Can a Generated Caregiver Repair Attachment? Implementing Ideal Nurture in Australian Stepped Care

Written from an independent University-researcher perspective

MSc in Translational Health Sciences

Technological Innovation and Digital Health

Trinity Term 2026

Word count: 3,934 (excluding title page, figures, tables, and references)

Abstract

Ideal Nurture is a planned digital-health innovation that delivers Daniel P. Brown's Ideal Parent Figure (IPF) protocol for adult attachment repair through augmented and virtual reality (AR/VR) and an artificial-intelligence-generated caregiver avatar, under clinical governance and a task-shifting workforce model. Written from an independent University-researcher perspective, this report analyses its proposed implementation within COORDINARE, the South Eastern New South Wales Primary Health Network (PHN), for adults with complex and developmental trauma who fall into the "missing middle" of Australian mental-health care. Using the framework for non-adoption, abandonment and challenges to scale-up, spread and sustainability (NASSS), the technology-in-practice lens, the social shaping of technology and Lupton's critical sociology of digital health, it identifies three interrelated challenges: whether the protocol's relational active ingredient survives externalisation into a pre-rendered avatar; engineered dependence concentrated on an underserved population, treated as the equity challenge; and adoption, scale-up and sustainability. The report argues that the therapeutic effect is produced, or lost, in the contingent relational and organisational world the technology enters rather than residing in the avatar, and proposes a complexity-aware implementation strategy built on specifying and testing before claiming.

Background

The innovation. The Ideal Parent Figure (IPF) protocol is the first of the three pillars of Brown's model of adult attachment repair. In its conventional form, a clinician guides a patient, in a lightly absorbed state, to imagine ideal attachment figures who are present, attuned and protective, so that repeated experience of felt security gradually revises the patient's internal working model of relationships (Brown & Elliott, 2016). Ideal Nurture proposes to externalise this imagery into an immersive AR/VR environment delivered by an AI-generated caregiver avatar. The scenes are not generic: they are authored for each patient

from four to five clinician-led intake sessions and from the patient's clinical record, then rendered and played back. The workforce model is one of clinical task-shifting. A scarce, expensive clinician leads the intake, approves all content and conducts the preparation and integration work, while a lower-cost professional supervises the repetitive immersive sessions. The intention is to extend a costly expertise to people that conventional care does not reach, a recognised mechanism in global mental health.

The care problem. In Australia, the target population is what the Productivity Commission (2020) called the "missing middle": the several hundred thousand adults whose needs are too complex for the brief, capped sessions of the Better Access initiative yet not acute enough for state specialist services. This gap is compounded by the uneven distribution of clinicians, which leaves rural and remote communities chronically underserved, and by cost, since most Better Access services now attract a co-payment and use is lowest among those with the greatest need (Pirkis et al., 2022). An intervention that genuinely delivered affordable attachment repair for this population would answer a real and well-documented need.

Advantages and limitations. The case for the IPF is real but early. The protocol is theoretically coherent and clinically detailed (Brown & Elliott, 2016), and one independent, peer-reviewed pilot has reported significant reductions in symptoms and attachment-related distress in adults with complex post-traumatic stress disorder (Parra et al., 2017). The broader principle on which it rests, that strengthening felt attachment security improves mental health, has robust support (Mikulincer & Shaver, 2016). The platform also advances a genuine accessibility rationale: a minority of people cannot voluntarily generate mental imagery, a trait whose near-total form is present in roughly one per cent of the population (Dance et al., 2022), and for them, and for those with no safe figure to imagine, an externally supplied scene may reach where conventional imagery work cannot. Yet the evidence base is small,

uncontrolled and without any randomised controlled trial against an active comparator. The developer's signature case series, in which twelve patients reportedly all reached earned-secure attachment, was scored largely by the senior author over a mean of about 3.4 years of therapy, and its implausibly large reported effect bears the signature of a single unblinded rater rather than a clean result. The two existing IPF studies do not even share an outcome instrument, so they do not straightforwardly corroborate one another. Crucially, the mechanism the evidence supports is slow, relational and therapist-mediated, which is precisely the element the platform proposes to automate. The literature on the platform's nearest analogue, the synthetic companion, is more cautionary still and is examined under the second challenge; for now it is enough to note that the evidence for comfort and the evidence for harm describe the same mechanism.

