

PERSPECTIVES ON CONNECTED CARE / 2026

Informed Consent and Shared Benefit

In AI-assisted consultations.

Why patient communication should begin before the appointment and continue afterwards. A practical model for understanding, choice and useful information on the patient's terms.

Echo Health AI

Visualisation Hub Pty Ltd trading as Echo Health AI

Revised public edition | 25 September 2026

Australian focus. Product direction and illustrative workflows; availability and deployment evidence are confirmed separately.

In this paper

Executive summary	03
1. What does the patient receive?	04
2. The Australian context	05
3. Beyond the doorway question	06
4. Information people can use	07
5. Two stages, one continuing choice	08
6. The patient's side of the exchange	09
7. Agreed actions and private reminders	10
8. The circle of care	11
9. Before the appointment: the flow	12
10. In the consultation: the flow	13
11. Practice controls and evidence	14
12. Accountability and regulatory scope	15
13. Evaluate understanding and benefit	16
Conclusion: a continuing conversation	17
References: consent and communication	18
References: patient benefit and scope	19

About this edition

This revised edition retains the supplied July 2026 paper's argument for two-stage consent, direct patient benefit and a patient-directed circle of care. Conceptual workflow diagrams replace screenshots containing unverified shipping, retention and training claims. The diagrams describe a proposed operating model, not a validated clinical deployment. This is a position paper, not a ready-to-use patient consent form or legal advice. References are numbered and linked.

01 / THE PATIENT'S SIDE OF CONNECTED CARE

Executive summary

When AI helps document a consultation, clinicians may receive a useful draft and relief from administrative work. The patient contributes the sensitive conversation from which that value is produced. A fair design should ask what the patient understands, what choices they have and what useful information they receive in return.

Consent is the precondition for the proposed use of the technology, not the benefit delivered to the patient. It should be informed, voluntary and specific to what will happen. An online checkbox is not proof of understanding, and a brief question in the consulting room is only useful if the patient has enough information and a genuine opportunity to choose.

This paper proposes two connected stages. Before the appointment, the practice explains the use of AI in plain language, makes fuller details available, invites questions and records relevant preferences. At the consultation, the clinician confirms the patient's current choice before listening or recording begins. A person who has not used the digital process must receive an equally meaningful conversation through another route.

The same workflow should deliver value after the encounter: a reviewed explanation, agreed next steps, appropriate resources and reminders that respect the patient's communication choices. Where a patient wants a family member or carer involved, the scope and duration of sharing should be explicit and changeable.

Echo One's direction connects the patient experience with clinical work and practice operations. The model described here is an argument for how that connection should work. It is not a claim that a product demonstration establishes clinical effectiveness, regulatory approval or the validity of every consent obtained through it.

The aim is a patient who understands the choice and can use the outcome of the consultation.

02 / THE PATIENT'S SIDE OF CONNECTED CARE

1. What does the patient receive?

The clinician's benefit from ambient documentation is visible: less effort turning a conversation into a note, and potentially more attention available during the encounter. Evaluations of scribe products report improvements in clinician experience and administrative burden, with results that depend on the product and setting. [1]

The patient's benefit should be equally concrete. A person may appreciate more eye contact, but still leave uncertain about what was agreed. They may not know whether an audio stream or transcript was processed by another organisation, whether it will be retained or how to correct an inaccurate summary.

An efficient note is not a sufficient answer to those concerns. The design should support both sides of the consultation: useful documentation for the professional and understandable information for the patient. The benefits need to be tested rather than assumed to follow automatically from time saved.

The proposed exchange is not a purchase of consent with better service. Patients should not have to accept optional AI processing to obtain a clear explanation, a care plan or access to the clinical information they are entitled to receive. Patient communication remains a responsibility whether AI is used or not.

Instead, the proposal is a design obligation. If technology helps organise the encounter, use that capability to support the person's understanding too. Keep the non-AI route practical. Explain the limits of the technology as carefully as its convenience. Measure whether the resulting communication helps the patient act on the agreed plan.

Distinction: permission to use a tool and the benefit of using it are separate questions. A better summary cannot repair uninformed or coerced consent.

2. The Australian context

Australian professional and privacy guidance provides a stronger basis for consent design than an assumption that patients are now familiar with AI. Ahpra's guidance says practitioners should involve patients in decisions about AI tools using their personal information, and highlights informed consent before recording consultations. [2]

The Office of the Australian Information Commissioner describes consent in terms of adequate information, voluntariness, currency and specificity, and the person's capacity to understand and communicate it. Health information also has particular protections. [3] These principles make the quality of the process matter, not merely whether a form exists.

