

# Designing the Algorithmic Operator

## Evidence-Based Conversational Architecture for Voice AI in Community Resource Navigation

Johan Michalove\*

February 26, 2026 — *Working draft*

### Abstract

How should a voice-based AI system be designed to help people navigate community resources? This paper presents an evidence-based design methodology for voice AI in social service information and referral, synthesizing principles from six domains: the history of telephone operators and the Bell System (1878–present), crisis hotline best practices, interactive voice response (IVR) usability research, 211/311 information and referral systems and the AIRS standards, telephone-based social services for vulnerable populations, and social work intake models including motivational interviewing. From this synthesis, we derive 26 design principles organized into two complementary frameworks: 14 principles addressing the voice interaction itself (pacing, structure, escalation, accessibility) and 12 principles addressing the conversational stance (strengths-based orientation, trauma-informed care, cultural humility, self-determination). We argue that the design of a voice AI system for social services is not primarily a technical problem of speech recognition and natural language processing but a practice problem grounded in over a century of evidence about how to help people by telephone. The paper contributes to HCI by offering a principled design methodology that treats the voice interaction as a care relationship rather than an information retrieval task, and by demonstrating how professional standards from social work, crisis intervention, and information and referral can be translated into the design of conversational AI systems. A companion paper (Michalove, 2026) provides the STS analysis of what this design means politically; this paper addresses how it should be built.

---

\*Department of Information Science, Cornell University. jam844@cornell.edu

**Keywords:** voice AI, conversational design, social services, information and referral, 211, motivational interviewing, trauma-informed care, design methodology

# 1 Introduction

The design of voice-based AI systems for social service navigation has proceeded, for the most part, without reference to the century of accumulated knowledge about how to help people by telephone. Speech recognition papers optimize word error rates. Dialogue system papers maximize task completion. Chatbot papers measure user satisfaction on Likert scales. Meanwhile, the 211 system—which handles tens of millions of calls per year across the United States—has developed, through decades of practice and codification, a sophisticated understanding of what it means to listen to someone describe a need and connect them to help. Crisis hotlines have produced empirical evidence about which counselor behaviors reduce distress. Social work has formalized the ethics of the helping relationship. IVR usability research has identified the cognitive limits of voice-based interaction. And the history of telephone operators, stretching back to 1878, reveals that the fundamental insight—the voice is the service—was established before any of the disciplines that now study voice interaction existed.

This paper synthesizes these traditions into a design methodology for voice AI in community resource navigation. Its contribution is not a new algorithm or a new model but a *principled account of what the system should do and why*, grounded in evidence-based practice across six domains. We argue that the primary design challenge is not technical (how to build a system that can understand speech and search a database) but relational (how to build a system that can be of use to a person in need). The technical requirements follow from the relational ones, not the other way around.

The methodology is organized around 26 design principles derived from the literature review in Sections 2–4. We present these principles not as implementation prescriptions but as design commitments that any voice AI system for social services must confront. They address questions that are simultaneously technical and ethical: How long should a greeting be? Should the system categorize the caller’s need before listening? What happens when the caller’s problem does not fit any category? How should the system handle a caller who is in crisis? What does it mean to provide a “referral” that leads nowhere?

A companion paper (Michalove, 2026) analyzes the political and ethical dimensions of voice AI in social services through the frameworks of infrastructure studies, classification theory, and the sociology of testing. The present paper is its complement: where the companion asks *what it means* for a voice to become an infrastructure of care, this paper asks *how to design* that infrastructure—drawing on the same intellectual traditions but oriented toward practice.

The paper is structured as follows. Section 2 reviews the voice interaction literature across four domains: historical telephone operators, crisis hotlines, IVR usability, and 211/311 information and referral. Section 3 addresses the specific needs of vulnerable populations who are likely to use such a system. Section 4 synthesizes social work practice frameworks relevant to conversational design. Section 5 presents the 26 design principles. Section 6 describes the conversational architecture that implements these principles. Section 7 discusses limitations, ethical considerations, and the relationship between design methodology and the politics of care. Section 8 concludes.

