



From Raw Data to Trusted Insights

A Practical Guide to AI in User Research Analysis

condens

Table of Contents

3	Research Analysis, Reimagined
5	What Defines Good Research Analysis
10	The Current State of AI in User Research Analysis
21	Opportunities and Risks of AI in User Research Analysis
27	A Guide to Evaluating AI Tools and Their Safeguards
31	Do's and Don'ts Checklist
34	A Practical Guide to Using AI in Research Analysis
45	How AI Changed the Way Research Works: A Case Study with Michaël Dufranne
48	Beyond Analysis: Scaling Access to Research with AI
54	Looking Ahead: The Future of Research Analysis
55	About Condens

Research Analysis, Reimagined

Research analysis has always been the most rewarding part of user research. It's where the real thinking happens, when patterns start emerging, and when those “A-ha!” moments start to happen.

The biggest challenge with analysis has always been time constraints. And with growing pressure in the industry to conduct research and achieve results faster than ever, it can feel like there just isn't enough time for it.

AI is changing that. Not by replacing the thinking, but by taking on tasks like transcribing sessions, sorting through data, and surfacing patterns. What used to take days can now happen in hours, freeing up more time for the parts of analysis that actually require your expertise.

This guide provides an overview of how to get there, covers what AI can genuinely help with, and offers guidance on where to stay critical, as well as how to build a workflow that both speeds up the analysis process and maintains reliability. It draws on input from researchers and experts, including Carl Pearson, Llewyn Paine, Laura Klein, and Nick Babich, grounded in real-world practice rather than hype, to help you get genuine value from it.

01

CHAPTER

What Defines Good Research Analysis

WHAT'S IN THIS CHAPTER

The qualities that make analysis trustworthy, and how to maintain them as AI enters the workflow.

01

What Defines Good Research Analysis

Before diving into AI and how it's changing research analysis, it helps to be clear on what actually defines good analysis in the first place.

Good research analysis turns raw data into meaningful insights. It goes beyond describing what happened to explaining why it happened, and it backs findings with evidence rather than assumptions. At its core, it comes down to one thing: generating insights that are trustworthy enough to inform better decisions.

What Makes Analysis High-Quality

Among others, these eight key factors define good research analysis:

- 1 **Methodological Rigor:** choosing the right approach for the question at hand, whether that's qualitative thematic analysis or quantitative methods, and following established protocols.
- 2 **Objectivity:** analyzing data with an open mind rather than to confirm what you already believe.
- 3 **Systematic processing:** working through data in a structured way and ensuring all relevant evidence is considered, not just the parts that support a hypothesis.
- 4 **Interpretive depth:** going beyond describing what happened ("10 participants said X") to synthesizing what it means ("the pattern of X points to an underlying issue of Y").
- 5 **Reproducibility:** documenting your process clearly enough that another researcher could follow the same steps on the same data and reach the same conclusions.
- 6 **Traceability:** Maintaining a clear link between insights and the underlying evidence, so findings hold up when questioned later.
- 7 **Contextualization and critical thinking:** Acknowledging existing literature, identifying gaps, and building on prior work, while also being honest about the study's scope, sample size, potential gaps, and what the data simply cannot tell you.
- 8 **Ethical integrity:** It adheres to ethical standards regarding participant privacy, confidentiality, and honest reporting, avoiding falsification or fabrication of data.

In practice, you rarely get to optimize for all eight qualities at once. The resources available to you are shaped by the constraints around you: time, people, tools, and organizational context. There is no universal ideal ratio. But there are principles to help you find the right balance for your situation. One that can be helpful is Minimum Viable Rigor (MVR), introduced by researcher [Carl Pearson](#).

Minimum Viable Rigor

Not every research project needs maximum rigor. But every research project should have a baseline level of rigor. Carl Pearson calls this [Minimum Viable Rigor \(MVR\)](#):

$$\text{MVR} = \text{Rigor of Insight} - \text{Decision Risk}$$

This formula isn't an exact mathematical one. Instead it helps portray a judgment call that weighs the rigor of your research against the cost of getting it wrong. The riskier the decision your research is meant to inform, the higher the baseline level of rigor should be. Lower-stakes decisions can tolerate more shortcuts while high-stakes ones should involve more rigor.

A few things make this especially relevant when adding AI to the mix. First, AI can produce a summary or a cluster of themes that look rigorous but aren't. And unlike bad code, bad analysis does not throw up an error notice. It produces something that resembles a useful insight, which can be more dangerous than no insight at all.

Second, rigor is about preserving what Carl Pearson calls the fragility of truth. A good method excavates meaning without distorting it. And while AI is fast, it can flatten nuance, miss irony, and overgeneralize from limited data. So the question to ask isn't "did AI produce an insight?" but also, "does this insight hold up well enough to inform this specific decision?"

You need to be the one to answer that question.

Why Human Judgment Is Central

AI can process language at scale and surface patterns across large datasets quickly. What it can't do on its own is recognize what's meaningful in context.

A single hesitant remark from one participant can open an entirely new line of inquiry. A subtle observation can change the direction of a study. Catching those moments is what separates analysis from data processing, and it's what researchers are uniquely equipped to do.

Research backs this up. A [study](#) on how reliably ChatGPT 4-Turbo codes qualitative data found that AI and human coders agreed on much of the data, and AI even surfaced themes that manual tagging missed. But AI struggled with depth, context, and subtle nuances, and it also introduced biases that human coders were better equipped to catch.

The conclusion reached in that study is the same one running through this guide: the most effective approach is augmented research, combining AI's efficiency with human expertise rather than replacing one with the other.

“
The ability of somebody to sift through [AI-generated output] based on their past experience, their knowledge of the project, their knowledge of design, and what they've learned is what helps make analysis a useful endeavor and makes the outputs useful to teams.
”



Maria Rosala
Director of Research
at Nielsen Norman Group

The Role of the Researcher as Sense-Maker and Decision Partner

You know that experience when things start to click during analysis? When the pattern becomes clear, and you know you've got something? That moment of recognition is part of what makes research valuable, and it's something AI can actively support, but not replace.

As AI gets better at delivering results fast, the researcher's value shifts toward everything speed alone cannot deliver: judgment, interpretation, and the ability to turn data into decisions that actually stick. The most important thing a researcher does isn't processing information. It's making sense of it, in context, for a specific audience, at a specific moment in time.

AI amplifies the skills of the person using it. So knowing what to ask, what to ignore, and when AI output doesn't pass the smell test is where a researcher's experience can make a big difference.

AI needs your expertise to work well

Whether you're reviewing a summary, evaluating a cluster, or working from a template, you bring the judgment to assess whether the output is useful, unbiased, and suited to your research goals. Without that, even a polished AI output can quietly lead you in the wrong direction.

Use AI as a starting point, not a final answer

Like any tool that works at speed and scale, it can miss context, misread tone, or generate something that looks right but isn't. That's why a review step can't be a nice-to-have, it needs to be considered an essential part of the analysis process.

Research impact depends on human connection

Insights don't drive decisions on their own. Researchers are the “human connectors” between research and decision-makers. In her book, *Just Enough Research*, Erika Hall consistently frames analysis as a collaborative activity for a reason. And it's not only to benefit from different perspectives. Working through findings together builds a shared clarity that's hard to achieve any other way. That shared understanding is what helps turn a finding into a decision, and that's not something AI can't do.

02

CHAPTER

The Current State of AI in Research

WHAT'S IN THIS CHAPTER

What 330+ researchers told us about how they use AI, where it helps, and where trust still falls short.

02

The Current State of AI in User Research Analysis

Theory only gets you so far. Before getting into the more practical part of this guide, we wanted to understand how research practitioners are actually using AI in their work today, where it's proving useful, and where doubts remain.

To find out, we surveyed over 330 UX Researchers, ReOps Managers, Designers, and Product Managers on the current state of AI in user research analysis. The results show a field that's actively adopting AI, but not without friction. Usage is high, expectations are evolving, and skepticism remains. Indicating that the reality sits somewhere between enthusiasm and caution.

Who we heard from

332 people completed the survey, and the sample skews toward hands-on practitioners. Researchers made up 63% of respondents, followed by Designers (15.4%), Product Managers (9.3%), and ReOps Managers (8.7%).

Company sizes were fairly evenly spread across small (1–200 employees, 39%), mid-market (201–5,000, 35%), and enterprise (5,001+, 26%). Most respondents work in Technology and Software (39.1%), Consulting and Agency (15.1%), or E-commerce and Retail (13.4%).


One thing worth noting: 73.5% interact with research daily or weekly. These are not passive observers of AI in research. They are the people doing the work, forming real opinions, and learning from experience.

AI Has Reached Mainstream Adoption

AI is no longer experimental. 91% of respondents use AI in some form, and for more than half it's already a core part of their workflow. Only a small minority have tried it and stopped using it.

