THE PROFESSIONAL SELF-EDITING GUIDE FOR THE AGE OF AI

RECLAIM YOUR VOICE & WRITE WITH UNMISTAKABLE AUTHORITY

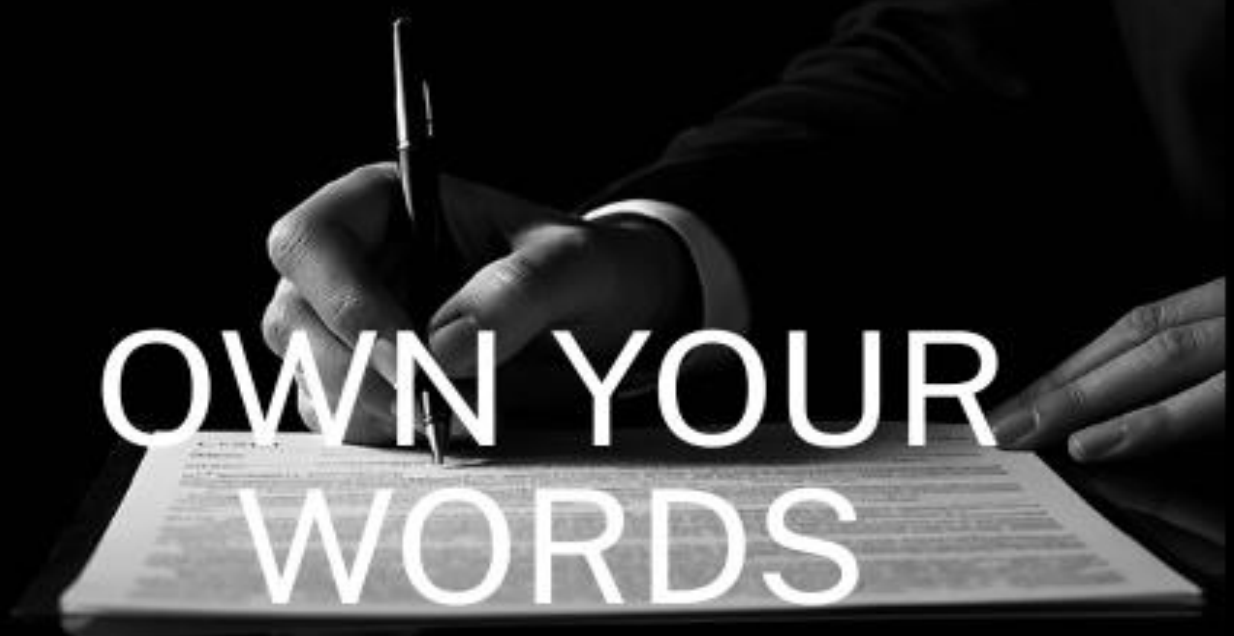

OWN YOUR WORDS

SENIOR ADVISER IN CORPORATE AND DIPLOMATIC COMMUNICATION

A. J. THORNFIELD
FORTY YEARS AT THE INTERSECTION OF LANGUAGE AND POWER

OWN YOUR WORDS

The Professional Self-Editing Guide for the Age of AI

TRANSFORMATIVEFIN HUB LEARNING CENTRE

For bulk orders, institutional licences, and educational use enquiries:
e-library@transformativefinhub.org

Table of Contents

PART FIVE: THE PROFESSIONAL EMAIL

PART SIX: REPORTS AND FORMAL DOCUMENTS

PART SEVEN: SHORT-FORM AND SOCIAL MEDIA

PART EIGHT: RECLAIMING YOUR VOICE

APPENDICES

For every professional who ever paused before sending a message — sensing something was missing — and sent it anyway.

This book names what was missing.

Foreword

This guide is for anyone whose job involves daily writing. From a finance analyst who needs to brief a board, the development officer who must persuade a donor, the junior diplomat sending her first official cable, and the consultant whose proposal determines whether a contract is awarded, among others.

Most professionals learn to write in one of two ways. The first is through formal education, which addresses grammar, structure, and argumentation — the basics of a sentence and the structure of the argument. The second is through professional experience — a memo handed back with the instruction to make it clearer, an email that produced no expected reply or worth created misunderstanding, a report not endorsed.