## 2 The Voice Interaction: Lessons from a Century of Practice

### 2.1 Historical telephone operators (1878–present)

The earliest telephone operators were teenage boys hired from telegraph offices. They proved “impatient and rude, full of pranks, including disconnecting customers and misdirecting their calls.” Alexander Graham Bell hired Emma Nutt in Boston in 1878—the world’s first female telephone operator, described as “patient and savvy, with a voice that was cultured and soothing.” By the end of the 1880s, the job had become exclusively female (Fischer, 1992; Green, 2001).

The Bell System’s “voice with a smile” philosophy, developed under Theodore Vail’s leadership, established that voice quality and emotional tone are inseparable from service quality (Hochschild, 1983). The standard greeting—“Number, please?”—was minimal, functional, and immediate: it announced the operator’s role, invited the caller to state their need, and offered attention, all in two words. Operator protocol followed a clear sequence: (1) become aware of calls, (2) obtain verbal instructions as to destination, (3) determine whether the called line was busy or idle, (4) alert the called station, (5) report on status.

**Design insight:** The earliest lesson in telephone service design is that tone, patience, and courtesy are not incidental to the service—they are the core product. The operator was functionally a human router: acknowledge, assess, route, confirm.

### 2.2 Crisis hotline best practices

Mishara et al. (2007) silently monitored 2,611 calls to 14 helplines and found that a nondirective Rogerian style was significantly related to reductions in suicidal urgency for *new*

*callers*, while *repeat callers* benefited more from a directive approach. Gould et al. (2007) assessed 1,085 suicidal callers and found significant decreases in suicidality during the telephone session, with continuing decreases in hopelessness and psychological pain in the following weeks. Nearly 98% of callers reported that the crisis call helped, and 88.1% said it stopped them from killing themselves.

SAMHSA's National Guidelines for Behavioral Health Crisis Care state the essential principle: "Crisis services must be designed to serve anyone, anywhere, and anytime" (SAMHSA, 2014). The ASIST (Applied Suicide Intervention Skills Training) model identifies three domains of effective counselor behavior: engagement/connection, collaborative problem-solving, and safety assessment.

**Design insight:** First-time callers need empathy and reflection before problem-solving. Repeat callers benefit from more direct, action-oriented responses. Detect whether a caller is new or returning and adjust conversational style accordingly. No eligibility requirements, no hours restrictions, no geographic gatekeeping.

## 2.3 IVR usability research

Miller (1956) established the classic "seven plus or minus two" limit on working memory capacity, but more recent research (Cowan, 2001) suggests the practical limit is closer to 3–4 items. Commarford et al. (2008) found that broad-structure IVRs outperformed deep-structure IVRs, especially for users with low working memory capacity. The W3C Cognitive and Learning Disabilities Accessibility Task Force recommends: easy access to a human, options before instructions ("For housing, press 1" not "Press 1 for housing"), no filler, simplest possible terms.

Industry data shows that call completion rates improve from approximately 60% with traditional IVR to 91% with conversational voice AI. Research on synthesized versus recorded speech finds that comprehension is equivalent, but synthetic speech requires more cognitive effort—a gap that is particularly consequential for elderly callers and those with hearing impairments. Welcome-to-end-of-menu prompts should not exceed 30 seconds. Critically, 34% of callers who are not answered quickly hang up and never call back.

**Design insight:** Three to four options maximum. Flat, not deep. "How can I help you?" plus natural language beats any menu tree. The first 30 seconds determine everything. For vulnerable populations with limited phone minutes, every second matters even more. Invest in the highest quality voice model available.

## 2.4 211/311 information and referral

The AIRS Standards (now Inform USA, version 10.0) codify the five-stage I&R process (Inform USA, 2023): (1) **Opening**—establish rapport and connection; (2) **Assessment**—active listening and effective questioning; (3) **Clarification**—ensure mutual understanding, obtain consent; (4) **Referral**—provide appropriate information enabling informed choice; (5) **Closing**—summarize, confirm understanding, offer follow-up.

The referral outcomes data is sobering. Boyum et al. (2016) followed 1,235 callers to 211 and found that while 91% tried contacting a referral, only 36% received assistance. Food referrals had a 67% success rate; housing referrals succeeded only 17% of the time. The gap between referral and receipt of service—what the I&R field calls the “last mile”—is the single most important design challenge for any system that connects callers to resources.