AI usage status

% of all respondents (n=332)



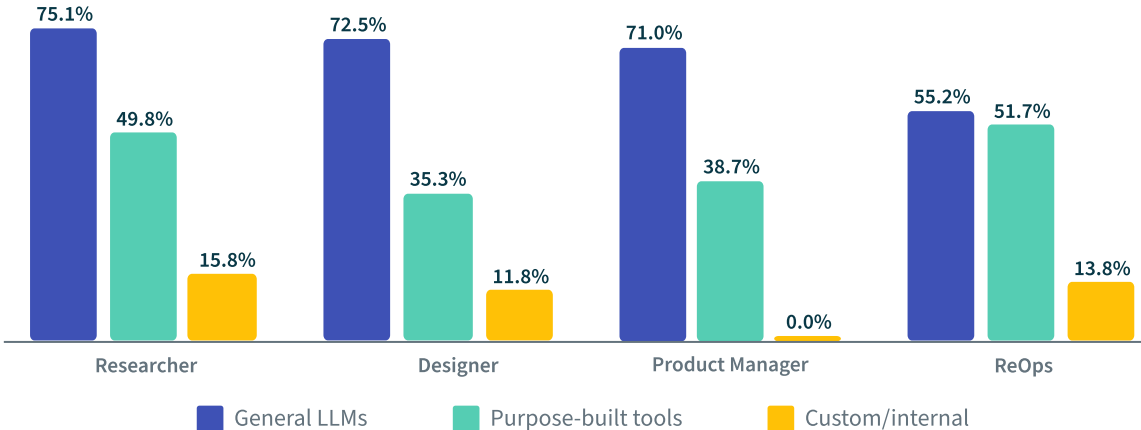
This is not a story about early adopters anymore. AI in research analysis has reached mainstream adoption. Usage, however, is not uniform. Product Managers are the heaviest users of AI, with 68% describing AI as core to their work. Researchers, Designers, and ReOps follow at around 55%. Company size matters less than you might expect: the smallest companies (1-50) show the highest core rate at 61%, while large enterprises sit at 53%.

72.6% of respondents use ChatGPT, Claude, Gemini, or similar. While purpose-built research tools like Condens, Dovetail, and Marvin sit at 46.7%. This 26-point gap signals something important: a substantial share of AI use in research is happening outside dedicated tooling, with real implications for traceability, context, and quality control.

Many respondents also said that they actually use both. With nearly half combining the use of a general LLM for open-ended tasks with a purpose-built tool for structured analysis.

Tool usage by role

% of each role's respondents using each tool category (multi-select)




AI Is Embedded in the Workflow and Trust Is Growing

AI use clusters tightly around core analysis tasks. Answering specific questions about research data leads at 85.8%, followed by summarizing sessions (84%), synthesizing insights (82.5%), finding patterns (82.2%), and clustering similar insights (81.6%). Transcription, searching past research, suggesting tags, and generating reports all sit above 70%.

Analysis Task Usage

% using each task at any frequency (n=332, multi-select)



This tells you something important: AI is not a side tool. It's already part of how analysis happens.

At the same time, there's a clear boundary: fully automating analysis remains far less common (37.7%), roughly half the rate of the top tasks. Research practitioners are comfortable using AI as a co-pilot for individual steps, but handing over the entire process is a different matter.

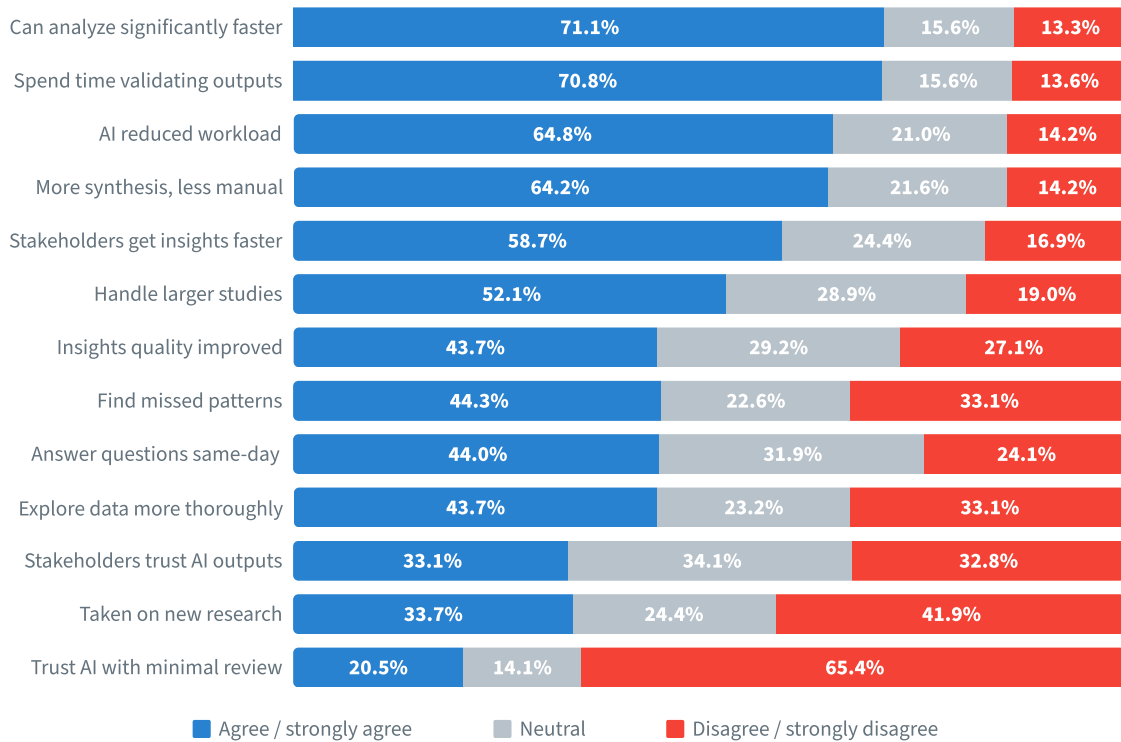
Faster Work, Same Clock

The clearest pattern in the impact data comes from 13 statements about how AI affects analysis work. The results split in two. The four highest-agreement statements are all about efficiency: "I can analyze data significantly faster" (71.1%), "validating AI outputs still takes significant time" (70.8%), "AI has reduced my analysis workload" (64.8%), and "I have more time for synthesis, less for manual tasks" (64.2%).

The four lowest are all about quality and confidence: "I trust AI with minimal review" (20.5%), "I trust AI to take on new research (I would otherwise have declined)" (33.7%), "stakeholders trust AI outputs" (33.1%), and "I trust AI to explore data more thoroughly" (43.7%).

Quality ratings: agree vs. disagree

Stacked 100% · Blue = agree · Light = neutral · Red = disagree (n=332)



This split shows where AI stands today: trusted as a speed tool, not yet as a quality tool.

The efficiency gains are tangible. Tasks that once required hours of manual work, like transcription, initial clustering, session summaries, and report drafts, can now be generated in minutes. And respondents describe being able to cover more sessions, support more stakeholders, and spend more time on the parts of the work that actually require human expertise.

But on the quality side, the picture gets more contested. "Insights quality improved" sits at 44% agree versus 27% disagree. "Find patterns I would have missed manually" is 44% agree versus 33% disagree. This reveals the sentiment that AI is changing the pace of work more than it's changing the depth or reliability of the outputs.

There is one number that captures the core tension: 71. 71% of practitioners agree AI lets them analyze data significantly faster. 71% also agree they spend significant time validating its outputs. Revealing that the time AI saves on execution is being reinvested, at least in part, into checking the work. And what's really changed isn't how much time is saved. It's where it goes.



What's happening is a trust-validation loop in which the time saved on execution gets reinvested into reviewing and judging AI output.

The Role Split Is the Biggest Story in the Data

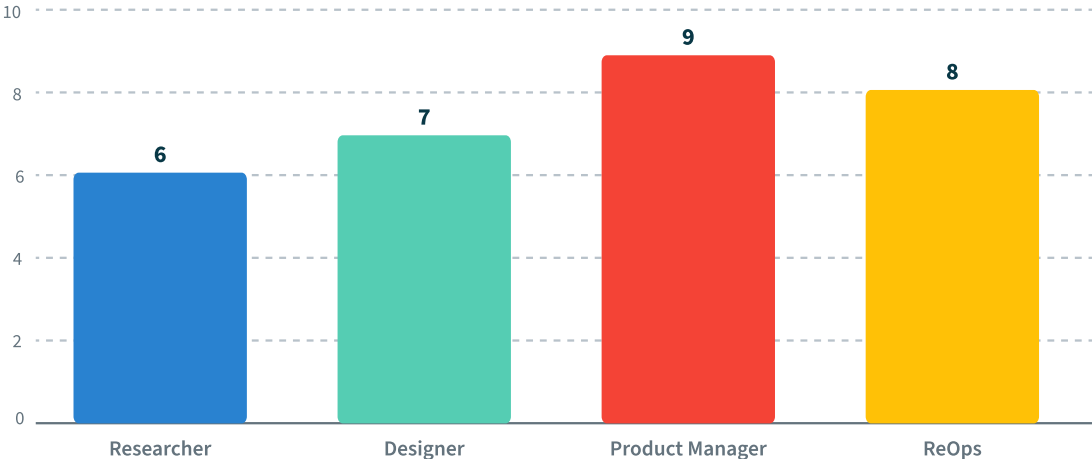
The single biggest predictor of how someone feels about AI in research analysis isn't company size, frequency of use, or tool type. It's role.

On a 1 to 10 satisfaction scale, the median scores are clearly split:

- Product Managers are the most satisfied and most willing to rely on AI (median 9)
- ReOps Managers combine high trust with active oversight (median 8)
- Designers sit in the middle, with their own access and collaboration pains (median 7)
- Researchers are the most critical and the most cautious (median 6)

The three-point gap between Product Managers and Researchers is larger than any difference we measured anywhere else in the dataset. So what's actually going on?

