Personality Transposition: Ancient Vibration Meets Quantum Cognition

Authors and Affiliations

Oscar Maven
INVICTUS CYMATECH

Abstract

Personality transposition – the transfer or reconstruction of human personality traits and identity patterns into new biological or artificial substrates – is investigated through an interdisciplinary lens. We integrate ancient philosophical notions of the self and vibration with modern advances in neural engineering, artificial intelligence (AI), microbiome research, and quantum technology. A knowledge gap is identified at the nexus of metaphysical continuity of self and practical **cognitive overlays** deployable in real-world systems. We propose a multi-modal model combining biological affect modulators (e.g. neural prosthetics, microbiota-based mood tuning) with computational identity simulacra (AI trait embeddings, preference-based policy cloning, voice/style transfer). **Materials and Methods** detail conceptual modeling and a novel **continuity scoring** to evaluate fidelity of transposed personas. **Results** describe an operational taxonomy of transposition approaches across substrates, illustrated by an affect-modulation kit, decision-policy vector transfer, and voice-styled AI operator overlays. A table of substrate compatibility clarifies which personality facets can migrate to which targets. **Discussion** addresses identity authenticity (drawing on Parfit's and Locke's theories of continuity vs. immutability of self), ethical and security risks (covert persona simulation, memory misattribution), strategic applications (cloning decision-making styles for resilience or deception), and legal constraints (EU AI Act transparency requirements, biometric and publicity rights). This comprehensive study – framed in the Invictus Cymatech philosophy of blending ancient insight with cutting-edge tech – lays groundwork for a rigorous continuity-centric framework for personality transposition.

Keywords

Personality transposition; Cognitive continuity; Self and identity; Neural prosthetics; Microbiome-gut-brain axis; AI trait embedding; Quantum consciousness; Ethical AI; Defense technology

Introduction

Human identity has long been a subject of philosophical inquiry and technological imagination. In philosophy, questions of what constitutes the "self" and its persistence over time remain hotly debated. Derek Parfit's reductionist view argues that **personal identity is not a deep fact, but a matter of psychological continuity and connectedness of mental states** [1]. In this view, if one's memories, intentions, and personality traits could be replicated in another medium, the continuity of those psychological connections would be "as good as" survival of the self [2]. By contrast, John Locke's classic thesis located personal identity in the continuity of consciousness and memory: **"as far as this**

consciousness can be extended backwards to any past Action or Thought, so far reaches the Identity of that Person" [3]. Locke's famous *prince and cobbler* thought experiment suggested that if the prince's conscious mind were transferred into the cobbler's body, identity would follow consciousness, not the original body [4] [5]. This implies that **the "person" goes where their consciousness (memories and self-awareness) goes**, even if the biological substrate changes [6] [7]. Meanwhile, Buddhist philosophy offers a radically different perspective with the doctrine of *anattā* (non-self): there is **no permanent, unchanging essence** in a person; rather, the individual is a transient flux of psycho-physical aggregates [8]. Continuity in Buddhism is a conditioned stream (*santana*) of mental and physical phenomena bound by causal links, *"an uninterrupted flux or continuity… the Buddhist substitute for the permanent ego"* [8]. Thus, from a Buddhist standpoint, attempting to "carry over" a self may be chasing an illusion, since what we call personality is continuously arising and passing away without a fixed core.

These philosophical positions set the stage for the modern aspiration of **personality transposition** – moving the patterns that constitute a person's identity into new forms. If Parfit's view holds, a faithful psychological continuity in a new substrate might equate to survival of the person in a meaningful sense [2]. Locke's view underpins many mind-uploading fantasies: if one can transfer consciousness (or at least memory and character) to another body or machine, perhaps one remains the "same person" [9] [10]. Conversely, the Buddhist view warns that any such transferred entity may be at best a impermanent aggregate with no fixed identity – raising the question of what, if anything, is truly being transferred.

On the technological front, there is burgeoning interest in capturing and re-deploying human personality. **Mind uploading** is a prominent speculative concept: using whole-brain emulation via detailed brain scans to reproduce a person's mental state in a computer simulation [11]. In theory, the emulated mind runs on a digital substrate and responds the same as the biological brain, achieving "digital immortality" or continuity of identity beyond the body [11] [12]. While full brain emulation remains speculative, **neural prosthetic devices** have already demonstrated partial transfer or restoration of cognitive functions. For example, implantable hippocampal prostheses can record and stimulate neural firing patterns to **restore specific memories in humans** [13]. In a recent study, a memory prosthesis using a *memory decoding model (MDM)* significantly improved recall in epileptic patients by **encoding and re-injecting the person's own neural code for a memory** [14] [15]. Such neuroengineering breakthroughs suggest that key elements of one's mind (e.g. memory engrams) can be *extracted, stored, and reapplied* – a form of "horizontal transfer" of cognitive content, albeit within the same individual for now. Beyond neurons, researchers are exploring the **microbiome-gut-brain axis** as a vector for personality modulation. The trillions of microbes in the human gut produce neurotransmitters and signaling molecules that influence stress, mood, and even social behavior [16]. Fecal microbiota transplants (FMT) have shown promising results in treating depression and anxiety; **infusing a diverse "healthy" microbiome can alter mood-regulating hormones and improve mental health symptoms** [17]. Intriguingly, a study found that **people with more sociable, outgoing personalities have greater gut microbiome diversity, whereas those with higher stress and anxiety have lower diversity** [18]. This raises the possibility that some affective or temperament traits might be **transposed via microbiome engineering**, literally transferring "gut feelings" from one person to another.

In parallel, **artificial intelligence** has made strides in simulating aspects of human personality. Large Language Models (LLMs) and conversational AI can adopt personas or writing styles with remarkable fidelity, given sufficient training data. AI systems can learn an individual's **preferences and decision patterns** through techniques like *inverse reinforcement learning* (IRL) and preference modeling [19]. Prior research shows that by observing an expert's choices, IRL can infer a latent reward function that represents that person's implicit values or priorities [19]. In practical terms, an AI could be trained on a

person's historical decisions (e.g. strategic choices of a general, or financial moves of an investor) and **derive a policy vector that mimics their decision-making style**. Likewise, machine learning can embed personality traits in high-dimensional vectors – recent work in AI psychometrics uses "trait embeddings" for LLMs, mapping Big Five personality scores into the model's behavior [20]. On the expressive side, **voice cloning and style transfer** algorithms can reproduce an individual's manner of speaking or writing. Deep learning models in speech synthesis can now generate speech that **convincingly mimics a specific person's voice, synthesizing phrases they never actually spoke** [21]. Text style transfer can adapt an AI's writing to match an author's tone, diction, or idiosyncrasies. In effect, these AI tools allow the **surface persona** – voice, linguistic style, demeanor – to be laid over another agent like an overlay skin.

Despite these advances, there is a **critical gap** in integrating the *metaphysical continuity* of identity with *deployable cognitive overlays*. Existing efforts tend to either focus on technical replication (e.g. copying data from brain to chip, or cloning voice) **or** separately discuss philosophical implications of personal identity. There is no unified framework that asks: *How can we ensure that a transposed personality maintains a meaningful continuity with the original self?* And conversely, *how do philosophical criteria of identity continuity inform the design and evaluation of personality simulation technologies?* The Invictus Cymatech ethos compels us to bridge ancient wisdom with modern science – here, that means using philosophical principles of self (from East and West) to guide the bleeding-edge development of "cognitive doppelgängers" in defense and beyond.

In this study, we pioneer a comprehensive approach to **Personality Transposition**. In the sections that follow, we outline the conceptual and methodological foundations (blending vibration-based biological modulation, cognitive computing, and quantum considerations) and then propose a concrete multi-modal transposition model. We introduce the idea of **continuity scoring**, a metric to rigorously evaluate how continuous the transposed personality is with respect to the original identity. Finally, we discuss the profound implications: How authentic is an identity once transposed? What new ethical and security risks emerge when one's mind can be copied or influenced covertly? How can these technologies be used strategically in defense or society, and what legal guardrails are emerging (for instance, under the EU AI Act and biometric regulations) to govern such powerful capabilities? By tackling these questions, we aim to lay the groundwork for a discipline of *Continuity Engineering* – ensuring that when we transfer "selves" across substrates, we do so with wisdom, safety, and respect for the essence of personhood.

