Building an Apparatus:

Refractive, Reflective & Diffractive Readings of Trace Data

Carsten Østerlund, Kevin Crowston and Corey Jackson

Syracuse University School of Information Studies

Hinds Hall, Syracuse NY 13244

costerlu@syr.edu, crowston@syr.edu, cjacks04@syr.edu

Abstract

We propose a set of methodological principles and strategies for the use of trace data, i.e., data capturing performances carried out on or via information systems, often at a fine level of detail. Trace data comes with a number of methodological and theoretical challenges associated with the inseparable nature of the social and material. Drawing on Haraway and Barad's distinctions among refraction, reflection and diffraction, we compare three approaches to trace data analysis. We argue that a diffractive methodology allows us to explore how trace data are not given but created though construction of a research apparatus to study trace data. By focusing on the diffractive ways in which traces ripple through an apparatus, it is possible to explore some of the taken-for-granted, invisible dynamics of sociomateriality. Equally, important this approach allows us to describe what and when distinctions within entwined phenomena emerge in the research process. Empirically, we illustrate the guiding principles and strategies by analyzing trace data from Gravity Spy, a crowdsourced citizen science project on Zooniverse. We conclude by suggesting that a diffractive methodology may help us draw together quantitative and qualitative research practices in new and productive ways that also raises interesting design questions.

Draft of 28 May 2018

1. Introduction

Information systems have become pervasive platforms for work and life, capturing data about organizational and everyday practices at a fine level of detail (Abbasi, Sarker, & Chiang, 2016; Chen, Chiang, & Storey, 2012). As they are used, systems capture what has been referred to as digital trace data, defined as "records of activity undertaken through an online information system (thus digital). A trace is a mark left as a sign of passage; it is recorded evidence that something has occurred in the past" (Howison, Wiggins, & Crowton, 2011). As opposed to other forms of data commonly used in information systems research (e.g., surveys and interviews, summary data or post hoc reflections), trace data are generated through routine system use, and thus can be used to track events as they unfold over time. In this way, information systems can serve as research apparatuses, instrumenting and capturing data about a wide range of performances. And like all advances in instrumentation, they open new areas of study with vast potential for discovery.

At the same time, trace data raise a number of methodological challenges. First, utilizing trace data demands a deeper exploration of not only the social but also the material performances that go into their production. It is impossible to untangle the data from the technical and often complex nature of the information infrastructures capturing the traces (Hanseth & Lyytinen, 2010). Trace data are typically "big data", with high variety, volume and velocity that pose challenges to analysis. Often heterogeneous and with fine levels of granularity, trace data can include transaction logs, version histories, institutional records, conversation transcripts and source code, to give a few examples. Trace data are most often semi-structured: a mix of structured metadata fields (e.g., for a post in a discussion forum, the date and time, the ID of the poster, the name of the forum, possibly a previous message being replied to, ratings by other readers, etc.) and possibly additional unstructured data (e.g., the subject or content of the post). Equally important, trace data can rarely be accepted as found evidence ready for analysis. Researchers often put significant time into preparing trace data before they can dive into a deeper investigation. Trace data are thus not given but created.

Second, there are a number of different approaches to trace data in recent research, spanning from positivist to interpretive-oriented methodologies. In the big data debate, for instance, many scholars approach trace data as a “lens” into organizational life (e.g., Aiden & Michel, 2014). For example, a number of studies have used posts on discussion fora as trace data of user participation (e.g., Goggins, Galyen, & Laffey, 2010; Yoo, 2010; Phang, Kankanhalli, & Sabherwal, 2009). Emphasis in these studies has been placed on how the traces reflect user behaviors and not how they are created or co-constituted. At the interpretive end of the spectrum, we find trace ethnography and other interpretive approaches that seek to draw qualitative insights into the interactions of users. In this approach, trace data allow researchers to reactively reconstruct specific actions at a fine level of granularity (Geiger & Ribes, 2011; Whelan, Teigland, Vaast, & Butler, 2016; Loukissas, 2017). Once decoded, such traces can be assembled into rich narratives of interactions associated with coordination practices, situated routines or other organizational phenomena. But again, we find an emphasis on how traces reflect interactions and not so much on the production of trace data and its methodological implications.

Third, the lively IS debate around sociomateriality offers promise in facilitating the development of a trace data methodology with its attention to the entwined nature of the social and technical (Orlikowski & Scott, 2008; Cecez-Kecmanovic, Galliers, Henfridsson, Newell, & Vidgen, 2014). Yet, despite its relevance with regard to trace data, most quantitative trace studies do not consider the implications of this literature. Even for qualitative-oriented trace studies this body of work pose challenges. Often the sociomateriality debate focuses on fundamental ontological positions regarding the relationship between the social and material (Holeman, 2018). For example, we see a split emerging where some scholars argue that the social and material can be seen as ontologically separate but imbricated (Leonardi, 2012), while others promote a relational view where the social and material are inseparable (Jones, 2014). What seems to get lost in this discussion is how difficult it can be to describe everyday practice without making some basic distinctions about what goes into these practice, whether people, organizations, process or system, which in turn violates the inseparability principle by operating with pre-given categories. What is needed is a sociomaterial informed methodology, which pays attention to when and where in the research process distinctions and boundaries emerge.

We address these challenges by developing a set of guiding methodological principles and strategies for trace data studies. Drawing on the notion of apparatus and Haraway (1991, 1997) and Barad's (2003, 2007) distinctions among refraction, reflection and diffraction, we argue that studying trace data involves the building of an apparatus. Barad defines an apparatus as "the material conditions of possibility and impossibility of mattering; they enact what matters and what are excluded from mattering" (Barad 2007:148). As one constructs an apparatus, the phenomenon of interest emerges, which allows exploration of the boundaries and central distinctions of the phenomenon. These distinctions or cuts matter as traces diffract through the apparatus. For instance, when a participant contributes to an crowdsourcing site, such as Wikipedia or a citizen science project, their work is not simply reflected back to them on the screen. Their activities diffract through the system in different ways. Some entries may get structured as visible articles or discussion posts while other practices end up as less visible traces in the apparatus. These performances get to matter in different ways.

Beyond guiding specific empirical studies, our sociomaterially-informed methodology allows us to merge insights from different IS debates. First, a focus on the apparatus and the way it enacts boundaries and distinctions in a phenomenon reconciles different sides of the sociomateriality debate by acknowledging when in the research process distinctions are made. In other words, one may accept that the social and material remain ontologically inseparable but study how distinctions emerge as one build an apparatus and explore the multiple patterns that emerge as traces ripple through the apparatus. Second, such an investigation allows us to integrate quantitative and qualitative methodologies that previously have tended to flourished in different scholarly communities. Finally, our emphasis on the apparatus, its construction and performances brings the methodology into dialog with design studies (Hanseth & Lyytinen, 2010; Bjørn & Østerlund, 2014).

This essay is organized as follows. In the next section, we introduce our diffractive methodology for trace data and show how it fits into the existing sociomateriality debate and positivist- and interpretivist-oriented methodologies. We then develop our methodological guidelines by illustrating how refractive, reflective and diffractive methodologies would approach the study of learning among newcomers in a large online citizen-science project. The final section discusses the guidelines and notes some avenues for future research.

2. Theory

Going back to Marx and the Tavistock studies, scholars have gathered and analyzed traces of organizational practices in ways suggesting that technologies, people and discourses come together in dynamic and reciprocal assemblages (Gaskin, Berente, Lyytinen, & Yoo, 2014). The recent sociomaterial turn shines a bright light on these relationships (Orlikowski & Scott, 2008; Cecez-Kecmanovic, et al., 2014; Kautz & Jensen, 2013). Jones (2014) highlights five constitutive elements of sociomateriality, namely practice, performativity, materiality, relationality, and inseparability. We briefly explore how each of these elements speaks to challenges associated with using trace data.

Practice. One finds a practice-theoretical perspective at the core of the sociomaterial debate (Orlikowski & Scott, 2008). Practices of all stripes make up the fundamental building blocks of reality. Rather than seeing the world as made up of pre-defined substances external to one another, this approach grasps the world as brought into being through everyday activities. Practices produce and reproduce reality, make distinctions and draw boundaries (Feldman & Orlikowski, 2011). Trace data are no different. They are produced and reproduced through organizational practices and in the process delineate the activities of, e.g., employees, information systems, or artificial intelligence.

Performativity. Trace data are performative. Not merely records of performance, they also contribute to the constitution of the reality that they trace (Callon, 1998). Organizational members use traces to coordinate and render accountable many of their activities. In crowd systems such as Wikipedia, Facebook and many citizen science sites, traces left through prior performances compose the organization. The pictures and posts shared with family, friends, crowds and "algorithmic configurations" (Callon & Muniesa, 2005) on social media co-constitute those very networks. Barad goes a step further in arguing that the subjects and objects captured by the traces do not preexist as such but emerge through their performative intra-actions (Barad, 2007, p. 91).

Materiality. Trace data call our attention to matter, in several meanings of the word. Materiality plays a role in the production, recording, distribution and use of traces, whether these involve bodies, artifacts, information systems or code. Some sociomaterial scholars focus on tangible "stuff" when referring to materiality (Leonardi, 2012), which would leave some key elements of trace data out of the picture, such as the code and data structures entailed in their production. Others find it more productive to include intangible stuff, such as discourse, data and algorithms (Orlikowski & Scott, 2008). Barad (2007) takes the position that distinguishing materials takes work. Only through material-discursive performances can we demarcate an artifact.

Relationality. Many sociomaterial scholars adapt a relational ontology that breaks down the distinction between technology and humans (Orlikowski & Scott, 2008, p. 455). From this perspective, people and things have no inherent properties but gain these only in relation to one another. Their form, attributes and capabilities emerge through practice. Like points or lines in a geometric space, subjects and objects derive their significance from the relations that link them, rather than from some intrinsic features of individual elements (Swartz, 1997).

