Adaptive Multi-User Smart Mirror with Identity-Aware Profiling and LLM-Driven Empathetic Health Coaching for Elderly Care

Annappa S. S.[1]; Asha S.[2]; Lohith C.[3]

[1]Guest Lecturer, Department of Computer Science, Sahyadri Science College, Shimoga, India
[2]Guest Lecturer, Department of Computer Science, Sahyadri Science College, Shimoga, India
[3]Assistant Professor, Department CSE (IoT & CSBT), East Point College of Engineering and Technology, Bengaluru, India

Publication Date: 2026/07/28

Abstract: Population ageing is increasing demand for continuous, low-burden monitoring of physical and emotional wellbeing in the home, yet most smart-mirror prototypes to date are single-user devices tuned to one fixed idea of a healthy expression or posture. This paper proposes an IoT-enabled Smart Mirror architecture built around three ideas that, to the authors' knowledge, have not previously been combined in a single mirror-based system: (i) identity-aware sensing that disambiguates several household or care-home residents sharing one device using fused facial and gait signatures; (ii) a per-user adaptive baseline that learns each resident's own resting facial-expression and posture pattern over time and flags deviations from that individual baseline rather than from a fixed population norm; and (iii) a large-language-model (LLM) coaching layer that turns the fused emotional-physical state into short, empathetic, context-aware spoken or on-screen guidance, with an explicit safety-guardrail stage before anything reaches the user or a caregiver. An edge-cloud split keeps the sensing-to-deviation loop local for low latency while the LLM coaching and longitudinal analytics run in the cloud. Because no hardware prototype has yet been built, this paper deliberately stops short of reporting experimental accuracy figures; instead it develops the mathematical formulation of the identity-aware fusion and adaptive-baseline mechanism and presents a small, clearly labelled synthetic simulation to illustrate how the deviation signal and the edge-cloud latency budget would behave. The paper is intended as a design and feasibility contribution that can guide a subsequent hardware implementation and user study.

Keywords: *Smart Mirror, Multi-User Recognition, Identity-Aware Sensing, Adaptive Baseline, Large Language Models, Health Coaching, Elderly Care, Emotion Recognition, Posture Recognition, Internet of Things, Edge-Cloud Computing.*

How to Cite: Annappa S. S.; Asha S.; Lohith C. (2026) Adaptive Multi-User Smart Mirror with Identity-Aware Profiling and LLM-Driven Empathetic Health Coaching for Elderly Care. *International Journal of Innovative Science and Research Technology*, 11(7), 1480-1486. https://doi.org/10.38124/ijisrt/26jul1281

I. INTRODUCTION

Interactive smart-mirror platforms sit at the intersection of ambient assisted living, affective computing, and human-computer interaction, and are increasingly proposed as a non-invasive way to fold health and emotion monitoring into a daily routine that people already perform — looking in the mirror. Early systems combined a camera and a display to surface mood, hydration, or posture cues to a single user [1]–[3]. More recent multimodal designs fuse facial expression with body posture using guided feature fusion and attention mechanisms so that occlusion or a fixed viewing angle in one modality is compensated for by the other.

Two gaps remain open in this line of work. First, almost all published smart-mirror prototypes assume a single, previously-enrolled user standing in front of the glass at a time; a family home or a shared care-home bathroom, by contrast, is used by several people across a day, and a system that silently pools their data or applies one population-level threshold to everyone risks both false alarms for a naturally expressive resident and missed warnings for a naturally reserved one. Work on shared voice assistants has already shown that pooling multiple residents into one profile produces a measurable failure mode — persona confusion — in which one person's history or preference leaks into the system's response to somebody else [11]. The same risk applies to a shared visual mirror, but it has not, to the authors' knowledge, been examined in that setting.

Second, the feedback loop in most smart-mirror systems is a fixed rule or a simple classifier label displayed on screen. Large language models have recently been explored as flexible, conversational health coaches that can turn structured physiological or behavioural signals into personalised, motivational language rather than a static label [9], and dedicated smart-mirror prototypes have begun to add a chatbot layer alongside facial-affect and physiological sensing [7], [10]. However, giving an LLM a caregiving voice inside the home also raises safety and trust questions that recent work on elderly-facing chatbots has begun to catalogue [13], and these have not yet been folded into a mirror-specific design.