A practice therefore needs to know what its proposed tool actually does. Does it process audio as it is spoken? Does it retain a recording or transcript? Which organisations receive the information? Where is it processed? What uses are allowed under the relevant contracts? Staff cannot provide a reliable explanation while those operational facts remain uncertain.

Recording laws and health-record requirements also differ across jurisdictions. A uniform screen can help deliver consistent information, but cannot by itself settle every legal question. Local policies and patient-facing statements need to be assessed for the service and jurisdiction in which they will be used.

The original paper described several news reports about adoption and consent disputes. This edition grounds the practical model in primary guidance and research instead of using those reports to estimate prevalence or reach a conclusion about individual providers.

Consent design begins with the facts of the proposed service. Wording must describe those facts, not the most reassuring version of them.

3. Beyond the doorway question

A short verbal confirmation is not automatically inadequate. It can be the right final step when the patient already understands the proposal, can ask questions and can decline. The problem arises when it is expected to deliver unfamiliar information, establish understanding and obtain agreement all at once.

Someone arriving for a medical appointment may be anxious, unwell or focused on what they need to tell the clinician. They may interpret a hurried request as a routine requirement rather than a choice. Familiarity with a smartphone or a consumer chatbot does not establish understanding of clinical data processing.

Research on consent for ambient documentation supports a transparent, flexible, multistep approach. Lawrence and colleagues describe preferences for information before the visit and a continuing opportunity to ask questions and make choices. [4] This supports the proposed two-stage design; it does not prove that any particular form produces valid consent.

Broader research on consent comprehension also favours interactive approaches, including discussion and checking understanding, while identifying weaknesses and variation in the evidence. [5] These studies concern medical and surgical consent more broadly, so their findings should inform design without being treated as a product validation for ambient AI.

Moving the same checkbox earlier is not enough. The additional time has value only if the information is understandable, the alternatives are real and the patient can change their mind. The clinician must be able to see when the earlier step was not completed and respond accordingly.

A useful opening is an invitation to questions and confirmation of the current choice. It must happen before recording; recording the consent conversation first would reverse that order.

4. Information people can use

The first explanation should be short enough to read and concrete enough to matter. It should name the purpose, describe what the tool receives and produces, explain the clinician's review, and say how the patient can ask questions or decline. Fuller information should be easy to find without forcing every person through dense technical text.

The detailed account needs to answer the questions that affect the decision. Identify the service and relevant providers; distinguish audio, transcripts, drafts and signed records; describe processing locations, access, retention and secondary uses; and explain how withdrawal affects future processing and information already retained.

Do not fill gaps with blanket assurances. Saying that information is secure does not describe who can access it. Saying that data is never used for training requires evidence across the actual providers and contractual chain. Saying that everything is deleted after the appointment may conflict with clinical record obligations and the system's real operation.

Understanding needs an accessible route: appropriate language, interpreters when needed, formats compatible with assistive technology and help for people who do not use digital forms. A family member's assistance must not automatically become authority to decide or receive the information.

The practice should test the explanation with the people who will use it. Can they describe what is proposed, identify their alternatives and explain how to stop? Completion rate alone is a poor measure if people accept because the easiest path appears to be agreement.

Plain language is a way to make important facts understandable. It is not a reason to leave them out.

5. Two stages, one continuing choice

The model separates preparation from confirmation while keeping them connected. The pre-visit step creates time for understanding. The consultation step establishes what the person wants now. Neither is an irrevocable permission for every later use.

01 Before the visit: explain and invite

Provide a plain-language account, fuller details and a way to ask questions. Describe the alternative to optional AI use. Record choices with the wording and version shown.

02 At the visit: confirm before starting

Check the patient's current preference, answer questions and resolve anything incomplete. Do not start recording on the strength of silence, a default selection or an old form alone.

03 During and after: honour changes

Provide a practical stop route, update preferences and explain what can change prospectively and what records must still be retained.

This is a proposed workflow, informed by the consent research and guidance discussed earlier. The point is not to require a smartphone. If the person has not completed the first stage, staff need time and a suitable alternative to provide the information and support a decision.

The system should distinguish a refusal, an unanswered item, a withdrawn choice and an expired or superseded statement. Treating all of these as simply missing consent can conceal what the patient actually said. Treating them as permission is worse.

Do not bundle the choices: AI documentation, reminder content and sharing with a support person serve different purposes. Agreeing to one should not silently authorise the others.

