

Towards Ethical and Effective **AI-Augmented** Research Writing

*A Conceptual Framework
and Practical Guidelines*

Prof. Junaid Qadir

College of Engineering · Qatar University · jqadir@qu.edu.qa

Two Mulla Nasruddin Stories to Start With: The Optimized Donkey



Mulla Nasruddin bought a donkey but found its food too expensive. "I will train it to eat less," he decided.

Each day, he reduced the ration. The donkey continued to work, appearing fine.

With Mulla giving it almost no food, the donkey eventually died.

Two Mulla Nasruddin Stories to Start With: The Optimized Donkey



Mulla Nasruddin bought a donkey but found its food too expensive. "I will train it to eat less," he decided.

Each day, he reduced the ration. The donkey continued to work, appearing fine.

With Mulla giving it almost no food, the donkey eventually died.

"What a pity," said Nasruddin. "If I had just a little more time, I had almost taught it to live without any food at all."

Two Mulla Nasruddin Stories to Start With: The Optimized Donkey



Mulla Nasruddin bought a donkey but found its food too expensive. "I will train it to eat less," he decided.

Each day, he reduced the ration. The donkey continued to work, appearing fine.

With Mulla giving it almost no food, the donkey eventually died.

"What a pity," said Nasruddin. "If I had just a little more time, I had almost taught it to live without any food at all."

The Parable of the Optimized Donkey

When we optimize away the productive struggle of writing, we optimize away the formation.

In education and AI: are we the analogue of this?

Two Mulla Nasruddin Stories to Start With: The Optimized Donkey



Mulla Nasruddin bought a donkey but found its food too expensive. "I will train it to eat less," he decided.

Each day, he reduced the ration. The donkey continued to work, appearing fine.

With Mulla giving it almost no food, the donkey eventually died.

"What a pity," said Nasruddin. "If I had just a little more time, I had almost taught it to live without any food at all."

The Parable of the Optimized Donkey

When we optimize away the productive struggle of writing, we optimize away the formation.

In education and AI: are we the analogue of this?

In chasing efficiency, we may be starving the very process that makes learning come alive.

Two Mulla Nasruddin Stories to Start With: The Subtle Art of Imitation

During an exam, Mulla Nasruddin sat next to a bright student who whispered, *“Mulla, let’s copy each other so we both pass.”*

Nasruddin nodded. After the exam, the results came:

The student got zero, and Nasruddin got full marks.

The furious student shouted, “But I copied every single answer from you!”

Two Mulla Nasruddin Stories to Start With: The Subtle Art of Imitation

During an exam, Mulla Nasruddin sat next to a bright student who whispered, *“Mulla, let’s copy each other so we both pass.”*

Nasruddin nodded. After the exam, the results came:

The student got zero, and Nasruddin got full marks.

The furious student shouted, “But I copied every single answer from you!”

Nasruddin smiled, “Yes, my friend—but you also copied my name.”

Two Mulla Nasruddin Stories to Start With: The Subtle Art of Imitation

During an exam, Mulla Nasruddin sat next to a bright student who whispered, *“Mulla, let’s copy each other so we both pass.”*

Nasruddin nodded. After the exam, the results came:

The student got zero, and Nasruddin got full marks.

The furious student shouted, “But I copied every single answer from you!”

Nasruddin smiled, “Yes, my friend—but you also copied my name.”

The AI Analogue: When people outsource research and writing to AI, they end up submitting **AI’s work—and AI’s name.**

Two Mulla Nasruddin Stories to Start With: The Subtle Art of Imitation

During an exam, Mulla Nasruddin sat next to a bright student who whispered, *“Mulla, let’s copy each other so we both pass.”*

Nasruddin nodded. After the exam, the results came:

The student got zero, and Nasruddin got full marks.

The furious student shouted, “But I copied every single answer from you!”

Nasruddin smiled, “Yes, my friend—but you also copied my name.”

The AI Analogue: When people outsource research and writing to AI, they end up submitting **AI’s work—and AI’s name.**

- **Arbitrary bold** mid-sentence, typical of ChatGPT-style formatting.
- A DOI that **does not exist** (hallucinated citations).
- URL with `utm_source=chatgpt.com` to **non-existent papers** (when authors explicitly claim no GenAI use).
- *More to come later in the presentation.*

Spot the AI: Which of These Two Excerpts Was AI-Generated?

Student Submission A

"This study **dives into** the **odyssey** of educational leaders, tracing their path from first forays into the field to **ascending the echelons** of leadership, all while **navigating** the swift currents of technological advancement. It examines how leaders **seamlessly integrate** recent technologies into **pedagogical** and administrative **frameworks**, thereby reshaping their roles... **unearthing** not just the mechanisms through which leaders **augment** their ability and knowledge, but also their **instrumental role** in **sculpting** the educational future... an in-depth examination of how educational leadership serves as a **cornerstone** for forging **resilient, forward-thinking** educational entities adept at confronting the **multifaceted** challenges of the twenty-first century."

Student Submission B

"The way leadership acts and behaves has a great influence on the behavior of students and the results they achieve. Research shows that successful leadership *is has* an important role in supporting and promoting students to obtain a great success. For example, a study made in Hong Kong spotted the light on the direct effect of primary leadership on school productivity, assuring the need for creating a strong leadership early among schools (Burton et al. 2005). Another study from the Philippines showed that leadership made a great transformation in learning environment and allowed students to have amazing results and satisfying performance (Estacio and Estacio 2022)..."

Source: Chaaban, Y. (2025). *Exploring research ethics through the lens of critical posthumanism in the age of Artificial Intelligence. Teaching in Higher Education*, 30(7), 1740–1755. <https://doi.org/10.1080/13562517.2025.2465995>

Spot the AI: Which of These Two Excerpts Was AI-Generated?

