Draft chapter from Baker, Ryan S.; Barany, Amanda (forthcoming) Artificial Intelligence in Qualitative Research: New Possibilities. Philadelphia, PA: Vassant Press.

Find the latest chapters online at
https://learninganalytics.upenn.edu/book/ai-qualitative-research.html

Copyright Ryan S. Baker and Amanda Barany, 2026, all rights reserved.

Cite this chapter as Baker, Ryan S.; Barany, Amanda (forthcoming) Artificial Intelligence in Qualitative Research: New Possibilities. Ch. 2: Introduction. Philadelphia, PA: Vassant Press.

---

2. Introduction
Version 1.01, 19 October 2026

What is qualitative research?

Before we explore the ways that artificial intelligence may (and may not) support qualitative research, we want to dig into what scholars actually *mean* when they say they are doing qualitative research. There is a rich and interdisciplinary history of the development and application of qualitative methods that we will touch on only lightly in this book, though if you are interested in a deeper dive, the second author recommends Denzin and Lincoln's discussion of eight historical moments of qualitative research (2003; 2005; 2018), and Jerry Willis and colleagues' exploration of the foundations of qualitative approaches for different paradigms and schools of thought (2007).

The first thing to establish is that qualitative research is very much not a single thing. Instead, the concept of "qualitative research" is what the second author often calls an "umbrella" term, or a set of commonalities that cover (fully or partially) a wide set of values, definitions, and approaches. What often unites these approaches is a commitment to research questions and methods that embrace and intentionally include complexity - or "messiness" depending on your point of view - as part of the research story (Packer, 2011). That complexity often means accepting and even emphasizing that a phenomenon is better understood by considering the impact of context. As a result, the practices and data sources that many tend to think of as "qualitative" in fact emerged as a way for researchers to explore aspects of complex phenomena that are not fully represented through the analysis of discrete variables and their relationships.

Qualitative research is also an approach that researchers undertake when they wish to explore and understand the meanings that individuals or groups ascribe to a social or human problem

(Creswell & Poth, 2016). Unsurprisingly, this focus then shapes every step of their research process, including what questions are asked (and valued), what data is elicited and collected (and how), what kinds of analyses are conducted, and what results are generated (and considered "good" or "useful"). Sometimes, though certainly not always, phenomena are studied in their natural settings (Denzin & Lincoln, 2005). In this kind of context-embedded work, researchers may choose to evolve their research questions and procedures as their understanding deepens and analysis proceeds inductively from context-driven insights toward broader themes. In this type of scholarship, the final results of the work are designed to help preserve and include the complexity of the qualitative "story," rather than flatten it into a set of more straightforward answers (Creswell & Creswell, 2017; Packer 2011). In this view, qualitative research is not just a technique for gathering and analyzing non-numeric data (which nowadays can be analyzed quantitatively as well), but a set of methodological decisions that prioritize interpretation, context, emergence, and complexity.

These priorities have important implications for study design that often start with the development of a research question. Maxwell (2013) describes qualitative research questions as built on "process theory" because they are best answered by understanding how events, situations, decisions, and actions unfold and influence one another over time. This helps explain why qualitative researchers tend to ask more descriptive *how* and *why* questions:

How did a person come to understand an experience in a particular way?

How did a specific practice emerge in a community?

Why did a particular decision make sense in context?

How did one event shape the course of another?

In contrast, quantitative-leaning research questions are more often guided by what Maxwell (2013) calls "variance theory" by identifying the most important variables and trying to understand their relationships as precisely as possible while eliminating distracting or "noisy" information. From this tradition, more predictive or explanatory questions may be more valued:

Are two variables related? How strong is their relationship?

Which variables might best explain or predict an outcome?

Qualitative research is also more likely to locate the observer (you, the researcher) in the world as a (hopefully intentional and reflective) influence on what is collected and how that shapes results (Denzin & Lincoln, 2011; Denzin et al. 2024). Because of this, data sources are often chosen and designed as a way to make a phenomenon visible by turning it into representations:

interviews, field notes, conversations, photographs, recordings, documents, performances, memos, and others. Quantitative research also does this, but the representations tend to be more standardized, abstracted, and designed to remove unnecessary "noise" from the exploration and ultimately enhance comparability.

Recognizing that every step of the research process is shaped by researchers' choices makes qualitative research an inherently *interpretive* practice, as opposed to, for example, a transparent window through which we might view a single shared reality. There is a long-standing body of scholarship arguing that no phenomenon can be represented through a single view or understanding (Goffman, 1956; Mead, 1934). Researchers in qualitative traditions are less likely to see themselves as passively measuring a phenomenon, but instead as active participants who engage with, represent, and ultimately make sense of a phenomenon through theoretically and socio-historically situated processes. In this way, the research process itself becomes a part of how understanding is constructed.

In sum, the qualitative research umbrella is formed from three general and overlapping principles:

(1) An emphasis on *social meanings*, or research phenomena as a set of human interpretations, perceptions, beliefs, and lived experiences.

(2) An understanding of phenomena as *complex* and multi-faceted, including how actions, relationships, and events unfold over time and within specific settings.

(3) Intentional consideration for the impact of *context*; the idea that human experience is socially situated, historically shaped, and that some questions are not fully answered by methods that reduce, simplify, or streamline the data.

These goals make qualitative research particularly valuable when the phenomenon of interest is complex, ambiguous, or context-dependent.

This does not mean that qualitative research is defined only by opposition to quantitative work (or vice versa), nor that qualitative research is unconcerned with the rigor, explanation, systematicity, or drive towards broader scholarly understanding that characterize the best quantitative work (Lakatos, 1978). Qualitative research can be (and typically should be) systematically theoretical, analytically precise, and deeply rooted in practices aimed at establishing the trustworthiness and quality of findings. But its rigor is typically directed toward understanding meaning, process, and situated human experience in ways that preserve some of the texture of lived reality and individual experience rather than abstracting too quickly away from it. For these reasons, qualitative research has become central across a wide range of

fields concerned with representing human life, social systems, culture, institutions, and experience, in fuller depth and richness.

In the sections that follow, we build from these goals to consider the values that often underlie qualitative inquiry, the major areas and traditions which have developed, and the range of methods through which qualitative researchers pursue these aims.