6. The patient's side of the exchange

Patients already seek ways to revisit clinical conversations. An Australian online survey by Prictor, Elwyn and Hyatt explored consultation recording and described motives including understanding, remembering and sharing information for care. [6] This is evidence of a need, not permission to assume every patient wants a recording or AI processing.

A useful patient-facing outcome starts with the agreed plan. What was discussed? What has the clinician decided? What does the patient need to do, and when? Who should they contact with questions? The answer should be tailored to the encounter and reviewed before release.

AI can help prepare simpler language, but simpler language can still be wrong. Research on transforming discharge summaries illustrates both the potential and the need to check completeness and accuracy. [7] A fluent rewrite must not change a dose, certainty, timeframe or responsibility, or remove a qualification that matters.

The patient should also be able to identify what is a confirmed plan, a proposed next step and general information. An approved resource can support understanding, but should not be presented as individualised clinical advice merely because it appears beside the visit summary.

The design goal is a useful record of the encounter that the patient can return to and discuss. Its quality should be tested with patients, including people with different language, literacy and accessibility needs. Clinician approval and patient comprehension are separate measures; both matter.

A summary is successful when it helps the person understand the reviewed plan, not merely when it is shorter than the clinical note.

7. Agreed actions and private reminders

Reminders can help a person complete an agreed action, but message delivery is not the same as care completion. Reviews of digital notifications report improved clinic attendance in studied settings. [8] That does not establish the effect of every reminder design, or justify disclosing more detail than the recipient wants.

A message saying that something is due exposes less than one naming a test or condition. Both can be useful in different circumstances. The patient should understand the available options and choose a channel and level of detail suitable for their situation.

The design needs to account for shared phones, changed numbers, lock-screen previews and messages seen by other people. Identity and destination checks matter before sensitive content is sent. A delivery receipt should not be interpreted as evidence that the intended person read or understood the information.

Reminders should reflect a reviewed plan or an approved practice process. The AI must not independently invent a clinical action and send it to the patient. If there is uncertainty about the plan, the destination or the permission to share, the workflow needs a visible exception and a responsible person to resolve it.

Withdrawal and changes need to reach queued communication as well as the preference screen. A system that displays a changed choice but still sends previously queued content has not honoured the choice in practice. Testing should follow the message to its destination and include failures and retries.

Illustrative choice: a neutral notification or more specific content through an agreed channel. The available options must reflect the actual service, with its limitations explained.

8. The circle of care

Some patients want a partner, family member or carer to help with the practical work of care. This paper uses circle of care to describe patient-directed sharing with those people. It does not give the recipient clinical authority, unrestricted record access or the right to make decisions for the patient.

The patient should choose the person and the categories of information to share. A summary, an action list and a reminder may have different sensitivity and purpose. The destination and recipient's willingness to receive the communication need to be established before sending.

The arrangement must be understandable and changeable. State how long it lasts, what is excluded, how the patient can stop it and what a recipient can do if they no longer wish to receive messages. Explain that stopping future sharing cannot retrieve information already delivered.

A private-looking digital form is not necessarily private. A person may share a device, depend on someone for access or be under pressure at home. Staff need a discreet way to confirm the arrangement with the patient and address concern. Do not assume that nomination from a home device proves an unpressured choice.

Capacity, guardianship and substitute decision-making require separate assessment. A named contact is not proof of legal authority. OAIC guidance also distinguishes the circumstances in which health information may be disclosed, including particular situations involving impaired capacity. [3, 9]

Make support possible without turning a helpful relationship into unrestricted access.

All examples in this paper are conceptual. No real patient, carer or practice record is reproduced.

9. Before the appointment: the flow

The following sequence replaces the earlier paper's product screenshots with an explanatory design. It shows the intended responsibilities without implying that the pictured controls have been accepted for a clinical deployment.

01 Reach the right person

Send through an agreed channel. Limit what a message reveals and use an appropriate identity process before displaying sensitive details. A date of birth alone is not a guarantee of identity.

02 Explain why the practice is asking

Invite the reason for the visit in the patient's words. Make clear who will receive it, when it is reviewed and where urgent clinical needs should be directed.

03 Offer separate, understandable choices

Explain AI use, communication preferences and any proposed support-person sharing independently. Provide fuller details, questions and alternatives.

04 Review before submission

Let the patient see and change the choices. Preserve the wording, version, scope and time associated with each response.