This paper addresses both gaps together. It proposes a Smart Mirror architecture in which (i) an identity-resolution stage recognises which of several enrolled residents is currently in front of the mirror from a fused facial-embedding and gait signature, (ii) each resident's facial-expression and posture features are compared against a running, per-user statistical baseline rather than a fixed population threshold, and (iii) the resulting personalised deviation score conditions an LLM prompt that generates short, empathetic coaching text, filtered through an explicit safety-guardrail stage before it is shown on the mirror or escalated to a caregiver. The remainder of the paper reviews related work, develops the mathematical formulation of the proposed pipeline, and presents a small synthetic simulation — clearly labelled as such throughout — that illustrates the qualitative behaviour of the adaptive-baseline signal and the edge-cloud latency budget the design implies. No hardware has yet been built and no human-subject data has yet been collected; the contribution of this paper is the design, its mathematical formulation, and the feasibility reasoning that supports building and evaluating it.

II. RELATED WORKS

Facial and body-language emotion recognition. Deep convolutional networks have been widely used to classify facial expressions into discrete affect categories from single images, achieving competitive accuracy on curated datasets but remaining sensitive to lighting, pose, and limited training data [and earlier reviews of the field have catalogued these constraints repeatedly]. Fusing facial features with skeleton-based posture cues, combined through guided or attention-based fusion, has been shown to recognise emotional and physical state more robustly than either modality alone, since posture compensates for facial occlusion or an unfavourable camera angle and vice versa.

Smart mirrors for home and elderly health monitoring. Silapasuphakornwong and Uehira proposed an IoT display device that screens for early signs of depression in elderly users through facial, voice, and posture cues while also functioning as a telemedicine communication tool Silapasuphakornwong and Uehira proposed a semi-electronic mirror device intended to screen elderly users for early signs of depression, combining facial recognition, speech recognition, and posture detection alongside a chatbot and telemedicine link [1]. Bianco et al. built a general-purpose smart mirror that adds short- and long-term emotion monitoring together with a double-authentication privacy protocol and a voice-assistant skill integration [2], which is one of the few prior systems to consider more than one identity at the mirror, although through explicit login rather than passive multimodal identification. Rachakonda et al. proposed iMirror, an IoMT stress-detection mirror for smart-city deployment that reported a high real-time stress-change detection rate in their evaluation [3]. A 2025 conference paper on “Reflective Wellness” likewise combines AI-based physiological and posture analysis in a mirror form factor for personalised health insight [5].

Kedanjoth et al. proposed an IoT multimodal framework that fuses facial, speech, and text emotion recognition at the decision level using ResNet-50, CNN-LSTM, and RoBERTa back-ends, reporting that decision-level fusion across modalities improved robustness over any single channel [4]. Santofimia et al. describe Sara, a virtual-assistant smart mirror developed within the EU SHAPES project that targets active and healthy ageing, and discuss the personalisation, privacy, and engagement barriers that must be overcome before such a system sees broad adoption by older adults [6]. Most recently, Kasno and Jung built a real-time smart mirror that streams a moving-average remote-photoplethysmography (MA-rPPG) signal from facial video alongside DeepFace-based affect detection and a GPT-4o mental-health chatbot, explicitly framing the result as a single-subject feasibility study rather than a validated clinical system [7]. Pu et al. separately proposed a transfer-learning based elderly-care emotion-recognition system that integrates multiple pretrained modalities to address the diversity of emotional expression in older adults, a population that standard single-modality classifiers trained on younger cohorts tend to under-serve [8].

LLM-based health coaching and elderly-facing conversational agents. Jörke et al. developed GPTCoach, a physical-activity coaching chatbot that implements the onboarding conversation of an evidence-based coaching programme and uses motivational-interviewing strategies together with wearable data to build a personalised activity plan, finding in a lab study that participants were comfortable sharing detailed context with the system [9]. Yang et al. built Talk2Care, an LLM-based voice assistant that mediates asynchronous communication between older adults and their healthcare providers, addressing a communication gap rather than sensing itself [10]. These systems demonstrate that an LLM can turn structured personal data into supportive, individualised language, which motivates using the same idea as the final stage of a sensing-driven smart mirror rather than only as a standalone chat interface.