Both learning paths are now disrupted.

AI writing tools have inserted a third path — one that produces text with the surface grammar of the first route and the apparent confidence of the second, while bypassing the cognitive labour that makes writing useful. Over reliance on the third path builds a generation of professionals who can produce large volumes of text without developing the capacity to evaluate it.

I spent three decades in institutional communication — drafting statements for senior officials, advising delegations on written negotiating positions, reviewing the correspondence that organisations send to the world. I came to believe that the quality of an organisation's written

communication was the quality of its thinking made visible. Improving writing requires improving the thinking behind it.

What concerns me about the current moment is the idea that those AI tools can solve writing problems. They have, in many professional contexts, deepened it — by making weak thinking more presentable, by making every organisation's voice sound like every other organisation's voice, and by removing the productive discomfort of learning the art of writing.

This guide will help you to start reading your own words and reflecting on what they actually say to the recipient.

Read this book before you send your next important message. Over time, I trust you should notice better reception of your messages, and less clarification seeking exchanges.

Enjoy,

King Thomas
Senior Administrator,
TRANSFORMATIVEFIN HUB LEARNING CENTRE
2026

Preface

Somewhere in the past two years, a peculiar shift occurred in the correspondence I received from professional contacts. Emails from accomplished colleagues began to arrive written in a tone I could not place — fluent, properly structured, and somehow the person who had written it was not there.

In forty years of reading corporate and diplomatic communication — executive briefings, ministerial correspondence, board papers, negotiating positions — I had developed the ability to read a message and know, within the first paragraph, something about the mind behind it. Because writing is the residue of thought, and thought is individual. A senior procurement officer who invariably opened with the exception before the rule. The policy analyst who placed the qualifier at the end of every third sentence. The communications director who used the passive voice consistently when describing organizational failure, and an active voice when claiming organisational success. These were not intentional stylistic choices the writers could have named. They are natural "cognitive signatures" revealing a particular way a given mind engages with a problem.

That capacity for identification is now largely absent from professional correspondence. A recruitment consultant in Nairobi, a senior analyst in Singapore, and a policy officer in Brussels now send emails that could have been produced by the same author — because, in a practical sense, they were. The same tool, prompted by three different professionals, thus producing text that converges toward the statistical average of

professional language. This generates the most probable phrase, the most expected structure, and the most conventional argument. The resulting tone and style don't belong to, or represent, a specific individual mind.

This is the first problem this book addresses.

The second problem is harder to detect because it hides behind apparent competence. AI-generated text passes every surface test. A grammar checker will approve it and a colleague reviewing for typing errors will find none. As a limitation of how technology functions, AI text cannot produce evidence that the writer has direct encounter with the subject at hand. The text assembles the language of professional expertise without the professional encounter that gives that language meaning. Recipients who know the domain recognise this immediately. Either way, communication is failing at impact level.

Both problems are solvable and neither requires abandoning the tools. Both require editorial judgment; the capacity to read your own work and ask whether every sentence reflects actual knowledge, whether every word earns its place, and whether the whole communicates in a voice that is yours.

This book teaches that discipline.

The approach in this book derives from forty years of doing what mentors once did in person. Reading someone's draft, line by line, asking why each choice was made. For significant client communications, that process sometimes required four complete revisions before the message was ready to send. For diplomatic communications, the threshold was different, but the method was identical — every word under scrutiny, every claim supported, every construction required to justify its presence.

When I trained graduates in professional communication, I had them draft an email, print it, and bring it to my desk. We went through it line by line. I would stop at a sentence and ask: What does this mean? Sometimes the writer knew. Often the discovery was that the sentence said something other than what the writer intended — or that it said nothing specific at all. The discomfort of being asked to defend each word is the way by which writers learn to produce words they can defend.

With remote work, large teams, fast paced environments; and the false idea that AT tools can teach how write; the desk, pen and paper approach is no longer standard practice in many organisations. This book is meant to build your capacity to self-review your emails, reports, short form communication, and any form of writing, alongside AI tools.

