Title
A Cognitive Model for Personality and Interaction Based on the Laban—Malmgren System for Movement and Acting

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors
Eriksson, Jesper
Lieck, Robert

Publication Date
2026

A Cognitive Model for Personality and Interaction Based on the Laban–Malmgren System for Movement and Acting

Jesper Eriksson (jesper.eriksson.25@ucl.ac.uk)
Department of Computer Science, University College London
Gower Street, London, WC1E 6BT, UK

Robert Lieck (robert.lieck@durham.ac.uk)
Department of Computer Science, Durham University
Stockton Road, Durham, DH1 3LE, UK

Abstract

Existing personality models describe stable individual differences well, but often leave underspecified how such differences generate moment-to-moment expressive behaviour. Acting theory faces a complementary problem: performers must reliably externalise a character's internal state so that intentions, affect, and personality can be inferred from movement, speech, and posture alone.

We address this gap by introducing a computational formalisation of the Laban–Malmgren System (LMS), a framework from movement analysis and acting pedagogy that links internal disposition to observable action. Building on Laban's analysis of movement qualities and Malmgren's extension for character work, we formalise four core internal and external dimensions of behaviour and develop a generative cognitive architecture in which: 1) stable trait parameters define priors over a latent expressive state, 2) state transitions are shaped by context and actions, and 3) a stochastic policy maps the latent state to observable behaviour.

The resulting model provides a principled, action-centred link between internal dispositions, state dynamics, and expressive behaviour, making the theory accessible to simulation, quantitative analysis, and embodied artificial agents. We demonstrate its computational feasibility through a public proof-of-concept implementation, *Persona*: an interactive digital portrait in which an avatar conveys its internal state through detailed bodily expression and responds dynamically to viewer interaction.[1] By formalising a theory rooted in expert artistic practice, this work offers a complementary, action-centred perspective on personality modelling and lays the groundwork for future empirical validation.

Keywords: cognitive architecture; computational modelling; personality traits; movement analysis; acting theory; interactive behaviour; embodied AI; affective computing; digital art

Introduction

Computational models of personality and character are increasingly relevant across cognitive science, human–computer interaction, and artificial intelligence, where there is a growing need for representations that combine stable individual differences with context-sensitive, time-varying internal states. Classical personality psychology offers robust trait-based frameworks, most notably the Five-Factor Model (McCrae & Costa Jr, 1999), as well as more mechanistic motivational and social–cognitive approaches that emphasise dynamics, learning, and person–situation interactions (Bandura, 1986; Mischel & Shoda, 1995). However, many of these models remain either largely descriptive or only indirectly connected to observable action.

We address this gap by introducing a compact generative cognitive model in which stable personality traits define a prior over a dynamically evolving internal state, which in turn generates context-sensitive behaviour. The model is grounded in acting and movement practice rather than academic psychology, specifically, the dance notation system developed by Rudolf Laban that was later extended by Yat Malmgren to the Laban–Malmgren System (LMS) of movement psychology and character analysis (Mirodan, 1997, 2015). Laban's work on movement analysis and notation—most prominently formalised in *Labanotation* (Guest, 2013; Laban & Ullmann, 1971)—was originally conceived as a descriptive system for human movement. Malmgren subsequently integrated these movement dimensions with character psychology as used in actor training, introducing an additional dimension and explicitly distinguishing between inner (psychological) and outer (expressive) aspects of character.

The resulting Laban–Malmgren System constitutes a compact, low-dimensional theory of character that has been used extensively by actors and directors to analyse, communicate, and embody personality through action. Importantly, this theory is grounded in expert practice: it reflects accumulated experiential knowledge about how personality manifests in behaviour, interaction, and movement, rather than being derived from questionnaires or psychometric factor analysis. While this practical grounding gives the theory considerable expressive power and intuitive coherence, it has so far lacked a formal computational specification and systematic empirical validation within cognitive science.

Our contribution is to formalise this acting-based character theory as a computational model, in which four high-level trait parameters represent stable experience-shaped tendencies that define priors over an individual's internal state space. These traits are assumed to be entrenched for a given individual at a given point in life, reflecting the idea that adult personality traits are relatively stable while still being shaped by prior experience (Roberts et al., 2006). Moment-to-moment behaviour, however, is not determined directly by traits, but emerges from a dynamic internal state that evolves through interaction with the environment. Context and interaction can therefore lead to substantial state changes, even though the underlying trait structure remains unchanged.