Multi-user identity awareness in shared assistants. Al-Ratrout et al. recently showed that a shared voice assistant used by multiple residents of the same household measurably conflates their conversational histories — a failure they term persona confusion — and proposed the Adaptive Friend Agent (AFA), which combines voice-based speaker identification with per-user memory stores, improving a persona-attribution-accuracy metric substantially over a no-routing baseline [11]. This finding is directly relevant to any shared home device, including a mirror used by more than one family member or care-home resident, and is one of the central motivations for the identity-resolution stage proposed in this paper.

Safety of LLMs in elder care. Systematic reviews of LLM-powered chatbots for older adults report growing patient-facing use for supportive conversation and symptom information, but also note that most studies remain small, feasibility-stage pilots with median technology-readiness levels well below deployment [12], and dedicated safety work has begun to build taxonomies and benchmarks specifically for elderly-chatbot interaction risk [13]. These findings motivate treating the LLM layer in the proposed design as a guarded, escalation-aware coaching assistant rather than an autonomous diagnostic agent, and they are the basis for the explicit safety-guardrail stage in Section III.

Taken together, the literature shows steady progress on (a) multimodal facial-posture emotion recognition, (b) mirror-specific health and wellness sensing including physiological proxies, and (c) LLM-based coaching and elderly-facing conversational support, but treats each of identity-aware multi-user disambiguation, personalised (rather than population) baselines, and LLM safety guardrails largely as separate concerns. The system proposed next brings these three elements together inside a single smart-mirror pipeline.

➢ *Proposed System*

The proposed Smart Mirror is organised into five cooperating stages, shown in Fig. 1: multi-user Sensing, Identity Resolution, Identity-Aware Feature Extraction and Fusion, a Per-User Adaptive Baseline and deviation estimator, and an LLM Coaching and Cloud Analytics layer that drives the mirror display, a caregiver-escalation channel, and long-term storage.

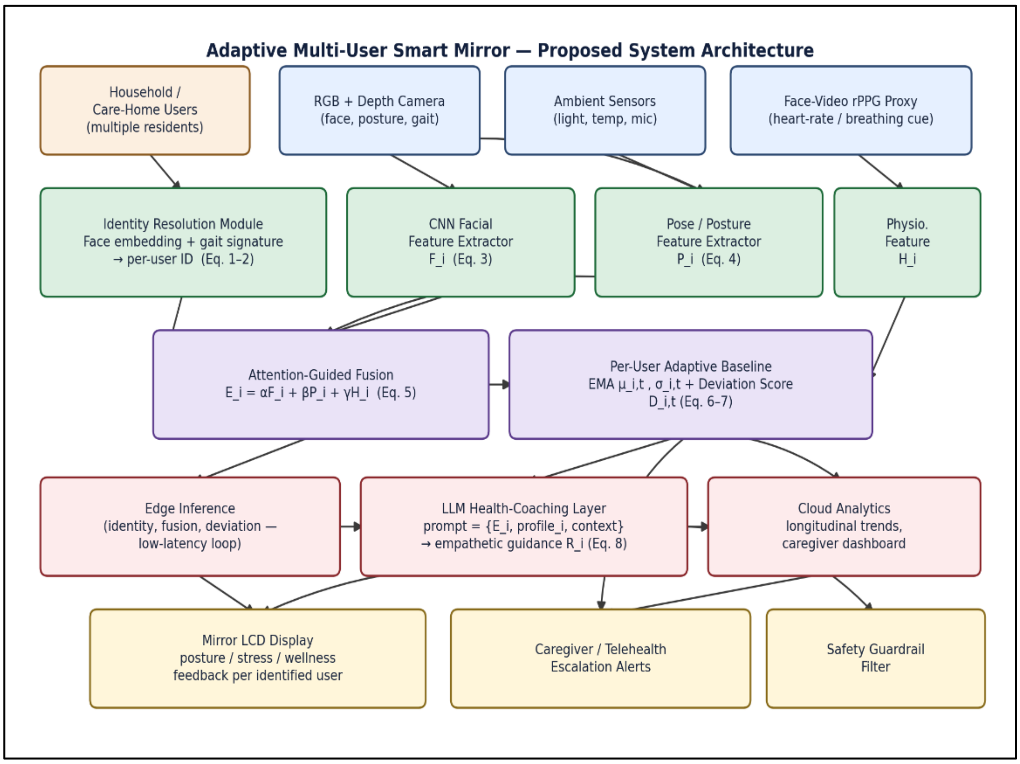

Fig 1 Proposed Architecture of the Identity-Aware, Multi-User Smart Mirror with LLM-Based Coaching.