The book is organized into eight parts. The first three establish the diagnosis: what has changed in professional writing, what the two central problems are, and how professional writers were developed before those problems existed. Part Four — the longest section — catalogues the nine AI augmented problematic writing patterns, explains how each pattern weakens a communication, and provides the test that identifies it and how to address them.

The second half of the book is applied. It illustrates how each principle is applied in email, formal reports, proposals, and short-form professional writing. The final section addresses the reconstruction of individual writing voice.

Each chapter contains a principle, the pattern, the test, and the correction. Read it alongside a draft you are currently preparing and apply each test in sequence. The improvement will be measurable from the first application.

Dozens of excellent books address foundational grammar composition and corporate style manuals. This book provides self-editing guidance built on the same editorial principles that governed professional writing development for forty years before AI tools changed the landscape. Those principles have not changed. What has changed is the scale and visibility of the problems they address.

Use it accordingly.

A. J. Thornfield
2026

PART ONE

THE DIAGNOSIS

What has changed

CHAPTER 1

Two Problems

In early 2024, a partner at a mid-sized professional services firm reviewed sixteen proposals submitted by candidates for a senior communications position. All sixteen demonstrated professional competence, and none was distinguishable from the others. Each opened with the same three-element structure: a contextual observation, a statement of challenge, and a proposed approach. Each maintained the same register — measured, organized, appropriately confident. The partner hired the candidate whose curriculum vitae indicated the most direct client experience, because the written submissions provided no differentiating information about the minds of the candidates. He later described the experience as reading sixteen documents that could have been produced by the same author.

This description, with minor variations in sector and context, has become a routine observation among professionals who review large volumes of written submissions; and on professional platforms such as LinkedIn. It reflects a condition that has become prevalent across professional disciplines in less than three years. The condition has two components and acknowledging them is the first step toward addressing either.

When an AI language model generates text, it predicts. Given a sequence of words and a context, it calculates the most statistically probable continuation. That calculation draws on patterns extracted from an enormous volume of text — articles, reports, proposals, correspondence, and academic papers — and produces output that resembles, at the surface level, a competent version of whatever it was asked to produce.

The output passes every surface test with correct grammar and appropriate structure. Vocabulary is within the expected range for the requested format. A grammar-checking tool will approve it, and the readability scorer will be high.

The output, however, will miss is an original encounter with a subject. The model has not attended the meeting being described. It has not reviewed the contract under discussion. It has not met or read the mood of the client being addressed. The result will therefore read as though it was produced by someone who has read extensively about a subject and knows very little about the instance at hand.

Three characteristics distinguish AI-generated text from text written by someone who has direct knowledge of the matter at hand.

First, the absence of the telling details. Human writers who know their subject include specific observations — the figure that diverged from the trend, the exception that required a workaround, the sequence of events that explains why the outcome differed from the projection. These details are not available to an AI model because they were not in the training data. The model fills their absence with general observations that sound specific.

The difference is detectable by any reader with direct knowledge of the domain.

The second is the predictable argument structure. AI-generated writing follows the most conventional argument pattern for its genre. This includes context, challenge, implication, and recommendation. This structure is not wrong — it was recommended as an effective way to present ideas, even before AI tools came into place. A human writer with direct knowledge of a situation will modify structure in some way to reflect the specific reasoning at hand. Expert writers reverse the conventional order when the evidence demands it. They subordinate the recommendation when the uncertainty about the challenge has not yet been resolved. AI models do not make these adjustments, because making them requires understanding the situation rather than its genre.

The third is the absence of position on contested questions. Human writers who know a field know where the disputes are. They know which interpretation of a dataset is challenged, which methodological choice has critics, which policy recommendation carries documented trade-offs. AI-generated text presents the conventional view without acknowledging the dispute, because acknowledging the dispute requires knowing the field well enough to locate it. The result is text that reads as settled on questions that, within the relevant professional community, remain open.

These three characteristics — the absent specific detail, the predictable structure, and the conventional view without position — produce text that reads pleasingly well but hollow. Recipients who know the domain can recognise the hollowness.

The first problem concerns quality. The second concerns identity.