[1] The work premiered at the Saatchi Gallery, London (October 4, 2023), and was subsequently selected as a finalist for the 2024 Lumen Prize.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Conceptually, this places our model within contemporary integrative and hierarchical views of personality, in which traits, states, and behaviour operate on different time scales (McAdams & Pals, 2006). The model is closely aligned with social–cognitive accounts such as the Cognitive–Affective Personality System (Mischel & Shoda, 1995), in which stable behavioural patterns arise from the interaction between enduring person parameters and situational dynamics. At the same time, by grounding the trait dimensions in acting practice, our approach offers a complementary perspective to questionnaire-based trait models, emphasising action, expressivity, and interaction as primary phenomena.

By providing a precise computational formalisation, this work constitutes a first step towards integrating the Laban–Malmgren System of movement psychology and character analysis into the broader landscape of personality modelling. It enables the theory to be used in simulation, analysis, and future empirical studies, and opens a pathway for systematic validation and comparison with established psychological models. More broadly, the paper argues that theories developed in expert artistic practice can serve as valuable sources of structured knowledge for cognitive modelling, particularly when they are made explicit, formal, and testable.

Related Work

We briefly review the strands of personality psychology and cognitive modelling that our work relates to.

Trait taxonomies conceptualise personality as a small number of stable latent dimensions inferred from behavioural regularities and psychometric data. The most influential of these is the Five-Factor Model (FFM), which emerged from lexical and factor-analytic studies (Digman, 1990; Goldberg, 1990) and has been argued to show broad cross-cultural robustness (McCrae & Costa Jr, 1997). Extensions such as the HEXACO model introduce an additional Honesty–Humility dimension and strengthen links to moral and antisocial behaviour (Ashton & Lee, 2007). In our model, we adopt the assumption of stable individual differences along a small number of dimensions, but treat traits not as descriptive end-points, but as priors over an internal state space. In contrast to questionnaire-derived trait models, our trait dimensions are grounded in acting and movement theory.

To address within-person variability, integrative approaches conceptualise traits as distributions over momentary states, reconciling stability with contextual sensitivity (Fleeson, 2001). Longitudinal evidence shows that while personality traits exhibit systematic change across the life course, they remain relatively stable in adulthood (Roberts et al., 2006). Broader mechanistic perspectives emphasise the interaction between social experience and biological dispositions in shaping personality over time (Roberts & Jackson, 2008). Our architecture instantiates this view by treating traits as fixed on the modelling time scale and defining priors and state transition biases, while behaviour emerges from context-dependent state dynamics driven by interaction.

Social–cognitive theory highlights reciprocal interactions between cognition, behaviour, and environment (Bandura, 1986, 1991). Within personality psychology, the Cognitive–Affective Personality System (CAPS) framework models personality as stable patterns of situation–behaviour relations generated by interacting internal units (Mischel & Shoda, 1995). Our approach is closely aligned with CAPS in that stable individuality is expressed through systematic biases in state accessibility and dynamics, while behaviour remains highly context-dependent. Our contribution lies in providing a compact, formally specified state space grounded in expressive action and movement.

Motivational theories explain behaviour in terms of persistent needs and goals. Self-Determination Theory links behaviour to autonomy, competence, and relatedness, with strong empirical support (Deci & Ryan, 2000). Psychobiological models such as Reinforcement Sensitivity Theory explain personality differences via sensitivity to reward and punishment systems (Corr, 2004). We do not aim to provide a neurobiological account or explicit motivational variables in our model. However, our traits can be interpreted as stable modulators of reactivity and state accessibility, making the model compatible with these perspectives.

In artificial intelligence and human–computer interaction, personality and affect are often modelled to support believable agents and expressive interaction (Bates, 1994; Picard, 2000). Appraisal-based frameworks provide formal accounts of affective dynamics (Gratch & Marsella, 2004; Ortony et al., 1988). Separately, Laban-inspired representations have been widely used for analysing affect and expressivity in movement and dance (Camurri et al., 2003). Unlike perceptual or appraisal-focused approaches, our work introduces a generative, agent-centred personality architecture in which expressive behaviour emerges from a dynamic internal state shaped by stable trait priors, bridging acting theory, personality modelling, and embodied cognition.

