.0: Next-Generation Chinese Astroturfing and Coping with Ubiquitous AI." *RAND*. September 7. <https://www.rand.org/pubs/perspectives/PEA2679-1.html>

Marczak, Bill, John Scott-Railton, Noura Aljizawi, Siena Anstis, and Ron Deibert. 2020. "The Great iPwn: Journalists Hacked with Suspected NSO Group iMessage 'Zero-Click' Exploit." *Citizen Lab*. December 20. <https://citizenlab.ca/research/the-great-ipwn-journalists-hacked-with-suspected-nso-group-imessage-zero-click-exploit/>

Marczak, Bill, Ali Abdulemam, Noura Aljizawi, Siena Anstis, Kristin Berdan, John Scott-Railton, Ron Deibert, and Bahr Abdul Razzak. 2021. "From Pearl to Pegasus: Bahraini Government Hacks Activist Group with NSO Group Zero-Click iPhone Exploits." *Citizen Lab*. August 24. <https://citizenlab.ca/research/bahrain-hacks-activists-with-nso-group-zero-click-iphone-exploits/>

McBrien, Tyler. 2021. "The Archive as Activism: Lessons from Mnemonic and the Syrian Archive." *Exposing the Invisible*. August 10. <https://exposingtheinvisible.org/en/articles/archive-activism-mnemonic/>

Mnemonic. 2020. "Methods." March 10. <https://mnemonic.org/en/about/methods/>

NetBlocks. 2022. "Internet disrupted in Iran amid protests over death of Mahsa Amini." September 19. <https://netblocks.org/reports/internet-disrupted-in-iran-amid-protests-over-death-of-mahsa-amini-X8qVEwAD>

Open Observatory of Network Interference (OONI). 2022. "Technical multi-stakeholder report on Internet shutdowns: The case of Iran amid Autumn 2022 protests." November 29. <https://ooni.org/post/2022-iran-technical-multistakeholder-report/>

Open Observatory of Network Interference. n.d. "OONI Explorer: Egypt." <https://explorer.ooni.org/country/EG>

Open Observatory of Network Interference. n.d. "Reachability of Censorship Circumvention Tools." <https://explorer.ooni.org/circumvention>

OpenAI. 2024. "Influence and Cyber Operations: An Update." October. https://cdn.openai.com/threat-intelligence-reports/influence-and-cyber-operations-an-update_October-2024.pdf

Peterson, Dahlia. 2021. "How China harnesses data fusion to make sense of surveillance data." *Brookings Institution*. September 23. <https://www.brookings.edu/articles/how-china-harnesses-data-fusion-to-make-sense-of-surveillance-data/>

Presight. n.d. "Presight AI-Policing Suite." <https://www.presight.ai/news/presight-launches-groundbreaking-ai-policing-suite-at-world-police-summit-2025-to-transform-law-enforcement-operations>

Rosson, Zach, Felicia Anthonio, and Carolyn Tackett. 2023. "Weapons of Control, Shields of Impunity, Internet Shutdowns in 2022: The #KeepItOn Report." *Access Now*. February 28. <https://www.accessnow.org/internet-shutdowns-2022/>

Reuters. 2023. "CCTV Cameras Will Watch Over Egyptians in New High-Tech Capital." January 4. <https://www.reuters.com/world/cctv-cameras-will-watch-over-egyptians-new-high-tech-capital--trfn-2023-01-04/>

Reuters. 2024. "Iran's Supreme Leader Calls for Regulation of Cyberspace." August 27. <https://www.reuters.com/world/middle-east/irans-supreme-leader-calls-regulation-cyber-space-2024-08-27/>

Reuters. 2026. "Digital blackout hits Tehran, other parts of Iran, NetBlocks says." January 8. <https://www.reuters.com/business/media-telecom/digital-blackout-hits-tehran-other-parts-iran-netblocks-says-2026-01-08/>

Sengupta, Neha, et al. 2023. "Jais and Jais-chat: Arabic-Centric Foundation and Instruction-Tuned Open Generative Large Language Models." *arXiv* (2308.16149). <https://arxiv.org/abs/2308.16149>

Shahbaz, Adrian and Allie Funk. 2019. *Freedom on the Net 2019: The Crisis of Social Media*. Washington, DC: Freedom House. <https://freedomhouse.org/report/freedom-net/2019/crisis-social-media>

Syrian Archive. n.d. "Methods and Tools." <https://syrianarchive.org/en/about/methods-and-tools/>

Tech & Justice Project (Oxford). n.d. "United Arab Emirates." *Oxford Institute of Technology and Justice*. <https://www.techandjustice.bsg.ox.ac.uk/research/united-arab-emirates>

Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov. 2020. "Political Effects of the Internet and Social Media." *Annual Review of Economics* 12: 415–438. <https://www.annualreviews.org/doi/10.1146/annurev-economics-081919-050239>

Techno-Orientalism: Artificial Intelligence and the “Other”

Arshin Adib-Moghaddam



Arshin Adib-Moghaddam is Professor in Global Thought and Comparative Philosophies at SOAS, University of London, a Chair that was exclusively created for him. Cambridge educated, he is the inaugural Co-Director of the SOAS Centre for AI Futures and the author of The Myth of Good AI: A Manifesto for Critical Artificial Intelligence (Manchester University Press, 2025).

This short paper interrogates the concept of Techno-Orientalism as a critical framework for understanding the sociopolitical and cultural implications of Artificial Intelligence (AI) development (see further, Adib-Moghaddam 2025). It situates the discourse within the historical construction of the “West” as a hegemonic entity, a narrative deeply entangled with the material and ideational presence of the “East.” By deconstructing the Eurocentric assumptions underpinning AI systems, the paper highlights how these technologies perpetuate racial and cultural hierarchies, treating race as an essentialized rather than socially constructed category. In a final step, I argue that AI ethics must be grounded in a recognition of shared humanity and the interconnectedness of local and global identities, with localized cultural symbolism.

Critical scholars have introduced the term “Techno-Orientalism” to accentuate techniques of exclusion engendered and galvanized by AI systems. In an important compilation of articles, some of these critical scholars link Techno-Orientalism to representing Asians and Asia itself in stereotypical fashion as technologically adept, yet intellectually underdeveloped and therefore in need of Western tutelage (Roh et. al. 2015).

In this way, Techno-Orientalism refers to the representational practice in which the East — particularly East Asia—is imagined through hyper-technological futures: roboticized societies, opaque digital authoritarianism, and alien forms of intelligence. Therefore, Techno-Orientalism continues older orientalist patterns, but now through the lens of technological imaginaries: The “Other” becomes a speculative laboratory for Western anxieties and desires about automation, surveillance, and posthuman futures.

Rather than capturing the realities of everyday life in non-Western societies, techno-orientalist narratives encode geopolitical fears—of Iran’s cyberwarfare, China’s AI dominance, Japan’s robotics economy, or South Korea’s digital culture—into fantasies that justify “Western” moral superiority, much in the same way traditional Orientalists deemed themselves a part of a civilizing mission (Adib-Moghaddam 2023). That cultural blindness towards the “Other,” which is an essential feature of the Western narrative, has not changed with the shifts in global power away from Europe and the United States.

Traditional and Techno-Orientalism

As the late Edward Said famously argued, traditional Orientalism depicted the peoples of the “East” as essentially backward because of their alleged racial inferiority (Said 1979). With the resurgence of Asian power centered around China, this attitude was very difficult to hold. Therefore, traditional orientalist attitudes transmuted into depicting Asians as “essentially robotic automata.” The process of othering typically relies on descriptions that deny the “Other” cultural agency or “civilizational” prowess. Asians may be successful, but they are only “capable of fine-tuned execution and coordination, but lacking the sort of individual creativity and spirit that defined the Western subject” (Kim 2022). Hence, a process of othering continues as the “West” represents “platforms in China as something inferior, different, or even morally depraved” (Yang 2023).

Techno-Orientalism thrives precisely through affect: it animates fantasies of threatening or technological futures that justify geopolitical anxieties. What I have called “the myth of good AI” relies on such civilizational claims—casting Western tech as rational and benevolent, while depicting non-Western tech as opaque or dangerous. Said would accentuate, via the French philosopher Michel Foucault, that such myths are systems of knowledge that naturalize power (see further Adib-Moghaddam 2012). In other words: Techno-orientalism and good-AI narratives mask who benefits from AI systems and who becomes vulnerable through them—elevating certain actors (big-tech, Western states) while rendering others (Global South users, marginalized communities) invisible. Wendy Hui Kyong Chun (2019) agrees: Such “high-tech Orientalism” safeguards a “white hegemony”

over the definition of what it means to be “human,” reiterating a racial hierarchy that rationalizes the exploitation of the “Other.” We may add that such techno-racism is largely a technique to safeguard Western dominance from the re-emergence of Asia as a locus for global power.

There is an important conceptual lesson to learn here. The idea that AI is a force for “good,” neutral, or civilizationaly beneficial operates as a myth that masks its political embeddedness. This myth requires three interdependent ideological moves that must be identified as essential techniques of obliteration: First, the purification of AI from its human, institutional, and corporate origins. Second, the cleansing of the dominant discourse about AI from the “Other,” and third, the erasure of structural harm, such as algorithmic discrimination, surveillance capitalism, and militarized automation. In this way, AI is framed as a pedagogical or civilizing tool—an extension of the “West’s” enlightenment project—while the real actors shaping AI systems, in particular private firms, security institutions, and states, remain abstracted away, safe and secure from public scrutiny and, most crucially, from legal accountability. In summary, then, the myth of good AI purifies a polluted and dangerous technology, creating the mirage that AI is apolitical, neutral, even benevolent, and an essential technology for human improvement. The harms that AI does on an everyday basis, conceptualized in the emerging critical AI studies corpus as algorithmic racism, surveillance capitalism, and labor exploitation, are treated as accidental glitches, rather than systemic features of AI technology.

Techno Co-Constitutions

So far, I have argued that “Western” AI is not

only a technical assemblage, but a cultural and (geo)political power project. When “Western” AI is imagined as “good,” it is repeatedly tethered to “Western” discursive and epistemological superiority. When the “East” (or the Global South) is imagined as hyper-technological, Asia becomes an object for Western fantasies and subjective imaginations, rather than a site with its own political, social, and cultural agency and complexity. Hence, the myth of good, “Western” AI reveals itself as yet another civilizational binary: Good AI belongs naturally to the “liberal,” “enlightened” rational societies of the “West.” The “East” is not to be trusted. When AI is harmful there, for instance, in Iran and China (but also in Russia), Western discourse frames it as a function of a retrograde political culture typical of the “Other,” thus rendered irrational. In this way, the myth of good AI allows Western institutions—including big tech companies—to claim a monopoly on ethical technological development, while externalizing harm to “illiberal” others.

Very much comparable to traditional Orientalism, the myth of good AI mobilizes desire—desire for progress, for mastery, for the rationalization of social life. At the same time, it displaces fear onto the “technological Other,” particularly to civilizations such as Iran and China, that have always had an exalted presence in the imagined “Western” narrative. Thus, the myth of good AI and techno-orientalism are co-constitutive. In this “Western” imagination, there is a fetish for “Asian efficiency” and “robotic perfection,” geared to a pathological fear of digital authoritarianism, framed with an increasing anxiety about Western decline. This angst about the resurgence of the “Other” animates cultural products, policy documents, think-tank reports, and media narratives.

This myth of good, “Western” AI relies on aspirational affects clearly discernible from the speeches of Elon Musk and others (for “Muskism” see Adib-Moghaddam 2025): A hysterical optimism about innovation, a religious belief in technological salvation, trust in the normative superiority of seemingly “Western” values such as capitalism and liberalism and a fundamentalist desire for the control of history, past, present, and future. These effects are incredibly dangerous for humanity, exactly because they constitute political reality and the fundamental friend-enemy distinction within societies and in international relations. The myths shape perceptions and, therefore, policies. They shape categories of threat and virtue. They determine which forms of violence become visible and which become normalized. In this sense, the myths surrounding AI are not peripheral discourses. They form the affective infrastructure of global technological politics.