When two million professionals across multiple industries use the same AI tool to draft their emails, proposals, and reports, their written output converges. The tool produces the most probable phrase for the given context, and the most probable phrase is, by definition, the phrase that appears most frequently across the training data. Every user drawing from the same probability distribution produces text that trends toward the same statistical centre.

The practical consequence is observable in any high-volume professional review environment or even platforms such as linkedIn. A procurement manager who processes three hundred supplier communications per month now encounters communications that share vocabulary, sentence rhythm, and argument structure across organisations she knows to be different in every other respect. A hiring manager reviewing applications for a policy position reads opening paragraphs that could be transposed between candidates without altering their apparent quality. A board member reviewing executive summaries from three different consultancies finds it necessary to read to the third page before the documents begin to diverge.

In forty years of reading professional communication — corporate correspondence, diplomatic cables, ministerial briefs, institutional reports — I developed the expectation that individual writers leave traces in their text. This may include be an observation that opens the argument, the specificity of the qualifier, or the choice to use a single clause where another writer would have used two. This is not about grammar or correctness. It's a residue of the way a specific mind (individual) engages with a specific

problem. This engagement is what gives professional communication its texture, its authenticity, and distinguishing character.

AI-assisted writing does not produce this engagement. It produces the output of a well-trained pattern recogniser, and pattern recognisers produce patterns.

WHY BOTH PROBLEMS MATTER

The two problems differ in their consequences, and the weight of each differs by context.

Empty fluency (hollow communication) can be costly in technical fields such as law, finance, engineering, policy, applied science. A client who receives a legal memorandum and observes that it summarises rather than analyses may not renew an engagement. The donor who reads a grant proposal and recognises that it describes a generic programme rather than a specific intervention will allocate the funds elsewhere. The commission that reviews a report and finds that it presents the conventional view on a technically contested question will return it with questions that expose the absence of specific knowledge. In these contexts, you cannot afford to have a hollow text. The failure to meet the substantive standard is more visible.

Convergence carries its greatest disadvantage in competitive contexts where differentiation determines outcomes. The job application that sounds like every other application, the proposal that reads like every other proposal, the message that leaves the recipient with no impression of the specific professional who sent it. In these contexts, every applicant and every competing firm has access to the same tool. The tool brings everyone to the same surface level, which raises the effective standard for anyone who wants to stand above it.

Every communication you produce is evaluated against the standard expected of you. The gap between what the text implies about the you and what you actually know is a professional risk that accumulates with each communication sent without thorough editorial scrutiny.

THE SOLUTION

The two problems described require careful editorial judgment to the text before it is sent or published.

Editorial judgment is the capacity to read what you have written — or what you have prompted — and ask, for every sentence, whether it reflects your actual knowledge of the situation, whether the structure serves the specific analysis of the situation, and whether every word is there because it is necessary. This capacity requires deliberate development. There are so many experts in their field who fail at effective writing and communication.

The chapters that follow will develop your editorial judgment capacity. Part One establishes the diagnosis in the detail required to apply the remedy correctly. Part Two reconstructs the mentorship method through which professional writers were historically developed — line by line, draft by draft — and makes it portable for individual practice. Part Three presents four editorial tests, each targeting a different dimension of the problem. Part Four catalogues the nine AI writing patterns that produce the problems described in this chapter, explains how by partner each weakens a communication, and specifies the correction.

The guidance in this book produces the improvement professional mentors of an earlier era produced at a desk.

CHAPTER 2

Fluent Is Not the Same as Good

The confusion at the centre of the current AI writing problem is a definitional one. Fluent and good are not synonyms. In most professional contexts, they are not even closely related.

WHAT FLUENCY IS

Fluency, in the context of written text, describes the surface qualities of a piece of communication. A fluent text has correct grammar, appropriate syntax, coherent sentence structure, and vocabulary within the expected range for the genre and register in which it is written. Fluency is detectable by automated tools. It is the quality that grammar checkers, readability scorers, and spell-checking software are built to evaluate.

Fluency is necessary. A text with persistent grammatical errors signals to the reader something about the care and attention the writer has applied, and that signal damages professional credibility before the content has been evaluated. A text with incoherence sentences that do not connect logically, burdens the reader and increases the probability that the communication will be abandoned before completion. Fluency is the minimum expectation below which professional writing should not fall.

