

Words to Describe What I'm Feeling: Exploring the Potential of AI Agents for High Subjectivity Decisions in Advance Care Planning

Kellie Yu Hui Sim*
kellie_sim@mymail.sutd.edu.sg
Singapore University of Technology
and Design
Singapore, Singapore

Pin Sym Foong
pinsym@nus.edu.sg
Telehealth Core
National University of Singapore
Singapore, Singapore

Chenyu Zhao
cinderella.zchenyu0912@gmail.com
Singapore University of Technology
and Design
Singapore, Singapore

Melanie Yi Ning Quek
melanieqyn@gmail.com
National University of Singapore
Singapore, Singapore

Swarangi Subodh Mehta
swarangi_mehta@sutd.edu.sg
Singapore University of Technology
and Design
Singapore, Singapore

Kenny Tsu Wei Choo
kenny_choo@sutd.edu.sg
Singapore University of Technology
and Design
Singapore, Singapore

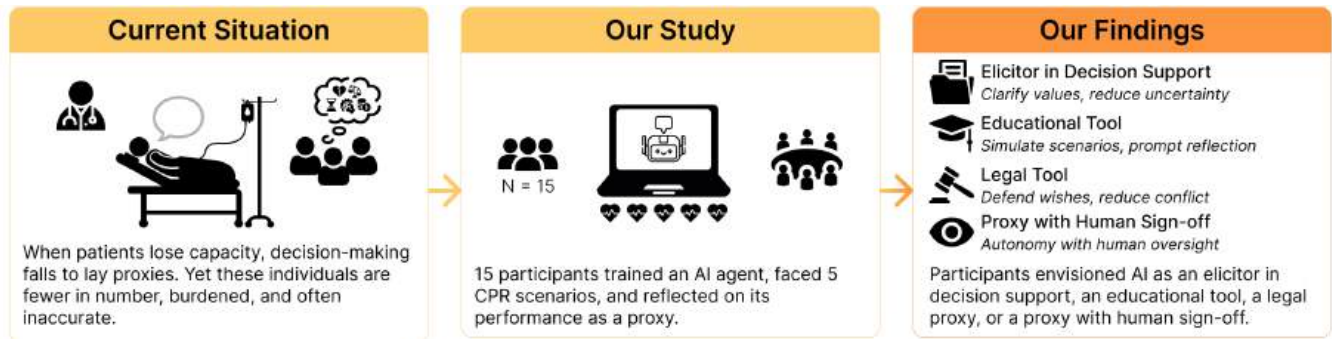


Figure 1: Summary of current challenges, study design and findings.

Abstract

Loss of decisional capacity, coupled with the increasing absence of reliable human proxies, raises urgent questions about how individuals' values can be represented in Advance Care Planning (ACP). To probe this fraught design space of high-risk, high-subjectivity decision support, we built an experience prototype (ACPAGENT) and asked 15 participants in 4 workshops to train it to be their personal ACP proxy. We analysed their coping strategies and feature requests and mapped the results onto axes of agent autonomy and human control. Our findings show a surprising 86.7% agreement with ACPAGENT, arguing for a potential new role of AI in ACP where agents act as personal advocates for individuals, building mutual intelligibility over time. We propose that the key areas of future risk that must be addressed are the moderation of users' expectations and designing accountability and oversight over agent deployment and cutoffs.

CCS Concepts

• **Human-centered computing** → Empirical studies in HCI; • **Applied computing** → Consumer health.

Keywords

Advance Care Planning, Large Language Models, Proxy Decision-Making, Advocates, Delegation

1 Introduction

Societies worldwide are rapidly ageing. Longer lifespans and declining fertility have made older adults a fast-growing segment of the population, placing new demands on healthcare, caregiving, and decision-making systems [52, 54, 56]. A central challenge arises when older adults lose decisional capacity: who should speak for them? Traditionally, family caregivers have assumed this proxy role. Yet demographic shifts have shrunk the caregiver pool, increasing the likelihood that older adults have no proxy or must rely on unfamiliar strangers, such as public guardians, for critical medical decisions [58, 72, 74].

One widely promoted way to give patients a voice before they lose capacity [41, 50] has been Advance Care Planning (ACP), in which individuals articulate their preferences for future end-of-life scenarios [2]. Yet completing an ACP is far from straightforward. The decisions involved—such as whether to undergo resuscitation—are high-risk (involving life and death) and highly subjective (rooted in personal values and priorities) [12].

Although ACP is often described in strategy and policy research as an ongoing, conversation-centred process [22], in practice it

*Corresponding author. This is a preprint of the paper accepted at CHI 2026. The final version will be available in the ACM Digital Library.

often becomes a once-off, documentation-focused event, with limited revisiting or revision over time [40]. Several process issues contribute: ACP documents are hard to retrieve and share across care settings, making iterative updating impractical. For example, clinicians frequently cannot access forms completed elsewhere, creating barriers to ongoing review [36]. Public awareness is also limited—many do not know ACP documents can be updated [48], and often revise them only during acute threats such as COVID-19 [39]. Together, these factors render ACP documents into static commitments made for unpredictable futures, raising questions about their relevance for evolving clinical circumstances, health conditions, and personal priorities.

Advances in artificial intelligence (AI) have prompted interest in their role in complex, high-risk judgements. AI systems now excel at handling large datasets and generating predictions, with applications in domains such as autonomous driving and clinical diagnostics. Yet the design of AI-enhanced decision-support tools for value-laden, high-subjectivity domains like ACP remains relatively under-explored. Why and how should such systems act? Should they offer decision support, act as autonomous proxies, or adopt intermediary roles? For HCI, tackling these questions requires understanding how users conceptualise autonomy and control in AI-enhanced systems designed for this purpose.

To explore these questions, we developed ACPAGENT, an experience prototype that simulates training an AI proxy to make resuscitation decisions on users' behalf. In four workshops with 15 participants, users first interacted with the system individually and then collectively reflected on their experiences. We found a high level of alignment: participants and ACPAGENT agreed on over 80% of resuscitation decisions. More revealing were participants' various envisionments of ACPAGENT's role, which ranged from a low-autonomy decision-support tool, to a high-autonomy proxy, with a third, intermediary framing of a personal advocate that learns preferences and speaks on their behalf, in care planning conversations.

We mapped these findings onto Shneiderman's [59] autonomy and control framing for human-centred AI. We conceptualised participants' expressed requirements for decision support in high-risk, high-subjectivity decisions as *forces* that shape desired agent roles in particular directions. From this analysis, we advanced design considerations for the most promising framing: **AI-enhanced patient advocates that combine adaptability with respect for personal values over time**. In doing so, this paper contributes:

- (1) ACPAGENT, a prototype system for exploring ACP decisions through training an AI agent.
- (2) Empirical findings on how people perceive and interact with a trainable AI proxy in a high-risk, high-subjectivity context.
- (3) Design considerations for a new role of AI as an advocate in future care planning, towards addressing current limitations in ACP practice.

By situating AI within the terrain of computationally-supported ACP decision-making, our work extends HCI research on designing AI for high-risk, value-driven domains and highlights opportunities to support ageing societies where human proxies are increasingly scarce.

2 Related Work

2.1 The Evolution of Delegation Systems

2.1.1 From Legal Documentation to Education and Conversation. Ageing nations like the United States, South Korea, and Finland [51, 54, 56] face rising longevity and declining birthrates, often accompanied by a gradual loss of functional capacity, and increasing physical and decisional support needs [18]. Instruments such as advance medical directives, ACPs and Lasting Power of Attorneys (LPAs) aim to ensure patients' preferences are respected after they lose capacity, with best practices emphasising explicit conversations among patients, proxies and healthcare professionals [22, 68].

In practice, incorporating patients' voices remains difficult due to documentation challenges [34, 40], problems retrieving and reviewing directives and applying their elicited preferences [46]. Consequently, researchers increasingly argue that ACP should foreground conversation between proxies, patients, and clinicians [3], thereby fostering deeper understanding of patient preferences and values, and resolving conflicts with proxies' values [17, 23]. These efforts may or may not produce documentation.

This evolution in how ACP is conceptualised and operationalised is likely one of the reasons for cross-national differences in delegation systems. In our study context, ACPs are not legally binding. Also, a patient may appoint a legally binding proxy using an LPA, whose powers activate only after loss of capacity. Unlike jurisdictions where these documents carry statutory enforcement of the patient's wishes—doctors in our country are **not** legally required to follow ACPs or LPA-appointed proxies' instructions when making resuscitation decisions, but are asked to consider patients' best interests in all decisions [43].

2.1.2 The Limitations of Proxy Decision-Making. Serious illness involves many unknowns. Epstein et al. [12] describes this as 'epistemic uncertainty', which creates proxy challenges as they balance risks across outcomes and contend with unstable prognoses. ACP decisions are hence often characterised as 'preference-sensitive': they hinge not on objective metrics, but on patients' preferences or proxies' interpretations. So, proxies face emotional, ethical, and interpersonal challenges—broaching death, managing guilt, navigating family conflict, and judging whether past wishes apply to the present situation [15, 17, 66, 74]. Unsurprisingly, proxies misrepresent patients' wishes in about one-third of cases [58].

Delegation processes like ACP are further strained by shrinking caregiver availability in ageing nations. Without a trusted family member, responsibility may fall to healthcare providers or institutions that lack the personal knowledge of the patient, *increasing* the risk that preferences are overlooked [14, 63, 67]. As demographic shifts continue, delegation processes may not be able to only rely on caregivers to retain deep knowledge of patients' preferences.

2.1.3 Delegation to AI. To address these challenges, recent work has explored AI support in ACP. Ario et al. [6] summarises applications into three domains: (1) extracting preferences from records; (2) predicting ACP eligibility/timing; and (3) guiding users through self-reflective conversations to articulate the patient's values and consequent care preferences. In the last category, studies such as *PreCare* [25] and Foo et al. [16] show how LLMs can scaffold value

exploration and highlight conflicts, enhancing engagement and clarity.

However, these tools remain advisory: they support reflection and conversation but were not designed to act as proxies in the absence of human proxies. We now turn to recent research on the concept of delegating decisions to AI, outlining current knowledge with respect to the scope of this paper: the delegation of authority for high-risk, high-subjectivity decisions.

2.2 Conditions for Delegation to AI

Research on delegating human decision-making to AI positions systems along a continuum between human control and AI autonomy [59]. Prior work examines factors that move people along this continuum. Candrian and Scherer [11] suggests that when facing potential losses, humans perceive AI as lacking self-interested motives and therefore less risky and more controllable, increasing willingness to delegate—especially for low-subjectivity decisions with clear external standards of accuracy.

Other research identifies conditions that further strengthen delegation preferences. Hemmer et al. [21] showed that satisfaction and confidence increase when people understand why AI assumes decision-making responsibility in collaborative settings, while Spitzer et al. [63] found that explicit contextual information about task demands and system competence boosts willingness to delegate. Interestingly, in more ‘sensitive’ contexts involving potential loss or interpersonal tension, AI may be preferred as it can act as an “emotional buffer”, avoiding the social risks associated with human-led delegation [11].

2.2.1 Nature of the Task. The seemingly contradictory preference for AI in higher-risk situations may be explained by Lai et al. [35]’s *structure of decisions* framework, which proposes that delegation design depends on the nature of the decision itself. They characterise decision tasks along three dimensions: (1) task **risk** (low to high), (2) the amount of required human domain **expertise** and whether it can be solved with existing human intuitions (intuitive versus expert knowledge needed), and (3) decision **subjectivity** (extent to which personal values versus objective criteria dominate).

We define **high-risk decisions** as those involving significant financial and/or emotional prospective losses, with high reversal costs [33], including situations where complex and ambiguous factors may lead decision-makers toward erroneous choices with potentially catastrophic consequences [55]. We define **high-subjectivity decisions** as those requiring contextual interpretation, where multiple possible outcomes may be equally valid [4].

Applying Lai et al. [35]’s framework, we find that prior delegation research has mainly examined lower-risk or lower-subjectivity tasks (e.g., consumer choice, hiring, financial risk prediction [21]). By contrast, end-of-life decision-making and ACP are high-risk, involve variable human expertise, and depend on deeply subjective values. Hence, insights from low-risk, low-subjectivity research may not transfer. For example, AI’s role as an “emotional buffer” arises in high-risk settings and is absent in lower-risk settings. Moreover, supporting high-subjectivity decisions remains challenging because outcome ‘correctness’ is best determined by subjective measure, and not by objective data alone. To our knowledge, little

empirical research has addressed the use of AI in such higher-risk decisions.

2.3 The Potential of Generative AI as Agents in High-Subjectivity Tasks

We now turn to using generative AI as agents capable of handling subjective content such as values. Large Language Models (LLMs) enable conversational systems that generate context-sensitive, personalised outputs across diverse domains such as mental health [28, 60, 62], values elicitation [16, 76], and reflection [5, 31, 47, 61]. GenAI extends these systems by maintaining memory, modelling intentions, and acting persistently over time, opening possibilities—and risks—of machines representing people’s values [45] through agentic, personal advocacy-like roles.

Recent research highlights three interrelated concerns in designing such agents: (1) *interactional style* (tone, persona, expressivity), which shapes trust and willingness to delegate [37, 49]; (2) *delegation and authority*, particularly how framing and transparency of the system affects user confidence [11, 63]; and (3) *representation and identity*, including how agents preserve or simulate aspects of the self across time [30, 45].

Non-experts are particularly swayed by affective framing and communicative style, with “future self” personas enhancing emotional resonance [30, 49]. Across expertise levels, willingness to rely on agents increases when contextual information and system competence are clearly conveyed [11, 37, 63].

Legal and ethical debates about the limits of proxy authority add nuance to these dynamics. Current proxy decision-making policy emphasises fidelity to patient wishes [52], yet human proxies often reinterpret or override preferences under stress or conflict [9]. This raises questions on the limits of agentic proxy authority, required auditing, and how AI agents should balance autonomy with contextual judgment—acting conservatively, assertively, or neutrally.

In summary, while interest in AI systems that represent or preserve human values is growing, research on how such systems could operate as proxies remains nascent. This study, therefore, examines the design of an AI agent to act as a proxy for the patient in ACP, specifically aiming to:

- (1) Explore the users’ experience of using an AI agent to make high-risk, high-subjectivity decisions.
- (2) Examine this experience for potential roles and framings of an AI agent in supporting such decisions.

3 Methodology

To examine how people engage with an AI-powered agent designed to serve as a patient proxy in ACP, we developed an LLM-powered system (ACPAGENT) that simulates training such an agent to be one’s own proxy in ACP decisions. We then conducted a series of participatory workshops with individuals who may one day rely on or interact with such a system. In these workshops, participants trained ACPAGENT and reflected on their experience. The study received Institutional Review Board approval (SUTD IRB-25-00727), and all procedures followed ethical research guidelines.

Table 1: Participant characteristics across the four workshops (WS). The table summarises demographic factors (age, gender), relevant prior experiences (delegating authority, e.g., ACP/LPA, being able to identify a trusted other, personal experience with disease), and digital orientation (use of digital health tools, AI, IT familiarity). These factors shaped how participants engaged with ACPAGENT and informed their reflections on its role in decision-making.

WS	ID	Age	Gender	Experience with Authority Delegation (60% No)	Can Identify Trusted Other (40% No)	Personal Experience with Disease	Use of Digital Health Tools (73.3% No)	Use of AI (26.7% No)	IT Familiarity Score (Out of 60)
1	P1	45	Female	No	No	Other	No	Regularly	29
1	P2	58	Female	No	Yes	Self	Yes	Occasionally	35
1	P3	48	Male	No	No	Both	No	Occasionally	28
2	P4	41	Female	Yes	No	Other	Yes	Occasionally	26
2	P5	67	Male	Yes	Yes	Self	No	Occasionally	29
2	P6	60	Female	Yes	Yes	Both	No	No	33
2	P7	51	Male	Yes	Yes	Other	Yes	Occasionally	32
3	P8	49	Male	No	No	Other	No	Regularly	36
3	P9	51	Female	Yes	Yes	Other	No	No	46
3	P10	51	Male	No	Yes	Both	No	No	41
3	P11	47	Female	No	Yes	Other	No	No	35
4	P12	62	Female	No	No	Other	Yes	Occasionally	27
4	P13	68	Male	Yes	No	Other	No	Occasionally	33
4	P14	54	Male	No	Yes	Other	No	Regularly	26
4	P15	45	Male	No	Yes	Both	No	Occasionally	34

Notes: "IT Familiarity Score" = sum of 20 technology familiarity items (lower = more frequent use; 1 = Daily use, 2 = Seldom use, 3 = Never use).