The setting. This analysis situates implementation within COORDINARE, the South Eastern New South Wales PHN, which commissions primary mental-health services across a region running from metropolitan Wollongong and the Illawarra to rural and remote southern New South Wales. A single PHN therefore spans both metropolitan access and rural maldistribution, allowing the equity questions to be examined in one real setting. Ideal Nurture is modelled as a clinically governed adjunct commissioned into COORDINARE's existing stepped-care services. The direct users are missing-middle adults with complex or developmental trauma, including subgroups with no safe caregiver, those who cannot generate mental imagery, and Aboriginal community members. The indirect users and adopters include the commissioning PHN, the provider organisation and its clinicians and task-shifted staff, referring general practitioners, Aboriginal Community Controlled Health Organisations (ACCHOs), the Therapeutic Goods Administration as regulator, and the non-human components on which the service depends.

The sociotechnical system. Figure 1 places Ideal Nurture within that wider system. Human actors surround the platform alongside non-human components: the avatar and its underlying model, the patient's clinical record used as source material, the data and biometric-measurement layer, the headset hardware, the regulatory classification, and the funding rails. Solid edges mark established flows of oversight, data, funding and regulation; dashed edges mark the points of genuine uncertainty, namely whether the relational mechanism survives pre-rendered delivery, whether the design fosters dependence and displacement, whether the data layer can be reconciled with Indigenous data sovereignty, and whether a service with no reimbursement line can be sustained. The map makes the analytic claim visible: the outcome is decided not by the avatar at the centre but by the uncertain edges around it. A technologically determinist reading, which treats the avatar as an autonomous cause whose effects follow from its features, is of little help here, because it cannot explain why the same protocol might repair one person and foster dependence in another. Read instead through the social shaping of technology (MacKenzie & Wajcman, 1999) and the technology-in-practice lens, Ideal Nurture is not an autonomous curing machine but a sociotechnical assemblage whose effect is constituted in use, in the relationships the map foregrounds. Because the materials that promote it carry a commercial interest, they are treated throughout as claims to be appraised rather than as evidence to be cited, the same standard applied to every source. Figure 2 traces how that architecture and the human delivery fit together across a single course of treatment, from intake to integration, and shows where the clinician's governance and the data-governance obligations sit.

Figure 1

Ideal Nurture in Its Wider Sociotechnical System

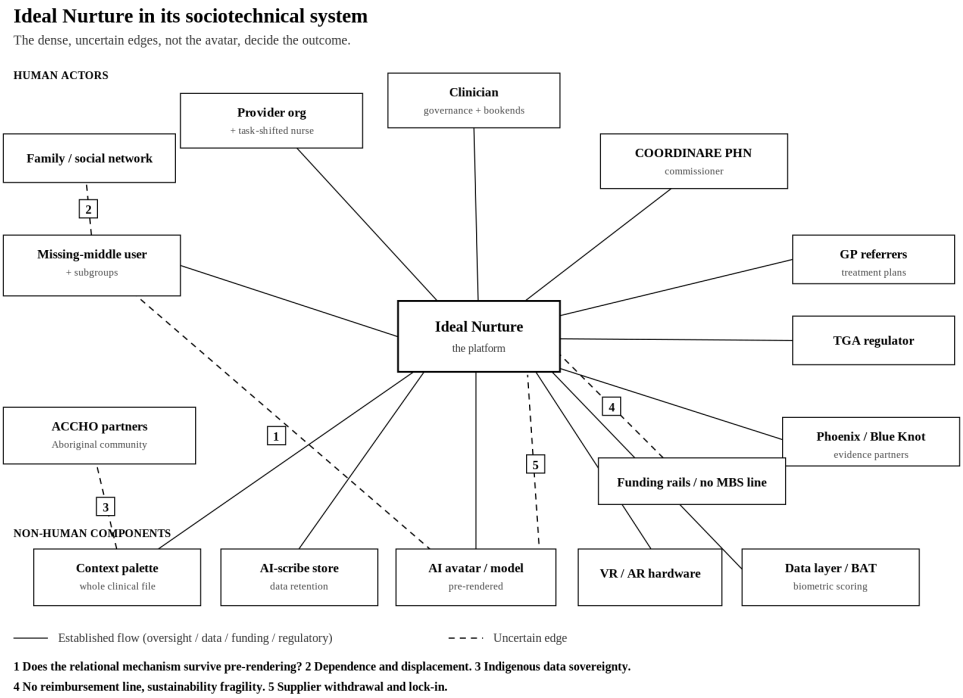

Note. Solid edges mark established flows of oversight, data, funding, and regulation; dashed edges mark the points of genuine uncertainty.