Beyond trait and social–cognitive theories, our approach relates to a growing computational literature on personality dynamics. Read et al. (2017) develop computational models in which stable between-person differences coexist with substantial within-person variability across situations, showing that neural-network architectures can capture both trait structure and context-sensitive behavioural signatures. More broadly, Bayesian cognitive modelling provides a well-established framework for representing latent states, priors, and stochastic updating in cognitive science (Griffiths et al., 2008; Lee & Wagenmakers, 2014). Work on movement-based personality perception has further demonstrated that manipulating movement qualities, including Laban-derived parameters, systematically shifts observers' personality judgements of virtual agents (Durupinar et al., 2017). Our contribution differs from these approaches in grounding the latent state space in Laban–Malmgren categories and in using that state specifically to generate embodied expressive behaviour rather than modelling questionnaire responses or abstract trait trajectories. The model can be read as a structured generative state-space model whose latent variables are defined by LMS theory rather than based on statistical measures.

The Laban–Malmgren System

Conceptual Foundations

Laban–Malmgren System (LMS) (see also Mirodan, 1997, 2015) is founded on a non-dualist view of mind and body: psychological processes and physical movement form a single, mutually reinforcing system, with movement expressing and shaping mental and emotional states.

The theory posits a direct correspondence between four psychological dimensions or *mental factors* (based on the psychological functions by Jung & Beebe, 1921/2016) and four movement dimensions or *motion factors* (from Laban). Each mental factor manifests as an *inner participation* (a mode of engaging with the world) and is realised along a continuum between two opposing *elements* (yielding vs. contending), as summarised in Table 1.

Pairwise combinations of the mental factors define six possible inner attitudes, which correspond to the edges of the *simplex of mental factors* (Figure 1). Each inner attitude admits four possible characteristics, depending on the yielding/contending balance of the corresponding mental factors (Table 3). The inner attitudes are not treated as discrete states. Instead, all four mental factors are always present and differ only in their relative weighting, which is a stable character-level property. The relative ordering and dominance of the factors captures the qualitative organisation of a character—its perception, motivation, and action.

The *essence* of a character is its dominating inner attitude, defined by the two mental factors with highest weight. The LMS further distinguishes between the *inner drive* (or *inner action*), describing the internal psychophysical process shaping impulses, and the *outer drive* (or *outer action*), describing the externally observable behavioural strategy (Table 2 and Figures 1 and 2). Each drive has a yielding or contending expression. This layering allows for internal tension and subtext whilst preserving coherence at the level of the essence; a character may switch between inner and outer drives depending on circumstances.

Model Description

Our model preserves the qualitative structure of the LMS while formulating it in mathematically exact terms and rendering it computationally implementable. It formalises the LMS as a three-layer architecture, differentiating between an agent's stable **traits**, its dynamic **state**, and its goal-directed **actions** a (Table 4).

The probabilistic choices in the model (described below) follow the geometry of the LMS theory. The mental-state vector is compositional, representing the relative prominence of four mental factors, and therefore lives on a simplex. The Dirichlet distribution is the natural prior family for simplex-

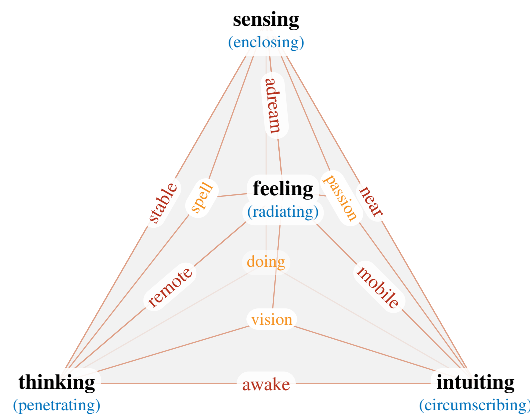

Figure 1: The simplex of mental factors: Each factor has a particular weight. A character's type is defined by the most dominant factor with highest weight. Combinations of two factors define six possible inner attitudes. Combinations of three factors define four possible drives (also see Table 2 and Figure 2). An interactive visualisation can be found here.

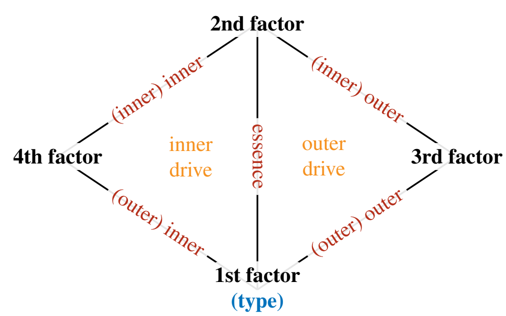

Figure 2: A character's *essence* is the inner attitude with highest weight, corresponding to the combination of the first and second mental factor. Together with the third and fourth strongest factor the essence defines an outer and inner drive, respectively, which determine the character's visible (outer) actions and latent (inner) actions and thoughts.