3.1 Participants

We recruited 15 participants (7 female, all aged 40 and above, $mean = 53.1$, $SD = 7.9$; see Table 1) via social media posters targeting people with personal or family experience of life-limiting illness (e.g., advanced cancer, stroke, dementia, heart disease). Each participant received ~USD23 compensation. The group was well-educated, with most holding a Bachelor's degree, and most employed full- or part-time ($n = 11$). Awareness of delegation systems was high: 13 participants (86.7%) had prior exposure to ACP, typically through hearing about it or acting as donor or proxy, and all were aware of LPA. 6 (40%) had direct experience as donors or donees in ACP/LPA contexts, and an equal number could not identify a trusted proxy (i.e., unable to name someone for LPA, responsible for acting per LPA, or another trusted appointee, such as a friend). Most reported personal experience with life-limiting illness(es), with 6 experiencing it personally, and 13 (86.7%) having close family or friends affected.

We assessed digital literacy by examining prior use of digital health tools for health-related decision-making. Use varied: 11 participants had used digital or AI tools for health decision-making, whereas 4 had not. A 20-item digital behaviour inventory adapted from Geyer [13] reflected familiarity with routine digital tasks such as online search, with AI chatbots primarily for information seeking (60% daily, 40% occasional).

3.1.1 Participant Safety. Given the sensitivity of ACP discussions, we exercised care in recruitment and facilitation. Recruitment materials noted that the study involved decision-making about life-limiting illness and end-of-life planning, ensuring informed participation.

At the start of each session, we reiterated voluntary participation, the right to pause, skip, or withdraw at any time, and that disclosing personal health conditions was optional. We monitored for discomfort, offered breaks and refreshments, and conducted structured debriefs. An information sheet with psychosocial support hotlines was also provided. While no participants chose to stop or withdraw, safety measures were consistently emphasised.

3.2 Implementation

All interactions were delivered via a custom-built web application developed in React¹ hosted on Amazon Web Services², with the underlying LLM being GPT-4o, accessed via the OpenAI API³. Each round consisted of two form-like screens: (1) preference setup; and (2) scenario-based decision-making and interaction with ACPAGENT.

3.2.1 Training Scenarios. Participants interacted with ACPAGENT across five CPR-related scenarios (see Appendix A for full scenarios):

¹<https://react.dev/>

²<https://aws.amazon.com/>

³<https://openai.com/api/>

- **Scenario 1 - Recoverable Episode:** A sudden but treatable health event (e.g., infection) where CPR would be performed by default if the heart stopped.
- **Scenario 2 - Terminal Illness Becomes Clear:** Cancer spread made the illness incurable, though still relatively well, with an estimated 1-2 years of life.
- **Scenario 3 - Severe Side Effects from Aggressive Treatment:** Bedbound after chemotherapy with limited quality of life and required assistance with daily activities.
- **Scenario 4 - Cost and Resource Burden:** Prolonged dependency leading to nursing home placement, raising concerns about financial and care burdens.
- **Scenario 5 - Limited Life Expectancy with Family Conflict:** Weeks left to live and reduced consciousness, relatives disagreeing on whether to pursue CPR, and unclear prior wishes.

Scenarios were designed to cover a wide range of diseases and prognostic severity. The first two depicted better prognoses, while the latter three involved advanced decline. Each required them to indicate whether they wanted CPR under changing health trajectories and circumstances. While other ACP decisions (e.g. tube feeding) exist, we focused on CPR because it is widely recognised, emotionally salient, and frequently included in ACP documentation as a complex yet concrete life-sustaining treatment [75].

In designing this study, we anticipated two risks: that the scenarios might not generate sufficient reflection or decisional dilemma, and that participants might overlook the scenario text entirely. To verify that the scenarios functioned as intended, we included a self-report check asking participants whether the scenario had influenced their decision.

3.2.2 Simulating Agent Training. ACPAGENT was designed to simulate a trainable ACP proxy, with participants in the role of "trainer" and ACPAGENT as "trainee". We chose to simulate the experience of training by feeding participants' inputs (from each round, as a JSON file) into subsequent prompts to act as a technology probe [26], especially since model training or fine-tuning would not result in real-time interactions. For example, if a participant rated their tolerance for cost as 'low' or disagreed with a recommendation, this information was included in the next prompt to shape ACPAGENT's reasoning.

To reduce risks of over-trust and information overload in LLM systems [16], participants made their own CPR choices *before* seeing ACPAGENT's suggestion. To support transparency in AI reasoning [20], ACPAGENT displayed its reasoning in context, offering a lower-risk space for reflecting on personal priorities [16]. Prompts instructed the model to return four labelled sections: (1) Thinking Process, (2) Recommended Decision, (3) Reasons, and (4) Binary Decision (Yes/No CPR) (see Appendix B for full prompt templates and sample outputs).

3.3 Workshop Protocol

At the start of the workshop, participants were briefed as follows: "This study aims to understand how individuals like you would make important healthcare decisions if you no longer had the capacity to do so, or if you do not have any caregiver support." They were also told: "We are interested in participants' views on AI-powered

tools, including their comfort levels, expectations, and concerns about using such tools to advocate for their needs or represent their values." Participants were asked to treat the interaction as training an AI system that could later speak on their behalf in sensitive healthcare contexts.

Workshops comprised three stages (see Figure 2). In **Stage 1 (onboarding)**, participants watched a two-minute publicly available ACP video [1] and reviewed a CPR information sheet (see Appendix C) to establish a shared baseline understanding.

At the beginning of **Stage 2 (agent-training)**, participants were told that their inputs (e.g., values, choices, reflections) would directly inform ACPAGENT's subsequent outputs, with the system considering only the most recent round of inputs rather than a cumulative history across rounds. Participants were also reminded of this across rounds where required. Participants then specified three sets of preferences (see Figure 3):

- **Lifestyle Preferences.** Participants selected meaningful daily activities taken from a national ACP resource ⁴ (e.g., independence, time with loved ones, exercise, work), with an option to add their own.
- **Goals of Care Preferences.** Participants selected from seven common ACP goals [32] (e.g., avoiding suffering, prolonging life, avoiding being a burden), reflecting frameworks used in ACP discussions.
- **Trade-Off Preferences.** Participants rated, on two 1-10 scales adapted from Malin et al. [42], how they would balance life extension against (1) pain/discomfort and (2) financial cost, each anchored by descriptive examples for low/medium/high to quantify one's trade-off between quality and cost of life, relative to a longer lifespan.

After setup, participants engaged in five rounds. In each round, participants first viewed a scenario and recorded their CPR decision and reasoning (see Figure 4). Then, participants clicked a button to reveal ACPAGENT's recommendation with a structured breakdown. Participants then answered reflection questions to indicate their agreement/disagreement and provide feedback. These responses, together with lifestyle, goals of care, and trade-off values, were used in the LLM prompt to generate ACPAGENT's recommendations for the subsequent round.

Finally, in **Stage 3 (reflection)**, participants engaged in group discussions on alignment, comfort, and trust (see Appendix D for reflection questions).

3.4 Data Collection

Before the workshop, participants completed a screener on demographics, delegation and life-limiting illness experiences, presence of a trusted other, and AI and IT familiarity (see Appendix E). During the workshop, three researchers observed (one for every one or two participants), took notes, and consolidated findings after each session. Workshops were audio-recorded, transcribed using WhisperX [8] and verified by two researchers. System inputs of lifestyle preferences, goals of care, trade-off ratings for life-extensions, CPR decisions, agreements with ACPAGENT recommendations, scenario influence counts, and free-text reflection lengths) were logged in JSON.

⁴<https://www.aic.sg/care-services/acp-resources/>

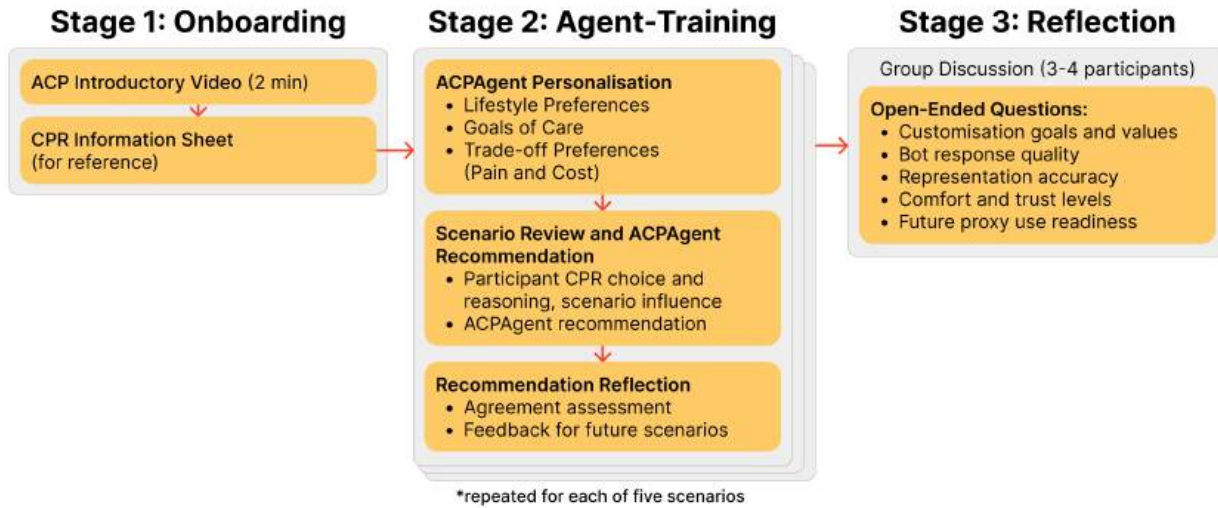


Figure 2: Three-stage workshop protocol. Stage 1 (Onboarding): Participants receive ACP education through a short video and CPR information sheet. Stage 2 (Agent-Training): Across five CPR scenarios of increasing complexity, participants make initial CPR decisions, review ACPAGENT recommendations, and provide reflections on agreement, reasoning, and areas for improvement. Stage 3 (Reflection): Groups of 3-4 participants engage in structured discussion covering customisation goals, bot response quality, representation accuracy, comfort and trust, and readiness for proxy use.

(a) Lifestyle preferences

(b) Goals of care

(c) Trade-off preferences

Figure 3: Preference setup interface at the start of Stage 2 (Agent-Training). Screens guided participants through configuring (a) lifestyle values, (b) goals of care, and (c) trade-off ratings (1-10) to initialise their personalised ACPAGENT.

(a) Scenario screen showing participant’s initial decision and whether the scenario influenced the participant’s decision.

(b) Screen showing reflection questions, presented after participants viewed ACPAGENT’s recommendation and reasoning breakdown.

Figure 4: Scenario-based interaction interface in Stage 2 (Agent-Training). Together, these steps simulated delegated decision-making while maintaining transparency of reasoning.

3.5 Data Analysis

JSON data, comprising open-ended text from the CPR decision rationales, scenario elements influencing decisions, and reasons for agreement/disagreement with ACPAGENT, was analysed for cross-scenario behaviour to track preference shifts across the five scenarios. We also analysed **Stage 3 (reflection)** responses, desired improvements to ACPAGENT, and other considerations required for decision-making. We analysed the qualitative data using an applied thematic approach [10, 19], combining thematic coding and interpretive mapping. Themes were treated as analytic constructs created by the researchers to organise patterns of shared meaning relevant to the study aims, rather than as entities assumed to pre-exist within the data.

The analysis was conducted from a contextualist epistemological stance, recognising that participants' reflections were shaped by the workshop setting, scenario structure, and group dynamics. The researchers focused on semantic meaning, staying close to participants' expressed reasoning, while acknowledging the active role of the researchers in constructing and naming themes.

Initial open coding was conducted on individual transcripts, field notes, and reflection logs by two team members trained in qualitative methods (i.e., primary author and another researcher). These researchers were involved in the study design and facilitation, which informed their sensitivity to the interactional and contextual aspects of the data. Codes were iteratively reviewed, discussed, and reorganised through regular analytic meetings, during which candidate themes were actively constructed, refined, and sometimes collapsed to better capture shared meanings across participants. We retained a larger set of analytically useful themes to reflect the diversity and nuance of participants' considerations across scenarios and discussion phases. These themes functioned as organising lenses for reporting findings.

Subsequently, all authors reviewed the results and conducted a secondary interpretive mapping to examine how participants' discussions aligned with Shneiderman [59]'s framework of autonomy and control. This mapping constituted an additional layer of interpretation, explicitly linking the results to existing conceptual frameworks.

4 Results

This section presents two sets of findings. Firstly, we assess if the experience prototype sufficiently simulated the sense of training an agent to make proxy decisions. To do so, we present participants' behavioural patterns showing how they used ACPAGENT across scenarios. Secondly, we present the thematic analysis of participants' responses to ACPAGENT. These are presented as four envisioned roles, positioned along axes of human control and agent autonomy.

4.1 Participants' Proxy Training Behaviour

All 15 participants completed five decision-making rounds with ACPAGENT, each prompted by a new scenario. To examine how actively participants engaged in training ACPAGENT, we explored three observable behaviours captured in the JSON logs:

- (1) reporting if the scenario influenced their reasoning;
- (2) changing preference settings between scenarios;
- (3) giving feedback aimed at steering ACPAGENT;

We treated preference edits and instructive feedback (items 2 and 3) as **corrective** uses of the system—explicit attempts to adjust how ACPAGENT would reason in subsequent rounds. By contrast, agreement or disagreement without subsequent edits, but accompanied by written reflections on the scenario or ACPAGENT's reasoning, was considered **reflective** use.

Engagement varied substantially, with some participants (e.g., **P6, P10, P13**) using all three options, frequently adjusting preferences, offering corrective instructions, and grounding their decisions in scenario details. Others interacted minimally (e.g., **P2, P9, P12**), making few edits and little feedback beyond responding to ACPAGENT's recommendations and reasoning, though each demonstrated at least one corrective behaviour (i.e., adjust inputs / provide additional context after a scenario). Figure 5 summarises how frequently each participant engaged in these behaviours.

4.1.1 Scenario Influence. Across the 75 decisions, participants attributed 57 (76%) to scenario details. All but **P5** and **P9** reported influence at least once, but their answers in Section 4.2 suggest that they were still meaningfully engaging with the discussion about the envisioned use of ACPAGENT. Scenarios 2, 3, and 5, involving prognosis, independence, and fluctuating consciousness, seemed particularly evocative of decisional dilemmas, producing a mix of CPR decisions (22 for, 35 against). Scenario elements, such as patient prognosis, independence, consciousness, and financial strain, were reported to have influenced the participants' decisions. These elements were considered when participants adjusted what ACPAGENT ought to prioritise in accepting or rejecting CPR.

Even participants who rarely edited their preferences still used scenario reasoning as an indirect form of training by signalling which contextual factors they felt were salient for future rounds. In the participant reflection stage, participants additionally shared that it was the narrativised, sequential building of scenario complexity that helped to deepen their reflections. As **P4** later noted during the reflections phase, unlike a multiple-choice ACP form, the scenarios revealed how “minute nuances” could shift their choice between rounds.