Student Submission A

"This study **delves into** the **odyssey** of educational leaders, tracing their path from first forays into the field to **ascending the echelons** of leadership, all while **navigating** the swift currents of technological advancement. It examines how leaders **seamlessly integrate** recent technologies into **pedagogical** and administrative **frameworks**, thereby reshaping their roles... **unearthing** not just the mechanisms through which leaders **augment** their ability and knowledge, but also their **instrumental role** in **sculpting** the educational future... an in-depth examination of how educational leadership serves as a **cornerstone** for forging **resilient, forward-thinking** educational entities adept at confronting the **multifaceted** challenges of the twenty-first century."

Student Submission B

"The way leadership acts and behaves has a great influence on the behavior of students and the results they achieve. Research shows that successful leadership *is has* an important role in supporting and promoting students to obtain a great success. For example, a study made in Hong Kong spotted the light on the direct effect of primary leadership on school productivity, assuring the need for creating a strong leadership early among schools (Burton et al. 2005). Another study from the Philippines showed that leadership made a great transformation in learning environment and allowed students to have amazing results and satisfying performance (Estacio and Estacio 2022)..."

Source: Chaaban, Y. (2025). *Exploring research ethics through the lens of critical posthumanism in the age of Artificial Intelligence. Teaching in Higher Education*, 30(7),

1740–1755. <https://doi.org/10.1080/13562517.2025.2465995>

- The ethical challenge is not the use of technology, but whether the writer retains their agency.
(*Writing is always co-shaped by tools—from search engines to reference managers.*)

Spot the AI: Which of These Two Excerpts Was AI-Generated?

Student Submission A

"This study **delves into** the **odyssey** of educational leaders, tracing their path from first forays into the field to **ascending the echelons** of leadership, all while **navigating** the swift currents of technological advancement. It examines how leaders **seamlessly integrate** recent technologies into **pedagogical** and administrative **frameworks**, thereby reshaping their roles... **unearthing** not just the mechanisms through which leaders **augment** their ability and knowledge, but also their **instrumental role** in **sculpting** the educational future... an in-depth examination of how educational leadership serves as a **cornerstone** for forging **resilient, forward-thinking** educational entities adept at confronting the **multifaceted** challenges of the twenty-first century."

Student Submission B

"The way leadership acts and behaves has a great influence on the behavior of students and the results they achieve. Research shows that successful leadership *is* has an important role in supporting and promoting students to obtain a great success. For example, a study made in Hong Kong spotted the light on the direct effect of primary leadership on school productivity, assuring the need for creating a strong leadership early among schools (Burton et al. 2005). Another study from the Philippines showed that leadership made a great transformation in learning environment and allowed students to have amazing results and satisfying performance (Estacio and Estacio 2022)..."

Source: Chaaban, Y. (2025). Exploring research ethics through the lens of critical posthumanism in the age of Artificial Intelligence. *Teaching in Higher Education*, 30(7),

1740–1755. <https://doi.org/10.1080/13562517.2025.2465995>

- The ethical challenge is not the use of technology, but whether the writer retains their agency. (*Writing is always co-shaped by tools—from search engines to reference managers.*)
- Beyond agency, there must be a *genuine intent* to act ethically—honoring community norms, meeting policy requirements, and never unduly desiring credit by taking it away from where it belongs.
- This requires transparency and wisdom.

A Wake-Up Call: The NeurIPS 2025 Hallucination Report

GPTZero Hallucination Check—NeurIPS 2025

- 100+ fabricated citations across 51 of 4,841 accepted papers
- NeurIPS = one of the world's most prestigious ML/AI venues
- Every fake reference **confirmed by a human expert**
- None flagged during peer review
- Source: gptzero.me/news/neurips^a

^aGPTZero (2025). Hallucinated citations in NeurIPS 2025. gptzero.me/news/neurips

What this reveals

- AI fabricates plausible-sounding references
- Authors submitted work they had **not read**
- Expert reviewers did not catch it

The moral principle

"When you cite something, you vouch for it. Your words carry moral weight. You must mean those things—and those things must be actual."

Part 1—The Crisis

15 min

AI slop, hallucinated citations,
the inside-outside paradox

Part 2—CRAFT Framework

20 min

Five dimensions of writer formation

Part 3—CRAFTED Framework

20 min

Seven pillars of effective output

Part 4—Hands-On Activity

20 min

Spot the slop; apply the framework

Part 5—Navigating Responsibly

15 min

QU policy, practical scaffolding,
your commitment

Rules of Engagement

[Poll] Five polls throughout

Questions after Part 1, 3 & Part 5. Send your
questions in Chat or ask directly.

Part 1:

The Crisis of AI-Augmented Writing

When Anyone Can Generate, What Does It Mean to Write?



Cambridge Dictionary & Merriam Webster both chose
"Slop" as their 2025 Word of the Year

When Anyone Can Generate, What Does It Mean to Write?



Cambridge Dictionary & Merriam Webster both chose
"Slop" as their 2025 Word of the Year

Today's Question:

**When AI can write for us,
what is the value of writing ourselves?**

When Anyone Can Generate, What Does It Mean to Write?



Cambridge Dictionary & Merriam Webster both chose
"Slop" as their 2025 Word of the Year

Today's Question:

When AI can write for us,
what is the value of writing ourselves?

Our Central Argument: The Multiplexity Imperative for AI-Augmented Writing

Excellence in the AI era requires **BOTH** becoming—cultivating authentic human capability—**AND** producing—creating valuable outputs.

Even in the AI era, authentic craft still matters.

The Inside-Outside Paradox

Inside—Writer Formation

Understanding, capability, authentic voice, and agency that develop **through** the writing process.

What the writer **becomes**.

Resists automation. Emerges only through direct engagement.

Outside—Output Quality

The quality, coherence, and effectiveness of the written artefact—what reaches the reader.

What the writer **produces**.

AI excels here: polished, grammatically correct, consistent.

Three Failure Modes

1. **AI Slop**—polished output, shallow comprehension
2. **Communication Failure**—deep expertise, inaccessible writing
3. **Compound Mediocrity**—neither developed

The Inside-Outside Paradox

Inside—Writer Formation

Understanding, capability, authentic voice, and agency that develop **through** the writing process.

What the writer **becomes**.

Resists automation. Emerges only through direct engagement.