Core forms of qualitative research

Discourse and Conversation Analysis

Discourse analysis and conversation analysis are two related but distinct areas of research that focus researchers' attention on the intricacies of language. Gee (1999) defines discourse analysis as the study of "language-in-use," which often involves a deep dive into the specific ways individuals used the content or structure ("grammar") of language to make meaning in social contexts. Conversation analysis focuses even more specifically on discourse as interaction, and the ways that multiple actors co-construct meanings through discourse (Sacks, Schegloff, & Jefferson, 1974).

Closely linked to these fields is work on semiotics (Nöth, 1990), which is oriented around the idea that there is no symbol of communication (e.g. spoken word, written symbol, gesture) that carries an inherent or universal meaning. We, as part of our social contexts, instead routinely define, negotiate and re-negotiate shared agreements about what symbols mean. In one context, a head shaken side to side could mean no or a negative response, but could indicate yes or affirmation in another context. By examining how individuals and groups use talk, text, interaction, and other visual symbols as sites of social action, scholars can answer questions about how people perform identities, negotiate relationships, and organize social realities. Conversation analysis often focuses on the finer-grained structure of interaction (e.g., turn-taking, repair, sequencing) (Sacks, Schegloff, & Jefferson, 1974), while broader discourse analysis can sometimes focus on larger systems of meaning such as narratives, ideologies, or institutional discourse (Gee, 1999, Ch. 2). These approaches may rely on naturalistic (naturally occurring) data collection, where data are gathered in real-world settings rather than environments designed or structured by the researcher. Discourse and conversation analysis may be used to answer questions like:

How do participants use language to construct meaning in a particular social context?

How are identities, relationships, or roles enacted and negotiated?

How do participants co-construct meaning through interaction?

How do broader systems of meaning manifest in shared narratives, norms, or ideologies?

Grounded Theory and Situational Analysis

In its purest form, grounded theory is the use of systematically collected and analyzed data to "discover" broader theory (Glaser and Strauss, 1967/1999). The use of a more specific data source to build or inform a more general theory (sometimes referred to as "bottom-up," data-driven, or inductive work) can be found in both qualitative and quantitative research (cf. Behrens, 1997); within this paradigm, doing so holds value as a way to help unify qualitative studies that are often context-rich, complex, and smaller-scale by nature around broader, cross-cutting lessons learned. Rather than beginning with a predefined framework, researchers engage in cycles of data collection and analysis, using processes such as coding, constant comparison, and memo writing to identify patterns and build toward increasingly abstract conceptual categories (Glaser & Strauss, 1967/1999; Charmaz, 2014). As analysis progresses, these categories are refined and related to one another to form a theoretical account that is grounded in the data.

Situational analysis, which is common in sociology, science and technology studies (STS), and public health research, extends grounded theory to focus on an entire "situation" as the unit of analysis, rather than focusing on individual actions and processes (Clarke, 2005). To broaden the focus of theory-building, researchers often create and iterate on (1) *situational* maps that lay out all relevant elements, (2) *relational* maps that examine connections among them, (3) *social worlds/arenas* maps that capture collective actors and sites of negotiation, and (4) *positional* maps that represent competing perspectives in the discourse (Clarke et al., 2015). This mapping process helps illustrate how people, institutions, discourses, material conditions, and nonhuman actors work together to shape a phenomenon. Researcher reflexivity on their perspectives, positions and decisions is particularly important in situational analysis because it shapes the researcher's interpretive maps.

Grounded theory and situational analysis may help answer questions like:
What patterns emerge from a participant's experiences?

How do actors, structures, and discourses interact to shape a particular situation?

How do individuals navigate, interpret, or respond to complex social situations over time?

How are different perspectives, positions, or voices represented, negotiated, or marginalized in a given context?

Thematic Analysis

Thematic analysis is a systematic approach to identifying themes that repeatedly appear in a qualitative dataset, and sometimes also the examination of their patterns of appearance in the dataset (Braun & Clarke, 2006). Unlike grounded theory, which emphasizes more purely data-driven (inductive) themes, thematic analysis guides researchers to develop only the themes that are relevant to their evolving, theoretically-grounded research questions (Braun & Clarke, 2013). Thematic analysts begin the process with a more direct understanding of what questions they wish to answer and what theories or bodies of research are guiding them, and let these perspectives inform the types of themes they construct in service of these goals. Grounded theorists, in contrast, may let the data "tell" them which themes are present so that these themes may later coalesce into a broader theoretical understanding of how a phenomenon happens. Qualitative themes that are developed using thematic analysis are often data-driven (inductive) but also framed around or informed by existing theories or frameworks (deductive). This hybrid integration of inductive and deductive practice is sometimes referred to as indo-deductive or abductive (Tavory & Timmermans, 2014).

The kinds of themes (sometimes called codes) that emerge from a thematic analysis can range from more explicit, concrete or semantic-based (e.g., "greetings" "gratitude") to more implicit, abstract, or latent patterns (e.g., "motivation" "confusion"), supporting analyses that range from describing participants' experiences to critically interpreting broader systems of meaning. The systematic ways that themes may be applied to the data (through processes like line-by-line coding) often make tests of analytical rigor such as inter-rater or test-retest reliability more possible than in other qualitative paradigms (Braun & Clarke, 2013). Researcher reflexivity and transparency around the ways themes were constructed are again important in this type of analysis to explain how and why each theme was developed (Braun & Clarke, 2019; 2021; 2023). Thematic analysis may help answer questions in a wide variety of research contexts, such as:

How do participants make sense of a particular phenomenon?

What recurring values/beliefs/concerns/etc. are reflected across different research contexts?

What themes characterize participants' perspectives, beliefs, or practices related to a phenomenon?

Qualitative Content Analysis

Qualitative content analysis is a set of techniques used to analyze textual, visual, or multimedia data by systematically identifying and interpreting patterns of meaning (Mayring, 2000; Hsieh & Shannon, 2005). While thematic analysis is concerned with identifying themes and patterns of

meaning within and across a dataset, qualitative content analysis typically places greater emphasis on the systematic development, application, and comparison of coding categories across data sources (Schreier, 2012). While some more quantitative forms of content analysis focus on counting frequencies, qualitative content analysis emphasizes the interpretation of meaning in context while still applying structured and transparent coding procedures (Krippendorff, 2019; Schreier, 2012). It is commonly used in fields such as education, communication, and health research to examine documents, transcripts, open-ended responses, and digital content.