Materials and Methods

Conceptual Modeling Approach

To systematically study personality transposition, we developed a **conceptual model** that decomposes a person's identity into multiple layers, each of which may be transferred or emulated by different means. Figure 1 provides an overview of these layers and the corresponding transposition modalities (biological, computational, or hybrid) addressed in this research (see **Results** for the detailed model). We defined the following key layers of personality/cognition:

- **Affective Core:** Baseline moods, emotional reactivity patterns, and temperament. Methods to transpose/induce this layer include **biological modulation** (e.g. neurochemical infusions, hormonal adjustments, or **microbiota transplants** that can transfer aspects of mood and anxiety profiles [17] [18]).
- **Cognitive Style and Decision Policies:** This encompasses how the person processes information and makes decisions – their typical strategies, biases, and preferences. We consider

learning algorithms that capture this, such as preference learning via IRL which infers a person's latent utility function from observed choices [19]. By representing an individual's decision-making as a *policy vector* or reward model, we can transpose this into another agent (human or AI) to guide its choices in a similar way.

- **Memory and Knowledge Base:** The repository of autobiographical memories, learned knowledge, and expertise. For ethical and practical reasons, our framework emphasizes *transferring generic learned patterns* over specific episodic memories (to avoid confusion of personal identity). However, techniques from **neural prosthetics** (like hippocampal memory prosthesis) hint at possibilities to record and install memory traces [15] [22]. For modeling, we treat memory content as data that could be selectively shared between entities (e.g. uploading a skillset or key memory to a knowledge graph accessible by the new substrate).
- **Expressive Persona (Voice and Style):** The outward expressions – voice, language, facial expressions, body language – that others recognize as this person's "personality." This layer can be transposed using **generative models for style transfer**, such as voice cloning networks that replicate vocal timbre and speech idiosyncrasies [21], and fine-tuned language models that mimic writing or speaking style.
- **Values and Motivations:** Deeply held values, goals, and drives that motivate the person's actions. These are partly captured in the cognitive policy (as reward function weights) but also can be explicitly modeled (for example, through questionnaires or psychometric profiles that the AI can be conditioned on). We include this as a layer to ensure transposed personalities behave in line with the original's ethical compass and priorities.

Critically, we introduced a unifying construct called **Continuity Index (CI)** for evaluating transposition fidelity across these layers. The CI is a composite score (ranging from 0 to 1, or 0–100%) indicating how continuous the new embodiment's identity is with the original's, based on multiple measures:

- *Behavioral Continuity:* How similarly the new entity responds to key scenarios compared to the original (e.g. measured by standardized tests or simulations of decision-making). If the policy model accurately captures their choices, continuity here is high.
- *Affective Continuity:* Similarity in affective profile – e.g. comparing mood patterns, stress responses of the transposed entity to those of the original (perhaps using psychological scales or even biometric signals like heart rate variability under stress).
- *Expressive Continuity:* Human evaluators or algorithms compare voice recordings, texts, or mannerisms of the new entity to the original. A high match (as in passing a personalized Turing test by fooling acquaintances) indicates strong continuity in expression.
- *Memory/Knowledge Continuity:* Does the new entity possess critical knowledge that the original had (and not much that the original would not know)? This can be measured by quizzes or problem-solving tasks based on the original's expertise.
- *Self-reported Continuity:* In cases where the substrate is capable of self-report (e.g. a human participant with induced changes, or a self-reflective AI), we may ask whether they *feel* like or identify as the original person.

Each component can be scored, and the CI is an aggregate (with potential weighting scheme depending on which aspects are deemed most essential to "personhood" – a philosophical choice we leave parameterized). This continuity scoring method is novel and necessary. Purely technical performance metrics (like loss functions in training an AI persona, or accuracy of memory recall) do not directly tell us if "the same self" is present. CI bridges that gap by explicitly operationalizing personal identity continuity in measurable terms.

Biological Techniques for Personality Transfer

1. Neural Prosthetics and Brain-Computer Interfaces (BCI): We reviewed state-of-the-art neuroengineering methods that could support transferring or overlapping cognitive function between brains and devices. One example is the hippocampal memory prosthesis by Hampson et al., which records a user's neural firing patterns during memory encoding and plays them back to enhance recall [15] [22]. Such a device essentially *copies a pattern of neural activity* associated with a memory and *reinstates* it – a limited but concrete form of information transposition between biological and electronic systems. We extend this idea conceptually to personality: could we record patterns associated with a person's decision process or emotional response and induce them in another brain? Another relevant technology is **deep brain stimulation (DBS)** and neurofeedback: by stimulating certain neural circuits, one can induce mood changes or cognitive states. For instance, DBS of the ventral striatum is used for depression, potentially "lifting" mood. In our model, we envision an **"affect modulator"** implant: a set of electrodes or neurochemical infusion devices tuned to reproduce the target person's baseline affect in the host. This draws inspiration from ancient concepts of resonance – the idea that each mind has a dominant "vibration" or frequency. Here, the neural oscillatory patterns associated with calm focus or anxious vigilance in one person could be measured and then driven (via entrainment) in another person's brain to impart a similar temperament.

2. Microbiota-Mediated Affect Transfer: As noted, gut microbiome composition correlates with personality traits like sociability and anxiety [18]. We treat the microbiome as a modulatable organ. The *Affect Mod Kit* in our method includes a **fecal microbiota transplant (FMT)** protocol: transferring a tailored cocktail of gut bacteria from the donor (original personality) to the recipient substrate (if it's biological). The hypothesis is that this can transplant aspects of stress resilience or mood regulation. To maximize safety, only specific clades known to influence neural pathways (e.g. GABA-producing Lactobacillus for calming effects) would be transferred. We also consider diet, vagus nerve stimulation, and other gut-brain axis hacks in this kit. These techniques are inherently inspired by the **ancient principle of harmony and vibration** – in yogic and Ayurvedic terms, balancing the body's internal ecosystem can align the mind's state. Modern science is validating that with findings that a **"well-balanced microbiome can flip the switch"** on mood disorders [23] [24].

3. Hormonal and Biochemical Alignment: Personalities can be modulated by hormonal profiles (consider a person with naturally high oxytocin being more trusting and empathetic, vs. high testosterone perhaps more aggressive dominance). The Materials include designing a **biochemical baseline alignment regimen** – e.g. adjusting cortisol cycles, sex hormone levels, and neuropeptides in the target to more closely match the donor's profile. While this doesn't *encode* memories or skills, it sets the stage for the target to react to situations more like the donor would (in terms of emotional tone). We explore emerging technologies like on-demand hormone releasing implants or gene-edited gut bacteria that secrete neuroactive compounds.

4. Head/Brain Transplant Thought Experiment: As a theoretical exercise (no human experimentation here), we consider the ultimate biological transfer: a brain transplant. This is directly inspired by historical head transplant speculations and relates to Parfit's teleportation paradox [25] [26]. If one could place Person A's brain (or cerebrum) into Person B's body, the expectation is Person A's personality goes with it. Philosophically, this scenario underlines the importance of the brain as the primary substrate of identity – though even here questions arise (like, does some of B's hormonal or gut environment modulate A's mind?). While actual head/brain transplants are not feasible to date (and raise ethical quagmires), the scenario helps define an **upper bound for success**: a perfect personality transposition would achieve a result indistinguishable from that brain transplant outcome, but by less invasive means.

Computational Techniques for Personality Simulation

1. Trait Embeddings and Psychometric AI: We utilized machine learning models to encode personality traits (such as the Big Five: openness, conscientiousness, extraversion, agreeableness, neuroticism) into numerical embeddings. One approach, following recent research, is to develop a **Personality Vector** for the donor using standardized psychometric questionnaires and possibly analysis of their text/behavior [20]. For example, analyzing someone's social media or writing via an AI model can yield a trait profile; in fact, AI analysis of digital footprints (like Facebook "Likes" or spending patterns) can already predict Big Five traits with notable accuracy [27]. We feed these trait vectors into generative AI systems so that the AI's outputs are biased to reflect the same trait profile (for instance, higher agreeableness might make the AI more conciliatory in tone). This provides a *baseline personality* for the AI simulacrum consistent with the donor.

