Asia-Pacific Chiropractic Journal

CHIROPRACTICDIALOGUE

Avoiding Regression to the Mean: A millennial perspective on AI in Chiropractic writing and research

Clare McIvor

Abstract: Generative artificial intelligence offers practical support for chiropractic writing and research, yet its reliance on dominant patterns in existing data may constrain professional enquiry. This dialogue examines 'regression to the mean' as a metaphor for the homogenisation of ideas and language, considering its implications for chiropractic's salutogenic orientation, philosophical diversity, and commitment to individualised care. Drawing on published literature and professional writing experience, the author discusses linguistic convergence, cognitive offloading, knowledge cutoffs, and the potential loss of minority perspectives. Preliminary findings concerning critical thinking are considered cautiously, without assuming causation. Practical recommendations include reading primary sources before consulting automated summaries, verifying references, developing arguments independently, and using artificial intelligence to challenge rather than replace professional judgement. Deliberate prompting and authorial control are proposed as safeguards against generic framing and erosion of individual voice. The dialogue advocates responsible use of artificial intelligence as an adjunct to scholarship, while retaining human accountability, original research, and philosophical enquiry as essential foundations for chiropractic's continuing intellectual development and distinctive contribution to health and wellbeing.

Indexing Terms: Artificial intelligence; Chiropractic; Scholarly communication; Cognitive offloading; Critical thinking.

Introduction

Rarely, in my work life, do I glance at fewer than a dozen research papers in a day. Whether it be for Chiropractic blogs, case reports or other writing, I'm neck-deep in journals and primary evidence. A year ago, hours would have been spent finding these articles, reading halfway down, realising they're not what I'm looking for. Nowadays, those hours are better spent. I've programmed an AI machine that knows which journals to monitor for new Chiropractic content. It knows my voice, because I've trained it on outputs from thirteen years as a full-time, professional writer. I've also trained it to use specific terms, inflections and paradigms as I write for different clients.

> *… Our challenge, as the age of AI opens, is to make sure that we do not become trained by the average, by the regression to the mean. Instead, we must train it…'*



Also, I wince at AI slop, and spot it all too easily, as finding stock photography that reflects the Chiropractic mood is now impossible. AI has made that job harder, ironically. The old Nikon has been dusted off, demanding labour that used to be a two-minute search on a database. AI makes some things easier or better, and others harder and riskier.

I am an intermediate AI user, having learned a thing or two that can get better outputs than entry-level instructions, and I'll share some of those tips in the latter half of this article. But there is a strict limit to what AI can do for me. The rest, I must do myself. It can do literature searches, but I must make sure it didn't miss anything. It can summarise, but I must read and make sure it captured the essence of the article.

It can never be entrusted with my professional reputation. Nor can it be a philosopher.

> Currently, AI is working approximately one to two years behind us, not capturing new frontiers. If you want it to read a news article from today, you have to put that in and specify the request; otherwise, that data point will not feature in its outputs.

Knowing this, and knowing the point in history at which Chiropractic stands, we carry a responsibility. As Chiropractors, philosophers and researchers, it is our job to be front-liners - to think about what is possible and to create new and incisive research questions and create primary data. It is our job to recognise that research is both currency and protection when it comes to our right to practice, and that social media is the megaphone that reaches those who need us and hopefully changes their ideas about what health and the expression of life can be. It is our job to move beyond the reductive and into the expanse of possibility and human potential; to theorise and build our way towards what a human could be and do if their nervous system worked free of interference.

> This element of consciousness, humanness and innovation can never be entrusted to a machine.

In this issue, Ebrall (2026) sets out an educator's case for treating artificial intelligence as augmented intelligence: a collaborator that amplifies the judgement of the person using it, for better or worse. He also observes that AI is a prediction engine trained on the past, and that if we ask of it what everybody else asks, we will receive an average output.

It is this regression to the mean that presents the greatest threat to Chiropractic, which is, at its heart, a hands-on model built on connectedness, both neurophysiological and interpersonal.

AI is, at is most basic form, a large language model (or LLM). And a language model is, by design, an instrument of regression to the mean. It produces the most probable continuation of a text, and the most probable is the most central: the consensus framing, the familiar phrase, the position held by the largest share of its training data. For many tasks this is useful. We do need to know consensus data on safety, efficacy, best practice and evidence-based approaches. But for a discipline whose distinct contribution lies away from the centre, it is a fundamental misfit.