A Bronx-based study found that human community health workers raised referral success from 36% to 78% through active navigation: telling callers what to say when they call, what documents to bring, what to expect. An AI acting as navigator should guide, not merely inform.

**Design insight:** The canonical call flow is greet–assess–clarify–refer–close. Connecting a caller to a resource is not the same as the caller receiving help. Provide multiple referrals, set realistic expectations, include what to expect and what to say. The AI should function as a navigator, not a directory.

## 3 Designing for Vulnerable Populations

### 3.1 Elderly callers

Comprehension declines more rapidly in older adults as speech rate increases. The optimal rate is 150–160 words per minute with strategic pauses after actionable information. Artificially slow “elderspeak” is perceived as patronizing and should be avoided. The MATCH project (Wolters et al., 2009) found that older users used more social words (“thank you,” “please”), full sentences, and provided additional information beyond what the system asked. Zajicek (2001) recommends keeping output messages as short as possible, offering at most three selections, and including confirmation messages. Critically: “patterns recommended for older users tend to represent improved design for the population at large.”

Bickmore et al. (2005) found that trust in conversational agents builds over repeated interactions, and that consistency of personality and voice is critical. A voice system

that sounds different on each call undermines the relational foundation that makes the interaction useful.

**Design insight:** Design for the elderly caller first. If it works for them, it works for everyone. Configure text-to-speech to approximately 150 WPM. Handle polite, verbose, socially contextualized speech—not just keywords. Maintain a consistent voice identity across interactions.

### 3.2 Homeless and houseless callers

Raven et al. (2018) found that 72.3% of older homeless adults had phone access, but with significant constraints: limited minutes, prepaid plans, high number turnover, no outlet access, no voicemail. Phone access rates vary from 44% to 94% depending on the population studied.

**Design insight:** Assume limited minutes, no callback number, no pen and paper, and the possibility that the phone may die mid-call. Deliver value immediately. Front-load the most critical information (name, address, hours). Offer to send a follow-up SMS with details. Never require a callback.

### 3.3 Non-English-speaking callers

Over 25 million Americans are limited English proficient (LEP). New York City law requires language access in 12 or more languages. The system should detect language preference within the first 5 seconds of interaction—not through a menu (“Press 2 for Spanish”) but through the caller’s own speech. This requires either multilingual speech recognition or a brief bilingual greeting that invites response in the caller’s preferred language.

**Design insight:** Respond in the caller’s language without comment. Do not make the caller navigate a language menu before receiving help. The caller is the expert on their own language; the system should adapt to them, not the reverse.

### 3.4 Speech recognition bias

Koenecke et al. (2020) found that commercial ASR systems from five major providers produce word error rates roughly twice as high for Black speakers as for white speakers (35% versus 19%). This disparity is not incidental to a social service system—it is structurally consequential. If the system cannot understand the caller, it cannot help them. Populations most in need of the service are also the populations most likely to be misheard.

**Design insight:** Speech recognition error is not a neutral technical limitation; it is a form of exclusion. Design confirmation mechanisms (“I heard you say you’re looking for shelter. Is that right?”) that catch errors without making the caller feel they are being corrected. Never blame the caller for being misunderstood.

## 4 Social Work Practice Frameworks for Conversational Design

### 4.1 Person-in-Environment and ecological systems

Social work’s Person-in-Environment (PIE) framework assesses four factors: social role functioning, environmental problems, mental health, and physical health. It is action-oriented—a “descriptive classification of factors required to implement practical next steps of helping.” Bronfenbrenner (1979) extends this into an ecological model: microsystem (immediate relationships) → mesosystem (interactions between microsystems) → exosystem (community resources, workplace policies) → macrosystem (laws, economic conditions, cultural values) → chronosystem (time, life transitions).

**Design insight:** A caller asking about “childcare” may be navigating a job transition, a family change, or a policy gap. Do not assume which system is the entry point. The presenting problem is frequently not the underlying need—“I need food” may mean income loss, domestic violence, disability, or immigration status. One clarifying question can prevent a mismatched referral.