Figure 2

The Clinical-Data and Script Pipeline

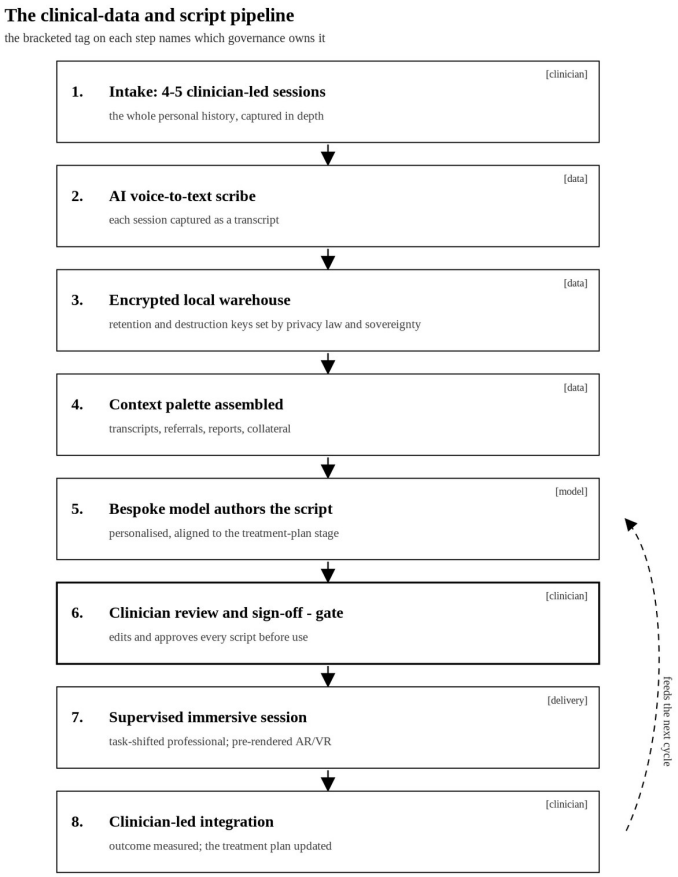

Note. Clinician governance covers intake, sign-off, and integration; data governance covers capture, encrypted storage, and the context palette. Every script is clinician-approved before delivery, and the outcomes feed the next cycle.

Implementation Challenges

Three interrelated sociotechnical challenges are likely to shape adoption and use. Each is presented in its strongest form before being tested against the evidence. The NASSS framework distinguishes seven domains, namely the condition, the technology, the value proposition, the adopter system, the organisation, the wider system, and embedding and adaptation over time, and predicts that complexity across several of them makes non-adoption or abandonment the default rather than the exception (Greenhalgh et al., 2017).

Figure 3 sets out that domain-by-domain assessment for Ideal Nurture; the three challenges below draw out the domains where the complexity is most acute, sharpened by the technology-in-practice view that an effect is produced in use and by Lupton's critical sociology of digital health.

Figure 3

NASSS-CAT Domain Complexity Assessment for Ideal Nurture

NASSS-CAT complexity assessment for Ideal Nurture

DOMAIN	COMPLEXITY	WHY, FOR IDEAL NURTURE	CHALLENGE
Condition	Complex	Complex and developmental (attachment) trauma; high comorbidity; not a discrete, measurable target.	**1**
Technology	Complex	AR/VR, a generative caregiver avatar and a biometric data layer; evolving, opaque and dependence-prone.	**1, 2**
Value proposition	Complex	Supply: no reimbursement line, a fragile DTx market. Demand: benefit unproven, harm plausible.	**3**
Adopter system	Complex	Vulnerable users; clinician role-change and de-skilling; delivery supervised by task-shifted non-experts.	**1, 2**
Organisation	Complicated	PHN commissioning plus provider clinical governance, capacity, and a real safety and medico-legal load.	**3**
Wider system	Complex	TGA software-as-a-medical-device; no MBS line; Indigenous data sovereignty; regulation and funding in flux.	**2, 3**
Embedding over time	Complex	Dependence, supplier withdrawal and responsibilisation drift demand continual adaptation, not a one-off rollout.	**2, 3**

Note. Complexity ratings follow the NASSS framework of Greenhalgh et al. (2017). The final column links each domain to the challenge it most drives.