Sensing. An RGB-depth camera pair captures facial video and skeletal joint positions for whoever is standing in front of the mirror; a microphone and ambient light/temperature sensors provide contextual signals; and the same facial video is reused, at low computational cost, to extract a remote-photoplethysmography (rPPG) proxy signal correlated with heart rate and breathing rate, following the feasibility approach reported for a related single-user mirror system [7].

Identity Resolution. Because the mirror is meant to be shared by several household or care-home residents, the first processing stage resolves which enrolled user is currently present. A facial embedding and a short gait/posture signature (extracted from the first few seconds of skeletal motion) are fused into a single identity vector, and the resident with the closest match above a confidence threshold is selected; below that threshold, the system falls back to an anonymous “guest” mode with no personalised baseline. This stage is the direct analogue, in the visual-sensing domain, of the voice-based speaker-identification stage that was shown to prevent persona confusion in shared conversational assistants [11].

Identity-Aware Feature Extraction and Fusion. Once an identity is resolved, per-user facial, posture, and physiological features are extracted and combined through an attention-guided fusion, following the same general fusion principle used in single-user mirror and multimodal-emotion literature [1], [4], but here the fused representation and everything downstream of it is keyed to the resolved user identity rather than to the device as a whole.

Per-User Adaptive Baseline. This is the central novelty of the proposed design. Rather than comparing a user's fused state against a fixed population-level threshold for “stressed” or “slouching”, the system maintains a slowly-updating

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1281

exponential moving average and variance of each resident's own fused state, and computes a personalised deviation score from that individual baseline. A naturally reserved resident and a naturally expressive one are therefore each compared only against their own typical pattern, which is expected to reduce both false escalations for expressive users and missed escalations for reserved ones — though this expectation still needs to be verified empirically once a prototype exists.

LLM Coaching and Safety Guardrail. When the deviation score for a resolved identity crosses a personalised threshold, the fused state, the resident's short-term profile, and the deviation score are assembled into a prompt for a large language model, which is asked to produce a short, empathetic, non-diagnostic coaching message in the style explored by prior LLM health-coaching systems [9], [10]. Before anything reaches the mirror display or a caregiver, the generated text passes through a safety-guardrail filter that checks for diagnostic claims, ensures an appropriate escalation tone, and enforces the interaction principles raised by recent elder-care chatbot safety work [12], [13]. Edge computation is used for the sensing-to-deviation loop to keep that loop low-latency, while the LLM call and long-term trend analytics run in the cloud, following the same edge-cloud split rationale used in prior IoT-based multimodal health-monitoring designs [4].

III. METHODOLOGY

The methodology below formalises the five stages of Fig. 1 for a single frame at time t and a candidate resident i, drawn from a set of U enrolled users.

➢ *Data Acquisition*

Let the captured facial image be I_f and the skeletal joint set be S_j, and let the short-window facial video used for the physiological proxy be V_f:

$$I_f \in \mathbb{R}^{(H\times W\times 3)} \quad (1)$$

$$S_j = \{(x_1,y_1), (x_2,y_2), \dots, (x_n,y_n)\} \quad (2)$$

Skeletal joints are obtained with a pose-estimation model such as OpenPose or MediaPipe, consistent with prior mirror and posture-recognition systems.

➢ *Identity Resolution*

A facial embedding e_f and a short gait/posture signature e_g are computed and fused into an identity vector, and the resolved identity i* is the enrolled user whose stored template has the highest similarity above a confidence threshold θ:

$$v\_id = \lambda\cdot e\_f + (1-\lambda)\cdot e\_g \quad (3)$$

$$i^* = \mathrm{argmax}\_i \ \ \mathrm{sim}(v\_id, \mathrm{template}\_i), \ \text{subject to sim} > \theta \quad (4)$$

If no enrolled template exceeds θ, the frame is processed in an anonymous guest mode with no personalised baseline update, which keeps an unrecognised visitor from contaminating any resident's profile.