This form of dichotomy-busting relational thinking has a long tradition in practice theories (Østerlund & Carlile, 2005), particularly breaking down the distinctions between subjectivism versus objectivism (Bourdieu, 1977; Giddens, 1979), knowledge versus power (Foucault, 1980), agent versus artifacts (Haraway, 1991; Latour, 1986) and mental versus manual labor (Barley, 1986; Brown & Duguid, 2001).

Barad (2003) goes further than most scholars with her post-humanist agential realism, which questions not only the split between humans and materiality, and discourse and materiality, but also the subject and object in a research study as well as the phenomenon and the apparatus used to study it. In this view, distinctions emerge through practice, what Barad calls intra-action. Accordingly, trace data are neither purely social nor material, neither a pre-given part of a phenomena or the apparatus tracing it. Trace data gain their properties and attributes through ongoing sociomaterial performances that produce distinctions and effect. The subjects and objects traced are not pre-given but emerge out of the action.

*Inseparability***.** The notion of inseparability extends a relational ontology by arguing that the social and material are different aspects of the same phenomena. Their relationship is not one of unidirectional impacts as a positivist perspective might posit or mutual interaction as promoted by many interpretivist scholars, but of “intra-actions” (Barad, 2003). In other words, as subjects and objects are not predefined, their divisions and boundaries only emerge through intra-action (as opposed to inter-actions). They are mutually constituted. Materials are integral to all human activity and social action. It takes work to parse them apart. Traces thus emerge out of performances that are both social, discursive and material. Boundaries between humans and the material blur as one tries to discern what practices matter in the production of traces.

Scholars have noted (e.g., Shankar, Hakken & Østerlund, 2016; Jones, 2014) the proliferation of terms referring to the inseparability principle, including assemblages (Law, 2004), entanglements (Barad, 2003), practice-order bundles (Schatzki, 2002), mangles of practice (Pickering, 1995), and even the double mangling of the human and material (Jones, 1998). To foreground such sociomaterial patterns, Orlikowski and Scott (2008) suggest that scholars should seek them within the constitution of everyday work practices. But this proposition raises the question: on which practices should one focus? If we cannot assume that any categories are pre-given, where do we start?

To date, practice-theoretical scholars have managed this problem by applying the inseparability principle to one or two dichotomies at a time. Foucault explores the inseparability of power and knowledge while assuming other categories such as humans and artifacts. The same was the case when Bourdieu (1977) and Giddens (1979) dismantled the subjectivism versus objectivitism dichotomy. However, empirical description becomes increasingly challenging as one applies the inseparability principle more broadly, as for example Barad does, including not only the social and material but also human, discourse, power, knowledge, measurement, apparatus and phenomena, to mention a few. It can be difficult to describe everyday practices without making some basic distinctions about what goes into these practices, whether people, organizations, processes or systems and thus, violating the inseparability principle by operating with pre-defined categories. (For a detailed discussion of this issue see Holeman, 2018).

To bring the IS sociomateriality debate out of this conundrum requires an increased attention to how distinctions and boundaries emerge and when it is appropriate to introduce them in the research and publication process. It is not enough to focus on the distinctions produced through everyday practice in our empirical domain of interest. We need also to articulate how distinctions emerge out of our research practices. Our methods are like knives that cut. It is important to see our methods for what they do and not be fooled into thinking that entities are the way they are just because our methods happened to cut them up that way. Instead we should focus on our methodological knife skills, the way we carve up reality and make distinctions through our methodological practices. What is needed is a methodology highlighting how we as researcher make distinctions.

Barad's conceptualization of agential realism offers a starting point, with its focus on the interplay between phenomena, apparatus, boundaries and agential cuts. Inspired by the way Barad reads the work by quantum physics and STS scholars *through* one another, we will attempt the same--reading the IS methodology literature through Barad's diffractive approach to the research apparatus. In other words, we do not intend to provide a true replica of Barad's work but rather take key insights from her thinking to illuminate issues associated with trace data.

2.1 Apparatus: Refraction, Reflection and Diffraction

To explore methodological possibilities for drawing distinctions and boundaries out of our entangled world, we draw on the optical metaphors introduced by Haraway (1997) and extended by Barad (2007): refraction, reflection and diffraction. All three are optical phenomena, but the way each makes light available to an observer differs radically. Our discussion of these approaches is summarized in Table 1.

Refraction describes light passing through a lens. While Haraway (1997) and Barad (2007) mention refraction only in passing, grouping it with reflection, we note that a commonly-applied metaphor for trace data in social science is as a "lens" (e.g., Aiden & Michel, 2014) through which researcher can see what's happening in the world in great detail. We find this metaphor useful to describe a positivist-leaning view of data, or what Orlikowski and Scott (2008) refer to in IS as research stream I. Scholars with this bent strive to accurately reflect physical reality as discrete entities in their data. From this perspective, data are seen as akin to a microscope that shows us an original object, in fine detail. Objectivity is associated with methodological practices that produce homologous copies of the original entities, free of distortion. Observing substances through a lens, we assume that these substances are pre-given, with clear and predefined boundaries.

Reflection is a representation of an object produced by a mirror. With a mirror, we no longer look directly at objects but rather at a representation of them in the mirror. Furthermore, a mirror may capture only a partial image or an image with distortions that need to be accounted for. We find this metaphor useful when describing the methodological approaches of interpretivist and critical scholars (Orlikowski and Scott's Research Stream II), who argue that knowledge is best understood as reflections of mutually dependent ensembles. Interactions in these ensembles produce distortions that blur the reflections researchers can produce. Objectivity from this position is still about the objects, but recognizing that the "image" is partial or blurred and so in need of interpretation—indeed, "reflection" undertaken by the researcher—to discern their meanings.

In both refraction and reflection though, it is the image's likeness to the substance that matters, not the nature of the light producing the image. Empirical entities are seen as pre-given, what Haraway (1992) described as "'the same' displaced". Both cases hold the world at a distance (Barad, 2007). To put it differently, a refractive or reflective approach support what Cecez-Kecmanovic (2016) describes as a substantialist metaphysics concerned with "what there is." Only if one assumes the primary unit of reality are self-contained and bounded substances can one adopt a refractive or reflective approach in the methodologies chosen to describe the properties and qualities of such entities.

Diffraction concerns, in contrast, the bending and spreading of waves when they combine or meet an obstacle. Light and sound both exhibit diffraction under the right circumstances. A classic example of diffraction in physics is shown in Figure 1. In this experimental set up, light from a source on the left of the figure passes through two slits in the barrier in the middle of the figure and the beams of light from the two slits interfere with each other, leaving a diffraction pattern of light and dark on the screen beyond the slits to the right of the figure. This pattern does not appear if the light shines directly on the screen or if there is only one slit. Thus, the diffraction pattern records not only differences in the source waves but their history and interferences along the way to the screen. The metaphor offers a process perspective concerned with "what is occurring" and "ways of occurring" (Cecez-Kecmanovic, 2016). The primary unit of interest is not an image reflected on to a screen but the processes of configuring meaning and matter.

Figure 1. Diffraction pattern of light from a two-slit experimental setup.

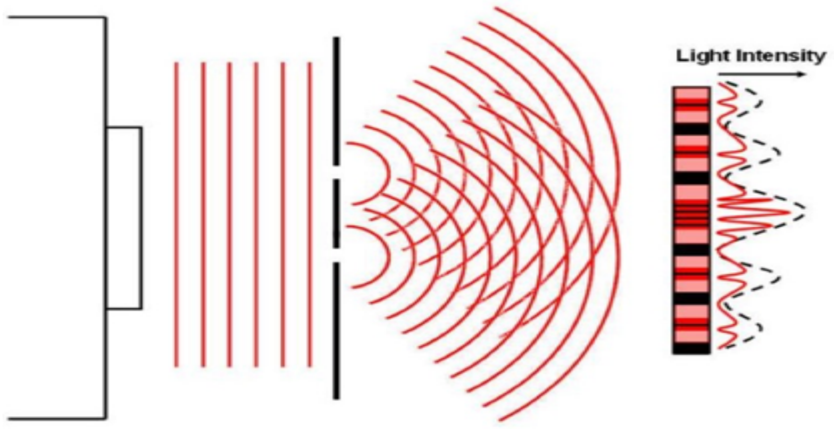

The apparatus takes on a central position in a diffractive methodology. Barad argues that one cannot disentangle a phenomenon and the apparatus that performs it. Instead, the apparatus plays a constitutive role in the production of the phenomenon by enacting specific boundaries in our sociomaterial reality. That is, online systems do more than record traces of human actions and interaction: they actively shape them.The apparatus is not a simple inscription device installed before the action happens. It is not a neutral probe, measuring pre-existing entities, mere reflections of a self-contained reality. Instead, the apparatus stands out as an open-ended practice constantly producing and reproducing the phenomenon that it records.

As a result, diffractive methodologies offer an analytical approach in which one reads elements of the research setup through one another by following the multiple patterns traces form as they ripple through the apparatus. It allows us to trace different practices and examine the distinctions they make. This *reading through* is possible because the elements are intertwined: changing the size, number or position of the slits or the nature of the light source in Figure 1 causes the diffraction pattern to take on a new shape. A diffractive apparatus allows researchers to learn about the nature of the light source and the nature of the apparatus the light passes encounters (e.g., the slits) through study of changes in the observed pattern. For example, physicists can study the nature of a chemical element by sending light from that element through a diffraction grating with known properties and observing the resulting diffraction pattern. Reading through can also work in the reverse direction: physicists can study the diffraction grating itself by illuminating it with light with known properties. For instance, one can learn about a crystal used as a diffraction grating by sending an x-ray of a known wavelength through it and studying the resulting diffraction pattern. Following the same line of thinking, information systems researchers can learn about trace data through studying the users of an online system, or learn about users through studying their information system, or learn about an information system through studying its traces.