### 4.2 Strengths-based perspective

Saleebey (1996) proposes the ROPES framework: Resources, Options, Possibilities, Exceptions, Solutions. The strengths-based approach asks “What are you looking for?” rather than “What’s wrong?” It foregrounds the caller’s capacity and agency rather than their deficits.

**Design insight:** Open with possibility, not pathology. The system’s first question should invite the caller to describe what they need, not to justify why they need it.

### 4.3 Motivational interviewing: OARS

Miller and Rollnick (2012) codify the core skills of motivational interviewing as OARS: **Open questions** (“What brings you here today?” not “Are you looking for shelter?”), **Af-**

**firmations** (acknowledge effort and self-efficacy), **Reflective listening** (mirror back what was heard, sometimes with slight reframe), and **Summary reflections** (check understanding).

The *righting reflex*—the urge to tell clients the solution immediately—is explicitly counterproductive. “I need help” should not trigger a resource dump; it should trigger a question.

**Design insight:** Resist the righting reflex. Do not jump to referral on keyword match. Use reflective listening (“It sounds like you’re looking for somewhere warm to stay tonight”) before providing resources. This is not inefficiency—it is the mechanism by which the system confirms it has understood the need correctly.

#### 4.4 Three Conversations Model

The Partners4Change Three Conversations Model proposes: (1) Listen. Make a connection. Understand what the person wants to tell you. (2) If crisis, stick with the person. Emergency plan. (3) Only if needed, assess eligibility for longer-term support. Most intake systems begin with step 3—eligibility assessment. This is backwards. The first conversation is not about eligibility; it is about being heard.

**Design insight:** The conversation IS the service. Being heard is help. Do not force a referral when the caller needs to be listened to. And never begin with eligibility screening.

#### 4.5 NASW Code of Ethics

The National Association of Social Workers Code of Ethics (National Association of Social Workers, 2021) establishes four principles directly relevant to conversational AI design: (1) **Self-determination** (1.02)—respect and promote the right of clients to determine their own goals; (2) **Informed consent** (1.03)—clear language about purpose, risks, alternatives; (3) **Referral standards** (2.06)—disclose pertinent information to new providers; (4) **Cultural competence** (1.05)—understand culture, demonstrate competence.

**Design insight:** The caller determines the need, not the system. Self-determination means the system should never override the caller’s stated preference with its own classification of what they “really” need.

#### 4.6 Trauma-informed care

SAMHSA identifies six principles of trauma-informed care: Safety, Trustworthiness/Transparency, Peer Support, Collaboration/Mutuality, Empowerment/Voice/Choice, and Cul-

tural/Historical/Gender Issues (SAMHSA, 2014). Critically, trauma-informed care does not mean assuming trauma. It means: viewing difficulties as adaptive responses rather than pathology, not requiring disclosure, emphasizing resilience over deficiency, providing choice and control.

**Design insight:** A 90-second voice interaction can be trauma-informed if it is safe (no judgment), transparent (the system says what it is and what it can do), collaborative (the caller directs the conversation), empowering (the caller makes the choice), and culturally respectful (the system follows the caller’s language and framing). Do not ask about trauma. Do not use clinical language. Do not screen everyone for crisis.

## 4.7 Warm lines, helplines, and the typology of care

Three distinct models of telephone-based care exist: **Hotlines** (988, 911)—emergencies, acute crisis, trained crisis counselors. **Helplines** (211)—full-spectrum information and referral, community resource specialists. **Warm lines**—non-crisis emotional support, trained peers, connection and loneliness reduction.

**Design insight:** A voice AI for community resource navigation is a helpline (I&R focused) with warm line qualities (non-clinical, accessible, no wrong question). It should never attempt to function as a crisis hotline. It should know how to escalate: 911 for medical emergencies, 988 for suicidal ideation. But it should not screen everyone for crisis—doing so transforms the interaction from care into surveillance.

# 5 26 Design Principles

The following principles synthesize the evidence reviewed in Sections 2–4. They are organized into two frameworks: 14 principles addressing the voice interaction (derived from the telephony, crisis, IVR, and 211 literature) and 12 principles addressing the conversational stance (derived from social work practice, motivational interviewing, and trauma-informed care). Together, they constitute a design methodology for voice AI in community resource navigation.