In practice, qualitative content analysis involves developing a set of categories or codes that represent key ideas, themes, or constructs in the data. These codes may be derived inductively from the data itself, deductively from existing theories or frameworks, or through a combination of both approaches. Researchers then apply these codes systematically across the dataset, often revising and refining them through iterative analysis. The process typically includes organizing coded data into categories, comparing patterns across cases, and interpreting how these patterns relate to the research question. Because of its structured approach, qualitative content analysis often supports transparency and consistency in coding, making it well-suited for studies that involve multiple coders or require careful documentation of analytic decisions.

Research questions may ask:

What patterns or themes are present within a set of texts or communications?

How are particular ideas, concepts, or issues represented across a dataset?

How do patterns of meaning vary across participants, groups, or contexts?

Theory-driven Analysis

Theory-driven analysis refers to qualitative research that uses existing theory, constructs, or frameworks to design, implement, analyze, or interpret a complex phenomenon. Rather than turning to the qualitative data to identify what themes are present, as in grounded theory or thematic analysis, theory-driven analysts use their own or other scholars' prior theory-building to guide data coding, interpretation, and explanation of their current work (Miles et al., 2014). Because of this, deductive (Boyatzis, 1998), or top-down approaches may offer qualitative themes that exist in the theory, but may not appear in the dataset (Chen, 1990), which might be useful to researchers who wish to assess how comprehensively a framework is represented in a specific context (for an example, see the second author's use of theory-driven codes to evaluate how well course design supported STEM identity development; Barany et al. 2021).

Theory-driven analysis also allows researchers to test, extend, or refine theoretical ideas in new contexts and with new data sources. This kind of work is also useful for design-based research studies, where researchers might iteratively design and refine an intervention for real-world contexts, while using theory to guide the evaluation of each design cycle and suggest design decisions for the next round of improvements (Design-Based Research Collective, 2003). While scholars may need to be cautious to avoid important insights about a specific data set or context being left out of a more targeted, theory-driven analysis, this approach offers conceptual clarity that may help answer research questions such as:

To what extent do existing theoretical constructs explain participants' experiences of a phenomenon in context?

How are key concepts from a theory reflected, challenged, or extended in the current data?

In what ways does applying a theory to this dataset support, refine, or complicate the theory's assumptions about a process or behavior?

Ethnography and Quantitative Ethnography

Ethnographic research emerged as one of the earlier forms of qualitative research, with foundations in the field of anthropology, as a way to systematically document and understand the cultural practices, beliefs, and everyday experiences of communities within their natural contexts (Malinowski, 1922/1978; Geertz, 1973). Given this goal, particular emphasis is placed on the methods for collecting data, and the types of data that are collected as documentation of important phenomena. Early ethnographers often immersed themselves in cultural settings for longer periods of time, engaging in participant observation and collecting field notes and artifacts that were created by participants in the context (Mead, 1934). These practices are also common in sociology and the learning sciences, where an understanding of how processes unfold naturalistically may be useful. In line with other qualitative practices mentioned in this chapter, ethnographic work prioritizes the importance of context, the situated nature of meaning, and the researcher's role in interpreting and representing lived experience (Geertz, 1973; Hammersley & Atkinson, 2007); it is distinguished by its commitment to making participants' own meanings and perspectives central (Malinowski, 1922/1978).

While traditional ethnography involves fieldwork, contemporary ethnography often applies the same principles and priorities to contexts in which immersed data collection may be less feasible. For example, many digital spaces now offer massive repositories of user artifacts (e.g., posts, writing, recordings) and detailed secondhand records of activity in context (e.g. log data of clicks or other virtual actions) (Hine, 2020; Kozinets, 2019). In these cases, researchers may

not be able to “witness” practices in real time, but can carefully reconstruct contexts of action to build their own ethnographic understanding.

As a result, new fields have built on these ethnographic traditions but incorporate some quantitative techniques to help craft these ethnographic reconstructions. Quantitative Ethnography (QE) emphasizes the systematic coding and structuring of qualitative data in ways that make patterns of connection visible and analyzable using computational techniques (Shaffer, 2017). Central to QE is the idea that meaning is reflected by modeling the complexity of a situated context, often using techniques such as epistemic network analysis to show how different culturally-embedded themes are connected in discourse across individuals, groups, or time periods (Shaffer et al., 2016). All data structuring and model building are informed by the researchers’ data and theory-driven understanding, with the goal of using quantitative approaches to enhance ethnographic understanding rather than to replace it.

Ethnographic and Quantitative Ethnographic research questions may look like:

How do complex social processes (e.g., learning, participation, problem-solving) unfold over time in naturalistic settings?

How do participants navigate and respond to challenges within a given context?

How do interactions between individuals and their environment shape those individuals’ meaning-making and involvement?

What patterns of connection among ideas, practices, or behaviors emerge across participants or over time?

Case Studies

For case studies, the scope of inquiry is more intentionally bounded and specifically defined: the in-depth exploration of a *case*. A case is a single unit, but that unit can be a variety of things: an individual, a group, a program, an event, or even a system. Within the case study method, researchers not only examine the characteristics of the case, but also the social, organizational, historical, or environmental contexts in which it is embedded. The selection of the case may be pragmatic (e.g., the only option available), intrinsically interesting or valuable, or instrumental for understanding the broader phenomenon of interest (Stake, 1995). Regardless of the motivation for sampling, the goal is to typically to build a detailed and holistic understanding of the case, often by triangulating (using overlapping evidence to confirm or expand findings) across multiple sources of data (e.g., interviews, observations, documents, artifacts, digital trace data) to capture different perspectives and dimensions (Yin, 2018; Stake, 1995). Case studies may be

(1) exploratory, for phenomena that are not yet well understood and to inform more targeted future work, (2) descriptive, to provide a rich account of the phenomena, or (3) explanatory, to understand how and why particular events, processes, or outcomes occurred (Yin, 2018); also, though the case study method is about studying cases in their individuality, a broader study may compare multiple cases to identify their similarities and differences. Case study research questions may ask:

How/why does a particular program, intervention, or experience unfold in a specific case?

What factors shape participants' experiences and outcomes within a case?

In what ways does the case reflect, challenge, or extend existing understandings of the broader phenomenon of interest?

Narrative Inquiry

Narrative inquiry centers on the study of stories as a way of understanding how individuals make sense of their experiences and identities over time. Researchers examine how narratives are constructed, told, and situated within broader social and cultural contexts, often attending to elements such as structure, voice, temporality, and meaning-making (Riessman, 2008; Clandinin & Connelly, 2000). Narrative inquiry treats personal experience as something that is actively organized and interpreted through narratives and the processes used to create them —what is included or omitted, how events are sequenced, and how meaning is shaped through storytelling.