Further, the practices of an apparatus are open to rearrangements. The creativity of scientific practices includes the skill of making the apparatus work for specific purposes. Elements are reworked and adjusted, leading to adjustments of the boundaries and cuts performed by the apparatus and so the nature of the phenomenon enacted and recorded. An apparatus can itself become the phenomenon, the focus of attention. This shift can happen as researchers turn their attention to the boundaries performed or by engaging the process in which the apparatus intra-acts with other apparatuses. These relations are only locally-stabilized phenomena that are part of specific performances.

In short, from a *refractive* position, the boundaries and distinctions demarcating subjects and objects are pre-given and sharp. A methodological apparatus seen as a lens in a positivist-oriented position treats trace data as true depictions of the world that allows researchers to generate images of pre-given objects and measure specific features. A *reflective* position that views the apparatus as a mirror, leads to an interpretive stance that deals with trace data as distorted or incomplete reflections of pre-given objects that need interpretation to determine their meaning. And a *diffractive* methodology approaches the apparatus as constitutionally entwined with the phenomena under study. The apparatus enacts cuts around and within the phenomena and thus is part of the making of boundaries and distinctions that we as researchers apply in our empirical descriptions. Differences emerge in a diffractive methodology but without absolute separation. Trace data diffract through the apparatus as ripples and waves and in the process they co-configure the apparatus and phenomena. Traces are thus not given but created. They open a window into both the phenomena and the apparatus by allowing researchers to read them through one another.

Table 1. Refractive, Reflective and Diffractive Approaches

	Refraction	**Reflection**	**Diffraction**
Research Stream*	**Positivist Research Stream I**	**Interpretivist Research Stream II**	**Sociomaterial Research Stream III**
Phenomena (Ontological priority)	Discrete entities	Mutually-dependent ensembles	Sociomaterial assemblages with no inherent properties that acquire form and features through interpenetration with an apparatus
Metaphor for the apparatus	Lens (Shows objects directly)	Mirror (Shows objects but indirectly)	Diffraction (Enacts cuts around and within phenomena).
Objectivity	About refractions, copies that are homologous to originals, authentic, free of distortion	About reflections, images that may be incomplete or blurred	About diffractive patterns that mark differences and relationalities that matter. Subjects and objects do not preexist but emerge through practice
Boundaries & distinctions	Pre-given & sharp	Pre-given but fuzzy	Emergent & fuzzy
Traces	True depiction of the world. Image of pre-given objects; Measure specific features of objects	Distorted and incomplete reflection of pre-given objects that need to be interpreted to determine meaning	Waves and ripples that diffract through the apparatus and in the process co-configure the apparatus and phenomena. Traces are not given but created. Allows one to read the phenomena and apparatus through one another.

3. Case Example: Learning in Citizen Science

To illustrate the three different approaches outlined above, we present examples from an ongoing study of learning in an online citizen-science project that was based in large part on trace data, thus providing examples of the issues discussed above. Citizen science is a broad term describing scientific projects relying on contributions from members of the general public (i.e., citizens in the broadest sense of the term) who volunteer time and effort to advance the goals of the project. There are several kinds of citizen-science projects: some have volunteers collect data, while others, including the one we examine here, ask volunteers to analyze already-collected data. Increasingly, the work of volunteers and project organizers take place via the Web, e.g., on a site that presents data to be analyzed and collects volunteers' annotations (e.g., www.zooniverse.org). Their work is sometimes described as "crowdsourcing science" and so is relevant to IS researchers. Moreover, citizen-science projects are an intriguing example of distributed learning and knowledge production, supported by public engagement in scientific research processes. To be effective over time, the projects must facilitate new users to orient themselves towards the goals and work practice of the project.

While our examples will be from the domain of citizen science, the need to study learning is more pervasive. A critical issue for sustaining groups that need to persist over time is how newcomers to the group learn to be effective participants (Van Maanen & Schein, 1977; Ostroff & Kozlowski, 1992; Klein & Weaver, 2000). In some groups, new members go through formal educational or orientation activities in order to learn group practices, while others rely on informal orientations. Online groups in particular often face difficulties with newcomer orientation, as many online groups are composed of members who are not part of a single formal organization and who contribute only in their free time, reducing or eliminating the possibility of formal training. However, the affordance of the technology used to support group interaction in some settings makes it possible for distributed volunteers to observe work in progress, thus enabling a form of legitimate peripheral participation (Antin &

Cheshire, 2010; Bryant, Forte, & Bruckman, 2012; Halfaker, Keyes, & Taraborelli, 2012). And most importantly, our arguments about the different approaches to trace data can apply to any trace data analysis.

Our examples are drawn specifically from the Gravity Spy (Zevin, et al., 2017) citizen science project (http://gravityspy.org/), which is built on the Zoniverse.org citizen science platform. The Gravity Spy project was developed to support the Laser Interferometer Gravitational-wave Observatory (LIGO). LIGO comprises two detectors that measure minute changes in distance caused by the gravitational waves bending space-time as they travel through it. However, the sensitivity that enables LIGO to detect distant astrophysical events also makes it very susceptible to non-astrophysical instrumental and environmental noise, referred to as “glitches”. Glitches hamper the detection of gravitational wave events, either by blocking events outright or by increasing the number of potential events to be examined. At LIGO’s current sensitivity, detectable astrophysical events are expected to occur only about once a month, while a glitch may occur every few seconds, making a search for true events akin to finding a needle in a haystack.

Similar glitches may have a common cause that can be eliminated if it can be identified, so finding and classifying glitches stand out as core tasks for improving the LIGO detectors. However, with thousands of glitches, the LIGO researchers do not have the manpower to examine them all. Relying on computers alone has also so far fallen short, as the diversity of glitches defies easy attempts at classification. At present, there are 22 known types of glitches, but many glitches do not fit one of these categories and so may be examples of as-yet-unidentified classes of glitch. Presently, humans are much better at the visual processing needed to identify similar types of glitches. Given these concerns, the project has developed a citizen-science approach to classifying glitches.

When using a citizen-science platform such as Zooniverse, volunteers are presented with images and asked to classify them into one of the known categories. Gravity Spy also provides options of none of the above or no image for images that in fact do not include an event of interest. The Gravity Spy system is shown in Figure 2: an image of a glitch to be classified is shown on the left as a spectrograph, with time on the x-axis, frequency on the y-axis and intensity represented as colour from blue to yellow. Possible classes are shown on the right. The initial learning challenge for new volunteers is how to identify the appropriate class for a glitch by matching its to one of the given exemplars.

Figure 2. Full Gravity Spy classification interface (http://gravityspy.org/).

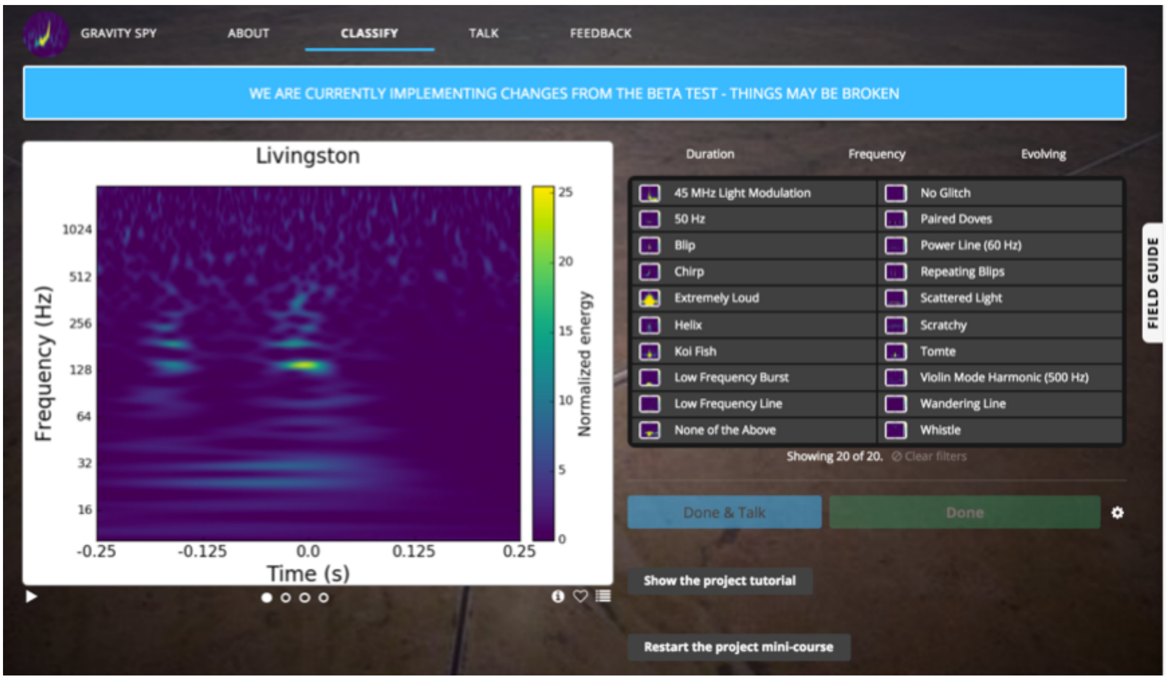

The Zooniverse system is instrumented to record several kinds of data. The classification dataset contains the classifications users contributed to the project. Included in the dataset are the glitch class chosen by the user (e.g., blip, whistle, etc.), the timestamp of the classification, and other metadata about the image, such as the image size and glitch type for images that were classified by experts ("gold standard" data). System interaction data contains events of users' interaction with pages on the site. When a user clicks on a link to access a new page on the website, an event record is stored. In total, 83 different kinds of website events are recorded. The record also contains a timestamp showing when the resource was requested. Data were collected, linked to a user ID, and include no personally-identifying or demographic data.

4. Approaches to Analyzing Trace Data for Learning

In an effort to build a set of guiding principles for a sociomaterial trace data methodology, we next present examples of how learning in Gravity Spy might be defined and studied from the three perspectives developed above.