➢ *Identity-Aware Feature Extraction and Fusion*

For the resolved user i, a CNN extracts facial features, a pose model computes posture-related joint angles and alignment metrics, and the physiological proxy is derived from the short facial-video window:

$$F\_i = \mathrm{CNN}(I_f) \quad (5)$$

$$P\_i = \mathrm{PoseModel}(S_j) \quad (6)$$

$$H\_i = \mathrm{rPPGProxy}(V_f) \quad (7)$$

The three per-user feature vectors are combined through an attention-guided fusion into a single fused state E_i,t:

$$E\_i,t = \alpha\cdot F\_i + \beta\cdot P\_i + \gamma\cdot H\_i, \quad \alpha+\beta+\gamma = 1 \quad (8)$$

Where α, β, and γ are attention weights learned during training, extending the two-term fusion used in single-modality-pair mirror systems [1] with a physiological term inspired by physiologically-aware mirror prototypes [7].

- *Per-User Adaptive Baseline and Deviation Score*

Each resolved user i maintains a running mean μ_i,t and standard deviation σ_i,t of their own fused state, updated by an exponential moving average with decay rate $\rho \in (0,1)$:

$$\mu\_i,t = \rho\cdot\mu\_i,t_{-1} + (1-\rho)\cdot E\_i,t \quad (9)$$

$$\sigma^2\_i,t = \rho\cdot\sigma^2\_i,t_{-1} + (1-\rho)\cdot(E\_i,t - \mu\_i,t)^2 \quad (10)$$

A personalised deviation (z-score-like) signal is then computed relative to that user's own baseline, rather than a fixed population threshold:

$$D\_i,t = |\, E\_i,t - \mu\_i,t \,| \,/\, \sigma\_i,t \quad (11)$$

A sustained value of D_i,t above a per-user threshold τ_i (initialised from population norms and refined as more of that user's own data accumulates) triggers the coaching-and-escalation path described next; sub-threshold values feed only a routine, low-urgency wellness tip.

- *LLM Coaching, Safety Guardrail, and Escalation*

When D_i,t exceeds τ_i, the system assembles a prompt from the fused state, the personalised baseline, the resident's short profile, and the deviation score, and queries an LLM for a short coaching response:

$$R\_i = \mathrm{LLM}(\, E\_i,t \,,\, \mu\_i,t \,,\, \mathrm{profile}\_i \,,\, D\_i,t \,) \quad (12)$$

R_i is passed through a safety-guardrail check — a rule-based and classifier-based filter that rejects diagnostic language, enforces a supportive tone, and decides whether the response alone is sufficient or whether a caregiver/telehealth escalation must also fire — before anything is shown on the mirror display, consistent with the caution urged by recent elder-facing chatbot safety research [12], [13]. Only the resolved user i, and no other resident whose data may also be stored on the same device, sees R_i, mirroring the identity-

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1281

routing principle shown to prevent persona confusion in shared conversational assistants [11].

Edge devices execute Sections IV-A through IV-D so the sensing-to-deviation loop stays low-latency; the LLM call in Section IV-E and the long-term, per-user trend analytics are executed in the cloud, following the edge-cloud division used in related multimodal IoT health-monitoring work [4].