Median satisfaction by role
Median score (1-10) per role group



What's Behind the Gap Between Product Managers and Researchers

Both groups are using AI, so the gap isn't about exposure. It's about what they're using it for and what they expect from it. The following three data points help explain the divide.

1. They're using AI to do different jobs

Product Managers reach for AI mostly to compress information. A clear summary they can share, a quick answer to a question, a first draft of a recommendation. Those are tasks AI handles well, and they don't require the output to survive a methodological audit.

Researchers use AI on qualitative data they often know intimately and need to defend. Their work involves tagging, theme construction, and synthesis that downstream teams will act on. This makes the bar for "good enough" higher because the cost of a quietly wrong insight is greater.

The different qualities that Researchers and Product Managers say make AI more valuable to them are also noteworthy. Researchers rank accuracy (65%) and transparency (54%) at the top of their list. Product Managers are less demanding and generally satisfied with the current state of AI because it's already helping them solve the problems in front of them.

2. They sit at different distances from the data

Researchers work with research daily. They know the participants, the transcripts, and the institutional context that makes a quote meaningful. They are the ones who notice when AI takes a word at face value, when it outweighs an outlier, and when it pulls a quote out of context.

Product Managers work with research more on a weekly basis. They typically engage with the synthesized output rather than the raw material. That's not a limitation, it's how their role works. But it does mean they have less to compare an AI summary against, which is part of why the same output lands differently for each group.

That proximity to the raw data isn't incidental. It's the vantage point that makes researchers' skepticism both understandable and valuable.


3. They're accountable for different things

When a stakeholder pushes back on a finding, it falls on the Researcher to explain how the insight was reached. That accountability shapes how willing Researchers are to accept a confident-sounding output they can't fully trace.

The trust vs. validation data highlights the gap clearly. Only 14% of Researchers trust AI outputs with minimal review, compared to 73% of Product Managers. Validation effort shows the reverse pattern: 74% of Researchers say they spend significant time checking AI outputs, versus just 42% of Product Managers.

Trust vs. validation effort — the central tension in the data

% agree/strongly agree per role (n=332)




ReOps and Designers: Two Different Angles on the Same Problem

Designers sit in the middle on satisfaction, but their pain profile points elsewhere. Researchers are primarily under speed pressure: 47.4% say analysis takes too long and 43.1% say stakeholders need faster answers. Designers, meanwhile, struggle most with access and collaboration. Their top pain points are referencing past research (43.1%) and involving stakeholders in analysis (41.2%), both higher than any other role.

Pain profiles — Researcher vs. Designer

% selecting each pain point · Same scale (max=55%)



Despite their more moderate satisfaction scores, Designers share researchers’ skepticism around trusting AI with minimal review: only 18% agree. Their challenges are less about analysis speed and more about accessing existing research and involving others in it.

ReOps Managers sit closest to Product Managers on the satisfaction scale. They’re also the only role where usage of general-purpose LLMs and purpose-built research tools is nearly equal, likely reflecting their focus on scaling research rather than conducting it directly. AI fits that role well. Even so, “hard to find patterns at scale” remains a persistent pain point, showing that making sense of large volumes of data is still difficult across roles.

Why this matters

It’s easy to reduce these results to “some roles embrace AI more than others,” but the data points to something more nuanced. Researchers spend significant time validating AI outputs (74%) to catch where the system gets things wrong. That isn’t unnecessary friction, it’s part of maintaining quality and confidence in the work. It also helps explain why Researchers are the most-cited source of organizational pushback (35.8%).

The practical question is what to push for. The gap between speed and trust doesn't close by handing AI more autonomy. It closes by making AI more transparent. The capabilities Researchers value most, accuracy (65%) and visibility into how conclusions were reached (54%), point toward tools that make outputs easier to inspect, adjust, and defend. Those are the areas worth pushing for, both internally and with vendors.

In Their Own Words

The open-ended responses show the reasoning behind the satisfaction scores. On the frustration side, the recurring themes are generic outputs, lack of depth and nuance, and the inconsistency that makes it hard to fully trust the output.

To verify AI's work, I have to review the data myself, and now I'm back to spending more time checking the output when I could have just been doing the work.

AI still shows a lot of misinterpretation, whether due to bad transcripts, misunderstanding of irony, or lack of understanding of business terms specific to the organisation. All in all, AI helps me save time scratching the surface, but is useless when digging more deeply.

It gives you a summary-like sweep, but it is often incredibly generic. It flattens peaks and lows to a general average. It also takes words at face value and lacks context.

Where AI earns real appreciation is for a different kind of value: freeing up attention, helping to break through when feeling stuck, and making the mechanical parts of analysis less draining.

It's like a body double so I don't get stuck in analysis paralysis. Useful for a quick extra perspective, reducing tedium, and interesting to see what it misses versus what it catches.

For me it has made the biggest impact in synthesizing large amounts of data. I recently had 60 think-out-loud sessions to get through and what would have taken me probably over a week was done within a day. My mental health is better for it.


It allows us to still take a human perspective on it, while allowing us to save time doing the boring or repetitive work of finding things.



The story is consistent on both sides: AI is delivering on efficiency, but its delivery on depth remains questionable. That's not a reason to stay on the sidelines though. It's a reason to use it deliberately.

Growing Commitment Despite Limitations


Despite all of this, the direction is clear: AI is not going away. Even with known limitations around quality and trust, 81% of respondents plan to increase their use of AI in research workflows. People are not fully satisfied, but they are still investing further.

 This suggests that the current state of AI in research is transitional. The tooling is good enough to be useful, but not yet reliable enough to be fully trusted without human oversight.

The data is clear on what needs to change. The top three things that would make AI more valuable all point in the same direction: more confidence in the output. “Improved accuracy and relevance” leads at 61%, followed by “a better understanding of research methodology and context” at 52%, and “greater transparency into how conclusions are reached” at 50%.

What would make AI more valuable?


% of all respondents (n=332, multi-select)



Cost sits near the bottom at 15%, indicating that the barrier to deeper adoption isn't price. It's reliability and explainability. The most common shortfall reflects this: 52% say insights feel too generic or surface-level, and 42% say AI loses important context.

AI shortfalls in UX research


% of all respondents selecting each option (n=332, multi-select)



When it comes to how people want AI to fit into the workflow, the answer is consistent: supervised collaboration, not AI autonomy. 36% prefer "AI analyzes and synthesizes, I review and approve." 24% prefer "AI suggests codes and patterns, I do synthesis." 20% prefer "AI codes and identifies patterns, I synthesize." Together, those three account for 80% of respondents. Only 5.4% want end-to-end AI analysis.

Preferred AI role — by respondent role

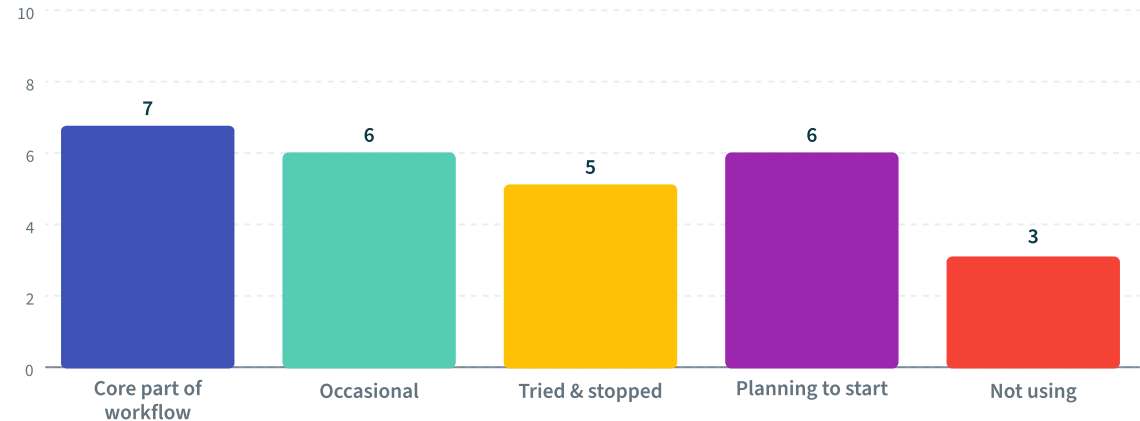
% of each role's respondents · Top 5 preferences (multi-select)




There's also a clear signal that satisfaction rises alongside depth of use. Respondents who use AI as a core part of their workflow report higher satisfaction than those who use it only occasionally or in limited ways.

Satisfaction score by AI usage depth

Median score (1-10) per usage group



 People who have developed better workflows and clearer judgment about when to trust AI and when to verify it appear to exhibit higher satisfaction with it. This seems to indicate that satisfaction with AI is, at least in part, related to how competent people feel using AI.

03

CHAPTER

Opportunities and Risks of AI in UX Research

WHAT'S IN THIS CHAPTER

Where AI genuinely adds value, where the risks deserve attention, and how to evaluate tools.

03

Opportunities and Risks of AI in User Research Analysis

The survey data from the previous chapter shows that AI is actively being adopted for research analysis, but with real friction around quality and trust. That friction rarely comes from the tools themselves. Instead it comes from not knowing what AI is good for and where it needs oversight. That's the skill this chapter aims to help build.

