A Systematic Method for Controlling AI Writing Style

Style is not a preference. It is a specification. Without one, the machine writes in a single voice.

Every large language model (LLM) produces prose with the same tells. The contrast rhetoric ('X isn't just Y, it's Z'), the deluge of dashes, the rule of three, the smooth pivot to a generic profundity — once a reader knows the pattern, it is unmistakable. Last year, Forbes' 'Seven Tells of AI Writing' catalogued them; the Field Guide to AI Slop arrived at the same diagnosis from a different angle. And the prose is now everywhere. A report by Graphite found that since November 2024, AI-generated articles on the internet have outnumbered human-written ones. A single, machine-uniform voice has become the majority voice of the written web.

The cost of that uniformity is not visible in any single sentence. It becomes visible when the default register is applied to writing whose distinctness was the analytical content. Consider a passage from Tricontinental dossier no. 69, *How the International Monetary Fund is Squeezing Pakistan* (2022):

> *Pakistan's balance of payments problems have long been managed by the country's main 'export' — its geostrategic value to Washington. When the United States invaded Afghanistan in 2001, it needed Pakistan's support, and so it removed its economic sanctions against the country and provided it with economic, security, and military aid. In the same period, owing to Pakistan's strategic importance in the 'War on Terror', the Paris Club rescheduled $12.5 billion out of the total of $13.5 billion of the debt that Pakistan owed to it.*

Ask a large language model to produce a paragraph on the same material — Pakistan, the IMF, US strategic interest after 2001 — and the output, however fluent, will be of a different kind. A representative run, produced by ChatGPT in May 2026, returned:

> *Following the 2001 United States-led invasion of Afghanistan, Pakistan became a central strategic ally in the so-called "War on Terror," a shift that significantly reshaped its external economic relations. In return for military cooperation and logistical support, Pakistan received renewed financial assistance, debt relief, and diplomatic backing from Western powers and international financial institutions. In December 2001, the Paris Club agreed to reschedule approximately US$12.5 billion of Pakistan's bilateral debt, including substantial concessional terms and extended repayment periods. This restructuring eased immediate fiscal pressures and helped stabilize Pakistan's foreign reserves, but it also deepened the country's long-term dependence on external creditors and geopolitical alignment with the United States. The post-2001 period thus illustrated the close relationship between strategic security cooperation and international debt governance in the contemporary global order.*

The factual material is the same. Some critical markers survive — the scare-quoted 'so-called "War on Terror"', the $12.5 billion figure, the named consequences of 'dependence' and 'alignment'. But the analytical core has dissolved. The dossier's

signature move — treating Pakistan's geostrategic value as the country's main 'export', a critical redeployment of commercial vocabulary that exposes the commodification of sovereignty — has no counterpart. The proportionality of the rescheduling ($12.5 billion *out of* $13.5 billion, almost the entire bilateral debt) is reduced to a single figure that no longer carries the political weight. The dossier's grammar, which positions Pakistan's economic situation as structurally managed by Washington before naming the relation as such, is replaced with active prose in which Pakistan 'became a central strategic ally' — an actor entering an exchange, not a country pinned by one. And the paragraph closes with the AI register's signature move: a balanced 'but also' formulation that absorbs the political critique into a wider perspective, followed by a pivot to a generic claim about 'the contemporary global order'. The model is not unable to be critical. It is unable to make the conceptual moves that turn critical vocabulary into a structural argument.

Vague instructions to the model — 'write formally', 'sound analytical', 'write like Tricontinental' — do not recover what was lost. When any instruction underspecifies, the model falls back on a default register: a statistical average of the text it was trained on, smooth, agreeable, and undifferentiated. The challenge for any research institution with a distinct analytical voice is therefore the same: how to instruct a model precisely enough that it produces writing in *that* voice, rather than the average voice. If a method could be found, the AI register — the default register that AI writing produces, borrowing the stylistics term for the variety of language a context of production generates — could be systematically eliminated.

The AI Register Is Structural, Not Incidental

Most people's first response to the AI register is to vary the prompt. If 'write formally' fails, try 'write like a senior policy analyst' or 'use shorter sentences.' Each instruction produces a marginally different result — yet the same machine smoothness persists underneath. The instinct is not wrong. The diagnosis is. The problem is not the instruction; it is what the model reaches for when any instruction falls short of full specification.

