Peel back

Can AI help health research reach the people who need it?

Findings from a discovery study with 152 researchers, patients, clinicians and decision-makers

Funded by Innovate UK. Delivered in collaboration with the University of Exeter.

July 2026

Foreword

Three years ago, someone very close to me fell ill with cancer. Like so many people in that situation, I found myself googling late into the night, wading through pages that were either too vague to help or too technical to understand. Somewhere out there was research that mattered to us.

A few years before that, I had joined the world of academia, and I was constantly amazed by the wealth of knowledge being produced there: careful, rigorous, hard-won research on almost everything, including the questions I was typing into a search bar at midnight. I couldn't fathom why the two didn't connect better.

I'm a science communicator, and I run a science communication agency that helps researchers and companies share their work. So I saw the other side of the problem too: researchers who cared deeply about their findings reaching people, and simply had no time, tools or support to make it happen.

Rather than assume I understood why, I decided to investigate it properly. With funding from Innovate UK and in collaboration with the University of Exeter, we spent six months asking 152 people on both sides of the gap: the researchers who produce health research, and the patients, clinicians, commissioners and policymakers who need it.

I went into this project thinking about simplification, about making complex research easier to read. I came out thinking about translation: helping researchers communicate the same finding differently to different audiences, for different purposes. That shift runs through everything in this report.

These findings are now shaping what we're building at Peelback. But we're sharing them openly because the problem is bigger than any one tool, and because the people who gave us their time, their frustrations and their hopes deserve to have what they told us heard.

Emilie Fontaine

Founder, Peelback

01 About this research

Between November 2025 and March 2026, Peelback and the University of Exeter ran a discovery project funded by Innovate UK to investigate a question both sides of the research system recognise: why does so much health research never reach the people who could use it?

We asked the people who produce research and the people who need it. 152 participants took part across six research activities: four qualitative workshops and two quantitative surveys, spanning health and medical researchers on one side, and patients, carers, clinicians, policymakers, commissioners and industry professionals on the other. Full details of the participants and methods are at the end of this report.

This report presents what they told us.

84%

of researchers agreed that who the audience is, and what the communication is for, should shape the output.

(n = 50)

73%

of researchers cite a lack of time as a barrier to sharing their work.

(n = 50)

77%

of decision-makers ranked researcher verification as the most important feature.

(n = 80)

02 The dissemination problem

Across all six research activities, a consistent picture emerged: health research dissemination is structurally under-supported on both sides.

Researchers want their work to reach beyond the journal. Every researcher we surveyed had attempted dissemination beyond academic publication. But 73% cite a lack of time as a significant barrier, a third have no dissemination resources at all, and only 16% have communications support. Nearly half spend four or more hours producing communication for a single paper, and that effort typically reaches one audience, when the audiences they want to reach include patients, clinicians, policymakers and commissioners.

The people who need research describe the same gap from the other side.

> *"Journal articles are written for an academic audience, even when they are discussing something that is potentially applicable to the wider NHS."*

The problem, in short, is not a shortage of research. It is that research is produced in one format, for one audience, and everyone else is left to translate it themselves or go without.

53%

say research papers are too technical for their needs.

(n = 80)

44%

say the research they find is rarely actionable.

(n = 80)

41%

say they waste time translating research into a usable format themselves.

(n = 80)

03 What researchers told us

Dissemination is wanted, but unsupported

Every researcher in our survey sample had tried to share their work beyond academic publication. A third do so regularly as part of their role, 43% occasionally, and a quarter have managed it only once or twice. The intent is there; the structure is not. Participants in the researchers workshop described dissemination as "audience-dependent and tailored" but noted that tailoring content is specialist work that is rarely available within research teams. The default is what the system rewards: peer-reviewed publication and the occasional press release, formats that do not reach most of the audiences researchers want to engage.

One researcher's post-it note, asked about the immediate outcome of disseminating their work, read simply:

> *"to me, the uni, the world: Nothing."*

That is not apathy. These are people who care deeply about their work reaching the right audiences. This is not a motivation problem; it is a structural one.

The audience is the point

The strongest agreement in the entire research programme came on a single idea: 84% of researchers agreed that who the audience is, and what the communication is for, should shape the output. A finding written for a patient is not a shorter version of the paper; it is a different piece of communication with a different job. This distinction, translation rather than simplification, was endorsed independently by researchers, by patients and by decision-makers, and it sits at the centre of this report's conclusions.

What researchers want to produce

Asked which output formats matter most, researchers prioritised infographics (92%), lay summaries (88%) and slide decks (65%), with video summaries valued by around half. These are, notably, the formats their audiences say they rarely receive.