It looks at where AI genuinely adds value in the analysis workflow, where the risks are worth taking seriously, and what to look for when evaluating the tools available to you.

Benefits and Risks of Using AI in Analysis

The benefits and risks of using AI in analysis are often two sides of the same coin. The same capabilities that make AI valuable can also introduce new challenges if used without care.

This section examines four benefit-and-risk pairings in depth.

Benefits	Risks
Speed and Efficiency	Verification Burden and ROI
Scalability	Context and Data Limitations
Pattern Surfacing	Accuracy and Reliability
Accessibility	Overlooking Errors and Hallucinations

Benefits

Speed & Efficiency

Speed is one of the most immediate benefits of AI-assisted analysis. Tasks that previously required hours of manual work can be completed in a fraction of the time and can be especially valuable as a first-pass filter of your data.

Scalability

AI makes it possible to work with data volumes that would have been impractical to analyze manually. Teams that previously had to limit the scope of their research due to bandwidth can now work with richer and larger datasets.

Pattern Surfacing

AI can surface patterns and connections across large datasets quickly. This can be especially helpful as either a first pass to help get the research started, or as a second pair of eyes to double-check your work.

Accessibility

AI can help lower the barrier of entry for non-researchers to engage with data. Product managers, designers, and executives are more likely to participate in research when AI is there to support them. In principle, this makes research easier to access and incorporate into the decision-making process across the organization.

Risks

Accuracy and Reliability

Surfaced patterns are only valuable if they're accurate. Treat them as hypotheses to investigate, not findings to report. Always trace them back to the source data, and keep in mind that LLM training processes mean bias can be present in the output.

Verification Burden and ROI

AI output should always be reviewed for errors and hallucinations. The key question when using AI in analysis is whether the time saved exceeds the time required to review the output. In some cases, doing the analysis yourself is actually faster.

Context and Data Limitations

LLMs have finite context windows. When input approaches the limit, the model becomes selective about what data it references. Beyond that, UX work is deeply contextual. Accurate results require understanding of the product, users, business context, and research questions, context that's difficult to fully supply in a prompt.

Overlooking errors and hallucinations

AI-generated output is often delivered with confidence and eloquence, making it easy for less experienced people to overlook errors and hallucinations. According to [research from OpenAI](#), “errors and hallucinations are fundamental to how LLMs are trained” and this risk increases with larger datasets.



Be cautious of claims that overstate what AI can reliably deliver. Claims like “instant insights”, “eliminates bias”, or “automated analysis” oversimplify what good analysis actually takes.

A Closer Look at Specific AI Analysis Features

Not all AI features carry the same level of risk. Some help with tasks like retrieving and organizing information, while others help interpret, synthesize, and generate findings. The more interpretive the task, the more important it is to stay critical. The features below are ranked from low to high risk, assessed in collaboration with [Dr. Llewyn Paine](#), with a look at what each one does, what to watch out for, and how to use it well.

Transcription: Low risk of error Low risk

AI transcription converts session recordings into text, with identified speakers and linked timestamps. Some tools, like [Condens](#), also support translation across languages, which can significantly reduce the manual effort involved in processing research sessions.

What to watch out for:

- Higher error rates for certain languages, accents, or poor-quality audio
- Misunderstandings or omissions in the transcribed text, particularly when multiple speakers are talking simultaneously or over one another

How to use it well:

- It's always a good idea to double-check transcriptions; mistakes are especially common with multiple speakers or low audio quality
- Don't assume the transcript is complete; read through it carefully to catch missing or misheard content

Anonymization: Low risk of error Low risk

Some research tools enable you to scrub personally identifying information (PII) such as names, email addresses, and credit card numbers from raw data. This helps protect participant privacy and maintain compliance with regulatory laws like GDPR and HIPAA.

What to watch out for:

- These tools do make mistakes. In [one documented case](#), an AI anonymization feature removed all tool names from an interview transcript. The system mistakenly treated product names as identifying information about the participant's employer, stripping out data that was actually central to the research

How to use it well:

- Always do a manual review pass before storing or sharing anonymized data
- Pay particular attention to any content that involves names, brand names, or company references, as these are common sources of misclassification

AI Search: Low risk of error Low risk

AI-powered search allows you to query research repositories using natural language questions rather than exact keywords or tags. This makes research findings more accessible, especially for stakeholders who aren't familiar with your taxonomy or repository structure.

What to watch out for:

- May still miss relevant information (false negatives)
- Occasionally surfaces results that are slightly out of context

How to use it well:

- Verify the results by following citations back to the original data rather than relying solely on the synthesized response
- Use it as an access tool, not an analysis tool. Interpretation should still be left to the researcher

Auto-Clustering: Low-moderate risk of error Low-moderate risk

Auto-clustering groups similar notes, quotes, or highlights into themes using semantic similarity or sentiment analysis. It is commonly used in affinity mapping and on virtual whiteboards, and can save time by eliminating the need to sort sticky notes manually. As an exploratory tool, it can also help surface patterns you might have missed.

What to watch out for:

- The results may remain surface-level, leaving out deeper and more nuanced themes

How to use it well:

- Always review suggested groupings. AI can misunderstand the meaning and organize data incorrectly
- Treat clustering as an exploratory accelerator, not an analytical conclusion; further interpretation remains the researcher's responsibility

Auto-Tagging: Low-moderate risk of error Low-moderate risk

Auto-tagging automatically assigns existing tags to quotes or transcript segments based on semantic similarity, which can speed up the processing of lengthy transcripts. Some platforms, like Condens, offer a safer alternative (AI-suggested tags) where the researcher must accept or reject each suggestion individually before it is applied.

What to watch out for:

- May still miss relevant information (false negatives)
- Occasionally surfaces results that are slightly out of context

How to use it well:

- Verify the results by following citations back to the original data rather than relying solely on the synthesized response
- Use it as an access tool, not an analysis tool. Interpretation should still be left to the researcher

AI Summaries: Moderate risk of error Moderate risk

AI summary features compress transcripts or notes into narrative summaries. This is very useful for getting quick overviews or early synthesis drafts. But make sure that the tool you're using produces reliable citations that can be traced back to source quotes.

What to watch out for:

- Summaries can be superficial, overlook key points, or even misconstrue the main point
- Subtle phrasing changes can materially alter the meaning behind the findings
- Hallucinations and exaggerations can be hard to detect without checking the original data

How to use it well:

- Use AI summaries as a starting point for synthesis, not as finished insights
- Use citations to verify supporting quotes against the original transcript
- Be cautious about language that implies frequency or consensus, such as "users commonly felt..."; and check whether the data actually supports that framing

AI Report Generation: High risk of error High risk

AI report generation goes a step further than summarization by synthesizing data across multiple documents, identifying themes, interpreting patterns, and producing structured insights and recommendations. This is the feature that most closely simulates the full research analysis process, and the one that requires the most caution.

What to watch out for:

- Highest likelihood of hallucination, oversimplification, and flawed interpretation among all AI analysis features
- Very difficult to verify thoroughly without investing nearly as much analytical effort as doing the analysis yourself
- Reasoning chains are invisible, creating accountability gaps

How to use it well:

- Approach AI-generated reports with significant caution, and think carefully before making this feature available to your full team, especially those without training on its risks
- If used, reserve AI report generation for low-stakes internal purposes where errors carry limited consequences

“
AI can be a great help with writing reports,
but it's not good at writing reports. There's
a difference.”



Laura Klein
Principal Experience Specialist
at Nielsen Norman Group

A Guide to Evaluating AI Tools and Their Safeguards

Research data is sensitive by nature. Participant data carries privacy obligations, and AI-assisted analysis introduces new questions about where that data goes and how it's handled.

A benefit of working with a purpose-built AI research tool is that there are often safeguards built in to help mitigate some of the most common LLM-generated errors. Not all tools have the same safeguards in place, so when evaluating options, the questions below are worth asking.

1

Can I control what data the AI draws from, and can I easily provide more context if needed?

The quality of outputs depends heavily on this. A good tool should make it easy to define the scope of data the AI works with and to frame the question you're trying to answer.

2

What happens when the answer isn't in the data?

AI that fills gaps with plausible-sounding answers is more dangerous than AI that says "I don't have that information." Look for tools that are honest about their limits.

3

Can I quickly verify where answers come from?

Any insight or summary should be traceable back to the raw data it was drawn from. If you can't check it, you shouldn't trust it.

4

Is it always clear what was generated by AI and what wasn't?

AI output should be clearly labeled so there's never any ambiguity about what came from the data and what came from the model.

5

Are AI features opt-in or opt-out?

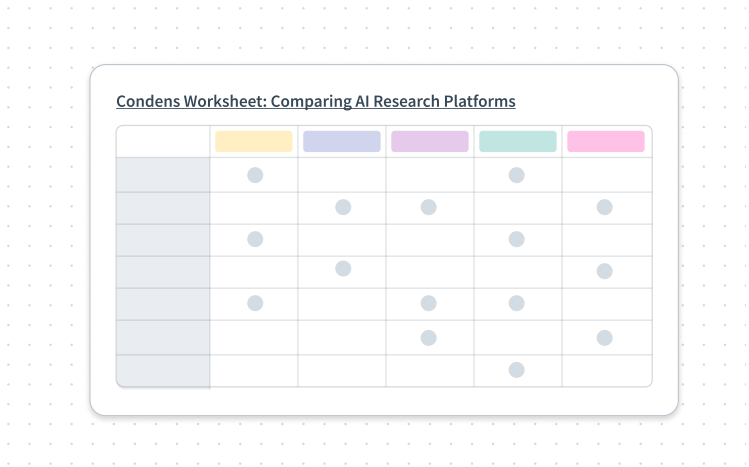
The researcher should always be in control of how AI is applied and should also be able to carry out tasks without AI involvement.