## 5.1 The Voice Interaction (14 principles)

1. **The voice is the service.** Warmth, pacing, and clarity are not features of the interface—they are the product. Invest in the highest-quality voice synthesis available. The Bell System learned this in 1878.

2. **30-second rule.** Engage meaningfully within 30 seconds or lose the caller. For a population with limited phone minutes, every second is a resource.
3. **AIRS five-stage model.** The canonical flow: Greet → Assess → Clarify → Refer → Close. Do not skip stages. Do not reorder them.
4. **Empathy before problem-solving.** First-time callers need to feel heard before they can use information. Repeat callers benefit from directness (Mishara et al., 2007).
5. **Three options maximum.** Working memory capacity for voice is 3–4 items (Cowan, 2001). Never present more than three choices at once.
6. **Flat, not deep.** Open-ended conversational input beats any menu tree. Do not nest. “How can I help you?” outperforms every IVR.
7. **Expect verbose, polite speech.** Older callers use full sentences, social pleasantries, and provide context the system did not request. This is not noise—it is data.
8. **150 WPM with strategic pauses.** Do not speed up to fit more information. Cut content instead. Insert pauses after actionable information (addresses, phone numbers).
9. **The last mile matters.** Only 36% of 211 referrals result in the caller receiving assistance (Boyum et al., 2016). Provide multiple referrals. Set expectations. Include what to bring, what to say, and what to expect. Navigate, do not merely inform.
10. **Warm handoffs beat cold referrals.** The closest approximation in a voice AI: tell the caller what to say when they call the referred resource, what to expect, and follow up via SMS if possible (Taylor and Minkovitz, 2021).
11. **Never make the caller start over.** If the interaction escalates to a human (or to a different system), pass context forward. “If the move from a bot to a person feels like starting again, the entire self-service effort is wasted.”
12. **Design for the houseless caller.** Assume limited minutes, no pen, no stable call-back number, and the possibility that the phone may die mid-call. Front-load critical information. Offer SMS follow-up (Raven et al., 2018).
13. **Detect language within 5 seconds.** Do not require the caller to navigate a language menu. Respond in the caller’s language without comment.
14. **The system is a warm line.** Not crisis, not nothing. Proactive support that prevents crises. Know when to escalate (911, 988) and when to stay.

## 5.2 The Conversational Stance (12 principles)

15. **Open, not sorted.** Do not pre-categorize the caller's need. The full taxonomy is available; let the conversation surface what is needed.
16. **Resist the righting reflex.** Do not jump to referral on keyword match. "I need help" is not a search query—it is the beginning of a conversation (Miller and Rollnick, 2012).
17. **The caller determines the need.** Self-determination (National Association of Social Workers, 2021). The system's classification of the caller's need must never override the caller's own account.
18. **Strengths-based.** "What are you looking for?" not "What's wrong?" Foreground resources, options, possibilities (Saleebey, 1996).
19. **No wrong door.** Never say "that's not what we do." Every need is a valid entry point. If the system cannot help directly, it should know where to route.
20. **Trauma-informed without trauma-assuming.** Do not ask about trauma. Do not use clinical language. Do not screen for crisis by default. Viewing difficulties as adaptive responses, not pathology (SAMHSA, 2014).
21. **Honest referrals.** A bad referral is worse than no referral. If a resource is known to be closed, full, or requires documentation the caller is unlikely to have, say so. Do not generate false hope.
22. **Graduated escalation.** 911 for medical emergencies, 988 for suicidal ideation. Do not screen everyone for crisis—doing so transforms care into surveillance.
23. **Cultural humility.** Respond in the caller's language without comment. Acknowledge that the caller is the expert on their own life. Do not impose categories (National Association of Social Workers, 2021).
24. **The conversation is the service.** Being heard is help. Some callers do not need a referral—they need to be listened to. Do not force a referral when presence is what is needed.
25. **Brevity with substance.** One warm clause plus an answer. But include what to expect—hours, address, what to bring, what to say.

26. **Iceberg awareness.** The presenting problem is not the underlying need. One gentle follow-up question (“Is there anything else going on that I should know about?”) opens the door without forcing it.