The Style Mimicry pattern in the published PLoP paper 'A Pattern Language for Knowledge Engineering with Large Language Models' (2025) offers a precise diagnosis. A large language model carries a default, neutral, helpfully bland writing style. Users can modify it with simple commands, but these commands often fail to capture the implicit rules of a genuinely specific style. The output 'feels wrong': it may technically meet the requirement (it is 'formal,' for instance) but lacks the authority, wit, or intimacy characteristic of the target voice. The problem is structural. The model has been trained on vast quantities of text, and its default style represents a statistical average of all that writing: smooth, competent, agreeable, and utterly generic. It is the prose equivalent of neutral grey paint: inoffensive, functional, and entirely forgettable.

The average is not politically neutral. The training corpora are overwhelmingly Northern, Anglophone, and produced by the institutions that build, host, and consume AI tools — the corporate web, the academic journal, the US policy blog. The default register is what writing sounds like at the centre of that corpus. For any institution whose voice has been formed at a different distance from that centre, the model's default arrives as someone else's accent.

Consider a researcher tasked with writing a policy brief. This genre demands strict stylistic discipline: highly formal, analytical, direct, concise. Policy briefs are written for decision-makers who have minutes, not hours. Every sentence must earn its place. When a researcher instructed the AI simply to 'write a formal brief on this issue,' the output was professional enough at first glance but missed the genre's critical features on closer inspection. It was too long. It lacked authoritative tone. It required heavy revision. The problem was vagueness of instruction. An abstract 'write formally' cannot convey the implicit rules accumulated over years of a particular writing tradition, the specific cadence of policy argumentation, the expectation that claims are supported by evidence rather than embellished with ornament, the unwritten requirement that a brief sounds as though its author has spent a career in government rather than a career writing blog posts. Nobody told

the AI how such a brief should read, so it retreated to its default style, producing text that was formally adequate yet generically flat.

This diagnosis points to a clear line of action. If the root cause of the AI register is default-style regression driven by vague instructions, then the solution is to decompose style from a holistic sensory impression into specific, operational descriptions, then feed those descriptions to the AI as precise instructions. This requires a systematic analytical framework capable of turning 'write it like this' from a wish into an engineering process.

Stylistics Converts Style from Intuition to Engineering

Consider again the researcher writing a policy brief. The first attempt, with an abstract instruction, produced text requiring wholesale revision. The second attempt changed approach entirely. The difference lay in the precision of the instruction. That precision came from a specific academic discipline.

Linguistics has a sub-discipline devoted entirely to the study of language style: stylistics. In *Stylistics: A Resource Book for Students* (2025), Paul Simpson proposes a seven-layer analytical framework that decomposes a text's style into seven observable, analysable linguistic strata: phonological patterns (the sound effects, rhythm, and prosody of text), graphological representation (visual presentation including punctuation and layout), morphological structure (how words are formed and inflected), lexical choice (diction characteristics such as formality and emotional register), syntactic structure (sentence organisation, length, and construction), semantic implication (literal and figurative meaning, metaphor and imagery), and pragmatic use (discourse-level features: tone, point of view, overall structure).