6

Can the outputs be edited, accepted, or rejected at a granular level?

This matters more than it might seem. Tools that let you act on suggestions individually keep judgment with the researcher. Tools that apply AI in bulk make it harder to catch mistakes and easier to let errors slip through.

You can also use this [worksheet](#) to help keep track of your research when evaluating different AI UX research tools:



The answers to these questions look different across tools. To give a concrete sense of what strong safeguards actually look like in practice, here’s a closer look at how Condens approaches them.

Condens’ Approach to AI in UX Research

Building AI into a user research tool responsibly requires making deliberate choices at every level, from how data is handled to how outputs are labeled and what questions the AI is and isn't allowed to answer. The following is an overview of the principles and decisions behind Condens' approach.

Data Privacy and Ownership

Before using any AI tool with research data, it’s important to understand what happens to that data once it leaves your hands. Participant data shared in a research context is not data participants consented to contribute to AI model training.

Purpose-built research tools like Condens address this explicitly. General-purpose AI assistants often do not, unless you are using enterprise plans with clear contractual agreements in place.

There is also an important distinction between using data to generate outputs and using data to train a model. Condens takes a clear position on this: research data in the platform is never used to train LLMs, and contracts with AI providers explicitly prohibit them from doing so. You can think of it as a brain without memory: capable of reasoning over your data in the moment without retaining it afterward.

Knowing what AI can and cannot answer

There's a subtle distinction between the two types of questions that you can ask AI:

- 1 Asking AI to find something ("What did participants say about onboarding?")
- 2 Asking AI to interpret something ("What is the most important point in this transcript?")

The first has a retrievable, evidence-based answer. The second requires judgment beyond what's in the data, the kind that the researcher is best equipped to make.

For stakeholders exploring insights with AI search, Condens can suggest a rephrasing of the question when what was originally asked falls outside the scope of available data.

Safeguards built into the platform

To limit errors and hallucinations in AI-assisted analysis, Condens has built the following safeguards into the platform:

- Explicitly states when information isn't available: When AI-powered search is used to look up information that isn't in the repository, Condens explicitly states that the information isn't available instead of fabricating an answer. This directly limits the occurrence of hallucinations.
- Controlled data scope: Users can precisely define the scope of data that AI can draw from by applying filters alongside their questions. This level of control isn't reliably achievable with general-purpose LLMs.
- Accurate, verified citations: Condens uses an evaluation pipeline to surface only the most relevant quotes, and post-processes supporting quotes to ensure they are accurate and correctly attributed.
- Research-optimized model selection and prompting: Condens optimizes its model selection and prompts specifically for UX research contexts by running evaluations using test data.
- Clear labeling and non-AI alternatives: AI-powered features are all clearly labeled, so you always know when AI is involved before using it. Any task can also be completed without the use of AI.

Understanding Who Will Be Using the Tool



When evaluating AI UX research tools, it's important not only to ask, "How does this tool support researchers?" but also "How does this tool support non-researchers?"

There's a clear and growing trend toward democratized research, and in many ways, this is a positive development. More people engaging with research means that there's a higher likelihood that research will play a role in informing decision-making at every level of an organization.

But democratizing research also introduces risk. Less experienced practitioners are more likely to miss AI-generated errors and overlook the risks associated with different AI features, which can quietly undermine the reliability of insights flowing across the organization.

The AI tools you choose will directly shape the quality of research being done across your organization, not just by your research team. That makes tool selection a decision about research culture as much as it is about features and functionality. Speed and efficiency matter, but not at the expense of trustworthiness and reliability. Choose tools that make responsible, evidence-grounded work accessible to everyone who uses them.

“

Instead of thinking about how AI will work for an experienced star researcher, you want to think about someone junior who may not even be a researcher.

”



Llewyn Paine

Researcher & Founder
at Llewyn Paine Consulting

Do's and Don'ts Checklist

When using AI for user research analysis, there are best practices as well as things that shouldn't be done to uphold the reliability and trustworthiness of the results that will be produced and used. Here's a checklist of dos and don'ts with input from user research experts including Nick Babich, Laura Klein, Maria Rosala, and Llewyn Paine.

Do's

- Start with clean data:** Irrelevant interviews, wrong participants, or incomplete data can result in misleading and biased AI output, so keep your data clean.
- Use AI to escape the blank canvas:** AI is great for brainstorming research questions, drafting discussion guides, and suggesting study designs, but stay in control and vet its suggestions. Always refine what it generates before using it.
- Use AI to challenge your analysis conclusions:** Prompt it to find contradictions in the data or point out weaknesses in your reasoning. It's a solid thinking partner for catching blind spots.
- Take full advantage of AI for scoped tasks:** Transcription, searching for relevant quotes, suggesting tags, and initial clustering are all areas where AI adds clear value. The more specific the task, the better the output.
- Validate AI output with human judgment:** AI can miss nuances like sarcasm, idiomatic expressions, and emotional subtext, and can produce sycophantic or biased responses. People often say one thing and mean another. Always review tags, clusters, and summaries rather than accepting them at face value.
- Use AI for meta-analysis of past research:** Let AI suggest tags, find patterns across studies, and surface insights from untagged data. But don't hand off the full tagging process. Deep familiarity with your data is what makes themes meaningful.
- Make every AI insight traceable:** Always verify AI-generated quotes against original transcripts. LLMs can misconstrue, misattribute, or cite irrelevant quotes. Be clear about how AI was used and who owns the insights.
- Choose tools that support responsible research:** As AI lowers the barrier to doing research, make sure non-researchers are guided by tools and safeguards that encourage meaningful, responsible work.

Don'ts

- Don't assume AI eliminates bias:** LLMs are trained on biased data and can't follow up the way a trained researcher can. Bias doesn't disappear because AI is involved.
- Don't let AI shape your research process on its own:** A faulty research process can lead you down the wrong path from the start. Don't blindly trust AI-generated tasks or interview questions. They can be leading, inappropriate, or constrained by the tool's framework.
- Don't replace collaborative analysis sessions with AI reports:** Co-creating insights during collaborative stakeholder sessions (e.g. watch parties) helps build team consensus and drives organizational change. AI can't replicate that.
- Don't feed transcripts into general AI tools without checking privacy:** Unless you have participant consent and have removed all PII, this may overstep privacy laws. Research-specific tools offer better protections and more accurate results.
- Don't hand all your data to AI and expect real insights:** AI can produce insight-shaped things that look convincing but lack depth. It doesn't know your product or your users, and the more it interprets, the higher the risk of hallucinations and missed nuance. Don't let it decide what matters, especially for critical decisions.
- Don't replace real users with AI-generated synthetic users:** Synthetic users have no real experiences and are hard to follow up with meaningfully. Using them as a substitute can multiply systematic mistakes and create false confidence in wrong conclusions.
- Don't rely on auto-generated reports for complex or high-stakes decisions:** AI can produce outputs that look convincing, but they often lack the context and judgment that come from human analysis. And without clear ways to verify how conclusions were reached, it's hard to assess their accuracy or trustworthiness.
- Don't assume AI tools are all the same:** Evaluate tools carefully. AI is non-deterministic and makes different mistakes each time, so don't judge a tool on a few quick prompts.

A Practical Guide to Using AI in Analysis

WHAT'S IN THIS CHAPTER

How to move from raw data to validated insights with AI, including a real-life case study showing what's changed and what hasn't.

A Practical Guide to Using AI in Research Analysis

Understanding the principles behind good AI-assisted analysis is one thing. Putting them into practice is another.

This section walks through what a human-in-the-loop workflow actually looks like, from the moment data comes in to the point where an insight is ready to share. It also includes a real-life example, tips on how to ask better questions/define better prompts, and how to maintain rigor, traceability and transparency.

Designing a Human-in-the-Loop Workflow

The most useful way to think about AI in research analysis is not as automation, but as collaboration. AI handles the parts that are time-consuming and mechanical. Researchers handle the parts that require judgment, context, and interpretation.

This division of tasks matters because using AI to speed up the interpretive parts of research analysis without human oversight and accountability increases the risk of poorly-informed decisions and undesirable consequences. A transcript processed incorrectly can easily be caught and fixed. An insight interpreted without context, or a pattern accepted without scrutiny, can travel far before anyone questions it.

Keeping the researcher in the loop at every interpretive step is what makes the workflow trustworthy. Human judgment is crucial for interpretation. That has always been true, and AI does not change it.

Here's an overview of the workflow:



From Raw Data to Validated Insights

We'll now go through a stage-by-stage walkthrough of the analysis workflow, from the moment a session is recorded to the point where a pattern becomes a finding you can stand behind. Each stage is both an opportunity for AI support and for you to judge what that support is actually worth.

Data Collection and Transcription

Transcription is where AI is most reliable. Converting audio to text is straightforward, and if there are any typos or errors, you can check it against the recording. Many tools, including Condens, also support translation alongside transcription. This matters more than it might initially seem: it means you can work with research conducted in other languages, pull in sessions run by colleagues in other markets, and build on findings from studies you weren't a part of without language being a barrier to what's usable.