Outside—Output Quality

The quality, coherence, and effectiveness of the written artefact—what reaches the reader.

What the writer **produces**.

AI excels here: polished, grammatically correct, consistent.

Three Failure Modes

1. **AI Slop**—polished output, shallow comprehension
2. **Communication Failure**—deep expertise, inaccessible writing
3. **Compound Mediocrity**—neither developed

Inside-Outside Paradox



The Inside-Outside Paradox (The Primacy of the Inside)

Inside—Writer Formation

Understanding, capability, authentic voice, and agency that develop **through** the writing process.

What the writer **becomes**.

Outside—Output Quality

The quality, coherence, and effectiveness of the written artefact—what reaches the reader.

What the writer **produces**.



The Grammarian (obsessed with form):
“Have you studied grammar?”

The Boatman: “No.”

The Grammarian: “Then half your life has been wasted.”

The Inside-Outside Paradox (The Primacy of the Inside)

Inside—Writer Formation

Understanding, capability, authentic voice, and agency that develop **through** the writing process.

What the writer **becomes**.



The Grammarian (obsessed with form):
“Have you studied grammar?”

The Boatman: “No.”

The Grammarian: “Then half your life has been wasted.”

Outside—Output Quality

The quality, coherence, and effectiveness of the written artefact—what reaches the reader.

What the writer **produces**.



The Boatman (as the boat sinks):
“Have you learned to swim?”

The Grammarian: “No.”

The Boatman: “Then ALL your life is lost.”

A Polarity to Navigate, Not a Problem to Solve

Problems have permanent solutions.

Polarities have interdependent poles
that need each other.

A Polarity to Navigate, Not a Problem to Solve

Problems have permanent solutions.
Polarities have interdependent poles
that need each other.

Goal: Both-And Navigation

Use AI to enhance efficiency **and simultaneously**
develop authentic human capability.

Not either-or. Both-and.

The Core Principle

*Excellence is found in the continuous
navigation between the **Writer** (Internal
Capability) and the **Writing** (External
Artifact).*

A Polarity to Navigate, Not a Problem to Solve

Problems have permanent solutions.
Polarities have interdependent poles that need each other.


Goal: Both-And Navigation

Use AI to enhance efficiency **and** simultaneously develop authentic human capability.


Not either-or. Both-and.

The Core Principle

*Excellence is found in the continuous navigation between the **Writer** (Internal Capability) and the **Writing** (External Artifact).*



Need for Polarity Management: Cognitive Debt and Irony of Automation




¹ Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.

² Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.

³ Kosmyna, N., et al. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. MIT Media Lab.

Need for Polarity Management: Cognitive Debt and Irony of Automation



Automation Paradox (*It gets harder, the easier it gets*):¹


The more automated a system becomes, the more critical human judgment is—yet the less we practice it.

¹ Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.

² Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.

³ Kosmyna, N., et al. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. MIT Media Lab.

Need for Polarity Management: Cognitive Debt and Irony of Automation



Automation Paradox (*It gets harder, the easier it gets*):¹

The more automated a system becomes, the more critical human judgment is—yet the less we practice it.

Matthew Effect (*Rich gets richer, poor get poorer*):²


Advantages compound over time: those who practice thinking become stronger thinkers, while those who offload it become increasingly dependent on tools.

¹Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.

²Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.

³Kosmyna, N., et al. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. MIT Media Lab.

Need for Polarity Management: Cognitive Debt and Irony of Automation



Automation Paradox (*It gets harder, the easier it gets*):¹

The more automated a system becomes, the more critical human judgment is—yet the less we practice it.

Matthew Effect (*Rich gets richer, poor get poorer*):²

Advantages compound over time: those who practice thinking become stronger thinkers, while those who offload it become increasingly dependent on tools.

Cognitive Debt (*You have to pay back later for using AI*):³

When we outsource cognitive effort to AI, we incur a hidden debt: the gradual erosion of our own thinking abilities.

¹ Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.

² Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.

³ Kosmyna, N., et al. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. MIT Media Lab.

The CRAFT(-ED) Framework: A Holistic Strategy

CRAFT Inside: Formation



Consciousness

Epistemic awareness of what you genuinely understand vs. AI-generated content

Relational

Genuine engagement with ideas, readers, and scholarly communities

Agentic

Ownership, judgment, and self-authorship in writing

Formative


Development over time, becoming capable of original contribution

Truthful

Authenticity, honesty, and epistemic integrity

Both-And Integration

Neither formation nor output alone suffices.
Excellence requires continuous navigation.




The Challenge:

Use AI to enhance efficiency
and output quality *while simultaneously*
developing authentic human capability.

Not either-or, but both-and.

CRAFTED Outside: Output Quality



The Core Principle

*Neither formation (**CRAFT**) nor output focus (**CRAFTED**) suffices alone.*

[Poll #1]

How do you currently use AI in your research writing?

- A) I do not use AI tools at all
- B) I use AI occasionally—grammar / editing only
- C) I use AI to help brainstorm or outline
- D) I use AI substantially (drafting, summarising sources)
- E) AI is central to most of my writing workflow.

[Poll #2 — Your Diagnosis of the Crisis]

Which of these dangers feels most real in your own context?

- A) Fake references slipping into published research undetected
- B) Researchers slowly losing skills they once had — without realising it
- C) Writing that sounds impressive but says nothing the author truly owns
- D) Using AI to go faster, while quietly going shallower
- E) The trap is already closing — and most people around me haven't noticed

We will return to your responses when we discuss scaffolding and institutional policy in Part 5.

Part 2:

CRAFT—Inside-Focused Writer Formation

CRAFT: Who You Become as a Writer

CRAFT = five dimensions of authentic writer development

These are qualities that **resist commodification** and cannot be delegated to AI. They emerge only through direct engagement with ideas.

CRAFT: Who You Become as a Writer

CRAFT = five dimensions of authentic writer development

These are qualities that **resist commodification** and cannot be delegated to AI. They emerge only through direct engagement with ideas.