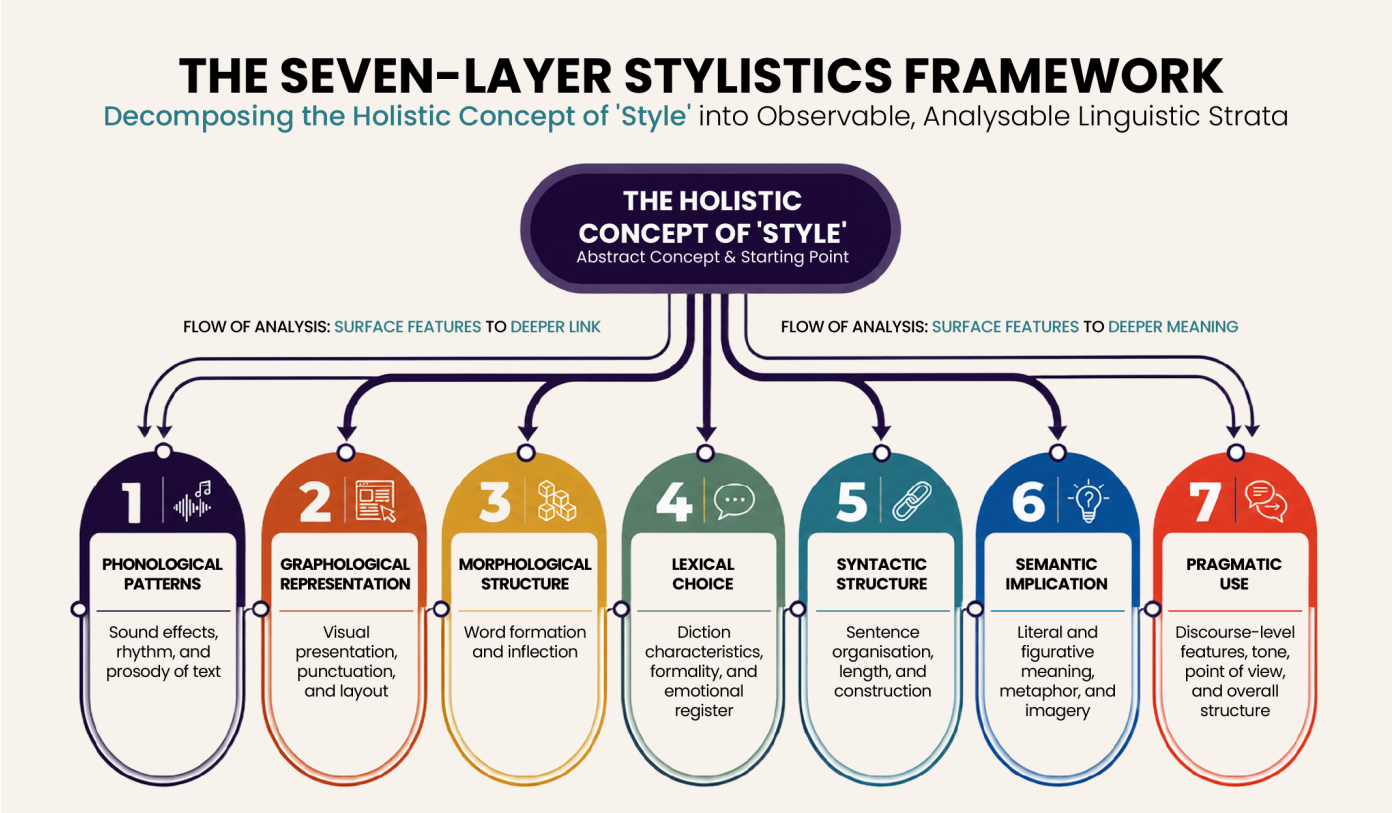

The seven-layer stylistics framework: decomposing the holistic concept of 'style' into observable, analysable linguistic strata

The framework's core value is this: it converts 'style' from a holistic, sensory impression into specific, operational descriptions. When applied to a target text, the seven layers yield a structured 'style profile' prescribing how each linguistic stratum should be selected. The transformation is decisive. Where previously a researcher could only say 'make it sound like *The Economist*,' now the instruction specifies exactly what that means at every linguistic level: what kinds of sentences, what register of diction, what rhetorical structures, what relationship between author and reader. The intuition 'I know it when I see it' is replaced by a specification that anyone, including a machine, can follow.

Imitating a style thereby becomes a two-step engineering process. Step one: analyse the target text with the seven-layer framework to generate a structured style profile. Step two: use the style profile as part of the instructions, directing the AI to produce new text according to the features prescribed at each layer.

