Editorial – The Moon Under Water

Daniel McDonald
Librarian, Darling Downs Health | Editor, JoHILA
Daniel.Mcdonald@health.qld.gov.au | https://orcid.org/0000-0001-8385-3671

In an essay published in the Evening Standard, 9 February 1946, George Orwell described his perfect pub, the "*Moon Under Water*". Among his ten stipulations were "You cannot get dinner at the Moon Under Water, but there is always the snack counter where you can get liver-sausage sandwiches, mussels (a speciality of the house), cheese, pickles and those large biscuits with caraway seeds in them which only seem to exist in public-houses". Each to their own, no doubt.

This issue of JoHILA is (mostly) focused on Artificial Intelligence. One wonders what more could possibly be said, or whether further commentary is even helpful. However, its ubiquity and its promise and its peril demand attention, analysis, maybe acclaim, maybe alarm. In the spirit of Orwell I prompted several LLM AI tools to adapt Orwell's task and to "Design the ideal hospital library".

CoPilot suggested "The ideal hospital library is not simply a room full of books. It is a clinical knowledge hub, learning centre, research support service, and wellbeing space that improves patient care, education, and innovation across the hospital." It concluded with a gold standard concept: "The truly ideal hospital library functions less like a traditional library and more like a Clinical Knowledge Centre, where librarians are embedded in healthcare teams, digital resources are available anywhere and anytime, research support is proactive, and every service contributes directly to safer, evidence-based patient care. For a modern health service, this model transforms the library from a support department into a strategic clinical asset." Solipsistically it also mentioned AI-assisted literature searching, AI literacy education, and an AI-powered discovery platform. Interestingly under key resources it listed UpToDate but not Dynamed, which will elicit no further editorial comment.

ChatGPT wanted to meld a patient library with a clinical library, including in its suggested layout a children's corner and a central reading garden. Helpfully it did suggest the most important feature were librarians. It described the ideal experience as "A patient enters → finds a welcoming reception → can browse independently → discovers trustworthy information → sits somewhere peaceful → gets help if needed → leaves feeling more informed and less overwhelmed.
Meanwhile, a clinician can enter from another part of the library, search the latest evidence, collaborate with colleagues, and return to work without disturbing patients.

In short: the ideal hospital library is simultaneously a knowledge centre, learning laboratory, sanctuary, and bridge between healthcare professionals and the people they care for."

Claude wanted plain-language health information organized by condition, not by Dewey decimal — think "Newly Diagnosed: Diabetes" not "616.4". It did include UpToDate AND Dynamed in its suggested point-of-care database access, which will elicit no further editorial comment. It was also keen to point out the one thing many hospital libraries get wrong (using a first-person pronoun no less): "They optimize for the collection and forget the human interface. The best hospital libraries I'd design around are staffed by someone whose job is literally to sit with a scared family member and help them find something true and readable about what they're facing — that single role does more for the mission than any amount of shelf space."

Gemini asked me straight up "What's the vibe, Daniel?", which seems entirely unnecessary. Nevertheless, it suggested "The ideal modern hospital library is no longer just a quiet room full of dusty medical texts; it is a dynamic, multi-purpose hub designed to serve three distinct audiences with vastly different needs: clinical staff (physicians, nurses, researchers), patients and their families, and administrative/educational staff." It also mentioned UpToDate and BMJ Best Practice (only), alongside AI search facilitation and an AI-powered knowledge discovery layer.

Perplexity was the only tool to include pictures… and references to sources. It summarised the ideal hospital library as "not primarily a book room. It is a clinical knowledge hub: a highly accessible, trusted service that helps clinicians make better decisions, supports education and research, and gives patients understandable health information—at the point and pace of care". It did suggest that librarians could help "identify hallucinations or outdated guidance" and was generally very bullish about the core role of the librarian: "The library's real collection is therefore both digital and human: licensed evidence resources plus librarians skilled at finding, evaluating, synthesising, and translating information." It concluded rather poetically: "A useful standard is this: a nurse on night shift, a registrar at the bedside, a researcher planning a review, and a family trying to understand a diagnosis should each be able to get trustworthy help quickly—without needing to know where information lives or how to search for it. In short, build a calm, flexible physical refuge—but fund it as a high-impact clinical information service. The shelves are optional; expert mediation, reliable digital access, inclusive learning space, and integration into care are not."

Grok was (honestly, not a joke) "experiencing issues. We are working on restoring service as quickly as possible." It also wanted me to "sign up to continue seamlessly with Grok's full power." I did not. I did wonder if Tesla cars and SpaceX rockets also "experience issues".