Data may include interviews, written accounts, and increasingly digital artifacts such as social media posts. Analysis focuses on how experiences are organized and interpreted through narrative form; in some approaches, the analytic process itself takes a narrative form, with the researcher constructing a coherent account that synthesizes and represents the data (Polkinghorne, 1995). Some approaches emphasize the structure of stories (e.g., how narratives are organized and sequenced), while others focus on the meanings conveyed across stories or the ways narratives are co-constructed in interaction (Riessman, 2008). Narrative inquiry highlights the role of storytelling in shaping both personal and collective understandings of the world, including how individuals position and think about themselves and others within broader cultural or institutional narratives (Bruner, 1990).

Research questions may ask:

How does an individual construct and revise their identities through narrative over time?

How do cultural or institutional narratives shape personal storytelling?

How do individuals position themselves and others within their stories?

How do individuals' stories about a particular experience change as they retell them over time, or to different audiences?

Critical Theory and Critical Approaches

Critical theory informs a set of qualitative approaches that use assessments of social interaction, context, and complexity for a targeted purpose: to illuminate and challenge systems of power, inequality, and social injustice as they are produced and reproduced in society (Horkheimer, 1972; Freire, 1970). The cornerstone of critical approaches is to move beyond description or interpretation toward an examination (and where needed, critique) of broader structural conditions and societal ideologies. As a result, critical theory views research as never neutral, but as inherently political, either reinforcing or disrupting existing power relations -- and critical theory work itself typically takes the frame of being intentionally political, towards disrupting unjust power relations.

Critical research often focuses on how systems of oppression (e.g., racism, sexism, classism, other forms of marginalization) are embedded into everyday practices, institutional structures, and dominant forms of knowledge. Instead of focusing on exploring and explaining what participants say and do in context, critical researchers also attend to whose voices are represented and excluded, and how the broader discourse and context shape what knowledge and practice is considered “legitimate.” Two practices are particularly emphasized in critical research as a way to avoid reifying the same systems of oppression. The first is participatory research practices that situate participants as stakeholders in the research process. The second is robust and ongoing reflexivity on the part of the researcher on their positionality in relation to their work (Denzin & Lincoln, 2018). While critical theory works to shift the goals and interpretive lens of qualitative inquiry, many of the same practices and approaches remain useful, including interviews, discourse analysis, and ethnography.

A study adopting critical approaches may ask:

How are systems of power and inequality produced and maintained within a particular context?

Whose perspectives are represented, marginalized, or excluded in this setting?

How do participants experience and navigate structures of power, privilege, or oppression?

Participatory and Community Action Research

In line with critical approaches that examine power, voice, and inequality, participatory and community action research emphasize collaboration between researchers and the individuals or communities that are the focus of study, positioning participants as active contributors in shaping research at a variety of stages. Community action research of this nature is often intended to generate outcomes that are directly useful to the participating community, such as informing practice, supporting advocacy, or improving local conditions. Participants and communities co-develop and co-enact (1) project conceptualization and the formation of research questions, (2) collaborative data collection and analysis, and (3) practical uses of findings (Lewin, 1946; Freire, 1970). This emphasis is further developed in related paradigms such as decolonizing methodologies, which center community members as co-leaders whose questions, methods, and purposes of inquiry are treated as central to the entire endeavor (Tuhiwai Smith, 1999). Including participants and communities as core members in the research process is done with the goal of enhancing the prospects of meaningful and systemic change by centering the knowledge, experiences, needs, and priorities of those who are most directly impacted.

Methods used by participatory and community action researchers vary widely but often prioritize jointly created research artifacts such as interviews, focus groups, participatory design sessions, co-analysis workshops, and community forums. Researcher reflexivity and awareness of the impacts of positionality are again treated as being essential to avoid reifying existing power dynamics.

Participatory and community action research teams may ask:

What issues, needs, or priorities do community members identify as most important?

How can researchers and community members work together to better understand and address a shared concern?

How do participants interpret and use research findings in ways that are meaningful to their context?

How do community or cultural frameworks and ways of knowing drive both the questions we ask, and how we answer them?

Phenomenology and Interpretative Phenomenological Analysis

Phenomenological approaches are used to examine how individuals experience and make meaning of phenomena (Husserl, 1970; van Manen, 1990). The approach first emerged in philosophy as a way to ground the study of knowledge and meaning in the lived and conscious experiences of those who directly experience it. Scholars in education, psychology, and health research later adopted phenomenological research to develop deep, richly textured understandings of how a phenomenon is experienced by and across individuals (Giorgi, 2009; Finlay, 2011).

Phenomenological research typically involves in-depth analysis of participants' descriptions of their experiences through interviews, written reflections, or other first-person accounts. Researchers look for similarities and differences in how participants describe and make sense of their experiences, with the ultimate goal of building a broader understanding of the phenomenon of interest. Similar to some of the atheoretical and data-driven analysis conducted by grounded theorists, phenomenological researchers often engage in "bracketing" and attempt to set aside any prior assumptions or biases in order to focus more closely on participant perspectives (Moustakas, 1994). While phenomenology shares some similarities with narrative inquiry, the two differ in emphasis: phenomenology seeks to understand the essential structure of an experience as it is lived, whereas narrative inquiry is more concerned with how individuals construct and sequence meaning through the stories they tell over time.

Phenomenological approaches are often used to build broader understanding from multiple individual deep-dives; Interpretative Phenomenological Analysis (IPA) by contrast maintains focus on individuals and their specific patterns of sense-making. As a result, IPA favors a more detailed analysis of a small number of cases, with attention to both what participants say and how they describe or reflect on experiences (Smith et al., 2009). Common research questions across both Phenomenology and IPA may ask:

What are participants' lived experiences of a particular phenomenon?

How do individuals make sense of and interpret their experiences in a specific context?

What common patterns or themes characterize how a phenomenon is experienced across participants?

How do participants' differences in perspective or context shape the ways they understand or describe their experiences?