Therefore, myths constitute an increasingly real delusion when they cease to be identified as false stories and transmute into symbolic architectures and techniques that organize society and political life. Techno-Orientalism functions exactly as a myth of Western technological virtue, producing a moral geography where the “West” remains the guardian of ethics and democratic progress, while the “East” is frozen as a site of either dystopian danger or monstrous hyper-modernity.

Techno-Orientalism as an affective myth is extremely potent. Such Techno-Orientalism legitimizes hybrid warfare, for instance, the recurrent attacks on Iran’s nuclear energy infrastructure. It prevents global technological governance, for example, the removal of legal barriers for AI research by Donald Trump and the US’s intransigence in governing AI internationally. Techno-Orientalism guides

export controls and sanctions such as the current US tariff regime and the comprehensive sanctions against Iran, Cuba, Venezuela, and other countries. It informs racialized policing and exploitative data extraction from marginalized communities within the “West” itself and in the Global South.

Decolonial Futures

What I have called critical AI studies suggests three pathways out of the current conundrum. First, the growing critical AI community needs to continuously challenge the myth that AI, and by implication technological innovation, belongs naturally to any one country or “civilization.” Our successful efforts in the humanities and the social sciences to dismantle notions of civilizational exceptionalism need to be transferred to the discourse about AI. We need to harass the “tech-bros” into compliance by the sheer weight of our intellectual argument in support of a form of “community AI” that serves human security.

Secondly, we need to “provincialize” the false moral authority inscribed into the Western narrative enveloping AI. The idea of the “West” has to be discarded in favor of a global reading of history, in general, and of the emerging AI universe, in particular. Only if we think of AI in universally human terms, rather than in a primitively tribal West-versus-the-rest dichotomy, can we tame the power excesses of private firms and states. Critical AI studies demand exposing this selective moral geography that the “West” has prided itself on for too long. Ancient cultures are reservoirs of a critical approach to AI, and these vast untapped sites of human knowledge can easily defeat ignorant delusions so typical of Orientalisms of any kind.

Such a movement, thirdly, foregrounds pluralistic human experiences and their technological cultures, epistemologies, and governance models, not as Ferrel deviations from the ideal “West” but as co-equal contributors to a common human experience. This means recognizing that ethical AI cannot emerge from a singular civilizational locus or template. Nor can AI be considered “good” simply because it originates from the geographical centers of global capitalism. As I have argued, Techno-Orientalism and the myth of good AI are not separate phenomena. They are intertwined mythologies that sustain global hierarchies, structure geopolitical anxieties, and shape perceptions of technological virtue and vice.

A critical engagement with AI requires precisely this kind of myth-diagnosis. We must understand how narratives produce power, how affects sustain hierarchies, and how the technological futures we imagine are inseparable from the political and social worlds we inhabit. I want to reiterate that our task is not merely to debunk these myths but to cultivate alternative imaginaries—ones that resist the civilizational binaries of Techno-Orientalism and the moral complacency of the “good AI” myth. Only then can we imagine technological futures that are genuinely plural, accountable, and humane. In this way, truly good AI emerges as a tool for global justice, rather than a mechanism of political domination and social exclusion.♦

References

Adib-Moghaddam, Arshin. 2012. “Can the (sub)altern resist? A dialogue between Foucault and Said.” In *Orientalism Revisited: Art, Land, Voyage*, eds. Ian R. Netton, 33-54. London: Routledge.

Adib-Moghaddam, Arshin. 2023. *Is Artificial Intelligence Racist? The Ethics of AI and the Future of Humanity*. London: Bloomsbury.

Adib-Moghaddam, Arshin. 2025. *The Myth of Good AI: A Manifesto for Critical Artificial Intelligence*. Manchester: Manchester University Press.

Chun, Wendy Hui Kyong. 2019. "Genealogy of the Posthuman." *Critical Posthumanism*. April 30. <https://criticalposthumanism.net/author/wendy-hui-kyong-chun/>

Kim, Leo. 2022. "Korean Illustrator Kim Jung Gi's Resurrection via AI Image Generator Is Orientalism in New Clothing." ARTnews, December 9. <https://www.artnews.com/art-news/news/kim-jung-gi-death-stable-diffusion-artificial-intelligence-1234649787/>

Roh, David S., Betsy Huang and Greta A. Niu (eds.). 2015. *Techno-Orientalism: Imagining Asia in speculative fiction, history and media*. New Brunswick: Rutgers University Press.

Said, Edward W. 1995. *Orientalism: Western Conceptions of the Orient*. London: Penguin.

Yang, Fan. 2023. "A Glitch in Translation: (Self-)Orientalism and Post-Orientalism in Platform Governance." *Yale-Wikimedia Initiative on Intermediaries & Information*: 1-12. https://law.yale.edu/sites/default/files/area/center/isp/documents/translation_ispesayseries_2023.pdf

Building AI Capacity Between Beijing and Washington: The Gulf's China Strategy

Guy Burton



Guy Burton is a Visiting Fellow on the SEPAD Project at Lancaster University. He has previously held teaching and research positions in Brussels, the Gulf and Malaysia. His research covers rising powers in the Middle East and he is the author of China and Middle East Conflicts (Routledge, 2020).

What defines cooperation between China and the Gulf Cooperation Council (GCC) states in high-value technologies such as artificial intelligence (AI)? What drives it, and what are its likely consequences? Set against intensifying US–China rivalry—between the world's two largest economies and leading AI powers—how are Gulf states, long-standing US security partners, managing the pressure?

The evidence points to strategic hedging. Gulf governments are not aligning exclusively with either Washington or Beijing when it comes to AI. Instead, they are seeking to maximize economic and technological opportunity while limiting political and security exposure. This reflects the national interests and motivations of these governments regarding AI, along with the extent to which they have advanced in developing a deeper and more expansive AI capacity. Moreover, these efforts do not operate in a vacuum, but must also take into account the two models of corporate autonomy and state strategy offered by the US and China. In the US case, it is a market-led approach in which big tech companies have scaled up innovations from venture capital-funded AI labs. Meanwhile, the Chinese model is one where the state has a greater role in directing and coordinating

national corporations in technology transfer and AI.

The United States remains the dominant external security actor in the Gulf and the wider Middle East. GCC states appear committed to maintaining this status quo. Qatar hosts US Central Command, the US Fifth Fleet operates from Bahrain, and Washington has designated several Gulf partners—including Saudi Arabia, Kuwait, and Qatar—as major non-NATO allies (Sentner 2025). Gulf states remain among the largest purchasers of US defense systems. High-level diplomacy reinforces this architecture: President Trump's regional visit in May 2025 and Crown Prince Mohammed bin Salman's White House meeting in November 2025 underscored deep political and economic ties, including Saudi pledges to invest up to \$1 trillion in the US economy (Ratney 2026).

Yet, a security partnership with the US has not translated into economic exclusivity. Gulf governments have diversified partnerships across Europe and Asia, including with Chinese firms (Mogielnicki 2022). Crucially, this engagement has not—at least so far—produced structural technological subordination. Gulf regimes have avoided becoming locked

into Chinese state or corporate control in the AI sector and have signaled ambitions to shape, rather than simply accept, the terms of cooperation.

Elite incentives help explain their visions and actions. Gulf rentier systems historically rested on hydrocarbon revenues that sustained regime durability through patronage, welfare provision, and guaranteed public employment (Beblawi and Luciani 1987). Such arrangements typically dampen reform. However, structural pressures—energy transition, price volatility, demographic and generational shifts, and technological disruption—have altered elite calculations. Economic models anchored solely in hydrocarbons are increasingly exposed.

This recognition has driven diversification programs over the past decade, including Saudi Vision 2030, UAE Vision 2031, and Qatar National Vision 2030. While differing in emphasis, all aim to reduce reliance on oil and gas rents and expand into various high-value sectors, including advanced manufacturing, digital infrastructure, and AI (Bazoobandi 2025). Initiatives to encourage AI development range from Saudi Arabia's \$100 billion Project Transcendence to the UAE's G42 partnerships with firms such as TSMC, Samsung, and Microsoft Azure.

Diversification attempts in the Gulf have a long history, dating back to the 1970s in Saudi Arabia and Oman, and later in the rest of the region. Earlier waves faltered for multiple reasons. Investment cycles tracked hydrocarbon revenues: when oil prices rose, diversification expanded; when they fell, those measures stalled. Among the reasons for this happening were tied to the heavy reliance on state support and insufficient independent forms of assistance through the private sector

and foreign investment. Additionally, these diversification efforts were not undertaken alongside a rationalization or improvement in efficiency gains in the public sector. There was also insufficient uptake from domestic capital and a continued reliance on migrant labor. As a result, governments retreated to subsidies and public hiring when social resistance intensified (Hvidt 2013).

The current diversification efforts may be different. Sovereign wealth funds in the Gulf have expanded significantly over the past decade, more than doubling from \$2 trillion in assets in 2015 to \$5.35 trillion by 2025. This constitutes more than 40% of all assets under management by sovereign wealth funds around the world (Touazi 2025). These assets provide capital buffers capable of sustaining long-term investment domestically and abroad. Reform has also proceeded incrementally to manage political risk. In Saudi Arabia, for example, this has included selective “authoritarian liberalization” measures such as expanded entertainment sectors and lifting restrictions on women driving in order to recalibrate the social contract without loosening political control (Bassiouni 2022).

China has become central to the current diversification trajectory as well. Previously marginal, China became a net energy importer in 1993, prompting closer regional ties to secure reliable energy for its rapidly growing economy (Huwaitdin 2022, Olmat 2023). While the initial focus was on energy, economic exchange and trade have expanded. Chinese contractors became prominent in infrastructure construction under the Belt and Road Initiative. More recently, cooperation has shifted toward telecommunications, 5G networks, fiber-optic infrastructure, data centers, and AI applications under the complementary Digital Silk Road framework

(Chaziza 2025; Aburaya 2025). Firms such as Huawei, ZTE, Alibaba, and Tencent have become embedded in Gulf digital ecosystems, offering competitively priced, integrated technological packages (Mogielnicki 2022).

Washington responded to these developments with varying degrees of caution. As the US-China rivalry intensified around the world, it impacted the AI sector as well, from semiconductor supply chains and critical minerals to data governance, as well as concerns about indirect technology leakage (Huq 2024). During the first Trump and Biden administrations, the US pressed Gulf partners to be wary of economic engagement with Chinese entities (Carcidi and Soliman 2024). Gulf governments, reliant on US security guarantees, sought to manage these sensitivities. The UAE's G42 reportedly divested from certain Chinese partnerships amid American scrutiny (Clemmensen, Redlich, and Rumley 2024). Since Trump's return to the White House, he has taken a less confrontational approach, authorizing advanced semiconductor exports to Saudi Arabia and the UAE despite earlier concerns over potential downstream exposure to Chinese actors (Soliman 2026). In May 2025, he was won over during his visit to the Gulf, where GCC leaders understood his preference for personal and transactional relations which seek to maximize commercial opportunities. Saudi Arabia pledged to invest up to \$600 billion in the US economy – a figure that Mohammed Bin Salman increased to \$1 trillion when visiting the White House later in the year (Unit for Political Studies 2025). In this regard, Trump's foreign policy during the first year of his second term was in line with that of his first term as whole: focused on countering perceived grievances against the rest of the world by advancing "America First" populist nationalism while challenging the precepts of the liberal

international order that the US had both constructed and maintained (Schortgen 2024, Reiss 2026). In the Middle East, this was echoed in ongoing uncertainty and unpredictability, including his attack on Iran with Israel in June 2025 and again in February 2026.

Chinese AI engagement in the Gulf is uneven and reflects the readiness gap between states rather than any preference to align with the Chinese or US approaches to AI – although, as will be observed later, much of the Gulf's current AI infrastructure has been built on the US model and by US providers, which presents challenges in operating independently.