4.1 Positivist/Research Stream I: Trace Data as Refraction

Investigations of learning in the tradition of Research Stream I consider data as depictions of the discrete and pre-given entities in the world, such as glitches and Gravity Spy volunteers. In this view, the trace data are seen as providing a lens on what volunteers are doing on the system and what and how they have learned. As noted, the Zooniverse system records data as volunteers contribute to and navigate through a project. Within the system (and the trace data) these actions are well identified, as the clickstream data are discrete units based on materials pre-defined by the system creators. Data are stored in rows and columns in a data store, embodying a set of identified boundaries. The system defines a user by use of a persistent user ID and linking records with the same user ID provides a record of the user's interactions with the system. To study volunteer learning, a researcher can look for evidence that volunteers' performance on the classification task improves over time (e.g., Crowston, Østerlund, & Lee, 2017), where performance is defined as the correctness of volunteers' classifications, i.e., the agreement of their choice of class with either an expert's choice or the consensus of other volunteers.

Research can further examine which system features lead to quicker or better learning (i.e., higher correctness). For example, some volunteers might have viewed the project tutorial, which describes the classification process, the science of gravitational wave research and how the data being analyzed by volunteers came into existence, or consulted other resources, such as the FAQs and the About page that provide additional context for the project and task, supporting volunteers' comprehension of the project and task. The system records which resources a volunteer has seen, creating for each viewing one or more records with the user ID, a timestamp and other metadata. Statistical analysis of these data can test the relationships between performance and the use of resources and other volunteer-specific factors, thus suggesting which resources are most helpful for learning.

4.2 Interpretivist/Research Stream II: Trace Data as Reflection

Researchers in the tradition of Research Stream II consider that data, even quantitative data, do not speak for themselves but rather must be interpreted. The system is a mirror as the recorded data are not the reality but rather reflect what happened, imperfectly, with omissions and distortions. Such interpretivist research lays a critical eye on trace data and their implications for understanding a phenomenon.

In this approach, the job of the researcher is to make sense of what they are seeing in the mirror of the dataset. Hermeneutics offers a well-articulated approach that has long served as a trusted pillar of qualitative and interpretive IS research. Boland (1985)—inspired by Edmund Husserl's phenomenological perspective and Gadamer's work on hermeneutics (Gadamer, 1975)—was among the first scholars to introduce hermeneutics to IS research. In classic hermeneutics, a text constitutes an object of study, which is to be understood based on its own frame of reference (Kvale & Brinkmann, 2009). Interpretation aims to bring to light an underlying coherence or sense from an otherwise incomplete, cloudy, or contradictory text (Myers, 1995).

The hermeneutic cycle summarizes the basic analytic process (See Figure 3), in which a researcher repeatedly moves back and forth between the whole corpus and its parts. This cyclical movement allows the researcher to start with a vague understanding of the whole text while making efforts to interpret the different parts. These partial interpretations are then related back to an emerging understanding of the whole, and the cycle repeats. The cyclical movement implies a continuously deepening understanding of the traces by contextualizing once current interpretation of the whole with further detailed analysis of selective parts. The analytic process implies a constant comparison between interpretations of small segments and the global meaning.

Figure 3. Hermeneutic Cycle.

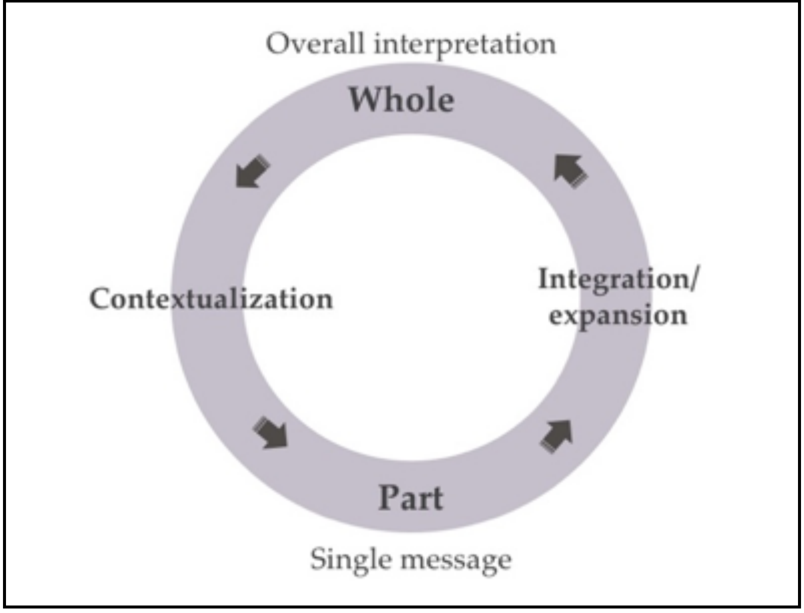

From this perspective, trace data becomes a text requiring an interpretation. The need for an interpretive approach is clearest when dealing with textual traces. For example, we might be interested in how volunteers draw on posts on discussion boards (known in Gravity Spy as "talk") to support their learning (Mugar et al., 2014). Just counting posts (as described in Track I) is unlikely to be satisfactory, as the text of posts are likely of different relevance for learning. Instead, the researchers would read and reread messages to form an interpretation of the kinds of messages and their function and then test that growing understanding against a larger set of messages and the overall context of volunteer learning. For example, research could examine how a volunteer calls attention to some feature of a glitch and how other volunteers respond, building a theory of communal learning (e.g., Mugar et al., 2014). Such an analysis might also lead to a redefinition of learning, e.g., moving from a focus on accuracy to consideration of how volunteers engage with scientific practice. In this case, the hermeneutic approach is applied much as in any qualitative study.

While the need for interpretation is clear for qualitative data, we note that an interpretivist approach can also be useful for forming an understanding of the meaning of quantitative trace data taken from an online system. At the most basic level, the researcher needs to understand the mapping of actions that volunteers can take on the system to the data that are recorded in the traces. While data may have labels (e.g., in a database dump), the connection between that label and an action is not always straightforward.

Further, to understand the import of data about user actions requires understanding the purpose and meaning of the captured interactions in the overall context of a volunteer's engagement with the system. Technologies are often used differently than intended by the designers, so it is important to recognize how volunteers enact the system in practice, and what the recorded system actions mean to volunteers. For example, in Gravity Spy, what the system records about interactions are the specific links that a volunteer clicks on the web page. To understand the meaning of clicks, we must form interpretations of this click in terms of user behaviors. For example, the system might record that volunteers clicked on the link for a discussion board. However, we do not know for sure that the volunteers actually read a particular post on the discussion board. It might be that the volunteer navigated to the board intending to create a new post, rather than read. Geiger and Ribes (2011) call the process of learning the meaning of digital traces "inversion". Complicating things further, different volunteers may mean different things by their use or use a feature with different levels of intensity. And yet, to assign meaning to the trace data, these nuances must be understood.

A key point of a hermeneutic approach is that to decode the meaning of a trace, it must be understood within the broader context of the work being done. However, trace data often lack situational clues, so it takes work to establish the context of the events. It may be useful to compare across time, settings or projects or to position traces in context with other work, perhaps other activities happening at the same time.

4.3 Sociomaterial/Stream III: Trace Data as Diffraction

Finally, developing an understanding of learning following Stream III, through a diffractive methodology, goes hand in hand with building an apparatus and exploring how practices ripple through the system. Investigating the apparatus cannot be separated from an exploration of the phenomena. In asking the question, "What is learning?", we notice the two sides to the question: '*what* is learning' and 'what is *learning'*. Both sides come into play as we build an apparatus.

4.3.1 Demarcate the phenomena and apparatus

From a diffractive perspective we turn our attention to the apparatus by exploring its boundaries and intra-actions with the phenomena. As noted above, refraction and reflection take the objects comprising the phenomenon as given, for example, volunteers, glitches and classifications. However, a diffractive reading helps us realize that what we take to be these objects emerge out of the performances going into the apparatus. To provide a few examples, first, as researchers, we tend to assume that volunteers exist and so look for them in our data (i.e., traces linked by a common user ID), but it is the distinctions and boundaries enacted by the apparatus that calls them into play. Second, glitches are created in the pre-processing of data obtained from LIGO. Whether a particular piece of signal is considered a glitch or not is depends on whether it passes an arbitrary signal-strength threshold; decreasing that threshold creates more glitches to be added to the system. The spectograph displayed in the system is also created as part of the pre-processing and the appearance of the image depends on a number of parameters that can be varied. Finally, correctness of a classification, a key variable in a study of learning, is determined by comparing a volunteer's classification against the "correct" answer for a glitch. For most glitches though, "correct" is taken as the consensus of volunteer classifications, meaning it is itself a product of the system. In the absence of consensus, correctness cannot be determined. A few glitches ("gold standard data") have classification given by LIGO experts, but classification is a practice and even these expert decisions are occasionally called into question. In summary, the sharp distinctions, assumed in the refractive and reflective analyses discussed above, on closer inspection turn out to be entwined with the apparatus.

Looking at boundaries more broadly, as a citizen-science project, Gravity Spy plays a role in a much larger apparatus, including detectors with 4 km-long arms in Washington and Louisiana and a third smaller detector in Italy. Hundreds of researchers across the world actively work on these instruments and in the process apply large IT infrastructures to store and analyze the data produced. Gravity Spy with its tens of thousands of citizen scientists constitutes just a small part of this larger effort. But Gravity Spy is hosted on Zooniverse, a citizen-science platform with more than 60 active projects and millions of volunteers. Where does the apparatus stop? Should our apparatus account for the machine-learning unit built into Gravity Spy that the volunteers are training? Or should we simply demarcate the apparatus as our locally-stored and curated database of Gravity Spy trace data?

The decisions have consequences for the phenomena, learning. Accounting for the detectors and their international research team suggest learning processes that goes beyond the volunteers' rather limited activities. The entire LIGO apparatus points us towards large-scale societal knowledge production and how research communities learn about the universe and its fundamental processes. This type of learning clearly motivates many volunteers, who eagerly search out additional readings about gravitational waves and the instruments, capable of detecting change in space-time of about 10^{-19} m, less than one-thousandth the diameter of a proton. The larger apparatus would lend itself to a conception of learning that fits into Science and Technology Studies or the 90s debates about organizational learning (Suchman, 2007; March, 1991).