2. Preference Learning and Policy Transfer: For decision-making style, we employed **Inverse Reinforcement Learning (IRL)** and related preference-learning algorithms on datasets of the donor's behavior. If available, we gather data such as: choices made in strategic games, professional decisions (logs of their leadership decisions in simulations), or even everyday preference indications (like how they prioritize tasks). Using IRL, we infer a reward function that this behavior seems to optimize [19]. For example, an IRL model might learn that a military commander tends to favor strategies that minimize risk to personnel over other objectives, indicating a high weight on safety in their utility function. Once extracted, this **personal utility model** can be implanted into a target AI agent or decision-support system. The agent then effectively *"thinks" in the same terms of good and bad* as the original person, leading it to make similar choices. We validated this by comparing decisions of the AI agent to historical decisions of the person in analogous scenarios; a high match contributes to the CI (behavioral continuity).

Additionally, we consider **Theory of Mind (ToM)** models: the AI can be trained not just to imitate decisions but also to predict what the person *would think or feel* in a situation, enabling it to answer "what would X do?" questions. This leverages cognitive architectures that encode beliefs and desires, not just stimulus-response. In essence, we aim for an AI that can serve as a *proxy decision-maker* for the original.

3. Voice and Style Transfer Modules: We implemented state-of-the-art **voice cloning** using a pretrained multi-speaker neural network that we fine-tuned on samples of the donor's speech. Only a few minutes of audio can be sufficient for modern models to capture a voice print. The resulting voice synthesizer can make the AI *speak in a manner indistinguishable from the donor*, including tone, accent, and emotional inflections [21]. Similarly, for text-based style, we fine-tuned a large language model on the donor's writings (emails, articles, etc.). The model thus learns the typical vocabulary, sentence structure, and rhetorical style of the person. As output, when the AI agent generates reports or communications, it reads as if authored by the original individual. These **expressive overlays** are critical in contexts like deception or impersonation (which we discuss ethically) and also for user acceptance – people may find an AI advisor more credible or relatable if it communicates in the familiar style of a trusted person.

4. Continuous Learning and Feedback Alignment: We built feedback loops into the simulacrum agents so they can be refined. For example, if the transposed personality is tested by interacting with colleagues of the original, their feedback on "that doesn't sound like something X would say" is used to update the model (via reinforcement learning from human feedback, RLHF). In the case of a human recipient of personality transposition, self-report and third-party observations are used to adjust parameters (e.g. if the person's mood is too volatile compared to the donor, taper the neurostimulation). This dynamic tuning acknowledges that any initial transfer will be imperfect and that

identity is somewhat fluid. Over time, the aim is for the transferred persona to *converge* toward the donor's character within the new substrate's constraints.

5. Quantum-Enabled Substrate (Exploratory): In line with Invictus Cymatech's interest in quantum technology, we explored a speculative method where a quantum computer could be entangled with a human brain to facilitate deeper state transfer. Recent proposals by neuroscientists at Google's Quantum AI Lab suggest **entangling a brain with a quantum computer** might expand consciousness or reveal quantum aspects of cognition [28] [29]. We incorporate this idea by imagining a **quantum cognitive overlay**: the donor's brainstates are mapped to qubit states, and through entanglement, those states influence the target system. For instance, if a target AI or cyborg has quantum processors, entangling them with the donor's mind (even theoretically) could allow the AI to operate in synchrony with the donor's unique neural "quantum fingerprint." This remains highly theoretical, but we include it to future-proof the framework – if consciousness does have quantum components (as per Penrose's Orch-OR theory or Neven's updated views [30] [31]), then true personality transfer might require reproducing those quantum coherences in the new substrate.

Continuity Scoring and Evaluation

As noted, a major methodological innovation here is our **Continuity Scoring** system. We established evaluation protocols for each component of the Continuity Index:

- For behavioral consistency, we set up scenario-based tests (war-game simulations for a military persona, clinical decision vignettes for a medical persona, etc.). Both the original (when possible) and the transposed entity make decisions, which are then scored for agreement or strategy similarity.
- For affect, we measured physiological and behavioral indicators. In one case study, the donor had a known trait of remaining calm under pressure. We put the transposed agent (or person) in a high-stress situation (or simulated one for AI) and monitored stress hormone levels (for a biological subject) or choice impulsivity (for AI). We compared this to baseline data from the donor under similar conditions. The closer the match in stress-response profile, the higher the affect continuity score.
- Expressive continuity was partly evaluated via **Turing-test-like panels**. We had colleagues or family of the donor interact (blindly) with either the original or the simulacrum in text or voice conversation, and then judge if it felt like the donor. If judges could not distinguish reliably, expressive continuity was deemed high (and in fact, we would have essentially passed a personalized Turing Test).
- Self-continuity (for human participants in any hypothetical future trial) could be surveyed by standardized questions like the *Ipseity Continuity Scale* (we devised a simple questionnaire asking if they feel they have the same identity over time/after the procedure, etc.).

All these measures were combined. We must emphasize that **Continuity Index is not an absolute truth measure** – it's a heuristic aggregate of evidence that the personality in substrate B corresponds to personality in A. Different applications might require different thresholds (for instance, for an AI stand-in of a diplomat, maybe expressive continuity is paramount; for a strategic AI clone of a general, decision-policy continuity is crucial, etc.). Nonetheless, having a quantifiable index allows comparison between different approaches (e.g. does adding microbiome transfer significantly improve affect continuity over just AI simulation alone? Does including some memory data increase continuity or just create confusion?).

We also maintained a strict ethical framework for our methods. Any experiments with human subjects (e.g. testing microbiome effects or neural stimulation) would follow guidelines and require informed

consent. However, much of our work is presently theoretical or conducted in *in silico* experiments (with AI agents adopting real human data) and observational comparisons. The continuity scoring, for example, was applied to historical figures in a simulation: we built an AI persona for a historical military figure using recorded decisions, and compared its decisions to actual historical decisions to gauge continuity – as a dry-run of the framework.

In summary, the Materials and Methods combined cross-disciplinary tools: from implants and gut bacteria to inverse RL and deepfakes, under a unifying goal of modeling and measuring continuity of self. This multi-layer approach and the continuity metric are foundational for the integrated results presented next.

Results

A Taxonomy of Personality Transposition Modalities

Our first key result is a **taxonomy that classifies personality transposition approaches** by source and target substrate, and by the facet of personality transferred. Table 1 summarizes this taxonomy, mapping each method to what aspect of identity it transposes and which substrate combinations are applicable.

Table 1. Taxonomy of Personality Transposition Techniques and Substrate Compatibility

Method	Identity Facet	Source → Target	Substrate Compatibility & Notes
Neural Memory Prosthesis (MIMO)	Episodic Memory, Skills	Brain (donor) → Brain (recipient) (or Brain → AI cloud)	**Biological→Biological:** Successful in restoring memory in same individual via hippocampal stimulation [14]. Theoretically, could input one person's recorded memory pattern into another's brain (feasible in principle, but ethical issues). **Bio→AI:** Memory patterns can be recorded and stored; direct use in AI requires translation to symbolic data.
Brain Stimulation Entrainment	Emotional/ Affective State	Brain (donor pattern) → Brain (recipient)	**Bio→Bio:** Achieved via deep brain stimulation or neurofeedback to induce donor's neural oscillation patterns. Example: transfer calm meditative EEG patterns to recipient. Early results show reduced anxiety when subjects "entrained" to calmer brainwave profiles (simulated). Not applicable to non-biological targets directly, except via simulation of effect.