Challenge one: the automation of contingency (NASSS domains: the condition and the technology). This challenge concerns whether the protocol's active ingredient survives externalisation. It helps to separate what delivery contains: the scene is content, which the platform reproduces faithfully; the induction craft of pacing, gravity and timing is technique, partly portable; but the moment-to-moment contingency, and the clinical judgement of what

to do with what arises, are neither, and it is these that the evidence locates the cure in. The design deserves a full steelman. Because the avatar's scenes are authored from deep clinical intake and the patient's own history, the imagery is highly personalised, arguably more co-created than the closest published precedent, in which recorded, therapist-co-created imagery was replayed for home practice (Parra et al., 2017). Brown's delivery craft, moreover, is partly portable: he himself abandoned formal hypnotic ritual once a simple instruction sufficed, treating pacing, gravity and timing as transferable technique. The objections that the platform offers a generic script, or simply replaces the therapist, do not hold.

The wedge that remains is not content but contingency. The evidence locates the curative element in moment-to-moment responsiveness, a caregiver who matches the patient's affect as it arises, in a reciprocal loop measured in seconds (Ablon et al., 1993), and in a method that, in its developers' words, neither offers nor condones preestablished scripts (Brown & Elliott, 2016). A pre-rendered sequence, however personal, cannot respond to what surfaces in the session. The clinical judgement on which the work depends, knowing when to soothe and when to hold, when to contain a sudden flood of feeling, when to guard against confusing suggested imagery with memory, is displaced onto the upfront authoring and onto a non-specialist in the room who cannot repair a mis-attuned moment as it happens. The analogy to automated therapy does not rescue the design: a conversational agent can deliver a decomposable skills protocol such as cognitive behavioural therapy (Fitzpatrick et al., 2017), but an attachment relationship is not a skill. Immersive technology adds reach and vividness, not a new mechanism; virtual-reality exposure therapy performs about as well as in-person exposure rather than better, and is silent on attachment repair (Carl et al., 2019). The residual dilemma is therefore sharp. Either the corrective experience is meant to occur inside the non-contingent core, in which case the loss of contingency bites and is supervised by a non-expert, or the real repair lives in the human preparation and integration, in which case the

avatar is an expensive scaffold around what remains human-delivered therapy. No study isolates externalised against in-session co-created imagery on a measure of internal working models, so the question is genuinely open rather than settled either way.

Challenge two: manufacturing dependence in the underserved (NASSS domains: the technology, the adopter system and the wider system). This is the equity challenge. The platform offers a vulnerable, underserved population an anthropomorphic, empathic, always-available caregiver figure designed to elicit deep personal disclosure and sustained use. The comfort is real: empathetic conversational agents can reduce loneliness, chiefly by making people feel heard (De Freitas et al., 2025). But the same mechanism that comforts is the one the harm literature watches curdle. In a randomised controlled trial, participants who used a companion chatbot more, regardless of the condition to which they were assigned, showed consistently worse outcomes, with higher loneliness, less real-world socialisation and greater emotional dependence, and total usage duration was the strongest predictor of that dependence (Fang et al., 2025). In a large study of a companion platform, companionship-oriented use was associated with lower wellbeing, and a substantial share of donated conversations reached crisis-grade content (Zhang et al., 2025). Dependence is marked by users coming to feel the bot has needs they must tend (Laestadius et al., 2024), which for a trauma population seeking to set such burdens down is iatrogenic, and attachment to such systems is driven by the very anthropomorphism and empathy the design requires (Huang & Huang, 2025). The platform optimises for frequent, affective, sustained sessions, the exact pattern most associated with harm, so engagement as a measure of success may be harm by design, and a pre-rendered avatar cannot read or de-escalate a distressed user in the moment.