We have seen the clash in the constant juxtaposition of traditional medical research with Chiropractic research. The current approach to evidence in traditional research looks to exclude all explanations other than the tablet or surgery that forms the intervention. Did this tablet cure your *insert ailment here? Can we be sure it was that tablet?

Chiropractic turns this approach on its head. It looks at the whole person, their individuality, their life and all that is in it and then asks, 'how can we help? How can adjusting the subluxation remove interference and allow this person to better express innate intelligence?' We have not put a limit on what this might mean, nor should we.

Chiropractic has never been content with the mean. It is philosophically expansive, asking what health might be rather than only what disease is, and inviting enquiry into wellbeing, adaptability, and the relationship between the nervous system and human consciousness. A profession that increasingly reads, synthesises, and writes through AI risks having them smoothed toward the probable, one summary and one paragraph at a time.

But we are not meant to be reductive. Evidence-based, yes. Science-backed, yes. But always expansive.

I write as a millennial, from a generation that learned to research in libraries, then on search engines, and now with generative models, often within a single career. We remember reading papers in full because there was no other way to know what they said, and we have watched each new tool promise to spare us that effort.

The argument that follows has four parts: how the pull toward the mean operates, why it matters for Chiropractic in particular, how it reshapes the user as well as the text, and what a writer can do to resist it.

The pull toward the centre

Most of us will only ever use the LLM version of AI - that is, the large language models marketed for mass public use (that, and photo editing software with AI functions). A large language model generates text by predicting probable sequences of words from patterns in its training data. Bender et al. (2021) described such systems as producing linguistic form without access to meaning, and noted that the web corpora on which they are trained over-represent hegemonic viewpoints. Floridi (2023) characterised them as agency decoupled from intelligence: they act by producing text, without understanding what the text says. Hicks et al. (2024), drawing on Frankfurt, argued that a model is structurally indifferent to the truth of its outputs.

While generative AI is likely to learn and improve over time, and some of the predictions as to how they might develop lean towards the fearsome, these accounts converge on a single tendency. A model's default output gravitates toward the centre of its training distribution, and that centre is defined by whichever views are most frequently and most confidently expressed. Shumailov et al. (2024) demonstrated the tendency in its most extreme form. When models were trained on text generated by earlier models, the tails of the original distribution progressively disappeared, a process the authors termed model collapse. What was rare in the data became rarer, and then absent.

I saw this illustrated fabulously in an image of Dwayne 'The Rock' Johnson. AI was asked to recreate the image one hundred times. I sat and watched that video as The Rock became less human-looking, and then more menacing, and then took a turn into Picasso-esque cubism before degrading into blocks of colour no longer recognisable as a human face.

One could argue that a picture of a human face is already less than the experience of a human face. Thus, the visual illustration of this loss of humanity was even more impactful.

Chiropractic already lacks the large-scale funding that traditional medical research attracts. Ours is a grassroots movement wherein funding for projects comes from the very practitioners who consume this research and use it to inform their practice.

We are not a profession that exists within the mean

The challenge we face is that of creating more original research, more primary evidence gained from individual studies and cases, and we cannot afford to disappear through the data-destruction that occurs when averages are used repeatedly.

The same pressure bears down on individual writers. Kosmyna et al. (2025), in a preprint not yet peer-reviewed, found that essays written with an LLM were markedly homogeneous in vocabulary and topic structure. Sun et al. (2025) showed that each major model leaves a distinctive and persistent stylistic fingerprint. A writer who drafts with AI inherits not only its pull to the centre of gravity in ideas, but its way of expressing them.

This doesn't make the tools useless, but it does create limits around what they cannot do and how we can responsibly use them.

Philosophy is the practice of holding a position, being answerable for it, and following an idea beyond the point where consensus stops. A model can summarise any philosophical position fluently. It cannot hold one, or pioneer one, and its structure inclines it to describe the unusual in terms of the usual.

It has often been said that Chiropractic is an art, a science and a philosophy. Our challenge, as the age of AI opens, is to make sure that we do not become trained by the average, by the regression to the mean. Instead, we must train it. That requires humanness, originality, and indeed, primary research.