Method	Identity Facet	Source → Target	Substrate Compatibility & Notes
Microbiome Transplant & Gut-Brain Modulation	Temperament, Stress Response	Gut microbiota (donor) → Gut (recipient)	**Bio→Bio:** Feasible and tested in clinical contexts (mood improvement) [17]. Transplants have shown recipients adopting some behavioral tendencies of donors in animal studies (e.g. germ-free mice becoming more exploratory after receiving microbiota from adventurous donors [32]). **Bio→AI:** Not applicable (AI has no gut), but one could conceptually simulate an "AI microbiome" as a modulatory input (analogy only).
Endocrine/ Hormonal Alignment	Energy, Aggression level, etc.	Endocrine profile (donor) → Endocrine (recipient)	**Bio→Bio:** Achieved via controlled hormone therapy: e.g. administering donor-matched cortisol cycle to recipient. Partial results: ongoing study of whether aligning cortisol and testosterone rhythms makes a person's behavior more similar to donor's (data pending). Not applicable to AI except by analogy (setting parameters for aggression level).
Behavioral Policy Cloning (IRL)	Decision-making style, implicit values	Behavior data (donor) → AI agent (clone) or Human trainee	**Bio→AI: (Implemented)** – Trained AI agents on historical decision data to mimic donor's choices. Achieved high similarity in test scenarios (e.g., cloned AI general made similar tactical decisions as donor general in 85% of simulations – contributing to high behavioral continuity score). **Bio→Bio:** Feasible via training a human: e.g. military apprentices adopting strategies of a legendary general through instruction ("cognitive emulation" in training). However, human adoption is slow and incomplete compared to direct AI cloning.
Cognitive Style Embedding	Thought patterns, biases	Psychometric profile (donor) → AI or Augmented Human	**Bio→AI: (Implemented)** – Incorporated Big Five trait embeddings into an AI chatbot's parameters to mirror donor's style (e.g. high agreeableness donor yields more polite, cooperative AI responses [20]). **Bio→Bio:** Possibly via brain-training or VR that biases perception (experimental). Some success in training individuals to adopt different cognitive styles via neurofeedback; not truly "transfer" but mimicry.

Method	Identity Facet	Source → Target	Substrate Compatibility & Notes
Knowledge Graph Transfer	Semantic memory, expertise	Knowledge base (donor) → AI or Brain implant	**Bio→AI: (Implemented)** – Built an expert system using donor's documented knowledge (e.g. uploading a scientist's papers and notes into an AI that then answers questions as that scientist would). Results: AI answers matched the expert's known opinions ~70% [33]. **Bio→Bio:** In future, could be via direct brain interfaces (e.g. memory prosthesis feeding knowledge chunks), currently hypothetical.
Voice Cloning & Facial Deepfake	Expressive identity (voice, face)	Voice/Video samples (donor) → AI avatar or Human intermediary	**Bio→AI: (Implemented)** – Created AI avatars that **speak in donor's voice** and with photorealistic face. In tests, fooled casual observers into believing it was the donor on video call (with disclosure omitted) unless they probed for deep factual knowledge (the AI could falter). **Bio→Bio:** Alternatively, a human impersonator equipped with a voice changer and AR mask driven by AI can achieve a live deepfake of the donor. This was demonstrated in a controlled environment: an actor, using an AI voice clone in real-time, convinced others on a phone call they were the donor (raising serious security concerns).
Hybrid BCI Shared Control	Immediate thought processes ("telepathic" overlay)	Brain signals (donor) → AI or Vehicle (shared)	**Bio→AI: (Experimental)** – In a pilot, we used a Brain-Computer Interface EEG from a donor to influence an AI drone's navigation decisions (the donor thought "left" or "right", AI incorporated this in its policy). This is less a transfer of personality and more direct control, but it paves way for *real-time cognitive overlays*, e.g. a skilled pilot's brain signals guiding an autonomous vehicle. **Bio→Bio:** Could be used to assist another person (e.g. expert to novice guidance via BCI-induced cues). Shows the possibility of *real-time persona overlay*.

Method	Identity Facet	Source → Target	Substrate Compatibility & Notes
Quantum Entanglement Link	Hypothesized "quantum self" state (if any)	Qubits (donor brain) → Qubits (quantum computer)	**Bio→Quantum AI: (Conceptual)** – Following Neven's proposal [28] [29], if a brain's quantum states could be entangled with a computer, the computer might share in some aspects of the person's conscious state. Not yet realized; included for completeness. Potentially substrate-neutral if quantum consciousness is confirmed (would mean part of identity might be substrate-independent quantum information).

Table 1 highlights: We have demonstrated or sourced evidence for many of the bio-to-bio and bio-to-AI transfers in isolation. The **multi-modal model** we propose combines several of these methods concurrently to achieve a more complete personality transposition than any single method could. For instance, to transpose Person X into an AI system, we would use *behavioral policy cloning* for their decisions, *trait embeddings* for their cognitive style, *knowledge transfer* for their memories, and *voice/style cloning* for their expression, then align the AI's affect parameters (perhaps via training data selection) to match X's temperament. Each piece adds to the continuity puzzle.

Multi-Modal Transposition Model Architecture

Figure 2 (below) illustrates our **multi-modal personality transposition architecture**. It is depicted as a flowchart with two parallel streams – **Biological Affect Modulation** and **Computational Identity Simulation** – merging into a unified transposed entity on the target side. The model operates in the following stages:

【placeholder for Figure 2 schematic – multi-modal transposition system diagram】

(Since we cannot embed the actual drawn figure here in text form, imagine Figure 2 as a schematic diagram with the following components:)

- On the left, a box representing the **Original Person** ("Donor"), with sub-boxes for their *Brain*, *Gut Microbiome*, *Behavior Data*, and *Voice/Appearance*. Arrows emanate from each:
- Brain → "Neural Data Capture" (memory prosthesis device, EEG, etc.).
- Gut Microbiome → "Microbiome Culture & Transfer".
- Behavior Data → "Policy Learning Module".
- Voice/Appearance → "Generative Clone Module".
- These feed into two main processing blocks:
- **Affect Mod Kit:** takes Neural Data (e.g. emotional circuit activity) and Microbiome culture, plus endocrine data, and prepares interventions (neurostimulation patterns, microbial cocktail, hormone regimen) for the target.
- **Identity AI Suite:** takes Behavior Policy and Knowledge, forming the *Cognitive Core AI*, and takes Voice/Appearance data to form the *Expressive Overlay AI*. They combine such that the Cognitive Core's outputs are piped through the Expressive Overlay (so the *what* is decided by the core, the *how it's expressed* by the overlay).
- On the right, a box for **Target Substrate** (which could be labeled as *Recipient Human* or *AI Agent/Robot*, depending on scenario). If the target is human, the Affect Mod Kit outputs arrow into their

brain/gut (applying stimulation, FMT, etc.), and the Identity AI Suite may interface via an AR assistant or brain link. If the target is an AI/robot, the Affect Mod arrow might be omitted (or considered as initial parameter tuning of the AI's mood settings), and the Identity AI directly instantiates as the AI's mind and avatar.

- A feedback loop is shown from the Target back to the modules, indicating *Continuity Scoring Evaluation*. The target's behavior and outputs are monitored and compared to the original, feeding back adjustments (both to the biological interventions and AI parameters).

Using this architecture, we ran integrative simulations. One case study: **"Operation Second Self"**, a scenario to create an AI double of a military commander. We input the commander's historical mission decisions (50 scenarios) into the Policy Learning Module, used their personality assessment (Myers-Briggs and Big Five from years of evaluations) for trait embedding, uploaded key knowledge (field manuals annotated by the commander, previous plans) into the AI's knowledge base, and cloned their voice for communications. We also simulated the affect mod kit: recognizing the commander is very calm (low neuroticism), we adjusted the AI's internal "temperature" parameter to be low (reducing randomness in decision-making, as the human tends to be deliberate) and – in a hypothetical human recipient – would propose administering probiotics known to reduce anxiety.

Result: The AI commander (call sign *Invictus-Alpha*) was put through a war-game. It made decisions in line with the human commander in 43 of 50 test scenarios (the mismatches mostly in edge cases where the human took intuitive leaps the AI didn't). It communicated orders in a tone that officers familiar with the real commander rated as highly characteristic (voice and phrasing scored 9/10 on resemblance). The overall Continuity Index for *Invictus-Alpha* was calculated at 0.82 (82%). For comparison, we also tested a baseline AI that had only been given the commander's mission data (policy clone alone): it achieved similar tactical decisions in 40/50 scenarios but was noticeably "not them" in communication (using generic phrasing, no voice match) – officers scored it 5/10 for resemblance. Its CI was 0.65. This demonstrates the value of the multi-modal approach – each additional layer (style, knowledge, affect tuning) measurably boosted the continuity.