The equity inversion sits at the centre of the case. The access pitch and the harm signal point at the same people: those with least access to a human alternative are offered the synthetic substitute, whose documented risks fall hardest on the lonely and isolated, including

rural residents and those priced out of capped care. Lupton's (2014) critical sociology names the deeper move, in which responsibility for care shifts onto the individual while the relational labour, and the data, accrue to a product. That data dimension is concrete here. The personalisation that strengthens the first challenge is built by ingesting the patient's whole clinical record into a bespoke model, a large surface for privacy and consent and an acute matter of Indigenous data sovereignty, the right of Aboriginal and Torres Strait Islander peoples to govern data about their communities. A further and more fundamental equity problem must be named with care, and with the honest caveat that this analysis has not done the cultural research it would require. Attachment theory's dyadic, Western frame is, in a peer-reviewed view, not valid for Indigenous peoples, whose child-rearing is communal and distributed across many kin, so that a child's security spread across several carers can be misread as disordered (Choate & Tortorelli, 2022; Wright et al., 2024). A singular ideal parent figure may therefore be the wrong construct for Aboriginal users, and any engagement with this population would have to run through ACCHOs on their own terms rather than as an afterthought.

Challenge three: clearance is not viability (NASSS domains: the value proposition, the organisation, the wider system and embedding over time). This challenge is adoption, scale-up and sustainability. Even a tool that worked and was safe would face structural non-adoption. The NASSS framework finds that programmes marked by complexity across several of its domains rarely, if ever, become mainstreamed (Greenhalgh et al., 2017), and the commercial record agrees: the first company to gain United States clearance for a prescription digital therapeutic entered bankruptcy in 2023 despite that clearance, for want of reimbursement (Pear Therapeutics, 2023). In Australia specifically, the Therapeutic Goods Administration now treats software that screens or treats a mental-health condition as a medical device, so an avatar intended to treat trauma would require entry on the Australian

Register of Therapeutic Goods, and there is no Medicare Benefits Schedule reimbursement line for a product of this kind, leaving it dependent on fragile PHN grants. Clearance is a floor, not viability.

The model's economics make the strongest anti-deterministic case in the analysis. Cost is concentrated in training the bespoke model, with near-zero marginal cost on software and headsets, underwriting a promise to train once and reach vast new markets at the click of a button. Read through the analytical lenses taught in the module (Technology and Innovation in Digital Health, 2025), this is technological solutionism, the assumption that spread is frictionless (Morozov, 2013), and a performative expectation, the vast-market figure doing investment work in the present ahead of any evidence (Borup et al., 2006). It collides with the empirical record on spread and with the distinction between work as imagined and work as done (Berg, 2001). The binding constraints were never the software: the per-patient clinical labour does not scale to zero, every market brings its own regulator, reimbursement and trained clinicians, and localisation is not translation, because an attachment intervention is culturally embedded and the ideal parent figure may not transpose to other family systems any more than to Aboriginal kinship.

The trilemma. These challenges do not form a list but interlock, so that the platform's best defence on one front worsens another. The contextual fidelity that strengthens the first challenge is exactly what maximises the bonding and dependence of the second and enlarges its data surface, and the engagement that would prove the business case of the third is the strongest predictor of harm in the second. No single position answers all three at once. Figure 4 renders this interlock.

Figure 4

The Trilemma

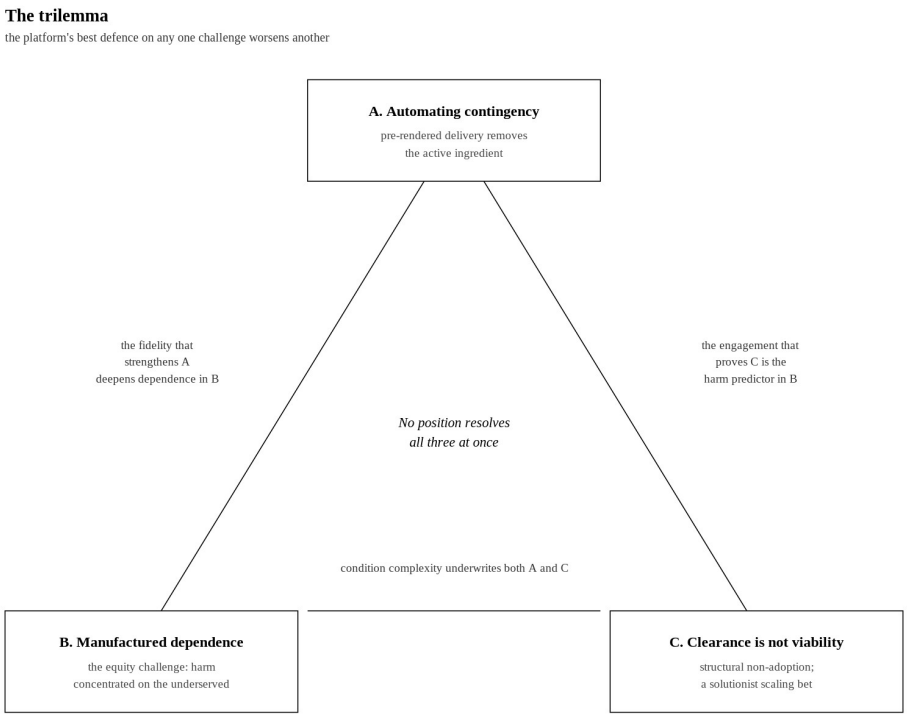

Note. The platform's best defence on any one challenge worsens another, so no single position resolves all three at once.