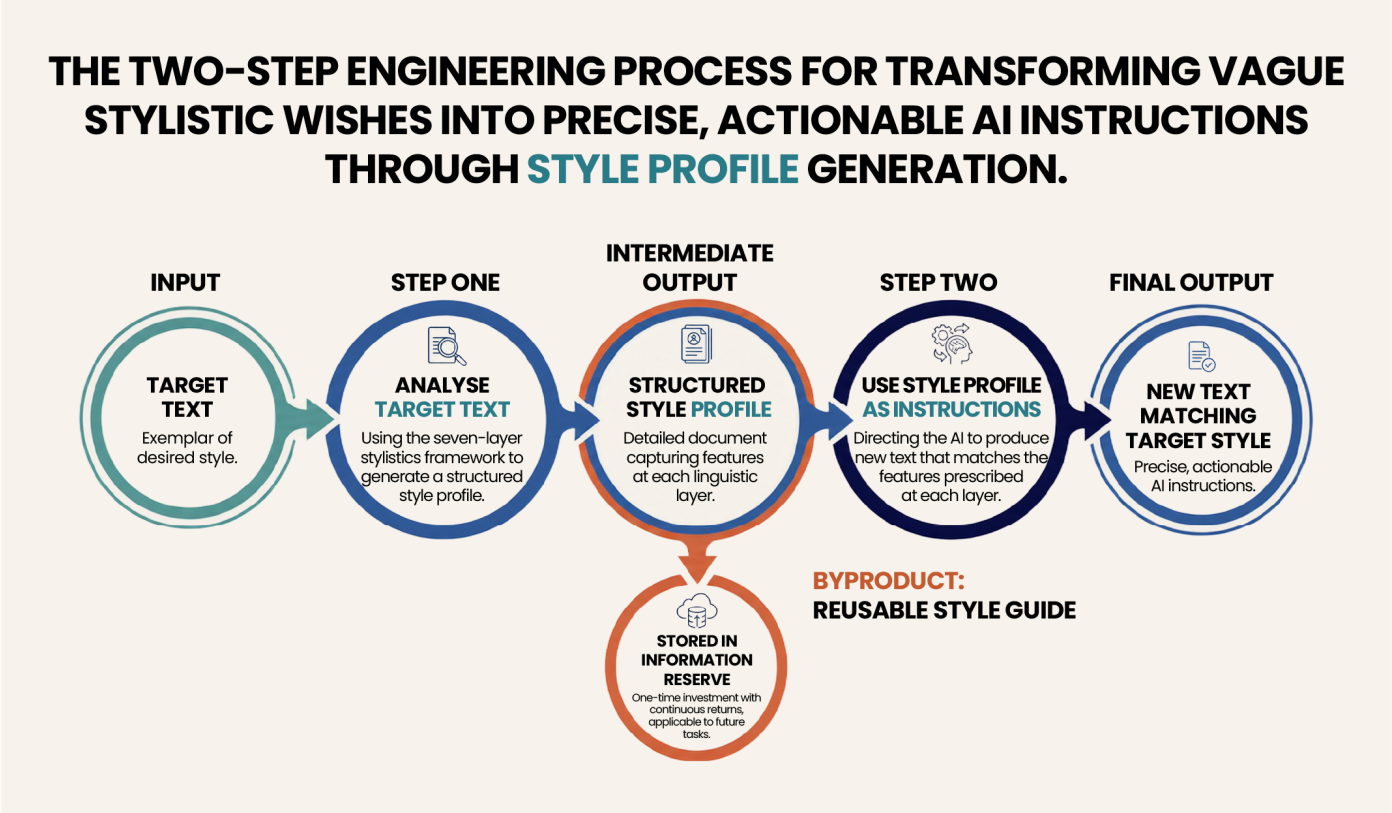

The two-step engineering process: transforming vague stylistic wishes into precise, actionable AI instructions through style profile generation

The policy brief experiment puts this framework to the test. The earlier attempt, using the abstract instruction 'write formally,' failed. The second attempt changed approach entirely. Two high-quality policy brief examples were provided, and the AI was asked to analyse them using the seven-layer stylistics framework and produce a detailed style guide. The guide made precise prescriptions at each layer. Under lexis, for instance, it noted that vocabulary is consistently formal and official, devoid of emotional language, preferring specific technical terms over generalities. Under graphology, it observed that numbered enumeration ('first, second, third...') imparts clarity, order, and a systematic impression. Under syntax, it reported that sentences are long and structurally complex, using formal logical connectors ('therefore', 'however', 'in view of') to pack considerable information into single sentences. This style guide was then fed back as part of the instructions for drafting a new brief. The result: structure, tone, and vocabulary closely matched the target genre. Editing time dropped significantly.