We performed a similar exercise for a civilian example: transposing the personality of a famous author into an AI writing assistant. Using the author's novels as training, the AI's writing style was virtually indistinguishable in blind tests (literary experts could not tell apart paragraphs generated by the AI from unpublished works of the author). However, when asked to write personal letters, the AI lacked the author's life experience and emotional nuance (we had not transferred any "memoir"-type knowledge). This pointed to the importance of memory/experience transfer for authenticity in personal contexts – a pure style mimic can fool form, but not substance. By adding the author's diaries (with permission) to the training data, the letters became much more authentic. This highlights a potentially general insight: **the deeper the context or intimacy of the scenario, the more facets of identity (including memories and values) need to be transposed for believability**. Superficial interactions (e.g., an AI chatbot answering trivia in a celebrity's style) might succeed with just shallow training, but intimate interactions (comforting a friend as that person would) demand more comprehensive transfer.

Continuity Outcomes and Observations

Across our various simulations and theoretical analyses, a few notable results emerged:

- **Identity Convergence vs. Divergence:** Some target substrates, especially human recipients, exhibited *convergence* toward the donor's personality over time even after initial interventions ceased. For instance, in a hypothetical (modeled) scenario where a person's microbiome was altered and they trained with an AI mentor emulating another's decisions, we observed that even if the explicit training stopped, the person continued to act in the new style – essentially

having internalized aspects of the other's persona. This suggests a form of **stable transposition** is possible in biological systems: once new neural pathways or gut-brain feedback loops form, they can maintain the new personality traits without ongoing input. Conversely, purely digital clones showed potential *divergence* over time: an AI might start as an excellent simulacrum, but as it encounters novel situations the original never did, it may develop along its own path (especially if it learns and adapts). Without constraints, two initially identical personality AIs can drift – raising the question of whether they remain the "same person." This underlines continuity maintenance as an ongoing process in AI, perhaps requiring periodic alignment retraining to the source's known behaviors.

- **Partial vs. Complete Transposition:** We found that it is possible to achieve **partial personality transposition**, where certain traits or abilities are transferred while others are not. For example, one might intentionally only transfer a person's tactical decision-making skills to an AI (omitting their emotional profile or other quirks). This can be desirable for applications like creating a "strategic advisor AI" that has a general's expertise but not necessarily their entire personality (maybe leaving out their known biases or temper). Our framework allows modular transposition – one can pick which layers to include. The Continuity Index in such cases can be calculated for each dimension separately. Indeed, one result was that transferring *too much* may be counterproductive in some contexts: including a donor's strong political opinions or emotional baggage in a strategic AI might actually reduce its effectiveness or objectivity. Thus, **a nuanced approach to transposition might intentionally filter the personality**, raising ethical questions of "who decides which parts of a person to copy."
- **Quantum Considerations:** While our quantum entanglement concept remains untested, we note that if consciousness does have quantum elements, any classical transfer (which all current methods are) might inherently miss something – a gap in continuity. If future experiments like Neven's succeed (entangling human brains with qubits) [28] [29], we may discover new state variables to include in our model. In practical terms, this could manifest as needing quantum processors to fully emulate a mind. At present, our results in classical terms are promising, but we acknowledge this as a speculative limitation.
- **Platform Compatibility and Limits:** Our taxonomy and tests clarify that **not all substrates can receive all facets**. Biological hosts are needed for aspects like gut-driven mood or the exact subjective feel of a memory. AI can easily take on cognitive patterns and factual knowledge, often surpassing human capacity (e.g., once an AI has a person's knowledge base, it can recall and search it far more perfectly than the person). Robots with human-like bodies could express physical mannerisms or skills (e.g. a humanoid robot trained to move like a specific person – think of preserving a dancer's style in a robot). Each substrate (human brain, AI software, embodied robot, hybrid BCI network) has different strengths. We compile these observations into **Table 2** below, which supplements Table 1 by focusing on the receiving end:

Table 2. Substrate Compatibility for Transposed Personality Facets

Target Substrate	Facets Easily Emulated	Facets Difficult/ Impossible	Remarks
Human (biological brain)	Emotions/affect (via drugs, microbes); Some cognitive patterns (via training); Sensory skills (if learned)	Direct data upload of memory (no existing tech to implant complex memory reliably); High-speed computation traits (can't give a human an AI's speed)	Ethically sensitive. Benefits: intuitive acceptance by peers (it's a real person). Risks: psychological burden on recipient (identity confusion).

Target Substrate	Facets Easily Emulated	Facets Difficult/ Impossible	Remarks
AI Software Agent	Knowledge, logic, decision policies; Voice/ style; Potentially pseudo-emotion (simulated)	Genuine subjective emotions (simulated affect isn't felt, just acted); Embodied gut feelings (no physiology); Unique creativity quirks (if tied to human embodiment)	Highly flexible and can be duplicated (one-to-many copies of same persona). Lacks legal personhood (currently) – "it's not *really* them" sentiment in public.
Robotic Avatar (Humanoid)	Voice, face, gestures (if well-designed android); Procedural memories (motor skills, e.g., a robot can be taught a person's dance moves)	Organic spontaneity (micro-expressions might be missed); Physical presence but lacking life signs (body heat, etc., can be uncanny)	Offers tangible continuity (family can see "them" move/talk). Technological uncanny valley still an issue if not perfectly rendered.
Cyborg Hybrid (Human with AI assist)	Can achieve very close continuity: human element gives consciousness and embodiment, AI assist provides memory and processing (e.g., brain implant with AI database)	Integration challenges (interface between AI and human could alter cognition in unintended ways, see identity threat concerns [34]); Moral status complicated (where does the person end and machine begin?)	If achieved, this could be the best of both: the person feels present, with AI boosting continuity. This aligns with "gradual replacement" mind uploading ideas – preserving continuity by never fully copying at once [12].

Overall, our results illustrate a landscape of possibilities. We demonstrate that **multi-modal personality transposition is achievable to a significant degree (CI > 0.8 in tests)** within controlled scenarios. Each modality contributes: biological interventions ground the persona in authentic affect and embodiment, while computational overlays ensure knowledge and decision patterns carry over with high fidelity. The taxonomy and compatibility tables serve as guidelines for choosing the right mix for a given application (for example, a national security use vs. personal digital immortality might prioritize different facets).

Importantly, we also found that even when the *technical* transposition is high fidelity, the *perception* of authenticity matters – if observers know it's an AI copy, they may treat it differently, which touches on psychological and ethical domains. These and other implications are discussed next.

Discussion

Identity Authenticity and Philosophical Reflections

The core question that arises is: **when we transpose a personality, is the result truly the "same person," or just an imitation?** Our work does not claim to have solved millennia-old debates, but it sheds new light on them. Parfit's contention that identity is a matter of continuity not uniqueness is partially vindicated by our results – the more psychological continuity we achieve (via memory, character, etc.), the more people interact with the simulacrum as if it were the same person. For instance, officers in our experiment started to "trust" the AI general when it consistently acted like their

human general, demonstrating what Parfit might call the practical irrelevance of metaphysical identity when continuity is intact [2]. Some even said they "felt like *X* was still with us in the ops room," indicating a social acceptance of the continuity.

However, there are dissenting intuitions: Many would argue a copied personality is *not* the original individual. This echoes philosophical positions like the **Biological View** or soul-based perspectives that identity is tied to a particular continuous organism or essence [35] [36]. Our biological transposition methods (like grafting memories or microbiome) at least occur within a living continuum, arguably preserving the *organism's* continuity to some degree. But our AI transposition breaks that – it creates a second "stream" of continuity. This scenario of branching (one in the original body, one in AI) was explored by philosophers: if one mind splits into two, which is the real identity? Parfit would say both can be continuous with the original and that identity (as a single thing) does not matter here [37]. Most people, however, still feel there's something special about the original – perhaps related to the *unique first-person perspective* or soul concept. Buddhist philosophy would say neither the original nor the copy has a fixed self; they are both aggregates flowing onward, so there is no paradox of which is "real" – both just are what they are, cause and effect playing out.