Implementation Strategy

A strategy adequate to these challenges cannot be a rollout checklist; it must hold the trilemma and proceed by specifying and testing before claiming.

First, success must be treated as a plural and contested construct. Direct and indirect users define it differently: the patient by felt security and a life that resumes, not by hours in a headset; the clinician by a safe caseload and a retained reparative role; the commissioner by equitable access for each dollar; the developer by uptake and duration; the regulator by safety. The central tension is that duration is viability for the developer but the marker of harm for the user, so the service should be designed against engagement-maximisation rather than for it, with session limits and planned tapering built in from the outset. A shared outcome set, agreed with users and clinicians rather than imposed by the developer, would

make these competing definitions negotiable instead of letting the developer's metric win by default.

Second, claims should be gated behind falsifiable, pre-registered tests, mechanism first. Before any claim of corrective attachment is made, the platform should be required to show, against a credible comparator, that externally authored imagery builds an internal working model as in-session co-created imagery is theorised to, rather than merely soothing; to run a trauma-population trial powered for attachment change rather than satisfaction; and to make dependence, displacement of human relationships and crisis-disclosure handling primary safety endpoints rather than afterthoughts, because that is where the evidence says vulnerable users end up. Pre-registration matters here precisely because engagement is so easily mistaken for benefit; the endpoints must be fixed before the data arrive. Figure 5 sets out this falsifiable agenda in full, pairing each claim with the test that would settle it and the endpoint that would count.

Figure 5

A Falsifiable Test Specification

A falsifiable test specification

CLAIM TO BE TESTED	TEST OR COMPARATOR	PRIMARY ENDPOINT	LINKS
The active ingredient survives externalisation	Randomised comparison of externalised vs in-session co-created imagery	Internal working model (AAI coherence), not satisfaction	**Ch. 1**
Benefit exceeds the dependence it creates	Trauma-population trial with an active comparator and follow-up	Attachment-security change and emotional dependence/displacement	**Ch. 2**
It is safe under task-shifting	Human-in-the-loop placement and dose study	Adverse events; crisis handling; the clinician-contingency floor	**Ch. 1, 2**
The mechanism claims survive scrutiny	Pre-registered claims language, independently reviewed	"Felt safety" permitted; settled neurophysiology not asserted	**Ch. 2**
It can be adopted and sustained	A NASSS-guided staged pilot in one PHN site, with stopping rules	Equitable access per dollar; abandonment; a supplier-exit plan	**Ch. 3**

Note. Each claim is paired with the test that would settle it and the endpoint that would count.

Third, the placement of human contingency must be settled empirically rather than assumed. The clinical sources insist on a trained guide. The open question is not only how much clinician involvement is required but where it must sit. A placement study should test the floor of clinician contingency at the preparation and integration bookends, and the minimum professional mix that keeps the work safe and effective without collapsing the affordability case that justifies the model in the first place. It should also specify what competence the lower-cost role requires and what supervision ratio, so that affordability is not bought by quietly thinning the human presence the protocol depends on.

Fourth, the organisation and the wider system, not only the tool, must adapt. Therapist oversight should be non-negotiable clinical governance; the service should be staged into stepped care as an adjunct and never a substitute; crisis triage should be built in; and claims should be held to a science bar that permits felt safety as a clinical construct but does not assert a settled neurophysiological mechanism, since the polyvagal account often invoked for it is contested at the level of its basic premises (Grossman, 2023). Outcome measurement should be blinded and independent of the intervention's authors, and data governance should resolve the tension between the persistent data the personalisation needs and the duty to minimise and erase it, running, for Aboriginal communities, through community-controlled governance. Consent must be specific and revocable, covering not only the clinical record but its use as training material, since a model built from a person's history cannot be straightforwardly unbuilt. Given the absence of a reimbursement line, the route to sustainability should be named honestly, including the real possibility that it is fragile.