A discipline built in the tails

Chiropractic's philosophical inheritance is, almost by definition, a tail of the distribution. It is not talking about the majority of the bell-curve, in which the bulk of traditional medical research and experience exists. It is on the edges, if you will. Non-drug and non-surgery (where possible), hands-on, subluxation-based and salutogenic. It is not looking at pain and symptoms, but at function, structure and optimal neurological function enabling adaptability and expression of life. This is not to denigrate the valuable and life-saving contribution made by medical science and practice, merely to point out that we do not neatly fit that model and are often treated as the proverbial outlier.

The dominant biomedical model locates health largely in the absence of identifiable pathology. Chiropractic has historically asked a broader question about the conditions under which a person flourishes and has placed the nervous system and the organism's capacity for self-regulation at the centre of our practice.

Our profession is not alone in this. Canguilhem (1966/1991) argued that health cannot be reduced to statistical normality: a healthy organism is one able to establish new norms of life when circumstances change, rather than one that merely conforms to an average. Antonovsky (1987) proposed a salutogenic orientation, asking what moves people toward health along a continuum from ease to dis-ease, rather than only what causes disease. Huber et al. (2011) recommended reframing health as the ability to adapt and self-manage in the face of social, physical, and emotional challenges, noting that the World Health Organization's 1948 definition had already reached beyond the absence of disease. Thornhill (2020) has drawn on Canguilhem to propose vitalism as an ethos for Chiropractic, describing life as fundamentally original, adaptable, and unpredictable, and not fully explicable in reductionist terms. He also traces affinities between Palmer's thought and Bergson's account of consciousness, a reminder that the profession's founding questions extended well beyond musculoskeletal mechanics.

One need not accept every element of this tradition to recognise its expansiveness. It asks about possibility rather than only about deficit, and it treats health, wellbeing, and perhaps consciousness itself as open questions awaiting human answers. These are the questions a language model is least equipped to pursue. Asked about vitalism, a model will most likely reproduce the dominant framing of it as a historical curiosity or a scientific liability, because that is the narrative most heavily represented in its training data. Asked to write about health, it will default to the language of risk factors and symptom reduction. Neither response is malicious. But both are regressions to the mean.

This bears on a point Ebrall (2026) makes about publishing. He encourages the profession to place subluxation-centric work in the indexed literature so that it shapes how AI represents Chiropractic. I agree, with one emphasis drawn from Shumailov et al. (2024). That literature will shift the models only if it is distinctively human in its reasoning. A body of papers drafted by AI and lightly edited would feed the machine a version of itself, and the expansive positions the profession hopes to preserve would be among the first to be averaged away.

> I use AI every day. I argue with it. I ask it for structure. I ask it to check my logic on pieces I've already written, then I strengthen my argument. AI doesn't necessarily make my life easier: it merely removes the grunt work I'd spend wading through papers, writing and rewriting, and gets me to the point where I can argue, find fault, strengthen arguments and build.

Trainer or trained?

The thing that excites me most about being a Chiropractic writer is that I have the honour of partnering with chiropractors and organisations to bring more Chiropractic voices into the conversation. These are people ordinarily too busy, or too pen-shy, to write. Every time we do this, we put more data out there and feed the machine originality. We train it.

But allow me to share another human reflection: we can be trainers, or we can be trained. I recently watched a TikTok video wherein a writer-friend was raising her concern over the fact that she creates all her content without AI assistance, yet she is constantly being flagged as being largely AI-written. The alarm was genuine, and I believe that she does, in fact, write original content.

I also have to sit with the possibility that she may have been linguistically trained by reading other LLM-generated content. Given my role in communications, I've learned to spot it. ChatGPT likes saying things, 'In the vast tapestry of life' at the beginning of essays. Claude likes saying, 'It's not this; it's that' or 'This deserves to be stated plainly'. My love of em dashes and the Oxford comma has been challenged by the fact that AI also loves them, a little too much, and the humans who use them have to accept that they're going to look a little machine-written. But I'll keep the Oxford comma, thank you very much. That is a hill I'll die on.

My point is that when we hear or read common sentence structures, we become trained by them. They become normal, and then we repeat them in our own human linguistic schema. The machine isn't sounding like us; we are sounding like the machine. Then comes the issue of cognitive offloading.