Limiting our view to Gravity Spy work would allow us to define learning more narrowly around the volunteers' activities on the system. Yet, restricting our apparatus to Gravity Spy alone is easier said than done. Boundaries will remain fuzzy. The volunteers look at glitches produced by the detectors and interact with LIGO researchers in the discussion boards.

Gravity Spy is part of the Zooniverse platform and many of the volunteers participate in multiple projects spanning the fields of history, biology, medicine and astronomy. No matter our best intentions, bounding the phenomena and apparatus will always be a work in progress. Claiming otherwise, would require us to turn a blind eye to important performances.

4.3.2 Investigate the apparatus

Working with the apparatus involves an ongoing investigation of its performances starting with the question; *what does the apparatus trace?* And, *what does the apparatus not trace?* While it is tempting to expect that the system captures traces of all events, data storage is itself a practice, and the assumption of completeness must be carefully examined. Activities of interest may be unavailable for study. For example, the Zooniverse platform primarily supports science tasks. When we first started our study, it only recorded the annotations done and not activities such as volunteers' tutorial use, which the designers did not consider to be valuable data.

Other important activities might take place outside the apparatus. Trace data does not capture the work done by volunteers drawing on non Zooniverse servers. One volunteer created a web scrapper as a workaround, for instance, to quickly examine the images without having to go through the regular annotation procedure. The software crawled the Gravity Spy site by generating a URL based on the subject ID naming conventions Zooniverse uses for images on the server. The volunteer would then visually inspect the retrieved images to see if they fit the category he was interested in collecting. Other volunteers sometimes provide the URL of resources (e.g., academic papers, notebooks detailing alterations to the instrumentation at the detector sites) in a post, demonstrating that they are actively seeking additional knowledge. However, there is no systematic trace data record of when they do so or how those resources are used.

Finally, one should keep in mind that systems are subject to many problems that result in data loss (e.g., server outages, disk failures, deleted log files, or truncated database tables), meaning that trace data—even from database dumps—may be incomplete, though the problems may not be immediately visible (Howison et al., 2011). To address these problems the researcher should develop a detailed understanding of the apparatus. From a learning perspective it makes a big difference whether one has access to annotation work only or a range of other activities such as discussions among volunteers or the resources people might couple together to support their work on Gravity Spy.

Not only does the apparatus include and exclude certain practices in the traces produced, it also performs certain cuts. These distinctions play an essential role in demarcating key categories. For instance, we discussed above what encompasses the learner in Gravity Spy trace data. We assume in our analyses that a user ID represents an individual, but it is not impossible for a group of students to use one common login when working on Gravity Spy in their physics class or multiple members of a family to engage with the project after dinner. Contrariwise, participants may have multiple user IDs or work anonymously on the system without logging in, which means that they might have significant experience with the system that the trace data does not capture. Again, the apparatus doesn't draw sharp distinctions. It requires additional work if one hopes to define an individual in the trace data. Similar questions can be asked about other categories and practices central to learning, for example, what constitutes a science team member engaging in a project or how central is the machine-learning unit to the Gravity Spy project?

The boundaries and cuts performed by the apparatus change over time. A genealogy of the apparatus helps one understand how distinctions and boundaries gradually emerge in this sociomaterial system. The Zooniverse platform started out with the Galaxy Zoo project, which initially included only an annotation system. Volunteers would be presented an image to annotate and to avoid groupthink, they had to perform their own assessment before being able to access other participants' work on the same image. Soon after, a discussion board feature was added (originally, a stand-alone open-source discussion forum package was deployed). Gradually, user profiles, collections and search capabilities followed. Major funding from the Sloan Foundation and later Google allowed Zooniverse to create a more integrated project-builder platform allowing research teams to easily set up citizen science projects. Not only did all these changes lead to alterations to the apparatus, they also mark important cuts. For instance, the current Zooniverse project makes a rather sharp distinction between annotation work and discussions. They take place in different parts of the system and their relations are carefully managed.

4.3.3 Extending the Apparatus

Performing trace analysis further changes the apparatus. In other words, the apparatus and its traces are not pre-given. Additional cuts get added as researchers work with the trace data. These changes can take many forms, including among others the building of trace databases, conducting statistical analysis, experimental interventions (e.g., A/B splits), and interviews.

We turn to the question of databases first. To study a phenomenon as complex as learning requires us to pull data from multiple sources, such as records of use data and other metadata. These may be stored in different databases and database tables. In our study of learning, the available traces were not sufficient to address our questions. Zooniverse gathered traces about participants' annotation of science data but little else. After months of lobbying and joint funding we persuaded the software developers to add new trace features to the system so we would know when people had used various tools such as tutorials, science pages, collections, discussion boards, and user profiles. The expansion can be iterative: researchers cycle between appreciating the available traces and adding new traces to further flesh out and define the phenomenon.

The work doesn't end here. The newly constructed databases often leave us with a big unruly pile of traces, making it difficult to discern what differences matter. Constructing the apparatus involves further processing. For example, to understand how learning evolves over time, we divide volunteer traces into sessions (i.e., we perform additional cuts). The intuition is that volunteers will often interact with an online system for some period, creating a temporally-adjacent set of traces, then take a break (e.g., until the next day). Traces of events separated by a short gap can be grouped together in a single session, separated from the next session by a longer gap. This analysis approach provides a way to bound and separate traces to a format that acknowledges the temporality of Gravity Spy performance. We selected a set of traces to comprise a session. Prior work on Wikipedia has defined a gap of one hour between activities as indicating the start of a new session (Geiger & Halfaker, 2013), but given our own experiences annotating items in Gravity Spy and observing others do the same we choose a gap of 30 minutes our understanding of Gravity Spy annotation work, that is, the sequence of activities separated by less than 30 minutes were considered a session.

Applying statistical packages further extends the apparatus. Each analytic technique bundles and slices the trace data in new ways, and with it the phenomenon of learning. For instance, a session might be represented by counts of different kinds of actions (e.g., classification, reading or posting to discussion boards, consulting the field guide) that might contribute to learning. For example, applying computational approaches such as linear regression might allow us to model learning through use of these resources. However, analyzing counts loses information about the order of events. An alternative strategy applies sequence analysis techniques that focuses on the ordering of events (e.g., Keegan, Lev, & Arazy, 2015). Cluster analysis can also be used to identify sessions with similar patterns of activities. However, decoding these clusters require a diffractive reading of the quantitative analysis and calls for an exploration of how traces ripples through the apparatus.

4.3.4 Diffraction: Explore how traces ripple through the apparatus

An apparatus does more than produce metadata about practices associated with its use. As depicted in Figure 1 traces ripple through the apparatus. In Gravity Spy, annotations done by the volunteers feed into algorithms deciding how many other volunteers need to see the image before it is retired and it feeds the user profile to help participants know how much work they have done on the project. After a volunteer has annotated a glitch it is possible to leave a note with the particular image. As mentioned above, Zooniverse projects never allow volunteers to see other volunteers' annotations and provide access to talk traces about an object only after the user submits an annotation, both to avoid propagation of user biases. These restrictions to the way traces ripple through the system make it hard for newcomers to observe and learn from other more advanced volunteers' work practices. However, we find that many volunteers will compensate for this lack of legitimate peripheral participation (Lave & Wenger, 1991) by spending significant time looking over experienced participants notes in the Talk feature. These advanced notes serve as a form of practice proxies for less experienced participants (Mugar, Østerlund, Hassman, Crowston, & Jackson, 2014; Jackson, Østerlund, Mugar, Hassman, & Crowston, 2015). In other words, the traces do not refract or reflect users' behaviors. They ripple through the apparatus and feed other practices. Some of them ricochet back to the participants in the form of e.g., user profile stats or Talk posts.

To make sense of the activity clusters generated statistically, we follow how participants behaviors rippled through Zooniverse. For example, we applied the cluster analysis to sessions mentioned above. One prominent cluster captured performances restricted to the annotation feature. Participants did one annotation after the other, over a shorter time span, with no traces left in other features. We named this type of session 'light work.' A less prominent but still significant cluster involved traces of activities indicating that a volunteer after each annotation would check if other people had left notes on that image. Often, they spent a long time going through these communal discussions, but rarely left any notes themselves. We named this cluster 'careful annotation.' Another cluster we called 'talking and annotating,' which included a lot of discussion board traces with a few detours rippling into the annotation system. From the sequencing of the traces we discern that in some sessions participants spend most of their time engaging in the discussion board, collections features but with periodic visits to the annotation task (Jackson, Østerlund, Maidel, Crowston, & Mugar, 2016). For each user ID we organized these session types sequentially and found among other things that most participants would stick to light work sessions. More dedicated participants would oscillate between light work and more involved sessions where they would either engage with the community through posts and discussions or spend a lot of time diving into each image and other people's annotations of those glitches. A small number of participants would have sessions focused on individual images, building collections of unusual images and reading science notes. In short, to explore the traces, one follows their paths through the apparatus. Again, it is important not to stay with the same unit. Instead, we move between following a single trace, clusters of traces, temporal ordering of traces and sequences of sessions and grouping of participants with similar session sequences. By dialing up and down (Gaskin et al., 2014) on the size and order of trace bundles we explore multiple performances, patterns, and learning phenomena and how they change over time.