Fields where Qualitative Research Happens

Education

Qualitative research in education typically focuses on understanding learning and teaching practices, and the ways that related systems and processes unfold. Some education research emphasizes themes such as understanding the ways students learn, including how students make sense of content, how learning identities and motivations develop, or how social interactions shape learning experiences (Vygotsky, 1978; Gonzalez et al., 2005; Gee, 2000). Other research in education explores best practices and the impacts of designed teaching and learning interventions (Bransford et al., 2000), including the design of curricula or educational tools and technologies (Collins, 1992; Barab & Squire, 2004). A third element of education research focuses on best practices for formal and informal teaching, including practices for training new teachers (Darling-Hammond, 2006). Data collection methods range from more naturalistic approaches such as classroom observations to more active data collection such as interviews or focus groups. Analysis of pre-existing artifacts of teaching and learning are also common, especially in online learning spaces where detailed records of reflection and action accumulate (Wise et al., 2014). Some education research focuses on the design of technological or pedagogical improvements, including participatory co-design, development, implementation and testing of new instructional designs, curricula, and educational policies (Design-Based Research Collective, 2003).

Given this diversity, research studies in education may draw on (and in some cases combine) elements of many different qualitative traditions (Creswell & Poth, 2016). Learning and its related processes (e.g., motivation, self-regulation) tend to be complex, context-dependent concepts that are hard to directly observe. As a result, researchers in the learning sciences may draw more heavily on relevant learning theories to frame what they are looking for, how they look for it, and how they interpret what they find over the course of a research study. In more recent work, particularly in the field of learning analytics, qualitative methods may be combined with computational and quantitative approaches to better understand learning processes at scale, while maintaining sensitivity to context and learner experience (Krumm et al., 2018).

Sociology

Qualitative research in sociology (perhaps unsurprisingly) asks questions about social phenomena, including how social structures, institutions, and cultural systems shape human experience and interaction, and how individuals and groups in turn construct, negotiate, and sometimes challenge the social order (Burawoy, 1998). While some sociological research emphasizes smaller-scale interactions and processes of meaning-making (e.g., social identities, roles, relationships) (Goffman, 1956; Blumer, 1969), other studies focus on larger-scale social phenomena such as education, work, family structures, and the ways these systems may produce and reproduce patterns of inequality, power, and opportunity (Bourdieu, 1984). Some

sociological research also explores the relationship between discourse and culture, revealing how shared sets of habits, skills and styles can help people craft strategies of action (Geertz, 1973; Swidler, 1986).

Data collection in sociology research, like education research, may include sources that are researcher-generated (e.g., observation field notes, interviews, focus groups), or naturally occurring and secondary sources (e.g., documents, media, cultural artifacts). Much sociological research relies upon naturalistic observation of real-world settings, which allow researchers to understand the ways social processes unfold over time and in context. Like many fields, sociological research increasingly examines digital spaces as well, most notably social media and other online communities (Lupton, 2015).

Public Health

Qualitative research in public health answers questions about health, illness, and well-being as concepts that are socially and culturally situated (Kleinman, 1980; Pope & Mays, 2006; Scheper-Hughes & Lock, 1987). These studies may stand alone, or may complement epidemiological assessments of health outcomes, focusing on individuals and community experiences with health and healthcare and the ways that broader social, cultural, and structural factors impact health behaviors and outcomes (Kleinman, 1980; Marmot, 2005). Examples include research on the ways people understand illness (Good, 1994), how they navigate healthcare systems (Dixon-Woods et al., 2006), and how they make decisions about prevention and treatment (Conrad, 1985). Other work examines larger systems, including healthcare institutions, policies, and social determinants of health, and how these systems contribute to disparities in access, care, and outcomes (Farmer, 2004; Braveman & Gottlieb, 2014).

Within public health and health services research, qualitative methods may be used to study health and to study to improve health. The former includes investigations of how concepts such as illness, risk, or care are constructed and experienced. The latter refers to applied studies that are designed to improve interventions, programs, and service delivery. Qualitative research in healthcare and clinical contexts focuses on understanding patient experiences, clinical interactions, and the social dynamics of care delivery. Researchers often examine how individuals experience illness, navigate healthcare systems, and interact with providers, using methods such as interviews, observations, and narrative analysis (Charmaz, 1991; Kleinman, 1988). This work complements population-level public health research by providing insight into how health and care are experienced at the individual level, helping to inform more patient-centered, context-sensitive approaches to treatment, intervention design, and healthcare delivery.

Political Science

Qualitative research in political science is often framed around the goal of building broader theoretical understanding of political processes, institutions and systems, using detailed contextual analysis to understand how and why political phenomena unfold in the ways that they do — questions that are difficult to answer through quantitative analysis alone (George & Bennett, 2005). Substantive areas of inquiry range from the workings of specific institutions and policy processes to broader questions about democratization, political mobilization, social movements, and conflict. Common approaches in qualitative political science research include case studies, historical analysis, comparative research, and process tracing, which uses case studies to reconstruct sequences of events that can reveal broader causal mechanisms, to understand how political outcomes emerge (King et al., 1994). Researchers may draw on archival and documentary sources, interviews, ethnographic fieldwork in political settings, or comparative analysis across cases selected for their theoretical leverage. Qualitative research within political science often emphasizes careful case selection and clear links between observations and broader theoretical claims. Scholars may examine how institutions evolve, how policies are implemented, or how political actors make decisions by tracing event sequences within those contexts (Pierson, 2004).

Human-Computer Interaction and Design Research

Qualitative practices in Human-Computer Interaction (HCI) are used to explore the design and use of a current or future technology (Blandford, Furniss, & Makri, 2016). HCI often overlaps with design research, which explores how new knowledge is constructed by building and studying the designed object (e.g., a tool, a system, a concept) as well as the process used to develop it (Koskinen et al., 2013). For example, a researcher might develop a tool for tracking patients' preventative health practices, study how patients and clinicians use it in practice, then use those interactions to refine the system and develop new insights about care practices and patient decision-making.

In both fields, design is treated as an ongoing, iterative process that should be shaped by users, contexts, and emerging qualitative insights. Researchers are typically interested in how technologies fit into everyday practices, how users engage with design features, and how design decisions shape user experiences (e.g., usability, engagement, accessibility). To guide the research cycles, qualitative methods are often embedded directly into the design process at multiple points. Design and usability researchers might interview participants as they conduct their regular work activities (Beyer & Holtzblatt, 1997), observe users interacting with a design prototype in a laboratory setting (Nielsen, 1994), or analyze artifacts such as interaction logs, prototypes, or user-generated content (Hutchinson et al., 2003). Participatory and co-design

approaches are also common in both fields to increase the likelihood that design meets user needs and goals and to ensure that solutions are evaluated on the right dimensions (Bødker et al., 2004). Qualitative studies in HCI are often considered to be “semi-structured” in the sense that they balance clear design goals with flexibility to explore unexpected findings as they emerge using a variety of techniques (Blandford, Furniss, & Makri, 2016).