Dimension	Core Focus	Key Question
Consciousness	Epistemic awareness	Do I know what I genuinely understand?
Relational	Genuine engagement	Have I authentically engaged with sources?
Agentic	Self-authorship	Am I directing AI, or is AI directing me?
Formative	Capability growth	Am I developing, or just assembling?
Truthful	Epistemic integrity	Can I stand behind every claim?

CRAFT: Who You Become as a Writer

CRAFT = five dimensions of authentic writer development

These are qualities that **resist commodification** and cannot be delegated to AI. They emerge only through direct engagement with ideas.

Dimension	Core Focus	Key Question
Consciousness	Epistemic awareness	Do I know what I genuinely understand?
Relational	Genuine engagement	Have I authentically engaged with sources?
Agentic	Self-authorship	Am I directing AI, or is AI directing me?
Formative	Capability growth	Am I developing, or just assembling?
Truthful	Epistemic integrity	Can I stand behind every claim?

Core question: Is the writer developing genuine understanding and irreplaceable capability?

C—Consciousness: Epistemic Awareness

Knowing the difference between what you **genuinely understand** and what AI generated.

Arabic concept of *shaur*—consciousness, awakening.

C—Consciousness: Epistemic Awareness

Knowing the difference between what you **genuinely understand** and what AI generated.

Arabic concept of *shaur*—consciousness, awakening.

The Ghazali Test



The Robber to Ghazali: "How can you claim to knowledge if it can be stolen from you so easily?"

Can you explain every major claim in your document **without referring to your notes?**

Could you restructure your argument if asked, maintaining its **conceptual integrity?**

C—Consciousness: Epistemic Awareness

Knowing the difference between what you **genuinely understand** and what AI generated.

Arabic concept of *shaur*—consciousness, awakening.

The Ghazali Test



The Robber to Ghazali: "How can you claim to knowledge if it can be stolen from you so easily?"

Can you explain every major claim in your document **without referring to your notes?**

Could you restructure your argument if asked, maintaining its **conceptual integrity?**

Developmental trajectory:

Novice: accepts AI output uncritically

Competent: checks AI against own understanding

Expert: uses AI as exploration tool; knows exactly what they know, what they verified, what remains uncertain.

C—Consciousness: Epistemic Awareness

Knowing the difference between what you **genuinely understand** and what AI generated.

Arabic concept of *shaur*—consciousness, awakening.

The Ghazali Test



The Robber to Ghazali: "How can you claim to knowledge if it can be stolen from you so easily?"

Can you explain every major claim in your document **without referring to your notes?**

Could you restructure your argument if asked, maintaining its **conceptual integrity?**

Developmental trajectory:

Novice: accepts AI output uncritically

Competent: checks AI against own understanding

Expert: uses AI as exploration tool; knows exactly what they know, what they verified, what remains uncertain.

The paradox

Using AI effectively requires **more** sophisticated consciousness—you must evaluate outputs against your own understanding.

R—Relational: Genuine Engagement

Authentic engagement with sources, ideas, readers, and scholarly communities—not extractive use of AI summaries.

Learning happens through **relationships**: with texts, peers, mentors, and the scholarly conversation you are joining.

R—Relational: Genuine Engagement

Authentic engagement with sources, ideas, readers, and scholarly communities—not extractive use of AI summaries.

Learning happens through **relationships**: with texts, peers, mentors, and the scholarly conversation you are joining.

Key question

Have you authentically engaged with key sources, or merely used AI summaries?

Does your work contribute to ongoing scholarly conversations?

R—Relational: Genuine Engagement

Authentic engagement with sources, ideas, readers, and scholarly communities—not extractive use of AI summaries.

Learning happens through **relationships**: with texts, peers, mentors, and the scholarly conversation you are joining.

Key question

Have you authentically engaged with key sources, or merely used AI summaries?

Does your work contribute to ongoing scholarly conversations?

NeurIPS connection

Those fabricated citations represent a **relational failure**: the authors had no relationship with the sources they cited.

They never read them. They outsourced that relationship to a tool that confabulated it.

AI can help you *find* sources.

Only you can *engage* with them.

A—Agentic: Ownership and Self-Authorship

You are the author. Every structural choice, every argument, every framing decision is **genuinely yours**.

The agency gradient

Low: Prompt AI for a full draft. Accept. Submit.

Medium: Use AI for a draft; revise sections; check logic.

High: Develop your own argument first; use AI to stress-test, polish, and check for gaps; retain authorial control throughout.

A—Agentic: Ownership and Self-Authorship

You are the author. Every structural choice, every argument, every framing decision is **genuinely yours**.

The agency gradient

Low: Prompt AI for a full draft. Accept. Submit.

Medium: Use AI for a draft; revise sections; check logic.

High: Develop your own argument first; use AI to stress-test, polish, and check for gaps; retain authorial control throughout.

Key question

Can you defend every major organizational and stylistic choice in your document?

Are you directing AI, or is AI directing you?

A—Agentic: Ownership and Self-Authorship

You are the author. Every structural choice, every argument, every framing decision is **genuinely yours**.

The agency gradient

Low: Prompt AI for a full draft. Accept. Submit.

Medium: Use AI for a draft; revise sections; check logic.

High: Develop your own argument first; use AI to stress-test, polish, and check for gaps; retain authorial control throughout.

Key question

Can you defend every major organizational and stylistic choice in your document?

Are you directing AI, or is AI directing you?

The authorship test

Could a colleague who knows your work recognise your intellectual fingerprint in this document?

F—Formative: Capability Development Over Time

Writing is a **formative practice**, not merely a task. The struggle is the learning.

Sennett (2008): Craftsmanship is *“an enduring, basic human impulse, the desire to do a job well for its own sake.”^a*

The Key Risk

When AI short-circuits the struggle, it enables **production without formation**: generating outputs without developing capabilities.

F—Formative: Capability Development Over Time

Writing is a **formative practice**, not merely a task. The struggle is the learning.

Sennett (2008): Craftsmanship is *“an enduring, basic human impulse, the desire to do a job well for its own sake.”*^a

The Key Risk

When AI short-circuits the struggle, it enables **production without formation**: generating outputs without developing capabilities.