Broadly, Kuwait, Bahrain, and Oman lag in digital infrastructure and AI adoption. By contrast, Saudi Arabia and the UAE have invested heavily in digital transformation and e-government, positioning themselves as regional leaders (Bi et al. 2025). When it comes to AI, Soliman (2025) notes that the UAE was the region's first mover, building regulatory frameworks, data center density, and AI partnerships. Saudi Arabia's advantage is its size and capacity to scale, which it is now doing by building up its institutions and computing power capacity. For the other GCC countries, like Qatar, Oman, Kuwait, and Bahrain, there are varying levels of ambition and interest in terms of what they wish to focus on when it comes to digital technologies and AI. Among them, it is Qatar, along with Saudi Arabia and the UAE, that have attracted the bulk of Chinese AI activity (Ayaz 2025).

Dubai was an early adopter of Chinese technology, embedding facial recognition and behavioral analytics systems in initiatives such as "Police without a Policeman" and "Oyoon" – "Eyes" in Arabic – alongside broader smart-

city and cybersecurity programs (Hasib and Shire 2022). Saudi Arabia has deployed AI tools for Hajj crowd management and proposed to incorporate Chinese-linked technologies into megaprojects such as NEOM, the planned 170-kilometer-long futuristic linear city that was to run from the sea into the desert. However, following delays and the technical and financial challenges of building such a city entirely from scratch in the middle of the desert, coupled with demands on government spending elsewhere, the more fanciful elements around NEOM have been downgraded. In its place, NEOM has been repurposed as a potential hub for large-scale computing infrastructure. The reasons for this shift revolve around its geographical location, from Saudi Arabia's potential to provide connectivity between Europe, Africa, and Asia, as well as its access to the sea and desert: the former providing water to support the cooling systems needed, and the latter, where solar and wind power can be accessed to provide the energy needed for large-scale computing (Chaudhary 2026). As for Qatar, its AI approach with China has been more research-focused, emphasizing smart-city development in Lusail and academic and cloud infrastructure partnerships (Ayaz 2025).

Risks remain in China-GCC AI cooperation. The deployment of surveillance-oriented AI within authoritarian governance frameworks raises human rights concerns. While these concerns may not apply to Trump himself, they certainly do to a large segment of the US political establishment, as well as the private technology sector. Some of these concerns were captured in the narratives presented by the National Security Commission on Artificial Intelligence, which operated between 2019 and 2021 (Retzmann 2024: 238, 241).

Technological transfer is also structurally

limited: Chinese firms typically export mature systems while retaining frontier innovation capacity, leaving the Gulf with capital and markets but little integration into China's core AI developments (Marks 2026a). This reflects Beijing's more state-oriented approach and concern with technological autonomy and risk management – an area that arguably appears less of a cause for concern following Trump's decision to scrap his predecessor's semiconductor export control regime. It may be characterized as a continuation of the "energy-for-industry" dynamic that has shaped China-GCC relations since 1993. In other words, the Gulf countries export hydrocarbons and China exports manufactured and industrial goods. However, such assumptions underestimate Gulf agency, as regional leaders are not content to accept a purely subordinate role (Cheng 2025).

GCC states seek to maximize economic opportunities by accommodating both Chinese and American interests. At the same time, they also pursue independent technological ambitions and are unwilling to be passive recipients of advanced technology. In particular, these states aspire to turn the Gulf into an AI hub with its own related infrastructure. Their objective is not only to serve their own residents, but also to act as a gateway to the wider Global South, including the rest of the Middle East, Africa, and Asia. In doing so, the Saudis and Emiratis envisage themselves as offering an alternative third option to the infrastructure provided and powered by American and Chinese entities (Marks 2026a, 2026b). Additionally, the Gulf states aim to use the comparative advantages they have to reach these goals. For instance, the UAE aims to become a regional hub for AI development and integration. Meanwhile, Saudi Arabia intends to invest in rare-earth extraction and processing. It also endeavors to exploit its

cheap energy to support high-performance computing, positioning itself as potentially the third-largest AI market globally (Alrajjal 2025; Ghosh 2025).

That said, Saudi and Emirati ambitions are not without their challenges and risks – all of which may only become apparent with time. For one, the Gulf AI “Third Option” may presume a degree of neutrality between the American and Chinese technological and AI superpowers. However, the reality remains that the bulk of the AI infrastructure of the GCC has been built on US models, chips, and cloud architecture – which means that it can never fully break with Washington and its concerns, including governance. Additionally, other countries in the Global South may be less interested in using the Gulf as an alternative and instead rely on either or both of the superpowers. Regardless of whether the American option is available, these countries turn to the Chinese one because it is cheaper and offers fewer restrictions on governance or regulation. Finally, the approach to AI development and innovation inside and outside the Gulf could change from the current big tech/scaled-up model reliant on large, centralized data centers to smaller and more regionalized hubs (Marks 2026b).

In sum, the Gulf presents a dynamic space within which wider US-China competition is playing out, especially in the AI sector. Gulf states—particularly Saudi Arabia, the UAE, and Qatar—are making use of the opportunities presented by the differences between Washington and Beijing to widen their economic and technological options. Whether it is the market-led AI model associated with the US or the more state-oriented and coordinated version offered by China, both have contributed to the forms and availability of technologies exported to the region. How-

ever, it is important to note that even as GCC states seek to become an alternative AI hub, they will not be able to do so without the input from the two technological superpowers. The efforts of these states to hedge to minimize risk and maximize opportunities remain contingent on the wider geopolitical climate. Should the US veer in a different direction and intensify pressure on GCC countries to disengage from China, coupled with Beijing resisting deeper integration, they could have less space to maneuver.♦

References

- Aburaya, Farah. 2025. “The Future Role of China in the GCC’s Tech Transition,” *Al Habtoor Research Centre*, October 20. <https://www.habtoorresearch.com/programmes/the-future-role-china-gcc/>
- Alrajjal, Tala. 2025. “Saudi Arabia Is Making a Massive Bet on AI,” *CNN*, November 2. <https://edition.cnn.com/2025/11/02/tech/saudi-arabia-ai-powerhouse>
- Ayaz, Behrouz. 2025. “China’s AI Push in the Persian Gulf Region,” *The Diplomat*, 22 Revised.November. <https://thediplomat.com/2025/11/chinas-ai-push-in-the-gulf-region/>
- Bassiouni, Moustapha Chérif. 2022. “Saudi Reforms: Change for Survival or for Progress?” *European Institute of the Mediterranean*. <https://www.iemed.org/publication/saudi-reforms-change-for-survival-or-for-progress/>
- Bazoobandi, Sara. 2025. *The GCC Countries’ Quest for Economic and Digital Transformation*. Hamburg: GIGA.

- Beblawi, Hazem, and Giacomo Luciani. 1987. *The Rentier States in the Arab World*. London: Croom Helm.
- Bi, Ran, Fozan Fareed, Jeong Dae Lee, Sidra Rehman, Yuan Gao Rollinson and Tongfang Yuan. 2025. *Digital Transformation in the Gulf Cooperation Council Economies*. Departmental Paper (April). Washington DC: International Monetary Fund.
- Carcidi, Vincent, and Mohammed Soliman. 2024. *The Role of the Middle East in the US-China Race to AI Supremacy*. Washington DC: Middle East Institute.
- Chaudhary, Isha. 2026. "Saudi Arabia Scales Back the Line Project to Build AI Data Centers," *Parametric Architecture*, March 9. <https://parametric-architecture.com/saudi-arabia-scales-back-the-line/>.
- Chaziza, Mordechai. 2025. "The Great Power Tech Race in the Persian Gulf: A Western Perspective." In *Technology Rivalry Between the USA and China*, ed. Peter Chow. 355-390. Basingstoke: Palgrave Macmillan.
- Cheng Wang. 2025. "Reflections on the Future Investment Forum," *China-MENA Newsletter*, November 19. <https://chinamenanewsletter.substack.com/p/cheng-wangs-reflections-on-the-recent>
- Clemmensen, Andrew, Rebecca Redlich and Grant Rumley. 2024. "G42 and the China-UAE-U.S. Triangle" *Washington Institute for Near East Policy*, April 3. <https://www.washingtoninstitute.org/policy-analysis/g42-and-china-uae-us-triangle>
- Ghosh, Indranil. 2025. "U.S. gains in AI race as Gulf nations ditch China for chips," *Rest of World*, November 21. <https://restofworld.org/2025/mideast-us-chip-deal-and-china/>
- Hasib, Bassant, and James Shire. 2022. "Cybersecurity in the GCC," *Middle East Policy* 29(1): 90-103.
- Huq, Aziz. 2024 "A world divided over artificial intelligence," *Foreign Affairs*, March 11. <https://www.foreignaffairs.com/united-states/world-divided-over-artificial-intelligence>
- Huwaidin, Mohamed. 2022. "China and the Gulf region: From strangers to partners." In *Routledge Handbook on China-Middle East Relations*, ed. Jonathan Fulton. 81-92. London: Routledge.
- Marks, Jesse. 2026a. "China's AI Cooperation in the GCC," *Coffee in the Desert*, January 6. <https://jessemarks.substack.com/p/chinas-ai-cooperation-in-the-gcc>
- Marks, Jesse. 2026b. "Can the GCC build the third AI option?" *Coffee in the Desert*, February 23. <https://jessemarks.substack.com/p/can-the-gcc-build-the-third-ai-option>.
- Mogielnicki, Robert. 2022. "Technological Dimensions of China-MENA Economic Relations." In *Routledge Handbook on China-Middle East Relations*, ed. Jonathan Fulton. 281-296. London: Routledge.
- Olimat, Muhammad. 2023. "China and the Middle East: An Overview." In *Routledge Companion to China and the Middle East and North Africa*, ed. Yahia Zoubir. 9-24. London: Routledge.
- Ratney, Michael. 2026. "Saudi Arabia's Mohammed bin Salman Got a Lot from Trump. What Did the United States Get?" *Center for Strategic & International Studies*, January 5. <https://www.csis.org/analysis/saudi-arabias-mohammed-bin-salman-got-lot-trump-what-did-united-states-get>

Reiss, Mitchell. 2026. "Trump's Foreign Policy After Year One: A Look Back, A Look Ahead," *Royal United Services Institute*, January 20. <https://www.rusi.org/explore-our-research/publications/commentary/trumps-foreign-policy-after-year-one-look-back-look-ahead>

Retzmann, Nike. 2024. "Winning the technology competition': Narratives, power comparisons and the US-China AI race." In *Comparisons in global security politics* eds. Thomas Müller, Matias Albert and Kerrin Langer. 237-256. Bristol: Bristol University

Press. Schortgen, Francis. 2024. "U.S. Foreign Policy in the Trump Administration: The Unmooring of America's Global Roles?" In *The Legacy of the Trump Administration: A Disruptive Presidency*. ed. Michael Grossman, Francis Schortgen, Ronald Matthews Jr and David Cohen. 121-145. Basingstoke: Palgrave Macmillan.

Sentner, Irie. 2025. "Trump elevates Saudi Arabia to 'major non-NATO ally' status," *Politico*, 18 November. <https://www.politico.com/news/2025/11/18/trump-saudi-arabia-ally-00658467>

Soliman, Mohammed. 2025. "From Crude to Compute: Building the GCC AI Stack," *Middle East Institute*, December 10. <https://mei.edu/report/from-crude-to-compute-building-the-gcc-ai-stack/#pt1>.

Soliman, Mohammed. 2026. "US Authorizes Chips for the UAE, Saudi Arabia," *Middle East Institute*, January 28. <https://mei.edu/policymemo/us-authorizes-chips-for-the-uae-saudi-arabia-2/>

Touazi, François-Aïssa. 2025. "Les fonds souverains du Golfe, acteurs majeurs de la finance mondiale" [Gulf sovereign wealth funds, major players in world finance], *Politique étrangère* 90(4): 25-38.