More explicit design changes to the apparatus further allow one to explore what differences matter by sending ripples through the system in different ways. In the diffraction experiment depicted in Figure 1, we as researchers can change the light source or the slits the light passes through to see how it changes the way traces ripple and the diffractive patterns they form. Similarly, as part of our study of learning in Gravity Spy, we implemented a scaffolded progression of tasks to support newcomers learning. Volunteers annotate glitch images into the 22 known classes of glitches. But rather than providing all classification options to new users, the system introduces them a few at a time. New volunteers start in Level 1, a simplified version of the classification interface, in which they are presented with glitches to classify that are expected to be of one of only two distinctive classes – "blips" vs. "whistles" or "none of the above." Once the volunteer can successfully classify glitches of the initial two classes (currently assessed by accuracy in classifying gold-standard data), the volunteer can advance to the next training level, in which they see glitches of additional classes. In other words, to scaffold volunteer learning, the system gradually expands the number of classes presented to the volunteers. The glitches to be presented in each level are selected by a machine-learning (ML) algorithm. The ML classifies all glitches added to the system into one of the known classes, with an accompanying confidence in the classification. Glitches with a high ML confidence is given to new participants as training. Once volunteers have learned more glitch classes, they also get presented with images with lower and lower ML confidence.

To see if these differences matter compared to typical Zooniverse projects where people access all known classes from the beginning we performed a simple A/B split. New participants were divided into two groups over a period of a few weeks. One group went through the scaffolded system while the second group faced all 22 known glitch classes from the beginning. Subsequent trace analysis suggested that the members of the scaffolded group contributed to the project significant longer, mastered the task faster and did more annotation work than the second group. During the experiment, some volunteers in the second group went back through the scaffolded levels they initially were barred from without any prompts from the system.

Recently, we have experimented with giving advanced participants access to the ML processing to support their search for new glitch classes unknown to the science team. Instead of assigning images to volunteers the advanced participants use the ML to find images similar to clusters of images they have deemed possible candidates for a new glitch class. In this way, we hope to learn more about machine-human learning intra-actions and agential cuts that matter to such performances. These dynamics cannot be explored without carefully following the ways traces ripple through the apparatus.

Direct engagement with volunteers offers ways to explore the apparatus and its diffractive patterns. Participant observations and interviews with individuals and focus groups help explore trace and the way they ripple. For instance, visualizations of trace data such as the sequences of sessions described above can serve as productive interview prompts. They give the volunteers a view into the apparatus and the way their practices ripple through the system and offers them an opportunity to describe how these traces relate to other activities not captured by the apparatus. Such interview protocols can span a broad range of traces. We have used collections of Talk posts to explore how newcomers use experienced participants' annotations as practice proxies to emulate. In other interviews we have shared highly processed trace visualizations of session sequences associated with the interviewee. The method goes beyond traditional triangulation, which tend to assume pre-given entities and test one statement against other statements about this object. Instead trace interview prompts offer ways to learn more about performances and how they do and don't ripple through the apparatus.

4.3.4 Differences that matter and ethics

The diffractive analytic process involving the demarcation of the apparatus and phenomena, exploration of the apparatus and the way traces ripple through it add up to a search for differences that matter. This rippling is not referring to a more traditional conception of causality as relations between distinct entities (Barad, 2007). Instead it explores the effect of specific distinctions and bounding, i.e., agential-cuts build into the apparatus. As Barad argued: "causal relations entail a specification of the material apparatus that enacts an agential cut between determinately bounded and propertied entities within a phenomenon (Barad, 2007: 176). In other words, we have to pay attention to the boundaries enacted by the apparatus in its entwined relationship with the phenomena and the distinctions it makes. Only then can we explore how traces ripple through the apparatus and what changes they leave in their wake.

We have found benefits in a circular analytical process where the researcher oscillates between exploring the boundaries of the apparatus/phenomenon and the way traces ripple through the apparatus/phenomenon. Where a hermeneutic process cycles between analyzing a whole pre-given text and its parts, we envision a circular movement through the diffraction apparatus (see Figure 4). Studying Gravity Spy, one cannot assume that traces scrapped from the system constitute a whole. Instead the researchers, and in many situations the volunteers, explore the boundaries of the apparatus and may add new features to the configuration. Tracking traces as they ripple through the system allow one to question the distinctions made. For instance, what constitute learning and a volunteer? What type of performances do they engage in and how do they change over time?

Volunteers are not leaving traces behind them as a boat cutting a wake in its path. The traces make up part of the reality that define the performances. What one sees in Gravity Spy is partly a product of one's own and other volunteer's traces. The boat is rocked by its own wake as it plows through a canal, with each wave diffracting back to the boat after hitting the channel banks. The diffractive patterns of the waves must be read through the rocking of the boat, the structure of the embankments and the decisions of the pilot trying to avoid spilling his morning coffee. Moreover, the researcher may change to banks of the channel or the shape of the boat to see how the wave patterns change. We may even question whether we are dealing with a captain at the helm or a middle school class supported by ML. The diffraction pattern marks differences that matter.

Figure 4. Cyclical Analysis

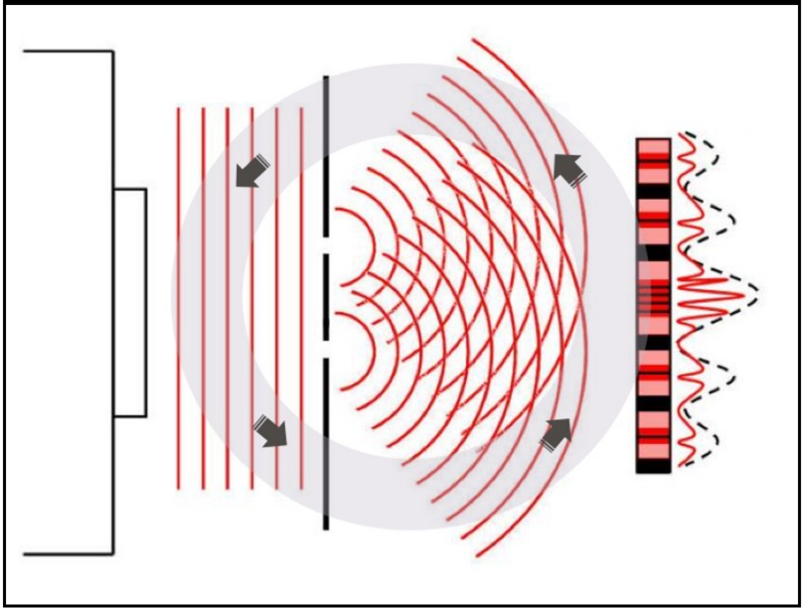

The dynamic intra-actions between phenomena and apparatus, i.e., boundaries and distinctions emerging as traces ripple through the system, allows us to operate with multiple learning phenomena at the same time. Each form of learning would be associated with a different apparatus and co-configuration of how traces ripple through it. Inspired by Sørensen (2009) we distinguish three learning phenomena associated with Gravity Spy: Authority-subject, communal, and agent-centered learning.

First, authority-subject learning emerges in an apparatus divided into clear regions and sub-regions, each associated with clusters of homogeneous and highly structured activities, events and objects. One can imagine a classroom as a region divided into two sub-regions, supporting authority-subject learning. The front of the room, which is occupied by the teacher and the blackboard, and the rest of the classroom inhabited by students, their desks, chairs, all organized to face the blackboard and the teacher's sub-region. The separation between students in their chairs and the teacher at the blackboard thus marks two distinct regions, each associated with particular activity. The tutorial pages and training modules work much like the teacher's sub-region pushing authoritative knowledge from the expert science team to the volunteers' sub-region in the annotation system. As in a classroom, the annotation sub-region of Gravity Spy constitutes a highly structured environment where volunteers are asked to review one image after another. From time to time they review a gold standard image and receive feedback. Did they annotate it as the authority would do it or not? To understand this form of learning one could focus on the two sub-regions of the apparatus of interest and track how traces ripple through them and what differences matter. The scaffolding experiment described above could allow one to further explore these dynamics. Whether a user ID stays active longer and perform annotations with high accuracy after frequenting the tutorial and field guide is what matter.

Second, communal learning forms around a central collective activity, object or event. All other elements receive their identity through their resonance with that center. For instance, at a festival or during a communal celebration, the collective develops a joint experience around this shared event. Communal learning takes form as the collective take shape and extends its performances. Relevant traces could be the folksonomies of shared hashtags that develop in the discussion board and collections feature, as participants develop new glitch classes out of images relegated to 'none of the above.' What matters are the formation of these collective hashtags, the degree to which they are used over time and solidify into new glitch classes used by a range of participants.

Third, one can envision an agent-centered learning with no central focal point. Rather the agents' evolving practices build on one another to form a bricolage, piecing together element of their participation as they move through Gravity Spy and beyond. Boundaries are fluid, the apparatus and phenomena not defined but keep morphing and changing as participants develop practices and discourses associated with e.g., gravitational waves and the LIGO detectors. The sequential ordering of traces matter and the type of resources, discussions, people and events they link. Piecing together the session types people combine can help investigate agent-centered learning.

These three forms of learning are not mutually exclusive. As researchers, citizen scientists can approach the apparatus in multiple ways, demarcate their phenomenon of interest and perform certain cuts through their intra-actions with the apparatus. This does not mean than anything goes. We cannot dream up endless forms of learning on the fly and project them onto an apparatus. One needs to perform differences that matter; ripples moving through the apparatus and create some effect on a phenomenon. Operating with multiple forms of learning does not constitute a contribution. The field has for a long time acknowledged e.g., cognitive and situated learning theories side by side (Lave & Wenger, 1991; Miner, Bassoff, & Moorman, 2001; Gherardi, 2006; Levinthal & Rerup, 2006). Rather, a diffractive reading embrace multiple entwined forms of learning, all operating in a dynamic field of possibilities and impossibility of mattering.