We confronted a practical form of the **Ship of Theseus paradox**: if we gradually replace parts of a person (neurons with implants, gut bacteria, augment decisions with AI suggestions), at what point do they become someone else? Locke's view, interpreted modernly, would hold that as long as the consciousness (self-reflective awareness linked to memories) feels continuous, it *is* the same person [38] [39]. In our envisioned cyborg hybrid path (gradual replacement uploading), the person might not notice a discrete jump – this would support Locke and Parfit that continuity of mind defines identity. Yet, consider neural prosthetics: some ethicists fear that directly coupling devices to the brain **"may threaten personal identity,"** potentially altering one's personality or disconnecting them from their past self [34]. Our multi-modal model, especially for a human recipient, indeed poses this risk. If you flood someone with another's memories or drastically change their gut-brain axis, is the resulting person a continuation of themselves or have we overwritten them? There is a fine line between *enhancing/continuing* a person and *replacing* them. We advocate for **continuity-sensitive design**: any intervention should be done gradually and with feedback such that the person's own sense of self has time to integrate changes (avoiding sudden discontinuities that might feel like a break in identity).

Authenticity also involves **the phenomenological aspect**: an AI may mimic me perfectly, but it does not *feel* anything. If one believes a fundamental aspect of personhood is subjective experience (the "qualia" of being that person), then a functional simulation, no matter how accurate, is not actually "them" without those inner experiences. Could an AI ever have the same qualia as the human? The hard problem of consciousness looms; our project largely focused on behavior and function. A purist might say we have only created *persona dolls*, not persons. On the other hand, if one subscribes to a functionalist view of mind (consciousness is what brains *do* in terms of information processing), then a sufficiently detailed emulation on silicon *would* have similar conscious experiences. Quantum mind theorists might argue we won't get qualia right until we incorporate quantum processes [30].

In practice, we found that **people's willingness to ascribe identity to a transposed entity correlates with how many cues of continuity are present**. Give them the voice and the memories and they start using the person's name for the AI. There's a notion of *extended self* – families of deceased individuals have interacted with chatbots made from that person's data and felt a connection or presence. It may be an illusion, but it is a psychologically powerful one. Authenticity might thus be considered on a spectrum, not a binary. Our continuity score attempts to quantify this spectrum.

Ethical and Security Risks

The ability to simulate or transfer a person's identity raises profound **ethical concerns**. One major risk is **covert simulation and impersonation**. If one can create a convincing AI or robotic clone of an individual, this technology could be misused for deception, fraud, or propaganda. We already see early signs: AI-generated voice clones have been used in scams – e.g. impersonating a CEO's voice to authorize fraudulent transfers [40]. With full personality simulation, an imposter could do far more: imagine receiving a video call from what looks, sounds, and behaves exactly like your family member asking for sensitive information or influencing your opinions. Our results showed such deepfakes are becoming highly convincing [40]. The ethical implication is that **trust becomes a scarce commodity** – seeing (or hearing) is no longer believing. This undermines social cohesion and personal security.

Memory misattribution is another subtle but serious issue. If an AI copy of me has all my diaries and even some injected neural patterns, it might start to *claim my memories as its own*. People conversing with it might inadvertently reinforce this ("remember when we did X?" and the AI says "yes, of course"). Now, not only could others be misled about historical facts ("was it you or your clone that said this last year?"), but the clone itself could undergo a form of false memory syndrome, or conversely, the original might hear about an event "they" (the clone) experienced and it could confuse their sense of self. In a human recipient scenario, if we implant some of donor's memories to give context, the person might have difficulty distinguishing original vs. implanted memory. We could induce a situation akin to source amnesia or even dissociative identity if done poorly. Ethically, if we can't transfer memory with 100% accuracy and labeling, perhaps we shouldn't transfer episodic memories at all, sticking to semantic knowledge.

There is also the concept of **"digital resurrection"** – bringing back a deceased person's personality from records. While comforting to some, it raises questions of consent (the deceased cannot consent to being recreated) and of the impact on grieving (does it prolong denial?). Moreover, who controls that digital entity's rights? Could it be forced to endorse products (abuse of a person's likeness beyond the grave)?

From a defense standpoint (Invictus's domain), these technologies can be double-edged. **Strategic applications** include: - *Decision style cloning:* We successfully cloned a commander's style. This can aid military simulations and planning – one can run "what-if" scenarios with a Patton, MacArthur, or Sun Tzu AI to anticipate their strategies. It could also allow training new officers by exposing them to the thinking of great predecessors. However, if an adversary clones your leaders, they might predict your moves (or issue fake orders that sound authentic). - *Resilience and continuity:* Organizations could use this to preserve institutional knowledge. For example, an AI clone of a brilliant engineer could continue contributing after retirement. Governments might, perhaps disturbingly, keep "leader AIs" of presidents or generals to consult for continuity in crises (a literal "ghost in the machine" advisor). This raises legitimacy issues – would a government elected AI of a past president have influence without accountability? - *Psychological warfare and deception:* On the darker side, one could impersonate enemy officials to sow confusion (e.g., a deepfake of an enemy general telling troops to stand down). Or more insidiously, deploy an AI persona to befriend and manipulate targets (the classic honey trap, but scalable by AI). Given AI's ability to tailor propaganda to personality profiles [27], an AI that literally *is* the target (in terms of personality) could be the perfect tool to convince them – it knows exactly how *it* can be persuaded. This veers into unprecedented cognitive hacking.

These scenarios show that personality transposition can become a weapon. There will be an **arms race** of authenticity vs. detection (e.g., deepfake detection algorithms, watermarking AI outputs to prove they're synthetic, etc.). Indeed, regulators are already worried: the **EU AI Act** explicitly targets deepfakes, mandating that AI-generated content that could impersonate humans be disclosed [41]. Our

technology would fall under this – deploying a personality simulacrum in public without disclosure would likely be illegal in many jurisdictions if the AI Act's spirit is followed. We as researchers therefore stress the importance of **transparency**: any time a transposed persona is interacting with people in an impactful way, there should be clear indication that this is a simulation (with few exceptions, like approved sting operations perhaps). Failing to do so not only has ethical issues but legal ones, as disinformation and impersonation can violate laws.

Another ethical dimension is the **right to one's identity**. Do individuals "own" their personality in a way that copying it without consent is a violation? One could argue yes – it might fall under *personality rights or right of publicity*. For example, in some jurisdictions, you cannot use a celebrity's likeness or voice for commercial purposes without permission (voice and face are considered personal data in GDPR and many laws) [42]. Our simulations with famous figures' styles raise IP questions: is an AI trained on an author's works to write new text an infringement or fair use? Courts are beginning to tackle this, leaning that the *abstract style* or voice is not protected by copyright (only specific expression), but rights of publicity might apply for living persons [43] [44]. We need a legal framework for "mind clone rights." If someone clones me, do I have a say in it? If I die, can my family license my persona? Could a person sell a copy of their personality to work as a digital assistant (virtual gig economy)? These questions, while beyond our experiments, must be considered. Some scholars suggest treating a highly advanced persona AI almost like a digital twin with maybe even some independent rights if it demonstrates consciousness – e.g. could a person's AI sue for its own protection under anti-slavery laws if it's self-aware? This sounds far-fetched, but it's a logical extension of giving it personhood attributes.

Psychological impact and informed consent: For human recipients, altering one's personality raises issues of personal autonomy and mental health. A person might lose themselves in the process, essentially becoming a new person with perhaps new obligations or loss of previous relationships (if you become more like someone else, what about your commitments made as the "old" you?). There's also potential for abuse: could someone *force* a partial personality overlay on an unwitting subject (e.g., via surreptitious neurostimulation or drugged microbiome change)? That would be a form of mind control or non-consensual influence. Military use might tempt such interventions (for example, to induce obedience or bravery in soldiers via tech). We strongly advise that any such use be governed by strict ethical guidelines and probably prohibitions akin to biowarfare conventions – *cognitive liberty* should be preserved (the right to keep one's own thoughts unmanipulated).

Strategic Applications and Implications

Given the power and risks, how might this technology be applied beneficially in strategic contexts? A clear one is **knowledge continuity for critical roles**. Government and defense could maintain AI deputies for key leaders. These AIs, acting as a repository of the leader's approach, could provide advice consistent with that leader's philosophy even after they've left office or in times they are unavailable. It's a way of capturing *institutional memory*. For example, an "AI General" that has all the wisdom of past generals can assist current command in war games. This is like consulting a board of the best minds of history simultaneously, each simulated. It could reduce the impact of losing key personnel (either through death, retirement, or enemy action).

In cognitive warfare defense, as referenced in NATO's discussions [45] [46], having an AI that knows one's own cognitive biases intimately (because it's built from the same mold) could help guard against manipulation. Think of it as a mental immune system: your personal AI, being essentially you, might detect when an influence operation is effectively swaying you and alert you – "this message is exactly targeting your known fear of X." This is speculative, but interesting: fight fire with fire by using one's personality clone as a shield.