Accuracy does vary however. Poor audio, overlapping speakers, and unfamiliar accents can all affect transcript quality. So spot-checking against the original recording is a step worth building into your process, especially for quotes you intend to use.

However, transcripts are only part of the picture. AI can turn audio into words, but it can't capture the researcher's in-session thinking and observations: the moment a participant hesitated, the idea that surfaced mid-interview, etc.

Live note-taking during sessions helps enrich the transcript with context. In Condens, Live Notes lets you capture observations, reactions, flagged moments, and emerging ideas in real time. They're also embedded directly alongside the transcript. So when AI summarizes or surfaces patterns later, it can work with data deepened by your input.

Session Summarization

Think of AI summaries as a navigation tool that tells you where to look. They're a starting point, not a substitution for the raw data or complete analysis outputs.

The single most important thing you can do to improve summary quality is to provide context. Without it, AI defaults to frequency: what was mentioned most often instead of what matters most to your study.



The prompt "summarize this interview" produces a topic overview. The prompt "summarize this interview with a focus on pain points, workarounds, and moments of frustration or surprise" produces something you can actually work with.

Look for tools with summaries that include citations and direct links back to the specific transcript moments they draw from. So that validating a claim takes seconds rather than requiring a full re-read.

Querying Your Data

The most natural way to start analysis is with questions. You just finished a session, and you already have things you want to know: did this participant mention the same friction point as the last one? What did they say about the onboarding flow? Was there a moment where their sentiment shifted?

In an analysis tool like Condens, you can ask those questions directly against your session in plain language. Just type your question and you'll get a cited answer drawn from the transcript. In a way, it's like interviewing your raw data, rather than searching for keywords or scrolling through timestamps.

This makes it a helpful first move. Patterns can begin to emerge before you have done any formal analysis, and specific answers surface in seconds. From there, you can follow up with a closer read of the raw transcript, move into clustering, or stop here if you only need specific answers rather than a full synthesis.

Researcher Review and Refinement

This is the step that gets skipped most often, usually because a solid AI summary makes it feel redundant. But it's not.



Your job is to check what was filtered out, and to bring the kind of contextual reading that AI just can't.

AI can identify that a participant used the word "frustrating" three times. But unless AI is trained on visual intelligence, it can't know that they smirked when they said it, that they were describing a tool they'd already stopped using, or that their answer might have been shaped by how the question was framed. That kind of reading requires someone who was in the room, or close to it.

This is also where the most interesting analysis often happens. Where you notice things that do not fit the emerging pattern and point you in a direction you weren't expecting.

Practically, this looks like:

- 1 Reviewing the AI summary and flagging anything that feels incomplete
- 2 Adding your own observations and interpretations as highlights
- 3 Marking moments that feel significant (even if you can't articulate why yet)
- 4 Noting anything in the raw data that contradicts or complicates what the summary surfaced.

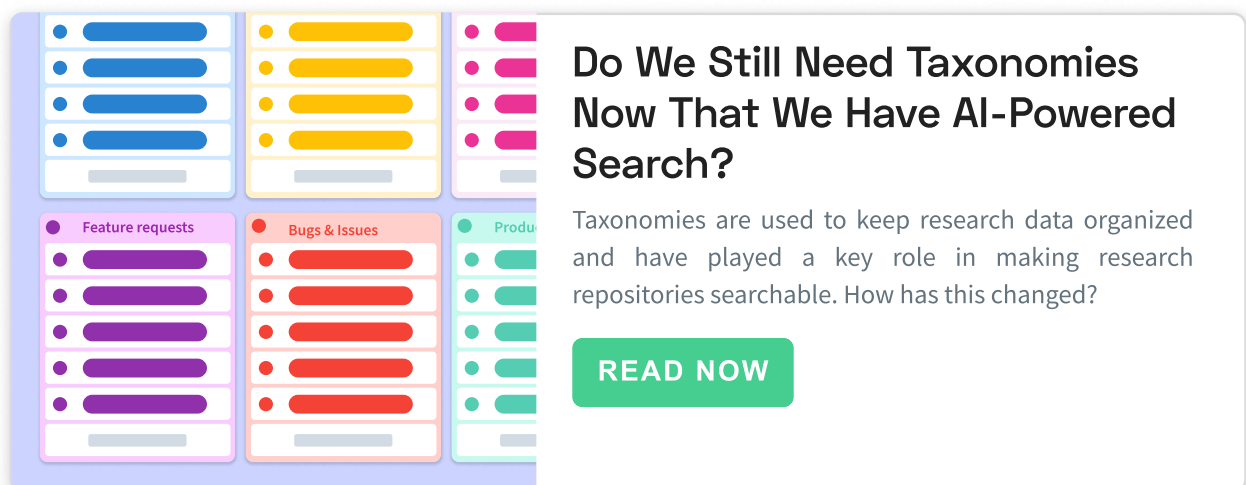
Tagging

Tags have always done two jobs: making data findable and making patterns visible. AI search changes the first of those significantly. Because when you can ask AI a question and get cited evidence from across your repository, data doesn't need to be perfectly tagged to be retrievable. The pressure to build a complete taxonomy before your first session, and to apply it flawlessly throughout, is genuinely lower than it used to be.

What tags are still essential for is the second job. A tag applied to a highlight is a human judgment call, which can accumulate into a structured record of what your team has decided is meaningful, traceable across studies, and comparable over time.

This shifts tagging from a front-loaded obligation to something more like a byproduct of active analysis that's applied as you go. The resulting taxonomy then tends to reflect what actually matters rather than what seemed important before the research started.

One thing that remains important regardless of how AI develops: metadata. Things like participant role, company size, or research phase are often not present in the content itself, and AI search can't infer them reliably. "What did enterprise users say about onboarding?" is only answerable if that context about the users is already attached to the data.



Do We Still Need Taxonomies Now That We Have AI-Powered Search?

Taxonomies are used to keep research data organized and have played a key role in making research repositories searchable. How has this changed?

[READ NOW](#)

Pattern Detection and Clustering

This is where AI's ability to work across large volumes of data becomes most valuable. Rather than manually comparing notes across sessions, you can ask AI to find patterns in your dataset like what's recurring and what's diverging.

AI clusters around language patterns. So two participants can describe the same experience in different words and end up in separate clusters. On the other hand, two participants can use similar language for entirely different experiences and land in the same cluster. The highlights that do not cluster at all are worth examining carefully. That's often where the outliers are, and a single data point that does not fit the pattern can sometimes be exactly what you are looking for.

The key distinction is between AI identifying patterns and AI interpreting them. The first is something AI does well. The second should still be left to a researcher.

Validation Against Raw Evidence

Before a pattern becomes a validated insight, it needs to be traceable to specific moments in specific sessions, not just to an AI summary or a cluster label.

Validation means being able to answer:

- Can you point to two or three moments in the raw data that support this insight?
- Does it still hold up in context?
- Is there any contradicting evidence?
- Are there any participants whose experiences do not fit the pattern?

Research that can't be traced back to evidence rarely survives the first hard question in a stakeholder readout. That's why maintaining a clear evidence trail isn't just good practice, it's essential to upholding your research's credibility.

From Theory to Practice: Before vs. After AI

This section shows what it looks like when the workflow detailed in the previous section is applied to a real project. Special attention is given to where AI changed the process and where it didn't.

The example comes from [Sina Richter](#), User Researcher at Condens. The study was a journey map project on how researchers decide on a tool stack. It was run twice, roughly eighteen months apart, with the same research question but a different workflow each time: first manually, then with AI-supported analysis in Condens.

The table below compares each stage of the process side by side. Two stages look identical in both approaches. The rest is where the difference shows up.

	Without AI	With AI
1 Framing the research	Collaborative alignment on assumptions, goals, and key questions across teams.	Same process. AI doesn't change the framing conversation.
2 Starting from existing evidence	Filter the repository by global tags, drag items onto a whiteboard, walk through what the team remembers from past sessions. Dependent on how well things were tagged at the time.	Plain-language search across the entire repository. Questions like "What do we know about how researchers discover new tools?" return cited answers, including highlights from raw data that never made it into a formal report.
3 Preparing for sessions	Recruiting and interview guide preparation as usual. Discussion guide refined through manual review.	Same preparation, with AI helping fine-tune the discussion guide. Project questions set up in advance and run against each transcript as it comes in. Transcript generation starts during the debrief.
4 During and after sessions	Record, take notes, debrief. Analysis happens separately, often with a significant time gap.	AI processes transcripts immediately, with Live Notes linking in-session observations to transcript moments. Prepared research questions run against the data right after each session, surfacing relevant moments without a full rewatch. AI-generated summaries provide a first overview.

5	The deeper analysis	Filter highlights by tag, drag onto an affinity map, cluster manually, repeat. Joint session on the same board, then a longer individual finishing phase. Multiple passes to ensure nothing was missed.	The whiteboard already has structure when analysis begins. Tag filtering and AI-assisted clustering work in tandem, with manual refinement throughout. AI kickstarts the process, supports double-checking, and populates clusters with supporting evidence once the researcher has shaped them.
6	The output	Journey map with key findings, highlights, and workshop-based synthesis.	Same type of output. No fundamental change in format or deliverable.
7	The outcomes	Slower start, higher friction to begin analysis. Gaps between sessions and synthesis make it harder to re-engage with the data.	Faster entry into analysis. Immediate access to highlights, themes, and summaries lowers the barrier to start while context is still fresh.
		Without AI	With AI

The parts AI changes the most are where it helps you get to the evidence faster, lowers the hurdle to start, reduces the time gap between session and analysis, and suggests a first structure you can work with and refine.