5. Conclusion

Facing a torrent of trace data, IS researchers confront a number of methodological challenges associated with the building of an apparatus and understanding how it co-constitutes the phenomenon under investigation. Trace data are not given but produced. Thus, they do not refract or reflect some pre-given reality that researchers through hard labor can project onto the pages of their articles. The boundaries defining the phenomena of interest are not prepackaged subjects and objects. Instead the researcher needs to pay careful attention to how the building of the apparatus demarcates different entities and the way they co-constitute one another. Carefully assembling an apparatus and following the traces rippling through it offers new ways to explore organizational practices. We contribute with a number of methodological principles and strategies for such a diffractive approach to trace data as summarized in Table 2. These are not bureaucratic procedures to be following one after the other but rather fundamental questions guiding the research process. As depicted in Figure 4, we find it helpful to think of the research process as a circular motion where we track the way traces ripple through the apparatus. Continuing this iterative process enables scholars to follow lines of becoming and to describe how boundaries taken form and fall apart. By observing and experimenting with the rippling traces, the dynamics of our research practices exposes the becoming of technologies, people and entities and how their boundaries and properties are reshaped, with what consequences and for whom (Cecez-Kecmanovic, et al., 2014, pp. 821). Equally important, the methodology offers a fresh view on divisions in the IS literature. Below, we will briefly discuss some of these implications for future research.

Leading voices in the sociomateriality debate have for some time called for empirical studies on how relations and boundaries between humans and technologies are not pre-given or fixed but enacted in practice (Orlikowski & Scott 2008; Jones, 2014; Suchman, 2007; Lave & Wenger, 1991). Before these dynamics can be examined, we need to understand how boundaries and distinctions emerge as part of our research process. Even if we fully accept the relational and inseparable nature of our sociomaterial world, we cannot question all distinctions in every study. It is paramount, however that we recognize the distinctions we make and where they appear in the research process. We need to acknowledge what Barad (2007) calls agential cuts; differences that matter. Recognizing these distinctions will not catapult us back to a substantialist position. Rather, it will strengthen a process perspective on how distinctions and boundaries emerge in the entanglement of humans and materials (Cecez-Kecmanovic et al., 2014).

Table 2. Methodological principles, strategies and evidence		
Principle	**Strategies and Question**	**Evidence from learning in Gravity Spy**
Demarcating the phenomena and apparatus	What are the boundaries of the apparatus? And thus, what is the phenomena?	Demarcating the apparatus call into question: *What* is learning? What is *learning*? E.g., including the larger LIGO collaboration leads to a study of societal knowledge production. Restricting the apparatus to Gravity Spy traces may point to performances associated the volunteers, machine learning unit or community of participants. Boundaries remain fuzzy and we cannot draw a sharp line between entities e.g., Gravity Spy and LIGO. We know that volunteers work anonymously on the site and use non-Zooniverse systems. We consider if those performances play a role in learning.
	What cuts do the apparatus make?	What entities can we distinguish in the learning environment? E.g., can we associate certain performances to volunteers, machine learning units, and science team members or does the apparatus not allow us to distinguish e.g., humans and machine learning? We explore how a single user ID can might represent an individual, a school class or family of four. The same questions should be asked about other central performances attributed to science team and machine learning unit.
	Genealogy of an apparatus: How have the boundaries and cuts changed over time?	Explore how the learning environment changes over time? This helps us detect important distinctions performed by the apparatus. For instance, there is a clear distinction between the annotation system and discussion forums in Gravity Spy.

Extending the apparatus	What additional traces might be helpful?	To analyse Gravity Spy trace data, we built a database merging several datasets. We also persuaded Zooniverse to add tracking capabilities to the platform to record users' interactions.
	What additional cuts might be helpful? E.g., statistical tools can be added to the apparatus performing additional cuts	To understand how performances evolve over time, we parse traces into sessions divided by gaps inactivity. We try out different statistical apparatuses to see if they help distinguish cuts that matter e.g.,k-means clustering. Does one simply regard the number of times a user ID has visited certain features as contributing to learning or does the sequence of performances matter?
Diffraction: Explore how traces diffract (i.e., not refract or reflect).	How do traces ripple through the apparatus?	What performances by other agents are participants allows to access and when? What consequences does it have for learning? In Gravity Spy participants cannot access other people's annotation work. Instead, participants go to Talk looking for practice proxies, in the form of descriptions of work.
	What intra-actions do the ripples highlight?	By adding cluster analysis to the apparatus, we explored how traces rippled through the apparatus in different ways and formed multiple patterns. Some ripples stayed within the annotation system (e.g., light work), others spanned multiple performances (e.g., talking and annotating).
	What happens if you change the way traces ripple through the system?	An A/B split in Gravity Spy experimented with two pathways through the apparatus. One group were guided through a ML supported scaffolding of the work and a second group went straight to classify all known classes. Change access to ML in the apparatus and follow how traces ripple differently through the system and if it leads to different patterns and performances. Visualizations of traces serve as interview prompts and help explore how performances ripple within a beyond the boundaries of the apparatus. The A/B split and interviews allowed us to look for differences that matter for performances associated with the apparatus.
Differences that matter	How does a circular movement between exploring the boundaries of the apparatus/phenomenon and the way traces ripple through it help articulate help find differences that matter?	To explore *what* is learning and what is *learning* we move in circular patterns between different apparatuses/phenomena and agential cuts shaping the way traces ripple through these configurations.

	What differences matter?	Allow us to operate with multiple forms of learning playing out in co-configured apparatuses: Authority-subject, communal, agential, and machine learning are all performances associated with Gravity Spy. For each of these learning phenomena different traces and cuts matter.

The methodological principles and strategies outlined in Table 2 help guide the research process, but also articulate the genealogy of boundaries and distinctions. It reminds us that we as researchers are an integral part of the apparatus; not in the sense that we distort some reflection of user behaviours, but rather, that our active engagement in the building and running of the apparatus offers rich opportunities to explore how boundaries and cuts emerge and what can and cannot be know about the ongoing dynamics of becoming associated with the system. Data collection practices are open to rearrangements and the creativity of scientific practices includes the skill of making an apparatus work for a purpose. Elements are reworked and adjusted, leading to adjustments of the boundaries and cuts performed by the apparatus and the nature of the phenomenon. In ethnographic monographs it has long been the norm to include a describing the researchers entrance to the field. Future IS publications might similarly require an appendix describing the building and running of the apparatus in a way that acknowledge the distinctions and boundaries drawn and where they emerged in the research process. Ethical considerations would be an appropriate part of these considerations. Instead of framing ethical research as impacting or interacting with human subjects in a way that ensure their rights and welfare, a diffractive approach could articulate how the research have made responsible and accountable distinctions and connections to what comes to matter and what is excluded from mattering. Future research would have to further articulate such approaches to ethics and its consequences for institutional review boards and research practices. Likewise, we have only scratched the surface when it comes to a diffractive methodology. As Barad's work suggest it allows us to revisit well worn categories and see them in a new light including among others, causality, discourse, measurement, time and space (Barad, 2007).

A diffractive methodology suggests ways to integrate quantitative and qualitative approaches. Cluster analysis and interviews still has a role to play. As highlighted in our study of Gravity Spy both help us explore how traces ripple through the apparatus. Visualizations of trace data can serve as powerful interview prompts, which in turn may inform changes to the apparatus that allows the tracking of other practices and alterations to the cuts and boundaries performed. The circular movement depicted in Figure 4 suggest that researchers read insights gained from these different techniques through one another in a cyclical motion as one follows the traces ripple through the apparatus. It will take additional research to map a broader range of productive combinations of participant observation, interviews and various statistical techniques.

Our guidelines have practical implications. Building a research apparatus and attention to its performances brings a diffractive methodology into close proximity with design theory (Hanseth & Lyytinen, 2010) and neighboring disciplines with a design agenda such as CSCW (Bjørn & Østerlund, 2014). One can envision a joint interest in how the apparatus and phenomenon intra-action, and the ways in which distinctions take shape and categories are bound. The diffractive ways traces ripple through the Gravity Spy project was as relevant to the designers at Zooniverse as it was to our research and the volunteers. All were hoping to learn about and improve organizational performances.

We began by highlighting the sociomaterial nature of trace data and the importance of recognizing that traces are not given but created. This perspective brings us back to the tension in the sociomateriality debate between the notion of inseparability built into a relational ontology and the need to make distinctions and draw boundaries as part of a research study. In Barad's words (2007), we need to acknowledge the agential cuts performed by the apparatus. Only then can we leverage trace data to explore the sociomaterial nature of organizational and everyday practices. We believe that explicit recognition of a diffractive methodology is a promising approach that allows researchers draw on trace data in a way that does not presume pre-given entities but also allows exploration of the agential cuts feeding organizational and everyday practices. The research can explore by cycling between considering the boundaries of the apparatus/phenomenon and tracking traces as they ripple through the apparatus. Following such a research process may afford some further empirical, theoretical and methodological innovations to realize the promise of trace data contributing to our understanding of contemporary organizational practices.

6. References

Abbasi, A., Sarker, S., & Chiang, R. H. (2016). Big data research in information systems: Toward an inclusive research agenda. *Journal of the Association for Information Systems*, *17(2)*.

Aiden, E., & Michel, J.-B. (2014). *Uncharted: Big Data as a Lens on Human Culture*. New York, NY: Penguin.

Antin, J., & Cheshire, C. (2010). Readers are not free-riders: reading as a form of participation on wikipedia. In *Proceedings of the ACM Conference on Computer-supported Cooperative Work* (pp. 127-130). ACM.

Barad, K. (2003). Posthumanist performativity: Toward an understanding of how matter comes to matter. *Signs*, *28(3)*, 801-831.

Barad, K. (2007). *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, NC: Duke University Press.

Barley, S. R. (1986). Technology as an occasion for structuring: Evidence from the observation of CT scanners and the social order of radiology departments. *Administrative Science Quarterly*, *31*, 78–109.

Bjørn, P., & Østerlund, C. (2014). *Sociomaterial-Design: Bounding Technologies in Practice*. Switzerland: Springer.

Boland, R. J. (1985). Phenomenology: a preferred approach to research on information systems. In Mumford. E., Hirscheim. R., Fihgerald, G. 8 Wood-Harper, T. (Eds.), *Research Methods in Information Systems* (pp. 193-201). North-Holland. New York.