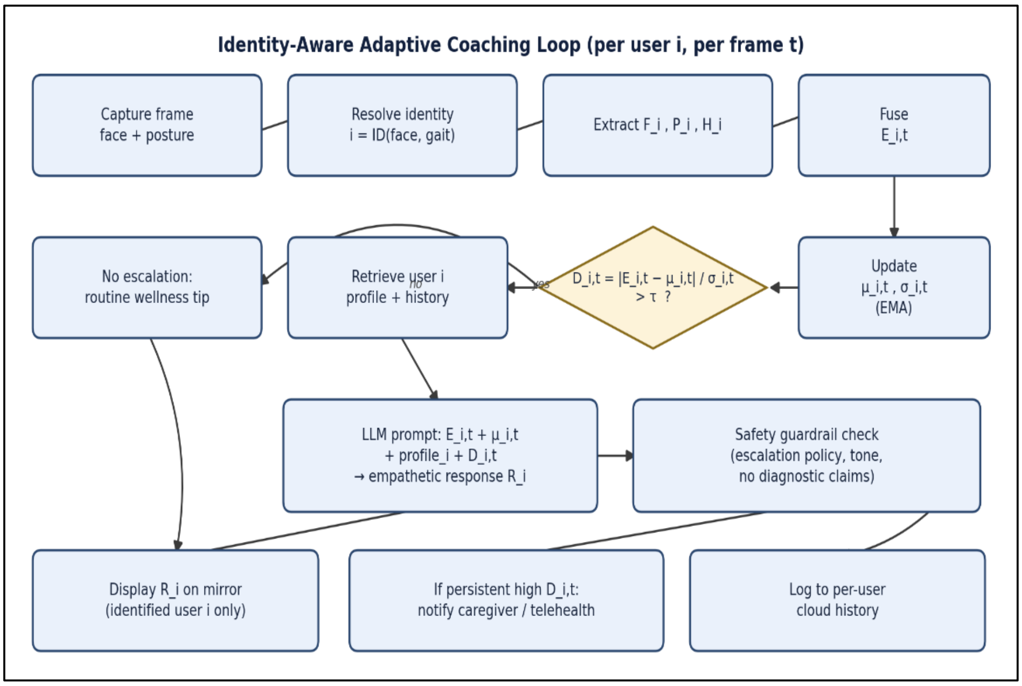

Fig 2 Identity-Aware Adaptive Coaching Loop Executed for Each Resolved User i at Each Processed Frame t.

➢ *Illustrative Simulation Study*

No hardware prototype of the proposed mirror has been built and no human-subject data has been collected at the time of writing; the figures in this section are therefore synthetic data generated purely to illustrate the qualitative behaviour the design is intended to produce, not experimental results, and no accuracy, precision, or detection-rate figures are reported anywhere in this paper.

Fig. 3 shows a synthetic 10-minute session for two simulated users, each given a different underlying noise level and baseline oscillation to represent two people with genuinely different resting expressiveness, with two artificially injected "events" for User A and one for User B superimposed on that noise. The purpose of the figure is only to show that a per-user deviation score of the form of Eq. 11 can, in principle, surface a distinct event against each user's own noisy baseline without the two users needing to share one threshold.

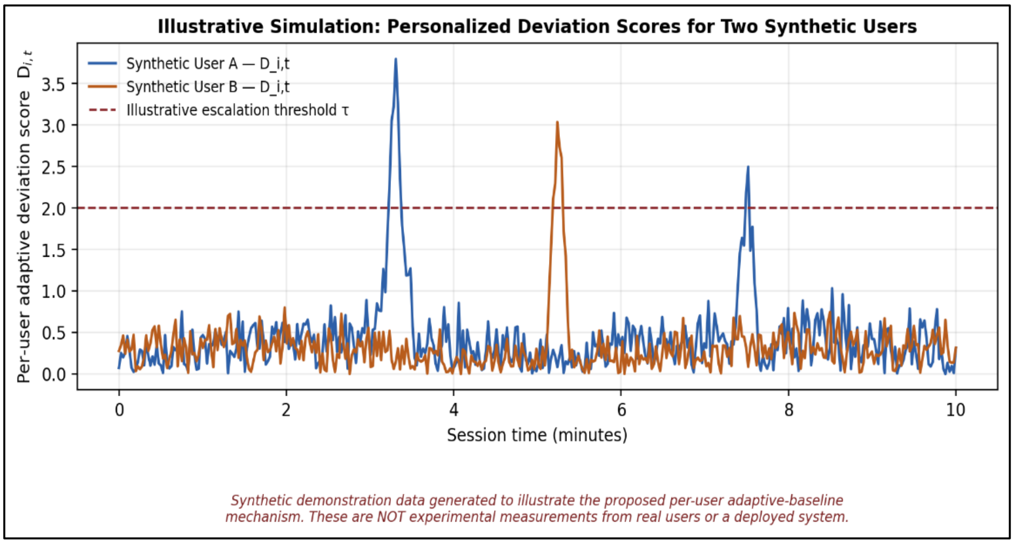

Fig 3 Illustrative Simulation of the Per-User Deviation Score $D_{i,t}$ for Two Synthetic Users (Synthetic Data; not An Experimental Measurement).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26jul1281

Fig. 4 sketches a hypothetical latency budget across the pipeline stages of Section IV, splitting each stage into an edge or cloud component. It is included only to reason qualitatively about where an LLM coaching call would sit relative to the low-latency edge loop — the numbers are assumed for illustration and were not measured on any hardware.