Unit for Political Studies. 2025. "Deals Without Guarantees: Trump's Gulf Visit and the Future of US-Gulf Relations," *Arab Center Washington DC*, May 22. <https://arabcenterdc.org/resource/deals-without-guarantees-trumps-gulf-visit-and-the-future-of-us-saudi-relations/>

Roundtable Discussion: AI in MENA Politics Research

Carolyn Barnett in conversation with Muhammad S. Abdo, Tariq Adely, Cinzia Bianco, Ashrakat Elshehawy, Robert Kubinec, and Michael Robbins

Muhammad S. Abdo is a dual-major PhD student in Computational Linguistics and Middle Eastern Languages and Cultures at Indiana University Bloomington, with a minor in Computer Science. Tariq Adely is a PhD candidate in the Department of Anthropology at George Washington University. Drawing on ethnographic methods, his research examines intersections of language, labor, and technology in the Arabic-speaking world. Carolyn Barnett is an Assistant Professor in the School of Government and Public Policy and School of Global Studies at the University of Arizona. Her research focuses on public opinion, social norms, and political behavior in the Middle East and North Africa, particularly on gender-related issues. Cinzia Bianco is a senior researcher at the European Council on Foreign Relations, where her work focuses on Gulf states and their relations with Europe. She holds a PhD in Gulf studies from the University of Exeter and lectures on Gulf studies at the LUISS School of Government. Ashrakat Elshehawy is an Assistant Professor in the Department of Political Science at University College London. Her research focuses on political economy, development, and historical political economy, employing methods of computational text analysis and AI across multiple languages. Robert Kubinec is an Assistant Professor in the Department of Political Science at the University of South Carolina. Michael Robbins is Project Director and co-Principal Investigator at Arab Barometer.

Introduction, by Carolyn Barnett

For this roundtable, I invited a range of emerging and established scholars who conduct research about and/or use AI in some capacity to share their perspectives on the implications of AI for the future of research on politics in the MENA region and academic research more broadly. I posed questions that asked participants to reflect on lessons learned from their own experiences, potential positive and negative impacts, and hopes and concerns about the future. Contributors submitted their responses in writing, and I have combined them into the following conversation which highlights where their views converge and diverge.

Broadly, most of the contributors are cautiously optimistic about the ability of AI

tools to facilitate working with and deriving insights from large quantities of textual data, ranging from historical records to contemporary media to survey data.

While in part these capabilities will speed up existing research, they also make it possible to undertake analyses that would not have been attempted before because they would have been too time- or labor-intensive. At the same time, contributors are variously wary of the energy demands and ethical implications of AI; the temptation to view LLMs as understanding and synthesizing the true meaning of texts, when this is not what they do; the possibility that existing biases in data will be reproduced and amplified; and the risk that the new relative ease of data analysis will

increase the volume of misleading or downright incorrect claims in circulation.

The conversation is accompanied by a stand-alone essay by Richard Nielsen which argues that Generative AI (GenAI), by streamlining the design and conduct of increasingly advanced quantitative analyses, both frees researchers to spend more time engaged in human-intensive, qualitative data collection and underscores the unique value of interpretive insights generated through direct interaction with the societies we study. While drawing inferences from large quantities of data will be technically easier than ever, deriving meaningful insights from that data will continue to require—and may require more than ever—the expertise and ethnographic sensibilities that researchers hone through personal experience and human interaction, raising the value of these experiences and commitments.

Barnett: What is important to understand about AI itself as we think through its impact on our field?

Kubinec: It is important to be clear about what AI tools are and are not. The building blocks of these models preceded the release of ChatGPT by a significant margin. One core technology, text embeddings, made its way into political science research several years earlier (Gupta et al. 2020; Rodriguez and Spirling 2022; Rodriguez, Spirling, and Stewart 2023). Embeddings are a form of text analysis that is much more nuanced than earlier iterations popular in political science, like structural topic models (Roberts et al. 2014), because embeddings account for both words themselves and their context from many examples of usage. ChatGPT's model combined text embeddings derived from a massive textual corpus and a model that predicts

natural language responses given prompts. As a result, these large language models (LLMs) can both interpret human input in quantitative terms (embeddings) and create responses by regurgitating relevant passages and words from its training data.

While these technical accomplishments are impressive, especially given the enormous scale of the training data, LLMs are not models of cognition. While many people feel as though they are talking to a sentient being when interacting with an AI bot, that is mainly due to the enormous scope of the data from which its embeddings are derived. As a result, while LLM providers like Anthropic and OpenAI have gone to great lengths to make their software present itself as sentient, these tools have no underlying intelligence. Rather, by projecting most human-written text available into a quantitative format (embeddings) and employing remarkable computational advances (Nvidia's graphical processing units or GPUs), they can identify relevant texts when prompted with such speed that to a human it appears as if the machine is conversing.

Abdo: Today, when people talk about artificial intelligence, they are often referring to large language models (LLMs). Here's a way to think about them: In 1941, the Argentine writer Jorge Luis Borges published his short story *The Library of Babel*. In it, Borges depicts the universe as an infinite library housing every possible arrangement of 25 characters. The central idea of the story emerges from this premise: if every possible arrangement of these characters does exist somewhere, then the library must contain not only countless meaningless books but also every meaningful book that has been written or could ever be written. The library is both omniscient and indifferent in such a way as

to leave its librarians wandering in perpetual uncertainty about where meaning lies.

LLMs are, in many ways, the technological realization of this premise. Trained on massive corpora that approximate the library's totality, even though they do not "know" in any human sense, they still can reproduce linguistic patterns with enough fluency to simulate understanding. These models move through *the library* far more quickly than humans, but they are still tied to the same deeper uncertainty about meaning, truth, and the limits of interpretation.

Barnett: What is one way you have used AI in your research to date, and what is one reason you have been cautious about or skeptical of applying AI tools in your work?

Abdo: Over the past decade, artificial intelligence has occupied a central place in my formal research agenda. My research typically involves curating datasets that support the development, evaluation, and critical analysis of language technologies while also fine-tuning LLMs to better handle domain-specific and linguistically complex tasks. And since the advent of LLMs, I have collaborated with colleagues on building several AI systems across different domains. For example, in AMWAL (Abdo et al. 2025), a project focused on Arabic financial news, we developed a first-of-its-kind model capable of automatically detecting and classifying key financial entities and events, including companies, financial instruments, and products, which is an essential task for structuring unorganized financial text into analyzable data. We also worked on a project designed to help AI systems distinguish between 18 Arabic dialects, a task that remains especially difficult because of Arabic's rich and complex morphology.

However, despite my extensive involvement in AI research, I have approached its applications with a degree of caution and skepticism. In addition to the many problems that I have witnessed first-hand while working with LLMs—such as the risk of bias and inequity in datasets, hallucination issues, ethical and privacy challenges, and lack of interpretability—I remain cautious about the sustainability and scalability of these technologies. The major concern hinges on the intensive and expensive computational demands of training and deploying such models, which raise significant environmental and accessibility issues. This is because training large models demands prohibitive amounts of Graphics Processing Unit (GPU) time that is equivalent to months of household energy use (Strubell et al. 2018; Champion 2023), which compounds the challenge for countries in the MENA region where the necessary technologies remain largely inaccessible. This escalating drain on resources increases the carbon footprint of AI, which is already on par with entire sectors like aviation (Jackson & Hodgkinson 2022) and presents numerous challenges to global sustainability goals.

These bottlenecks become even more pronounced in low-resource environments where deploying these models on mobile devices yields real-time performance that often falters due to scalability-related deficiencies. For instance, when deploying a lightweight transformer-based model on a mid-range Android device, response delays can exceed several seconds per query, which renders the tool unusable for real-time applications. This forces developers to either compromise on model capability or depend on cloud connectivity, an assumption that remains unrealistic in areas with limited or unstable Internet infrastructure. These challenges may exacerbate economic disparities in multiple ways,

including favoring well-funded institutions and limiting access for researchers in developing countries. AI also introduces opportunity costs. Efforts to build AI capabilities may divert attention from viable alternatives to AI, such as human-led data curation, which may offer more equitable and less resource-intensive trajectories.

Adely: My primary engagement with artificial intelligence (AI) technologies has come in the context of fieldwork. In my research, I use ethnographic methods to examine how differently situated workers produce, adapt, and deploy AI models for Arabic, focusing on the information technology (IT) sector in Amman, Jordan, a regional outsourcing destination for technical and linguistic labor. While conducting fieldwork from 2023 to 2025, I regularly tested AI language models and their updates, carried out participant observation at startups utilizing AI technologies, and took courses about building and fine-tuning AI models. Experimenting with these models—and reflecting on those interactions—was central to understanding how my interlocutors made, and made sense of, these emerging technologies. My work shows how ideological claims about Arabic and long-standing linguistic hierarchies shape AI models, political economies, and labor arrangements.

Outside of these fieldwork contexts, I have largely minimized my use of AI technologies, aside from a small number of proofreading tools (e.g., Grammarly) and scholarship search engines (e.g., Research Rabbit) billed as “AI-powered.” Despite—or, more likely, because of—my research on these models, I remain skeptical about their applications for data collection, analysis, and text generation. AI technologies raise a variety of ethical questions: the environmental impacts of

training and maintaining models (see Wu et al. 2022 for an overview); the exploitative labor practices underpinning their production (Gray 2019; Kirk 2025); and the wholesale scraping of academic knowledge to train and monetize these models (Gibney 2024). The key source of my skepticism, however, relates to the technological logics underpinning today’s generative AI models, including large language models (LLMs) like OpenAI’s GPTs, Gemini, and Claude. Put simply, the capacity of these AI technologies to generate text and speech that mimic human communication resides in their ability to predict what “token,” a unit of textual data, will most likely follow in a sequence, given the surrounding context. As AI engineers often put it during interviews: models are very good at guessing the next word.

The fact that AI’s “generative” capabilities rest in predicting the most likely response, reflecting countless past examples in its training dataset, casts a shadow over its potential for scholarly work. I would also suggest that this concern is especially pronounced for scholarship on the Middle East and North Africa (MENA), given the long history of essentializing frameworks and narratives applied to political, cultural, and social life in the region (see Wedeen 2002 on cultural essentialism; also Lockman 2010; Said 1979). This is not to say that we ought to turn away from artificial intelligence—as a technology, a discourse, and a driver of shifting political-economic relations—in our research on the region. It is reshaping labor dynamics, state-society relations, policing, surveillance and countless other domains of social, political, and economic life. But as a technology whose outputs reproduce an archive of previous utterances, I am not optimistic about its potential to advance social scientific scholarship at the current moment—particularly from and about

the MENA.

Elshehawy: One fast-growing use of LLMs in political science is text classification and annotation (Gilardi, Alizadeh, and Kubli 2023; Törnberg 2025). In recent ongoing work, my coauthors and I used a large language model (GPT-4o) to classify approximately one million police press releases issued by local police stations across Germany over a decade (Elshehawy et al. 2025). One of our main goals was to identify whether each press release contained out-group cues: explicit or implicit references to a suspect's foreign nationality or ethnicity. This is the kind of classification task where LLMs can be useful. The corpus was massive and unstructured, authored by different press officers, and the information we needed to extract was often conveyed through subtle, contextual language rather than consistent keywords. Dictionary-based approaches and some of the conventional Natural Language Processing tools would likely miss part of what mattered, especially where cues were implicit and context-dependent. LLMs allowed us to capture implicit meaning at scale.

The classification challenge we faced (large-volume, unstructured, context-dependent text) is one that many MENA researchers will recognize. Whether the source material is Arabic-language media archives or historical documents, the core problem is the same: extracting structured information from messy text where meaning is often embedded in context rather than stated outright. LLMs can be well-suited to this kind of task, especially when the target concept is context-de-

pendent.

But, of course, the promise of LLM tools has to come with careful thinking about how to validate their output. It is important to think carefully about the prompt engineering phase. Small changes in prompt wording can meaningfully shift classification output. Sensitivity to prompt variation should be tested and carefully considered (Halterman and Keith 2025). A good practice is for researchers to validate the LLM output against a hand-coded gold standard and report full performance metrics: F1 scores, precision, and recall, not just overall accuracy.¹ Overfitting is a risk when prompts are iteratively refined. Validation on a hold-out set not used during prompt development guards against this. Finally, as part of validation, it is useful to assess inter-coder reliability among human coders and agreement between human coders and the LLM. This allows researchers to check if the LLM judgments align with those of trained annotators, while recognizing that agreement alone is not sufficient evidence of validity. These are some of the common validation practices, but the methodological literature in this area is expanding fast and worth keeping up with.

In sum, LLMs make it increasingly feasible to extract structured information from the large, often unstructured corpora that are common in MENA research, but their use also makes careful validation especially important.