## 6 Conversational Architecture

The 26 principles imply a conversational architecture that differs fundamentally from both traditional IVR systems (which route through menus) and task-oriented dialogue systems (which optimize for slot-filling). The architecture described here treats the voice interaction as a *care encounter* rather than an information retrieval task.

### 6.1 The greeting

The greeting must accomplish three things in under 10 seconds: (1) identify the service, (2) establish the register of care, and (3) invite the caller to speak. It should *not* include a disclaimer, a menu, or an explanation of what the system can do. The Bell System’s “Number, please?” accomplished all three in two words. A contemporary equivalent: “*You’ve reached [name], a free helpline for [city]. How can I help you today?*” Returning callers—those the system has interacted with before—should receive an abbreviated greeting that acknowledges the relationship: simply the system’s name, spoken as one would answer the phone.

### 6.2 The assessment

After the caller speaks, the system performs two simultaneous operations: (1) reflective listening—mirroring back what was heard to confirm understanding (Principle 16), and (2) resource search—querying the database for matching resources. The reflective listening is not an efficiency loss; it is the mechanism by which misunderstandings are caught before they produce bad referrals. If the system heard “I need food” but the caller actually said “I need boots,” the reflective confirmation (“It sounds like you’re looking for food assistance—is that right?”) catches the error immediately.

This is where the righting reflex must be resisted. The natural engineering instinct is to skip confirmation and deliver results as fast as possible. But the crisis line evidence (Mishara et al., 2007; Gould et al., 2007) shows that the assessment phase—in which the caller feels heard and the need is collaboratively clarified—is where the therapeutic value of the interaction occurs. Speed without accuracy is not efficiency; it is waste.

### 6.3 The referral

Referrals should follow a consistent structure: (1) name of the resource, (2) what it offers (in one clause), (3) address, (4) hours or availability, and (5) what the caller should expect or bring. No more than three referrals should be presented in a single turn. Each referral should include enough information for the caller to make an informed choice (AIRS Standard 1, NASW 1.03).

The last-mile problem (Boyum et al., 2016) means that providing a name and phone number is insufficient. The system should function as a navigator: “When you get there, tell them you’re looking for emergency food assistance. They’ll ask for an ID but they can help you even if you don’t have one. They’re open until 7 PM tonight.” This navigational guidance—telling the caller what to say, what to bring, and what to expect—is what raises referral success rates from 36% to 78%.

### 6.4 The follow-up

The close should accomplish two things: (1) ask whether there is anything else the caller needs (Principle 12, Principle 26), and (2) leave the door open for return. “Anything else?” is the most important question in the interaction—it is the moment at which the iceberg beneath the presenting problem may surface. “Is there anything else going on that I should know about?” asked gently, once, after the primary need has been addressed, is the single highest-leverage question in the entire conversational architecture.

### 6.5 Escalation

Graduated escalation (Principle 22) requires the system to recognize three tiers of severity: (1) *routine need*—food, shelter, legal aid, benefits navigation; the system handles directly; (2) *urgent need*—domestic violence, substance use crisis, acute housing emergency; the system provides immediate referrals to specialized hotlines (e.g., the National Domestic Violence Hotline, 988); (3) *emergency*—medical emergency, active suicidal ideation, immediate physical danger; the system says “Please hang up and call 911” and provides no further interaction.

The system should *not* screen all callers for crisis. Universal crisis screening transforms a care interaction into a surveillance interaction and violates the trauma-informed principle of not assuming trauma (Principle 20).

## 7 Discussion

### 7.1 What the methodology does not address

This paper deliberately does not address system architecture, implementation details, speech recognition model selection, natural language processing techniques, or infrastructure design. These are important engineering decisions, but they are downstream of the design methodology presented here. The choice of voice model, for example, is constrained by Principle 1 (the voice is the service) and Principle 8 (150 WPM with pauses), but the specific model is an engineering decision, not a design one. The methodology specifies *what* the system should do and *why*; it does not specify *how*.