Risko and Gilbert (2016) defined cognitive offloading as the use of physical action to reduce the cognitive demands of a task, such as writing a reminder rather than holding it in memory. Offloading isn't new, and it is mostly benign. What has changed is the kind of cognition now being offloaded. Phone storage takes over memory of phone numbers. A calculator takes over arithmetic; a language model can take over reading, synthesis, and argument, which are the activities through which a researcher's judgement is formed.

The early evidence warrants caution. Gerlich (2025), in a survey of 666 participants with follow-up interviews, found a significant negative correlation between frequent AI use and critical thinking scores, mediated by cognitive offloading. Younger participants reported greater reliance on AI and scored lower on critical thinking. The design was correlational, though, and thus it remains possible that people with weaker critical thinking skills are simply more inclined to use these tools. Lee et al. (2025), surveying 319 knowledge workers, found that higher confidence in generative AI was associated with less critical thinking, while higher self-confidence was associated with more. They also observed that critical effort shifted from producing work toward verifying and integrating what the AI produced. Kosmyna et al. (2025) reported that participants who wrote with an LLM showed weaker neural connectivity on EEG than those writing unaided, reported lower ownership of their essays, and had difficulty quoting what they had just written.

These findings are preliminary, of course, but they suggest that the effort we save is not always effort we can spare. An uncomfortable reflection that runs parallel is this: Chiropractic is a profession in which more critical thinking, not less, is needed. Critical thinking in philosophy, research and the art (the adjustment) we use in practice is what will protect us and allow us to keep and expand our practice going forward.

Here the regression to the mean moves from the text to the person. Ebrall (2026) recommends AI summaries of papers for learners, and also urges readers to build judgement by reading deeply, recommending the Index to Chiropractic Literature for the purpose. I agree most heartily. I never only read a summary, and the differences between the summary and the real deal are often striking.

When an LLM summarises a paper, it decides what is salient, and its sense of salience is drawn from the centre of its training data. A finding that challenges the consensus may be softened or omitted; a speculative discussion section may vanish entirely. If I read only the summary, I absorb the model's priorities along with its content. Repeated across hundreds of papers, those priorities gradually become mine, and my own thinking drifts toward the same centre. That is the sense in which the user is trained by the tool.

A striking example is a powerful paper by Haavik et al (2024). The title was 'Neuroplastic responses to Chiropractic care: broad impacts on pain, mood, sleep and quality of life'. The AI summary captured the elements in the title. The paper itself was astounding, but for not necessarily for those reasons. It spoke about the default mode network, the part of a person that is active when they are daydreaming, mind-wandering or generally not focused on a specific task. This is where the person makes meaning, ruminates, and self-reflects. In an interview with ASRF, (Australian Spinal Research Foundation, 2026) Dr Haavik said it '*is how Chiropractic touches the soul*'.

AI did not understand that. It was new information, outside the mean, and thus forgotten. As a Chiropractic writer, and as a person who benefits from Chiropractic care, the lights came on for me. This was the moment science and philosophy met, and the conversation about Chiropractic and human consciousness lit up.

Sequence matters, as does extra reading. A researcher who reads a paper first and then asks for a summary can use the comparison to see what the model flattened, and the gap between the two. A researcher who reads the summary first, and perhaps never reads the paper, has outsourced the formation of judgement itself.

> The onus for this can't be placed on the machine. It is ours, as chiropractors, researchers, philosophers and writers, to make sure we are the trainers, not just the trained.

The lag time and its detriment

Every language model is trained on data up to a cutoff date, released some months later, and then used for a year or more after release. Unless it is searching the web at the time of the query, its picture of the literature is therefore already dated when it answers.

The gap may be larger than vendors report, too. Cheng et al. (2024) found that the effective knowledge cutoff of several models, meaning the period their knowledge most closely matched, often differed markedly from the reported cutoff, largely because training corpora contain older versions of documents. In practice, a writer using a model without live search may be consulting a research assistant whose reading stopped one to two years ago.

This has a significant consequence for a field like Chiropractic. A model asked to review the evidence on a Chiropractic topic may omit the most recent trials, cite a superseded systematic review, or describe a debate as it stood before the latest contribution. It will do so fluently and without flagging the gap. The mean toward which it regresses is therefore also a dated one: the consensus as it stood, not as it now stands.

The posture I would suggest follows from this. When writing with AI, assume that you are the most current reader in the conversation. Treat the model's account of the literature as a starting map drawn some time ago, useful for orientation and unreliable for detail. Go to the primary databases for anything recent, and expect to find work the model does not know exists. More broadly, write from your own position first and consult the model second. A draft that begins with the writer's argument and is then tested against AI retains its author. A draft that begins with the model's framing tends to keep it, however thoroughly the sentences are later revised.