All of these answers (apart from Grok) were perfectly perfect in their summation of what health librarians would typically want in an ideal space and service. None of them has the elan of a George Orwell essay, unable to suggest the equivalent of not simply a snack counter but a snack counter where you can get liver-sausage sandwiches, mussels (a speciality of the house), cheese, pickles and those large biscuits with caraway seeds in them which only seem to exist in public-houses.

Therein lies the conundrum. The technology is obviously extraordinary, with profound implications for the way all of us generate, organise, retrieve, and understand information. It also traduces vast swathes of (non-reimbursed) human creativity to a polished but soulless blancmange, a very convincing parrot, but a stochastic parrot nonetheless. Maybe, probably, none of this matters. It is an information technology, and it will find its place, which is probably at the centre and at the extremes. Gutenberg's printing press and Turing's computer and Jobs' iphone all contain multitudes. Technology absolutists like Elon Musk and Peter Thiel and Marc Andressen spout forth on the enormous existential implications of AI. While on Substack Sam Kriss posts "If you let AI do your writing, I will come to your house and kill you. Did you think I wouldn't be able to tell? I can tell. I hate it. I find it viscerally disgusting; a cold shudder like someone's poured jelly down the back of my neck. I hate that it's everywhere; I hate that when I read basically anything now I'm constantly on alert, is this thing really what it says it is? Is this person actually a robot in disguise?". On the Red Hand Files Nick Cave posts "ChatGPT has no inner being, it has been nowhere, it has endured nothing, it has not had the audacity to reach beyond its limitations, and hence it doesn't have the capacity for a shared transcendent experience, as it has no limitations from which to transcend. ChatGPT's melancholy role is that it is destined to imitate and can never have an authentic human experience, no matter how devalued and inconsequential the human experience may in time become." At the New Yorker Jay Kaspian King writes of the despair of the professor in the age of A.I., asking "was it always the case that half of our students would cheat if it were easy enough?"

Meanwhile in health, clinicians are finding the middle ground, adopting these new tools and adapting to them, finding efficiencies and confirming treatment plans and opening up new patterns of thought and possibilities for diagnosis and intervention. All the while confronting the intractable present realities of bed block and resource scarcity, of perinatal and palliative care, of disease and death. In health libraries tools are being used in all sorts of ways, many described in the articles in this issue, while caution and discernment are also preached and practiced.

My son is five years old. Of late when visiting his grandparents he bounds in to his grandmother's bedroom to say hello. Initially I thought this was charming and a sign of good parenting resulting in intergenerational bonhomie. Until, that is, I cottoned on that his grandmother reads abundantly on her Surface tablet, and also lets him

watch dinosaur videos (and who knows whatever else is algorithmically served up), so he is actually chasing the screen and not (only) his grandmother's affection, which is probably a sign of very bad parenting. And yet, how are screens and their attendant information streams to be avoided? The dining table was maybe once sacrosanct, but screens are necessary to adjust the volume of music playing through the fridge and to talk about the photos on the app from that day's kindy summary. He, like me, like we, will have to navigate a world where the tools are ubiquitous, but there are very blurry lines between clean and clear information and mis and disinformation, where the boundaries between hallucinations and hard scholarly labour are stretched and warped, where creativity and imagination and thinking are at once enhanced and greatly challenged.

A recent report of MIT's Ad Hoc Committee on AI use in Teaching, Learning, and Research Training, published August 13 2026, is perhaps helpful in this navigation. It lists its guiding principles to: be humble; be bold; put humanity front and center, lean into learning; teach with intentionality; no one size fits all; augmentation not automation; think beyond the classroom and campus. Its three recommendations are:
1. Adapt educational processes for an AI-aware world
2. Center people, community, and the residential experience
3. Build processes, teams, and tools for continuous reflection, iteration, and improvement

And perhaps the last word needs to go to Orwell, who did so much in his writing to alert the world to the power, and danger, of words and ideas unmoored from morality and humanity, of artifice masquerading as intelligence. "In prose" he says in his influential essay *Politics and the English Language*, "the worst thing you can do with words is to surrender to them."

I commend the articles in this issue to you, words not surrendered but rather deliberately chosen in order to better inform and educate all of us. And this is far from the end of the conversation, or the generation of an evidence base. If you have opinions or quality improvement activities or research projects with respect to AI, either within the library or in collaboration with others in health, submit an article to an upcoming issue of JoHILA. Orwell would be most honoured, even if you use (a little) help from our new overlords.