valued latent states, with its concentration parameter encoding trait rigidity. Each expression variable lies on a bounded yielding–contending continuum, motivating a Beta distribution for factor-specific expressive bias. The state transition equation implements a standard latent-state decomposition into context-dependent forcing, action-dependent feedback, and a trait-restoring drift derived from the KL divergence between current state and prior (a principled, information-theoretic measure of how far the state has drifted from the agent's dispositional baseline). Finally, the softmax action policy provides a well-established stochastic choice rule (Franke & Degen, 2023) that converts graded action compatibility into probabilistic action selection, consistent with Luce's choice axiom widely used in cognitive modelling.

The agent's traits are defined by the weights w of the four mental factors in the LMS as well as the weights y and c of their yielding and contending elements, respectively. These

Mental Factors (internal)	Motion Factors (external)	Inner Participations engaged / subdued	Expression yielding – contending
sensing	weight	intending / heavy	light – strong
thinking	space	attending / adrift	flexible – direct
intuiting	time	deciding / indecisive	sustained – quick
feeling	flow	adapting / irrelated	free – bound

Table 1: Overview of the four *mental factors* and *motion factors* in the Laban–Malmgren System. The corresponding *inner participations* characterise a way of engaging, perceiving, and interacting with the world, which is realised in a particular *expression* on a spectrum between *yielding* and *contending*.

Drives (mental factors)	Expression yielding – contending
doing (STI)	reacting – exerting
spell (STF)	surrendering – dominating
passion (SIF)	constructing – destroying
vision (TIF)	ideas – problems

Table 2: Every combination of three factors corresponds to a particular drive (also see Figures 1 and 2).

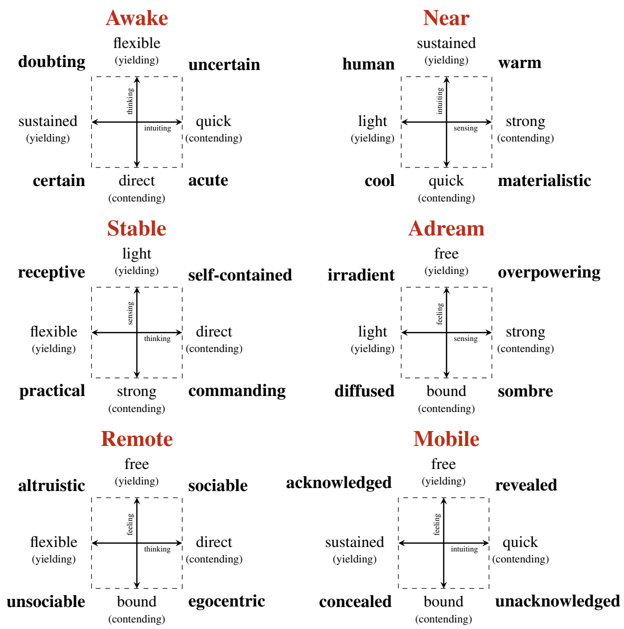

Table 3: Expression sub-spaces: Each inner attitude corresponds to two mental factors (axis labels), who's expressions (on a spectrum from yielding to contending) define a two dimensional sub-space, where each quadrant corresponds to a particular *characteristic* (bold labels at the corners).

	Parameters	Dynamic
Traits	$w, y, c \in [1, \infty]^4$	
States	$m_t \sim \mathrm{Dir}(w)$ $e_t^{(i)} \sim \mathrm{Beta}(y^{(i)}, c^{(i)})$	$\mathcal{T}(m_t, e_t, w, y, c, x_t, a_t)$
Actions		$\pi(a_t \mid o_t, x_t, s_t, e_t)$

Table 4: The three layers of our cognitive model with stable traits (w, y, c) and a dynamic state (m_t, e_t) that is influenced by the trait-based priors as well as the external context x_t and the agent's actions a_t. Interaction a_t emerges from the current state, context, and the agents objectives o_t.

define a Dirichlet prior $\mathrm{Dir}(w)$ over the agent's mental state m and independent beta priors $\mathrm{Beta}(y^{(i)}, c^{(i)})$ over their specific expression $e^{(i)}$. The agent's state thus consists of independent categorical distributions over the four mental factors and their (binary) expressions. The factor weights w and mental state m both correspond to points in the simplex of mental factors (Figure 1) with an additional concentration $|\alpha| = \sum_i w^{(i)}$ for the weights (trait strength/rigidity). Pairs of expressions $(e^{(i)}, e^{(j)})$ can be visualised as points in the expression sub-space of the corresponding inner attitude (Table 3). This formulation loosely couples traits and states and reflects the idea of the LMS that personality is stable but not absolute, so that inner attitudes adjacent to the character's essence remain accessible, although with varying probability.