Not every input needs a model. Identity checks, expiry conditions and consent-state handling belong to explicit rules. A model may help structure free text for a clinician, but it cannot turn an unanswered question into agreement or an uncertain identity into permission.

Missing, expired or incomplete intake should remain visible as such. The practice needs an assisted route, not a forced acceptance path.

10. In the consultation: the flow

The clinician needs a clear account of what the patient has been told and what they have chosen, together with any unresolved questions. The interface should help the clinician distinguish a current choice from a historical one.

01 Review the patient's preparation

Show the patient's own words, the status of each choice and the statement version. Keep model-generated summaries identifiable and linked to their source.

02 Confirm before listening begins

Ask whether the person remains comfortable, invite questions and record the current decision. If they decline optional AI documentation, use the agreed alternative.

03 Review the clinical output

Check the draft against the encounter and relevant records. Correct or reject it before signing or authorising a patient-facing version.

04 Agree the next step

Confirm actions, responsibility and communication preferences. Release reviewed information only to the intended, authorised recipients.

The review must be capable of changing the outcome. A notice saying that AI can make mistakes does not compensate for a workflow that makes correction difficult, hides source information or encourages hurried acceptance.

A draft can influence judgement before it is saved. Evaluation must therefore examine the content and its effect on the reviewer, not only whether the final write requires a click.

11. Practice controls and evidence

A practice needs more than editable wording. It needs a process for approving statements, keeping versions and ensuring that the live configuration matches what patients were told. Changes in providers, processing, purpose or retention can require a fresh explanation and review of the applicable consent arrangements.

A versioned record helps establish what statement was shown and what response was recorded. It is useful evidence of the interaction, but cannot establish understanding or voluntariness on its own. Staff should have a way to record questions, assisted communication and changes of mind without burying the relevant choice in free text.

AI documentation consent can be a required item for the practice to present without agreement becoming compulsory for the patient. These are different meanings of required. The design must preserve a refusal path and a workable alternative for optional AI use.

Configuration should cover identity attempts, link expiry, access roles, message sending windows and the handling of incomplete forms. Each setting needs an operational owner and a tested effect. A displayed retention field does not prove that downstream records, backups and provider copies are deleted on that schedule.

Before deployment, obtain evidence for the actual audio, transcript, draft and record handling. Confirm model-provider terms, secondary uses, locations, retention and deletion processes. This paper does not establish a blanket no-training promise or a universal deletion period for Echo One.

A consent statement and a running system must describe the same service.

12. Accountability and regulatory scope

AI assistance does not remove the need for professional judgement. The treating clinician must review the clinical material they rely on and the information they authorise for the patient. Practices and suppliers also retain responsibilities for their own policies, design, data handling and service operation.

The intended architecture separates a model's proposal from authority to act. Typed validation and deterministic checks sit between output and a permitted write, task or message. Clinical review adds a different layer: a structurally valid draft can still be incomplete, misleading or clinically wrong.

Patient benefit does not determine regulatory status. A plain-language summary, a code, a flag or a reminder needs assessment according to its intended purpose and the applicable rules. The TGA's guidance makes that intended purpose central to software regulation. Human review does not create a blanket exemption. [10]

The same distinction applies to privacy. Consent is one part of lawful handling; it does not make an unnecessary disclosure or insecure process acceptable. Information about the actual service must be accurate, and the practice needs procedures for access, correction, complaints and incidents.

Echo One is presented publicly as a product preview. This paper neither certifies a deployment nor replaces the assessment of a health service's local obligations. The workflows describe how responsibility should remain visible as information moves from patient input to clinical review and onward communication.

Implementation question: who can change, approve, stop and investigate each part of the workflow, and what evidence shows that those controls work?

13. Evaluate understanding and benefit

The model deserves evaluation on both sides of the consultation. For consent, test whether people understand the proposed processing, can find the detailed information and can decline or withdraw without being lost in the workflow. For patient benefit, test whether the reviewed summary and agreed actions are understandable and useful.

Measure what happens to people who do not complete digital intake. Examine language, disability, literacy, age and access barriers. A high digital completion rate can conceal exclusion if the people who struggle are simply absent from the denominator.

Operational evaluation should include statement changes, stale choices, wrong destinations, revoked sharing, failed messages and correction of inaccurate information. Follow the state through to the actual downstream process. Do not stop at a screen that looks correct.

Clinical evaluation needs representative cases, meaningful professional review and attention to omissions as well as invented content. Record the time spent correcting and explaining outputs. A workflow that transfers hidden work to clinicians or patients may not deliver the benefit it appears to promise.