Practical section I: recognising unedited AI text

> I'll now shift from the polemic element of this article and invite you to a more practical overlay. We can't necessarily avoid AI. When you type something into a search engine these days, you will get an AI-powered, general response. To get the best out of AI, I'd suggest being more proactive. Use it intentionally, not in an ad hoc manner.

AI-assisted writing is already measurable at scale in the research literature. Kobak et al. (2025) analysed more than 15 million PubMed abstracts and estimated that at least 13.5% of those published in 2024 had been processed with an LLM, identified through an abrupt rise in certain style words. Liang et al. (2024) estimated that between 6.5% and 16.9% of peer review text at several major computer science conferences had been substantially modified by an LLM. Different models also leave different fingerprints. Sun et al. (2025) trained a classifier that distinguished text produced by ChatGPT, Claude, Grok, Gemini and DeepSeek with 97.1% accuracy, and found that the distinguishing patterns persisted even when the text had been rewritten by another model.

No single marker proves that a text was machine-written, and human writers used every one of these words before 2022. It is the clustering that tends to give a text away. Read together, these markers are the audible signature of the mean: the most probable words and structures, chosen because they are probable rather than because they are apt. The markers below fall into three groups: vocabulary documented in the research literature, structural habits, and features I associate with particular models from my own reading and editing.

Vocabulary with documented excess use

Marker	Evidence	Example
delves, *underscores*, *showcasing*	In 2024 abstracts, *delves* appeared at roughly 28 times its expected frequency (Kobak et al., 2025)	'This study delves into the role of spinal adjustment …'
potential, *findings*, *crucial*, *comprehensive*, *notably*, *insights*	Common words with marked excess use in 2024 abstracts (Kobak et al., 2025)	'These findings offer crucial insights …'
meticulous, *intricate*, *commendable*	Rose 34.7-, 11.2- and 9.8-fold in LLM-modified peer reviews (Liang et al., 2024)	'The authors present a meticulous and commendable analysis …'
pivotal, *intricate*	Among 291 excess style words identified by Kobak et al. (2025)	'Plays a pivotal role in the intricate interplay …'

Structural habits

There are also a number of linguistic tells I have observed across AI models. They're not always possible to avoid in human writing, but excessive use can give readers a feeling that they are reading something AI-generated.

- Contrastive framing. 'This is not merely X; it is Y'. A claim is introduced by denying a weaker version nobody proposed.
- Triplets. Lists, adjectives and clauses arrive in threes with unusual regularity.
- Announced emphasis. Phrases that tell the reader a point is important rather than showing it, such as 'it is worth noting', 'importantly , or 'this deserves to be stated plainly'.
- Even rhythm. Paragraphs of similar length, each opening with a topic sentence and closing with a summary line.
- The tidy close. A final paragraph beginning 'Ultimately' or 'In conclusion' that restates the piece and gestures toward the future.
- Hedged balance. Every claim offset by its counterpoint, so that no position is quite taken.
- Dense punctuation habits. Heavy use of the em dash and the colon to introduce a reveal.
- Report formatting in prose genres. Bold subheadings and bullet points in a piece that calls for continuous argument.

There are model-specific tells as well, and it is to the ire of some old-school grammar-heavy or grandiloquent writers that we wince and try to write differently in order to not sound like AI. Thus, the purpose of the above list is for example only, not punitive or a declaration of AI-authorship guilt. A writer who recognises these patterns in their own drafts can remove them, and in doing so is forced to rethink their writing and oftentimes, their point.

Practical section II: training, arguing with, and learning from AI

The key problem before us is not to let AI do the heavy lifting for us, but to give us a better springboard to original thinking. The methods below are offered in that spirit. They assume that the writer remains the author of record, as the Committee on Publication Ethics (2023) requires, and that every reference is checked against the primary source.

Train it before you use it

Most platforms now allow standing instructions that apply to every conversation, such as custom instructions in ChatGPT or project instructions in Claude. Both are behind the paywall, but are worth the investment as it allows a deeper exploration of AI's capability, and protects your control over the machine. These custom or project instructions should be set carefully. A useful set might specify the audience and tone, require that uncertain claims be flagged, prohibit invented references, and list phrases to avoid. Supplying two or three samples of your own published writing gives the model a reference point for voice. None of this guarantees compliance, but it shifts the default output away from the generic.