¹ Precision measures the proportion of all positive classifications of the model that are in fact positive
 $\text{Precision} = (\text{True Positive}) / (\text{True Positive} + \text{False Positive})$

Recall measures the proportion of all actual positive cases that were classified by the model accurately as positives

$\text{Recall} = (\text{True Positive}) / (\text{True Positive} + \text{False Negative})$

The F1 score is the harmonic mean of precision and recall.

Kubinec: I have mainly employed these tools in my own research to take care of technical clerical tasks and machine translation. For my research, this involves deriving code for a variety of projects from cleaning data to running complex online surveys. The code generation capabilities of LLMs are quite useful because they can respond to natural language prompts; this means that I can express what code I want in colloquial language that the model can then interpret as code. Similarly, most LLMs seem to excel at translation—which is in a sense what they are doing when converting a text prompt into a programming language. In fact, some of the most useful tasks have involved prompting an LLM to input Arabic text directly into computer code, which is a very difficult thing to do given the right-to-left versus left-to-right differences in Arabic typography and code.

I put LLMs to the test in a recent project with Aseem Mahajan (King Fahd University of Petroleum and Minerals) in which we collected a dataset of over 1 million advertisements on Facebook that contained climate change-related content. We prompted the LLM to identify any content that could be described as “greenwashing”, i.e., including misleading statements about an organization’s impacts on climate. We employed a set of six open-source LLMs for this task to compare performance. Again, thanks to the LLM’s remarkable translation capabilities, we were able to conduct this exercise in more than ten languages (including Arabic)—which would be beyond the capabilities of most research teams.

While we found that the LLMs did identify useful information, in total their performance relative to humans was mediocre (Kubinec and Mahajan 2025). The LLMs classified about one-third of the ads as containing

greenwashing content whereas our human coders found only 2-3%. However, by combining our LLM coding with other indicators derived from the ads, we were able to use a quantitative measurement model to extract information about greenwashing from the noisy responses. The LLMs did improve our project, especially as human coding for a database of this size was not trivial. However, robust assessment was necessary to avoid making vastly over-confident coding decisions from the LLMs alone.

Robbins: The clearest use to date of AI in the research of Arab Barometer, where I serve as Project Director, has been to help us write code and summarize data. Each wave of Arab Barometer requires us to write thousands of lines of code to conduct quality control and produce data visualizations in our standard branding, for example. Much of this work, such as the code needed to produce a topline, is fairly rote in nature but can take a significant amount of time. We have therefore explored using AI to produce initial code in some cases.

Second, we have used AI to generate preliminary summaries of findings on specific topics or questions. AI can take results and quickly produce an overall summary or compare findings to previous waves without requiring us to manually enter all of the results.

The key limitation we have found so far is that AI frequently makes minor errors. None of the code it has written runs without corrections or adjustments by our researchers. Additionally, AI has not always reported results fully accurately in its summaries, requiring careful fact-checking of its findings. In short, any task we have asked AI to perform still requires detailed and engaged review. Still, it has been helpful in increasing the efficiency with which we use and dig into the

project's findings.

Barnett: What is one exciting opportunity you think AI presents for Middle East researchers? How can researchers responsibly take advantage of this opportunity?

Abdo: When considering what AI has to offer researchers in the MENA region, we often think first about its promise for accelerating large-scale data analysis, surveying the literature, and streamlining time-intensive research tasks. Yet for humanities-oriented research, AI may also present an even more consequential possibility: affirming and amplifying the institutional relevance of the humanities by making humanistic expertise central to the design, evaluation, and governance of language technologies. There is a need for MENA-focused humanities scholars to play such a role: current LLMs are not culturally neutral. Studies have shown that widely used models often reproduce value patterns associated with English-speaking and European contexts, while Arabic-focused and multilingual models can also display bias toward Western-associated entities and stereotypes (Navigli, Conia, and Ross 2023; Naous et al. 2024; Tao et al. 2024).

Meanwhile, the humanities have been under pressure not just rhetorically but institutionally. Humanities disciplines in the Middle East face declining prestige, reduced funding, deteriorating infrastructure, and worsening labor conditions. In Egypt, the different fields of humanities have been suffering from poor student-to-faculty ratios and lower social valuation than disciplines such as medicine or engineering. These pressures paint a concrete picture of a compounding crisis, where institutional strain and social devaluation operate in tandem (Fahmy 2017; Hodapp 2022). These concerns have not been voiced by

researchers alone. In Egypt, they have also appeared at the level of state discourse. In April 2025, President Sisi questioned the value of fields such as arts, commerce, and law and urged families toward programming, communications, and AI-related paths. Later, in March 2026, he called for an end to majors with no clear employment prospects. The policy was later clarified by the Chair of the Parliamentary Committee on Education and Scientific Research, and former Egyptian Minister of Higher Education, as not a full abolition of humanities fields, but a reduction in admissions to disciplines seen as misaligned with labor-market demand (Ahram Online 2023; Mamdouh 2026; Ahmed 2026). This may open a space for interested humanities scholars, such as linguists, historians, philosophers, and area specialists, to demonstrate how their expertise can contribute to the development of AI. They can work on defining what culturally valid evaluation benchmarks look like, curating archives and corpora, and developing new curricula in areas such as AI safety and AI governance.

Elshehawy: The opportunity I find exciting is the potential of multimodal AI to unlock sources, such as archival records, historical documents, maps, and administrative data, that have long been difficult to analyze at scale. Multimodal AI refers to systems that can work across multiple data types, such as images and text. Conventional optical character recognition has struggled with aspects of Arabic-script text for a long time (especially for handwritten or degraded documents); the cursive nature of the script, variation in diacritical marking, and font diversity have made reliable automated transcription elusive. Recent advances in multimodal LLMs represent a step forward. This is a rapidly expanding area of research. Models are increasingly capable of reading not just printed

Arabic, but also some handwritten manuscripts, although limitations remain across domains, layouts, and complex document formats (Heakl et al. 2025).

For MENA researchers, this matters because vast quantities of source material exist only as physical documents or image scans. Newspaper archives, court records, and administrative files across the region remain locked in formats that have long been difficult to analyze at scale. AI-powered Optical Character Recognition (OCR) can change that.

The implications for historical political economy research, for example, are profound. Multimodal AI could make it increasingly feasible to build larger and more systematic datasets from archival sources, tracing patterns of landholding, endowment activity, and institutional change across locations and over time. The methods can advance research on questions related to land inequality, institutional persistence, state-building, taxation, and state-society relations; especially as multimodal LLMs continue to extend the reach of computational tools to handwritten and archival materials that have so far been difficult to process at scale.

Of course, this also raises questions about the verification of these methods. The use of multimodal AI for Arabic-script OCR and archival data extraction is an expanding field. Sample-based human verification against original documents is important, particularly for degraded or ambiguous passages where models are most likely to err. Rather than treating all model output as equally reliable, workflows should flag low-confidence extractions. Important metrics depend on the use case, but may include character error rates and word error rates, among others. Researchers should also report limitations,

documenting, where possible, error rates by script type, document condition, and historical period. Ideally, domain experts—such as historians, paleographers, and linguists with deep knowledge of the source material—would be part of the verification loop. Scholars trained in the relevant traditions can catch the kinds of errors that would aggregate into misleading datasets.

AI tools can open archival documents and resources that have not been accessible. For MENA researchers, this is a genuinely exciting moment, but like any other empirical method, the quality of what we build with these tools depends on how carefully we validate the outputs.

Robbins: As someone engaged in a project involving a large amount of data, I think one of the most exciting opportunities AI presents is the possibility of improving how we process, compare, and learn from the data we collect. A major challenge for Arab Barometer is managing this large volume of information. Survey research in the Middle East often involves substantial complexity. One of the Arab Barometer's central goals is to enable cross-national comparison while taking individual country contexts into account. Its ability to process vast amounts of data could help us more efficiently collect and analyze survey material that is often underused, especially open-ended responses, interviewer notes, and other paradata collected during interviews. To date, we have limited the number of open-ended response items in part because of the challenge of coding them on the back end. AI tools could also help us identify latent themes across countries and survey waves.

But these possibilities also come with challenges. The data we collect are produced through careful instrument design,

interviewer training, sampling decisions, and context-sensitive fieldwork. AI can assist with some parts of this process, but it cannot replace the detailed knowledge necessary to conduct survey research, use survey data, and accurately interpret the nuances of public opinion surveys. Doing so requires highly specialized knowledge, a detailed understanding of both the region and public opinion, and a strong grasp of the survey research process. My recommendation would be for those engaged in survey research or similar projects to use AI cautiously as a support for quality control and exploratory analysis. The key principle is that AI may help scale parts of the survey research process and make them more efficient, but it cannot replace the careful methodological labor that makes survey inference valid in the first place.

Bianco: One important opportunity AI presents for Middle East researchers lies in the ability to work at scale while remaining attentive to context. Research on the region has long been shaped by the dispersion of relevant material across languages, institutions, and extended time horizons. Political communication, legal texts, media coverage, and official statements are historically difficult to analyze systematically. As a consequence, research agendas have frequently followed what is most visible or readily accessible rather than what is most analytically significant (Anderson 2006).

Large language models (LLMs) make it possible to assemble and navigate large textual corpora without requiring early narrowing of the research question. They can summarize long documents, support semantic search across multilingual corpora, and identify recurring themes across thousands of texts. For policy archives, legal documents, and media databases, this substantially lowers the cost of

entry and enables more systematic exploration. It also creates space for more transparent workflows, as researchers can retain clear links between analytical claims and specific documents or passages. This allows scholars to begin with a broader view of the informational environment and then determine where close reading and deeper interpretation are most needed. Rather than relying primarily on small samples or emblematic cases, researchers can trace patterns of repetition, framing, and omission across time and across institutions. This is particularly valuable for the study of incremental change. For example, shifts in policy language or elite priorities often unfold gradually and can be difficult to detect through conventional sampling strategies.

Responsible use depends on treating AI as an assistive instrument rather than as an interpretive authority. Models can help identify where patterns appear stable and where they shift, but they cannot determine why those patterns matter. Responsibility lies in ensuring that interpretation and judgment remain grounded in regional knowledge, triangulation, and transparent documentation of how models were used. AI can expand the range of material that can be examined, but credible research still requires clear reasoning about what the data can and cannot show (Anderson 2006).

Kubinec: Arguably, the most useful aspect of LLMs for research in the MENA region is its ability to easily handle multilingual texts. As many researchers know, Arabic/Farsi/Hebrew/Turkish translation, especially at any scale, is no easy feat, and as a result, scholars tend to rely more on English-language sources about the region (Fahmy 2026; Nielsen and Zhou 2025; Di Bitetti and Ferreras 2017). LLMs' facility with translation, not just with

readable text but also more difficult formats like PDFs and images, can help render much of the non-Western world legible to those without a native-level language facility. Similarly, and perhaps more importantly, LLMs permit researchers without English fluency to prepare monographs for submission to English-language journals much more easily. Ideally, LLMs can provide a boost to the Arabic, Hebrew, Persian, and Turkish languages by reducing pressures to read and write exclusively in English—though I think scholars should still write in their own language first and translate the content rather than relying on LLMs to produce an article wholesale.

Of course, as mentioned previously, there is danger in relying on LLMs for the purposes of cognition. LLMs do not think and consequently cannot produce anything truly original. The risk of over-reliance on translation tools could degrade linguistic fluency, which is particularly worrying for scholars from the Anglophone world. My advice for scholars is to consider investing more time, not less, in reading and writing in the spoken languages of the region. As AI tools hopefully take over time-intensive tasks like data cleaning and rote translation, consider investing that time into more fieldwork, interviews, and exploratory reading. In other words, by employing LLMs to take on tasks that require little thought, scholars can (hopefully) find more time for open-ended research that involves creativity and higher-level thinking.