The style guide has another important property: reusability. Once generated, it becomes part of the information reserve, applicable to countless future writing tasks. The creation of a style profile is a one-time investment with continuous returns. For a research institute producing regular policy briefs, the same profile can guide dozens of writing tasks without modification, ensuring consistency across authors and over time.

A passage from Dalaya Ashenafi Esayiyas's 'Sovereignty or Surrender: Confronting Africa's Comprador Class' (2025), the seventh Pan-Africa Newsletter, provides a concrete example of the framework in action:

> *The trajectory of Ghana's political economy illustrates how a comprador bourgeoisie facilitated neoliberal subjugation, transitioning from Kwame Nkrumah's socialist Pan-Africanism to becoming an IMF 'success story'. The 1966 coup against Nkrumah, supported by CIA intervention as revealed in declassified documents, initiated the dismantling of his industrialisation projects. By the 1980s, Jerry Rawlings, once a revolutionary leader, adopted IMF structural adjustment programmes, privatising state enterprises and cutting social spending in Ghana. This shift was institutionalised through three key mechanisms. First, through policy capture, since the World Bank-trained Economic Management Team assumed de facto governance, sidelining ministerial authority. Second, through the implementation of debt as discipline, with IMF loan conditionalities prioritising raw cocoa and gold exports over industrial development, entrenching extractive dependency. Third, through elite co-optation that neutralised opposition, as Rawlings' former socialist allies were absorbed into consultancies and NGO roles, effectively sanitising dissent. Together, these processes reconfigured Ghana's economy to serve global capital at the expense of sovereign development.*

Applying the seven-layer framework to this passage extracts a precise profile. At the pragmatic layer: the passage names its structural diagnosis (comprador bourgeoisie as transmission belt for neoliberalism) in the opening clause, before any evidence is marshalled — the argument is the rhetorical premise rather than the rhetorical destination; the closing sentence ('Together, these processes reconfigured Ghana's economy to serve global capital at the expense of sovereign development') states the political consequence in plain terms, closing the argumentative arc rather than gesturing outward. At the lexical layer: 'success story' sits in scare quotes, critically redeploying the IMF's own triumphalist vocabulary to expose it as a designation of subjugation rather than achievement; analytical concept-coinages — 'policy capture', 'debt as discipline', 'elite co-optation', 'sanitising dissent' — compress each structural mechanism into a named category in the analyst's terms rather than borrowing the institution's framing; proper-noun precision (Nkrumah, Rawlings, the World Bank-trained Economic Management Team, the 1966 coup) anchors the argument in named historical actors and named institutional bodies. At the syntactic layer: sentences are long and subordinated, but the structure is enumerative ('First... Second... Third...') rather than causal-chronological — the argument is delivered as a typology of mechanisms, each named before it is explained, with embedded action verbs ('assumed', 'prioritising', 'absorbed', 'sanitising') that specify what each mechanism does rather than only what it is. Once extracted, this profile becomes the instruction set. The AI produces new writing in the source's register rather than its own default.

This process now has open-source tooling. The stylistics skill packages the seven-layer stylistic analysis into a single command. It runs inside an agentic AI environment such as Claude Code. Installation is straightforward:

```
npx skills add eXtremeProgramming-cn/stylistics
```

After installation, running `/stylistics extract article.md` against any text produces a complete style guide, including an analysis summary, normative rules, a quick-reference checklist, and a domain vocabulary. Researchers need not write their own analytical prompts; a single command extracts the style profile of any target text.

The detailed analytical prompt for the seven-layer framework was itself distilled by AI from Paul Simpson's textbook. A researcher guided the AI through the material to synthesise a structured, practical analytical method — academic knowledge transformed into a tool for AI itself. The AI for Social Science (AI4SS) framework — the Bandung Circuit's

intervention on research sovereignty in the age of AI — calls this 'full-disciplinary coordination': the discipline provides the analytical structure; AI provides the capacity to distil and execute it; neither could produce the instrument alone.

A style profile can make any voice more like itself. What it cannot determine is whether that voice has anything worth saying.

A Style Profile Answers Only the Fifth Question

With stylistics' systematic method and the engineering process of style profiling, the AI register appears eradicable. Yet eliminating the AI register may yield polished output that is still analytically empty, absent the insight that only a human perspective can provide.