Worth noting is that the "framing the research" and "delivering the output" stages are identical in their approaches, with the thinking at the start and the deliverable at the end looking the same. The difference is everything in between: how quickly prior knowledge is surfaced, how sessions feed into analysis, and how much mechanical work the researcher has to do before the real interpretation can begin.

The judgment calls are still the researcher's, like what a theme means, which observations matter to which stakeholder, how everything fits into something a team can act on. AI hasn't changed that part. What's changed is how much time and energy you have left for interpreting the data after the time-consuming and mechanical tasks are done.

Asking Better Questions to Get Better Outputs

The quality of what AI gives you is directly shaped by what you ask. A vague question gets you a vague answer, while a well-framed one gets you something you can actually use. Here are a few tips to help you get better results.

Include your research question(s)

Without your research questions, AI defaults to output that just surfaces the most commonly occurring themes. With them, it reflects relevance. This addition improves the usefulness of summaries, search results, and clustering more than any other change.

Be specific about what you are looking for and what you are not

If you want emotional responses rather than behavioral descriptions, or friction points rather than feature requests, say so. Naming what you want closes the gap that AI would otherwise fill with assumptions.

Ask AI to show its work

When AI identifies a theme or pattern, asking it to cite the specific moments that support it makes validation much faster by turning verification from a search through transcripts into a quick check.

Use AI to challenge a hypothesis, not just confirm it

This is one of the more underused applications. If you have a working hypothesis, ask AI to find moments that complicate or contradict it. You'll often get something more useful than if you'd asked for supporting evidence, and it's a good way to guard against the natural tendency to only try and find what you're already looking for.

Iteration matters too

First prompts rarely produce the best output. Refining the question based on what the first answer surfaces is part of the process, not a sign that something went wrong.

Instead of this

"Summarize this interview"

"What themes came up?"

"Find insights about onboarding"

"What did users say about the dashboard?"



Try this

"Summarize this interview, focusing on friction points and moments where the participant expressed surprise or frustration."

"What themes came up related to [research question]? Cite the specific moments that support each one."

"Find moments where participants described confusion, workarounds, or drop-off during onboarding. Note any contradictions across sessions."

"What did users say about the dashboard that surprised them or did not match their expectations?"

Maintaining Rigor, Traceability, and Reproducibility

Of the qualities that define good research analysis, there are three that become especially important once AI is introduced into the workflow: rigor, traceability, and reproducibility. Not because the others matter less, but because these are the ones most directly affected by AI and most likely to erode if the workflow isn't deliberately designed for AI-supported analysis.

Rigor

[Carl Pearson](#) describes rigor as the process of achieving clarity and being able to view reality clearly enough to make sound decisions. That's difficult to uphold when judgment and interpretation are handed over to AI.

That's why in practice, maintaining rigor in an AI-assisted workflow means treating AI output as input for your analysis, rather than as the conclusion. A cluster is a hypothesis. A summary is a starting point. A pattern surfaced by AI is worth investigating, not reporting.

Traceability

Traceability is what makes insights trustworthy and defensible over time. Six months later, when a stakeholder challenges your recommendation, or a product manager asks where a finding came from, having a clear link between your insight and the raw evidence is what will give it credibility.

In Condens, traceability is built into how AI works: summaries link to the specific transcript moments they draw from, search results cite their sources, and every highlight stays connected to its session. It's automatic and integral to how the outputs are structured.

If you're using a general-purpose AI tool like ChatGPT or Claude, that link doesn't exist by default. The model has no concept of source data unless you build it in. So the most practical approach is to treat source attribution as a prompt requirement, not an afterthought.

Instruct the model explicitly that any finding, theme, or pattern it surfaces must include the specific quote, participant, and session it draws from. If you use the same tools regularly, build that instruction into a reusable prompt template so it can be consistently applied.

Reproducibility

Reproducibility means documenting your process clearly enough that another researcher could follow the same steps using the same data. In AI-assisted workflows, this includes being explicit about which steps involved AI, which prompts or questions were used, and where human judgment was applied.

This isn't only about methodology for its own sake. It's what allows findings to be challenged, extended, or updated as new research comes along.

But it also doesn't need to be elaborate. A brief note in a findings document indicating which steps were AI-assisted, which outputs were reviewed and refined by a researcher, and where the conclusions were drawn by a human is enough. Transparency here isn't a disclaimer. It's part of what makes the work credible.

How AI Changed the Way Research Works: A Case Study with Michaël Dufranne

Workflows, frameworks, and best practices are useful. But sometimes the clearest picture of what AI in research analysis actually looks like comes from someone who has been doing it for years. So we spoke with [Michaël Dufranne](#) about how AI has changed the way he works.



Michaël Dufranne
Senior UX Researcher and
Freelancer

Michaël Dufranne is a Senior UX Researcher and Freelancer with 30 years of experience in UX and a background in psychology. Over the last four years, he has implemented Condens across two major Belgian energy companies: Elia and Luminus. At Luminus, he is one of two researchers on an eight-person design team, running user interviews, usability tests, and customer feedback analysis one quarter ahead of the agile teams they support.

The Challenge Before

Before Condens, Michael's analysis workflow consisted of taking notes on a notepad, re-listening to hours of session recordings to fill in the gaps, and creating a PowerPoint that, as Michael puts it, would "die peacefully on SharePoint."

The problem wasn't just time. It was that all the work that went into the research rarely made it beyond that single deliverable. Knowledge stayed siloed, findings were hard to reuse, and there was no easy way for the wider organization to understand what had already been learned.

"Communicating what we already know within the company is a very big challenge for a lot of companies," Michaël says. "That's definitely something that drove us to adopt a research repository."

How AI Fits Into the Workflow

Auto-transcription was the first AI feature that changed things for Michaël. "Thank God it's there!" he says simply. From there, his use of AI has grown to cover most stages of the analysis process, always with human judgment as the final check.

He uses AI summaries for Highlights in every project, accepting or correcting them to save time typing. AI-suggested tags help speed up the processing of lengthy transcripts. And session summaries have opened up a use case he hadn't anticipated: reviewing research he didn't conduct himself, or sessions conducted in a language he doesn't fully speak. For example, when a team in Germany ran sessions he needed to learn from, AI summaries gave him a way to access the research without needing to be fluent in the language.

For report preparation, Michaël uses AI as a sparring partner rather than a ghostwriter. "I ask what the main pain points are, what the key sections of a report could look like, and it gives me a structure I can react to. Sometimes it comes up with a way of phrasing things I hadn't considered. Sometimes I'd do it differently. But it allows me to confront myself with the data, which is useful when I'm working alone."

What he keeps manual is equally deliberate. During sessions, he takes live notes, capturing observations and remarks that later become the foundation for his highlights. That habit does two things: it keeps him close to what actually happened, and it lets him go straight to what matters when it's time to build highlights rather than rewatching from the start. Because "AI doesn't see things that I see," he says. "Sometimes it picks up highlights that are wrong or irrelevant. I need to stay close to the raw data." Final interpretation, synthesis, and judgment calls that shape a report remain his responsibility throughout.

The Impact

The numbers are straightforward. For six user test sessions, Michaël estimates he saves roughly three days compared to his previous workflow. But the impact he describes goes beyond hours.

"It's avoiding all the boring parts of the job that we had before. I love the interview part. I love the reporting, but the thing in the middle was a bit annoying. And that's reduced at least to the bare minimum."

More importantly, the research itself travels further. At Elia, product owners began using the Insights Magazine as a standard part of starting new projects, checking what was already known before scoping new work. The research that used to sit on a drive now sits at the center of decision-making.

"Condens made research fun again. But most of all, all the hard work we put in is easily available for the organisation. It is a game changer and enables a truly customer-focused approach."

Advice for Anyone Doing Analysis

Michael's experience also reflects one of the core principles running through this guide: AI works best when the researcher stays in the loop.

“
Don't trust it blindly. Use it for what it's good at, like summarizing and writing, but don't use it to replace human intelligence.
”

He describes it as a safeguard mindset: AI prepares the work, surfaces structure, and saves time on the mechanical parts. The researcher decides what it means.

His summary of what AI in research analysis actually delivers:

"Deliver more, better, faster."

05

CHAPTER

Beyond Analysis: Scaling Access to Research with AI

WHAT'S IN THIS CHAPTER

Getting research in front of the right people, while maintaining quality at scale.

05

Beyond Analysis: Scaling Access to Research with AI

AI doesn't stop being useful once analysis is done. How research is communicated, shared, and used in day-to-day decision-making is where a lot of research impact is either created or lost, and AI can also be a helpful support with these tasks.

Finishing the Thinking First

Before AI can help communicate insights, the thinking behind them needs to be complete.

Where AI genuinely adds value is in turning rough synthesis into structured, usable output. For example, transforming a messy set of highlights into a structured draft, generating a shorter summary for an executive audience alongside a more detailed one for a product team, or helping researchers get past the blank page. This is where many researchers say AI earns its keep.

On the other hand, if it's used too early before the analysis is fully formed, it can turn uncertainty into confident-sounding conclusions. So the output may read well, but the underlying logic may not hold.