Another note of caution for scholars is to remember the risk of dependence on commercial tools as core parts of the research process. Part of the reason that we employed open-source LLMs for our study of Facebook ads was due to reproducibility problems with commercial providers like ChatGPT, which can suddenly disappear models (Barrie,

Palmer, and Spirling 2024). Scholars should have backup strategies for prompts, conversations, and output derived from any closed-source LLMs. Furthermore, it should go without saying that all LLM output should be checked—these are at bottom statistical models, and their flexibility means that they will always have some probability of providing incorrect output (i.e., hallucinations). Rigorous research still requires rigor, and only human beings can decide whether an analysis' claims are in fact valid.

Barnett: Some researchers, both within MENA and beyond, are exploring the possibility that “synthetic participants” (AI agents developed based on information about local populations' demographics) could be used to pilot surveys, test behavioral interventions, or otherwise facilitate research where access to populations is expensive, dangerous, or logistically difficult. To what extent do you think this is a promising or concerning development?

Adely: I consider “synthetic participants” a concerning development, especially given the opacity of AI models. The most prominent language models in circulation today (e.g., OpenAI's GPT models, Gemini, Claude) are built on closed-source, proprietary training data. Even some ostensibly “open-source” models disclose their code but not training datasets (e.g., Meta's LLaMA). Model makers might disclose the types of data used, such as large-scale scraping of text from the Internet. While there is little available data on Internet traffic and content produced from the MENA, it is telling that Arabic, Farsi, and Turkish only constitute around 3 percent of the Internet's linguistic landscape (W3Techs 2026). Second, the risks associated with this opacity are further accentuated by the common practice of using synthetic data for

contexts in which data is scarce or cost-prohibitive (Fitzgerald 2024). This would mean constructing synthetic participants on a foundation of synthetic data to design research instruments and projects to capture the experiences and perspectives of communities in the MENA. Finally, even in an (unlikely) scenario where a researcher could review and evaluate training data, the mechanisms through which models train are not fully explainable. This is not to suggest that AI technologies are a complete black box; as Salvaggio (2025) notes, “they associate words in a vast vector space, and we can trace likely word pairings. What remains opaque is how specific correlations are inferred from vast training data.” Together, these compounding layers of opacity make the prospect of integrating synthetic participants into research alarming.

Emerging technologies do offer possibilities for conducting research among populations and locales where fieldwork raises safety and ethical concerns. The COVID-19 pandemic prompted further development and adoption of digital methods from remote interviewing to online ethnographies. While they expand possibilities for research, these methods certainly bring with them ethical, political, and methodological issues; yet, I consider them more promising than “synthetic participants” as a way of connecting with and adapting to the realities of our field sites and communities.

Bianco: The use of synthetic participants is promising in limited and clearly defined ways. As tools for piloting surveys, testing question wording, or identifying implicit assumptions in research design, they can help researchers avoid basic errors before engaging real populations. In this role, they can improve research design by revealing ambiguity, cultural misalignment, or unintended leading

effects in prompts. They may also help anticipate how different framings could be received before significant resources are committed to fieldwork.

Concerns arise when synthetic participants are treated as substitutes for empirical engagement. These agents do not represent lived experience. They reproduce patterns present in existing data. In the Middle East, where documentation is uneven and often filtered through official or elite channels, this limitation is particularly consequential. Groups that are hardest to access are often those least represented in training data, meaning that synthetic participants may appear most confident precisely where their grounding is weakest. Using synthetic participants to generate substantive findings, therefore, risks reinforcing existing blind spots. Simulation can reproduce the assumptions embedded in available texts rather than the realities of populations whose voices are limited or strategically absent. For this reason, synthetic participants are best confined to instrument development rather than treated as evidence.

Robbins: The use of synthetic respondents might have some value for tasks such as checking the logic of a pilot survey or running other basic simulations. However, I would not advise using such agents for the other purposes mentioned here. Arab Barometer’s experience underscores that public opinion in the region cannot be reduced to demographic profiles. What respondents say is shaped not only by age, gender, education, or place of residence, but also by political context, trust, timing, interviewer effects, local conditions, and perceptions of risk. In any setting, but especially in MENA, respondents may conceal preferences, offer strategically cautious answers, or interpret politically sensitive questions in context-dependent ways.

Synthetic participants are poorly equipped to capture these dynamics precisely because they are built from generalized patterns rather than actual realities on the ground.

Having worked on Arab Barometer for the last two decades, I continue to be struck by the extent to which the data still produce major surprises. Commonly, correlations we expect to find do not appear, while other unexpected correlations emerge, or a long-standing trend reverses. Frequently, this happens not only in a single country but across the entire wave. As we dig more deeply into the data, we arrive at a much more nuanced understanding of political behavior in MENA and how ordinary people understand their worlds. Without measuring these concepts with human participants, we would surely miss these important findings. This reality points to the methodological danger that synthetic participants will inherit the blind spots of the data used to build them. Results from Arab Barometer are collected in a specific context at a specific time. We design questions to be less sensitive to day-to-day events, but we cannot fully prevent current events from influencing responses of those interviewed. Building synthetic populations from existing survey data or other measures of public opinion will likely replicate existing distortions. Moreover, given how limited the available data on populations living in MENA are compared with many other regions, existing biases may be especially misleading and could yield results that are counterproductive for researchers. At this point, I do not believe that the potential benefits outweigh the potential costs of using synthetic participants to try to measure public opinion in MENA.

Barnett: To what extent does the use of AI risk biasing inferences about the region toward more “data-rich” and AI-forward con-

texts like the GCC? What could be done to mitigate this bias?

Abdo: Aggregation bias could mean that the MENA region’s AI-driven future may be disproportionately shaped by a few data-rich hubs. This disparity of resources is confirmed by The Government AI Readiness Index, which ranks countries based on their capacity to deploy AI in public services, which notes that MENA has the widest performance gap of any global region. For example, the United Arab Emirates scores 72.36 and Saudi Arabia 71.57, while Syria and Yemen trail at 21.72 and 14.48, respectively. This clear divide reflects profound inequalities in infrastructure and digital investment (Oxford Insights 2025). This unevenness is further reinforced by the sheer scale of investment flowing into the Gulf Cooperation Council (GCC). Saudi Arabia’s \$40 billion HUMAIN initiative and the UAE’s backing of ventures such as MGX and G42, reflect a level of capital commitment that few other states in the region can match.

Bianco: Evaluating AI systems trained on GCC-related data gives away their limitations immediately to the expert eye. These systems tend to privilege environments that are digitally legible and well-documented. Within the Middle East, this creates a structural pull toward the Gulf, where official communication is abundant, standardized, and frequently available in global languages. As a result, AI-assisted research risks overrepresenting these contexts while underweighting others. This shows up most clearly in their interpretation of stability and risk. A system exposed primarily to large volumes of consistent, institutionally produced text may learn that stability is the baseline and that political change is unlikely, not because it has understood political dynamics, but because the

data it encounters is unusually uniform. Models trained in these environments often underestimate the possibility of change.

Preventing false confidence requires interventions at both the training and evaluation stages. Training data should be diversified to include sources that differ structurally from official narratives and that introduce variation across time, topic, and institutional voice. Evaluation should explicitly test for overconfidence, including whether models can flag ambiguity and avoid treating absence of evidence as evidence of absence. In high-stakes applications, human review remains essential when model confidence is high, but the underlying data stream is narrow or tightly controlled. In such settings, confidence itself should be treated as a potential risk signal rather than reassurance (King, Pan, and Roberts 2013; Xia et al. 2025).

Furthermore, researchers can deliberately incorporate lower-visibility sources and treat gaps and silences as analytically meaningful. What remains crucial is to be explicit about the scope of inference. When conclusions are driven primarily by data-rich contexts, this should be acknowledged as a limitation and, where relevant, interpreted as part of the region's political economy of information (Anderson 2006).

Barnett: What is one thing that concerns you about how AI might impact research in our field? What can be done to address this concern?

Adely: In the past several years, I have noticed a concerning trend of conflating technology's multilingual capabilities with its capacity to reflect diverse viewpoints. If a model can produce Arabic (or Farsi, Kurdish, and Turkish for that matter), the logic goes,

then the model can speak to cultural, social, and political life in the Arab world. In many ways, this trend reflects a broader issue with generative AI: the ability to generate language that convincingly mimics human communication becomes a proxy for claims about a model's intelligence, understanding, and reasoning (Bender and Koller 2020).

The conflation of language with meaning, understanding, and/or knowledge is especially fraught in the context of "low-resource" languages, linguistic varieties for which there exist comparatively few datasets and computational resources. For example, in the early months of my fieldwork in fall 2023, a recurrent topic of conversation with my interlocutors was the recent release of an LLM by a UAE-based company, which promoted their product as the most advanced model to date for Arabic. At tech forums and conferences across Amman, a regular refrain was the need for "Arab AI" that reflected Arab values, culture, and political-ethical commitments, as opposed to the "Western" perspective perceived to be embedded in AI tools like ChatGPT.

This Emirati model seemed poised to address these calls. Trained on a bilingual dataset (Arabic and English), the company's promotional materials boasted about training the model on over 100 billion Arabic tokens, or roughly one-third of the total training data. As I tested the model—on my own and alongside engineers and other workers with AI expertise—I could appreciate the model's ability to produce grammatically sound and contextually relevant statements in Arabic. (In 2023, just a year after ChatGPT's initial public release, the ability to effectively generate Arabic text was still quite novel). However, the Arabic-language data on which the model was trained was likely—at least par-

tially—machine-translated (see Sadallah et al. 2025 on the pitfalls of machine-translated data for Arabic). The anecdotes, references, and names in the model's outputs, while communicated in Arabic, suggested a reliance on data that was originally in English, reflecting a different set of social, cultural, and political assumptions. In sum, this conflation of localized language capabilities with localized understanding and meaning does not reflect how AI tools function, and this disparity is especially pronounced for the language contexts in which scholars of MENA work.

In terms of addressing this concern, there is no simple solution. At one level, it entails a critical eye toward information produced by AI models, even those with the capacity to produce the language varieties used in our field sites. More broadly, as AI technologies become more thoroughly imbricated in our field sites, I see the need for a renewed commitment to methods and methodologies that consider how utterances and symbols can reflect a wide range of meanings (Wedeen 2002). Just as we remain attentive to the gap between what research participants say and what they do or believe, we must bring a critical eye toward the possibility that participant responses, institutional statements, and publicly available data may be mediated by AI tools in ways that obscure social and political realities.

Bianco: A central concern is that AI may increase the risk of false confidence, as I noted above. LLMs are optimized to produce fluent and internally coherent outputs. When applied to political text, that fluency can be mistaken for accuracy or completeness. Consistency in the available data does not necessarily correspond to genuine consensus. This concern echoes long-standing debates in Middle East politics. Apparent stability in

authoritarian systems has often reflected active management rather than deep agreement, and periods of robustness have obscured underlying vulnerabilities (Bellin 2012). Moreover, political meaning frequently emerges from euphemism, selective emphasis, or omission. Dialectal variation and context-specific references further complicate automated parsing. Models may process such material smoothly while missing what experienced readers recognize as the central signal. In addition, many regional corpora are shaped by uneven digitization, inconsistent metadata, and shifting platform practices, which can introduce hidden measurement error.

Addressing these concerns requires methodological caution. Automated text analysis requires validation and human judgment to support credible inference (Grimmer and Stewart 2013). In practice, this involves assessing whether results hold across different source types, languages, and time periods, and being explicit about uncertainty. Recent AI research reinforces this point by showing that models frequently produce confident outputs even when they are wrong, and that verbal confidence should not be equated with reliability (Xia et al. 2025; Groot and Valdenegro-Toro 2024). For politically sensitive research, a useful discipline is to treat model confidence as a signal for further scrutiny rather than as a conclusion. Embedding these practices into research design can allow AI to support more careful scholarship rather than more persuasive errors (Grimmer and Stewart 2013).

Robbins: My major concern is that we will see an overreliance on AI to do work that is best left to a careful researcher. Good research is difficult and tedious and requires extremely detailed attention. There is a

natural tendency to try to speed up this process, particularly when working with large and complex datasets. In the case of public opinion research in particular, the use of AI could be especially troubling because much of what appears straightforward in a dataset is actually fragile and conditional. Responses depend on question wording, interview context, social desirability pressures, political conditions, and whether respondents trust that their answers are confidential. These issues are important in any context, but in many parts of the Middle East and North Africa, they are major considerations and are central to how survey responses should be interpreted.