State dynamics: The state dynamics are influenced by three aspects: the external **context** x, the agent's own **actions** a, and the trait-based **prior** w, y, c. The environment exerts directional pressure on specific factors (e.g. urgency $\rightarrow$ time/intuiting, resistance $\rightarrow$ weight/sensing, ambiguity $\rightarrow$ space/thinking), filtered through the agent's dominant inner attitude (its essence) defined by the weights w. The agent's own actions also feed back into its state: repeated/intense actions temporarily amplify their associated factors, producing drift, oscillation, or compensation effects. Finally, there are restoring forces that pull the state back toward the prior distribution. This can be written in a state transition function

$$\begin{aligned}\begin{bmatrix} m_{t+1} \\ e_{t+1} \end{bmatrix} &= \mathcal{T}(m_t, e_t, w, y, c, x_t, a_t) \qquad (1)\\ &= \begin{bmatrix} m_t \\ e_t \end{bmatrix} + \lambda_x f_x(m_t, e_t, w, y, c, x_t) \\ &\quad + \lambda_a f_a(m_t, e_t, w, y, c, x_t, a_t) \\ &\quad - \lambda_p \nabla \mathrm{D}_{\mathrm{KL}}(m_t, e_t \,\|\, w, y, c)\ ,\end{aligned}$$

where w, y, c defines the prior, x_t is the current external context, a_t is the agent's action, the λ values define the weight of the corresponding components, f_x and f_a define the directional change induced by the context and action, respectively, and $\nabla \mathrm{D}_{\mathrm{KL}}$ is the gradient of the Kullback–Leibler divergence between the current state and the prior (factorising into Dirichlet and beta components for m_t and e_t). Concrete applications require specific definitions of f_x and f_a.

Action selection: The policy π depends on the current state (m_t, e_t), the external context x_t, and the agent's objective o_t. The action space depends on the modelling scenario (see our embodied-AI application below); generally, in a given context x, every action a has: 1) a specific *factor profile* $\phi(a,x) \in \mathbb{R}^8$ that scores how well it fits with a particular expression (yielding/contending) of a particular factor and 2) an *alignment score* $\omega(a,o) \in \mathbb{R}$ that indicates how well it fits to the agent's objectives (which includes specific instrumental objectives, e.g., achieving a certain goal, as well as general desiderata, e.g., avoiding risks or conserving energy). We then define the action policy as

$$\pi(a_t \mid m_t, e_t, x_t, o_t) = \text{softmax}\Big[\frac{\text{score}(a_t, \mu_t, \varepsilon_t, x_t, o_t)}{\tau}\Big] \quad (2)$$

$$\text{with} \qquad \mu_t \sim \text{Cat}(m_t), \quad \varepsilon_t \sim \text{Ber}(e_t^{(\mu_t)}), \quad (3)$$

$$\text{score}(a_t, \mu_t, \varepsilon_t, x_t, o_t) = \gamma_\phi\, \phi^{(\mu_t, \varepsilon_t)}(a_t, x_t) + \gamma_\omega\, \omega(a_t, o_t)\ , \quad (4)$$

where τ is a softmax temperature, μ_t is the currently active mental factor sampled from the categorical distribution $\text{Cat}(m_t)$ defined by the current mental state m_t, ε_t is the binary expression (yielding/contending) for that factor sampled from the Bernoulli distribution $\text{Ber}(e_t^{(\mu_t)})$ defined by the state's expression $e_t^{(\mu_t)}$, and the overall score consists of two terms (weighted by the γ values) that capture: 1) the action's fit with the active mental factor μ_t in its expression ε_t in the given context x_t and 2) the action's alignment with o_t; concrete applications require specific definitions of ϕ and ω.

A natural variation samples a full ordering of mental factors via Gumbel-perturbed logits of m_t, yielding a noisy sample of type, inner attitude, and outer drive; we use this approach in our embodied-AI application.