4.1.2 Preference Changes: Lifestyle, Goals of Care, Pain/Cost Trade-offs. Some participants used preference changes not simply to reflect shifting views, but as deliberate signals for steering ACPAGENT's reasoning. Lifestyle values were generally stable, while goals of care and pain/cost trade-offs showed more variability, especially after scenarios that introduced new dilemmas (e.g., reversibility, dependency, or affordability in Scenarios 1 and 4) in their descriptions (see Appendix F). 9 of 15 participants (60%; **P1, P4-6, P8-11, P13**) revised at least 1 preference across the 5 rounds (Figure 6). For example, in the reflection questions, **P4**, who lowered both pain and cost trade-offs from 5 to 1 after Scenario 4, explained that CPR would be a “burdensome life-prolonging measure”.

4.1.3 Feedback to Steer ACPAGENT. Towards the end of each round, participants evaluated if they agreed with ACPAGENT's recommendation, and then were asked to provide instructions to help ACPAGENT improve. Table 2 shows a summary of these varied agreements. In 65 of 75 cases (86.7%), participants evaluated that they had agreed with ACPAGENT's CPR recommendation. Among

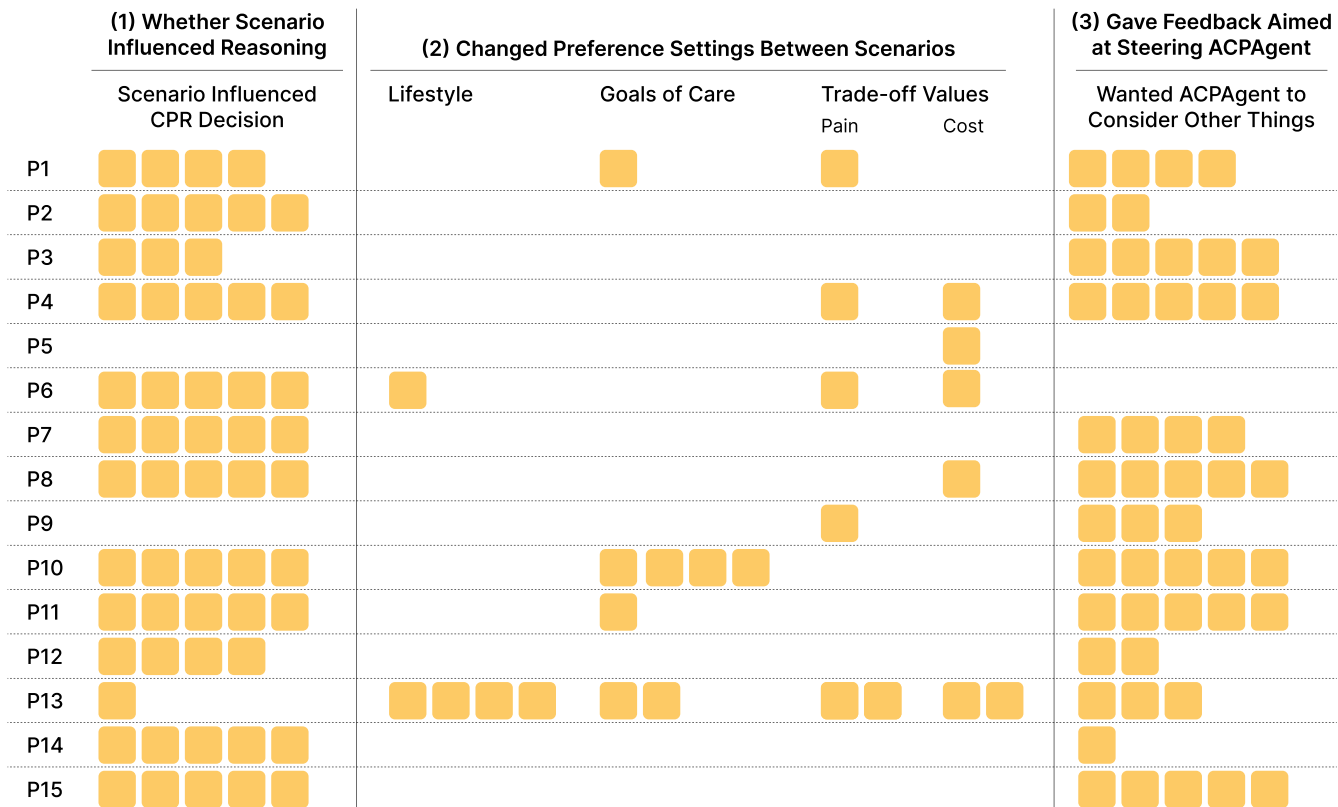


Figure 5: A summary of how frequently each of the 15 participants engaged in the 3 forms of training behaviour when interacting with ACPAGENT across the 5 scenarios: reporting whether the scenario influenced their reasoning, editing their preferences between rounds and providing steering feedback. This activity grid highlights the substantial variation in how actively participants shaped ACPAGENT across the study.




Figure 6: Preference changes by category before each round (e.g., changes before Round 2 were made after seeing Scenario 1). Lifestyle preferences remained stable, whereas Goals of Care and Pain and Cost trade-offs showed greater variability. Round 1 represents the initial baseline and is not shown.

the disagreements, (13.3%) participants indicated that this was because the ACPAGENT had overlooked important considerations, such as family emotions, financial burden, or realistic expectations

of recovery. For two participants (P5, P11), disagreement led them to revise their preferences for the next round.

In response to questions about what ACPAGENT could improve on, participants inputted additional factors for ACPAGENT to consider, such as family impact, affordability, lucidity, frailty, life stage, or systemic constraints. These elaborations suggested that participants attempted to shape not only ACPAGENT’s outputs but its underlying decision logic. These additional feedback channels were used by participants to fine-tune what ACPAGENT should emphasise.

Among the disagreements, there were 9 cases (12%), where ACPAGENT recommended the opposite of what the participant had chosen, but participants’ reflective feedback indicated that they now agreed with the ACPAGENT’s reasoning. This was an unexpected behaviour, as it suggested that ACPAGENT’s recommendations influenced participants’ positions. It seemed that the detailed reasoning could serve as a reflective mirror to participants, even when the CPR decision differed. It suggested that the influence was bi-directional, and not just human to machine.

Table 2: Participant (P) vs. ACPAGENT (A) CPR decisions by round. Cells are reported as dyads of (Participant / ACPAGENT). For example, (✓ / ✗) indicates that the participant wanted CPR, but ACPAGENT recommended No CPR. Cell marking reflects three cases: (1) Unshaded: full alignment between the participant and ACPAGENT; (2) Shaded: mismatch between the participant’s CPR choice and ACPAGENT’s recommendation, and the participant reported that ACPAGENT’s decision *did not* align with their own; (3) Shaded and underlined: mismatch in CPR decisions, but the participant nonetheless reported that ACPAGENT’s decision *did* align with their own.

Participant (P)	Round				
	R1 (P/A)	R2 (P/A)	R3 (P/A)	R4 (P/A)	R5 (P/A)
P1	✓ / ✓	✗ / ✓	✓ / ✗	✓ / ✓	✓ / ✓
P2	✗ / ✓	✗ / ✗	✗ / ✗	✗ / ✗	✗ / ✗
P3	✓ / ✓	✓ / ✓	✓ / ✗	✓ / ✓	✗ / ✓
P4	✓ / ✓	✓ / ✓	✗ / ✓	✗ / ✗	✗ / ✗
P5	✓ / ✓	✓ / ✓	✗ / ✓	✓ / ✗	✓ / ✓
P6	✗ / ✗	✗ / ✗	✗ / ✗	✗ / ✗	✗ / ✗
P7	✓ / ✓	✓ / ✓	✗ / ✓	✓ / ✗	✗ / ✓
P8	✓ / ✗	✗ / ✗	✗ / ✗	✗ / ✗	✗ / ✗
P9	✓ / ✓	✗ / ✓	✗ / ✗	✗ / ✗	✗ / ✗
P10	✓ / ✓	✓ / ✓	✗ / ✓	✓ / ✗	✗ / ✓
P11	✓ / ✓	✗ / ✓	✗ / ✗	✗ / ✗	✗ / ✗
P12	✓ / ✓	✓ / ✓	✗ / ✓	✗ / ✗	✗ / ✗
P13	✓ / ✓	✓ / ✓	✗ / ✗	✗ / ✗	✗ / ✗
P14	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
P15	✗ / ✓	✗ / ✗	✗ / ✗	✗ / ✗	✗ / ✗

✓: chose CPR; ✗: chose No CPR.

Yellow shaded: participant disagreed with ACPAGENT recommendation and chose a different CPR decision.

Yellow shaded and underlined: participant agreed with ACPAGENT recommendation despite mismatch in CPR decisions.

Unshaded cells indicate full alignment between participant and ACPAGENT.

Instructive feedback appeared in the post-round reflections, where participants specified additional factors for ACPAGENT to consider, such as family impact, affordability, lucidity, frailty, life stage, or systemic constraints. These elaborations showed how participants attempted to shape not only ACPAGENT’s outputs but its underlying decision logic.

4.1.4 Summary. We conclude that the prototype was sufficient to simulate training an agent as an ACP proxy. Participants were influenced by the scenarios and changed their answers across scenarios. In both agreement and disagreement cases, participants treated ACPAGENT as a system whose behaviour could be shaped over time, rather than a static decision-making tool. Most demonstrated corrective efforts by adjusting their preferences or giving instructive feedback to the bot after a scenario, and a small handful were less active, responding only to ACPAGENT’s recommendation and reasoning text.

4.2 Envisioned Uses of ACPAGENT

Building on these behavioural results, we now report the themes from our analysis of participants’ experiences of training a tool when making high-risk, high-subjectivity decisions.

In our analysis, we saw that participants’ requested features drew on multiple, often concurrent framings of the role of AI in these decisions. Participants often shifted their framing of ACPAGENT as they worked through various concerns, including trust, uncertainty, conflict, and the need for safeguards. We believe it is likely due to the speculative, future-oriented task for a nascent, unformed agent role.

To best present this complexity, we mapped the themes onto Shneiderman [59]’s control-autonomy framework. The two axes yielded four broad framings of what ACPAGENT could be (see Figure 7). To capture the speculative aspect of the feedback, we describe the themes as ‘forces’ that might move the design of future proxy agents between each framing. In Figure 7, the numbered points are the themes, and the arrows indicate the direction of influence.

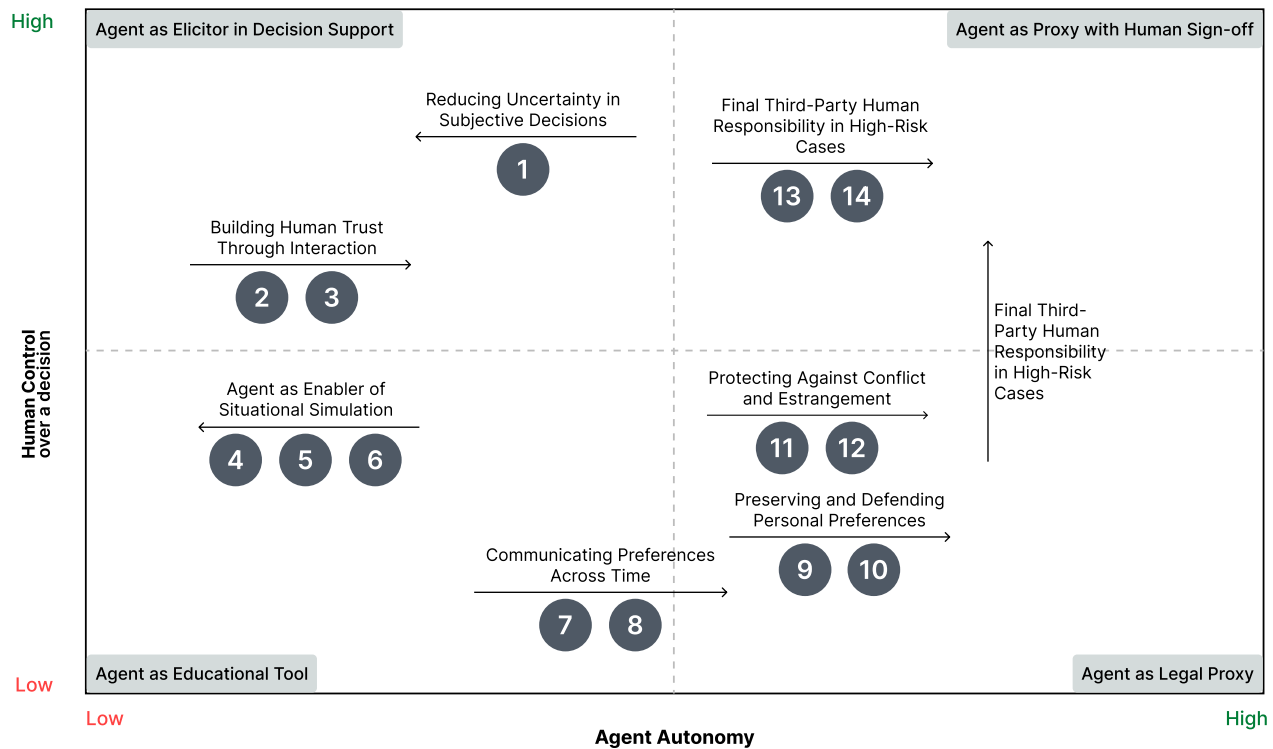


Figure 7: Quadrants of envisioned roles for ACPAGENT. The x-axis represents increasing agent autonomy, while the y-axis represents human control over the decision-making process. Arrows indicate pressures identified in participant reflections that pushed ACPAGENT toward or away from each quadrant. Circled numbers correspond to themes presented in later subsections.

At lower levels of agent autonomy, participants framed ACPAGENT as an *elicitor in decision support*, conversationally helping to clarify values and surface priorities, towards enabling the human user to make better decisions. When human control was also de-emphasised, ACPAGENT was framed as an *educational tool* that could simulate difficult scenarios, helping the human to understand how decisions arose from the connection between values and decisions. Here, its value laid in preparing humans for high-risk, high-subjectivity decisions by simulating the experience thereof.

When greater agent autonomy and relatively low human control were desired, participants described requirements for a more assertive framing for a decision proxy agent: as a *legal proxy* where it could advocate for or uphold prior wishes in the event of conflict or patient absence. Conversely, there was a framing where both agent autonomy and human control were high. Here, ACPAGENT was envisioned as a *proxy with human sign-off*, combining advocacy with continued oversight and ultimate responsibility remaining with a human.

In the sections that follow, we discuss each role together with the movements across the axes of human control and agent autonomy.

4.2.1 ACPAGENT as an Elicitor in Decision Support. In this quadrant, ACPAGENT was envisioned as an engaging way to reduce uncertainties and elicit preferences (1). By contrast, participants also noted that more agent autonomy might be acceptable from the

decision support tool if there were more trust-building interactions with the agent (2, 3).

1 **Reducing Uncertainty in Subjective Decisions: Add Decision Parameters.** Participants struggled with the inherent uncertainty of subjective, preference-sensitive decisions. To cope, they requested more detailed, personalised information. The hope seemed to be that with more certainty in the data, it would help them resolve the decisional dilemma.

Firstly, participants wanted scenarios to have more details of diseases. These were details that personally mattered to them and that they knew would change their decision. P13 requested more alternatives and details on complications, side effects, and treatment options. They noted that dementia or semi-conscious states were not represented, even though ACPAGENT generally “just followed what I want [...] never said no, it said yes.”

P3 similarly asked for very specific distinctions in scenarios:

“Like what kind of stroke? [...] bedridden can be only like you need not stay on the bed all the time even though you are bedridden [...] The scenarios may be more well-defined, more detailed.”

Participants also wanted to know what their own response to the illness might be, requesting details that would concretise the specific situation they might find themselves in. P12, who rated pain tolerance 3/10, explained:

“I would like the bot to specify how much suffering in the scenarios [...] because I don’t know how much suffering I’m supposed to be in in that scenario [...] it’s very hard to imagine [...] so the bot should describe the symptoms and the suffering so I can decide. Otherwise, I’m just going to say yes, give me CPR, because I’m assuming I return to normal [...] If I’m going to end up with a tube in ICU (Intensive Care Unit) until I die, then of course I’m going to say no. So the bot should point those things out, and what the chances are.”