Psychology

Qualitative research in psychology aims to understand individual experiences, thoughts, and meaning-making processes. This emphasis on understanding the individual draws heavily on phenomenological research and case study research (Smith, 2004). In some cases, individual experiences may be aggregated or compared to understand broader patterns around why people think, feel and act in the ways that they do; grounded theory approaches may be used to support this data-driven theory-building. Qualitative studies in psychology often help psychologists answer more exploratory questions about the ways people understand their own psychological experiences or the ways that social, cultural, and relational contexts shape psychological processes (Willig, 2013).

Some qualitative psychology research emphasizes lived experiences, such as the ways individuals navigate mental health or major developmental events (Smith, 2004), while other work examines how specific psychological constructs such as motivation, self-regulation, and emotions are experienced and enacted in everyday settings (Hidi & Renninger, 2006; Zimmerman, 1989). Qualitative research practices are particularly common in clinical, counseling, and social psychology, where researchers tend to explore phenomena that are harder to directly observe or assess through more controlled clinical or laboratory studies.

Anthropology

Building on traditions of long-term ethnographic immersion in a context (Geertz, 1973), qualitative research in anthropology focuses on analyzing cultural practices and meanings using rich, context-specific qualitative data. Anthropological research, like the closely related practice of ethnography, typically includes explicit dialogue with members of the group being studied, including practices such as checking interpretations (Spradley, 1979; Hammersley & Atkinson, 2007). Anthropological practices also extend beyond immersion alone to include the collection and examination of a variety of data sources, including observations (Mead, 1928; Malinowski, 1922; Evans-Pritchard, 1940), interviews and oral histories, and visual or written artifacts. These data sources are used to construct detailed accounts of participants’ social beliefs, values, rituals, and norms.

A key goal of anthropological research is to develop what Geertz (1973) describes as “thick description,” in which specific practices are interpreted in relation to the broader systems of meaning and social organization in which they occur. Anthropological findings tend to be highly contextual and resist broader generalization, with the goal of highlighting underlying cultural logistics, social relationships, and systems of value at play.

a. What is artificial intelligence?

Artificial Intelligence (AI), broadly defined, is used to refer to machines or other technologies that can perform tasks that are thought to require human intelligence (Russell & Norvig, 2021). This encompasses a wide range of capabilities — from recognising patterns in data, to understanding language, to making decisions. AI is not a single technology but rather an umbrella term covering a diverse and rapidly expanding set of approaches, each with its own strengths and limitations. The field developing these technologies has evolved significantly over the decades in both its ambitions and its methods.

The history of AI can be understood through three broad eras, each representing a significant shift in how machines are taught to "think." The first era, often associated with expert systems (Feigenbaum, 1977) in the 1970s and 1980s, involved researchers manually encoding human knowledge into rigid rule-based programs — if a patient presents with symptom X, then consider diagnosis Y (Shortliffe, 1976). These systems worked well within narrow, well-defined domains, but struggled to handle the messiness and ambiguity of the real world. Other models in this area attempted to reason logically over well-defined operators, in order to solve unseen problems (Newell & Simon, 1961), but were difficult to translate to real-world problems. The second era, machine learning, involved machines learning patterns directly from large quantities of data. A machine learning system is not explicitly told *what* to look for, but is instead exposed to examples and gradually learns to identify regularities on its own (Mitchell, 1997). The third and most recent era centres on foundation models — vast models trained on enormous datasets that can be adapted to a wide range of tasks (Bommasani et al., 2021). The most prominent examples are large language models (LLMs), such as the GPT and Claude families of models, which have been trained on billions of words of text and can generate, summarise, translate, and reason about language with applicability to problems never considered by the developers of the first LLMs. The most recent large language models build a reasoning layer that harks back in many ways to earlier problem-solving AI, on top of massive statistical models, to reason in complex fashions (Guo et al., 2025).

b. The explosion of data

Recent advances in artificial intelligence have been driven in large part by an explosion of data. Large language models got their start (somewhat obviously) with textual data. AI platforms built their models by ingesting the text on the public internet, as well as a great number of books. Textual data in general has exploded in quantity in recent years, including considerable amounts of data of interest to qualitative researchers: human-human conversations, human-AI conversations, posts on social media platforms, government meeting transcripts, and many more sources. Audio, visual, and video data have also become massively more prevalent and accessible, with sources such as classroom recordings, video platforms, and born-digital cultural artifacts. More traditional data files have also become available at scale for many types of data that were previously unavailable, such as school information data systems, survey open response data, structured field observations, and administrative records.

This expanded availability of data has proven to be a boon for quantitative researchers (cite EDM, another cite), enabling a broad range of projects investigating research questions that were unimaginable a generation ago. Massive data sets allow researchers to conduct in-depth studies of rare behaviors and phenomena, to track changes in individuals over time, and find uncommon combinations of factors -- all of which are difficult to study in small-scale data.

The growth in data sources and emergence of new analytical methods has created tensions around what constitutes "qualitative" research. Traditionally, in-depth interpretive study of unstructured data—data from interviews, documents and transcripts, audio and visual data, images, field notes—has defined qualitative research, to the point where these data types themselves become labeled by many researchers as "qualitative data". However, many of these types of data have long been possible to analyze quantitatively, an approach which has recently become much easier with large language models and multimodal large language models (as the upcoming chapters will discuss).

A narrower conception of “qualitative data” defines it not by modality, but by its provenance or analytic treatment (Miles et al., 2014; Saldaña, 2016): data generated through qualitative research processes (e.g., ethnography, memoing), or data analyzed using qualitative research methods. This perspective does not claim that particular data modalities are inherently qualitative; instead, it foregrounds how data are produced and interpreted. At the same time, it opens analytic space for materials that have traditionally been treated as quantitative—such as software interaction logs or administrative records—to be examined qualitatively as well.