Reach is the desire; distortion is the fear

The workshops surfaced an emotional dimension that the survey numbers alone would miss. Researchers spoke about wanting their work to be seen and used, and about the disappointment of research that lands nowhere, particularly for the patient groups who contributed to it. Alongside that desire sat a consistent fear: that communication produced quickly, or by someone else, or by a machine, might misrepresent the findings. Any approach to research communication that does not answer that fear is unlikely to be adopted.

04 What decision-makers told us

An audience that already runs on research

This is not an audience that needs persuading to use evidence. 76% need to find, interpret or apply health research at least weekly; 44% do so daily. They work from specialist databases, commissioned evidence reviews and direct contact with researchers. A quarter already use AI tools to help.

The formats they receive are not the formats they need

76% most commonly receive full academic papers. Asked what would actually be useful, they chose evidence briefs (80%), plain-language summaries (65%) and policy summaries (62%). Yet plain-language summaries are currently received by just 10% of them, and infographics by 8%. The gap between what researchers produce and what decision-makers can use is one of the clearest findings in this research.

These are not people who need research explained to them. They read a paper, assess whether the methodology holds up, and then spend hours rewriting it into a format their team can act on. That rewriting is invisible, manual, and repeated across every organisation that uses health research.

> *"It is difficult to translate health research into real-life policy decisions and advice that are practical, workable and have a positive impact on public health."*

Their biggest frustrations, in their own words, cluster around access and paywalls, time, the labour of translating academic outputs into briefings and summaries, and academic language itself.

Trust is the condition of adoption

When we asked what would need to be true for them to rely on AI-assisted research summaries, one answer dominated: verification by the researcher. 77% ranked researcher verification as the most important feature, and it was the single most cited requirement in the open-text responses (36%). Concern about accuracy or misrepresentation of findings was the leading barrier, raised by 79%.

> *"Trying to use AI for health research is a nightmare because it so often hallucinates things. Having a verified status would be very reassuring."*

A notable minority (32%) said they would distrust AI-generated summaries even if verified, and some would always want a route back to the original paper. Trust in AI alone is low. The human verification layer, and clear traceability to the source publication, appear to be the conditions on which everything else depends.

How they judge robustness

Asked how they assess whether research is strong enough to act on, decision-makers pointed to sample size (60%), study design (54%), statistical significance (50%) and peer review status (41%). Research communication that hides these details behind a narrative loses this audience; they need the robustness signals visible.

05 What patients, carers and clinicians told us

Research arrives in emotionally charged moments

Patients and carers do not encounter research neutrally. They meet it at diagnosis, during uncertainty, while searching for hope or answers. Participants described feeling empowered ("knowledge is power") and motivated, but also overwhelmed, confused and frustrated when the format or timing was wrong. One participant described the experience of finding research at exactly the moment it mattered:

> *"Too much text, worrying and technical language, no summary, no results."*

The same information can strengthen or alienate depending on how, and when, it is communicated. Patients do not need less research. They need research that arrives in a form they can use, when they need it.

The gap between good and bad communication is respect, not reading level

The strongest negative reactions in the workshop were not about difficult vocabulary. They came from being excluded from feedback loops, treated as a tick-box exercise, or presented with content that showed no awareness of the reader. One PPIE contributor described being handed academic documents throughout multi-year projects and having to insist on plain-language updates. Communication that signals care and accountability was valued far above communication that is merely easy to read.

Where content lives signals who it is for

Participants trusted NHS and NIHR branding, and read university websites as institutional spaces not designed for them. Public spaces, from GP surgeries to libraries, signal genuine public intent. The channel is part of the message.

AI as capacity-builder, not shortcut

The workshop produced one of the most important framings in the entire project. Participants drew a clear line: if AI frees researchers to do better engagement, it is welcome; if it frees researchers to stop caring, it is counterproductive. AI in research communication earns its place only as an enabler of human connection, not a replacement for it.

Clinicians, discussing research within clinical culture, valued journal clubs, conferences and professional publications as trusted routes, raised concerns about AI errors and wrong citations, and stressed the need for rigour: whether samples apply to their patients, whether research is representative, and whether limitations are visible.

06 Where all sides agree

Six research activities, two sides of the research system, and five points of convergence.

Translation, not simplification.

Researchers (84%), decision-makers and patients independently endorsed the same idea: different audiences do not need a dumbed-down version of the same text. They need the same finding genuinely reframed for who is reading it and what they need it for. A physiotherapy paper shared with a patient to help them manage their condition is a completely different thing from the same paper shared with a commissioner to make the case for funding. Same research. Different audience. Different purpose. Different output.

Verification is the trust architecture.