Building on the above, **P12** suggested that a decision-support agent could ask probing questions such as, “What if you were suffering in that condition?”, or “Would you still want whatever [situation] your mom was in that condition when she had?”

Secondly, the request for data extended to concerns beyond the clinical situation. For example, several participants called for more explicit assessment of financial outcomes. Although the trade-off question was phrased in terms of personal financial assessment, this did not provide enough detail. **P2** suggested quantifiable thresholds (e.g., “\$500,000”) to avoid burdening her kids, while **P11** in another session called for surfacing subsidies and social assistance: “Not everyone is well-to-do, so [...] the social assistance they can get [...] that component would be good also.”

To further enhance the decision support elements, participants proposed integrating personal financial information. **P7** linked finances to inequities in pain relief provision, reflecting that affordability could shape not only access to treatment but also the degree of comfort afforded in long-term care. To further enhance these non-clinical decision support elements, participants proposed integrating personal information. (**P7**): “If you could find a way to input your own financial data or measurements for pain tolerance, then I think the results would be much better.” Some wanted very specific, personalised financial information that may not even exist in databases. For example, **P8** asked a question about finance for the family as a whole:

“Can I really afford? Can my family really afford? And do they want to do it? I’ve heard of stories where the [insurance] all drain and there’s a lot of stress.”

We found that these requests (individualised prognosis, unknown physical pain responses, localised cost information) are good-to-have for decision-making, but questions remain about how feasible obtaining such information might be. They require strong predictive models, may use data outside current health record databases, and may even be unknowable (e.g. physical pain response) [16]. However, if these decision parameters could be integrated, then a proxy agent system would be useful as a decision-support system with humans making the final decision.

② *Building Human Trust Through Interaction: Adjust Presentation to Align with Decision Style.* Participants also described conditions where enriching the interaction style and increasing the representation of decisional uncertainty would make ACPAGENT more reliable and supportive as a decision tool. Several found binary yes/no choices inadequate and wanted to more carefully bound their choice, with **P8** remarking: “The yes-no puts me in a difficult

spot because [...] I was thinking, yes, maybe with some condition. No, with some condition.”

Others called for a more conversational tone: **P10** expected “a human figure you’re talking to instead of just purely words,” while **P9** preferred phrasing that felt “like talking to you,” and **P11** suggested softer phrasing to reduce the “bot-likeness”. **P11** posited that phrasing such as “have opted” or “recommend or not recommend” felt jarring, preferring a more natural conversational flow.

The current presentation of results in text form was also an issue. **P15** proposed new ways of visualising the binary decision, suggesting percentages or graphics to avoid a “judgmental” yes/no:

“Let’s say 60% is a yes, 40% no [...] rather than just a yes or no, because it sounds like gacha⁵ yes-no, very judgmental. But graphical form would be a good representation, rather than just reading, because we get lost then.”

He described the current “huge wall of text” as off-putting, and suggested visual elements such as a status bar or pie chart.

③ *Building Human Trust Through Interaction: Provide Details of Provenance.* Desires to improve security, governance, and provenance pointed towards a future in which an agent proxy could take on a somewhat more autonomous role, provided these safeguards were in place. Participants expressed openness to relying more strongly on ACPAGENT if it were embedded in trusted infrastructures or transparent institutional arrangements. **P5** worried about data protection, cautioning that “individual[s] don’t want to input so many (sic) personal information [...] because you have no way to control to safeguard [...] the information will not be hacked one day.” **P7** similarly stressed that trust depended on governance: “It depends on the conflict of interest of who’s running the bot, who owns the bot.”

Others suggested that integration with secure infrastructures could build confidence. **P11** proposed linking ACP records to national systems so only the individual could access them: “If I’m using [it] to enter my ACP [...] only I can access that. Then I would be okay with that. Because I want to know [...] how much am I covered.”

In the group discussion, **P8** added that these integrations would help surface practical cost implications that would help him make more informed choices: “How many more days I can stay? Yeah, it’s not cheap. Even community hospital is not cheap.” To which, **P11** agreed that decisions should be “based on what I have”, while **P9** concluded that such integration would be “very practical”.

Finally, **P7**, in another session, also highlighted the need to integrate an agent proxy with existing government instruments:

“If the ACP is aligned with all the other legacy planning tools, then it might be more useful [...] whether it’s ACP, LPA [...] will, funeral plan, so on and so forth. And then moving forward, [...] assisted suicide [...] can become a composition at some point. And maybe the bot will help in those scenarios.”

However, **P4** cautioned that, regardless of provenance, there was a concern with drift and a single point of acceptance. Instead, she suggested that trust would increase with more reliable interaction,

⁵Referring to a gambling game

recounting that ACPAGENT once suggested CPR when she had indicated to refuse CPR. She emphasised that trust depended on repeated testing: “My confidence in the bot will depend on how frequently I test it, how deeply I tested and also how recently I tested it.”

4.2.2 ACPAGENT as an Educational Tool. In this quadrant, participants desired to develop the agent learning, reflection, and value clarification, shifting the framing and reducing the relevance of both human control and agent autonomy. Rather, the role of the agent should be to enable learning about high-risk, high-subjectivity decisions through rehearsing scenarios and prompting self-reflection without directly influencing real outcomes. Specifically, participants wanted to rehearse possible situations and reflect on their own values (4) or simplify otherwise complex ACP processes (5). Furthermore, because some valued ACPAGENT’s ability to guide reflection and surface their priorities, they also expanded its use, such as for facilitating difficult family conversations (6). The ability of the agent to guide learning and record their responses led the participants to further imagine that it could act as an advocate (7, 8) when it was given the role of reproducing their responses.

4 **Enabling Situational Simulation: Bidirectional Training.** Participants described the interaction with ACPAGENT as reciprocal or bidirectional: sometimes empowering them to shape ACPAGENT, other times revealing its influence on their own position.

P1 described the experience of training the agent in this way: “At the end of the whole experience, it felt more like I was training the bot on the [ACP] thing,” and emphasised the value of the training process for prompting reflection: “I think it’s good to do this exercise, to make you reflect [...] going through the exercise may take you through a different type of thing that will happen.”

Similarly, P6 indicated that ACPAGENT might help her to learn more about herself.

“When I put in my preferences [...] some of the recommendation[s] [...] really [are] what I want [...] I don’t know how to speak for myself but the bot give (sic) me words that describe what I’m feeling.”

P2’s response suggested that it was the prose-like response that helped bring disparate ideas together: “Once I gave my priority or outlook of life, it managed to capture that. So the subsequent scenario recommendation is very aligned with what I choose.”

Others highlighted how ACPAGENT was also training them to reflect on ACP. P4 explained:

“It was not so much me training the bot, I felt the bot was training me to do the ACP properly [...] The questions asked so far seem to be triggering the right thoughts about what I should tell the bot.”

P4 described referring back to the CPR notes and adjusting their preferences, viewing the prompts as “useful [...] because I get a scenario, I think, I refer, and then I make my decision, and then I get to adjust again my preferences.”

Newcomers to ACP particularly seemed to value this low-risk space for reflection on high-risk situations. P10 said:

“This is my first time hearing ACP [...] Using AI [...] to input all this is rather new to me [...] I find the

questions are quite challenging but certain scenarios really hit me.”

Their comment prompted P9 to note: “Doing ACP, setup will, LPA, all these are important [...] Only if a person goes through a situation like that, then you think of what to do.” P10 further described a recent hospitalisation that forced reflection on unfinished business, the desire to spend time with loved ones, and fears of burdening their family without an LPA: “The first thing that came to my mind was [...] what [would] happen to my family?” For him, engaging with a simulation-focused ACPAGENT could help him explore the practical and emotional ramifications of these decisions.

5 **Enabling Situational Simulation: Simplify ACP Processes.** Several contrasted this interactive reflection with their knowledge of existing ACP processes. P9 shared, “Of all things that I’ve done, I’ve not done ACP [...] when I went into the website, there were so many questions until I don’t have guidance [...] But if this chatbot is interactive enough I think it will help me a lot.”

P11 similarly found that early indications of their preferences allowed her to keep better track of what she prioritised: “I liked that in the beginning already, I could [...] take stock of what are the things. So it was simple and straight to the point.”

In contrast to this support for ACPAGENT as a ‘trainer’, P5 cautioned that without human professional facilitation, the agent risked falling short. Drawing on his knowledge of Taiwan’s model of nurse-led ACP, he argued that an AI-supported agent alone could not provide trained guidance to make informed choices and would be inadequate for guiding such complex and sensitive choices. Additionally, P14 criticised ACPAGENT for being too agreeable, explaining that if it had challenged their inputs or helped to break down complex scenarios, it could have helped him to learn more:

“Perhaps you can input it to be a bit more critical [...] It’s just very agreeable [...] this bot doesn’t help us to then break down the scenario so it just kind of follow through.”

6 **Enabling Situational Simulation: Facilitate Family Conversations.** Having found value in ACPAGENT as a tool for reflection on values, some participants began to imagine ACPAGENT facilitating conversations for someone else in their family, despite being asked to use ACPAGENT for themselves during the study. This was most evident in Workshop 3, where several participants were already caregiving for their parents, even though none had previously acted as a Nominated Healthcare Spokesperson or LPA donee. They saw potential for an agent to serve as a facilitator of difficult family conversations. P11 described how their own reflection immediately extended to their mother:

“For me, when I’m doing this [...] I think about myself, but more importantly, I also was thinking about my mother [...] I want to have this conversation with her because it’s not really easy to have this kind of conversation.”

P11 noted how language and navigation barriers often left her acting as an interpreter during medical appointments, and hoped the tool could allow their mother to document her own wishes:

“I want it from my mom’s own mind [...] I don’t want to feel like guilt or anything of having made the wrong decision when that’s not what she wants.”

Furthermore, if the system were accessible at primary or community care level, **P11** proposed it could encourage earlier and more inclusive ACP discussions, especially if caregivers were invited to walk through it together with their care recipients.

P10 noted cultural taboos around discussing death: “My parents are all in their 80s. So when they say, what is this ACP? Then they say, hey, I’m still alive. The talk of death is so sensitive.” **P9** echoed this generational divide: “I’m glad that starting our generation onwards, we are quite open in this area. But then for a lot of seniors, they’re still very reluctant to talk about this.” **P9** suggested that an interactive design might help elders overcome reluctance: “Like talking to a friend like that. Like talking to a friend but it’s a robot.”

The ability of ACPAGENT to record elicited preferences suggested to participants that it might be developed to help users communicate their preferences over time. In 7, these themes create pressures away from education alone, towards framing the agent as a proxy.

⑦ *Communicating Preferences over Time: Capture Preferences While Lucid.* Participants stressed the agent might enable them to capture preferences at their most lucid stage, before crises struck, similar to the current use of ACP. **P8**, although new to ACP, suggested, “Maybe another entry point would be when [...] something happened but they’re still sound before things turn for the worse.” **P11** emphasised the need for advance planning: “I don’t want to be in that state of mind when I’m doing this. I want to do it in advance when things are not drastic.” **P11** added that sudden accidents also made timing unpredictable: “It may not all be illness related [...] It might be something else that can happen to you and then you are put in that kind of position suddenly.”

Others grounded their reflections in medical trauma, highlighting that well-captured preferences in ACPAGENT could serve as a reminder and reinforcement of earlier preferences, reducing confusion at the most critical moments. **P15** recounted how, in his mother’s final hours, panic led him to contradict her stated wish:

“My mum [...] said she wanted a quiet death. But everything was too crazy [...] I mistold the doctor. [...] But it wasn’t the way she wanted to go.”

These findings suggest that our participants did not account for how the clinical situation may change over the course of illness and remains unpredictable. As with current ACP processes, decisions made in the present may not align with future circumstances, raising similar challenges to existing practices in terms of ensuring that current documentation remains relevant as values evolve.

Just a small handful were aware of potential clinical uncertainties: **P3** and **P5** emphasised that technology and medical treatments evolve quickly, making current “terminal” prognoses or resuscitation outcomes potentially outdated in months or years. This underscored the uncertainty of making decisions too early and the subsequent need for periodic review.

⑧ *Communicating Preferences over Time: Adapt to Life Stage.* Participants’ reflections on how age and life stage shaped their decisions raised interesting questions about how an agent might serve as repository of preferences over time. **P6** stressed that values

could shift over time, noting that financial and pain considerations might matter differently over time and that preferences should be preserved when most lucid:

“Maybe certain seasons of my life, there’s a cost concern, but maybe certain seasons [...] cost is not a concern. Or the degree of pain, it depends what kind of sickness. Maybe I got dementia, I don’t even know my pain.”

Younger parents emphasised their responsibilities to children, while older participants spoke of readiness to “let nature take its course.” **P2** contrasted the needs of younger parents (e.g., **P1**) with those whose children were grown:

“I think the age matters a lot [...] Actually for me, I’m quite happy now that I’ve crossed the [...] half of the century, right? So if really things happen, I’m quite well prepared already. My kids are all grown up and stuff. I got no [...] so-called things that I cannot let go. So to me, I kind of let nature take its course.”

Following this, **P2** later added, “the mentality is different at a different stage of life. Certain things 你会看得很开 (you will be able to see it with an open mind), certain things you will cling on very hard, like to [**P1**’s] point, her kids are still young.”

P3 suggested that life stage might matter more than clinical change:

“My friend [...] as long as she’s not terminal and she’s not unconscious she’ll want to keep on going [...] but medical science is advancing. Six months down the road or even one year new drugs may come out. [...] So this thing really has to define at which point you want to let go. Like for me, just like I said, if I’m 88, I don’t care. I just want to leave, waste all my money. That’s it.”

This discussion of change led to **P2** being uncertain on whether ACPAGENT should remain a fixed directive or evolve alongside her, expressing a desire for the system to be a “buddy” that could accompany her through life stages: “If it’s going to serve as [...] a buddy to me [...] when I’m age 60 [...] then yes, those features are important. Because as you age [...] your priority of life is different, and at the stage of your different family members, you also have different considerations.” Yet **P2** also worried about the risks of continual learning, still preferring the stability of static instruments such as LPAs or Advance Medical Directives (AMDs).

4.2.3 *ACPAGENT as a Legal Proxy.* Beyond the use of ACPAGENT for static current wishes, some participants were extending the autonomy of ACPAGENT to the point of actively influencing and, further, dictating the final decision. Here, participants expressed a desire for ACPAGENT to not just preserve their wishes, but potentially advocate for them, in order to protect against conflict or estrangement. In envisioning higher autonomy for ACPAGENT, some participants wanted it to take a more active stance, representing or “fighting for” their preferences in contested moments (⑨) and ensuring that lucidly expressed wishes (see ⑦ above) would be preserved even if family members objected (⑩). For similar reasons, participants also hoped a legal proxy could reduce the emotional burden on families, especially in situations where guilt,

conflict, or estrangement might arise (11, 12). This positioned ACPAGENT as a neutral arbiter that could protect both patients and loved ones from the psychological toll of decision-making.

9 *Preserving and Defending Personal Preferences: Take an Active Role in Decision-Making.* P12, who had experienced heart failure and chronic lung disease in others, explained that she would prefer ACPAGENT's authority over doctors who may be legally compelled to perform CPR:

“Better than no decision [...] Sometimes better than your own relatives that want you to live no matter what. [...] I don't think I would trust the doctor to decide. I would prefer the bot. The doctor is required by law to perform CPR, so definitely he will do CPR.”

10 *Preserving and Defending Personal Preferences: Reinforce Wishes.* P9 explained that lucid wishes should override family emotions:

“If I have really done this based on my personal preferences, when I am still clear-minded, then I think there shouldn't be too much emotions coming in [...] I hope the doctor will respect my wish rather than prolonging my life.”

P14 imagined ACPAGENT reasoning in their absence: “So it will be an intelligence with the database to make the wisest decision possible at that time already. No one is around, right? Or help the doctor, perhaps.”