This is the perspective we take in our book -- that a wide range of data types are amenable to qualitative analysis, and that contemporary tools make it increasingly feasible to apply interpretive, context-sensitive qualitative inquiry approaches even to large-scale datasets.

c. The potential of artificial intelligence for qualitative research: an initial take

In recent years, there has been great enthusiasm about the potential of AI to revolutionize a wide range of areas of daily life, work, and scholarship. Qualitative research is no exception, with scholars waxing enthusiastic about new possibilities to conduct inquiry more efficiently, more comprehensively, more effectively. Qualitative research, for all of its virtues and successes over the last century, is not perfect. It can be slow (Bazeley, 2013). It can be difficult for qualitative researchers to find less common patterns in data (Friese, 2014). Findings tend to be difficult to generalize across projects (Yin, 2018). While a source of unique perspective, researchers' subjectivities and positionality can also limit what they can discover (Berger, 2015). AI can help qualitative researchers to overcome these challenges, and many more that we are still discovering.

At the same time, many reservations have been expressed about the risks and downsides of AI. For example, over 400 scholars and practitioners of qualitative inquiry signed on to a 2025 open letter asserting that "1. GenAI as simulated intelligence is incapable of meaning-making; 2. Qualitative research should remain a distinctly human practice; 3. The established manifold harms of GenAI, especially to the environment and workers in the Global South." (Jowsey et al., 2025, p.1). Perspectives beyond the qualitative research community have produced treatments of concerns about potential negative impacts of AI on the future of work (Eloundou et al., 2024), the future of the environment (Strubell et al., 2019), the future of humanity (Bostrom 2015).

In this book, we take an optimistic view towards the potential of AI to enhance qualitative research, while attempting to take seriously the limitations and negatives of AI as well. Yes, there are limitations and challenges -- but limitations and challenges emerge with all new tools, with all new approaches. We ask the question of how we can mitigate and work around these limitations and negatives, rather than rejecting this new technology wholesale as Jowsey and colleagues demand. In doing so, we join other scholars in calling for thoughtful consideration of the pluses and minuses of generative AI in qualitative research, such as the signers of a second open letter (Friese et al., 2025). Our perspective is shaped by some of the discussion in Friese et al. (2025) but ultimately sees a broader role for the use of AI than they do; we will return to their perspectives and critiques throughout this book, in some places through direct citation and in some places influenced more indirectly.

Within this book, we will refrain from in-depth discussion of the risks of AI that are beyond our expertise or the focus of this book, such as concerns about the future of humanity, or about the future of the environment (but see Laughing & Ruvalcaba, 2025 for critique of the claims around water use made in Jowsey et al., 2025). By contrast, our discussion in Chapter 4 about the potential ways that humans and AI might collaborate to conduct qualitative research inherently

connects to deeper consideration about the future of work. We return to such themes throughout the remainder of this book. We disagree with the view that AI must be rejected wholesale lest it replace us. Instead, we view the potential uses of AI as a tool, as a partner, as a critic or provocateur, to present enormous possibilities for the thoughtful qualitative researcher -- we view them ultimately as enormous possibilities for the human race.

Why this book?

We decided to write this book because there has been a great deal of interest in the use of generative AI within qualitative research, both enthusiastic and condemning. However, this interest just scratches the surface of what we believe is possible. Most of the enthusiastic work and writing invests its attention in a single goal: automating qualitative coding (Chapter 13). Most of that work uses AI to code, but most of that work remains idiosyncratic and project-specific (as we will discuss in that chapter in more detail), and does not deeply investigate where and how it works, much less *why* it sometimes works and sometimes fails to work.

Several communities are actively discussing the potential of AI for qualitative research. Researchers from NLP (natural language processing) and AI backgrounds have developed a literature that refers to automated coding as LLM-as-a-Judge (Zheng et al., 2023). Some of the best work in that tradition is rigorous, methodical, and thoughtful (see review in Gu et al., 2024) -- the term LLM-as-a-Judge has however often become used to mean using an LLM to label data with no attention to validity or rigor.

There is an active debate within the existing qualitative research community around whether and how AI should be used for qualitative research (Davison et al., 2024; Chatzichristos, 2025; Friese et al., 2025; Jowsey et al., 2025). At its best, this work imagines new potentials and deeply interrogates the trustworthiness of LLMs for this purpose. However, significant portions of the qualitative community's discussion of this appears uninformed about how AI works, adopting metaphors for how AI works developed based on older generations of AI, such as the stochastic parrot metaphor (i.e. Bender, 2022). The last few years have seen fascinating scholarship in the AI-adjacent modern rationalist community around metaphors for understanding AI (Janus, 2022; Cleo Nardo, 2023). It is disappointingly ironic to see a community that prides itself on deeper inquiry basing its recommendations about AI on metaphors from the GPT-2 era, when the debate within the modern rationalist community draws on postmodernist philosophy to more deeply explain how generative AI simulates reasoning (e.g. discussion of Derrida, 1967 by Cleo Nardo, 2023). We discuss this area of work and how it can influence qualitative researchers in depth within chapter 4.

Additionally, beyond these communities, there has been considerable interest in AI for qualitative and mixed methods research within the quantitative ethnography, learning analytics/educational data mining, and other disciplinary communities. These communities have led efforts to concretely interrogate the wheres, hows, and whys of AI in qualitative research, even as this work has been broadly ignored by both the LLM-as-a-judge and qualitative research communities.

As inter-disciplinary scholars whose work cuts across these communities (though we each have our own positionalities within them), our goal in writing this book is to bring these fields together. We aim to inform our readers with metaphors for AI and its use, leavened by the deepest scholarship in the modern rationalist and qualitative research communities. In doing so, we discuss what we already know, based on the best empirical work in the LLM-as-a-Judge, quantitative ethnography, and disciplinary research communities. Ultimately, our goal is to sketch out a vision for what AI can be in qualitative research and in our work. In a sense, this book can be seen as a proposed research and development agenda for a field of methodologists that have not yet come together -- but should come together -- to use this exciting new opportunity to produce work that is trustworthy and rigorous and takes full advantage of what is possible.

Our goal in doing so is to explore many of the places -- as yet underexplored -- where AI can enhance qualitative research. This may first take the form of solving common limitations in qualitative research by leveraging the different strengths of humans and AI. Certainly, much of the focus on AI in qualitative research thus far has attended primarily to feasibility and speed (as critiqued in both Jowsey et al., 2025 and Friese et al., 2025). But our deeper goal is much more ambitious -- to discover new possibilities and new directions for 21st century qualitative research, a vision where humans and AI work together to do work that surpasses what humans have been able to do unaided. We argue not for abandoning hard-won standards for rigor, trustworthiness, and reflexivity that are already well-established in qualitative practice. Instead, we argue for bringing those commitments to new work, respecting differences in those values across qualitative research paradigms, and ultimately better fulfilling our deep commitments to what is good and what is high-quality than we could, unaided.