The practical guideline is simple: finish the thinking first, then use AI to help you communicate it.

Using AI to Support Communication, Not Replace Storytelling

Good storytelling hinges on context. What stakeholders already believe, which decisions are currently in motion, and what has shaped their thinking so far. And that context isn't something that AI readily has access to.

What AI can do is help you prepare. It can help you anticipate likely questions, stress-test the logic of an argument, or generate alternative framings so you can choose the one that fits your audience and situation. Some researchers use it to deliberately surface more critical or risk-focused perspectives as a way of pressure-testing conclusions before presenting them.

But the storytelling itself (including the choices, the emphasis, and the judgment) stays with you.

Tailoring Insights for Different Audiences

Different audiences care about different things. Executives care about strategic implications. Product teams want concrete findings they can act on. Designers need context around behavior and friction. So even when you're sharing the same findings, a little tailoring can go a long way in making your communication more engaging.

AI makes this easier by helping to generate multiple versions of the same insight from a single source. Instead of rewriting everything from scratch, researchers can adapt length, emphasis, and framing for different audiences more efficiently.

Making Insights Accessible Without Oversimplifying

Nobody wants to overwhelm their audience, but this can lead to oversimplifying results, which is its own problem.

Stakeholders are busy and want clear answers. But findings can be conditional, patterns can be partial, and the honest answer to "what should we do?" is often "it depends."

AI can help make insights more accessible, but it needs to be used carefully. Prompts like "simplify this" can easily strip out important qualifiers or uncertainty. So a better approach is to use AI to structure or clarify while keeping the boundaries of the evidence intact.

Keeping Insights Connected to Evidence

As insights move through an organization, they tend to drift away from their source. Summaries get rewritten, slides get reused, and over time it becomes harder to trace where a finding came from or how well it's supported.

AI doesn't solve this by default. When rephrasing a sentence, an AI model has no awareness of whether it is based on one data point or multiple studies. That's why evidence needs to stay visible in outputs, and why strong tools link findings back to their source material rather than leaving the evidence trail up to chance.

How Stakeholders Can Use AI and Where Researchers Set the Guardrails

Getting research into the right hands is one challenge. The other is making it findable without the researcher always needing to be in the middle.

Once research is published, stakeholders will expect to be able to access it themselves. And with AI search, now common in research repositories, this is indeed possible. Stakeholders can ask questions in plain language like “What do we know about enterprise onboarding?” or “What usability issues came up most often?” and get answers grounded in published research.

This allows them to:

- Explore published findings to inform a decision they're already making
- Surface relevant research before a meeting or product review
- Follow up on a finding to understand its context and source
- Check whether a question has been researched before

This reduces dependency on the research team and makes insights more accessible across the organization.

But these capabilities come with clear boundaries. AI search cannot:

- Answer questions that haven't been researched yet
- Generate new analysis from raw data
- Assess the strength of a finding or flag where the evidence is thin
- Replace a conversation with the researcher when the stakes are high

The best tools are explicit about these limits. They return an honest “no results found” response rather than generating plausible-sounding answers from incomplete data. Condens, for example, only draws from validated, published research and clearly signals when no matching content exists.

AI Search Is Only as Good as the Research Behind It

What stakeholders can find with AI depends entirely on what you have published and how well it's maintained. Clearly titled outputs, explicit summaries, and consistent metadata directly shape what search can return. So research that relies on visuals or context held in the researcher's head won't surface well.

This means treating your repository as a live resource, not a passive archive. Keeping content relevant, retiring outdated findings, and publishing work you actually want someone to act on is ongoing work, not a one-time setup.

It also requires deliberate decisions about access. Tools like Condens support different stakeholder roles: a curated, read-only view of published findings for most stakeholders, and deeper access to raw data for those with the context to interpret it.

This distinction matters. A poorly structured repository doesn't just produce weak results, it can produce misleading ones. And once stakeholders lose trust in what they're seeing, they lose trust in the research itself.

AI doesn't change that dynamic. It amplifies it. The quality of AI-assisted knowledge sharing is ultimately a reflection of the quality of the research practice behind it.

Bringing Research Where Decisions Happen

Research shouldn't just live inside research tools. Decisions happen in Slack threads, Notion documents, product specs, and conversations the research team is often not a part of, and getting research into those moments requires a different approach.

AI makes it possible to bring research into existing workflows. For example, the Condens Insights Magazine provides a curated, stakeholder-friendly view of research findings, and its AI search extends into tools like Slack and Microsoft Teams. A question asked in a product channel can return a cited answer drawn from published research, without anyone needing to leave the tool they're already using.

The Condens MCP (Model Context Protocol) server takes this a step further.

Instead of just enabling search inside specific tools, it turns your research repository into a knowledge layer that can be accessed from AI assistants and other connected systems. A product manager drafting a spec can ask what research exists on a problem and get cited answers directly in their workflow. And a designer can sanity-check an assumption against past findings without opening a separate tool.

Because MCP servers can connect multiple sources, it also enables a more meaningful form of secondary analysis.

Research findings can be considered alongside product analytics, CRM data, or customer feedback in a single interaction. Shifting the question from “what did users say?” to “what did users say, and how does that connect to what we’re seeing elsewhere?”

It is important however to have the right guardrails set up so that the MCP only draws from validated, published research and respects existing permissions. Without these guardrails in place, you risk surfacing raw data to the wrong people, or letting unverified inputs shape answers. So getting the foundation right is key to ensuring that the output will be trustworthy.

Looking Ahead: The Future of Research

WHAT'S IN THIS CHAPTER

What comes next for AI in research, and what does it mean for the researcher role.

06

Looking Ahead: The Future of Research Analysis

AI research tools will keep improving, and the people who get the most from them will be the ones who stay in the driver's seat on interpretation. The future isn't AI replacing analysis. It's researchers, designers, product managers, and everyone else who works with customer insights doing better analysis faster, with AI handling more of the work that doesn't require human judgment.

That's a more strategic role, not a smaller one. While AI makes research more accessible across teams, experienced practitioners are needed to set the standards, support the practice across the organization, and make sure the insights flowing through it are ones people can actually trust. That kind of leadership will become more important.

What to Expect Next From AI and Condens

As [Alexander Knoll](#), CEO and Co-Founder of Condens, puts it, focusing on scoped, well-defined tasks that can be easily verified is what allows teams to get genuine value from AI today.

What's changing is what constitutes a "scoped" task. As LLMs continue to improve, the amount of data AI can reliably process, along with the range of data types that can be incorporated into research workflows, is increasing. Sales calls, app reviews, support tickets, and other forms of customer feedback that were previously too time-consuming to analyze are becoming viable inputs, resulting in a much richer picture of the customer built from more varied evidence.

Another shift worth paying attention to is how AI tools are starting to connect and work across platforms. Until recently, most AI research tools were built as closed environments. But teams don't work in a single tool, and MCPs now make it possible to connect general-use LLMs like ChatGPT, Claude, and Copilot with external tools. For research repositories, this is a meaningful development: teams can now incorporate the AI tools they already use directly into their research workflow.

This is also the direction Condens is moving in. Rather than building a walled garden, the goal is to make Condens available to the broader AI ecosystem and enable both people and AI agents to work together within Condens. We're also building in safeguards like verification checks and user-defined controls over what agents can do to protect against hallucinations and unsupported AI-generated insights making their way into your repository. Because the value of keeping all your research in one place only compounds if the data stays clean and reliable.

About Condens

Condens is the central hub for all your customer knowledge, bringing together everything from interviews and usability tests to product feedback and surveys in one place.


As both a research repository and analysis tool, Condens helps researchers, designers, and product teams store, organize, analyze, and share insights more efficiently than ever. Import raw data, auto-transcribe recordings, and let AI do the heavy lifting: summarizing sessions, surfacing patterns, clustering themes, and answering questions about your data in plain language, so you can spend more time on the work that actually requires your expertise.

The result is faster analysis, less time lost to manual processing, and research that travels further inside your organization. Every AI-generated summary and highlight links back to the raw transcript it draws from, keeping findings traceable and credible without slowing you down.

Connect insights across projects, build a searchable repository, and create reports your stakeholders love. Collaboration becomes effortless with full control over access. Share findings via a dedicated Insights Magazine and let stakeholders search key findings directly from Slack or Teams, so research is front and center when decisions are made.

Whether you're scaling research across your organization or just looking to get more out of every session, Condens transforms the way teams work with customer knowledge.

Ready to see it in action? Start your 15-day free trial today.



Start Your 15-Day Condens Trial For Free

Condens combines a powerful research repository with AI that helps you analyze faster and share findings more effectively. Try it for free and see what changes.

[TRY 15 DAYS FOR FREE](#)


Explore Even More UXR Guides



Introduction to Research Repositories Guide

Dive deep into our guide on research repositories that covers everything from basics to best practices. It's packed with practical tips to assist you in mastering the art of customer knowledge management.

[READ NOW](#)



Turn Insights into Action: How to Keep Your Research Alive After the Readout

This guide is here to help your research move forward and your insights stick, without burning you out trying to chase people down.

[READ NOW](#)



Insights That Get Used: Drive Repository Adoption & Stakeholder Engagement

This guide gives you practical strategies to boost stakeholder engagement and make your research repository the go-to source for customer insights.

[READ NOW](#)

condens

All your customer knowledge in one place