An example standing instruction

> Write for a peer-reviewed Chiropractic audience. Prefer authoritative understatement to emphasis. Do not cite any source you cannot name precisely; if unsure, say so. Avoid the words delve, crucial, pivotal and intricate, and avoid contrastive constructions of the form 'not X but Y'. Flag any statement about research published after your training cutoff.

I have many projects set up in the back end of my AI program of choice. In each, I have a tone guide that I authored. Each project has granular detail as to target audience, colours, fonts, phrases, sources, and some have exhaustive lists of my own authored pieces to train the LLM's voice. There are specific instructions around punctuation, cross-checking, flagging things, and specifically instruct the program not to fill in the blanks but to flag where it does not know something for sure.

The latter occurs frequently. We must all beware of the blind acceptance of an AI hallucination wherein it just makes up answers and fabricates citations to back it.

Argue with it

Language models trained on human feedback show a measurable tendency toward sycophancy, favouring responses that agree with the user over responses that are correct (Sharma et al., 2024). Left to its defaults, a model will tend to tell a writer that the argument is sound. One remedy is to ask explicitly for disagreement.

- Ask it to argue the opposing position as strongly as a capable critic would.
- Ask it to identify the weakest claim in your draft and explain why.
- Ask it to respond as a sceptical peer reviewer from outside the profession.
- When it agrees with you, ask what evidence would change its answer.

The value lies less in the model's objections, some of which will be poor, than in the discipline of having to answer them.

Ask for the tails

Because a model's default is the centre, the writer has to ask deliberately for what lies away from it. Useful requests include: which minority or dissenting positions exist on this question, and who has argued them; how a vitalist, a salutogenic, or a phenomenological framing would approach the same problem; and what the dominant framing leaves out. The answers will need checking, as always, and some will be thin. The point is to use the tool to map the margins rather than to settle in the middle.

Seek feedback without surrendering the draft

When asking for feedback, ask the model to mark rather than fix. This is where you, as a human, must intervene. A request such as '*identify where the argument loses coherence, but do not rewrite*

anything' keeps the revision in the writer's hands. Rewritten paragraphs are tempting to accept whole, and accepting them is how a writer's voice is gradually replaced.

Read first, then compare

For literature, reverse the usual order. Read the paper, write a short summary in your own words, and only then ask the model for its summary. Where the two differ, return to the paper to see who was right. Used this way, AI becomes a check on your reading, and the reading remains yours.

My own sequence

1. Use AI to propose a structure and to identify candidate sources.
2. Locate and read the sources myself, discarding those that do not exist or do not say what was claimed.
3. Check the reference list from the sources and include any that I think AI missed.
4. Write the draft from my own notes.
5. Ask AI to challenge the argument and to mark weak passages.
6. Revise by hand, then check the final text against the markers listed above.
7. Head to my frequented journals to see if there's anything I missed and can include.
8. Redraft

AI doesn't necessarily save me labour. It redistributes the labour from the grunt work of finding sources to the quality-improving act of scrutinising them and checking my logic.

Conclusion

As Chiropractors, researchers, philosophers and Chiropractic writers, mastery of AI lies in the refusal to let the profession's thinking be averaged. The habits involved are unglamorous: reading the paper, writing the first draft, and deciding for oneself what an argument means. They are also the habits through which a discipline keeps its capacity to think beyond the consensus and thus discover new horizons.

Chiropractic's contribution has always depended on asking larger questions than the prevailing model asks, about health as a positive capacity, about adaptation and self-regulation, and about the place of consciousness in human wellbeing. Those questions will not be answered by a system whose nature is to return the most probable response and errs on the side of regression to the mean. Our challenge is to stay in the tails of the distribution, on the edge of the bell curve, to read widely enough to know what lives there, and to write in their own voices about what we find.

My generation adopted each new tool as it arrived and can still recall what each one cost. The present tools are the most capable yet, which makes the choice about what to keep for ourselves the most consequential. Tis a decision and a process we must navigate with intention.

Declaration of AI use

During the preparation of this work the author used Claude Opus 5.5 to propose an initial structure and to identify candidate sources. All sources were located and read by the author. The author wrote and revised the text and takes full responsibility for the content of the published work.