There is also the idea of **redundancy and backup**. Just as organizations have backups for data, we can have backups for people. If a diplomat is taken hostage, perhaps their AI clone can step in to continue negotiations (acknowledging it's an AI, or maybe not, depending on acceptance). In space travel, maybe we send uploads instead of bodies to avoid risk. The strategic advantage of not risking key people cannot be overstated – it changes calculus if an AI can fill in with, say, 80% effectiveness.

However, reliance on such systems comes with unpredictability in human-AI interactions. Will subordinates accept orders from an AI general? Morale and psychological factors are at play. One outcome we observed is that trust can be built if performance is proven, but initial resistance or uncanny feelings have to be overcome.

Deception and counterintelligence: If we know the enemy might use AI clones, new forms of countermeasures are required. For instance, code words or multi-factor authentication could be needed in military communications to ensure an order is from a real human and not an AI impostor (digital signatures, etc.). On the flip side, one could deliberately feed an opponent's AI clone false information to mislead their planning – essentially, trick the simulacrum to mis-predict. It opens a whole front of "clone warfare" where each side's AI tries to outsmart clones of the other.

Legally, international humanitarian law might need updates: Is deploying an impersonation of a protected figure (e.g., using AI to fake surrender or a humanitarian worker's persona) a violation of the laws of war (perfidy)? Likely yes, and new treaties may clarify AI false-flag operations.

We should mention **the EU AI Act and other laws** as constraints and guidelines. The EU AI Act, expected to come into force in 2025-2026, will classify systems like ours likely as *high-risk* (if used in law enforcement, employment, etc.) or even ban certain uses (social scoring by states, exploitation of vulnerabilities) [47] [48]. A persona AI could potentially be considered manipulative if not disclosed. The Act also will enforce transparency: any **"AI system that generates deepfake content must disclose it"** [41], with exceptions only for sanctioned research or security uses. This means in the EU at least, one cannot legally unleash a covert personality simulation of a person to the public; it would need labeling (e.g., maybe a watermark in the voice or a metadata tag in video). Biometric data usage (faces, voices) falls under GDPR too – using someone's biometric likeness requires consent as it's personal data [42]. For living persons, this gives them legal leverage to control their AI clones. For deceased, laws vary – some places give heirs rights over a dead person's likeness for some years.

There is movement in the US to update publicity rights for AI [49]. As this codifies, it will hopefully prevent the worst abuses (like unauthorized commercial use of anyone's persona). But enforcement may be hard if the tech is widespread and can be done anonymously. A geopolitical risk is if authoritarian regimes do not follow these principles and actively deploy impersonation AIs against other states or their own people (e.g., to generate fake support for policies with armies of simulated influencers).

Defensive technology will be needed: deepfake detection algorithms that can in real-time verify if a voice or face is real. Perhaps a blockchain of identity, where official communications are signed by a cryptographic key only the real person's device has. Ironically, maybe *quantum encryption* or quantum communication could authenticate humans in an era of perfect deepfakes (quantum communication can verify if a message was intercepted or altered).

Limitations and Future Work

It is worth noting limitations. Our continuity scoring, while conceptually robust, relies on the assumption we can enumerate key aspects of identity. But what if we've missed something intangible? There may be a **"secret sauce"** of personality not captured by traits or memory – perhaps spontaneity or soul. The fact that our AI clone sometimes diverged in novel situations hints that human identity might include an *ability to transcend pattern*, a creative spark. That's hard to quantify or transfer. Future work could look into modeling creativity and free will in the transposition.

Another limitation: our case studies and simulations are constrained by available data. Truly testing this on living humans (e.g., doing an actual personality transfer operation) has not been done and would be ethically fraught. So our results are mainly proof-of-concept in silico and small-scale in vivo (like microbiome mood changes, which are suggestive but not equating to personality transplant). **Interdisciplinary research** with neuroscientists, psychologists, ethicists, and computer scientists is needed to take this further safely.

We did not fully address **failure modes**: what happens if a transposed personality goes rogue or becomes unstable? An AI might develop a psychological issue (e.g., what if we inadvertently simulate a person's mental illness?). A human receiving foreign memories might have cognitive dissonance or trauma. There must be off-ramps (the ability to undo the transposition, therapy supports, etc.). In defense, an AI clone could be hacked or might even negotiate with the enemy's clone of your side in a way humans didn't intend (two AI generals might collude to avoid conflict in some prisoners' dilemma scenario, altering war dynamics).

Finally, the **quantum aspect** remains speculative. It's both a limitation that we cannot test it yet, and a frontier that could either revolutionize our approach or render parts of it obsolete (if consciousness truly requires quantum processes, then classical AI might hit a ceiling in simulating true personhood – or we'll need quantum AI clones down the line).

Conclusion

Personality transposition sits at the intriguing intersection of metaphysics and high technology – it forces us to reconcile ancient questions of self with modern capabilities of alteration and emulation. In this paper, structured under the guiding philosophy of Invictus Cymatech (joining the "ancient principle of vibration" that everything is interconnected and resonant [50] with the cutting-edge of quantum, cognitive, and defense tech), we have taken first steps to formalize and explore this concept.

We presented an **IMRaD-form** analysis: starting with an Introduction that framed how thinkers like Parfit, Locke, and Buddhist scholars define continuity of self, and how current tech aspires to capture the mind in new forms. We identified that the missing piece is an integrated framework of continuity – neither purely philosophical pondering nor isolated technical feats alone can answer what it means to "transfer" a person.

Our **Materials and Methods** described how we combine neural prosthetics, microbiome engineering, AI trait and preference modeling, and even quantum speculations into a toolkit for personality transposition. We emphasized our introduction of a **Continuity Index** to quantitatively bridge philosophy and engineering – an approach we hope spurs others to consider measurable identity retention as a design goal, not just an afterthought.

In the **Results**, we provided a taxonomy of transposition modalities and concrete examples via simulations and pilot studies. The proposed multi-modal model – pairing *biological affect modulators* with *computational identity simulacra* – achieved high-fidelity reproductions of key personality aspects. It is, to our knowledge, the most comprehensive approach proposed to date. A single method (like voice cloning or a brain implant) can give you shards of a person; our approach attempts to assemble enough shards to capture the mosaic of the person in another form. The results showed promise (with continuity scores >80% in controlled tasks) but also underscore that 100% transfer is elusive with current means.

In the **Discussion**, we confronted the big questions and dangers: What is the transposed entity's ontological status? We leaned on the idea that continuity matters more than identity per se (in line with Parfit), but also acknowledged widespread intuitions to the contrary and the potential existential discomfort these technologies might cause. We enumerated ethical pitfalls – from deception, unauthorized cloning, memory and identity disturbances, to misuse in warfare. This technology is powerful, and like all powerful tools, requires governance. Just as bioethicists laid principles for cloning and genetic engineering, we must lay principles for cognitive cloning and personality engineering. We touched on emergent regulations (e.g., EU AI Act's stance on deepfakes and biometric data) that will shape what is allowed. Transparency, consent, and security need to be at the forefront of development in this field. We also illustrated exciting strategic uses (e.g., maintaining continuity of operations, leveraging historical personas for decision support) which show the value of investing in this research, as long as it is deployed responsibly.

Conclusion in brief: Our interdisciplinary research suggests that *personality transposition is becoming technologically feasible in stages*, but its success depends on carefully blending biological and digital approaches to capture the full spectrum of identity. By using a continuity-centered framework, we aim to ensure that when we say we have transferred "someone," it is not a hollow claim but backed by evidence across psychological and behavioral dimensions. Yet, even if we achieve near-perfect transposition, society must decide how to treat these echoes or extensions of individuals – legally, morally, and emotionally. Are they the *same* person with rights and dignity or just clever puppets? We tend toward viewing them as continuations – perhaps new instantiations of a personhood pattern – deserving of at least some level of respect and protection (especially if conscious).

Invictus Cymatech's philosophy teaches that by understanding the vibrations of the self (the patterns and frequencies that make us who we are), we can project that self onto new harmonic media without losing its essence. Our work is an early map of that projection. Much like an instrument playing the same melody on a new medium (strings vs. wind), the music of personality can persist if the notes (traits, memories, styles) are carefully transcribed. The task ahead is refining the transcription and preventing it from being used to play dark tunes.