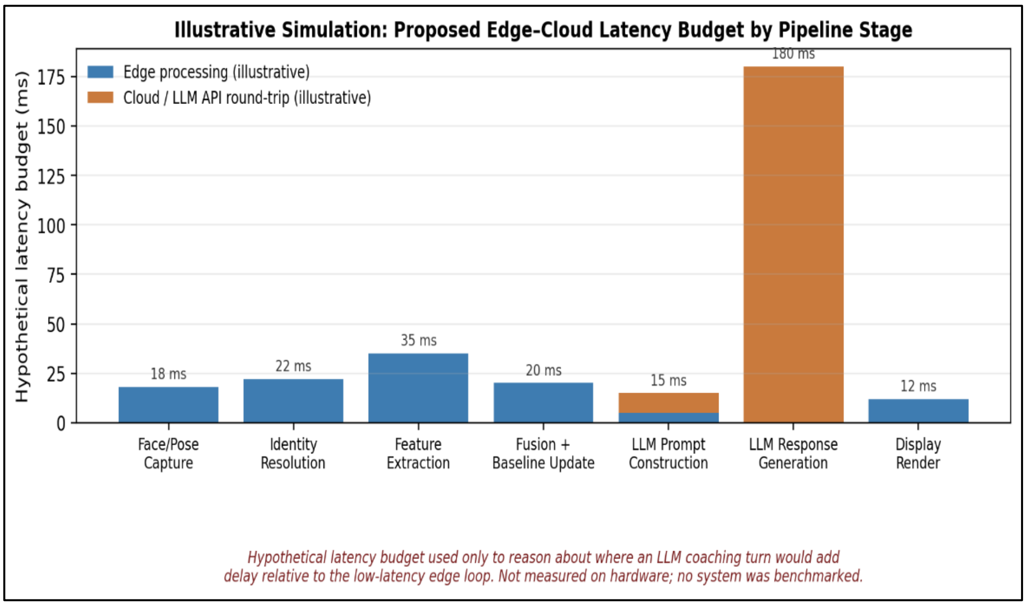

Fig 4 Illustrative, Hypothetical Edge–Cloud Latency Budget by Pipeline Stage (Assumed Values; not a Benchmark).

Because both figures are synthetic, they should be read only as a sanity check on the mathematical formulation of Section IV, not as evidence of the system's real-world accuracy, latency, or clinical usefulness. A hardware prototype and an Institutional-Review-Board-approved user study with enrolled household or care-home participants would be the appropriate next step to obtain genuine performance figures, and is discussed further in the Conclusion.

IV. DISCUSSIONS

The proposed design tries to close two gaps identified in Section II: it treats the mirror as a genuinely multi-user device by resolving identity from fused facial and gait signatures before any feature is attributed to a person, and it evaluates each resident against their own adaptive baseline rather than a single population norm, in the same spirit as identity-aware routing was shown to prevent persona confusion in shared voice assistants [11]. Feeding the resulting personalised deviation score into an LLM coaching layer, rather than a static rule-based message, follows the direction taken by recent LLM health-coaching and elder-facing communication systems [9], [10], while the explicit safety-guardrail stage responds directly to the caution raised in recent reviews and safety benchmarks for elderly-facing chatbots [12], [13].

The design also has clear limitations that a future prototype will need to confront directly. Identity resolution from face and gait alone may be less reliable for residents with very similar builds, for wheelchair users whose gait signature is unavailable, or under poor lighting; a fallback to voice or a wearable tag may be needed in practice. The per-user adaptive baseline needs enough historical data before it becomes personally meaningful, so a new resident will initially be compared against a generic population prior, which reduces but does not eliminate the cold-start problem. Finally, deploying an LLM as a coaching voice inside someone's home raises privacy, consent, and over-reliance questions that are only partly addressed by a safety-guardrail filter and will require a proper qualitative and clinical evaluation with real participants before any such system is used unsupervised in elder care.

V. CONCLUSION

This paper proposed an IoT-enabled Smart Mirror architecture that extends prior single-user, rule-based mirror designs in three directions: identity-aware sensing that resolves which of several household or care-home residents is currently at the mirror; a per-user adaptive statistical baseline that personalises what counts as an emotional or postural deviation for that specific resident; and an LLM-based coaching layer, gated by an explicit safety guardrail, that turns the resulting personalised deviation signal into short, empathetic guidance rather than a static label. The mathematical formulation of the identity-resolution, fusion, adaptive-baseline, and LLM-conditioning stages was developed in full, and a small, clearly-labelled synthetic simulation was used only to sanity-check the qualitative behaviour of the deviation signal and the edge-cloud latency budget the design implies — not to claim any measured accuracy or performance. The logical next steps are to build a hardware prototype, enrol a small panel of household or simulated care-home participants under appropriate ethical oversight, and evaluate identity-resolution accuracy, personalised-baseline calibration time, and the acceptability of the LLM coaching messages to real users and their caregivers before any claim of real-world effectiveness can be made.