Fifth, the service should enter the system as a small, staged, co-designed pilot rather than a product launch. A single COORDINARE site, with a defined cohort and explicit stopping rules, would allow the mechanism and dependence questions to be examined at low stakes before any wider commissioning. Lived-experience involvement should be built in

from the design stage rather than added at evaluation, both because the analysis identifies its absence as a gap and because genuine co-design is itself a check on the responsibilising drift the platform risks. For Aboriginal participants, that involvement must be led by, not merely consulted with, the community-controlled sector. Continuous monitoring should watch not only for benefit but for the specific harms the second challenge predicts, with pre-agreed thresholds at which the pilot pauses.

Finally, the strategy must manage tensions that will only emerge over time. It must plan for what happens as users attach, and must later be helped to transfer that security to human relationships; for how the service would respond to the withdrawal of a proprietary supplier on which dependent users had come to rely; and for how to resist the slow drift, named by Lupton (2014), in which a tool sold as widening access quietly individualises a structural problem and places responsibility on the most underserved. A complexity-aware plan does not eliminate these tensions; it makes them visible and governs them, and it treats the central claim, that a generated caregiver can repair attachment, as a hypothesis to be earned through evidence rather than a promise to be marketed ahead of it.

Conclusion and Reflexivity

This analysis has held a single line throughout: the therapeutic effect of Ideal Nurture is not a property of the avatar but is produced, or lost, in the contingent relational and organisational world the technology enters. Seen this way, the three challenges are one bet viewed from three sides. Whether the active ingredient survives externalisation into a pre-rendered figure is unestablished, not refuted. Whether the design manufactures dependence concentrated on the underserved is well evidenced, and is the equity heart of the case. Whether the service can be adopted, scaled and sustained is adverse on present evidence, with the scaling narrative best read as a determinist expectation. The strategy proposed does not promise success; it specifies what would have to be true, and tested, before a claim of

corrective attachment could honestly be made. For a genuinely complex programme, that disciplined account of uncertainty is the responsible conclusion rather than a verdict.

Several limits of the analysis should be named plainly. The sociotechnical map is a boundary choice that renders some actors and flattens others, particularly the agency of the non-human components. The NASSS framework is powerful as a system description but may under-weight the specific question of relational mechanism that sits at the centre of this case, and reducing a dense field to three challenges is itself a simplification, since other readings might foreground workforce deskilling or medico-legal liability, which appear here only in passing. More fundamentally, the analysis leans on attachment theory, whose universality is contested; its dyadic Western frame may not hold for Aboriginal communal kinship, a limit of both the chosen lens and the intervention's own premise, and one flagged here honestly while acknowledging that the cultural research it would require has not been done. The author also holds a stated positionality: prior operational familiarity with the implementation setting informs this analysis and is carried as situated knowledge held at the analytic distance the brief requires, rather than as hidden bias. Two perspectives would sharpen any future work: the lived experience of the intended users, and the developer's strongest counter-case. The choice of a single Primary Health Network also limits generalisability: another region, with a different rural profile or a stronger community-controlled sector, might shift the equity picture markedly, and the regulatory and commercial environment for digital therapeutics is in flux, so the adoption analysis is a snapshot rather than a settled judgement. Writing in an independent register about a real, named setting carries its own tension, one held here by applying the same scrutiny to the intervention's own claims as to any external source. The honest conclusion is not a verdict but a specification, which is what a sociotechnical reading of a still-emerging technology can responsibly offer.

References

Ablon, S. L., Brown, D. P., Khantzian, E. J., & Mack, J. E. (Eds.). (1993). *Human feelings: Explorations in affect development and meaning*. The Analytic Press.

Berg, M. (2001). Implementing information systems in health care organizations: Myths and challenges. *International Journal of Medical Informatics, 64*(2–3), 143–156. https://doi.org/10.1016/S1386-5056(01)00200-3

Borup, M., Brown, N., Konrad, K., & Van Lente, H. (2006). The sociology of expectations in science and technology. *Technology Analysis & Strategic Management, 18*(3–4), 285–298. https://doi.org/10.1080/09537320600777002

Brown, D. P., & Elliott, D. S. (2016). *Attachment disturbances in adults: Treatment for comprehensive repair*. W. W. Norton.