The Litmus Test

Are you developing capabilities through this work, or merely assembling outputs?

Would your abilities remain if AI disappeared?

^aSennett, R. (2008). *The Craftsman*. Yale University Press.

F—Formative: Capability Development Over Time

Writing is a **formative practice**, not merely a task. The struggle is the learning.

Sennett (2008): Craftsmanship is *“an enduring, basic human impulse, the desire to do a job well for its own sake.”*^a

The Key Risk

When AI short-circuits the struggle, it enables **production without formation**: generating outputs without developing capabilities.

^aSennett, R. (2008). *The Craftsman*. Yale University Press.

The Litmus Test

Are you developing capabilities through this work, or merely assembling outputs?

Would your abilities remain if AI disappeared?

Metanoia (Inqilāb):

Writing serves a hierarchy:

- **Writing to Remember**—building memory.
- **Writing to Inform**—clear communication.
- **Writing on the Heart**—internalizing wisdom.

T—Truthful: Epistemic Integrity

You must mean what you write. Your words carry moral weight. Claims must be actual.

You cannot ethically present AI-generated content as your own understanding. You cannot cite sources you have not read. You cannot make claims you cannot explain.

Key question

Can you stand behind every claim—explaining, defending, and taking responsibility?

Have you been transparent about AI's role?

T—Truthful: Epistemic Integrity

You must mean what you write. Your words carry moral weight. Claims must be actual.

You cannot ethically present AI-generated content as your own understanding. You cannot cite sources you have not read. You cannot make claims you cannot explain.

Key question

Can you stand behind every claim—explaining, defending, and taking responsibility?

Have you been transparent about AI's role?

The moral weight of citation

When you include a reference, you declare this source exists, says what you claim, and supports your argument.

The NeurIPS scandal is not just a technical failure. It is a **truthfulness failure** at scale.

T—Truthful: Epistemic Integrity

You must mean what you write. Your words carry moral weight. Claims must be actual.

You cannot ethically present AI-generated content as your own understanding. You cannot cite sources you have not read. You cannot make claims you cannot explain.

Key question

Can you stand behind every claim—explaining, defending, and taking responsibility?

Have you been transparent about AI's role?

The moral weight of citation

When you include a reference, you declare this source exists, says what you claim, and supports your argument.

The NeurIPS scandal is not just a technical failure. It is a **truthfulness failure** at scale.

Transparency is also truthfulness:

Disclose what AI tools you used, where, and how. This is not weakness—it is scholarly honesty.

CRAFT Self-Assessment Checklist

Dimension	Ask yourself. . .
Consciousness	Can I explain every major claim without notes? Do I know what I genuinely understand vs. what AI generated?
Relational	Have I authentically engaged with key sources—not just AI summaries? Does my work contribute to the scholarly conversation?
Agentic	Can I defend every organizational and stylistic choice? Am I directing AI, or is AI directing me?
Formative	Am I developing capabilities through this work, or merely assembling outputs? Would my skills survive without AI?
Truthful	Can I stand behind every claim? Have I verified all facts and citations? Have I disclosed AI use transparently?

Which CRAFT dimension do you find hardest to maintain when using AI?

- A) Consciousness—knowing what I genuinely understand
- B) Relational—engaging with sources rather than summaries
- C) Agentic—keeping ownership of the argument
- D) Formative—ensuring I am actually developing my skills
- E) Truthful—verifying everything and disclosing AI use

Part 3:

CRAFTED—Outside-Focused Communication


CRAFTED: What You Produce

CRAFTED = seven pillars of output quality

While CRAFT addresses writer formation (the *inside*), CRAFTED addresses effective expression (the *outside*): creating documents that are clear, accessible, and valuable to readers.

Pillar	Core Focus	Key Question
Coherence	Unified narrative	Does it tell one integrated story?
Reader Focus	Empathetic communication	Have you genuinely considered your reader?
Accuracy	Evidence-grounded claims	Have you personally verified every claim?
Flow	Seamless connectivity	Do transitions reflect real logical connections?
Top-Down	Strategic framing	Can you state your core message in one sentence?
Espirited	Distinctive voice	Does this sound authentically like you?
Directed	Intentional oversight	Are you directing AI—or vice versa?

CRAFTED—Writing Short Course (The Mechanics of Writing)



On Writing Well
6 Tips on Bringing Clarity to Writing

1. Top-Down (Big Picture)
2. Flow (Connected)
3. Coherence (United)
4. Empathy (Reader Focus)
5. Design (Consistency)
6. Style (Flair)

JUNAID QADIR

Writing Course - June 2020
by Junaid Qadir
Playlist • Public • 4 videos • 795 views
Technical Writing Short Course...more

Play all

Manual ▾

On Writing Well
6 Tips on Bringing Clarity to Writing

1. Top-Down (Big Picture)
2. Flow (Connected)
3. Coherence (United)
4. Empathy (Reader Focus)
5. Design (Consistency)
6. Style (Flair)

JUNAID QADIR

19:38

Writing Course Intro
Junaid Qadir • 904 views • 5 years ago

Part I—Writing Top Down (Writing Short Course)
Junaid Qadir • 388 views • 5 years ago

Part 2—Writing With Flow (Writing Short Course)
Junaid Qadir • 174 views • 5 years ago

Part 3—Writing With Coherence (Writing Short Course)
Junaid Qadir • 149 views • 5 years ago

<https://www.youtube.com/junaidq>

Coherence: Unified Narrative

Makes a document feel like a **unified whole**, not disconnected parts. Operates at: sentence → paragraph → section → whole document.

AI problem: AI generates competent paragraphs locally, but the global story arc requires human understanding.

Coherence: Unified Narrative

Makes a document feel like a **unified whole**, not disconnected parts. Operates at: sentence → paragraph → section → whole document.

AI problem: AI generates competent paragraphs locally, but the global story arc requires human understanding.

Test

Can you explain *why* each section connects as it does — not just that it does?

Coherence: Unified Narrative

Makes a document feel like a **unified whole**, not disconnected parts. Operates at: sentence → paragraph → section → whole document.