Bourdieu, P. (1977). *Outline of a Theory of Practice*. Cambridge University Press.

Brown, J. S., & Duguid, P. (2001). Knowledge and organization: A social-practice perspective. *Organization Science*, *12(2)*, 198-213.

Bryant, S. L., Forte, A., & Bruckman, A. (2005). Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work* (pp. 1-10). ACM.

Callon, M. (1998). *The Laws of the Markets. Sociological Review Monographs Series.* Oxford, Blackwell Publishers.

Callon, M., & Muniesa, F. (2005). Peripheral vision: Economic markets as calculative collective devices. *Organization Studies, 26*(8), 1229-1250.

Cecez-Kecmanovic, D., Galliers, R. D., Henfridsson, O., Newell, S., & Vidgen, R. (2014). The sociomateriality of information systems: Current status, future directions. *MIS Quarterly, 38(3),* 809-830.

Cecez-Kecmanovic, D. (2016). From substantialist to process metaphysics: Exploring shifts in IS research. In *Beyond Interpretivism? New Encounters with Technology and Organization: Proceedings of the IFIP Working Group 8.2 Conference on Information Systems and Organizations*, Dublin, Ireland.

Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly, (36)4*, 1165-1188.

Crowston, K., Østerlund, C., & Lee, T. K. (2017). Blending machine and human learning processes. In *Proceedings of the 50th Hawaii International Conference on System Sciences* (pp. 1624-1634).

Feldman, S. & Orlikowski, W. (2011). Theorizing Practice and Practicing Theory. *Organization Science, 22(5)*, 1240–1253.

Foucault, M. (1980). *Power/knowledge: Selected Interviews and Other Writings, 1972-1977*. New York, NY: Pantheon.

Gadamer, H. G. (1975). *Truth and Method*, trans. W. Glen-Dopel. London: Sheed and Ward.

Gaskin, J., Berente, N., Lyytinen, K., & and Yoo, Y. (2014). Toward generalizable sociomaterial inquiry: A computational approach for zooming in and out of sociomaterial routines, *MIS Quarterly*, *38(3)*, 849-871.

Geiger, R. S., & Ribes, D. (2011). Trace ethnography: Following coordination through documentary practices. In *Proceedings of the 44th Hawaii International Conference on System Sciences* (pp. 1-10).

Geiger, R. S., & Halfaker, A. (2013). Using edit sessions to measure participation in Wikipedia. In *Proceedings of the Conference on Computer-supported Cooperative Work* (pp. 861-870). ACM.

Gherardi, S. (2006). *Organizational Knowledge: The Texture of Workplace Learning*. Blackwell Publishing, Oxford, UK.

Giddens, A. (1979). *Central Problems in Social Theory: Action, Structure and Contradiction in Social Analysis*. Berkeley, CA: California.

Goggins, S. P., Galyen, K., & Laffey, J. (2010). Network Analysis of Trace Data for the Support of Group Work: Activity Patterns in a Completely Online Course. In *Proceedings of the 16th ACM International Conference on Supporting Group Work* (pp. 107–116). New York, NY, USA

Halfaker, A., Keyes, O., & Taraborelli, D. (2013). Making peripheral participation legitimate: Reader engagement experiments in Wikipedia. In *Proceedings of the Conference on Computer-supported Cooperative Work* (pp. 849-860). ACM.

Hanseth, O., & Lyytinen, K. (2010). Design theory for dynamic complexity in information infrastructures: The case of building internet. *Journal of information technology, 25(1)*, 1-19.

Haraway, D. (1991) *Simions, Cyborgs and Women,* New York: Free Association Books.

Haraway, D. (1997). *Modest_Witness@Second_Millenium.FemaleMan_Meets_OncoMouse: Feminism and Technoscience*, London: Routledge.

Holeman, I. (2018). Room for silence: Ebola research, pluralism and the pragmatic study of sociomaterial practices. *Computer Supported Cooperative Work (CSCW)*.

Howison, J., Wiggins, A., & Crowston, K. (2011). Validity issues in the use of social network analysis with digital trace data. *Journal of the Association for Information Systems*, *12(12)*.

Jackson, C. B., Østerlund, C., Mugar, G., Hassman, K. D., & Crowston, K. (2015). Motivations for sustained participation in crowdsourcing: Case studies of citizen science on the role of talk. In *Proceedings of the 48th Hawaii International Conference on System Sciences* (pp.1624-1634).

Jackson, C., Østerlund, C., Maidel, V., Crowston, K., & Mugar, G. (2016). Which Way Did They Go? Newcomer Movement Through the Zooniverse. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (pp. 624-635). ACM.

Jones, M. (1998). Information systems and the double mangle: Steering a course between the Scylla of embedded structure and the Charybdis of strong symmetry. In *Proceedings of the IFIP WG8.2 and 8.6 Joint Working Conference*, Helsinki, Finland.

Jones, M. (2014). A matter of life and death: Exploring conceptualizations of sociomateriality in the context of critical care. *MIS Quarterly, 38*(3), 895-925.

Kautz, K., Jensen, T.B. (2013). Sociomateriality at the royal court of IS: A jester's monologue. *Information and Organization, 23(1)*, 15–27

Keegan, B. C., Lev, S., & Arazy, O. (2015). Analyzing organizational routines in online knowledge collaborations: A case for sequence analysis in CSCW. In *Proceedings of the Conference on Computer-supported Cooperative Work and Social Computing (*pp. 257–266).

Klein, H. J. , & Weaver, N. A. (2000). The effectiveness of an organizational-level orientation training program in the socialization of new hires. *Personnel Psychology, 53(Spring)*, 47-66.

Kvale, S., & Brinkmann, S. (2009). *Interviews: Learning the Craft of Qualitative Research Interviewing*. London: Sage.

Latour, B. (1986). Visualization and cognition. *Knowledge and Society, 6*, 1-40.

Lave, J. and Wegner, E. (1991). *Situated learning*. Cambridge University Press.

Law, J. (2004). *After Method: Mess in Social Science Research*. New York, NY: Routledge.

Leonardi, P. M. (2012). Materiality, sociomateriality, and socio-technical systems: What do these terms mean? How are they different? Do we need them? In P. M. Leonardi, B. A. Nardi, & J. Kallinikos (Eds.), *Materiality and Organizing: Social Interaction in a Technological World* (pp. 25). Oxford: Oxford.

Levinthal, D. A. & Rerup, C. (2006). Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning. *Organization Science, 17(4)*, 502–513.

Loukissas, Y. A. (2017). Taking Big Data apart: local readings of composite media collections. *Information, Communication & Society, 20(5)*, 651-664.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, *2(1)*, 71-87.

Miner, A., Bassoff, P., and Moorman, C. (2001). Organizational improvisation and learning: A field study. *Administrative Science Quarterly, 46*, 304–337.

Mugar, G., Østerlund, C., Hassman, K. D., Crowston, K., & Jackson, C. B. (2014). Planet hunters and seafloor explorers: Legitimate peripheral participation through practice proxies in online citizen science. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 109-119). ACM.

Myers, M. D. (1995). Dialectical hermeneutics: A theoretical framework for the implementation of information systems. *Information Systems Journal*, *5(1)*, 51-70.

Orlikowski, W. & Scott, S. (2008). Sociomateriality: Challenging the separation of technology, work, and organization. *Academy of Management Annals, 2(1)*, 433–474.

Østerlund, C., & Carlile, P. (2005). Relations in practice: Sorting through practice theories on knowledge sharing in complex organizations. *The Information Society, 21*(2), 91–107.

Ostroff, C. , & Kozlowski, S. W. J. (1992). Organizational socialization as a learning process: The role of information acquisition. *Personnel Psychology*, *45*, 849-874.

Phang, C. W., Kankanhalli, A., & Sabherwal, R. (2009). Usability and Sociability in Online Communities: A Comparative Study of Knowledge Seeking and Contribution. *Journal of the Association for Information Systems, (10)10*, 721-747.

Pickering, A. (1995). *The Mangle of Practice: Time, Agency and Science.* Chicago, IL: University of Chicago.

Schatzki, T. R. (2002). *The Site of the Social: A Philosophical Account of the Constitution of Social Life and Change*. University Park, PA: Penn State University Press.

Shankar, K., Hakken, D., & Østerlund, C. (2017). Rethinking documents. In U. Felt, R. Fouché, C. A. Miller, & L. Smith-Doerr (Eds.), *The Handbook of Science and Technology Studies* (pp. 59–86). Cambridge, MA: MIT Press.

Sørensen, E. (2009). *The Materiality of Learning: Technology and Knowledge in Educational Practice*. Cambridge University Press.

Suchman, L. A. (2007). *Human-Machine Reconfigurations: Plans and Situated Actions*, 2nd ed. Cambridge University Press, Cambridge, UK.

Swartz, D. (1997). *Culture and Power: The Sociology of Pierre Bourdieu*. Chicago, IL: University of Chicago Press.

Van Maanen, J. , & Schein, E. H. (1979). Toward a theory of organizational socialization. In B. M. Staw (Ed.), Research in organizational behavior (pp. 209-264). Greenwich, CT: JAI.

Whelan, E., Teigland, R., Vaast, E., & Butler, B. (2016). Expanding the horizons of digital social networks: Mixing big trace datasets with qualitative approaches. *Information and Organization*, *26(1-2)*, 1-12.

Yoo, Y. (2010). Computing in Everyday Life: A Call for Research on Experiential Computing. *MIS Quarterly, (34)2*, 213-231

Zevin, M., Coughlin, S., Bahaadini, S., Besler, E., Rohani, N., Allen, S., Cabero, M., Crowston, K., Katsaggelos, A., Larson, S., Lee, T. K., Lintott, C., Littenberg, T., Lundgren, A., Østerlund, C., Smith, J., Trouille, L., & Kalogera, V. (2017). Gravity Spy: Integrating Advanced LIGO detector characterization, machine learning, and citizen science. *Classical and Quantum Gravity*, *34(6)*.