Researchers fear their findings being distorted; decision-makers will not act on AI-assisted summaries without a researcher's sign-off; patients want communication that shows accountability. One mechanism answers all three: the researcher who did the work confirms that each version is faithful to the findings, visibly.

AI earns its place as a capacity-builder.

Patients drew the line precisely: if AI frees researchers to engage more effectively, it is welcome; if it frees them to disengage, it is not. AI is not a replacement for care. Decision-makers drew a parallel line with verification. Both sides will accept AI in research communication only where the human relationship and the human accountability remain visible.

The format mismatch is structural.

What researchers produce (papers, press releases) and what their audiences need (evidence briefs, plain-language summaries, policy summaries) barely overlap, and neither side has the time or support to close the gap alone.

Time is the binding constraint.

73% of researchers cite a lack of time as a barrier to sharing their work. 41% of decision-makers waste time translating what they receive. The same hours are being lost on both sides of the same gap.

07 Summary of findings

1. **Health research dissemination is structurally under-supported on both sides of the system.** Researchers lack time, tools and institutional backing; the professionals who need research face access barriers, translation burdens and formats that do not match their needs.

2. **Audience-aware translation is the core validated concept.** The distinction between audience and purpose received the strongest agreement of any idea tested (84% of researchers), and the framing of translation rather than simplification was endorsed by decision-makers and reinforced by patient and PPIE feedback about respect and care in communication.

3. **Researcher verification is the condition of trust.** 77% of decision-makers ranked it as the most important feature of AI-assisted research communication, and it was the most cited requirement in their own words. Trust in AI alone is low; the human verification layer appears essential.

4. **Output needs diverge by audience.** Researchers prioritise infographics and lay summaries; policy and commissioning professionals need evidence briefs and policy summaries; patients need accessible, layered formats delivered with care. One format cannot serve them all, which is precisely the problem.

5. **Robustness must be visible.** Decision-makers judge research by sample size, study design, statistical significance and peer review status, and need these signals surfaced, not buried.

6. **The emotional dimension is not peripheral.** Researchers' fear of distortion, patients' experience of research arriving in difficult moments, and the shared frustration of effort that lands nowhere shape whether communication succeeds as much as its accuracy does.

Taken together, these findings answer the question on the cover of this report. Can AI help health research reach the people who need it? Yes, conditionally: when it translates rather than simplifies, and when the researcher's verification and human accountability remain visible.

08 Participants and method

The discovery project comprised six research activities conducted between November 2025 and March 2026, designed to capture both the supply side (researchers who create and disseminate research) and the demand side (the patients, clinicians, policy, commissioning and industry professionals who need to find and use it).

Activity	Side	Method	Sample	Date
In-person researchers workshop	Supply	Qualitative	n = 10	November 2025
In-person audiences workshop (patients, carers, PPIE, clinicians)	Demand	Qualitative	n = 7	December 2025
Online researchers workshop	Supply	Qualitative	n = 3	November 2025
Online policy and industry workshop	Demand	Qualitative	n = 3	February 2026
Researcher survey	Supply	Quantitative	n = 50	March 2026
Policy and industry survey	Demand	Quantitative	n = 80	March 2026

The workshops used facilitated discussion, journey mapping, post-it exercises and concept review, and were recruited through the University of Exeter's networks and existing PPIE relationships. The surveys were administered online, with respondents recruited via Prolific and screened for relevant roles: the researcher survey required health or medical research roles with published work, and the policy and industry survey required health policy, commissioning or industry roles (all UK-based). Respondents who did not meet the screening criteria were excluded from analysis.

This was a discovery study. The samples are appropriate for exploratory research of this kind, but they are not large, and the findings should be read as directional rather than definitive: consistent signals across six activities and two sides of the research system, not a final word. We intend to keep learning about this topic as Peelback develops. If you are a researcher working on research communication and dissemination, or you would like to build on these findings, we would be glad to hear from you at hello@peelback.ai.

We are grateful to the 152 people who gave this project their time, their experience and their honesty.

All quotations in this report are drawn verbatim from workshop discussions and open-text survey responses, and are reproduced with participants' consent under the project's research protocols.

Peel back

ABOUT PEELBACK

An AI platform for research communication

Peelback is an AI platform for research communication, built in response to these findings. The AI platform generates the translation, the researcher verifies it, and the decision-maker receives something they can rely on.

This research was conducted as part of an Innovate UK funded discovery project (Ref: 10165051), delivered by Peelback in collaboration with the University of Exeter. Research activities were conducted by Emilie Fontaine (Peelback) with Dr Fay Manning (University of Exeter) co-facilitating the workshops.

peelback.ai · hello@peelback.ai

UKRI Innovate UK

University of Exeter