Even before the rise of AI, I saw many analyses that did not sufficiently take these factors into account. I am deeply concerned that AI will make this problem more common by making data processing and analysis appear more straightforward than they really are. Researchers may be tempted to use models to classify attitudes, infer ideology, or summarize open-ended responses without adequately considering the conditions under which those responses were produced. The result could be a larger number of analyses using public opinion data, but also a larger number that misinterpret the results or are based on spurious correlations.

Ultimately, researchers cannot simply outsource research to AI. AI will not be able to compensate for poor questionnaire design, a weak sample design, translation problems, or a lack of understanding of the survey research process. As a result, researchers should be cautious in their use of AI and use it only to facilitate research, as opposed to expecting it to do the research itself. Most importantly, researchers must think carefully through research design and the patterns most likely

to account for the distributions they observe in the data. For example, if AI identifies what appears to be a novel finding in public opinion, that pattern must be checked very carefully against established measures, country expertise, and, where possible, subsequent survey rounds or qualitative evidence, just as it should be if AI had not been used. However, given the vast amounts of data that AI can process, the risks are far greater than before of generating meaningless correlations that could be advanced as new theories to explain attitudes and behavior in MENA. ♦

References

Abdo, Muhammad S., Yash Hatekar, and Damir Cavar. 2025. "AMWAL: Named Entity Recognition for Arabic Financial News." In *Proceedings of the Joint Workshop of the 9th Financial Technology and Natural Language Processing (FinNLP), the 6th Financial Narrative Processing (FNP), and the 1st Workshop on Large Language Models for Finance and Legal (LLMFinLegal)*, 207-213. Abu Dhabi, UAE: Association for Computational Linguistics. <https://aclanthology.org/2025.finnlp-1.20/>

Ahmed, Abdel Rahman. 2026. "Ashraf El-Shihy told Al-Masry Al-Youm: Sisi's directives regarding university specializations aim to link education with the labor market." *Al-Masry Al-Youm*, March 6. <https://www.almasryalyoum.com/news/details/4215356>

Ahram Online. 2023. "University Education Needs to Be Integrated with Job Market Needs: Sisi." *Ahram Online*, September 26. <https://english.ahram.org.eg/News/509069.aspx>

Anderson, Lisa. 2006. "Searching Where the Light Shines: Studying Democratization in the Middle East." *Annual Review of Political Science* 9: 189–214.

Barrie, Christopher, Alexis Palmer, and Arthur Spirling. 2024. "Replication for Language Models Problems, Principles, and Best Practice for Political Science." http://arthurspirling.org/documents/BarriePalmer-Spirling_TrustMeBro.pdf

Bellin, Eva. 2012. "Reconsidering the Robustness of Authoritarianism in the Middle East: Lessons from the Arab Spring." *Comparative Politics* 44(2): 127–149. <https://doi.org/10.5129/001041512798838021>

Bender, Emily M., and Alexander Koller. 2020. "Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data." In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. Association for Computational Linguistics. doi:10.18653/v1/2020.acl-main.463

Borges, Jorge Luis. 1941. *The Library of Babel*.

Calderon, N., A., et al. 2024. "Measuring the Robustness of NLP Models to Domain Shifts." *Findings of EMNLP* 2024.

Champion, Zach. 2023. "Optimization Could Cut the Carbon Footprint of AI Training by Up to 75%." *Michigan Engineering News*. April 17. University of Michigan. <https://news.engin.umich.edu/2023/04/optimization-could-cut-the-carbon-footprint-of-ai-training-by-up-to-75/>

Di Bitetti, Mario S., and Julián A. Ferreras. 2017. "Publish (in English) or Perish: The Effect on Citation Rate of Using Languages Other Than English in Scientific Publications." *Ambio* 46(1): 121–27. <https://doi.org/10.1007/s13280-016-0820-7>

Elshehawy, Ashrakat, Arun Frey, Violeta Haas, Sascha Riaz, and Tobias Roemer. 2025. "The Police as Gatekeepers of Information: Immigration Salience and Selective Crime Reporting." Working Paper.

Fahmy, Khaled. 2017. "The Crisis of the Humanities in Egypt." *Comparative Studies of South Asia, Africa and the Middle East* 37(1): 142–148.

Fahmy, Shahira S. 2026. "The Orphan Scholars: Communication Research in the Middle East and North Africa." *Journalism & Communication Monographs* 28(1): 12–17. <https://doi.org/10.1177/15226379251409514>

Fitzgerald, Andrew. 2024. "Why Synthetic Data Can Never Be Ethical: A Lesson from Media Ethics." *Surveillance & Society* 22(4): 477–482. <https://doi.org/10.24908/ss.v22i4.18324>

Gibney, Elizabeth. 2024. "Has Your Paper Been Used to Train an AI Model? Almost Certainly." *Nature* 632(8026): 715–16. doi:10.1038/d41586-024-02599-9

Gilardi, Fabrizio, Meysam Alizadeh, and Maël Kubli. 2023. "ChatGPT Outperforms Crowd Workers for Text-Annotation Tasks." *Proceedings of the National Academy of Sciences* 120(30): e2305016120.

Gray, Mary L. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. New York: HarperCollins Publishers.

- Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3): 267–297.
- Groot, Tobias, and Matias Valdenegro-Tooro. 2024. "Overconfidence Is Key: Verbalized Uncertainty Evaluation in Large Language and Vision-Language Models." *arXiv*.
- Gupta, Shloak, Sarah Bolden, Jay Kachhadia, A. Korsunskaya, and J. Stromer-Galley. 2020. "PoliBERT: Classifying Political Social Media Messages with BERT." <https://news.illuminating.ischool.syr.edu/2020/11/24/polibert-classifying-political-social-media-messages-with-bert/>
- Halterman, Andrew, and Katherine A. Keith. 2025. "Codebook LLMs: Evaluating LLMs as Measurement Tools for Political Science Concepts." *Political Analysis*. <https://doi.org/10.1017/pan.2025.10017>
- Heakl, Ahmed, Muhammad Abdullah Sohail, Mukul Ranjan, Rania Elbadry, Ghazi Shazan Ahmad, Mohamed El-Geish, Omar Maher, Zhiqiang Shen, Fahad Shahbaz Khan, and Salman Khan. 2025. "KITAB-Bench: A Comprehensive Multi-Domain Benchmark for Arabic OCR and Document Understanding." In *Findings of the Association for Computational Linguistics: ACL 2025*, 22006–22024.
- Hodapp, James. 2022. "The Crisis in the Humanities: A Perspective from the Middle East." In *Asian English: Histories, Texts, Institutions*, 165–184. Springer.
- Jackson, Tom, and Ian R. Hodgkinson. 2022. "What Is 'Dark Data' and How Is It Adding to All of Our Carbon Footprints?" *World Economic Forum*. October 5. <https://www.weforum.org/stories/2022/10/dark-data-is-killing-the-planet-we-need-digital-decarbonisation/>
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107(2): 326–343.
- Kirk, Richard. 2025. "For a Political Geography of Artificial Intelligence: Fighting Ghost Work, Exploitation, and the Making of a Global Digital Underclass." *Antipode Online*, July 30. <https://antipodeonline.org/2025/07/30/for-a-political-geography-of-ai/>
- Kubinec, Robert, and Aseem Mahajan. 2025. "Fifty Shades of Greenwashing: The Political Economy of Climate Change Advertising on Social Media." *arXiv*. <https://doi.org/10.48550/arXiv.2511.14930>
- Lockman, Zachary. 2010. *Contending Visions of the Middle East: The History and Politics of Orientalism*. Second Edition. New York: Cambridge University Press.
- Mamdouh, Mohamed. 2026. "Sisi calls for a review of university majors not in demand in the job market: Years should not be wasted without a professional future." Cairo 24, March 5. <https://www.cairo24.com/2383314>

Naous, Tarek, Michael J. Ryan, Alan Ritter, and Wei Xu. 2024. "Having Beer after Prayer? Measuring Cultural Bias in Large Language Models." In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, Volume 1: Long Papers, 16366–16393.

National Institute of Standards and Technology. 2023. *AI Risk Management Framework (AI RMF 1.0) Playbook*. U.S. Department of Commerce. <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>

Navigli, Roberto, Simone Conia, and Björn Ross. 2023. "Biases in Large Language Models: Origins, Inventory, and Discussion." *ACM Journal of Data and Information Quality* 15(2): 1–21.

Nielsen, Richard A., and Annie Yiwen Zhou. 2025. "Integrating Social Science Research Across Languages with Assistance from Artificial Intelligence." *Daedalus* 154(2): 51–67. https://doi.org/10.1162/daed_a_02140

Oxford Insights. 2025. *Government AI Readiness Index 2025*. <https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/>

Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58(4): 1064–82.

Rodriguez, Pedro L., and Arthur Spirling. 2022. "Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research." *The Journal of Politics* 84(1): 101–15. <https://doi.org/10.1086/715162>

Rodriguez, Pedro L., Arthur Spirling, and Brandon M. Stewart. 2023. "Embedding Regression: Models for Context-Specific Description and Inference." *American Political Science Review* 117(4): 1255–74. <https://doi.org/10.1017/S0003055422001228>

Sadallah, Abdelrahman, Junior Cedric Tonga, Khalid Almubarak, Saeed Almheiri, Farah Atif, Chatrine Qwaider, Karima Kadaoui, et al. 2025. "Commonsense Reasoning in Arab Culture." In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, Volume 1: Long Papers, 7695–7710. Association for Computational Linguistics. <https://aclanthology.org/2025.acl-long.380/>

Said, Edward W. 1979. *Orientalism*. First edition. New York: Vintage Books.

Salvaggio, Eryk. 2025. "The Black Box Myth: What the Industry Pretends Not to Know About AI." *Tech Policy Press*, June 17. <https://techpolicy.press/the-black-box-myth-what-the-industry-pretends-not-to-know-about-ai>

Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. "Energy and Policy Considerations for Deep Learning in NLP." *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. <https://aclanthology.org/P19-1355.pdf>

Tao, Yan, Olga Viberg, Ryan S. Baker, and René F. Kizilcec. 2024. "Cultural Bias and Cultural Alignment of Large Language Models." *PNAS Nexus* 3(9): 346.

Törnberg, Petter. 2025. "Large Language Models Outperform Expert Coders and Supervised Classifiers at Annotating Political Social Media Messages." *Social Science Computer Review* 43(6): 1181–1195.

UNESCO. 2022. *Recommendation on the Ethics of Artificial Intelligence*. Paris, France: United Nations Educational, Scientific and Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>

W3Techs (Web Technology Surveys). "Usage Statistics and Market Share of Content Languages for Websites, March 2026." *Technologies, Content Languages*. https://w3techs.com/technologies/overview/content_language

Wedeen, Lisa. 2002. "Conceptualizing Culture: Possibilities for Political Science." *American Political Science Review* 96(4): 713–28. <https://doi.org/10.1017/S0003055402000400>

Wu, Carole-Jean, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, et al. 2022. "Sustainable AI: Environmental Implications, Challenges and Opportunities." In *Proceedings of Machine Learning and Systems*, eds. D. Marculescu, Y. Chi, and C. Wu., 795–813. <https://doi.org/10.48550/arXiv.2111.00364>

Xia, Z., Y. Wen, J. Duan, and Y. et al. 2025. "A Survey of Uncertainty Estimation Methods on Large Language Models." *Findings of the ACL* 2025.

AI Will Drive the Coming Integration of Interpretive and Positive Methodologies¹

Richard A. Nielsen



Richard A. Nielsen is an Associate Professor in the Department of Political Science at the Massachusetts Institute of Technology, where he studies and teaches on Middle East and North African politics, International Relations, religion, gender, political violence, and research methods. Nielsen is the developer of free software tools for Arabic text analysis, causal inference, and qualitative case selection. At MIT, Nielsen is the faculty director of the MIT-MENA program and holds affiliations with the Center for International Studies, the Security Studies Program, the Institute for Data, Systems, and Society, the Political Methodology Lab, and the IPE group.