AI problem: AI generates competent paragraphs locally, but the global story arc requires human understanding.

Test

Can you explain *why* each section connects as it does — not just that it does?

Reader Focus: Empathetic Communication

Writing is not for you—it is for your reader. AI optimises for generic readability, not your *specific* reader.

Key question

Have you genuinely considered your readers' needs, background, and constraints?

Coherence: Unified Narrative

Makes a document feel like a **unified whole**, not disconnected parts. Operates at: sentence → paragraph → section → whole document.

AI problem: AI generates competent paragraphs locally, but the global story arc requires human understanding.

Test

Can you explain *why* each section connects as it does — not just that it does?

Reader Focus: Empathetic Communication

Writing is not for you—it is for your reader. AI optimises for generic readability, not your *specific* reader.

Key question

Have you genuinely considered your readers' needs, background, and constraints?

Exercise

Complete this sentence before drafting:

“My reader is [...], and after reading my paper they will be able to [...].”

A — Accuracy: Evidence-Grounded Claims

Every factual claim, statistic, and citation must be **personally verified** by the author.

The NeurIPS lesson

100+ hallucinated citations at a top AI conference.
Peer review did not catch them.

AI hallucination is not a bug to be fixed — it is an inherent property of LLMs to manage.

gptzero.me/news/neurips

A — Accuracy: Evidence-Grounded Claims

Every factual claim, statistic, and citation must be **personally verified** by the author.

The NeurIPS lesson

100+ hallucinated citations at a top AI conference.
Peer review did not catch them.

AI hallucination is not a bug to be fixed — it is an inherent property of LLMs to manage.

gptzero.me/news/neurips

Self-assessment

Have you *personally* verified every factual claim, statistic, and citation?

A — Accuracy: Evidence-Grounded Claims

Every factual claim, statistic, and citation must be **personally verified** by the author.

The NeurIPS lesson

100+ hallucinated citations at a top AI conference. Peer review did not catch them.

AI hallucination is not a bug to be fixed — it is an inherent property of LLMs to manage.

gptzero.me/news/neurips

Self-assessment

Have you *personally* verified every factual claim, statistic, and citation?

Practical guidance:

- Click every reference. Read what it actually says.
- Re-enter citations manually in Mendeley/Zotero — never copy-paste from AI output.
- For literature search, prefer grounded tools over general LLMs:
 - elicit.com
 - typeset.io (SciSpace)
 - consensus.app
- Run GPTZero's Hallucination Check before submission:
gptzero.me/hallucination-detector

Flow: Seamless Connectivity

The reader moves effortlessly from idea to idea because the connections are *real*, not cosmetic.

AI problem: Fluent transitions

(“Furthermore...”; “Additionally...”) can paper over logical gaps.

Flow: Seamless Connectivity

The reader moves effortlessly from idea to idea because the connections are *real*, not cosmetic.

AI problem: Fluent transitions

(“Furthermore...”; “Additionally...”) can paper over logical gaps.

Test

Remove all transition words. Does the logic still hold?

Yes = real flow. **No** = AI-glue.

F—Flow & T—Top-Down Thinking

Flow: Seamless Connectivity

The reader moves effortlessly from idea to idea because the connections are *real*, not cosmetic.

AI problem: Fluent transitions (“Furthermore...”; “Additionally...”) can paper over logical gaps.

Test

Remove all transition words. Does the logic still hold?

Yes = real flow. **No** = AI-glue.

Top-Down: Strategic Framing

State the conclusion first, then arguments, then evidence (Minto Pyramid Principle).

AI problem: AI naturally writes bottom-up — accumulates details, concludes late. Reads like a journey, not a declaration.

F—Flow & T—Top-Down Thinking

Flow: Seamless Connectivity

The reader moves effortlessly from idea to idea because the connections are *real*, not cosmetic.

AI problem: Fluent transitions (“Furthermore...”; “Additionally...”) can paper over logical gaps.

Test

Remove all transition words. Does the logic still hold?

Yes = real flow. **No** = AI-glue.

Top-Down: Strategic Framing

State the conclusion first, then arguments, then evidence (Minto Pyramid Principle).

AI problem: AI naturally writes bottom-up — accumulates details, concludes late. Reads like a journey, not a declaration.

Key question

Can you state your core message in **one sentence** and explain how each section serves it?

Espirited: Distinctive Voice

The signature failure of AI prose: competent but unmemorable. No personality, no specific angle, no authentic scholar.

The personality test

Could a colleague who knows your work recognise this document as yours, based on style alone?

Use AI for initial drafts; revise *extensively* for voice. Add metaphors, examples, phrasings that are distinctively yours.

Espirited: Distinctive Voice

The signature failure of AI prose: competent but unmemorable. No personality, no specific angle, no authentic scholar.

The personality test

Could a colleague who knows your work recognise this document as yours, based on style alone?

Use AI for initial drafts; revise *extensively* for voice. Add metaphors, examples, phrasings that are distinctively yours.

Directed: Intentional Oversight

The meta-pillar. You are the director; AI is the instrument. You set purpose, evaluate outputs, make final decisions.

Key question

Are you directing AI as a tool, or passively accepting whatever it generates?

Detection connection

Directed, well-revised AI use is precisely what detection tools *cannot* flag—and that is the point.

CRAFTED: Integration with CRAFT Dimensions

CRAFTED Pillar	Polarity Navigated	Capability Developed
Coherence	AI structure ↔ human logic	Epistemic awareness
Reader Focus	AI polish ↔ human empathy	Audience engagement
Accuracy	AI efficiency ↔ verification	Critical evaluation
Flow	AI bridges ↔ real connections	Logical integration
Top-Down	AI detail ↔ human framing	Strategic ownership
Espirited	AI formula ↔ human creativity	Distinctive voice
Directed	AI automation ↔ human oversight	Meta-cognitive control

[Poll #4]

Which CRAFTED dimension is hardest to achieve when AI is involved in your writing?