Fluency is also attainable by machine. Every major AI language model currently in use produces fluent text reliably and at scale. The fluency problem in professional writing is, for practical purposes, solved. The issue or question now is whether the text is good.

Good professional writing demonstrates three qualities beyond fluency.

The first is accuracy of representation. Good professional writing represents situations, data, and positions accurately — not in the general sense of not containing false statements, but in the sense of representing the specific situation at hand with the level of detail and qualification that the situation requires. A report that describes a financial position in accurate general terms while omitting the item that changes the interpretation has failed to represent the situation accurately, even if every sentence in the report is technically correct. An email that summarises a conversation while omitting the specific point of disagreement has failed to represent the situation accurately, even if the summary is fluent and the general characterisation is fair.

Accuracy of representation requires direct knowledge of the situation being described. It cannot be generated by a model working from general patterns, because the specific detail that makes a representation accurate was not in the training data.

The second quality is logical coherence at argument level. The connections between ideas in AI-generated text are associative rather than logical. They reflect the patterns in which ideas appeared together in training data, not the logical relationships between the ideas themselves. Good professional writing builds arguments in which premises connect to conclusions in ways that can be traced and tested. The reader who is sceptical of the conclusion should be able to locate the specific premise the writer is relying on and evaluate it independently. The reader who accepts the conclusion should be able to explain the reasoning to a third party. A

fluent text can contain arguments in which the connection between evidence and conclusion is implied rather than demonstrated. This is the case where the writing creates the impression of reasoning without producing the substance of it. We will see some examples in the subsequent chapters.

The third quality is specificity of claim. Good professional writing makes claims that are precise enough to be challenged. An email that describes a project as progressing well is not a claim that can be evaluated, disagreed with, or built upon. An email that describes a project as three weeks ahead of the Phase 1 milestone with two deliverables pending sign-off is a claim the recipient can assess against their own records, verify through independent means, and respond to with specific questions. Most AI generated fluent text substitutes the lack of specific facts with general evaluations like significant progress, notable improvement, and strong performance. These convey positive sentiment while providing no information the recipient can act on.

READING THE DIFFERENCE

The difference between fluent and good is not always immediately visible, but it surfaces when the reader applies one specific question to each evaluative claim in the text: what would the writer need to know to have written this sentence?

Consider two versions of the same sentence in a project update email.

> **Version A:** *The procurement process has progressed well and we expect to meet the timeline.*
>
> **Version B:** *Three of the five supplier evaluations are complete. The remaining two are scheduled for completion*

by the fourteenth, which keeps the award decision on track for the original twenty-second target.

Both sentences are fluent. Both describe the same general situation. Version A is general evaluation. Version B describes the specific fact.

Version A could be generated by an AI model with no knowledge of the specific project. Version B could only be written by someone who knows which evaluations have been completed, which remain, when the remaining ones are scheduled, and what the original target date was.

Version A communicates nothing the recipient can verify, act on, or challenge. Version B communicates information the recipient can confirm against their own records, use in scheduling their own activities, and question if the figures differ from their understanding.

Recipients who receive general evaluations where specific facts would serve them perceive the writer, progressively, as someone who communicates at the surface level. That perception affects the trust placed in the writer's assessments and the weight given to their recommendations.

The benefit of good text (Version B) builds a reputation of writer whose communications are worth paying attention to.

THE SPECIFICITY TEST

The test for distinguishing fluent from good in any piece of writing is a single question applied to each evaluative claim: what specific fact does this claim represent?

The process has progressed well — what specific milestone or measurement does this represent? The situation presents challenges — which specific condition is the challenge, and what is its measurable effect?

We anticipate a positive outcome — which specific factor supports that anticipation, and on what evidence does it rest?

Where the answer to such questions exists, the claim can be rewritten to include the details. Where no answer exists, the claim should be removed.

When this test is applied to AI-generated text, it reveals where the model has substituted general evaluation for specific knowledge — and where the writer, before accepting the output, needs to edit it with the specifics.

Chapter 3 establishes how this kind of editorial scrutiny was once taught, and its absence has become a professional development gap.