In the past few months, from December 2025 to March 2026, I have come to believe that generative Artificial Intelligence, colloquially called Gen-AI, will fundamentally change much of how we produce human knowledge. I've been using machine learning – the basis for artificial intelligence – in my research for years, and I've been skeptical that it will transform the social sciences. No longer. Catastrophic success² of colleagues executing years-long quantitative research projects in hours has changed my mind (Messing and Tucker 2026, Kustov 2026).

Generative AI is generative. It can be trained to work on improving itself and its analyses, so it is gaining function faster than our norms and institutions can adapt. If Gen-AI is currently at a level where it can perform the tasks of a sophisticated research assistant, it will very soon be able to perform most well-defined analytical tasks that I am able to perform. This kind of automation has been a goal of quantitative research all along. For

good reasons, we don't trust human intuition about probability, so when we want to make decisions that involve probability, we incorporate a series of algorithms: rules for how one ought to update one's beliefs given new information, and how one should make decisions, given those updated beliefs. Our graduate training in the quantitative social sciences has been oriented towards helping colleagues learn a common set of algorithms (and invent new ones based on the same normative principles, when we uncover problems for which no satisfactory algorithm has yet been created). Now, Gen-AI makes so much more of it automatable that the future of quantitative research will not be coding, but rather prompting AI agents to collect data and analyze it according to a comprehensive synthesis of a statistics literature so vast that no human can ever read it in a lifetime and produce answers that would have previously taken years or lifetimes to achieve in mere days, or even minutes.

¹ Acknowledgements: Thanks to Bernardo Zacka, In Song Kim, Naoki Egami, Ella Pace, Elizabeth Wright, Dominic Skinnion, Diana Fu, Ed Schatz, and Andrew Miller for conversations that shaped my thinking, and to my students for showing me a way forward. I wrote this without use of generative AI.

² Catastrophic, because the short-term success wildly exceeds expectations, but the long-term implications may be very negative. See Downes (2021).

Augmenting researcher time with algorithms isn't new. In my first book, *Deadly Clerics*, I collected 150,000 documents and had a rudimentary AI model "read" them because it would have taken decades to read them myself, synthesize them, and test my argument that Sunni Muslim clerics are more likely to preach anti-system jihad on the Internet when their academic ambitions are blocked (Nielsen 2017). It still took me two years to do a pilot study, and another 3 years to build the full study, write it, and publish it. Could this project be done overnight now? I don't know, but we are close. If data is on the open web, an AI agent can access it, collect it, analyze it, and produce findings in minutes. What is left for me to do? Why should someone pay me to do it? We are facing a crisis of economics in knowledge-making.

The theoretical claim of *Deadly Clerics*, then, is so relevant that it's a bit on the nose. When would-be academics have their ambitions blocked, they reject the system. If we use AI to eliminate many entryways to a research career, the junior scholars whose academic ambitions are inadvertently blocked will reject our discipline and take their talents elsewhere. We must think as a discipline about a fair transition to this new research world. Automating much of the research process will affect us all, but the most abrupt and painful consequences may hit junior researchers and those in precarious positions. Research workflows are becoming obsolete every month right now. How can a graduate student write a dissertation prospectus this Spring if the project will take three years with today's technology but will take only days to complete once they are on the job market? Will search committees reward them for time well spent, or ask, "shouldn't you have more?"

What, then, should a graduate student, writ-

ing a dissertation proposal this year, propose to do with their time? Rather than emphasizing the most algorithmic parts of what we do, the future work of this graduate student needs to be focused on what they, the human researcher, are making, creating, and contributing to new knowledge that is not reducible to an algorithm. I believe that human research has an aesthetic quality in addition to its technical correctness. I've come to understand myself and my own research as interpretivist, even as it draws deeply and faithfully from the canon of positivist methodology in many respects.

I recommend that a student starting now should adopt what Ed Schatz calls an ethnographic sensibility (Schatz 2009). This ethnographic sensibility is unique, not reproducible with AI, and developed from close, embodied interaction with the world, at the "nearest possible vantage-point" (Schatz 2009, p. 307). This is what is uniquely human about research. Human researchers come to their projects embodied and endowed with subjective agency: making meaningful choices about what questions to ask, how to ask them, and how to interpret the meaning of the answers for themselves and for the world.

When algorithms can do all the "objective" work, what is left is *interpretive* inference. I have encountered this most obviously in quantitative text analysis as I deploy machine learning models to infer the meaning(s) of a certain text, intended or not, by their authors and their audiences. These are problems of interpretive inference, and I believe that interpretive inference will remain a fundamentally human enterprise. I will trust a series of AI agents to "read" the statistics literature and design a causal study to test a correlation, long before I will trust an AI agent to tell me how humans make meaning out of what they

write, even when I use algorithms to summarize that writing in service of my interpretation. This thought experiment reveals, to me at least, that the part of research I am most confident will remain important in the age of AI is the interpretive. Because of the causal revolution's success at standardizing causal analyses, very few applied researchers will spend much of their time on solving causal inference problems. (We will still need methodologists to help us understand new types of causal inference problems.)

A third type of inference, descriptive inference, will remain important for Political Science. To be analyzed by Gen-AI, data that is not connected to the Internet must be connected. Here, humans will have an advantage for a long time. The archives will resist full digitization; there will be value in going in person to find the boxes that were missed in the rush to put "everything" online. Despite advances in using video oral histories (Milliff 2024) and virtual reality surveys (Miller 2024) to collect the experiences of humans, I believe there will still be an important role for human-centered interviewing and survey enumeration, following an interpretive, relational approach (Fujii 2017).

Other forms of description and measurement with an ethnographic sensibility, such as Cramer's gas-station and diner focus groups in Wisconsin (Cramer) or Wood's mapmaking exercises in El Salvador (Wood 2003), will increase in importance, even when they support quantitative findings, as in Jones' (2017) palace ethnography (paired with experiments) to understand social order in the UAE or Parker-Magyar's (2024) ethnographic observation to understand teacher protest in Jordan (paired with a network survey).

With coauthors, Diana Fu and Ed Schatz (2026), I have argued that in the age of AI, an ethnographic sensibility is more important, not less. Don't misunderstand me. I think causal inference is crucial, but automatable. Given fixed meanings and defined sets of relevant variables, an algorithm can (and should) produce the estimates we use for causal inference. This is the core of the "credibility revolution." When Don Rubin referenced "knowing the science,"³ he meant that the algorithms, which he pioneered, will only work for estimation if a researcher is willing to say, "here's the structure of my problem." However, that prior knowledge of the structure cannot be specified by Rubin's algorithms. It comes from a human understanding of the problem, what is previously known about the problem, and what we desire to know now about the problem. That's "the science." And so, I think, the task for human researchers in the age of AI will be to know the science. Quantitative analyses will become democratized. As almost anyone can access quantitative tools via adequate prompting, scholars will spend years honing their sensitivity to the structure of the social world to ask better questions and offer new interpretations.

Let me return to our ethical obligations as individuals and a discipline. My work has been so disrupted by advances in AI that I am rethinking the entire workflow I've relied on to remain gainfully employed for the last 15 years. I have the luxury of doing so with a secure academic position. Most researchers are rethinking, or will rethink, their workflow from a place of precarity. I feel the responsibility to think hard about what we owe to future generations of scholars and take concrete steps to create the discipline we want them to inherit. As in many industries, agentic AI is

³ Personal communication.

on the cusp of replacing entry-level research assistance, which will sever the pipeline that got me and so many dear colleagues into Political Science research. The project that made me fall in love with social science was scanning and coding Development Bank documents, a task that is completely automatable now. If we pull up the ladder to new entrants in a short-sighted, self-interested arms race to use AI to publish more articles in top journals for the next decade, then our field will cut itself off at the knees. Each researcher has a unique capacity to see the world through their own eyes, to hone that into a sensibility unique to them, that they use to teach us all about the infinite facets of human societies. We must step back from an AI-fueled publication arms race to intentionally choose discipline-wide institutions and norms that make the social sciences an enterprise worth entering, no matter how much we must transform for that to be true. We must do this so that we humans can keep learning about ourselves.♦

References

Cramer Katherine J. 2016. *The Politics of Resentment: Rural Consciousness in Wisconsin and the Rise of Scott Walker*. Chicago: University of Chicago Press.

Diana Fu, Richard A. Nielsen, Edward Schatz. 2026. "Ethnography and Ethnographic Sensibility in Political Science," *Annual Review Political Science*. 29.

Downes, Alexander B. 2021. *Catastrophic Success: Why Foreign-Imposed Regime Change Goes Wrong*. Ithaca: Cornell University Press.

Fujii, Lee Ann. 2017. *Interviewing in Social Science Research: A Relational Approach*. New York: Routledge. 1st ed.

Jones, Calvert W. 2017. *Bedouins into Bourgeois: Remaking Citizens for Globalization*. Cambridge: Cambridge University Press.

Kustov, Daniel. 2026. "Academics Need to Wake Up on AI." By *Popular Design*, March 02. <https://www.popularbydesign.org/p/academics-need-to-wake-up-on-ai>

Messing, Solomon, and Joshua Tucker. 2026. "The train has left the station: Agentic AI and the Future of Social Science Research." *Brookings*. March 3. <https://www.brookings.edu/articles/the-train-has-left-the-station-agentic-ai-and-the-future-of-social-science-research/>

Miller, Andrew C. 2024. *Silencing Citizens: How Criminal Groups Create Vacuums of Justice*. Cambridge: Cambridge University Press.

Milliff, Aidan. 2024. "Making sense, making choices: How civilians choose survival strategies during violence." *American Political Science Review*, 118, 3: 1379-1397.

Nielsen, Richard A. 2017. *Deadly Clerics: Blocked Ambition and the Paths to Jihad*. Cambridge: Cambridge University Press.

Parker-Magyar, Elizabeth. 2024. *Workplace Networks and Autonomous Organizations in Contemporary Jordan*. PhD diss. Massachusetts Institute of Technology.

Schatz Ed, ed. 2009. *Ethnographic Immersion and the Study of Politics*. Chicago: University of Chicago Press.

Wood Elizabeth J. 2003. *Insurgent Collective Action and Civil War in El Salvador*. Cambridge: Cambridge University Press.

MENA POLITICS

Newsletter of the Middle East and North Africa Politics Section of APSA

VOL. 9 ISSUE 1, SPRING 2026

ISSN: 2996-4911

MISSION STATEMENT

The Middle East and North Africa Politics Section is an academic society at the American Political Science Association (APSA), established in 2018 to support, develop and publish research on the politics of the MENA region utilizing interdisciplinary methodological, theoretical and empirical tools. It seeks to fully integrate the rigorous study of the politics of the Middle East with the broader discipline of Political Science, to serve as an institutional home for the community of political scientists dedicated to the Middle East, and to incorporate both scholars from the MENA region and diverse scholars from the United States into the global study of Middle East politics. The section's bylaws and diversity statement can be found on its website, <https://apsamena.org>.

MENA Politics is the official newsletter of the APSA-MENA section. It is a biannual publication devoted to publishing new research findings, hosting productive debates, and highlighting noteworthy developments among the scholarly community. It is managed by the chief editors with the joint confidence of the editorial board and the APSA-MENA section.

NEWSLETTER STAFF

EDITORS:

Diana B. Greenwald, Sebnem Gumuscu, Samer Shehata

EDITORIAL BOARD:

Carolyn Barnett, Sharan Grewal, Selin Bengi Gümrükçü,
Ian M. Hartshorn, Rola al-Husseini, Eric Lob,
Amanda Rizkallah

EDITORIAL ASSISTANT:

Sussan Saikali-Wehbe

MENA SECTION OFFICERS

Marwa Shalaby (Chair)
Sarah El Kazaz (Vice Chair)
Adria Lawrence (Treasurer)
Nermin Allam (At-Large)
Lisel Hintz (At-Large)

CONTACT

Email: apsamena@gmail.com

APSA Section Website:

www.apsanet.org/section49

Public Website: www.apsamena.org

Twitter: @APSAMENA

Facebook: www.facebook.com/APSAMENA

Blue Sky: <https://bsky.app/profile/apsamena.bsky.social>