- A) Coherence — the argument holds together as a whole
- B) Reader Focus — writing for a real audience, not for the AI prompt
- C) Accuracy — personally verifying every claim and citation
- D) Flow — transitions that reflect real logic, not AI-glue
- E) Top-Down — stating your conclusion first, then building the case
- F) Espirited — writing that still sounds distinctively like you
- G) Directed — a piece that serves a clear purpose and lands with impact

Part 4:

Hands-On Activities

AI Detection: A Signal, Not a Verdict

What detection tools **CAN** do:

- Reliably flag **single-turn, minimally-edited** AI output
- Provide a probabilistic signal—not a verdict—with a confidence score
- Catch the most egregious cases of outsourced writing
- Verify whether cited sources actually exist (**Citation Check**)

What detection tools **CANNOT** do:

- Catch well-revised, human-directed AI use
- Replace institutional culture and judgment
- Eliminate bias risk: non-native writers flagged more often (Liang et al., 2023)—good tools actively correct for this^a

^aLiang et al. (2023). GPT detectors are biased against non-native English writers. *Pat-terns*.

Used responsibly, detection starts a conversation

A high AI probability score is not grounds for immediate accusation. Ask students to produce evidence of their writing process—drafts, revision histories, brainstorming notes.

The honest conclusion

The real solution is a **culture of authentic scholarship**, not a detection arms race. Detection is one tool among many.

Activity: Identify the Type of Issue

20 minutes

What you will do:

Step 1 (5 min, individual): Read the four real NeurIPS 2025 citations on the next slide. For each: identify *what* is fabricated and *which* CRAFT dimension it violates.

Step 2 (10 min, breakout rooms): Compare findings. Agree on the single most alarming example. Why did peer reviewers miss it?

Step 3 (5 min, whole group): Each group shares their finding. Submit one key word to the Poll word cloud.

Guiding questions

- Can the writer explain every claim?
- Are the citations real and actually read?
- Is there a discernible argument—or just a tour?
- Does any authentic voice come through?
- Are transitions substantive or cosmetic?

After breakouts

Submit **one word or phrase** describing the paragraph's core weakness to the Poll word cloud.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Placeholder author names. The “to be updated” URL not updated. Published in an accepted NeurIPS paper.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Placeholder author names. The “to be updated” URL not updated. Published in an accepted NeurIPS paper.

[3] Author Song and AnotherAuthor Zhang. “Consistency in Diffusion Models: Improving Noise Embeddings.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023). URL: <https://arxiv.org/abs/2304.08787>.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Placeholder author names. The “to be updated” URL not updated. Published in an accepted NeurIPS paper.

[3] Author Song and AnotherAuthor Zhang. “Consistency in Diffusion Models: Improving Noise Embeddings.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023). URL: <https://arxiv.org/abs/2304.08787>.

No author or title match. The arXiv ID leads to a completely unrelated paper.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Placeholder author names. The “to be updated” URL not updated. Published in an accepted NeurIPS paper.

[3] Author Song and AnotherAuthor Zhang. “Consistency in Diffusion Models: Improving Noise Embeddings.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023). URL: <https://arxiv.org/abs/2304.08787>.

No author or title match. The arXiv ID leads to a completely unrelated paper.

[4] M. Garcia and A. Thompson. Applications of LLMs in legal document analysis. *Journal of Legal Technology*, 7(1):50–65, 2024.

Real Examples: Verified Hallucinations from NeurIPS 2025

These citations appeared in accepted NeurIPS 2025 papers—verified as hallucinations by GPTZero + human expert

[1] John Doe and Jane Smith. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.00001*, 2024.

The paper exists, but authors are fabricated (common names). The arXiv ID links to a different article entirely.

[2] Firstname Lastname and Others. Drivlme: A large-scale multi-agent driving benchmark, 2023. *URL or arXiv ID to be updated.*

Placeholder author names. The “to be updated” URL not updated. Published in an accepted NeurIPS paper.

[3] Author Song and AnotherAuthor Zhang. “Consistency in Diffusion Models: Improving Noise Embeddings.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023). URL: <https://arxiv.org/abs/2304.08787>.

No author or title match. The arXiv ID leads to a completely unrelated paper.

[4] M. Garcia and A. Thompson. Applications of LLMs in legal document analysis. *Journal of Legal Technology*, 7(1):50–65, 2024.

No author or title match. The journal does not exist.

Source: GPTZero Hallucination Check, NeurIPS 2025. Each citation confirmed by a human expert. None flagged during peer review. gptzero.me/news/neurips

Part 5:

Navigating Responsibly

Know Your Institution: QU's AI Policy

Qatar University—Instructions on the Use of Generative AI Tools

<https://www.qu.edu.qa/en-us/about/policy/Documents/Instructions%20onThe%20Use%20of%20Generative%20Artificial%20Intelligence%20Tools%20at%20Qatar%20University-Approved.pdf>

Know Your Institution: QU's AI Policy

Qatar University—Instructions on the Use of Generative AI Tools

<https://www.qu.edu.qa/en-us/about/policy/Documents/Instructions%20onThe%20Use%20of%20Generative%20Artificial%20Intelligence%20Tools%20at%20Qatar%20University-Approved.pdf>

Key principles for researchers and graduate students:

- **Disclosure is mandatory.** Any GenAI use must be acknowledged transparently.
- **Responsibility stays with you.** AI cannot be an author; you are accountable for every word.
- **Academic integrity applies fully.** Fabricated citations and undisclosed AI use violate QU policy.
- **Context determines what is permitted.** Check your program, supervisor, and target journal.

Know Your Institution: QU's AI Policy

Qatar University—Instructions on the Use of Generative AI Tools

<https://www.qu.edu.qa/en-us/about/policy/Documents/Instructions%20onThe%20Use%20of%20Generative%20Artificial%20Intelligence%20Tools%20at%20Qatar%20University-Approved.pdf>

Key principles for researchers and graduate students:

- **Disclosure is mandatory.** Any GenAI use must be acknowledged transparently.
- **Responsibility stays with you.** AI cannot be an author; you are accountable for every word.
- **Academic integrity applies fully.** Fabricated citations and undisclosed AI use violate QU policy.
- **Context determines what is permitted.** Check your program, supervisor, and target journal.