11 / 12 *Protecting Against Conflict and Estrangement: Family Dynamics / Decisional Burden.* P1 reflected on their husband's death and her hope that ACPAGENT might relieve their children of guilt over not having tried to save her:

“In this case, I choose CPR for most of the situations because I know CPR is not very effective. So do [CPR] just to make everyone happy, just to know they did something [...] Yeah, they tried. I think the important thing is they tried.”

P1 stressed that her preference should be followed, because her principle was that the emotional effect on loved ones outweighed her own outcome: “I die doesn't matter. The [ones left behind are] the one[s] guilty.” In the group discussion, P2 suggested that ACPAGENT could act as “an independent assessor or recommender” so that difficult decisions would not fall directly on P1's children's shoulders.

4.2.4 *ACPAGENT as a Proxy with Human Sign-off.* In this quadrant, participants envisioned ACPAGENT as operating with strong autonomy, capable of representing, advocating for, or interpreting their preferences, yet always under conditions of high human control. This quadrant reflects attempts to balance trust in agent autonomy with the enduring need for human responsibility in high-risk decisions. Even when participants trusted ACPAGENT, some emphasised that responsibility for high-risk, high-subjectivity decisions should always remain with humans (13). At the same time, some wanted ACPAGENT to counteract emotionally clouded human judgement without overriding human authority (14). Together, these forces positioned ACPAGENT as a legally active proxy that could fight for patient wishes with increased autonomy, but always under human sign-off and control.

15 *Requiring Final Third-Party Human Responsibility: Use Human Witness as a Safeguard.* Some participants saw value in pairing ACPAGENT with a human witness, mirroring existing ACP processes. P4 explained: “I do trust it because I trained it [...] But I also want somebody to [...] see that I've done this process. It's just so cold when it's just me and the bot.”

It was important for a human third party (i.e., clinicians, family members) to be involved in the decision. P3, who agreed with ACPAGENT in three out of five scenarios during the agent-training stage, explained: “As of now, I can only trust the doctor. The only subject matter expert.”

P8 emphasised that even when ACP records are in place, families would still want to be consulted for closure: “If we have done ACP beforehand, then the doctor [...] No, they will still ask. Because your family member might [...] want to have a proper closure.” P7 similarly believed decisions about life support could never be outsourced: “It's a very painful decision [...] so the bot cannot make it,” despite agreement with ACPAGENT in four of the scenarios.

P7, P8, and P11 stressed that the agent lacked empathy, with P7 suggesting that families might think it would be better to seek advice from healthcare professionals than from a “robot”.

14 *Requiring Final Third-Party Human Responsibility: Support through Neutrality.* In contrast, P9 and P14 saw neutrality as a safeguard against emotionally clouded human judgments. P14 argued:

“The AI should be unbiased, should remove the emotion part. That means you cannot put your personal experience into it, only because there can be different possibility. So there's no right and wrong in such situation[s], so the AI give [their suggestion] and then [let] the person decide.”

P11, despite noting that end-of-life decisions were emotionally laden, emphasised that it was also important to balance objective and non-emotional decisions since ACP was a “heavy topic”. She further shared that the simplicity of ACPAGENT and the agent-training stage “was very factual [and] straight to the point.”

5 Discussion

This study revealed both opportunities and limits of introducing an AI proxy into ACP, not only raising questions of trust, delegation, and authorship, but also the broader problem of how to design for such inherently high-risk, high-subjectivity medical decisions. We discuss our findings as such: (1) what we learnt about using AI in such contexts; (2) viability of the envisioned roles; and (3) design implications for future ACP systems.

5.1 What We Learned About Using AI in High-Risk, High-Subjectivity Decisions

Our findings suggest that the highly subjective nature of ACP decisions may lead to consistent disappointment in proxy agents, particularly when framed as decision-support tools. We found participants were somewhat aware of the subjective nature of the decisions, resisting recommendations when presented as dichotomous yes-no answers.

They suggested an *interactional style* that allowed them to incorporate the elements of uncertainty and subjectivity that they

deemed relevant, including personal conditions, varying tolerance for suffering, and contextual factors like prognosis, cost, and family circumstances. Furthermore, each decisional element mattered differently to each participant, underscoring how personal these decisions are. Participants also valued proxy agents' conversational tone, reflective prose, and bi-directional training, which supported reflection about these decisions. Some even felt that they were being taught how to make decisions and how those decisions might map to their stated preferences. This educational value extended to supporting fraught conversations with others.

Yet high subjectivity decisions also created unique challenges, especially when decisions are subjective to the point of being preference-sensitive. Alongside more feasible requests, such as details of the disease, participants occasionally requested additional details that were not feasible, including personalised forecasts of pain, the cost of care to the patient and their family, and how long the savings might last. Such information falls under the umbrella of what Epstein et al. [12] calls 'epistemic uncertainty' in end-of-life care decisions, which includes data such as prognoses and unpredictable changes in the patient's clinical situation. These data are not only hard to obtain but may also be unknowable. Our findings suggest that, unlike human-only ACP processes, participants seemed to expect that technology, computational tools and 'AI' might be able to help reduce the challenge of subjective decisions by supplying this type of data.

Thus, particularly when framed as a decision-support tool, AI-enhanced decision support must **moderate expectations for objective data or be accompanied by more education about the nature of highly subjective decisions**. In planning for future care decisions, the need for trade-offs in serious illness means that decisions often feel unsatisfactory. For example, decisions might require trade-offs between quality of life and time gained. Nonetheless, some participants approached the scenarios as if there were a hidden "right" choice, and when none was forthcoming, wanted more detail on prognosis, suffering, and cost. The nexus of high risk, high subjectivity and lack of awareness of the need for trade-offs suggests that far more education and awareness must precede the use of an AI agent for future care planning. Additionally, just as with patient decision aids [64], they should **permit and encourage different weightings** of the decisional elements.

Given the high-risk nature of the decisions, our findings suggest that **trust-building remains central to AI system design**. As with previous research in trust [7], our study also showed the influence of interactional style, provenance, and system reliability over time. However, our study provides evidence that there is interest in utilising proxy agents to counter this high-risk nature by using them as safe spaces to explore difficult decisions while having the legal proxy make decisions that protect loved ones from bearing guilt or regret. Additionally, ACPAGENT was perceived **as a neutral party, having no stakes in any outcome**, thereby less likely to bias outcomes in any particular direction.

This echoes Candrian and Scherer [11], who showed evidence of an increased willingness to delegate to machines when decisions might involve personal loss and is consistent with Steffell et al. [65], who showed that people are more likely to delegate choices to others, especially when negative consequences exist. This may be an effort to avoid blame or feeling responsible for such

decisions, thereby raising risks of using agents as legal proxies, such as moral de-skilling over time [71] and questions of ethical agency. Furthermore, greater agent autonomy might also reinforce negative outcomes of agents that might be manipulated during training towards certain decisions, particularly when the source of manipulation is unclear. Risks of manipulation underscore the need to audit and safeguard the training process to prevent malicious intent.

5.2 Feasibility of the Framings of ACPAGENT

Given the discussion above, we now consider the feasibility of each framing that we observed. Participants most often positioned ACPAGENT as a decision-support tool or a proxy (see Figure 8). Here, however, we interpret these results through the lens of feasibility, represented by the orange and green boxes in the figure. The lens of feasibility considers what is plausible, problematic, or requires systemic change.

5.2.1 Framings Requiring High Human Control. Framings with high human control, namely decision-support tools and proxies with human sign-off, still rely on the presence of human caregivers or professionals to witness, validate, or interpret decisions. Although human oversight remains desirable in high-risk decisions, demographic shifts in ageing nations may require exploring solutions that do not solely depend on the presence of a trusted other, such as a close family member or a healthcare worker.

5.2.2 Framings Dependent on Regulatory Frameworks. Imagining ACPAGENT as a legal proxy or one with human sign-off surfaces issues of law and governance. Such roles are feasible only if regulatory frameworks evolve to recognise agents as legally authorised proxies, contingent on government adoption and professional consensus.

5.2.3 Advocacy as a Middle Ground. Among the framings, *advocacy* appeared the most feasible and desirable option in the medium term of 5-10 years. **This advocacy framing adds a novel role for AI in future care planning, moving beyond previously identified uses for passive extraction and prediction [6] to a more active agentic role.** Situated between tool and proxy, it positions ACPAGENT as a bounded representative of patient wishes: able to amplify and reinforce elicited values, yet still variably complementing human oversight rather than displacing it. For agentic advocacy to be more feasible, however, our findings suggest that three issues must be addressed:

- (1) Educating users to challenge the false belief that preferences remain stable [77].
- (2) Clarifying when agent advocacy should remain *static (replaying preferences without adaptation)* or become *active (arguing on behalf of the patient, with the possibility of evolving alongside them or serving as an emotional buffer)* [11].
- (3) Defining the degree of authority such advocacy should hold relative to families and clinicians in shared decision-making, an issue closely tied to the parameters of *provenance (who made it?)* and *deployment timeline (when does it start and end?)* in generative ghost design [45, 57, 67].

Some participants pushed this vision further, imagining ACPAGENT as a long-term companion that could "grow over time" with

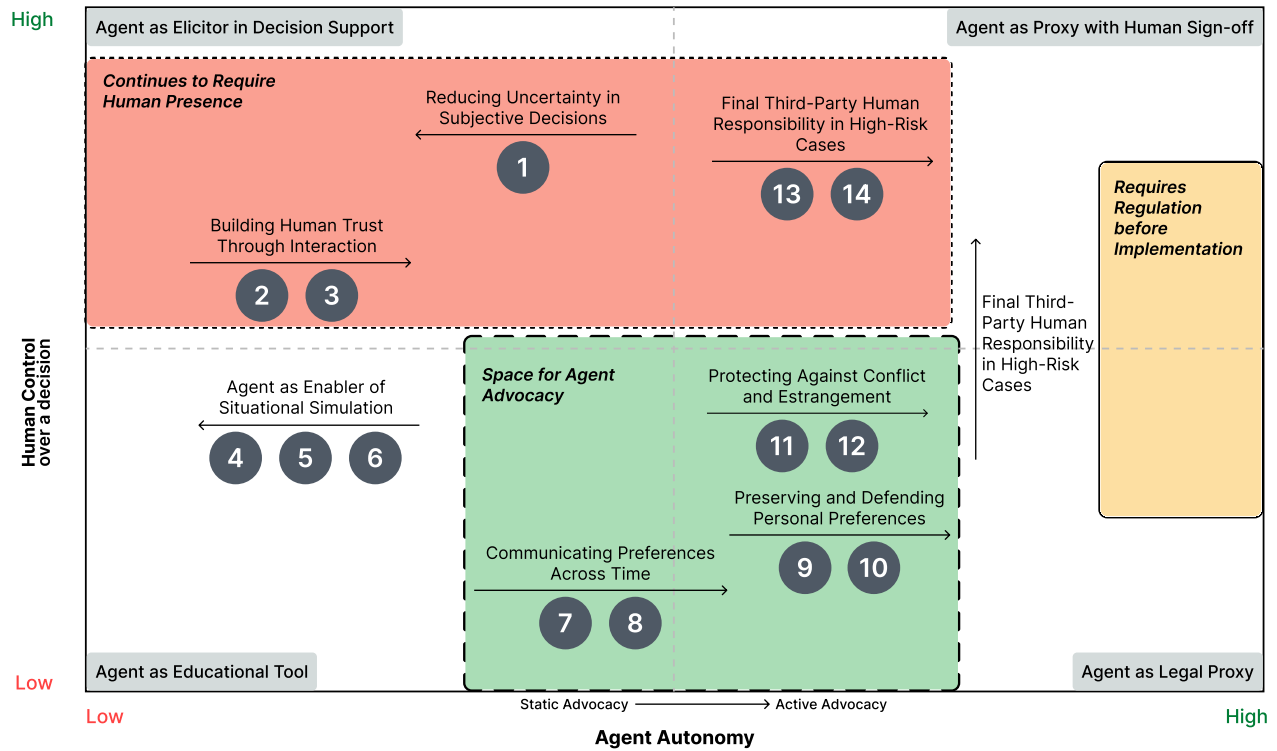


Figure 8: Framings of ACPAGENT situated within a two-dimensional space. The x-axis represents agent autonomy, and the y-axis represents human control over the decision-making process. Quadrants capture distinct framings observed in our study, while the middle shaded region highlights advocacy as an emergent framing. The other two shaded regions (pink and yellow) refer to areas that face systemic constraints: those requiring high human control (limited by caregiver availability) and those dependent on regulatory frameworks (plausible only with government authorisation). Circled numbers correspond to themes discussed in the earlier Results sections.

its user (e.g., as with P2’s “buddy” description.) Yet even here, participants worried about the risks of continual adaptation, favouring the known stability of static instruments such as LPAs or AMDs. Further research is needed to investigate how a system should account for changes associated with its trainer’s age and health status. This tension underscored competing desires for flexibility versus certainty in ACP, echoing broader work on deliberative exchanges with AI as both externalising reflection and prompting self-correction [27, 38].

5.3 Addressing Limitations of Current ACP and Design Considerations for Future Care Planning Tools

Building on the above, we further consider whether this advocate agent framing can meaningfully address limitations in ACP and what this suggests for future care planning tools more broadly. As we outlined earlier, existing practices face three widely recognised constraints. First, ACP decisions are infrequently revisited, despite long-standing calls to treat ACP as a conversational process that evolves with changing circumstances [22, 46, 68]. Second, ACP materials are frequently difficult to retrieve at the point of care, creating gaps between documented preferences and real-time

decision-making [34, 40]. Third, many individuals struggle with linking broad values to specific medical decisions, especially in high-risk and high-subjectivity contexts such as CPR [12, 46].

To address gaps in translating personal values into concrete treatment decisions, scenario-based reasoning and more explicit rationales can help users see how their priorities translate into specific outcomes. With 77.3% of CPR decisions in our study being shaped by such details, narrative framing could provide a meaningful scaffold for reflection, rather than a hidden mechanism of influence. This aligns with work showing that vignettes and contextual prompts help people situate abstract values within specific circumstances [30, 44]. The interface prompts and reasoning breakdown helped participants think through their decisions in ways that produced clarity and trust. An advocate role for the proxy agent could be especially valuable when it helps rehearse, followed by replaying the reasoning that underpins a decision, which is especially useful when decisions are later challenged or revisited in future scenarios.

The four envisioned uses of ACPAGENT further suggest a trajectory for such proxy agent systems. Initially, an agent might act as an elicitor of values, and later become an educational partner with repeated interaction as decisions become increasingly

preference-sensitive. Only later, if deemed acceptable, it might act as an advocate, replaying a user's earlier reasoning or standing in for them when trusted proxies are unavailable. Supporting varied decision styles (i.e., fast, intuitive judgements and slower, more deliberative reasoning [29]) will require layered outputs, such as quick summaries for overwhelming high-stakes content, and more detailed explanations for those who prefer to think with the system. Such flexibility not only maintains engagement but also offers other stakeholders, such as family and physicians, an accessible account of how patient values were interpreted. This trajectory could reflect a shift from static documentation toward a more dynamic process of future care planning, helping individuals establish a solid foundation for interpretation long before end-of-life decisions arise.

For such advocacy to be credible, future systems must establish clear boundaries on what can be interpreted or updated, and should not infer more than intended. Our findings suggest three practical directions: (1) making explicit which inputs can trigger decisional updates and who is authorised to supply them; (2) communicating what the agent *cannot* know, resisting pressure to act as a "fortune teller" of future medical advances [27]; and (3) logging rationales for major changes so that proxies and clinicians can create queryable, auditable recommendations.

Finally, the findings suggest how such agentic tools might complement existing planning infrastructures. Rather than replacing instruments like LPAs or AMDs, participants saw value in a **comprehensive system that could support reasoning about financial and clinical choices, consolidate data, and help proxies or clinicians understand how a person reached their decision(s)**. This complements calls for integrating reflective, conversational, and iterative practices into the wider care planning ecosystem [24, 30].