We do so now -- writing this book as fast as we can, with a little help from AI and a little help from our friends -- because this moment matters. AI is already being used in qualitative research. It is already being used to supplant qualitative research. Unreasonable extremes -- the categorical rejection of AI, its use as a full substitute for humans, and perhaps worst, covert use to avoid critique -- are emerging. The term LLM-as-a-Judge is being used to refer to practices that actively reject the thoughtfulness and rigor of canonical papers in that area. We

need transparency. We need community norms, negotiated by each community, for what is appropriate practice. We need to understand the full scope of possibility and risk. This is an important moment, when expectations, norms, and possibilities are still open, before we all retreat back to our methodological and intellectual silos. No one book can address all possible perspectives and concerns, but we hope with this book to bring together ideas from a range of communities, and bring that broader picture back to those communities -- so that the work we do, all of us, is the most thoughtful, rigorous, and trustworthy it can be.

Who are we, and why would you want to read a book on this topic by us?

Many of our more senior readers will never have heard of us, and will wonder why they should read a book by the two of us. We would like to briefly introduce ourselves.

Your first author is Ryan S. Baker. I am Professor of Artificial Intelligence and Education at Adelaide University, in South Australia, and have worked at a range of other institutions in the USA and UK. I was involved in founding the educational data mining community (a still very active predecessor to the better-known learning analytics community), which used machine learning and data mining to study complex educational phenomena. Within that community, I have been particularly known for methodological work on how to validly detect complex constructs (such as aspects of engagement and student knowledge), as well as more theory-building and discovery-oriented work. My work has, unusually for that community, been highly mixed-methods -- using field observation to underpin detector validity, and using algorithms to select cases for open-ended interviews. Arguably, over the last two decades, I have been more of a mixed-methods scholar than a quantitative scholar, and have published with several hundred co-authors, mostly educational researchers, co-writing in English, Portuguese, and Spanish. In parallel to this research, I developed a massive online open course on educational data mining that has enrolled well over 100,000 users and persisted for 14 years as of this writing, and conducted practical work that has changed (hopefully for the better) the learning experiences of hundreds of thousands of learners in areas ranging from scientific inquiry skill, to mathematics, to high school dropout prediction. My philosophical positionality is pluralist and pragmatic (as opposed to pragmatist). My scholarly career originated with an orientation towards the critical rationalism of Lakatos (1970) (which, please note, is quite distinct from the modern rationalist community referred to elsewhere in this book) but I have moved towards the belief that many research questions and goals for what we do with that research call for a broader range of theoretical orientations, particularly when different viewpoints and lived experiences bring importantly different lenses.

Your second author is Amanda Barany. I am the Assistant Director for the Penn Center for Learning Analytics in the University of Pennsylvania Graduate School of Education, where my

work combines qualitative and quantitative analytical approaches, and artificial intelligence. I first participated as a researcher in the learning sciences and educational psychology traditions at the University of Wisconsin-Madison, with a focus on game-based learning and student motivation, respectively. Now, I also study the ways that large language models may be integrated into rigorous qualitative and mixed-methods research, which includes applying and testing different hybrid human–AI workflows for developing qualitative codebooks, automating qualitative data coding, and expanding the ways we assess the quality of our results.

Across my work, I find that Quantitative Ethnography (QE) has served as both an intellectual home and a methodological bridge. I treat QE not simply as a set of analytic techniques, but as a vehicle for integrating qualitative nuance with quantitative systematization, supporting the distillation of complex patterns (often in datasets that are too large to deeply study otherwise) without stripping them of meaning. As one of the first-ever QE instructors and QE course designers, and a leader in that research community, much of my scholarship has been collaborative.partnership-driven and often cross-disciplinary. I have collaborated in QE research on digital learning platforms, AI-supported tutoring systems, and game-based environments, but also on projects spanning fields as diverse as computer science, urban planning, environmental science, teacher education, animal behavior, and the biological sciences.

Philosophically, my positionality is pragmatic and pluralist, grounded in the belief that methods are tools to be shaped in service of meaningful educational questions rather than defended as ends in themselves. This means that I am particularly attentive to issues of epistemic fit—whether analytic tools meaningfully align with the cultural, linguistic, and experiential contexts of the learners whose data we study. I love teaching and mentorship as a way to support students and emerging scholars as they navigate diverse and complex methodological choices. My pragmatism is not only theoretical but pedagogical: I care about what works for particular people, projects, and contexts. To that end, I hope this book serves as a practical and empowering resource for researchers—especially students—seeking thoughtful, principled, and adaptable ways to integrate AI into their research goals.

Given both of our deep roots in educational research, it is perhaps unsurprising that many of the examples in this book will involve topics from educational research. Where possible, we have found examples from other fields. We believe that the relatively high educational focus of this book is not unreasonable, however, given what a large proportion of the scholarship on AI for qualitative research has involved educational scholars, particularly when going beyond automated qualitative coding/LLM-as-a-Judge. To our readers in other fields, we beg your patience and indulgence; we'll do what we can to reach out to you given our ambitious writing timeline, and ask you to fill in examples from your own head, where you can.

What the rest of the book is going to look like

Before we conclude this chapter, we'd like to offer a quick guide-post to the rest of the book. In our next chapters, we will discuss foundations: the foundations of qualitative methods and values (ch. 3), the foundations of AI and metaphors for AI (ch. 4), and the foundations of qualitative processes (ch. 5). Chapters 6-17 will each address a single step in the qualitative process and the current and potential uses of AI in that area: what was done before AI, where AI could potentially become useful, what is needed to maintain rigor, trustworthiness, and reflexivity, known concerns, considerations, and traps to beware of, and finally a case study representing some of the best empirical/analytical work conducted thus far.

Not all readers will perhaps be interested in all chapters. Overall, many readers new to this type of work may want to read the book from beginning to end; we've designed it with that in mind. Whichever approach you choose, thank you for joining us! We hope you find the discussion in the upcoming chapters useful, or at least thought-provoking.

Find the most up-to-date version of this chapter's citations at
https://learninganalytics.upenn.edu/book/ai-qualitative-research.html