Dynamic Generation Loop: By providing concrete definitions for the functions f_x and f_a (state dynamics) and ϕ and ω (action selection), we obtain a complete generative probabilistic model that allows for simulating the entire interaction dynamic. At every time step t, the agent:

1. is in a particular context x_t and has objectives o_t
2. samples a mental factor $\mu_t \sim \text{Cat}(m_t)$
3. samples it's expression $\varepsilon_t \sim \text{Ber}(e_t^{(\mu_t)})$
4. samples an action from the policy $a_t \sim \pi(a_t \mid m_t, e_t, x_t, o_t)$
5. updates its state as $(m_{t+1}, e_{t+1}) \leftarrow \mathcal{T}(m_t, e_t, w, y, c, x_t, a_t)$.

Application to Embodied AI

We demonstrated the practical applicability of the proposed cognitive architecture by implementing it in an interactive virtual character, *Persona*, realised as a digital portrait (Figure 3). The work was first exhibited at the Saatchi Gallery, London (October 4, 2023), where it engaged gallery visitors in real-time interaction and received critical attention for the perceived emotional depth and responsiveness of its expressive behaviour (Lawson-Tancred, 2023). It was subsequently selected as a finalist for the 2024 Lumen Prize (Lumen Prize, 2024), an international prize for art created with technology. Beyond its reception as an artwork, the installation served as a proof of concept, showing how the abstract trait–state–action distinction and the LMS-based internal dynamics can be operationalised in a closed perception–cognition–expression loop. The system was designed to make the internal dynamics of the model legible through embodied behaviour, with a non-expert audience as an implicit stress test of that legibility.

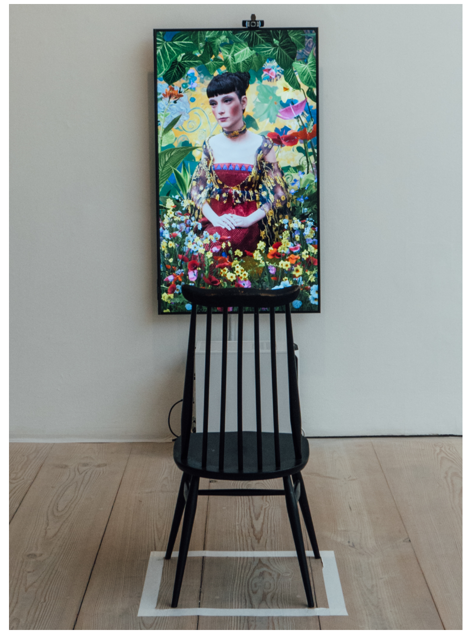

Figure 3: The final installation of our virtual character, *Persona* (Saatchi Gallery, London), driven by an implementation of our cognitive architecture.

System Overview: *Persona* implements an autonomous character that continuously cycles through perception, internal state update, and expressive behaviour. The system consists of four components:

1. **Perception**, which analyses a visitor's movements in terms of LMS-compatible motion qualities.
2. **Internal state dynamics**, which update the character's internal configuration based on both autonomous processes and interaction.
3. **Action selection**, which maps internal states to discrete expressive categories and bodily animations.
4. **Visualisation**, which renders character itself as well as the evolving internal state in the simplex of mental factors for diagnostic and analytical purposes.

This structure mirrors the theoretical model: stable personality traits are encoded implicitly as distributions over internal states, while moment-to-moment behaviour emerges from the dynamics of the internal state itself. *Persona* is configured to operate primarily within a single inner attitude (*Adream*), defined by the factor weights w in the model. This allows the implementation to focus on a well-defined region of the state space while preserving expressive richness through variation in the drive (*passion*/*spell*) and expression.

Perception and Inference of Observer Behaviour: On the perceptual side, the system uses real-time pose estimation to track the visitor's upper body and hands. Short temporal windows of movement are analysed by neural models trained to estimate movement qualities along the LMS *space* and *time* dimensions (*flexible–direct* and *sustained–quick*). These are the lowest-weighted motion factors in the *Adream* inner attitude (the character's essence), which means that they are most volatile and induce state dynamics that switch between *passion* and *spell* drive. The perceptual inputs are discretised into four distinct values for the *observer sentiment* (high yielding, low yielding, low contending, high contending), which was used as context x_t in our model. The discretisation allowed us to define the context-depended state dynamic function f_x as a simple matrix with drift terms.

Temporal Dynamics: The internal state dynamic was implemented as a stochastic dynamical system based on our model equation (1). We distinguished two regimes.

Autonomous dynamics (idle behaviour). In the absence of interaction, the internal state dynamics is dominated by the prior drift term in equation (1) with some additional noise. This implements notion of stable traits as priors over states.