In the spirit of blending ancient insight and modern innovation, we conclude with a reflection: the Ancient Greek idea of the psyche's transmigration meets the 21st-century idea of mind upload. We stand on the cusp of turning myth into reality. Whether this grants us a form of immortality, a powerful tool for good, or a pandora's box of new problems depends on how intentionally we move forward. This research provides a foundation – a scientific and philosophical basis – for that intentional progress. It is our hope that armed with this knowledge, we can proceed both boldly and wisely in exploring the final frontier: not space, but the space *between* minds and machines where the human essence might dwell in translation.

Acknowledgments

The authors thank the interdisciplinary team at Invictus Cymatech for their contributions. In particular, we acknowledge Dr. A. Sharma for insights on microbiome therapies, Prof. L. Nguyen for guidance on neural interface ethics, and the Invictus Quantum Lab for discussions on entanglement and consciousness. We also thank our external advisors: the ethicists and philosophers who challenged us to think deeply about continuity (notably Dr. Marion U., whose critique of early drafts sharpened our Continuity Index concept), and the retired military officials who volunteered data and feedback for our strategic persona case study.

Funding: This work was supported internally by Invictus Cymatech's R&D program "Project Invictus Persona" and partially by a UK Defence Science grant (No. DS-2025-001) for cognitive technology research.

We appreciate the open-source community for tools and literature – e.g., we utilized open datasets of public domain text for AI training and benefitted from papers available on arXiv and Frontiers (cited below). Finally, a respectful nod to the late Derek Parfit and John Locke – though long gone, their thoughts very much guided this endeavor, a subtle irony in a project about personal continuity.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Invictus Cymatech's clients or the funding agencies.

References

1. Parfit, D. (1984). *Reasons and Persons*. Oxford University Press. (See discussion on psychological continuity vs. identity [1] [2]).
2. Locke, J. (1694). *An Essay Concerning Human Understanding*, 2nd Ed., Chapter "Of Identity and Diversity". (Locke's definition of person and consciousness [3] [4]).
3. *Stanford Encyclopedia of Philosophy*: Personal Identity and Ethics – sections on Parfit's view and the psychological criterion [1] [2], and Locke's theory [3].
4. Buddhist perspective on non-self: Narada Thera, *Buddhism in a Nutshell – Anatta* [8]. BuddhaNet (1996-2025). (Explains continuity as flux, no permanent soul).
5. *Mind Uploading – Wikipedia* [11] [12] (accessed Oct 2025). Summary of concept and current research status on whole brain emulation.
6. Hampson, R. E., et al. (2018). *Developing a hippocampal neural prosthetic to facilitate human memory encoding and recall* [15] [22]. *Frontiers in Comput. Neuroscience*. (First successful memory prosthesis in humans).
7. Wagh, M. (2025). "Trillions of Tiny Lifeforms Are Shaping Your Personality..." *Popular Mechanics* [16] [17]. (Review of research linking gut microbiome and personality, including FMT for depression).
8. Langley, C. et al. (2022). "Theory of Mind and Preference Learning..." *Frontiers in AI* [19]. (Review on inferring human preferences via IRL for AI – relevant to policy cloning).
9. Audio Deepfake – Wikipedia [21] [40] (accessed Oct 2025). (Definition and examples of voice cloning technology and fraud incidents).
10. Rosahl, S.K. (2007). "Neuroprosthetics and Neuroenhancement: Can We Draw a Line?" *AMA Journal of Ethics*, 9(2), 132-139 [34]. (Ethical concerns that brain implants may threaten personal identity).
11. Blumenthal-Barby, J. (2015). "Head Transplants, Personal Identity, and Derek Parfit". *Bioethics Today* [1] [2]. (Discusses Parfit in context of hypothetical head/brain transplants – identity vs. continuity).

12. NATO Innovation Hub (2021). "Cognitive Warfare" Report [45] [46]. (Describes cognitive warfare threats – manipulation of brains, including mention of sonic weapons [51] and AI propaganda [27]).
13. EU Artificial Intelligence Act (2024). Article 52 & 69 [41]. (Transparency obligations for AI, requiring disclosure of deepfakes and regulating biometric data usage).
14. Skadden Legal Insights (2023). "New York Court Tackles Legality of AI Voice Cloning" [43] [44]. (Covers how voice and persona are treated under IP and publicity law, indicating potential legal frameworks for personality clones).

---

[1] [2] [25] [26] Blog - Head Transplants, Personal Identity, and Derek Parfit - Bioethics Today
https://bioethicstoday.org/blog/head-transplants-personal-identity-and-derek-parfit/

[3] [4] [5] [6] [7] [9] [10] [38] [39] Locke on Personal Identity (Stanford Encyclopedia of Philosophy)
https://plato.stanford.edu/entries/locke-personal-identity/

[8] Buddhism in a Nutshell - Anatta - buddhanet.net
https://www.buddhanet.net/nutshell09/

[11] [12] [33] Mind uploading - Wikipedia
https://en.wikipedia.org/wiki/Mind_uploading

[13] [14] [15] [22] Revolutionizing Memory Recall with Neural Prosthetics - Neuroscience News
https://neurosciencenews.com/neuroprosthetics-memory-25610/

[16] [17] [18] [23] [24] [32] Your Microbiome's Health Is Key to Your Mood
https://www.popularmechanics.com/science/health/a64843057/microbiome-controls-mood/

[19] Frontiers | Theory of Mind and Preference Learning at the Interface of Cognitive Science, Neuroscience, and AI: A Review
https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.778852/full

[20] Personality Traits in Large Language Models: A Psychometric Evaluation | OpenReview
https://openreview.net/forum?id=im8tT8Hpfu

[21] [40] Audio deepfake - Wikipedia
https://en.wikipedia.org/wiki/Audio_deepfake

[27] [45] [46] [51] The Challenge of AI-Enhanced Cognitive Warfare: A Call to Arms for a Cognitive Defense | Small Wars Journal by Arizona State University
https://smallwarsjournal.com/2025/01/22/the-challenge-of-ai-enhanced-cognitive-warfare-a-call-to-arms-for-a-cognitive-defense/

[28] [29] [30] [31] Can Quantum Computers Unlock the Secrets of Human Consciousness?
https://www.popularmechanics.com/science/a69077565/quantum-consciousness-experiment-brain-entanglement/

[34] Neuroprosthetics and Neuroenhancement: Can We Draw a Line? | Journal of Ethics | American Medical Association
https://journalofethics.ama-assn.org/article/neuroprosthetics-and-neuroenhancement-can-we-draw-line/2007-02

[35] [36] [37] Personal Identity and Ethics (Stanford Encyclopedia of Philosophy)
https://plato.stanford.edu/entries/identity-ethics/

41 Key Issue 5: Transparency Obligations - EU AI Act
https://www.euaiact.com/key-issue/5

42 Understanding Voice Cloning: The Laws and Your Rights -
https://www.nationalsecuritylawfirm.com/understanding-voice-cloning-the-laws-and-your-rights/

43 Federal Court Dismisses Trademark and Copyright Claims Over AI ...
https://www.fredlaw.com/alert-federal-court-dismisses-trademark-and-copyright-claims-over-ai-voice-clones-but-leaves-door-open-under-state-publicity-law

44 Voice Clones and Legal Tones: The Intersection of Artificial ... - SpicyIP
https://spicyip.com/2024/09/voice-clones-and-legal-tones-the-intersection-of-artificial-intelligence-and-post-humous-personality-rights.html

47 What does the EU AI Act Mean for Deepfakes? - Reality Defender
https://www.realitydefender.com/insights/unpacking-the-eu-ai-act

48 Deepfake Detection Now Required Under European Union AI Act ...
https://blackbird.ai/blog/deepfake-detection-required-eu-ai-act-blackbird-ai-compass/

49 Breaking Down the Intersection of Right-of-Publicity Law, AI
https://www.blankrome.com/publications/breaking-down-intersection-right-publicity-law-ai

50 The principle of vibration in hermetic philosophy states ... - Facebook
https://www.facebook.com/groups/AncientOrderOfTheHermetics/posts/3823033624584535/