The earlier paper referred to proprietary Echo research. Those instruments and their methods were not established as a published evidence base for this edition, so they are not used to support clinical or population claims here. The proposal rests on an explicit design argument, the cited literature and the need for local evaluation.

Count understanding, usable information and completed care processes. Do not mistake acceptance clicks or messages sent for those outcomes.

Conclusion: a continuing conversation

The question behind this paper is simple: when AI helps a health service document care, what does the patient receive? The answer should be more than a request to permit recording at the moment the consultation begins.

A better process provides information early, makes choice practical and confirms it before the technology starts. It gives the person opportunities to ask questions and change their mind. It recognises that an unanswered form is not refusal, agreement or evidence of understanding, and that digital convenience must not narrow access to care.

After the encounter, a reviewed explanation and agreed next steps can make the conversation more useful. Communication preferences should govern the channel and detail. A support person should receive only what the patient has authorised within an appropriate arrangement, with a practical route to stop future sharing.

These are design commitments to evaluate. A product screen, a version number or an AI-generated summary does not demonstrate that the commitments have been met. The practice and supplier need evidence that the actual system, the information provided and the operational process agree.

Echo One's connected-care direction is an opportunity to put that discipline into the patient journey from the beginning. The measure of success is not that the AI has listened. It is that the person has been heard, understands the choice and can use the reviewed outcome of their care.

AI should help make care understandable and actionable for the patient.

Read the companion paper, *After the Scribe*, at echo-one.health/resources. The two papers connect the workflow argument with the patient's experience of it.

References: consent and communication

[1] **Olson KD et al. (2025).** Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout. JAMA Network Open, 8(10), e2534976. Evidence from evaluated scribe use; not an Echo One outcome study. [Read the study.](#)

[2] **Ahpra and National Boards.** Artificial Intelligence in healthcare. Professional guidance including informed consent, privacy and practitioner responsibilities. [Read the guidance.](#)

[3] **Office of the Australian Information Commissioner.** Guide to Health Privacy: Introduction and key concepts. Describes health information and the elements of consent. [Read the guidance.](#)

[4] **Lawrence K et al. (2025).** Informed Consent for Ambient Documentation Using Generative AI in Ambulatory Care. JAMA Network Open, 8(7), e2522400. Research informing the proposed multistep process; not proof of the validity of a particular consent form. [DOI: 10.1001/jamanetworkopen.2025.22400.](#)

[5] **Glaser J et al. (2020).** Interventions to Improve Patient Comprehension in Informed Consent for Medical and Surgical Procedures: An Updated Systematic Review. Medical Decision Making, 40(2), 119-143. Broader consent evidence, with limitations in measures and populations. [DOI: 10.1177/0272989X19896348.](#)

Sources support the statements associated with them. The two-stage model, operational examples and evaluation questions are Echo Health AI's synthesis and proposal, not a validated legal template.

References: patient benefit and scope

[6] **Prictor M, Elwyn G, Hyatt A (2024)**. How often are patients recording their healthcare consultations in Australia and why? An online survey. *PEC Innovation*, 5, 100355. [DOI: 10.1016/j.pecinn.2024.100355](https://doi.org/10.1016/j.pecinn.2024.100355).

[7] **Zaretsky J et al. (2024)**. Generative Artificial Intelligence to Transform Inpatient Discharge Summaries to Patient-Friendly Language and Format. *JAMA Network Open*, 7(3), e240357. Evidence about draft transformation, with accuracy and completeness concerns. [Read the study](#).

[8] **Robotham D et al. (2016)**. Using digital notifications to improve attendance in clinic: systematic review and meta-analysis. *BMJ Open*, 6, e012116. [DOI: 10.1136/bmjopen-2016-012116](https://doi.org/10.1136/bmjopen-2016-012116).

[9] **Office of the Australian Information Commissioner**. Guide to Health Privacy: Chapter 7, Disclosing information about patients with impaired capacity. [Read the guidance](#).

[10] **Therapeutic Goods Administration**. Understanding how we regulate software-based medical devices. Intended purpose and regulatory scope. [Read the guidance](#).

Publication details

Visualisation Hub Pty Ltd trading as Echo Health AI. Revised public edition, 25 September 2026. Guidance checked for this edition. This is a position paper, not a patient consent form, legal opinion or claim of clinical validation. Conceptual diagrams replace the earlier screenshots; no real patient information is included.

Website: echo-one.health. Privacy and security contact: security@echo-health.ai. Copyright 2026 Echo Health AI.