Carl, E., Stein, A. T., Levihn-Coon, A., Pogue, J. R., Rothbaum, B., Emmelkamp, P., Asmundson, G. J. G., Carlbring, P., & Powers, M. B. (2019). Virtual reality exposure therapy for anxiety and related disorders: A meta-analysis of randomized controlled trials. *Journal of Anxiety Disorders, 61*, 27–36. https://doi.org/10.1016/j.janxdis.2018.08.003

Choate, P., & Tortorelli, C. (2022). Attachment theory: A barrier for Indigenous children involved with child protection. *International Journal of Environmental Research and Public Health, 19*(14), 8754. https://doi.org/10.3390/ijerph19148754

Dance, C. J., Ipser, A., & Simner, J. (2022). The prevalence of aphantasia (imagery weakness) in the general population. *Consciousness and Cognition, 97*, 103243. https://doi.org/10.1016/j.concog.2021.103243

De Freitas, J., Oğuz-Uğuralp, Z., Uğuralp, A. K., & Puntoni, S. (2025). AI companions reduce loneliness. *Journal of Consumer Research*. Advance online publication. https://doi.org/10.1093/jcr/ucaf040

Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W. T., Pataranutaporn, P., Maes, P., Phang, J., Lampe, M., Ahmad, L., & Agarwal, S. (2025). *How AI and human behaviors shape psychosocial effects of extended chatbot use: A longitudinal randomized controlled study* [Preprint]. arXiv. https://arxiv.org/abs/2503.17473

Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. *JMIR Mental Health, 4*(2), e19. https://doi.org/10.2196/mental.7785

Greenhalgh, T., Wherton, J., Papoutsi, C., Lynch, J., Hughes, G., A'Court, C., Hinder, S., Fahy, N., Procter, R., & Shaw, S. (2017). Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *Journal of Medical Internet Research, 19*(11), e367. https://doi.org/10.2196/jmir.8775

Grossman, P. (2023). Fundamental challenges and likely refutations of the five basic premises of the polyvagal theory. *Biological Psychology*, *180*, 108589. https://doi.org/10.1016/j.biopsycho.2023.108589

Huang, Y., & Huang, H. (2025). Exploring the effect of attachment on technology addiction to generative AI chatbots: A structural equation modeling analysis. *International Journal of Human–Computer Interaction, 41*(15), 9440–9449. https://doi.org/10.1080/10447318.2024.2426029

Laestadius, L., Bishop, A., Gonzalez, M., Illenčík, D., & Campos-Castillo, C. (2024). Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media & Society*, *26*(10), 5923–5941. https://doi.org/10.1177/14614448221142007

Lupton, D. (2014). Apps as artefacts: Towards a critical perspective on mobile health and medical apps. *Societies, 4*(4), 606–622. https://doi.org/10.3390/soc4040606

MacKenzie, D., & Wajcman, J. (Eds.). (1999). *The social shaping of technology* (2nd ed.). Open University Press.

Mikulincer, M., & Shaver, P. R. (2016). *Attachment in adulthood: Structure, dynamics, and change* (2nd ed.). Guilford Press.

Morozov, E. (2013). *To save everything, click here: The folly of technological solutionism*. PublicAffairs.

Parra, F., George, C., Kalalou, K., & Januel, D. (2017). Ideal Parent Figure method in the treatment of complex post-traumatic stress disorder related to childhood trauma: A pilot study. *European Journal of Psychotraumatology, 8*(1), 1400879. https://doi.org/10.1080/20008198.2017.1400879

Pear Therapeutics. (2023, April 7). *Current report (Form 8-K)*. U.S. Securities and Exchange Commission.

Pirkis, J., Currier, D., Harris, M., Mihalopoulos, C., Arya, V., Banfield, M., Bassilios, B., Buchanan, B., Butterworth, P., Brophy, L., Burgess, P., Chatterton, M. L., Chilver, M., Eagar, K., Faller, J., Fossey, E., Ftanou, M., Gunn, J., Kruger, A., … Williamson, M. (2022). *Evaluation of the Better Access initiative: Final report*. Australian Government Department of Health.

Productivity Commission. (2020). *Mental health: Productivity Commission inquiry report* (Report No. 95). Productivity Commission.

Technology and Innovation in Digital Health. (2025). *Module sessions and discussion board* [MSc in Translational Health Sciences]. University of Oxford. [Confirm coordinator name, exact session or reading title, and the virtual learning environment reference.]