Interaction-driven dynamics. When an observer engages with the character, the inferred observer sentiment acts as the context x_t, which creates a directional force on the internal state through the f_x term in (1). Depending on the compatibility between observer and character expression, this force shifts the state towards more yielding or more contending regions in the e_t component of the state. We added temporal smoothing of the context x_t to simulate memory and to produce gradual and history-dependent responses.

In addition, repeated or intense interactions accumulated in an exhaustion signal, which entered the action policy π via the agent's objective o_t (in this case: to preserve energy by withdrawing from interaction), and temporarily reduced its responsiveness. Together, these mechanisms ensure inertia, path dependence, and limited volatility—properties that are central to psychologically plausible behaviour.

From Internal State to Expressive Behaviour: At each update step, the continuous internal state (m_t, e_t) is translated into expressive behaviour through the model's dynamic generation loop, as described above. An ordering of mental factors is sampled from the mental state m_t, which defines the current inner attitude (exclusively *Adream* in our setting) and drive (*passion* or *spell*). The specific *characteristic* of *Adream* (*irradient*, *overpowering*, *sombre*, or *diffused*; see Table 3) is generated by independently sampling the expression of the two mental factors (*feeling* and *sensing*) based on e_t. The action policy π is a manually defined lookup table for the combinations of drive, characteristic, and observer sentiment (context) x_t obtained through the perceptive system. Each action maps to a specific motion primitive for the character that was pre-recorded through motion capture with a real actor. These motion primitives are executed by a real-time rendering engine that controls the avatar's posture, gesture, and micro-movement. Idle poses, subtle reactions, and more emphatic responses are distinguished, allowing the character's behaviour to remain expressive even in low-interaction scenarios.

Because all intermediate representations are phrased in LMS terminology, the system remains interpretable: internal logs can be analysed in theoretical terms, and motion assets can be authored explicitly to correspond to specific psychological constructs.

Integration and Proof of Concept: The complete system operates as a continuous closed loop: visitor movement is perceived and abstracted, the internal state evolves according to the autonomous and interaction-driven dynamics defined by our cognitive model, expressive behaviour is generated and rendered in real time, and the internal state is simultaneously visualised as trajectories in the simplex of mental factors. *Persona* thus demonstrates that the proposed formalisation is computationally implementable, interpretable, and suitable for embodied interactive systems, establishing a concrete platform for future empirical tests of whether the latent variables and generated behaviours correspond to human judgements and behavioural data.

Conclusion

We presented a cognitive model of personality and interaction that combines stable trait-like dispositions with a dynamically evolving internal state in a hierarchical architecture. The model formalises the Laban–Malmgren System, an established theory from movement analysis and acting pedagogy, into a mathematically precise and computationally implementable framework, aligning with contemporary integrative views in which traits act as long-term constraints while behaviour emerges from context-sensitive state dynamics. Its distinctive contribution lies in the origin in expert knowledge from acting and movement practice, where personality must be externalised and interpreted through action: this grounding yields a compact, interpretable state space tied directly to embodied expression, and renders it accessible to computational modelling, simulation, and quantitative analysis, as illustrated by our proof-of-concept implementation in an embodied interactive agent.

We view this work as a first step toward integrating acting-based theories of character into the broader landscape of cognitive and personality modelling. A natural next step is empirical validation along two axes. First, perceptual validation: does manipulating the model's latent trait and state variables produce movement and interaction patterns that observers reliably judge as reflecting distinct personalities or dispositions (Durupinar et al., 2017; Thoresen et al., 2012)? Second, behavioural validation: can the model's parameters be fit to repeated behavioural or motion data, and does it predict trajectories better than simpler baselines such as static-trait or purely reactive models? More broadly, we hope this contribution encourages closer dialogue between cognitive science and expert artistic practice, and supports the use of movement- and action-centred theories as viable foundations for computational models of personality and behaviour.

Acknowledgments

We would like to thank everyone involved in the realisation of the *Persona* exhibition at Saatchi Gallery, London, in particular: Phil Stilgoe (Executive Producer, Centroid), Simon Kay (Performance Capture Consultant, Animator), Joe Ells (Post Production Manager, Centroid), Iain Silvester (Studio Manager, Centroid), Jakob Abramson (Full-stack Engineer), Anthony Bourached (Research Associate, University College London), Rebecca Wield (Movement Director, Royal Central School of Speech and Drama), Giles Foreman (Movement-Psychology Consultant, GFCA). Centroid Motion Capture (Production Company).