Spirit AND letter of the law

Compliance with policy is necessary but not sufficient. The spirit demands authenticity, truthfulness, and genuine contribution.

A Practical Guide: The ETHICAL Protocol

Alduais, Qadhi, Chaaban, & Khraisheh (2025),^a
AI Magazine—a nine-item checklist built from
analysis of 74 policy documents.

- E** Establish your purpose
- T** Thoroughly explore options
- H** Harness the right tool
- I** Inspect and verify output
- C** Cite and reference accurately
- A** Acknowledge AI use transparently
- L** Look over publisher guidelines

A Practical Guide: The ETHICAL Protocol

Alduais, Qadhi, Chaaban, & Khraisheh (2025),^a
AI Magazine—a nine-item checklist built from
analysis of 74 policy documents.

- E** Establish your purpose
- T** Thoroughly explore options
- H** Harness the right tool
- I** Inspect and verify output
- C** Cite and reference accurately
- A** Acknowledge AI use transparently
- L** Look over publisher guidelines

^a Alduais et al. (2025). The ETHICAL Protocol for Responsible Use of Generative AI Tools. *AI Magazine*, 46, e70047.

Item 5: Inspect & Verify

“Copying and pasting content from any generative AI tool without scrutiny is the essence of irresponsible use.”

Click every quotation. Verify it appears on the specified page. Re-enter references manually.

Item 9: Finalize

Before submitting: remove any AI sign-off phrases left in accidentally.

“This is the list you asked me to create. . .”
Editors notice. Desk rejections follow.

Operationalizing the Framework: Laws of Intelligence Augmentation

From Qadir (2024), *IEEE Potentials*^a

The core operationalization

- Aim for proficiency in **BOTH** formation AND communication
- Prioritize the higher values: inside-focus, quality, holistic, intrinsic
- Watch for *fixes that fail*: skill atrophy is slow and invisible
- **Separate times for NO AI** (development) and **Ethical AI** (polish, critique, enhance)

Operationalizing the Framework: Laws of Intelligence Augmentation

From Qadir (2024), *IEEE Potentials*^a

The core operationalization

- Aim for proficiency in **BOTH** formation AND communication
- Prioritize the higher values: inside-focus, quality, holistic, intrinsic
- Watch for *fixes that fail*: skill atrophy is slow and invisible
- **Separate times for NO AI** (development) and **Ethical AI** (polish, critique, enhance)

^aQadir, J. (2024). Learning 101 reloaded: Revisiting the basics for the GenAI era. *IEEE Potentials*, 43(6), 6-12.

Develop your own judgement

The key meta-skill is knowing *when* to use AI and when not to.

This judgement is itself a capability that must be developed—and it develops only through practice, reflection, and honest self-assessment.

The Ghazali Test (revisited)

Can you explain it without the notes?

Ensure that over time, your skills and capabilities are **growing**—not quietly atrophying.

The Both-And Navigation in Practice

Fast & Slow

Slower than zero-engagement AI slop generation.

Faster than unaided, unoptimised human writing with all its inefficiencies and blocks.

The sweet spot: AI as scaffold—you as builder.

The Both-And Navigation in Practice

Fast & Slow

Slower than zero-engagement AI slop generation.

Faster than unaided, unoptimised human writing with all its inefficiencies and blocks.

The sweet spot: AI as scaffold—you as builder.

Learning & Performance

Not only performance (submitting outputs) but also learning (developing your craft).

It is futile to increase quantity without increasing genuine understanding.

The Both-And Navigation in Practice

Fast & Slow

Slower than zero-engagement AI slop generation.

Faster than unaided, unoptimised human writing with all its inefficiencies and blocks.

The sweet spot: AI as scaffold—you as builder.

Learning & Performance

Not only performance (submitting outputs) but also learning (developing your craft).

It is futile to increase quantity without increasing genuine understanding.

The enduring question

“The central question of AI-augmented writing is not ‘Should we use AI?’ but ‘How do we maintain dynamic equilibrium between capability development and productivity, between authentic human contribution and effective tool use, between becoming and producing?’ ”

What is your one commitment for your next writing project?

Open word cloud—type your commitment in a few words

Examples:

“Verify every citation personally”

“Write my argument before prompting AI”

“Disclose AI use in my methods/ acknowledgement section”

“Apply the Ghazali Test to my next draft”

“Share your own please ..”

Closing: When Anyone Can Generate, Learn the CRAFT

The central insight

When anyone can generate grammatically correct text, the challenge shifts from *production* to *discernment*.

Text generation is commodified.

Authentic writing—reflecting genuine understanding and disciplined thought—remains irreplaceable.

Take away from today:

- ✓ The CRAFT(-ED) framework as a self-assessment tool
- ✓ GPTZero Citation Check—verify every reference
- ✓ QU's AI policy—know it, follow it
- ✓ Your one commitment (Poll #3)

Remember Nasruddin

The student copied every answer perfectly.
He just forgot to make them *his own*.

Don't submit AI's work—or AI's name.

GIFT 1: Free GPTZero Access—One Semester for All Participants

Exclusive offer for today's participants

As a GPTZero Ambassador, I am pleased to offer everyone here **free access to GPTZero for one full semester.**

GPTZero gives you:

- AI writing detection with confidence scores
- **Citation Check**—flags hallucinated references instantly

Scan the QR code to record your attendance & receive 1 semester of GPTZero for free!

You will also receive a certificate of completion from GPTZero.



Teaching Responsibly with AI

Teaching Responsibly with AI Attendance form

This form is used to record your attendance at the Teaching Responsibly with AI workshop.

We will follow up with instructions on how to receive your semester of GPTZero for free!
You will also receive a certificate of completion for the webinar.

Please contact hriday@gptzero.me with any questions.

* Indicates required question

GIFT 2: Participation Certificate—Please Fill the QU Feedback Form



Questions & Discussion

`jqadir@qu.edu.qa`