REFERENCES

[1]. P. Silapasuphakornwong and K. Uehira, "Smart Mirror for Elderly Emotion Monitoring," 2021 IEEE 3rd Global Conf. on Life Sciences and Technologies (LifeTech), Nara, Japan, 2021, pp. 356–359, doi: 10.1109/LifeTech52111.2021.9391829.

[2]. S. Bianco, L. Celona, G. Ciocca, D. Marelli, P. Napoletano, S. Yu, and R. Schettini, "A Smart Mirror for Emotion Monitoring in Home Environments," Sensors, vol. 21, no. 22, p. 7453, 2021, doi: 10.3390/s21227453.

[3]. L. Rachakonda, P. Rajkumar, S. P. Mohanty, and E. Kougianos, "iMirror: A Smart Mirror for Stress Detection in the IoMT Framework for Advancements in Smart Cities," 2020 IEEE Int. Smart Cities Conf. (ISC2), Piscataway, NJ, USA, 2020, pp. 1–7, doi: 10.1109/ISC251055.2020.9239081.

[4]. C. S. T. Kedanjoth, M. Thomas, D. H. Pohren, A. D. S. Roque, and E. P. D. Freitas, "An IoT-based Multimodal AI System for Emotional and Behavioral Analysis," 2025 12th Int. Conf. on Future Internet of Things and Cloud (FiCloud), Istanbul, Turkiye, 2025, pp. 151–158, doi: 10.1109/FiCloud66139.2025.00029.

[5]. L. V, K. K, K. S, G. D. K, H. R, and M. V, "Reflective Wellness: An AI-Powered Smart Mirror for Personalized Health Monitoring and Insights," 2025 3rd Int. Conf. on Artificial Intelligence and Machine Learning Applications (AIMLA), Namakkal, India, 2025, pp. 1–5, doi: 10.1109/AIMLA63829.2025.11041129.

[6]. M. J. Santofimia, X. del Toro, C. Bolaños, J. Dorado, and S. Colantonio, "A Smart Mirror to Your Health: Personalized Virtual Coaching for Active and Healthy Ageing," in Digital Health and Informatics Innovations for Sustainable Health Care Systems, Springer, 2025, doi: 10.1007/978-3-031-84158-3_10.

[7]. M. A. Kasno and J.-W. Jung, "Feasibility of an AI-Enabled Smart Mirror Integrating MA-rPPG, Facial Affect, and Conversational Guidance in Realtime," Sensors, vol. 25, no. 18, p. 5831, 2025, doi: 10.3390/s25185831.

[8]. Y. Pu, J. Zhang, M. Li, and Y. Mou, "Design of smart elderly care emotion recognition and elderly care system based on transfer learning," Health Informatics Journal, 2025, doi: 10.1177/14727978251337893.

[9]. M. Jörke, S. Sapkota, L. Warkenthien, N. Vainio, P. Schmiedmayer, E. Brunskill, and J. A. Landay, "GPTCoach: Towards LLM-Based Physical Activity Coaching," in Proc. 2025 CHI Conf. on Human Factors in Computing Systems, ACM, New York, NY, USA, 2025, doi: 10.1145/3706598.3713819.

[10]. Z. Yang, X. Xu, B. Yao, E. Rogers, S. Zhang, S. Intille, N. Shara, G. G. Gao, and D. Wang, "Talk2Care: An LLM-based Voice Assistant for Communication between Healthcare Providers and Older Adults," Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., vol. 8, no. 2, pp. 1–35, 2024.

[11]. M. Al-Ratrout, P. Ravva, S. Sharmin, A. Raikwar, J. Y. Shin, and L. Barmaki, "AFA: Identity-Aware Memory for Preventing Persona Confusion in Multi-User Dialogue," arXiv preprint arXiv:2604.25022, 2026.

[12]. "The role of LLM-powered chatbots in assisting elderly people: systematic review," Network Modeling Analysis in Health Informatics and Bioinformatics, Springer, 2025, doi: 10.1007/s13721-025-00698-9.

[13]. "GrandGuard: Taxonomy, Benchmark, and Safeguards for Elderly-Chatbot Interaction Safety," arXiv preprint arXiv:2605.20203, 2026.