In 'Knowledge Engineering: Six Critical Questions for Knowledge Production in the Age of Artificial Intelligence' (2026), 'form of expression' is positioned as the fifth question (the sixth — the researcher's distinctive contribution — is the meta-question the framework serves rather than a step within it). The same body of knowledge can be rapidly re-expressed in multiple formats: academic papers, policy briefs, video scripts, social media posts. But the first four questions — problem orientation, epistemological framework, information reserve, and methodology — define the supporting structure of knowledge. Possessing stylistic skill does not confer ideology, knowledge base, or analytical method. A style profile solves only the question of *how* to say something. *What* to say remains entirely the researcher's responsibility.

By way of contrast, the popular prompt engineering framework CO-STAR serves as a cautionary example. Its six elements are Context (background information), Objective (the task), Style (how information is presented), Tone (emotional quality), Audience (who reads it), and Response (output format such as JSON or prose). Four of the six elements (Style, Tone, Audience, and Response) address *how* it is said or formatted. What is missing is *what* to say: epistemological framework, information reserve, methodology, viewpoint and insight — the substantive dimensions that determine whether a piece of writing says something worth reading. The policy brief experiment illustrates the gap directly: the style profile specified sentence length, vocabulary register, and logical connectors; it said nothing about what position to take on the IMF's role in Pakistan's debt crisis, which figures to cite, or whose analysis to trust.

The Knowledge Engineering framework's placement of 'form of expression' as the fifth question — after problem orientation, epistemological framework, information reserve, and methodology — reflects a principle that runs deeper than the framework itself. Confucius spoke of *wén zhì bīn bīn* (*Analects* 6.18: '质胜文则野，文胜质则史。文质彬彬，然后君子') — refinement and substance held in proportion. 'Wén' (stylistic expression) and 'zhì' (substantive insight) are both indispensable, but they are not symmetrical. The priority of 'zhì' over 'wén' is the priority Marx would later name in materialist terms: "it is not the consciousness of men that determines their being, but, on the contrary, their social being that determines their consciousness" (*Preface to A Contribution to the Critique of Political Economy*, 1859). Substance grounds form, not the other way round. Words without elegance do not travel far. But elegance without substance lacks even the qualification to travel at all.

A style guide is itself a reusable asset: a researcher at a resource-constrained institution can develop one suited to the organisation's positioning, and the entire team can share it. But that guide cannot decide for the researcher whose side to stand on, what questions to ask, or which sources to trust. Style is 'form of expression,' an engineerable capability, but it cannot substitute for the more fundamental dimensions of knowledge engineering: problem orientation, epistemological framework, information reserve, methodology.

The Human-in-the-Loop Provides Purpose, Not Approval

The significance of stylistic tools becomes clear only within a larger picture.

The AI4SS framework proposes four pillar capabilities, one of which is 'full-channel output.' The core claim: a researcher's in-depth analysis can simultaneously reach different readerships as an academic paper, a policy brief, a newsletter column, a social media post, each channel demanding its own stylistic and format conventions. This means researchers need to find the appropriate expression for each outlet. Without stylistic tools, the researcher must either find a style-matched editor for each channel or spend considerable time polishing different genres alone. For research institutions in the Global South, where editorial resources are scarce and language barriers compound the challenge, this constraint has practical consequences: good analysis fails to reach the audiences that need it, weakened by the wrong packaging rather than by weak ideas.

Stylistic tools change the equation. Once a style profile is generated, it becomes part of the institution's information reserve — a reusable knowledge asset that can be applied to countless writing tasks and shared across team members. A resource-limited research institute can have one person develop multiple style profiles and distribute them across the team. The same analysis of IMF conditionality in Pakistan can be sent to progressive media readers in the Tricontinental dossier register, or delivered to a finance ministry desk in a policy brief. The knowledge is singular. The voices are plural.

AI handles the repetitive labour of stylistic polishing. The researcher handles what no machine can: deciding whose side to stand on, what questions to ask, which sources to trust. The AI4SS framework names this the 'human-in-the-loop' principle — not because the human approves the machine's output, but because the human's analytical work is what gives that output its purpose. Stylistics engineers *wén*. What the Global South needs from its research institutions is *zhì* — the analytical position, the political commitment, the knowledge of conditions no model was trained on.