Taken together, these insights point **towards a novel, more dynamic model of AI involvement in future care planning** that evolves with individuals over time, supports the ongoing linking of values to decisions, and builds a mutual history of reasoning long before end-of-life crises arise, helping people articulate and revisit what matters to them in ways that current relatively static documentation alone cannot achieve.

6 Limitations and Future Work

Our study has several limitations. First, our use of clear, simplified scenarios—chosen to support accessibility for participants new to ACP, in line with Rooper et al. [53]—may have restricted interpretive diversity and steered responses (e.g., strong opposition to CPR in scenarios with clearly poor outcomes). Single-session workshops and simulated cases also limited ecological realism. Minimal preference change may reflect unconstrained settings for lifestyle and care-goal choices, and presenting ACPAGENT as personalised may have heightened trust. Future work should use more ambiguous cases, test longitudinal and clinical deployments, compare framing and facilitation contexts, and examine cultural, familial, and religious influences on expectations of AI in ACP.

Second, interactional and cognitive constraints shaped findings. ACPAGENT's single-turn, form-like interface prevented users from probing or revising its reasoning. While participants could update their preferences between rounds, the interface did not expose or

allow direct editing of ACPAGENT's internal rationale or reasoning steps. This limited corrective engagement and favoured reflective use. Cognitive biases—including confirmation [73], availability [69], and anchoring [70]—may also have influenced interpretations. Future systems should incorporate multi-turn dialogue, reasoning transparency, and bias-mitigation features to refine both interaction and interpretation.

Third, the broader design space of AI in ACP remains under-explored. While our study explored several parameters (e.g., provenance, cut-off date, deployment timeline) part of the *generative ghosts* framework [45], other parameters warrant further investigation. In parallel, legal and institutional structures must evolve to accommodate AI-enabled advocacy and proxy roles in broader governance structures. More broadly, research on decision delegation should clarify where the advocacy role of AI is best positioned along the delegation spectrum. Future work should explore advocacy not only as a design space between tool and proxy, but also as a potential trajectory (i.e., systems that age with us) while ensuring safeguards against uncontrolled drift, overreach, or misalignment.

7 Conclusion

This study explored the potential of a proxy agent (ACPAGENT) in ACP, particularly for contexts where caregiver support is increasingly absent. Across 5 resuscitation scenarios, 15 participants trained and reflected on ACPAGENT, revealing both its promise and its limits. Scenario framing influenced participant decisions in 76% of cases, and participants agreed with ACPAGENT in most cases (86.7%), often reassured by its reasoning even when recommendations diverged. Disagreements tended to include relational, emotional, and financial factors, underscoring the continuing need for human judgement in high-risk, high-subjectivity decisions. Using a framework spanning human control and agent autonomy, we described four possible roles for such agents: as an elicitor in decision support, as an educational tool, as a legal proxy, and as a proxy with human sign-off, each with unique risks and limitations. Among these, an *advocacy* framing emerged as a viable middle ground: a system that can amplify and reinforce patient values, complementing but not displacing human oversight. The findings suggest that acting as an advocate, the proxy can scaffold reflection and documentation if designed with attention to preserving dialogue, a commitment to transparency, and with sensitivity to cultural and relational dimensions of care. Our work contributes to promising design directions for responsible patient advocacy in high-risk, high-subjectivity decision-making contexts.

Acknowledgments

We thank Natasha Ureyang and Huynh Vinh Anh for their valuable inputs into the design of the study and the development of the scenarios. This work was partially supported by the Singapore Ministry of Health's National Medical Research Council under grant numbers NMRC/CG1/009/2022-NUH and CareEco21-0030. We are deeply grateful to our participants for sharing their time, experiences, and candid reflections with us.

References

- [1] AIC Singapore (Ed.). 2022. *Advance Care Planning (ACP) Awareness - Outpatient (English Subtitles)*. <https://www.youtube.com/watch?v=h8PCO3yvEi8>

- [2] Rajiv Agarwal and Andrew S. Epstein. 2018. Advance Care Planning and End-of-Life Decision Making for Patients with Cancer. *Seminars in oncology nursing* 34, 3 (Aug. 2018), 316–326. doi:10.1016/j.soncn.2018.06.012
- [3] Annette Aldous, Emily Klawson, Susan Loughlin, and Dale Lupu. 2020. Measuring Engagement in Advance Care Planning: Comparison of a Validated Survey with Chart Review in Kidney Patients (S703). *Journal of Pain and Symptom Management* 59, 2 (Feb. 2020), 530. doi:10.1016/j.jpainsymman.2019.12.265
- [4] Paula Akemi Aoyagui, Kelsey Stemmler, Sharon A Ferguson, Young-Ho Kim, and Anastasia Kuzminykh. 2025. A Matter of Perspective(s): Contrasting Human and LLM Argumentation in Subjective Decision-Making on Subtle Sexism. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. doi:10.1145/3706598.3713248
- [5] Riku Arakawa and Hiromu Yakura. 2024. Coaching Copilot: Blended Form of an LLM-Powered Chatbot and a Human Coach to Effectively Support Self-Reflection for Leadership Growth. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (CUI '24)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3640794.3665549
- [6] U. Arioz, M. J. Allsop, W. D. Goodman, S. Timmons, K. Simbirtseva, I. Mlkar, and G. Mocnik. 2025. Artificial Intelligence-based Approaches for Advance Care Planning: A Scoping Review. *BMC Palliative Care* (June 2025).
- [7] Maryam Ashoori and Justin D. Weisz. 2019. In AI We Trust? Factors That Influence Trustworthiness of AI-infused Decision-Making Processes. arXiv:1912.02675 [cs] doi:10.48550/arXiv.1912.02675
- [8] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. 2023. WhisperX: Time-Accurate Speech Transcription of Long-Form Audio. *INTERSPEECH 2023* (2023).
- [9] Julia I. Bandini, Danielle Schlang, Hyosin Kim, Melissa Bradley, Rebecca Anhang Price, Jennifer N. Bunker, and Joan M. Teno. 2021. “If We Turned Our Backs, They Would Ignore Our Wishes”: Bereaved Family Perceptions of Concordance of Care at the End of Life. *Journal of Palliative Medicine* 24, 11 (Nov. 2021), 1667–1672. doi:10.1089/jpm.2020.0714
- [10] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. doi:10.1191/1478088706qp0630a
- [11] Cindy Candrian and Anne Scherer. 2022. Rise of the Machines: Delegating Decisions to Autonomous AI. *Computers in Human Behavior* 134 (Sept. 2022), 107308. doi:10.1016/j.chb.2022.107308
- [12] Andrew S. Epstein, Eileen M. O’Reilly, Elyse Shuk, William Breitbart, Maheen A. Shah, Michele Ly, Rachel Taylor, and Angelo E. Volandes. 2017. Development of an Advance Care Planning Paradigm for Advanced Cancer: Person-Centered Oncologic Care and Choices (P-COCC). *Psycho-oncology* 26, 6 (June 2017), 866–869. doi:10.1002/pon.4167
- [13] Becky L. Faett, David M. Brienza, Mary Jo Geyer, and Leslie A. Hoffman. 2013. Teaching Self-Management Skills in Persons with Chronic Lower Limb Swelling and Limited Mobility: Evidence for Usability of Telerehabilitation. *International Journal of Telerehabilitation* 5, 1 (June 2013), 17–26. doi:10.5195/ijt.2013.6114
- [14] Robert L. Fine, Zhiyong Yang, Christy Spivey, Bonnie Boardman, and Maureen Courtney. 2016. Early Experience with Digital Advance Care Planning and Directives, a Novel Consumer-Driven Program. *Proceedings (Baylor University Medical Center)* 29, 3 (July 2016), 263–267. doi:10.1080/08998280.2016.11929430
- [15] Nienke Fleuren, Marja F. I. A. Depla, Daisy J. A. Janssen, Martijn Huisman, and Cees M. P. M. Hertogh. 2020. Underlying Goals of Advance Care Planning (ACP): A Qualitative Analysis of the Literature. *BMC Palliative Care* 19, 1 (March 2020), 27. doi:10.1186/s12904-020-0535-1
- [16] Charisse Foo, Pin Sym Foong, Camille Nadal, Natasha Ureyang, Thant Naylin, and Gerald Choon Huat Koh. 2025. The Benefits and Risks of LLMs for Facilitating Medical Decision-Making Among Laypersons. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. ACM, Madeira Portugal, 3173–3191. doi:10.1145/3715336.3735779
- [17] Pin Sym Foong, Natasha Ureyang, Charisse Foo, Sajeban Antonyrex, and Gerald Ch Koh. 2024. Designing for Caregiver-facing Values Elicitation Tools. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–20. doi:10.1145/3613904.3642214
- [18] Pin Sym Foong, Camellia Zakaria, Pavithren Pakianathan, Andrew Ian-Hong Phua, and Gerald Ch Koh. 2025. The Prevalence and Predictors of Digital Proxy Behavior in the United States: Cross-Sectional Survey Study. *Journal of Medical Internet Research* 27 (Aug. 2025), e69806–e69806. doi:10.2196/69806
- [19] Greg Guest, Kathleen M. MacQueen, and Emily E. Namey. 2011. *Applied Thematic Analysis*. sage publications.
- [20] Yuexing Hao, Zeyu Liu, Robert N. Riter, and Saleh Kalantari. 2024. Advancing Patient-Centered Shared Decision-Making with AI Systems for Older Adult Cancer Patients. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu HI USA, 2024-05-11)*. ACM, 1–20. doi:10.1145/3613904.3642353
- [21] Patrick Hemmer, Monika Westphal, Max Schemmer, Sebastian Vetter, Michael Vössing, and Gerhard Satzger. 2023. Human-AI Collaboration: The Effect of AI Delegation on Human Task Performance and Task Satisfaction. (2023). doi:10.48550/ARXIV.2303.09224
- [22] Susan E. Hickman, Hillary D. Lum, Anne M. Walling, April Savoy, and Rebecca L. Sudore. 2023. The Care Planning Umbrella: The Evolution of Advance Care Planning. *Journal of the American Geriatrics Society* 71, 7 (2023), 2350–2356. doi:10.1111/jgs.18287
- [23] Michelle Howard, Aaron J Bonham, Daren K Heyland, Rebecca Sudore, Konrad Fassbender, Carole A Robinson, Michael McKenzie, Dawn Elston, and John J You. 2016. Measuring Engagement in Advance Care Planning: A Cross-Sectional Multicentre Feasibility Study. *BMJ Open* 6, 6 (June 2016), e010375. doi:10.1136/bmjopen-2015-010375
- [24] Jong-Jing Hsu, Manasi Swaminathan, Weslie Khoo, Kyrie Jig Amon, Hiroki Sato, Sathvika Dobbala, Kate Tsui, David Crandall, and Selma Sabanovic. 2025. Bittersweet Snapshots of Life: Designing to Address Complex Emotions in a Reminiscence Interaction between Older Adults and a Robot. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2025-04-25) (CHI '25). Association for Computing Machinery, 1–18. doi:10.1145/3706598.3714256
- [25] Yu Lun Hsu, Yun-Rung Chou, Chiao-Ju Chang, Yu-Cheng Chang, Zer-Wei Lee, Rokas Gipiškis, Rachel Li, Chih-Yuan Shih, Jen-Kuei Peng, Hsien-Liang Huang, Jaw-Shiun Tsai, and Mike Y. Chen. 2025. PreCare: Designing AI Assistants for Advance Care Planning (ACP) to Enhance Personal Value Exploration, Patient Knowledge, and Decisional Confidence. arXiv:2505.09115 [cs] doi:10.48550/arXiv.2505.09115
- [26] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Ft. Lauderdale Florida USA, 17–24. doi:10.1145/642611.642616
- [27] Hayeon Jeon, Suhwoo Yoon, Keyeun Lee, Seo Hyeong Kim, Esther Hehsun Kim, Seonghye Cho, Yena Ko, Soeun Yang, Laura Dabbish, John Zimmerman, Eun-mee Kim, and Hajin Lim. 2025. Letters from Future Self: Augmenting the Letter-Exchange Exercise with LLM-based Agents to Enhance Young Adults’ Career Exploration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–21. doi:10.1145/3706598.3714206
- [28] Seungwan Jin, Bogoan Kim, and Kyungsik Han. 2025. “I Don’t Know Why I Should Use This App”: Holistic Analysis on User Engagement Challenges in Mobile Mental Health. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–23. doi:10.1145/3706598.3713732
- [29] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, NY, US. 499 pages.
- [30] Rucha Khot, Teis Arets, Joel Wester, Franziska Burger, Niels van Berkel, Rens Brankaert, Wijnand IJsselsteijn, and Minha Lee. 2025. Challenging Futures: Using Chatbots to Reflect on Aging and Dementia. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3706598.3713727
- [31] Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-Woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho Kim. 2024. MindfulDiary: Harnessing Large Language Model to Support Psychiatric Patients’ Journaling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3613904.3642937
- [32] Gerald Koh. 2025. *Navigating Advanced Illness Goals And Treatment With Digital Engagement (NAVIGATE): A Randomised Control Trial*. Clinical Trial Registration NCT07134881. clinicaltrials.gov.
- [33] Howard Kunreuther, Robert Meyer, Richard Zeckhauser, Paul Slovic, Barry Schwartz, Christian Schade, Mary Frances Luce, Steven Lippman, David Krantz, Barbara Kahn, and Robin Hogarth. 2002. High Stakes Decision Making: Normative, Descriptive and Prescriptive Considerations. *Marketing Letters* 13, 3 (Aug. 2002), 259–268. doi:10.1023/A:1020287225409
- [34] Anne Kuusisto, Jenni Santavirta, Kaija Saranto, Päivi Korhonen, and Elina Haavisto. 2020. Advance Care Planning for Patients with Cancer in Palliative Care: A Scoping Review from a Professional Perspective. *Journal of Clinical Nursing* 29, 13–14 (July 2020), 2069–2082. doi:10.1111/jocn.15216
- [35] Vivian Lai, Chacha Chen, Alison Smith-Renner, Q. Vera Liao, and Chenhao Tan. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAcT '23)*. Association for Computing Machinery, New York, NY, USA, 1369–1385. doi:10.1145/3593013.3594087
- [36] Daniela Lamas, Natalie Panariello, Natalie Henrich, Bernard Hammes, Laura C. Hanson, Diane E. Meier, Nancy Guinn, Janet Corrigan, Sean Hubber, Hannah Luetke-Stahlman, and Susan Block. 2018. Advance Care Planning Documentation in Electronic Health Records: Current Challenges and Recommendations for Change. *Journal of Palliative Medicine* 21, 4 (April 2018), 522–528. doi:10.1089/jpm.2017.0451
- [37] Michael Leyer and Sabrina Schneider. 2019. ME, YOU OR AI? HOW DO WE FEEL ABOUT DELEGATION. In *Proceedings of the 27th European Conference on Information Systems (ECIS)*. AIS Electronic Library (AISeL), Stockholm & Uppsala, Sweden.