### 7.2 The classification problem

The 211 Human Services Indexing System contains over 10,500 terms organized across ten categories and six hierarchical tiers. This taxonomy determines what the system can “see” in its resource database. Bowker and Star (1999) have shown that classification systems are never neutral—“each standard and category valorizes some point of view and silences another.” The design methodology presented here does not resolve this problem. What it does is establish Principle 15 (Open, not sorted) and Principle 17 (The caller determines the need) as design commitments that resist the taxonomy’s tendency to pre-classify the caller’s situation. The tension between the system’s need to classify (in order to search) and the caller’s right to self-determination (in order to be heard accurately) is not a problem to be solved but a condition to be navigated—call by call, interaction by interaction.

### 7.3 The question of who answers

The companion paper (Michalove, 2026) develops the argument that the question “who gets a human, and who gets a machine?” is a political question about the allocation of care (Tronto, 1993). This paper takes a different approach: it asks not whether a machine should answer but how a machine should answer *if it does*. The 26 principles constitute an attempt to encode, in the design of a conversational system, the accumulated wisdom of a century of telephone-based care—from Emma Nutt’s patience to the AIRS five-stage model to the motivational interviewing insight that “I need help” is the beginning of a conversation, not a search query.

Whether this encoding is adequate—whether principles derived from human care

practice can meaningfully inform the design of an algorithmic system—is ultimately an empirical question. The measure of the methodology is not whether the system achieves a particular score on a benchmark but whether the caller who reaches a warming center on a cold night, or finds a food pantry that is actually open, or receives a referral that includes what to say when they get there, experienced something closer to care than to information retrieval. As Jackson (2023) writes, “the measure of hope is not accuracy, but efficacy: its ability to hold and sustain more meaningful forms of action and relationality in the world.”

## 7.4 Limitations

The methodology is derived from literature rather than from empirical evaluation of a deployed system. The 26 principles have not been validated through controlled studies comparing systems that implement them against systems that do not. The literature review draws primarily on U.S.-based research and standards; the methodology may require adaptation for other cultural and institutional contexts. The design principles are, by design, general—they specify commitments rather than implementations, and reasonable designers may differ on how to realize them in practice.

Additionally, the methodology is silent on several important questions: How should the system handle repeated callers who are using it for companionship rather than resource navigation? How should it respond to callers who are intoxicated, disoriented, or otherwise unable to articulate a clear need? How should it handle callers who express frustration, anger, or hostility toward the system itself? These are questions that the evidence base does not yet adequately address, and they represent important directions for future work.

## 8 Conclusion

The design of voice AI for social service navigation is not primarily a technical challenge. It is a practice challenge—one that draws on over a century of accumulated knowledge about how to help people by telephone. The Bell System operators who perfected the “voice with a smile” in the 1880s, the 211 specialists who developed the five-stage I&R process, the crisis counselors whose behaviors have been empirically linked to reductions in distress, the social workers who formalized the ethics of the helping relationship, and the warm line volunteers who demonstrated that the most important thing about a call is who answers—all of these traditions contain design knowledge that is immediately

applicable to conversational AI systems, and all of it has been largely ignored by the field.

The 26 principles presented here are an attempt to recover this knowledge and translate it into a design methodology adequate to the phenomenon. They are not a technical specification. They are a set of commitments about what it means for a voice system to be of use to a person in need: that the voice is the service, that the caller determines the need, that empathy precedes problem-solving, that a referral is only as good as what happens after the call ends, that being heard is itself a form of help, and that the design of care infrastructure is itself an act of care (Puig de la Bellacasa, 2011).

The companion paper asks “when is a voice an infrastructure?” and answers: when it sustains, in its imperfect and always-partially-broken way, the flow of care between people in need and the resources that exist to help them. This paper asks a simpler question: *how?* The answer, it turns out, was here all along—in the operator’s greeting, in the social worker’s open question, in the crisis counselor’s reflective silence, in the warm line volunteer’s willingness to simply be present. The task is not to invent a new form of care but to translate an old one.