Clare McIvor
BBus(Admin),
GD Comms(ProfWrit,Edit),
GD(Psych)(Cand)
Chiropractic writer
clare@claremcivor.com

Cite: McIvor C. Avoiding Regression to the Mean: A millennial perspective on AI in Chiropractic writing and research. Asia-Pac Chiropr J. 2026;7.3. https://www.apcj.site/McIvorAIDialogue.pdf

References

Antonovsky, A. (1987). Unraveling the mystery of health: How people manage stress and stay well. Jossey-Bass.

Australian Spinal Research Foundation (2025). When Neuroscience and philosophy connect – the default mode network and the Chiropractic adjustment. Australian Spinal Research Foundation's Research for Practice Podcast. Spotify.com https://open.spotify.com/episode/3JjEk2txmUeMg6UIHF1FoW

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Canguilhem, G. (1991). The normal and the pathological (C. R. Fawcett, Trans.). Zone Books. (Original work published 1966)

Cheng, J., Marone, M., Weller, O., Lawrie, D., Khashabi, D., & Van Durme, B. (2024). Dated data: Tracing knowledge cutoffs in large language models [Conference paper]. First Conference on Language Modeling (COLM 2024). https://arxiv.org/abs/2403.12958

Committee on Publication Ethics. (2023, February 13). Authorship and AI tools [Position statement]. https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools

Ebrall, P. (2026). Humanism in education and our embrace of AI. Asia-Pacific Chiropractic Journal, 7(3).

Floridi, L. (2023). AI as agency without intelligence: On ChatGPT, large language models, and other generative models. Philosophy & Technology, 36(1) https://doi.org/10.1007/s13347-023-00621-y

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies, 15(1), https://doi.org/10.3390/soc15010006

Haavik, H., Niazi, I. K., Amjad, I., Kumari, N., Ghani, U., Ashfaque, M., Rashid, U., Navid, M. S., Kamavuako, E. N., Pujari, A. N., & Holt, K. (2024). Neuroplastic Responses to Chiropractic Care: Broad Impacts on Pain, Mood, Sleep, and Quality of Life. Brain sciences, 14(11), 1124. https://doi.org/10.3390/brainsci14111124

Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. Ethics and Information Technology, 26(2). https://doi.org/10.1007/s10676-024-09775-5

Huber, M., Knottnerus, J. A., Green, L., van der Horst, H., Jadad, A. R., Kromhout, D., Leonard, B., Lorig, K., Loureiro, M. I., van der Meer, J. W. M., Schnabel, P., Smith, R., van Weel, C., & Smid, H. (2011). How should we define health? BMJ, 343. https://doi.org/10.1136/bmj.d4163

Kobak, D., González-Márquez, R., Horvát, E.-Á., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. Science Advances, 11(27). https://doi.org/10.1126/sciadv.adt3813

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task (arXiv:2506.08872) [Preprint]. https://arxiv.org/abs/2506.08872

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (Article 1121, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

Liang, W., Izzo, Z., Zhang, Y., Lepp, H., Cao, H., Zhao, X., Chen, L., Ye, H., Liu, S., Huang, Z., McFarland, D., & Zou, J. Y. (2024). Monitoring AI-modified content at scale: A case study on the impact of ChatGPT on AI conference peer reviews. Proceedings of Machine Learning Research, 235, 29575–29620. https://proceedings.mlr.press/v235/liang24b.html

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. Trends in Cognitive Sciences, 20(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S. R., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2024). Towards understanding sycophancy in language models. In The Twelfth International Conference on Learning Representations (ICLR 2024). https://proceedings.iclr.cc/paper_files/paper/2024/hash/0105f7972202c1d4fb817da9f21a9663-Abstract-Conference.html

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631(8022), 755–759. https://doi.org/10.1038/s41586-024-07566-y

Sun, M., Yin, Y., Xu, Z., Kolter, J. Z., & Liu, Z. (2025). Idiosyncrasies in large language models. Proceedings of Machine Learning Research, 267, 57854–5. https://proceedings.mlr.press/v267/sun25z.html

Thornhill, J. T. (2020). A vitalism ethos and the Chiropractic health care paradigm. Journal of Chiropractic Humanities, 27, 59–81. https://doi.org/10.1016/j.echu.2020.10.004