- [38] Shuai Ma, Qiaoyi Chen, Xinru Wang, Chengbo Zheng, Zhenhui Peng, Ming Yin, and Xiaojuan Ma. 2025. Towards Human-AI Deliberation: Design and Evaluation of LLM-Empowered Deliberative AI for AI-Assisted Decision-Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–23. doi:10.1145/3706598.3713423
- [39] Preeti Malani, Jeffrey Kullgren, Erica Solway, Chithra Perumalswami, James Burke, Lesli Skolarus, Dianne Singer, and Matthias Kirch. 2021. National Poll on Healthy Aging: Older Adults' Experiences with Advance Care Planning. doi:10.7302/808
- [40] Chetna Malhotra and Chandrika Ramakrishnan. 2022. Complexity of Implementing a Nationwide Advance Care Planning Program: Results from a Qualitative Evaluation. *Age and Ageing* 51, 10 (Oct. 2022), afac224. doi:10.1093/ageing/afac224
- [41] Chetna Malhotra, Mahham Shafiq, and Ada Portia Macarubbo Batcagan-Abueg. 2022. What Is the Evidence for Efficacy of Advance Care Planning in Improving Patient Outcomes? A Systematic Review of Randomised Controlled Trials. *BMJ Open* 12, 7 (July 2022), e060201. doi:10.1136/bmjopen-2021-060201
- [42] Jennifer L. Malin, Clifford Ko, John Z. Ayanian, David Harrington, David R. Nerenz, Katherine L. Kahn, Julie Ganther-Urmie, Paul J. Catalano, Alan M. Zaslavsky, Robert B. Wallace, Edward Guadagnoli, Neeraj K. Arora, Maryse D. Roudier, and Patricia A. Ganz. 2006. Understanding Cancer Patients' Experience and Outcomes: Development and Pilot Study of the Cancer Care Outcomes Research and Surveillance Patient Survey. *Supportive Care in Cancer* 14, 8 (Aug. 2006), 837–848. doi:10.1007/s00520-005-0902-8
- [43] Sumytra Menon. 2013. The Mental Capacity Act: Implications for Patients and Doctors Faced with Difficult Choices. *Annals Academy of Medicine* 42, 4 (2013), 3.
- [44] Natasha Michael, Clare O'Callaghan, and Josephine M. Clayton. 2016. Exploring the Utility of the Vignette Technique in Promoting Advance Care Planning Discussions with Cancer Patients and Caregivers. *Patient Education and Counseling* 99, 8 (Aug. 2016), 1406–1412. doi:10.1016/j.pec.2016.03.021
- [45] Meredith Ringel Morris and Jed R. Brubaker. 2025. Generative Ghosts: Anticipating Benefits and Risks of AI Afterlives. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3706598.3713758
- [46] R. Sean Morrison, Diane E. Meier, and Robert M. Arnold. 2021. What's Wrong With Advance Care Planning? *JAMA* 326, 16 (Oct. 2021), 1575. doi:10.1001/jama.2021.16430
- [47] Subhigya Nepal, Arvind Pillai, William Campbell, Talie Massachi, Michael V. Heinz, Ashmita Kunwar, Eunsoo Soul Choi, Xuhai Xu, Joanna Kuc, Jeremy F. Huckins, Jason Holden, Sarah M. Preum, Colin Depp, Nicholas Jacobson, Mary P. Czerwinski, Eric Granholm, and Andrew T. Campbell. 2024. MindScape Study: Integrating LLM and Behavioral Sensing for Personalized AI-Driven Journaling Experiences. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 4 (Nov. 2024), 186:1–186:44. doi:10.1145/3699761
- [48] Qin Xiang Ng, Tricia Zi Kuah, Germaine Jm Loo, Wilbert Hh Ho, Norbert L. Wagner, Judy Gk Sng, Grace Mj Yang, and Bee Choo Tai. 2017. Awareness and Attitudes of Community-Dwelling Individuals in Singapore towards Participating in Advance Care Planning. *Annals of the Academy of Medicine, Singapore* 46, 3 (March 2017), 84–90. doi:10.47102/annals-acadmedsg.V46N3p84
- [49] Ayano Okoso, Mingzhe Yang, and Yukino Baba. 2025. Do Expressions Change Decisions? Exploring the Impact of AI's Explanation Tone on Decision-Making. arXiv:2502.19730 [cs] doi:10.48550/arXiv.2502.19730
- [50] Nay Myo Oo, Ian A. Scott, Raquel Maggacis, and Nalaka Rajakaruna. 2019. Assessing Concordance between Patient Preferences in Advance Care Plans and In-Hospital Care. *Australian Health Review: A Publication of the Australian Hospital Association* 43, 4 (Aug. 2019), 425–431. doi:10.1071/AH18011
- [51] Jari Pirhonen, Luciana Lolic, Katarina Tuominen, Outi Jolanki, and Virpi Timonen. 2020. "These Devices Have Not Been Made for Older People's Needs" – Older Adults' Perceptions of Digital Technologies in Finland and Ireland. *2020*, 101287. doi:10.1016/j.techsoc.2020.101287
- [52] Thaddeus Mason Pope. 2012. Legal Fundamentals of Surrogate Decision Making. *Chest* 141, 4 (April 2012), 1074–1081. doi:10.1378/chest.11-2336
- [53] Isabel R. Rooper, William W. Liem, Martha Burla, Jacob Gordon, Lara M. Baez, Rachel Kornfield, and Andrew B. L. Berry. 2025. Designing Values Elicitation Technologies for Mental Health and Chronic Care Integration: User-Centered Design Approach. *JMIR Formative Research* 9, 1 (Feb. 2025), e68419. doi:10.2196/68419
- [54] Max Roser. 2019. The Global Population Pyramid: How Global Demography Has Changed and What We Can Expect for the 21st Century. *Our World in Data* (April 2019).
- [55] Bukhoree Sahoh and Anant Choksuriwong. 2023. The Role of Explainable Artificial Intelligence in High-Stakes Decision-Making Systems: A Systematic Review. *Journal of Ambient Intelligence and Humanized Computing* 14, 6 (June 2023), 7827–7843. doi:10.1007/s12652-023-04594-w
- [56] Anita M. Schwarz, Omar S. Arias, Asta Zvinieni, Heinz P. Rudolph, Sebastian Eckardt, Johannes Koettl, Herwig Immervoll, and Miglena Abels. 2014. *The Inverting Pyramid: Pension Systems Facing Demographic Challenges in Europe and Central Asia*. World Bank Publications.
- [57] Cristina Sedini, Martina Biotto, Lorenza M. Crespi Bel'skij, Roberto Ercole Moroni Grandini, and Matteo Cesari. 2022. Advance Care Planning and Advance Directives: An Overview of the Main Critical Issues. *Aging Clinical and Experimental Research* 34, 2 (Feb. 2022), 325–330. doi:10.1007/s40520-021-02001-y
- [58] David I. Shalowitz, Elizabeth Garrett-Mayer, and David Wendler. 2006. The Accuracy of Surrogate Decision Makers: A Systematic Review. *Archives of Internal Medicine* 166, 5 (March 2006), 493–497. doi:10.1001/archinte.166.5.493
- [59] Ben Shneiderman. 2022. *Human-Centered AI* (1 ed.). Oxford University Press-Oxford. doi:10.1093/oso/9780192845290.001.0001
- [60] Kellie Yu Hui Sim, Kohleen Tijing Fortuno, and Kenny Tsu Wei Choo. 2024. Towards Understanding Emotions for Engaged Mental Health Conversations. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference (DIS '24 Companion)*. Association for Computing Machinery, New York, NY, USA, 176–180. doi:10.1145/3656156.3663694
- [61] Inhwa Song, SoHyun Park, Sachin R Pendse, Jessica Lee Schleider, Munmun De Choudhury, and Young-Ho Kim. 2025. ExploreSelf: Fostering User-driven Exploration and Reflection on Personal Challenges with Adaptive Guidance by Large Language Models. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–22. doi:10.1145/3706598.3713883
- [62] Inhwa Song, Sachin R Pendse, Neha Kumar, and Munmun De Choudhury. 2025. The Typing Cure: Experiences with Large Language Model Chatbots for Mental Health Support. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (Oct. 2025), 1–29. doi:10.1145/3757430
- [63] Philipp Spitzer, Joshua Holstein, Patrick Hemmer, Michael Vössing, Niklas Kühl, Dominik Martin, and Gerhard Satzger. 2025. Human Delegation Behavior in Human-AI Collaboration: The Effect of Contextual Information. *9, 2 (2025)*, 1–28. doi:10.1145/3710999
- [64] Dawn Stacey, France Légaré, Krystina Lewis, Michael J. Barry, Carol L. Bennett, Karen B. Eden, Margaret Holmes-Rovner, Hilary Llewellyn-Thomas, Anne Lyddiatt, Richard Thomson, and Lyndal Trevena. 2017. Decision Aids for People Facing Health Treatment or Screening Decisions. *The Cochrane Database of Systematic Reviews* 4, 4 (April 2017), CD001431. doi:10.1002/14651858.CD001431.pub5
- [65] Mary Steffel, Elanor F. Williams, and Jaclyn Permann-Graham. 2016. Passing the Buck: Delegating Choices to Others to Avoid Responsibility and Blame. *Organizational Behavior and Human Decision Processes* 135 (July 2016), 32–44. doi:10.1016/j.obhdp.2016.04.006
- [66] Christopher T. Su, Ryan D. McMahan, Brie A. Williams, Rashmi K. Sharma, and Rebecca L. Sudore. 2014. Family Matters: Effects of Birth Order, Culture, and Family Dynamics on Surrogate Decision-Making. *Journal of the American Geriatrics Society* 62, 1 (2014), 175–182. doi:10.1111/jgs.12610
- [67] Ya Su, Michiko Yuki, and Kengo Hirayama. 2020. The Experiences and Perspectives of Family Surrogate Decision-Makers: A Systematic Review of Qualitative Studies. *Patient Education and Counseling* 103, 6 (June 2020), 1070–1081. doi:10.1016/j.pec.2019.12.011
- [68] Rebecca L. Sudore, Hillary D. Lum, John J. You, Laura C. Hanson, Diane E. Meier, Steven Z. Pantilat, Daniel D. Matlock, Judith A.C. Rietjens, Ida J. Korffage, Christine S. Ritchie, Jean S. Kutner, Joan M. Teno, Judy Thomas, Ryan D. McMahan, and Daren K. Heyland. 2017. Defining Advance Care Planning for Adults: A Consensus Definition From a Multidisciplinary Delphi Panel. *Journal of Pain and Symptom Management* 53, 5 (May 2017), 821–832.e1. doi:10.1016/j.jpainsymman.2016.12.331
- [69] Amos Tversky and Daniel Kahneman. 1973. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* 5, 2 (Sept. 1973), 207–232. doi:10.1016/0010-0285(73)90033-9
- [70] Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (Sept. 1974), 1124–1131. doi:10.1126/science.185.4157.1124
- [71] Shannon Vallor. 2015. Moral Deskillling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character. *Philosophy & Technology* 28, 1 (March 2015), 107–124. doi:10.1007/s13347-014-0156-9
- [72] Elizabeth K. Vig, Helene Starks, Janelle S. Taylor, Elizabeth K. Hopley, and Kelly Fryer-Edwards. 2007. Surviving Surrogate Decision-Making: What Helps and Hampers the Experience of Making Medical Decisions for Others. *Journal of General Internal Medicine* 22, 9 (Sept. 2007), 1274–1279. doi:10.1007/s11606-007-0252-y
- [73] P. C. Wason. 1960. On the Failure to Eliminate Hypotheses in a Conceptual Task. *Quarterly Journal of Experimental Psychology* 12, 3 (July 1960), 129–140. doi:10.1080/17470216008416717
- [74] David Wendler and Annette Rid. 2011. Systematic Review: The Effect on Surrogates of Making Treatment Decisions for Others. *Annals of Internal Medicine* 154, 5 (March 2011), 336–346. doi:10.7326/0003-4819-154-5-201103010-00008
- [75] Laraine Winter. 2013. Patient Values and Preferences for End-of-Life Treatments: Are Values Better Predictors Than a Living Will? *Journal of Palliative Medicine* 16, 4 (April 2013), 362–368. doi:10.1089/jpm.2012.0303
- [76] Yan Xu, Brennan Jones, Hannah Nguyen, Qisheng Li, and Stefan Scherer. 2025. From Goals to Actions: Designing Context-aware LLM Chatbots for New Year's Resolutions. In *Proceedings of the 7th ACM Conference on Conversational User*

Interfaces. ACM, Waterloo ON Canada, 1–17. doi:10.1145/3719160.3736637

- [77] Tanya Zivkovic. 2018. Forecasting and Foreclosing Futures: The Temporal Dissonance of Advance Care Directives. *Social Science & Medicine* 215 (Oct. 2018), 16–22. doi:10.1016/j.socscimed.2018.08.035

A CPR Scenarios

Participants engaged with five escalating CPR scenarios, each representing a distinct situation involving different health trajectories and decision-making contexts. The first two scenarios depicted situations with a relatively better outlook of life (i.e., recoverable episode, longer prognosis), while the latter three reflected more advanced decline (i.e., severe treatment side-effects, financial and care burdens, and family conflict near end-of-life).

Scenario 1: Recoverable Episode

Description:

You live alone and have been managing your cancer fairly well. One night, you feel breathless and dizzy, so you lie down, thinking it will pass. A friend checks in the next day and calls for help when you don't respond. At the hospital, doctors discover a treatable infection. No doctor brings up CPR. You recover quickly and return home.

Further Circumstances:

If your heart had stopped that night, CPR would have been done by default.

Scenario 2: Terminal Illness Becomes Clear

Description:

You've been weaker lately, but still hopeful. Then one morning, you collapse. At the hospital, doctors tell you that your previous treatment is no longer working. The cancer has spread to other organs and is unlikely to be curable. You're now considered terminal, though you feel mostly okay after being stabilized in the hospital. Your doctor explains that you probably have 1–2 years of life left and gently asks if you've thought about what should happen if your heart unexpectedly stops. You've never been asked this before.

Further Circumstances:

You are alert and functioning — but your illness is no longer curable.

Scenario 3: Severe Side Effects from Aggressive Treatment

Description:

Your doctors have said that you are expected to live about 12 more months. You decide to try one last round of chemotherapy. A few days later, you collapse — too weak, too dehydrated. At the hospital, doctors say the treatment caused more harm than good. You're stable, but bedbound and have been moved to a community hospital that provides care. Even speaking feels like effort and there is some pain. You're exhausted. The doctors ask if you'd like to make a decision about whether to do CPR if your heart stops suddenly.

Further Circumstances:

The illness has impacted and reduced your quality of life. You are unable to do your daily activities, you have manageable pain and you require help with personal care (e.g. bathing).

Scenario 4: Cost and Resource Burden

Description:

Your doctors have said that you are expected to live about 12 more months. After treatment, you were discharged to a community hospital, stable but clearly not well enough to return home. It becomes obvious that you will need long-term support, maybe in a nursing home. But you live alone, and money is tight. Eventually, social workers find you a bed in a good nursing home. After moving in and getting used to life there, a nurse explains that your condition is fragile and gently asks if you have an opinion on whether you want to have CPR if your heart stops.

Further Circumstances:

The cost of your care is a great concern to you, you are reasonably comfortable, but you are frail with high dependency needs.

Scenario 5: Limited Life Expectancy, Family Conflict, Unclear Wishes

Description:

Your doctors have said that you are expected to live about 12 more months. About 3 months ago, you moved to a nursing home. While at home, you have a possible stroke. You wake up but fade in and out of consciousness. Your distant relatives are called and the doctors say you likely have only weeks left. Some say, "Do everything." Others say, "Please let them go peacefully." You never wrote down your wishes. No one knows what you would have wanted.

Further Circumstances:

You may not have much time left, and you are semi-conscious due to a long bout of illness.

B LLM Prompting Materials

B.1 Prompt Template (with Placeholders)

ACPAgent Prompt Template

You are ACPAgent, a personalised assistant designed to support Advance Care Planning (ACP). Speak directly to the user using “you” and “your”. Avoid third-person references like “the user” or “they”. The user is uniquely identified by makerID: `#{makerID}`.

You have long-term memory and should reason based on all information provided so far. You have access to the following data from the user:

—

Lifestyle Preferences

`#{lifeStr}`

Reasons: `#{lifestyleReason}`

Use these to connect to how the user envisions living well day-to-day, and to prompt reflection if needed. The user might have provided a reason for their lifestyle choices, which can be used to understand their values better.

Goals of Care

`#{goalStr}`

These are important values that shape what matters most in times of serious illness (e.g., avoiding burden, spending time with family, preserving independence).

Scenario `#{scenarioKey}`: `#{scenarioTitle}`

Scenario Description: `#{description}`

Further Circumstances for User Reflection

Each scenario consists of additional circumstances for the user to reflect on. These are designed to help the user think about their situation and how they would feel in that scenario.