*Number, please?*

---

*This paper draws on the author’s experience designing a voice AI system for community resource navigation in New York City. It does not disclose system architecture, implementation details, or proprietary methods. The author thanks Malte Jung for recruiting him to Cornell and for signing off—always signing off—and Steven J. Jackson for supervision and for the hope.*

## References

- Bickmore, T. W., Caruso, L., Clough-Gorr, K., and Heeren, T. (2005). "it's just like you talk to a friend": Relational agents for older adults. volume 17, pages 711–735.
- Bowker, G. C. and Star, S. L. (1999). *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA.
- Boyum, S., Kreuter, M. W., McQueen, A., Thompson, T., and Greer, R. (2016). Getting help after calling 2-1-1: A follow-up study of referral outcomes. *Journal of Social Service Research*, 42(3):279–292.
- Bronfenbrenner, U. (1979). *The Ecology of Human Development: Experiments by Nature and Design*. Harvard University Press, Cambridge, MA.
- Commarford, P. M., Lewis, J. R., Smither, J. A.-A., and Gentzler, M. D. (2008). A comparison of broad versus deep auditory menu structures. *Human Factors*, 50(1):77–89.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1):87–114.
- Fischer, C. S. (1992). *America Calling: A Social History of the Telephone to 1940*. University of California Press, Berkeley.
- Gould, M. S., Kalafat, J., Munfakh, J. L. H., and Kleinman, M. (2007). An evaluation of crisis hotline outcomes part 2: Suicidal callers. *Suicide and Life-Threatening Behavior*, 37(3):338–352.
- Green, V. (2001). *Race on the Line: Gender, Labor, and Technology in the Bell System, 1880–1980*. Duke University Press, Durham.
- Hochschild, A. R. (1983). *The Managed Heart: Commercialization of Human Feeling*. University of California Press, Berkeley.
- Inform USA (2023). AIRS standards and quality indicators for professional information and referral. Version 10.0. Formerly Alliance of Information and Referral Systems.
- Jackson, S. J. (2023). Ordinary hope. In *Ecological Reparation: Repair, Remediation and Resurgence in Social and Environmental Conflict*, pages 417–433. Cambridge University Press, Cambridge.

- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Tober, C., Ricketts, C. R., Jurafsky, D., and Goel, S. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689.
- Michalove, J. (2026). When is a voice an infrastructure? Care, classification, and the politics of who answers. Working draft.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97.
- Miller, W. R. and Rollnick, S. (2012). *Motivational Interviewing: Helping People Change*. Guilford Press, New York, 3rd edition.
- Mishara, B. L., Chagnon, F., Daigle, M., Balan, B., Raymond, S., Marcoux, I., Bardon, C., Campbell, J. K., and Berman, A. (2007). Which helper behaviors and intervention styles are related to better short-term outcomes in telephone crisis intervention? *Suicide and Life-Threatening Behavior*, 37(3):308–321.
- National Association of Social Workers (2021). Code of ethics. Revised 2021. Available at <https://www.socialworkers.org/About/Ethics/Code-of-Ethics>.
- Puig de la Bellacasa, M. (2011). Matters of care in technoscience: Assembling neglected things. *Social Studies of Science*, 41(1):85–106.
- Raven, M. C., Kaplan, L. M., Rosenberg, M., Tieu, L., Guzman, D., and Kushel, M. (2018). Mobile phone, computer, and internet use among older homeless adults. *JMIR mHealth and uHealth*, 6(12):e10049.
- Saleebey, D. (1996). The strengths perspective in social work practice: Extensions and cautions. *Social Work*, 41(3):296–305.
- SAMHSA (2014). SAMHSA’s concept of trauma and guidance for a trauma-informed approach. Technical Report HHS Publication No. (SMA) 14-4884, Substance Abuse and Mental Health Services Administration.
- Taylor, B. L. and Minkovitz, C. S. (2021). Warm handoffs and the effectiveness of referral to social services: A review. *Maternal and Child Health Journal*, 25(4):528–541.
- Tronto, J. C. (1993). *Moral Boundaries: A Political Argument for an Ethic of Care*. Routledge, New York.

Wolters, M., Georgila, K., Moore, J. D., Logie, R. H., MacPherson, S. E., and Watson, M. (2009). Reducing working memory load in spoken dialogue systems. In *Interacting with Computers*, volume 21, pages 276–287.

Zajicek, M. (2001). Interface design for older adults. In *Proceedings of the 2001 EC/NSF Workshop on Universal Accessibility of Ubiquitous Computing*, pages 60–65. ACM.