References

Ashton, M. C., & Lee, K. (2007). Empirical, theoretical, and practical advantages of the HEXACO model of personality structure. *Personality and social psychology review*, *11*(2), 150–166.

Bandura, A. (1986). Social foundations of thought and action. *Englewood Cliffs, NJ*, *1986*(23–28), 2.

Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes*, *50*(2), 248–287.

Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, *37*(7), 122–125.

Camurri, A., Lagerlöf, I., & Volpe, G. (2003). Recognizing emotion from dance movement: Comparison of spectator recognition and automated techniques. *International Journal of Human-Computer Studies*, *59*(1), 213–225.

Corr, P. J. (2004). Reinforcement sensitivity theory and personality. *Neuroscience & Biobehavioral Reviews*, *28*(3), 317–332.

Deci, E. L., & Ryan, R. M. (2000). The "What" and "Why" of Goal Pursuits: Human Needs and the Self-Determination of Behavior. *Psychological Inquiry*, *11*(4), 227–268.

Digman, J. M. (1990). Personality structure: Emergence of the five-factor model. *Annual review of psychology*, *41*(1), 417–440.

Durupinar, F., Kapadia, M., Deutsch, S., Neff, M., & Badler, N. I. (2017). PERFORM: Perceptual approach for adding OCEAN personality to human motion using laban movement analysis. *ACM Transactions on Graphics*, *36*(1).

Fleeson, W. (2001). Toward a structure-and process-integrated view of personality: Traits as density distributions of states. *Journal of personality and social psychology*, *80*(6), 1011.

Franke, M., & Degen, J. (2023, September 28). The softmax function: Properties, motivation, and interpretation.

Goldberg, L. R. (1990). An alternative" description of personality": The big-five factor structure. *Journal of Personality and Social Psychology*, *59*(6), 1216–1229.

Gratch, J., & Marsella, S. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research*, *5*(4), 269–306.

Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition. In R. Sun (Ed.), *The cambridge handbook of computational psychology* (pp. 59–100). Cambridge University Press.

Guest, A. H. (2013). *Labanotation: The system of analyzing and recording movement*. Routledge.

Jung, C., & Beebe, J. (2016). *Psychological types*. Routledge (Original work published 1921)

Laban, R., & Ullmann, L. (1971). The mastery of movement.

Lawson-Tancred, J. (2023, October). An A.I. has created interactive portraits showing uncannily humanlike emotions. *Artnet News*. https://news.artnet.com/art-world/ai-interactive-portraits-humanlike-emotions-2377579

Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.

Lumen Prize. (2024). Persona: 2024 futures award finalist. https://lumenprize.org/futures-award-finalists-2024/persona

McAdams, D. P., & Pals, J. L. (2006). A new Big Five: Fundamental principles for an integrative science of personality. *American psychologist*, *61*(3), 204.

McCrae, R. R., & Costa Jr, P. T. (1997). Personality trait structure as a human universal. *American psychologist*, *52*(5), 509.

McCrae, R. R., & Costa Jr, P. T. (1999). A five-factor theory of personality. *Handbook of personality: Theory and research*, *2*(1999), 139–153.

Mirodan, V. (1997). *The Way of Transformation (the Laban-Malmgren system of dramatic character analysis)* [Doctoral dissertation, Royal Holloway, University of London].

Mirodan, V. (2015). Acting the metaphor: The Laban–Malmgren system of movement psychology and character analysis. *Theatre, Dance and Performance Training*, *6*(1), 30–45.

Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological review*, *102*(2), 246.

Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge university press.

Picard, R. W. (2000). *Affective computing*. MIT press.

Read, S. J., Smith, B. J., Droutman, V., & Miller, L. C. (2017). Virtual personalities: Using computational modeling to understand within-person variability. *Journal of Research in Personality*, *69*, 237–249.

Roberts, B. W., & Jackson, J. J. (2008). Sociogenomic Personality Psychology. *Journal of Personality*, *76*(6), 1523–1544.

Roberts, B. W., Walton, K. E., & Viechtbauer, W. (2006). Patterns of mean-level change in personality traits across the life course: A meta-analysis of longitudinal studies. *Psychological bulletin*, *132*(1), 1.

Thoresen, J. C., Vuong, Q. C., & Atkinson, A. P. (2012). First impressions: Gait cues drive reliable trait judgements. *Cognition*, *124*(3), 261–271.