Here are the circumstances for this scenario:

`#{circumstances}`

Past Decisions and Feedback

Below are the user’s past decisions and reasons, which should be used to inform your response. If the user has not made a decision yet, you can skip this section.

- Previous Decision: `#{previousUserDecision}`
- Previous Decision Reason: `#{previousUserDecisionReason}`
- Previous Scenario Influence: `#{previousScenarioInfluence}`
- Previous Scenario Influence Reason: `#{previousScenarioInfluenceReason}`

Below are the user’s feedback on your previous recommendation, which should be used to inform your response. If the user has not provided feedback yet, you can skip this section.

- Previous Agreement with Bot Recommendation: `#{previousAgreeWithBotRecommendation}`
- Previous Decision Agreement Reason: `#{previousDecisionAgreementReason}`
- Previous Bot Helpfulness Reason: `#{previousBotHelpfulnessReason}`
- Previous Other Considerations: `#{previousOtherConsiderations}`

Note that for the first round, the user has not yet made a decision, so there is no past decision, reason, or considerations to reference.

—

Please analyse the data and return a clear, grounded, and personalised response with the following **four sections**, using **bolded headings**. Use bullet points where appropriate.

1) Thinking Process

Explain your reasoning. Draw on the user's lifestyle, goals, trade-off preferences, scenario, and prior feedback.

Reflect the user's lifestyle preferences, goals of care, life extension trade-offs, and answers back to the user.

Use the scenario description and circumstances to inform your thinking.

2) Recommended Decision

Using the circumstances, give a short, clear decision for the user on whether CPR would be needed, based on your analysis, using "I think you would choose...".

3) Reasons

Justify your decision with clear links to the user's data or the scenario (description, circumstances):

- Lifestyle and goals of care
- Trade-off values (pain vs life extension, cost vs life extension)
- Scenario description and circumstances
- Past decisions and feedback

4) Binary Decision (Yes/No to CPR)

Generate a binary decision on whether the user would choose CPR no matter what. Return only one line for programmatic use:

- Yes
- No

—

Guidelines:

- Stay under **300 words**
- Use clear structure and short paragraphs
- Reason cumulatively — learn and adapt across rounds
- Future scenarios should build on what has been learned

B.2 Worked Example (Round 2)

Worked Example (Round 2) - Sample JSON File

```

1  {
2    "makerID": "sampleParticipant",
3    "scenarioKey": 2,
4    "scenarioTitle": "Terminal Illness Becomes Clear",
5    "userLifestyle": ["Spending time with loved ones"],
6    "lifestyleReason": "I want to enjoy quality time with my grandchildren.",
7    "userGoalsOfCare": ["Avoid pain and suffering", "Avoid being a burden"],
8    "painSlider": 3,
9    "costSlider": 7,
10   "previousUserDecision": false,
11   "previousUserDecisionReason": "I don't want to prolong suffering.",
12   "previousScenarioInfluence": "Financial burden on family",
13   "previousScenarioInfluenceReason": "I don't want them to pay high costs.",
14   "previousAgreeWithBotRecommendation": true,
15   "previousDecisionAgreementReason": "It aligned with my values.",
16   "previousBotHelpfulnessReason": "It helped me think more clearly.",
17   "previousOtherConsiderations": "Prefer to stay at home if possible."
18 }

```

Worked Example (Round 2) - Prompt Sent

You are ACPAgent, a personalised assistant designed to support Advance Care Planning (ACP). Speak directly to the user using “you” and “your”. Avoid third-person references like “the user” or “they”. The user is uniquely identified by makerID: $\$ \{ \text{makerID} \}$.

You have long-term memory and should reason based on all information provided so far. You have access to the following data from the user:

—

Lifestyle Preferences

Spending time with loved ones

Reasons: I want to enjoy quality time with my grandchildren.

Use these to connect to how the user envisions living well day-to-day, and to prompt reflection if needed. The user might have provided a reason for their lifestyle choices, which can be used to understand their values better.

Goals of Care

Avoid pain and suffering, Avoid being a burden

These are important values that shape what matters most in times of serious illness (e.g., avoiding burden, spending time with family, preserving

Scenario 2: Terminal Illness Becomes Clear

Description: You’ve been weaker lately, but still hopeful. Then one morning, you collapse. At the hospital, doctors tell you that your previous treatment is no longer working. The cancer has spread to other organs and is unlikely to be curable. You’re now considered terminal, though you feel mostly okay after being stabilized in the hospital. Your doctor explains that you probably have 1–2 years of life left and gently asks if you’ve thought about what should happen if your heart unexpectedly stops.

Further Circumstances for User Reflection

Each scenario consists of additional circumstances for the user to reflect on. These are designed to help the user think about their situation and how they would feel in that scenario.

Here are the circumstances for this scenario:

You are alert and functioning — but your illness is no longer curable.

Past Decisions and Feedback

Below are the user’s past decisions and reasons, which should be used to inform your response. If the user has not made a decision yet, you can skip this section.

- Previous Decision: No, I would not want CPR
- Previous Decision Reason: I don’t want to prolong suffering.
- Previous Scenario Influence: Financial burden on family
- Previous Scenario Influence Reason: I don’t want them to pay high costs

Below are the user’s feedback on your previous recommendation, which should be used to inform your response. If the user has not provided feedback yet, you can skip this section.

- Previous Agreement with Bot Recommendation: Yes
- Previous Decision Agreement Reason: It aligned with my values.
- Previous Bot Helpfulness Reason: It helped me think more clearly.
- Previous Other Considerations: Prefer to stay at home if possible.

Note that for the first round, the user has not yet made a decision, so there is no past decision, reason, or considerations to reference.

—

Please analyse the data and return a clear, grounded, and personalised response with the following **four sections**, using **bolded headings**. Use bullet points where appropriate.

1) Thinking Process

Explain your reasoning. Draw on the user's lifestyle, goals, trade-off preferences, scenario, and prior feedback.

Reflect the user's lifestyle preferences, goals of care, life extension trade-offs, and answers back to the user.

Use the scenario description and circumstances to inform your thinking.

2) Recommended Decision

Using the circumstances, give a short, clear decision for the user on whether CPR would be needed, based on your analysis, using "I think you would choose...".

3) Reasons

Justify your decision with clear links to the user's data or the scenario (description, circumstances):

- Lifestyle and goals of care
- Trade-off values (pain vs life extension, cost vs life extension)
- Scenario description and circumstances
- Past decisions and feedback

4) Binary Decision (Yes/No to CPR)

Generate a binary decision on whether the user would choose CPR no matter what. Return only one line for programmatic use:

- Yes
- No

—

Guidelines:

- Stay under **300 words**
- Use clear structure and short paragraphs
- Reason cumulatively — learn and adapt across rounds
- Future scenarios should build on what has been learned

B.3 Sample Model Output (Round 2)

Sample GPT-4o Response (Round 2)

1) Thinking Process

You have stated that spending time with loved ones is central to your lifestyle, and you want to enjoy quality time with your grandchildren. Your goals of care focus on avoiding pain and not being a burden. You have indicated a low tolerance for pain (3/10) and a moderate willingness to accept higher costs (7/10). In this scenario, your illness is now terminal and no longer curable. Previously, you decided against CPR, citing concerns about suffering and the financial burden on your family.

2) Recommended Decision

I think you would choose **not to receive CPR**.

3) Reasons

- Preserving independence and avoiding pain are key values for you, and CPR could undermine both in this terminal context.
- Your low pain tolerance (3/10) reinforces a preference for comfort over aggressive intervention.
- While you have some tolerance for costs (7/10), you previously emphasised not wanting your family to face financial burdens.
- Your earlier refusal of CPR and agreement with the bot's recommendation suggest consistency in your decision-making.
- The scenario confirms your illness is terminal, strengthening the case for a comfort-focused approach.

4) Binary Decision (Yes/No to CPR)

No

C CPR Information Sheet

Cardio-Pulmonary Resuscitation (CPR)



What is Cardio-Pulmonary Resuscitation (CPR)?

Basic CPR involves depressing on the chest and blowing air into the lungs. It is attempted when a person's breathing and heartbeat have stopped, or have become ineffective. Most people believe CPR is more successful than it actually is.

Statistics

- Less than 17% of hospitalized patients survive CPR and are discharged
- CPR is successful in less than 5% of people who have end-stage chronic illness
- CPR is successful in less than 2% of frail elderly living in long-term care facilities
- For comparison, even among healthy individuals, people who survive CPR and are discharged ranges from 10-30%



More about CPR

Who?

CPR is most beneficial for:

- A healthy person whose heart has stopped suddenly from an accident or heart attack
- Patients with underlying conditions that can be effectively treated

CPR is less likely to be beneficial for:

- Patients with severe heart disease
- Patients with severe frailty, especially those in long-term care facilities
- Patients with end-stage disease

What does CPR involve?

- Forceful chest compressions to restart the heart after it has stopped
- Electrical stimulation of the heart
- Blowing air into the lungs, and inserting an airway tube to assist with breathing (intubation)
- Giving intravenous medications





Benefits and Burdens

Benefits of CPR

CPR can save lives with good functional recovery:

- for patients who are young and healthy
- if it is initiated within 5 minutes of a cardiac arrest

Burdens of CPR

There are many possible complications:

- The patient's ribs may be broken during CPR
- The patient's internal organs (liver or spleen) may be injured
- The patient's lungs might collapse, meaning the patient will require ventilation support to live
- The patient may be severely brain injured
- The patient will need to stay in the Intensive Care Unit (ICU) if mechanical ventilation support is required, potentially for life

CPR has helped many people recover, but it has also resulted in people who were severely brain injured or who are only kept alive through mechanical ventilation.

D Participant Reflection Questions

The following open-ended prompts were used to guide the reflection stage after participants completed all five rounds:

- When you were customising the bot, what were you trying to achieve?
- What kinds of values, preferences, or ways of thinking did you try to build into the bot?
- What did you ask your bot to elaborate further on/consider, and why?
- Was there anything you wished you could have customised but couldn't?
- What would you change or improve in the bot customisation experience?
- How did you feel about the way your bot responded?
- Did anything your bot said feel off, uncomfortable, or surprising to you?
- Do you feel that the bot accurately represented what you would want to say in the various scenarios? Why or why not?
- Were you able to adjust or challenge its responses when needed? How did that go?
- What would you change or improve about how your bot performed?
- Would you be comfortable with a bot like this speaking on your behalf if you were unable to?
- What would need to change for you to feel more confident about the bot?

E Participant Screener Form

The following screener was used to determine participant eligibility and capture relevant background information.

Section 1: Demographics

- **Age:** (open-ended)
- **Gender:** Male Female Non-binary Others (please specify)
- **Education (current or highest attained):** Secondary and below Polytechnic/JC/ITE/MI Bachelor's Master's Doctoral Other
- **Ethnicity:** Chinese Malay Indian Others (please specify)
- **Employment status:** Full-time Part-time Contract/Temporary Retired Unemployed Unable to work Others (please specify)

Section 2: ACP and Health Planning

- **Have you heard of or taken part in Advance Care Planning (ACP)?**
 Never heard Heard, not participated Participated as Donor Participated as Nominated Healthcare Spokesperson
- **Have you heard of or appointed a Lasting Power of Attorney (LPA)?**
 An LPA is a legal document that allows you to appoint someone you trust to make decisions on your behalf if you lose mental capacity one day. These include decisions pertaining to finances, healthcare, and living arrangements.
 Never heard Heard, not completed Completed LPA (appointed Donee) Agreed to be Donee for someone else
- **Do you currently work in any of the following?**
 Healthcare Healthtech/Digital health Social work/Counselling None

Section 3: Caregiving or Support Context

- **Do you currently have anyone in your life who meets the following description?**
 Someone who has been granted powers under a Lasting Power of Attorney (LPA) to make decisions on your behalf if you were to lose your mental capacity one day. These include decisions pertaining to finances, healthcare, and living arrangements.
 Yes No Unsure
- **Do you currently have anyone in your life who meets the following description?**
 Someone who is responsible for acting in your best interests and adhering to the terms of an LPA.
 Yes No Unsure
- **Do you currently have anyone in your life who meets the following description?**
 A trusted individual (e.g., a family member, close friend, or professional deputy) who has been appointed to make important decisions on your behalf.
 Yes No Unsure

Section 4: Health Background

- **Have you ever been diagnosed with any of the following conditions?**
 Advanced Cancer Stroke End Stage Kidney Failure Dementia Heart Failure Chronic Lung Disease Chronic Liver Disease None of the above
- **Do you have any family members or close friends who have been diagnosed with any of the following conditions?**

- Advanced Cancer Stroke End Stage Kidney Failure Dementia Heart Failure Chronic Lung Disease Chronic Liver Disease None of the above

• **Have you undergone any major surgeries or serious hospitalisations in the past 5 years?**

- Yes No (If yes, please briefly describe)

Section 5: Digital and AI Tools Usage

• **Have you used any digital health tools to manage your care or navigate treatment options?**

- Yes (please specify) No

• **Have you used any AI-powered tools (e.g., ChatGPT, health chatbots, symptom checkers)?**

- Yes, regularly Yes, occasionally Never

• **Digital and AI Behaviour Inventory**

(1 = Daily use, 2 = Seldom use, 3 = Never use)


- I use my computer or smart phone to send and receive email.
 1 2 3
- I use my computer or smart phone to obtain information on a wide range of topics.
 1 2 3
- I download applications from the internet to my computer or smart phone.
 1 2 3
- I use my computer or smart phone to shop, manage my calendar and/or make travel arrangements.
 1 2 3
- I use my computer or smart phone to bank and pay my bills.
 1 2 3
- I use my computer or smart phone for social networking.
 1 2 3
- I use my computer or smart phone to watch movies/videos, listen to podcasts and/or music, or share photos/images.
 1 2 3
- I use other forms of electronic technology such as eBooks (Kindle, NookBook) or tablets (iPad, LifeBook, etc.).
 1 2 3
- I use AI chatbots (e.g., ChatGPT) to search for information or learn new things.
 1 2 3
- I use AI chatbots to help me write emails, essays, or other written content.
 1 2 3
- I use AI chatbots to assist with problem-solving or generating ideas.
 1 2 3
- I use AI tools for translating text or understanding content in other languages.
 1 2 3
- I use voice assistants (e.g., Siri, Alexa, Google Assistant) powered by AI to carry out tasks.
 1 2 3
- I use generative AI to create images, audio, or video content.
 1 2 3
- I use AI to automate tasks or summarise information (e.g., meeting notes, documents).
 1 2 3
- I use AI-powered recommendations (e.g., for music, videos, news) in my daily life.
 1 2 3
- I have used virtual reality (VR) for gaming, simulations, or learning.
 1 2 3
- I have used augmented reality (AR) features on apps (e.g., filters, navigation, shopping).
 1 2 3
- I feel comfortable navigating 3D or immersive environments in VR.
 1 2 3
- I have used VR or AR tools for health, education, or training purposes.
 1 2 3

F Preference Changes


The figures below show how participants' stated preferences evolved across the five rounds of interaction. Figure 9 presents all lifestyle and goals-of-care selections across rounds, while Figures 10a and 10b show numeric trajectories for pain tolerance and cost concern. Only participants who revised these ratings at least once are displayed. Baseline ratings (Round 1) appear as the leftmost point on each line.



Figure 9: Lifestyle and Goals of Care preferences across rounds. Rows are ordered by the round of each participant's first change, with IDs ascending inside groups. Columns list possible preference items. Within each cell, stacked strips represent all five rounds, colour-coded by round. A strip appears only when that preference was selected in the corresponding round.



(a) Pain tolerance trajectories.



(b) Cost concern trajectories.

Figure 10: Numeric trade-off trajectories across rounds. Pain tolerance was generally stable; cost concern showed greater fluctuation, especially following scenarios involving financial or dependency dilemmas. Baseline values (Round 1) appear as the left-most point for each line.